Assignment A6: 3D from Stereo Image Pairs (adapted from Prof. Gerig, Spring 2013)

CS 6320
Spring 2014

Assigned: 3 February 2014

Due: 19 February 2014

For this problem, handin a lab report pdf (include name, date, assignment and class number in pdf) which studies camera calibration. You should handin the report pdf as well as the Matlab code used in the study. The code should conform to the style requested in the class materials.

In addition, please turn in a hardcopy of the report in class before the start of class on February 19, 2014.

The objective of this assignment is to take a stereo pair of images, and explore the epiploar gemeotry and the ability to recover depth from a pair of images. For this assignment, you will need two stereo pairs of your choice, one with arbitrary translation and rotation of the two camera views, and one where you simulate a translated camera by shifting one camera horizontally and by recording the distance of the baseline.

4a: Epipolar geometry from F-matrix

We discussed the essential and fundamental matrices and their properties to establish a relationship between corresponding points in the left and right cameras, and relationship of a points and their associated epipolar lines. Following the examples discussed in the slides (CS6320-CV-F2013-chap11-multiple-views-part-II-animated.pptx) and chapter 7.1 in the textbook, you will explore the following steps:

- Shoot a stereo-pair of a scene of your choice with your camera, where you take a left and a right picture where you translate and rotate the camera.
- Specify a set of corresponding pixel landmark pairs in the left and right camera views.
- Calculate the F-matrix using instructions as discussed in the class and the slides.
- Choose a point in the left image and calculate the epipolar line in the right image. Display the point and the line as overlays in your images.
- Do the same for a point in the right image and epipolar line in the left image.
- Calculate the position of the epipole of the left camera (see last two slides for instructions). Discuss if the calculated position seems reasonable.

4b: Dense Distance Map from Dense Disparity

Search for corresponding points along same scan lines (rectied image pairs). Given optimal correspondence based on the match criteria (here use normalized cross correlation), calculate disparity and associated depth Z for each pixel. The result is a Z-buffer, or distance image.

Given the high computational expense, you are strongly encouraged to choose small size images for this experiment. Please note that images obtained by your camera can be easily subsampled using standard image display programs.

Use a rectied image pair to simplify correspondence search along horizontal lines. You can create two images by simply shifting a camera horizontally along its image plane (e.g., placing a camera against a wall and shifting it horizontally along this wall). Remember to get a precise estimate of the baseline length of this shift.

- Write a Matlab program to scan two images in a TV-scan like fashion, i.e., a line-by-line, row-by-row algorithm where at each location, the nxn neighborhood of each pixel is available for calculations. An algorithm is needed for simultaneous calculations on two images with different scan locations.
- For each nxn window in the left image, calculate normalized cross-correlation (see course slides) for each window in the right image. Remember that for each window in the left image, you need to search the whole horizontal line for the best matching window. Choose the best corresponding location as the location with the highest correlation. It is suggested to write a version for 3x3 and for 5x5 windows. Please

note that the image border cannot be processed, i.e., you need to leave a frame of 1 or 2 pixel width, respectively.

- Calculate disparity for each corresponding pair of pixels. The disparity in pixel units needs to be converted into mm-units using the known pixel size from camera calibration.

- Using the disparity, calculate the depth Z at each pixel location.

- Display the stereo pair and the resulting depth image side by side and discuss your result.

- A nice way to display depth images is the calculation and display of the slope, i.e., the image gradient. This can be obtained by calculating intensity differences of neighboring pixels (horizontally or vertically) and then display those as an image. Please note that differencing results in negative and positive values, so that double-type image matrices need to be shifted into the positive range of integer images to be displayed properly.