Image Landmarks: Application

For a recent application of image landmarks, consider the Mars Rover Teach and Repeat task.

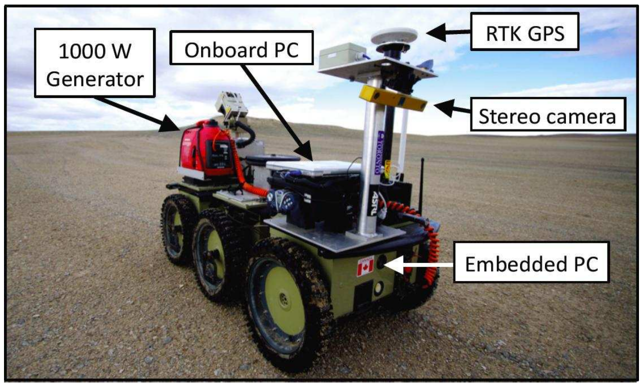

Rover @ UTIAS.

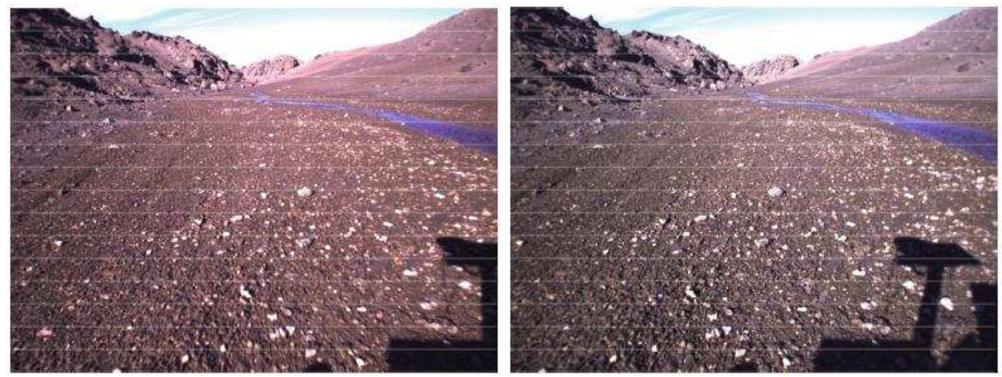

Images from Furgale, 2011.

Teach. Manually drive a rover over a path on the surface of Mars.

1. SURF keypoints (aka image landmarks) are extracted from stereo cameras on the rover.
2. The resulting keypoints are compiled into local 3D maps.

Repeat. The rover autonomously traverses the same path, using the stereo SURF points to:

1. Perform visual odometry (infer its 3D motion through the environment).
2. Match its position and pose with respect to the stored 3D maps.

Matching SURF Keypoints

SURF keypoints are extracted and matched in real time. These are used to build a 3D stereo point cloud, with each point associated with a keypoint descriptor. Collections of SURF keypoints prove to serve as reliable landmarks (modulo strong lighting changes).

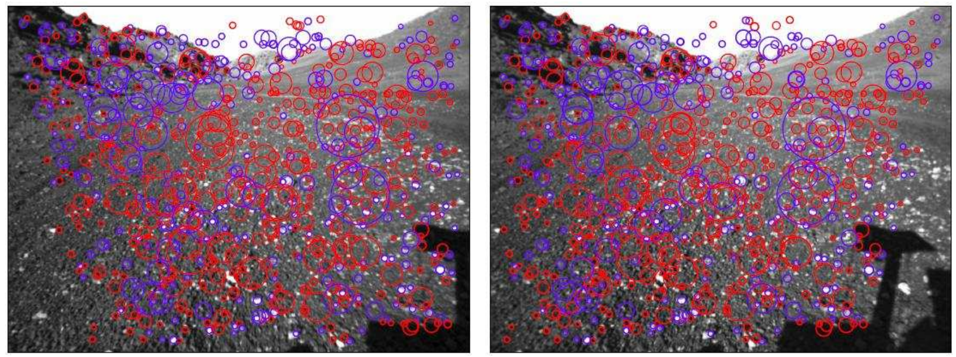

SURF points.

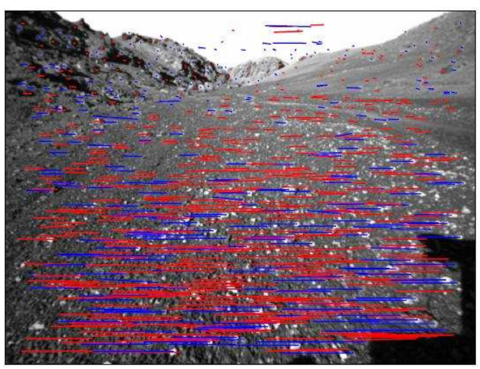

(Left-Right) Matched Stereo Points.

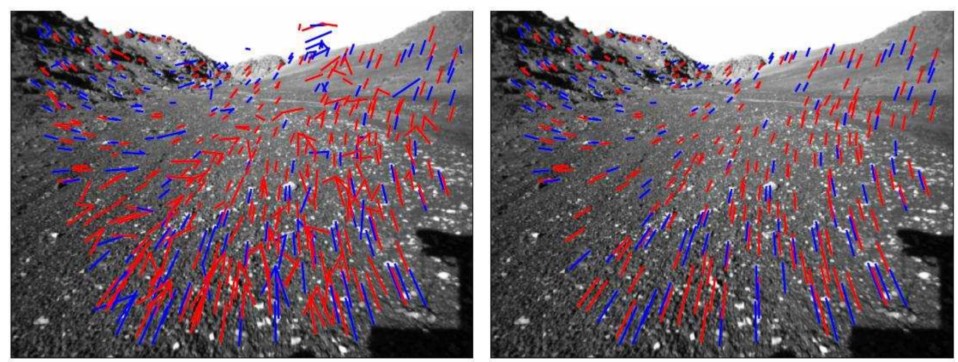

Tracked 3D Points (left) and Inliers (right) (all images from Furgale, 2011).

See http://www.youtube.com/watch?v=5bcKwrL_1As for a video of the UTIAS robot following a pre-trained path over rough terrain.

Discussion Question

Significant changes in lighting lead to failures in matching SURF points between the current image and the stored 3D map. Are we lost yet?

Sample lighting changes for the same location. Recognize it? (From Furgale, 2011.)

Lighting changes and lack of distinguishing landmarks. (From Furgale, 2011.)

How would you build image landmarks for this type of surface, so that they tolerate the range of lighting conditions illustrated above?

Google Goggles

A large scale commercial application of local image features is Google Goggles:

From http://www.google.com/mobile/goggles/#landmark

A key problem is the immense size of the image database being searched, with the current goal in the billions of images. For recent results on 100 million images in the database, see Mohamed Aly, et al, BMVC, 2011.

Motivation for Parameter Estimation

Parameter estimation is critical to following tasks, which were all utilized in the Visual Teach and Repeat task (it is less important for Google Goggles):

- match keypoints across two stereo views for the purpose of calibrating the stereo system, or for inferring 3D depth for the keypoints;
- match keypoints across multiple views of a stationary/rigid 3D object (e.g., the Mars landscape) for the purpose of visual odometry;
- match keypoints between an image of a stationary scene and a stored 3D keypoint cloud-map for the purpose of self-localization within the map.

We discuss parameter estimation starting now.

Later in this course we will apply what we learn about parameter estimation to many tasks, including stereo and multiview reconstruction.

References

For SURF features, see:

Herbert Bay, Andreas Ess, Tinne Tuytelaars, Luc Van Gool, SURF: Speeded Up Robust Features, Computer Vision and Image Understanding (CVIU), Vol. 110, No. 3, pp. 346–359, 2008.

The following papers are from UTIAS on the Teach and Repeat Task:

P.T. Furgale, Extensions to the Visual Odometry Pipeline for the Exploration of Planetary Surfaces, Ph.D. Thesis, Dept. of Aerospace Science and Engineering, Univ. of Toronto, Sept. 2011.

P.T. Furgale, and T.D. Barfoot, Visual Teach and Repeat for Long-Range Rover Autonomy. Journal of Field Robotics, special issue on Visual mapping and navigation outdoors, 27(5): 534-560, 2010.

P.T. Furgale, and T.D. Barfoot, Stereo Mapping and Localization for Long-Range Path Following on Rough Terrain. In Proceedings of the International Conference on Robotics and Automation (ICRA). Anchorage, Alaska, USA, 3-8 May 2010.

P.T. Furgale, and T.D. Barfoot, Visual Path Following on a Manifold in Unstructured Three-Dimensional Terrain. In Proceedings of the International Conference on Robotics and Automation (ICRA). Anchorage, Alaska, USA, 3-8 May 2010. ICRA 2010 Kuka Service Robotics Best Paper Award.

For a recent paper on Google goggles, see:

Mohamed Aly, Mario Munich, and Pietro Perona, Distributed Kd-Trees for Retrieval from Very Large Image Collections, British Machine Vision Conference (BMVC), Dundee, UK, August 2011.