ENGN 2502 3D Photography / Spring 2018 / SYLLABUS

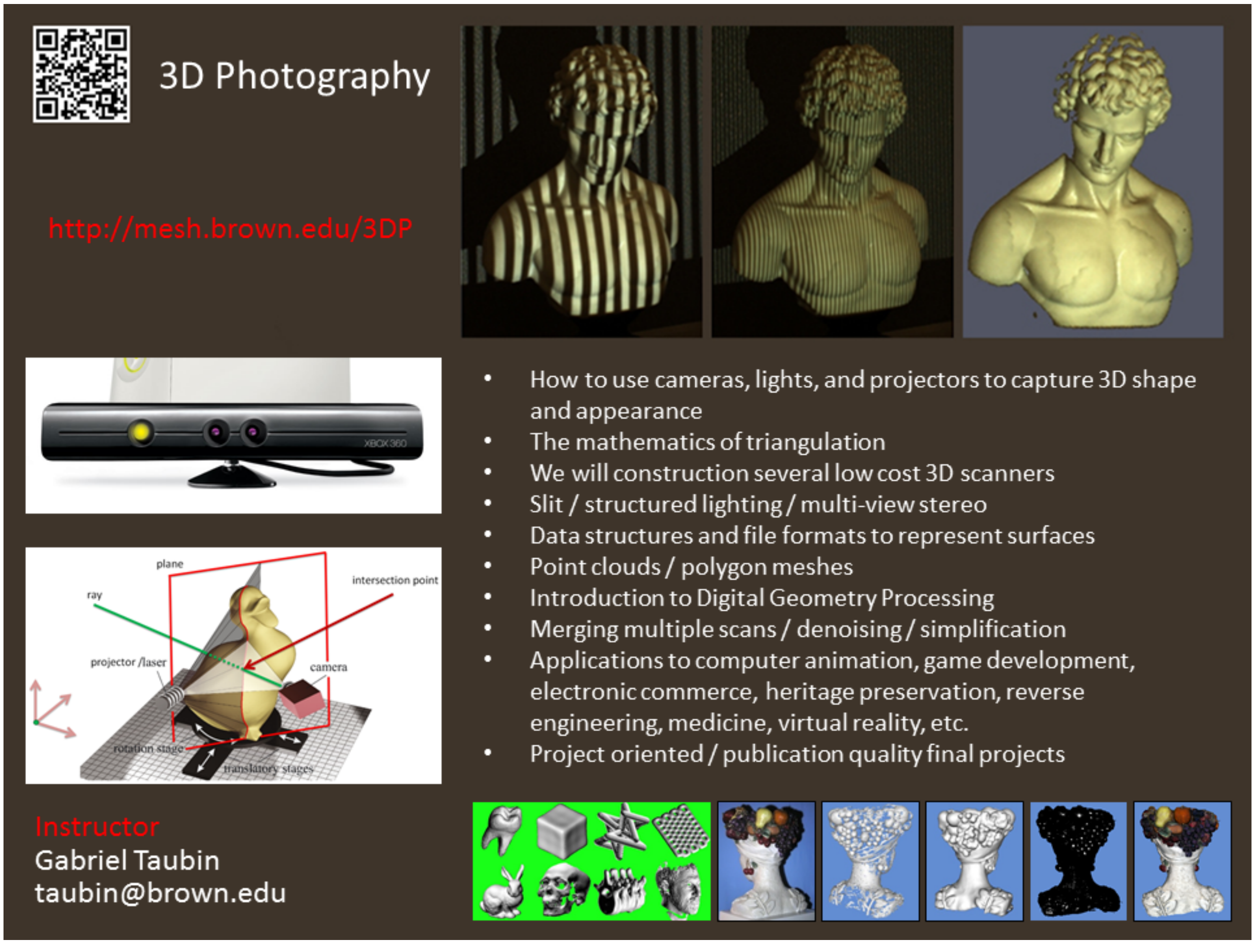

Description of the proposed course

Over the last decade digital photography has entered the mainstream with inexpensive, miniaturized cameras routinely included in consumer electronics. Digital projection is poised to make a similar impact, with a variety of vendors offering small form factor, low-cost projectors. 3D television in the home is now a reality. As a result, active imaging is a topic of renewed interest in computer graphics and computer vision. In particular, low-cost homemade 3D scanners are now within reach of students and hobbyists with a modest budget. This course provides the students with the necessary mathematics, software, and practical details to leverage projector-camera systems in their own 3D scanning projects. An example-driven approach is used throughout, with each new concept illustrated using a practical scanner implemented with off-the-shelf parts. First, the mathematics of triangulation is explained using the intersection of parametric and implicit representations of lines and planes in 3D. The particular case of ray-plane triangulation is illustrated using a scanner built with a single camera and a modified laser pointer. Camera calibration is explained at this stage to convert image measurements to geometric quantities. A second example uses a single digital camera, a halogen lamp, and a stick. The mathematics of rigid-body transformations are covered through this example. Next, the details of projector calibration are explained through the development of a classic structured light scanning system using a single camera and projector pair. A minimal post-processing pipeline is described to convert the point-based representations produced by the example scanners to watertight meshes. Key topics covered in this section include: surface representations, file formats, data structures, polygonal meshes, and basic smoothing and gap-filling operations. The course concludes by detailing the use of such models in 3D printing and for web-based applications.

Instructor

Prof. Gabriel Taubin, taubin@brown.edu, Barus & Holley 328

Teaching Assistant

Peter Walecki, peter_walecki@brown.edu, Barus & Holley 317

Meeting Times & Room

Tue-Thu 1:00-2:20, ERC 125.

Learning goals for students

This course has two main goals: 1) to learn the theory supporting various techniques used to acquire and process 3D data; and 2) to develop practical skills necessary to implement such systems.

Concrete learning objectives

Upon completion of the course students will be able to demonstrate knowledge of basic techniques to recover 3D shape from images, awareness of state-of-the-art commercial systems, as well as practical knowledge in the implementation of 3D shape capture and processing systems.

Student assessment and evaluation criteria

The course is project oriented. The evaluation is based on a number of software development assignments, the construction of a number of 3D scanning systems, and a final projects. As new concepts and techniques are introduced in the lectures, the students gradually implement various components of full 3D data capture, processing, optimization, and display systems. In the final project the students, working in groups, work on novel problems. They have to produce a working implementation, a publication-quality report, and present their work to the class in a conference setting, with an ultimate goal of submitting the work for publication.

Course calendar/overview

Introduction

The course topics, goals, and motivation are presented. A brief review of existing commercial and academic 3D scanners is provided, including active and passive systems. The specific 3D scanners explained and built using the material in this course are introduced. The unifying concepts of ray-plane and ray-ray triangulation are used to link existing systems and those presented in this course. The limitations of these 3D scanners are explained, particularly the restriction to solid surfaces. A general 3D scanner pipeline is presented.

1. Overview of 3D Scanning
   a) triangulation using active and passive systems
   b) commercial scanners
   c) seminal academic publications
   d) unifying concept of ray-plane and ray-ray triangulation
2. 3D Scanners Described in this Course
   a) laser-striping
   b) planar shadows [Bouguet and Perona 1998]
   c) structured light using Gray codes
3. Concepts Described in this Course
   a) mathematics of triangulation
   b) camera, planar light source, and projector calibration
   c) point cloud alignment, surface reconstruction, and mesh processing
   d) point cloud and surface visualization
4. The 3D Scanning Pipeline
   a) controlled illumination sequence
   b) correspondence assignment (ray-to-plane or ray-to-ray)
   c) triangulation (ray-to-plane or ray-to-ray)
   d) multi-view point cloud registration
   e) surface reconstruction
   f) mesh processing (gap-filling and filtering)
   g) visualization

The Mathematics of 3D Triangulation

The general mathematics of triangulation are presented. Cameras and projectors are treated as devices used to measure geometric quantities (e.g., points, line, and planes), thus the details of calibration are left to later sections. Key topics include parametric and implicit representations of lines and planes in 3D. Least-squares solutions for the intersection/reconstruction of 3D points are presented.

1. Representation
   a) parametric representation of lines/rays
   b) implicit representation of planes
2. Triangulation/Reconstruction
   a) ray-plane intersection

    b) ray-ray intersection

3D Scanning with Swept-Planes: Laser Striping and Planar Shadows
This section will present the practical implementation details for two specific 3D scanners using ray-plane triangulation: (1) a "classic" laser stripe scanner consisting of a single digital camera and a laser pointer modified to project a single line, and (2) the "3D Desktop Photography" system originally proposed by Bouguet and Perona in 1998. For both methods, a spatio-temporal method is presented to assign ray-to-plane (i.e., pixel-to-plane) correspondences.

1. The "Classic" Laser Stripe Scanner
    a) hardware components and assembly
2. "3D Desktop Photography" [Bouguet and Perona in 1998]
    a) hardware components and assembly
    b) similarity to laser-stripe scanners
3. Assigning Correspondences
    a) spatio-temporal processing to assign pixel-to-plane correspondences

Camera and Swept-Plane Light Source Calibration
This section covers the mathematics and software required to calibrate the camera and illumination used in the previous pair of "swept-plane" scanners. Intrinsic and extrinsic calibration is achieved using planar checkerboard patterns following the well-used method of Tsai [1987]. Coordinate transformations between the world and the camera systems are defined.

1. The Pinhole Camera Model
    a) intrinsic parameters
    b) extrinsic parameters
    c) lens distortion model
2. Calibration
    a) intrinsic calibration with planar patterns
    b) extrinsic calibration (i.e., pose estimation) with known intrinsic parameters
3. Applications
    a) live demonstration of camera calibration (MATLAB Camera Calibration Toolbox)
    b) mapping camera pixels to optical rays

Reconstruction and Visualization using Point Clouds
Point clouds are reconstructed using the swept-plane scanners from a single viewpoint. File-formats and methods to visualize point clouds are presented.

1. Point Cloud Reconstruction
    a) results for laser striping
    b) results for planar shadows
2. File Formats
    a) VRML, SFL, and related formats
3. Visualization
    a) point splatting

    b) software (Java-based viewer and Pointshop3D)

Structured Lighting
This section will present structured lighting as a popular method to overcome some limitations of swept-plane scanners. Emphasis will be placed on the possible illumination patterns that can be used to minimize the acquisition time, especially the well-known example of Gray codes. Methods for decoding the sequences, to establish per-pixel ray-to-plane correspondences, will also be covered.

1. Building a Structured Light Scanner
    a) hardware components and assembly
2. Structured Light Sequences
    a) swept-plane sequence (i.e., single column/row per-exposure)
    b) binary encoding of swept-plane sequence
    c) Gray code optimization of binary encoding
3. Decoding
    a) method for decoding binary/Gray codes
    b) introducing code redundancy

Projector Calibration and Structured Light Reconstruction (Lanman, 15 minutes)
This section will describe how to calibrate projectors as "inverse" cameras. Extensions to existing camera model and calibration methods will be presented. The basic method will require projecting a checkerboard, or structured light sequence, on a plane will known extrinsic calibration (e.g., a plane with a printed checkerboard also present).

1. The Pinhole Projector Model
    a) the projector as the "inverse" of a camera
2. Projector Calibration using a Calibrated Camera
    a) intrinsic calibration with planar printed and projected patterns
    b) extrinsic calibration with known intrinsic camera/projector parameters
3. Applications
    a) live demonstration of projector calibration (MATLAB Camera Calibration Toolbox)
    b) mapping projector rows/columns planes
4. Structured Light Reconstruction
    a) reconstruction results for structured light sequences

Combining Point Clouds Recovered from Multiple Views
This section will describe how to reconstruct a complete object model by merging reconstruction from multiple viewpoints. A manually-assisted Iterative Closest Point algorithm will be presented for aligning multiple point clouds.

1. Merging Multiple Views
    a) capturing multiple viewpoints (e.g., manual rotation, turntables, time-multiplexing)
    b) aligning point clouds using Iterative Closest Point (ICP)
    c) live demonstration of manually-assisted alignment procedure

Surface Reconstruction from Point Clouds
This section will describe methods and software for extracting a polygonal mesh from merged point clouds.

Elementary Mesh Processing
This section will describe basic data structures and algorithms for mesh processing. Key topics include the half-edge data structure.

Conclusion
This section will summarize the key mathematics, software, and practical details presented in the course. The general 3D scanning pipeline will be reviewed. A brief discussion of more advanced topics and recent academic publications will be provided. This section will conclude with a brief discussion of rapid prototyping, entertainment, cultural heritage, and web-based applications for 3D scanning systems. Audience questions will be taken afterward.

1. Course Summary
    a) 3D scanning using ray-plane triangulation
    b) camera and projector calibration
    c) point cloud alignment, surface reconstruction, and mesh processing
    d) point cloud and surface visualization
2. Recent Progress in 3D Scanning
    a) scanning without triangulation
    b) progress in passive systems
    c) progress in active systems
    d) projector-camera systems
3. Application Areas
    a) 3D scanning for entertainment, cultural heritage, and web-based applications

Listing of associated readings/materials

The instructor has taught a short version of this course at Siggraph 2009 and Siggraph Asia 2009, two of the top conferences in Computer Graphics and Interactive Techniques. A web site was developed to complement these short courses http://mesh.brown.edu/byo3d with material, such as detailed course notes http://mesh.brown.edu/byo3d/notes/byo3D.pdf, that will be covered in much more detail in this semester-long course.