

Intro to Image Understanding (CSC420)

Projects

Proposal Deadline : Nov 2 (Sunday), 11.59pm, 2014

Report Submission Deadline : December 7 (Sunday), 11.59pm, 2014

Max points for report: 30, max points for presentation: 20

Projects can be done individually or in pairs. Most projects are quite big, so pair work is actually encouraged. If a project is done in a pair, each student should still hand in his/her own report and defend the project on his/her own. From the report it should be clear what each student contributed to the project. By December 7 you will need to hand in the project report including code. Make the code organized and documented, possibly including scripts that run your pipeline. In the oral defense you'll need to run some of your code and be able to defend it. The tentative schedule for the oral defense is December 16. The grade will evaluate a project report (30% of the grade) and an oral presentation (20% of the grade).

Whenever you use ideas, code or data from a paper, forum, webpage, etc, you need to cite it in your report. Whenever you use code available from some paper or method, you need to include a short description of the method showing your understanding of the technique. If there were multiple options for techniques or code, please also explain why you chose a particular one.

The grade will take into account the following factors:

- Your ideas to tackle a problem: how appropriate the techniques you chose are for the problem. Coming up with novel ideas is obviously a big plus.
- Your implementation: the accuracy of the solution, speed, and partly also how organized the code is (scripts that run the full pipeline or specific subtasks, documentation). The grade will take into account also the performance of your solution with respect to other students' results.
- Whether you implemented a technique yourself or used code available online
- How you present your results in the report: ideally you would show your results for each task by including a few pictures in your report. Even more ideally, you would show good cases where your method works and also bad cases where it doesn't. Providing some intuitive explanation of the failure cases is a plus.
- Thoroughness of your report: How well you describe the problem and the techniques you chose, how well you analyzed your results and whether your citations to the techniques you used are appropriate.

It may be that the project has too many questions and you might not be able to solve them all. Don't worry. The goal is to teach you how to solve cool problems and hopefully you have fun doing it. We will evaluate relative to how everyone does. Think of it as a competition.

You may discuss the project with your colleagues. Of course, be also careful since it's a competition, after all. ;) Do **not** put the video clips we provide online or share with anyone.

Please submit a short project proposal by Nov 2. A project proposal is your commitment to a particular project. Write down what you chose, who your project partner is (if any), and a few sentences describing your plan how to tackle the problem. If you are working on a non-listed project (given that you were approved by the instructor), please write a longer project proposal, outlining your general ideas and difficulties you think the project has.

Project 1

Everyone has a large movie or series collection at home these days. This project is about efficient browsing of such video content. You will be given a few video clips from a series Buffy the Vampire Slayer. Here are the tasks:

- (a) Detect shots in the videos. A shot is a set of consecutive frames with a smooth camera motion. For this you may need to classify the frame into day or night and perhaps use different thresholds for day and night scenes.
- (b) How would you evaluate how well you are detecting the shots? Please compute the accuracy of your method.
- (c) Detect faces in each frame.
- (d) Perform face tracking by correctly associating a face detection in the previous frame to a face detection in the current frame.
- (e) Find face tracks that belong to Buffy. You will need to train a classifier for Buffy. To do this you will take a few images of Buffy's face and faces of other characters (the images will be provided to you), compute image features, e.g., HOG, and train a Buffy-vs-rest classifier, e.g., SVM. Once trained, you will predict the identity (Buffy or rest) of each face detection in the video. That is, you'll take each face detection (a crop in the image specified by the face box), compute appropriate image features, and use your classifier to predict who the face belongs to. How would you decide whether a full face track is Buffy or not?
- (f) **Extra credit:** Predict whether Buffy (the face tracks predicted to be Buffy) is talking. How would you do that?

Based on the above, please find and show the following shots in the video:

1. It's a night scene
2. It's a close-up shot of someone (a person's face is taking most of the image height)
3. It's a close-up shot of Buffy

Project 2

Autonomous driving is one of the major research venues these days. A lot of effort is devoted to it by both the academics as well as industry. In this project you'll familiarize yourself with some of the most important problems that arise in the field of autonomous driving.

The input to your algorithm is a stereo image pair and the camera parameters. You will also have available a set of training images where the cars have been annotated with 2D bounding boxes as well as viewpoint. Furthermore, you'll have a few images where the road has been annotated. Here are the tasks to solve:

1. Compute disparity between the two stereo images. We do not mind if you use existing code as long as you include a description of the algorithm you used, showing you understand what it is doing.
2. Compute depth of each pixel. Compute 3D location of each pixel.
3. Train a road classifier on a set of annotated images, and compute road pixels in your image. Which features would you use? Try to use both 2D and 3D features.
4. Fit a plane in 3D to the road pixels by using the depth of the pixels. Make sure your algorithm is robust to outliers.
5. Plot each pixel in 3D (we call this a 3D point cloud). On the same plot, show also the estimated ground plane.
6. Detect cars in the image. You can use the pre-trained models available here: http://kitti.is.tue.mpg.de/kitti/models_lsvm.zip, and detection code available here: <http://www.cs.berkeley.edu/~rbg/latent/>
7. Train a classifier that predicts viewpoint for each car. The viewpoint labels are in 30° increments, thus train 12 classifiers. Which features would you use?
8. Show a test image with the detected car bounding boxes and show the estimated viewpoints by plotting an arrow in the appropriate direction.
9. **Extra credit:** Given the ground plane, estimated depth, and the location of the car's bounding box, how would you compute a 3D bounding box around each detected car? Add the 3D bounding boxes to your plot from 5.

Project 3

This project is about analysis of news broadcast. You will be given a news video clip. Here are the tasks to solve:

- (a) Detect shots in the videos. A shot is a set of consecutive frames with a smooth camera motion.
- (b) (Manually) Annotate shot boundaries in the video. How would you evaluate how well you are detecting the shots? Compute your performance.
- (c) Detect the news company's logo.
- (f) Detect faces in the video.
- (g) Perform face tracking by correctly associating a face detection in the previous frame to a face detection in the current frame.
- (i) You will be given a dataset of female and male faces. Train a classifier that can predict whether a face is female or male. For each face track in the news video predict whether it is female or male. To do this you will take a few images of faces, compute image features, e.g., HOG, and train a male-vs-female classifier, e.g., SVM. Once trained, you will predict the gender of each face detection in the video. That is, you'll take each face detection (a crop in the image specified by the face box), compute appropriate image features, and use your classifier to predict the gender of the face. How would you decide whether a full face track is female or male?
- (j) Visualize your results: produce a video in which you show a bounding box around the detected company logo, and bounding boxes around the detected faces. Each face bounding box should have text indicated whether the face is male or female.

Project 4

This is a project for a single student (no pair work allowed). In this project you will reconstruct a car in 3D and fit a 3D CAD model to the point cloud. You will be given a sequence of images of a car.

- (b) Compute structure-from-motion to reconstruct the car in 3D.
- (c) Remove the points that are far from the camera (outliers) as well as ground pixels. How?
- (d) Fit a collection of CAD models of cars to your point cloud available here: <http://www.cs.toronto.edu/~fidler/projects/CAD.html>. You can do this using the ICP algorithm, http://en.wikipedia.org/wiki/Iterative_closest_point. There is a lot of code online, but it is desirable that you implement your own algorithm. Find the CAD car that best fits your data.
- (e) Visualize your result: plot the 3D point cloud you obtained from structure-from-motion as well as your best fitting CAD model.

Project 5

Imagine a phone app where you take a picture of a DVD and the app tells you which movie it is. In this project, the goal is to localize and recognize a DVD cover in an input image given a database of a large set of DVD covers. You will need to implement the following paper:

Nister, Stewenius, *Scalable Recognition with a Vocabulary Tree*, CVPR 2006, http://www.vis.uky.edu/~stewe/publications/nister_stewenius_cvpr2006.pdf

You will be given a set of images of DVD covers. You will also be given a set of test images, each containing one DVD in a non-frontal viewpoint. Here are the tasks:

1. Implement homography estimation with RANSAC to match one DVD cover image to a test image.
2. Implement the efficient retrieval approach by Nister and Stewenius.
3. For a test image, retrieve top 10 matches (DVD cover images from the database) returned via the implemented approach. Compute a homography with your implementation from (1) for each retrieved DVD cover image and the test image.
4. Find the DVD cover image from (3) with the highest number of inliers. Plot the test image with the localized DVD cover and as well as the best retrieved DVD cover.

Useful Code and Techniques

Shot Detection. A simple way of detecting shot boundaries is to look at differences in color histograms between two consecutive frames. An even better way is to look at Displaced Frame Distances, described in Section 2.1 of this paper:

Makarand Tapaswi, Martin Baeuml and Rainer Stiefelbogen, “*Knock! Knock! Who is it*” *Probabilistic Person Identification in TV Series*, CVPR 2012, <https://cvhci.anthropomatik.kit.edu/~mtapaswi/papers/CVPR2012.pdf>

More options are discussed in this paper:
Y. Yusoff, W. Christmas, and J. Kittler, *A Study on Automatic Shot Change Detection. Multimedia Applications and Services*, 1998, <http://www.cs.utoronto.ca/~fidler/slides/CSC420/papers/shots.pdf>.

Face Detection. A few options:

- The first paper that did great on faces was Viola-Jones face detector:
Paul Viola and Michael Jones, Rapid object detection using a boosted cascade of simple features, CVPR, 2001, <https://www.cs.cmu.edu/~efros/courses/LBMV07/Papers/viola-cvpr-01.pdf> There is lots of code online.
- Paper:
P. Felzenszwalb, R. Girshick, D. McAllester, D. Ramanan, *Object Detection with Discriminatively Trained Part Based Models* IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 32, No. 9, Sep. 2010 <http://cs.brown.edu/~pff/papers/lsvm-pami.pdf>
Code: <http://www.cs.berkeley.edu/~rbg/latent/>
with a trained model called `dpm_baseline.mat` available here:
https://bitbucket.org/rodrigob/doppia/src/tip/data/trained_models/face_detection/?at=preparing_v2. This detector may work a little better. You can also check a detailed analysis and code for different face detectors here: http://markusmathias.bitbucket.org/2014_eccv_face_detection/
- Paper:
X. Zhu, D. Ramanan. *Face detection, pose estimation and landmark localization in the wild*, CVPR 2012, <http://www.ics.uci.edu/~xzhu/face/>.
This detector however also gives you the keypoints for the facial landmarks.

Tracking Here are some options (papers), choose the one it suits you best:

- Tracking via dynamic programming: <https://engineering.purdue.edu/~qobi/papers/acs2012b.pdf>
- Tracking via Hungarian method: <http://www.cvlibs.net/publications/Geiger2014PAMI.pdf>, first 3 paragraphs of Section 4.1
- Tracking via Kalman filter: lots of tutorials online

Available code:

- Check <http://www.ics.uci.edu/~dramanan/> for paper (and code): ‘H. Pirsiavash, D. Ramanan, C. Fowlkes. “Globally-Optimal Greedy Algorithms for Tracking a Variable Number of Objects”, Computer Vision and Pattern Recognition (CVPR), 2011
- <http://www.cvlibs.net/software/trackbydet/>
- <http://research.milanton.de/dctracking/>

Stereo For autonomous driving, it’s best to check the stereo challenge of the KITTI dataset (<http://www.cvlibs.net/datasets/kitti/>). It has links to code and gives you running time of each algorithm. Some options:

- <http://userpage.fu-berlin.de/spangenb/>

Flow For autonomous driving, it’s best to check the stereo challenge of the KITTI dataset (<http://www.cvlibs.net/datasets/kitti/>). It has links to code and gives you running time of each algorithm.

- <http://people.csail.mit.edu/celiu/OpticalFlow/>
- Matlab has some functions to compute flow: <http://www.mathworks.com/help/vision/ref/opticalflow.html>
- OpenCV has flow and trackers, and other useful stuff: <http://opencv.org/documentation.html>

Structure-from-motion There is lots of SFM code on the web, choose one that is closest to the programming language you are using. We mention two here

- Bundler: <http://www.cs.cornell.edu/~snavely/bundler/>
- While this code doesn’t work as well, it is very good for learning about SFM: <http://vision.princeton.edu/courses/SFMedu/>

Other For features, super pixels, classifiers, etc, look under Resources of class webpage: <http://www.cs.utoronto.ca/~fidler/CSC420.html>