

Advanced Vision Practical 2

Bob Fisher
School of Informatics

February 2016

Abstract

This describes the second assignment for assessment on Advanced Vision. The main goal is to register and fuse 16 range images and then extract some structures from the fused data. The assignment is due: **4pm Thursday 17 March**. You must do this practical in teams of 2, and submit 1 PDF report only. Your practical partner must be a different person from the first practical. There will also be an assessed live demonstration of your practical.

Task Background

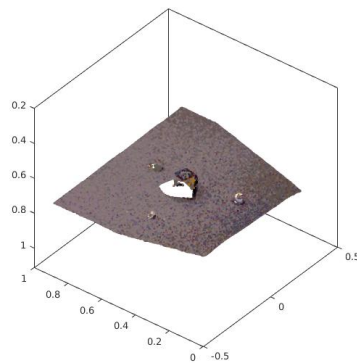

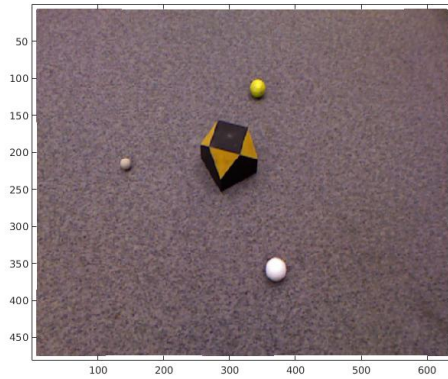
At the URL:

<http://homepages.inf.ed.ac.uk/rbf/AVDATA/AV216DATA/>

you will find 5 files:

1. `av_pcl.mat` (92 Mb)
2. `fscatter32.m` - A support file for `plotpcl.m`. We use a modified version of the file at www.mathworks.com/matlabcentral/fileexchange/2993-fscatter3-m
3. `plotpcl.m` - displays the data for each frame
4. `sphereFit.m` - fits a sphere to a set of points
5. `sphereFit_test.m` - demonstrates `sphereFit.m`

`av_pcl.mat` contains the data for 16 views captured by a Kinect depth sensor. It contains a cell structure with 16 entries, for the 16 views. You can extract the i^{th} $480*640*6$ image from this by: `I=pcl_cell{i}`; Values `I(:, :, 1:3)` are an RGB image, and values `I(:, :, 4:6)` are registered XYZ points. `plotpcl(pcl_cell{i})` will display the 3D point cloud, the RGB image and a registered depth image. Examples from frame 1 show the colour image, the depth image and a 3D point rendering of the data:



The overall task for this assignment is to fuse the individual XYZ point clouds, extract the 9 planes, and create as much of a complete model as possible. To do this you need to write a set of programs that can:

1. Extract and remove the ground plane from the image data
2. Extract and describe the 3 spheres
3. Use the 3 spheres to register and fuse the XYZ data from all images
4. Extract the 9 planes from the data
5. Build as complete a 3D model as possible and characterise the surfaces.

Each of these is described in more detail below.

Background plane extraction

The textured grey background is a flat planar surface. A plane containing most of the background data points should be extractable using one of the techniques from the lecture examples.

All the 3D points from the background plane can be discarded. The plane extraction might leave some small isolated 3D points that can also be removed.

Sphere extraction

There are 3 spheres, each with different sizes (diameters yellow 40 mm, large white 40 mm and small yellowish 25 mm) and colours. After you have removed the background plane, the test object and spheres should be the main regions that remain.

You can use the intensity image to help isolate the points from each sphere, and then possibly fit a sphere to the 3D data using `sphereFit`. What we need from this stage is the 3D position of the centres of the 3 spheres. If the sphere fit is bad, i.e. produces an unrealistic sphere, then use the centre of mass of the extracted data points.

Registration

Choose one of the views to be the baseline view that all of the views is to be registered to. E.g., the one with the largest view of the target object. Registration and fusion consists of 3 steps:

1. Pairing the extracted spheres between the baseline and current images.
2. Estimating the transformation $\mathbf{E} = [\mathbf{R} \mid \vec{t}]$ that maps the 3D sphere centres in the new image to the baseline image. Estimate the rotation \mathbf{R} mapping the current 3 sphere centres into the orientation of the baseline image. Use the lecture algorithm with the 3 vectors formed by the difference between each pair of sphere centres. Estimate the translation \vec{t} by averaging the translation vector from each of the rotated sphere centres to the baseline sphere centres. The quality of the registration can be partially evaluated by computing the normal vector of the floor and evaluating how parallel they are, similarly for the black square on top of the object to register.
3. Use \mathbf{R} and \vec{t} to transform all the points from the block in the current image into the baseline view.

Object plane extraction

Once you have transformed all of the 16 point sets into a common coordinate system, use methods like that given in the lecture to extract the 9 planes. If the registration is too noisy to determine which points belong to each plane, you might need to use some of the colour information associated with each data point. If a point set is badly registered, then ignore it. Report which were deleted. Explain why. Show images of the misalignment.

This gives 9 infinite planes.

Complete model building

Once the 9 infinite planes are extracted, compute 9 bounded 3D planar patches that correspond to the 9 faces of the cube. You could represent these extracted faces by 100 coloured 3D points on each face (or a set of coloured triangles if you are feeling ambitious).

Create a table showing the angle in degrees between the outward facing normals of each surface. This should be a 9x9 table. Include the table in your report.

Your Report

Each team writes a single report that describes:

- The algorithms that you used for each stage of the process.
- How well the algorithms performed on the supplied test data. Show the statistical results and images requested above.
- Show example images of each processing stage, including a few examples of detected balls, background planes and final planar surfaces.
- Show examples of successful and unsuccessful point set registrations.
- Discussion on performance: successes and failures, causes of failures and potential remedies.
- Your code. Do not include code that was downloaded from the AV or IVR web sites.

Other Comments

1. You can use the lecture example code from:
<http://www.inf.ed.ac.uk/teaching/courses/av/MATLAB/>

Assignment Submission

Submit your report in PDF online by 4pm Thursday 17 March. The online submission line is:

```
submit av 2 FILE
```

where FILE is the name of your PDF file. Submit 1 file for each team.

Live Demonstration

There will also be a demonstration session assigned between 9:00-13:00 on Friday March 18, where you will have to demonstrate your assignment (using the same dataset). We'll email you about the location and schedule.

You will need your matlab program to show, for each of 16 frames result images showing

1. The extracted background plane.
2. The located balls.
3. The located block data.
4. The full set of registered 3D point cloud after fusing each new image's 3D points.

After fusing all the 3D points, show

1. All of planes extracted to make a 3D block viewable in 3D.
2. The angle between the top plane and each of the 4 diagonal planes.

The live demonstration will use the same dataset as specified above.

The assignment is estimated to take 10 hours coding/test and 5 hours report writing per person, resulting in a 5 page report plus the code appendix. You must do this assignment in teams of 2. You must find your partner and email Bob Fisher (rbf@inf.ed.ac.uk) the name of your partner. A single, joint, report is to be submitted. Split the work so that each partner does most work independently (i.e. share the work rather than duplicate it).

The assignment will be marked as follows:

Issue	Percentage
1. Clear description of sensible algorithms used	40%
2. Performance on supplied data set	20%
3. Clear Matlab code	10%
4. Discussion of result quality and causes of any failures	10%
5. Live demonstration performance on data set	20%

Publication of Solutions

We will not publish a solution set of code. You may make public your solution **but only 2 weeks after the submission date**. Making the solutions public before then will create suspicions about why you made them public.

Good Scholarly Practice: Please remember the University requirement as regards all assessed work for credit. Details about this can be found at:

<http://www.ed.ac.uk/schools-departments/academic-services/students/undergraduate/discipline/academic-misconduct>

and at:

<http://web.inf.ed.ac.uk/infweb/admin/policies/academic-misconduct>

Furthermore, you are required to take reasonable measures to protect your assessed work from unauthorised access. For example, if you put any such work on a public repository then you must set access permissions appropriately (generally permitting access only to yourself, or your group in the case of group practicals).

Plagiarism Avoidance Advice

You are expected to write the document in your own words. Short quotations (with proper, explicit attribution) are allowed, but the bulk of the submission should be your own work. Use proper citation style for all citations, whether traditional paper resources or web-based materials.

If you use small amounts of code from another student or the web, you must acknowledge the original source and make clear what portions of the code were yours and what were obtained elsewhere. You can ignore this condition for the AV lecture examples, which can be used freely.

The school has a robust policy on plagiarism that can be viewed here:

<http://web.inf.ed.ac.uk/infweb/admin/policies/guidelines-plagiarism>.

The school uses various techniques to detect plagiarism, included automated tools and comparison against on-line repositories. *Remember: a weak assignment is not a ruined career (and may not reduce your final average more than 1%), but getting caught at plagiarism could ruin it.*

Late coursework policy

See: <http://web.inf.ed.ac.uk/infweb/student-services/ito/admin/coursework-projects/late-coursework-extension-requests>