Course Announcement

EE290T: 3D image processing and computer vision
Instructor: A. Zakhor, avz@eecs.berkeley.edu,
Class times: Tuesdays, 9 to 11 am, 212 Cory Hall
Pre-requisite: Signals and systems at the level of EE120, basic linear algebra

3D image processing and computer vision techniques are used in a variety of applications including gaming e.g. Kinnect, photo-tourism e.g. Photo-synth, mapping e.g. streetview, scene reconstruction and modeling for cultural heritage, consumer electronics e.g. mosaicing of camera images, mobile augmented reality e.g. in cell phones, virtual reality, and entertainment industry e.g. special effects. Many of these applications use other sensors besides visible range cameras in order to improve performance or facilitate the required processing. These include projectors, laser scanners, inertial measurement units. For example Kinnect uses infrared projectors in addition to cameras in order to perform gesture recognition.

This course is intended to give students at the advanced undergraduate or introductory graduate level, and researchers a self contained introduction to the geometry of three dimensional image processing and computer vision. The course concentrates on the analysis of scenes that contain a number of rigidly moving objects captured by one or two cameras or other sensors, which could be either stationary or moving. We seek to answer the following questions: to what extent and how can we estimate the 3D shape and position of each object? To what extent can we recover the motion of each object relative to the camera or the motion of the camera itself? To what extent can we recover a model of the geometry of the camera itself? Traditionally these questions are referred to as structure from motion problem. The course deals with algorithms designed to address these questions, namely, to estimate 3-D structure, motion, and camera calibration from a collection of images. In this sense, the course deals with how to go from 2-D images to 3-D models of the geometry of the scene.

The course will cover the following topics: camera model and calibration, single-view metrology, triangulation, epipolar and multi-view geometry, two-view and multi-view stereo, structured-light, feature tracking and matching, structure-from-motion, shape-from-silhouettes and 3D modeling and applications. The grade is determined by a series of written and programming assignments, as well as a class project of student's choosing and instructor's approval.