Multi-view stereo for 3D reconstruction

Carlos Hernández
Toshiba Research Europe Ltd,
Cambridge, UK

Abstract

Acquiring 3D shape from images is a classic problem in Computer Vision occupying researchers for at least 20 years. Only recently however have these ideas matured enough to provide highly accurate results. We present a complete algorithm to reconstruct 3D objects from images using the stereo correspondence cue. The technique can be described as a pipeline of four basic building blocks: camera calibration, image segmentation, photo-consistency estimation from images, and surface extraction from photo-consistency. In this tutorial we will put more emphasis on the latter two: namely how to extract geometric information from a set of images without explicit camera visibility, and how to combine different geometry estimates in an optimal way. This algorithm is currently the top performer in the recent evaluation of multi-view stereo algorithms by Seitz et al.

Syllabus: Multi-view stereo, photo-consistency, camera visibility