

Due: Monday March 21 at 6:00am

Assignment 4—Depth from Stereo

In a nutshell

Implement an application to recover a 3-D point cloud from a pair of images of a real world scene and to render the scene from novel viewpoints.

Details

1. Scene
 - a. Print out a black and white checker pattern.
 - b. Tape the checker pattern to a table top (i.e. a flat surface).
 - c. Put an object about 10cm x 10cm x 10cm in size on the checker pattern.
2. Input photos
 - a. Take at least 2 photographs of your object.
 - b. Each photograph should capture the entire checker pattern minus what is occluded by the object.
 - c. Each photograph should capture the entire object.
 - d. There should be a baseline of 10cm between the acquisition viewpoints.
3. Calibration
 - a. Intrinsic—OK to use results of previous assignment
 - b. Extrinsic—use checker pattern visible in each image
4. Depth extraction
 - a. For each pair of images, for each pixel in the left image find the corresponding pixel in the right image.
 - b. Only work with pixels that are on the checker or on the pattern (disable “background” pixels beforehand).
 - c. For each correspondence found, triangulate a 3-D point.
5. Rendering
 - a. Show the recovered point cloud as 3-D points with color
 - b. Make a 10 second 30 frames per second 640x480 video illustrating the recovered points
6. Extra credit
 - a. Texture mapped rendering: triangulate the points in 2-D on the left image, OK to use off-the-shelf Delaunay triangulation software; texture the resulting mesh with the left image (5%)
 - b. Fill in holes using pull-push (see Lumigraph, SIGG 1996 paper) (2%)
 - c. Use multiple image pairs to obtain a complete model of the object (3%)

Turn in

Upload archive on Blackboard containing:

- Your source code and binaries
- Your input images
- Your movie
- README.txt file with GUI description and a list of extra credit features attempted

© Popescu 2011