Course Announcement

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

This course is intended to give students at the advanced undergraduate or introductory graduate level, and researchers in computer vision, robotics, and computer graphics, a self contained introduction to the geometry of three dimensional image processing and computer vision. This is a study of the reconstruction of 3-D models of objects or scenes from a collection of 2D images. The sense of vision plays an important role in the life of primates; it allows them to infer spatial properties of the environment that are necessary to perform crucial tasks for survival. A visual system in broad terms is a collection of devices that transform measurements of light into information about spatial and material properties of a scene. Among these devices, we need photosensitive sensors such as camera or retina, as well as computational mechanisms such as computer or a brain that allows us to extract information from the raw sensory readings. The factors that affect our visual measurements include photometry, i.e. the illumination and material properties of objects, geometry, i.e. shape of the objects in a scene, and dynamics of the environment, i.e. movement of the objects in the scene. In this course, our goal is to infer the latter two aspects of objects and scenes from a collection of images; to the extent that illumination and material properties affect visual measurements, we have to take their influence into account in designing algorithms capable of inferring the shape and motion of objects in a scene.

The course concentrates on the analysis of scenes that contain a number of rigidly moving objects; we seek to answer the following questions: to what extent and how can we estimate the 3D shape of each object? To what extent can we recover the motion of each object relative to the camera? 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 text book for the course is “An Invitation to 3-D Vision” by Ma, Soatto, Kosecka, and Sastry. There grade is determined by a class project of students’ choosing and instructor’s approval.