CSC420— Introduction to Image Understanding, Fall 2011

Project: Scene Properties

Due: 3:00pm, Wed., Dec. 7, see below for intermediate milestones.
This project is worth 60 marks towards your grade in this course.

The image sequence Flowers_03.avi will be available on CDF and from the course home page. Due to its size, do not make a separate copy on CDF. The handout code `proj.m` reads selected frames from this sequence, you are free to modify this code in any way. I have installed a software package from `VLFeat.org` on CDF which can be used for extracting SIFT features. In general, you can use any freely available software for this project, providing that you provide a clear acknowledgement in your reports.

Individual projects can involve any combination (including just one) of the following topics. Projects which are done by a group of n students, $n \leq 3$, should involve a selection of n different topics from the list below.

1. **Detect Flowers.** Detect the bright yellow flowers in the image sequence, and identify them as distinct individuals throughout the sequence. This means that your code keeps track of unique labels or names for each flower, and maintains their names across the sequence. For efficiency, you do not need to process every frame of the sequence, although you can if you want to. But try labelling only every n^{th} frame, where n is a variable, say $n = 20$ or 50.

2. **2D Motion Estimation.** For short subsequences (say, $T = 20$ to 50 frames long) use image motion estimation to track several hundred automatically selected image patches through every frame of the subsequence. This is roughly similar to the code in `motionTutorial.m`, although here you are automatically selecting the patches to track. Also, you must track each individual patch using similarity transforms (see Szeliski, p.36), which require the estimation of four linear degrees of freedom for the pose of a patch in each new frame. Alternatively, you could try using affine transforms (Szeliski, p.37) for the frame to frame motion of a patch. You can start with the `motionTutorial` code, however you must modify it to estimate the appropriate four or six degrees of freedom per patch.

   Also consider the problem of automatically identifying the patches that have been tracked correctly throughout the T frames. That is, a correctly tracked patch should be in correspondence with a particular scene location throughout the T frames.

3. **Projective Reconstruction (Part I).** As a first step in projective reconstruction, consider two successive frames separated by roughly $T = 20$ to 50 frames, that is, consider I_t and I_{t+T}. Estimate the F matrix for this image pair. In addition, estimate a homography H which accounts for a "dominant" or reference plane in the scene (e.g., it accounts for the positions of as many feature points as possible). This homography H is related to F by $F = [\vec{e}]_{\times} H$, where $\vec{e}$ is the left epipole of F (see Szeliski, eqn 7.29, p.353). Rewarp one of these images by the homography H, so that the dominant plane is stationary in the resulting image pair. Describe precisely the epipolar lines, and the displacement of corresponding points, for this new pair of images.

4. **Projective Reconstruction (Part II).** Given the left epipole $\vec{e}$ and the homography H from part I above, compute a projective 3D reconstruction of corresponding patches in the scene. For example, you could use the synthetic camera matrices described in equation (7.34) on p.354 of Szeliski's book, and then triangulation for the projective depth. Alternatively, you could use the pair of images from part I for which one image has been warped by the estimated homography. In either case the projective depth will be related to the distance from the reference plane identified in part I.

In all cases choose a solution strategy which is not overly specific to the particular video sequence, Flowers_03.avi. We have several sequences of that flower bed taken that same day, and your approach should be expected to work on those. In addition, it should be applicable to roughly similar sequences. In your report, explain what would need to be done to apply your solution to a new sequence of a different flower bed.

What to hand in. On Tues., Nov. 15, at the beginning of class, hand in a short description of your proposed project. If you will be working in a group (group sizes are limited to at most three people), indicate the initial responsibilities of each group member. Include a clear description of what you propose to have done by the date of the intermediate milestone, Dec. 2nd, and your goals for the overall project. This initial report will be worth 5% of your final mark. I will use it to give you feedback on your project.

For Dec. 2, at 11am (at the beginning of class), write a short report describing the progress to date. Include a description of important papers, or sections of the text, that you have found useful for your project. Clearly describe the algorithms you have experimented with, along with both the successes and failures of these approaches. Briefly describe what is left for you to do on the project. In addition, in class on that date, you will be asked to present a short demo of the current state of your project. This milestone report and presentation will be worth 15% of your final mark.

Your final project report is due on Wed., Dec. 7 at 3pm (extensions up to Dec. 16th will be considered upon request). This report is to be an updated version of the milestone report, without the future work section. Moreover, for a group project, it must have a page describing the contributions of each member of the group. We will be marking both the content and style of your written report. Package the report as a pdf file, ProjReport.pdf, and submit it to CDF. In addition, submit your project code as ProjCode.zip. (Include URLs for any freely available software that your project depends on, but not the software itself.) This portion of the project will be worth 40% of your final mark.