Structure from Motion

Marc Pollefeys[1] and Richard Szeliski[2]

[1] ETH, Zurich
[2]Microsoft Research Ltd, Redmond, USA

Abstract

Structure from motion is the essential ingredient in today's large scale image-based modeling and rendering systems, which can take a collection of photos or videos and turn them into photo-realistic 3D models. In this tutorial, we review all of the essential stages in a typical structure from motion pipeline. We begin with a quick introduction to feature detection and matching, followed by two-frame relative motion estimation, pose estimation and triangulation. These elements are then combined into a sequential structure from motion algorithm. We then briefly describe self-calibration approaches. This is followed by a tutorial on bundle adjustment, the non-linear least squares formulation of structure and motion recovery that can deal with arbitrary kinds of measurements and acquisition configurations, and can produce statistically optimal estimates. Finally, we discuss issues in large-scale bundle adjustment, such as exploiting sparsity, using reduced (skeletal) representations, and using sparse iterative solvers for extremely large problems. Applications to domains such as robotics and Photo Tourism will also be presented.

Note: Marc Pollefeys will teach the first half of the course (features, sequential structure from motion, and self-calibration methods), while Richard Szeliski will teach the second half (bundle adjustment).

Syllabus: feature matching, epipolar geometry, pose estimation, triangulation, self-calibration, large scale bundle adjustment