

CMSC 426, Computer Vision

Project 1: Shake my Boundary!

Prof. Yiannis Aloimonos,
Jack Rasiel and Kaan Elgin

February 8, 2018

1 Due Dates

This project is divided into two "milestones", one week apart:

- **MILESTONE 1 DUE DATE:** 11:59:59PM on Thursday, **Feb 15th** 2018
 - Steps 4.1 (Filter Bank) and 4.2 (Half-disk Masks).
- **MILESTONE 2 DUE DATE:** 11:59:59PM on Thursday, **Feb 22nd** 2018
 - Steps 4.3 through 4.10.

Students who reach both milestones on time will receive 10 points of extra credit on this assignment (and the TAs' undying admiration). Each milestone will be submitted separately on ELMS; for the full submission guidelines **see section 7**.

2 Introduction

Boundary detection is an important, well-studied computer vision problem. Clearly it would be nice to have algorithms which know where one object stops and another starts. But boundary detection from a single image is fundamentally difficult. Determining boundaries could require object-specific reasoning, arguably making the task "vision hard".

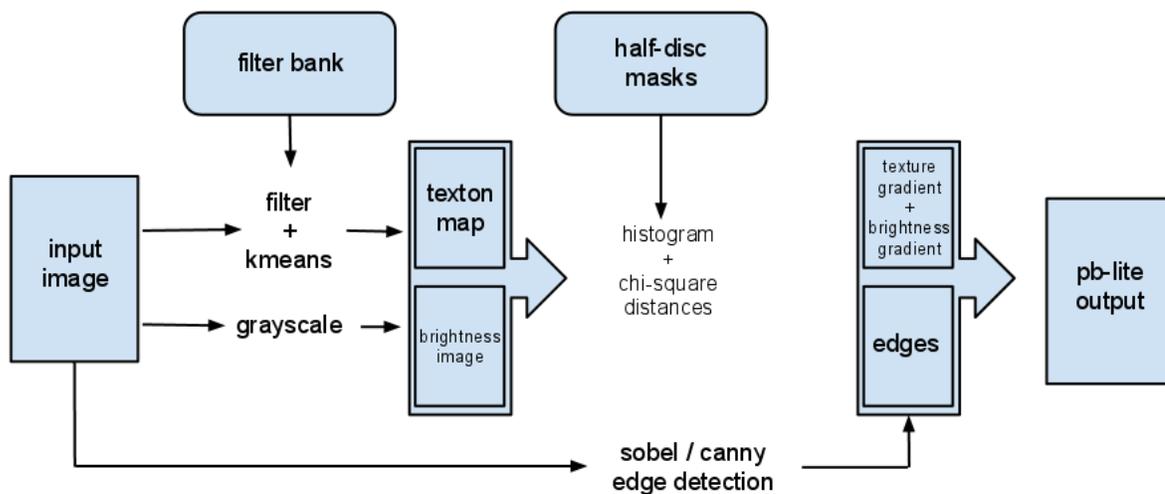
Classical edge detection algorithms, including the Canny and Sobel baselines we will compare against, look for intensity discontinuities. The more recent pb (probability of boundary) boundary detectors significantly outperform these classical methods by considering texture and color gradients in addition to intensity. Qualitatively, much of this performance jump comes from the ability of the pb algorithm to suppress false positives that the classical methods produce in textured regions.

In this project, you will develop a simplified version of `pb`, which finds boundaries by examining brightness, color, and texture information across multiple scales. The output of your algorithm will be a per-pixel probability of boundary. Several papers from Berkeley describe their algorithms and how their methods evolved over time. Their source code is also available for reference (don't use it). Here we investigate a simplified version of the recent work from Ref. [1]. Our simplified boundary detector will still significantly outperform the well regarded Canny edge detector. Evaluation is carried out against human annotations (ground truth) from a subset of the Berkeley Segmentation Data Set 500 (BSDS500).

3 Overview

The main steps in Ref. [1] are:

- Low-level feature extraction: (1) brightness, (2) color, and (3) textures
- Multiscale cue combination with non-maximum suppression
- Spectral clustering



The focus of `pb-lite` will be on the representation of brightness, color and texture gradients (bullet point 1). This is covered (briefly) in Section 3.1 of Ref. [1]. We will trivialize the multi-scale cue combination. We will get a simple form of non-maximum suppression by combining our `pb-lite` estimates with a classical edge detector. We will skip spectral clustering. You will implement the following pipeline:

The major steps are to pre-define:

- a filter bank of multiple scales and orientations.
- half-disc masks of multiple scales and orientations.

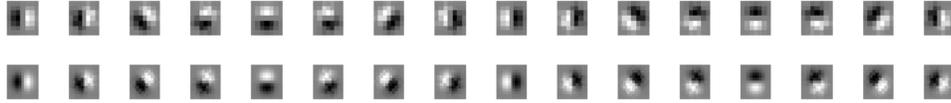


Figure 1: Oriented First Derivative Gaussian Filter Bank

And then for every image:

- create a texton map by filtering and clustering the responses with kmeans.
- compute per pixel texture gradient ($\tau\mathbf{g}$) and brightness gradient (\mathbf{bg}) by comparing local half-disc distributions.
- output a per-pixel boundary score based on the magnitude of these gradients combined with a baseline edge detector (Canny or Sobel).

Finally the output for all of the evaluate using the the Berkeley Segmentation Data Set 500 (BSDS500) (code will be given soon).

4 Texture Representation

The key distinguishing element between `pb-lite` and classical edge detection is the ability to measure texture gradients in addition to intensity gradients. The texture gradient at any pixel should summarize how quickly the texture is changing at that point. The key technical challenges are how to represent local texture distributions and how to measure distances between them. As in `pb`, we will represent texture as a local distribution of textons, where textons are discrete texture elements generated by clustering filter bank responses (more on this later). Texture and brightness gradients will be measured by comparing the distributions of textons or brightnesses within half-discs centered on a pixel of interest.

4.1 Filter Bank - 25Pts

Filtering is at the heart of building the low level features we are interested in. We will use filtering both to measure texture properties and to aggregate regional texture and brightness distributions. A simple but effective filter bank is a collection of oriented Derivative of Gaussian filters. These filters can be created by convolving a simple Sobel filter and a Gaussian kernel and then rotating the result. Suppose we want o orientations (from 0 to 360°) and s scales, we should end up with a total of $s \times o$ filters. A sample filter bank of size 2×16 is shown in Fig. 1.

4.2 Half-disc Masks - 25Pts

The half-disc masks are simply (pairs of) binary images of half-discs. This is very important because it will allow us to compute the chi-square distances using a filtering operation,

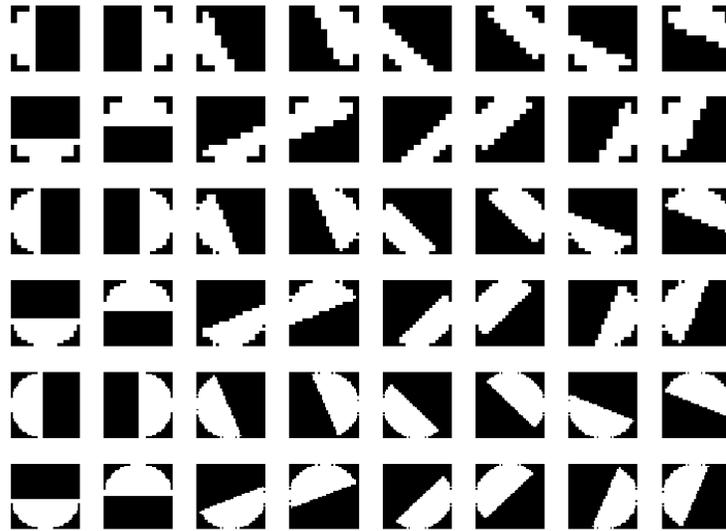


Figure 2: Half disk masks

which is much faster than looping over each pixel neighborhood and aggregating counts for histograms. Once you get the above filter bank right, forming the masks will be trivial. A sample set of masks (8 orientations, 3 scales) is shown in Fig. 2.

The filter banks and masks only need to be defined once and then they will be used on all images. You have some discretion to experiment with different filter banks and masks.

Hint: some useful MATLAB functions include: `imrotate`, `conv2`, `imfilter`, `cell`, `reshape`.

4.3 Generating a Texton Map

Filtering an input image with each element of your filter bank results in a vector of filter responses centered on each pixel. For instance, if your filter bank has 16 orientations and 2 scales, you would have 32 filter responses at each pixel. A distribution of these 32-dimensional filter responses could be thought of as encoding texture properties. We will simplify this representation by replacing each 32-dimensional vector with a discrete texton id. We will do this by clustering the filter responses at all pixels in the image in to K textons using `kmeans` (use MATLAB's `kmeans` function). Each pixel is then represented by a one dimensional, discrete cluster id instead of a vector of high-dimensional, real-valued filter responses (this process of dimensionality reduction from 32 to 1 is called "Vector Quantization"). This can be represented with a single channel image with values in the range of $[1, 2, 3, \dots, K]$. $K = 64$ seems to work well but feel free to experiment. To visualize the a texton map, you can

```
try imagesc(tmap); colormap(jet);
```

4.4 Local Texton Distributions

We will represent local texture distributions by building K-dimensional texton histograms over regions of interest. These histograms count how often each texton is observed. The regions of interest we will use are the half-disc masks previously discussed.

Analogous local histograms could be built for brightness or color, after brightness and color have been quantized in to some number of clusters.

4.5 Computing Texture Gradient (**tg**) and Brightness Gradient (**bg**) - 10Pts (5Pts each)

The local texton gradient (**tg**) and brightness gradient (**bg**) encode how much the texture and brightness distributions are changing at a pixel. We compute **tg** and **bg** by comparing the distributions in left/right half-disc pairs centered at a pixel. If the distributions are the similar, the gradient should be small. If the distributions are dissimilar, the gradient should be large. Because our half-discs span multiple scales and orientations, we will end up with a series of local gradient measurements encoding how quickly the texture or brightness distributions are changing at different scales and angles.

We will compare texton and brightness distributions with the chi-square measure. The chi-square distance is a frequently used metric for comparing two histograms. It is defined as follows:

$$\text{chi_sqr}(g,h) = .5 * \text{sum}[i=1:K ((g_i - h_i)^2 / (g_i + h_i))],$$
 where g and h are histograms with the same binning scheme, and K indices through these bins. Note that the numerator of this expression is simply the sum of squared difference between histogram elements. The denominator adds a “soft” normalization to each bin so that less frequent elements still contribute to the overall distance.

To efficiently compute **tg** and **bg**, filtering can be used to avoid nested loops over pixels. In addition, the linear nature of the formula above can be exploited. At a single orientation and scale, we can use a particular pair of masks to aggregate the counts in a histogram via a filtering operation, and compute the chi-square distance (gradient) in one loop over the bins according to the following outline:

```
chi_sqr_dist=img*0
for i=1:num_bins
    tmp = 1 where img is in bin i and 0 elsewhere
```

```

    g.i = convolve tmp with left_mask
    h.i = convolve tmp with right_mask
    update chi_sqr_dist
end

```

The above procedure should generate a 2D matrix of gradient values. Simply repeat this for all orientations and scales, you should end up with a 3D matrix of size $n \times m \times (o \times s)$, where n, m are dimensions of the image.

Hint: you might want less than 256 bins when computing \mathbf{bg} .

4.6 Rich Filter Banks - LM, S and MR - 30Pts (10Pts each)

Instead of using filter bank of Oriented First Derivative Gaussian filters, we can use richer (more complex and hopefully better) filter banks like Leung-Malik, Schmid and Maximum Response for rich feature descriptors. For detailed explanation refer to <http://www.robots.ox.ac.uk/~vgg/research/texclass/filters.html>.

Note that, the equations are not given to you for this part and hence we want you guys to figure it out on your own. However, feel free to discuss with upto 2 other students and/or come to office hours for some tips.

It's interesting to see how each filter bank affects the response on \mathbf{tg} . One can use all the filter banks at once to generate a single \mathbf{tg} or you can try to generate 3 different \mathbf{tgs} , one for each of the filter banks, i.e., \mathbf{tg}_1 for LM, \mathbf{tg}_2 for S and \mathbf{tg}_3 for MR (you will be running K-means three times here). See which performs better. You can also come up with your own way of combining the filter banks. The end goal is to get your precision-recall curve above the canny response curve.

4.7 Color Gradient - 10Pts

You need to use color information in L^*a^*b color space to use Color Gradient in addition to Texture Gradient and Brightness Gradient as used in Ref. You can bin the values of L and then compute chi-square distance to get brightness gradient \mathbf{bg} . Similarly we use a^* and b^* channels to compute color gradients $\mathbf{cg}(a^*)$ and $\mathbf{cg}(b^*)$. Now you can combine these two to get one unified color gradient $\mathbf{cg}(ab^*) = \mathbf{cg}(a^*) + \mathbf{cg}(b^*)$ check Ref. [1]

4.8 Sobel and Canny baseline

You need to run `canny_pb` and `sobel_pb` functions to generate canny and sobel baseline edges which we further use for `pb` edges

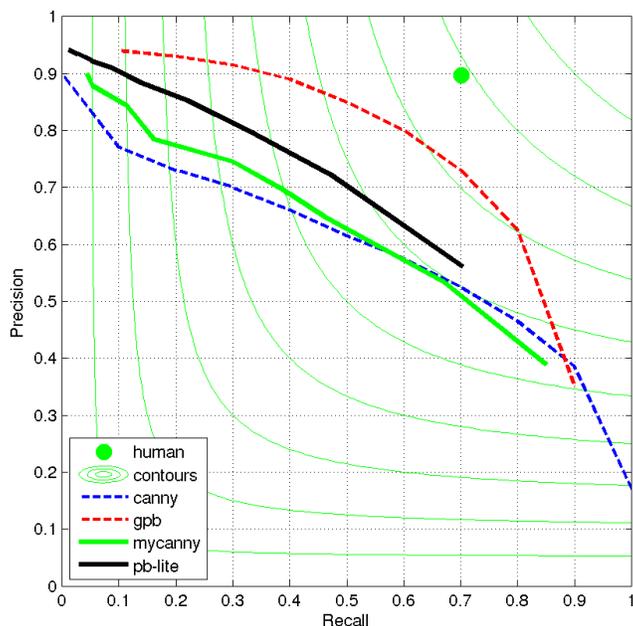


Figure 3: Precision Recall curve from Ref. [1]

4.9 Output pb-lite

The final step is to combine information from the features with a baseline method (based on Sobel or Canny edge detection) using a simple equation

$$\text{PbEdges} = (\text{tg} + \text{bg}) .* (\text{cannyPb} + \text{sobelPb})$$

The magnitude of the features represents the strength of boundaries, hence, a simple mean of the feature vector at location i should be somewhat proportional to pb . Of course, fancier ways to combine the features can be explored for better performance and extra credit. As a starting point, you can simply use an element-wise product of the baseline output and the mean feature strength to form the final pb value, this should work reasonably well.

4.10 Evaluating Boundary Detection

The goal of this project is to have a boundary detector that beats the baseline method provided, which is based on Sobel/Canny edge detection. The evaluation is based human annotation (ground truth), collected as part of the BSDS500 dataset. A detailed description of the evaluation scheme is presented in Ref. [1]. A good performance measure is the F score of the precision-recall curve. The stencil code does automatic evaluation for you. A sample precision-recall curve is shown in Fig. 3.

(dotted lines are copied from figure 17 in Ref. [1], solid lines come from actual implementations in MATLAB)

The above diagram is generated using all of the 200 test images in the dataset. If you are using a subset of the test set, you might get better/worse performance because some images are “easier” or “harder”. But if you do better than gpb (dotted red line), something is either wrong or you can publish a top paper.

5 Extra Credit

5.1 Dynamic weights - 20Pts

In this boundary detection algorithm we fuse ‘Pb-lite output’ with Canny or Sobel baseline edges. Instead of weighing them equally i.e.,

$$\text{PbEdges} = (\text{tg} + \text{bg}).*(\text{cannyPb} + \text{sobelPb})$$

You have to come up with different weights for Canny and Edge depending upon image.

5.2 Color Segmentation - 40Pts

Come with some approach to use Color prior to detect Object boundaries which can help further boost the quality of boundary detection.

5.3 Do Something Cool! - upto 20Pts

Add any other heuristic which will help improve the over all boundary detection.

6 Starter Code

The full pipeline has been given in the `Starter.m` and you need to fill in the code in each section. Use the images from `../TestImages` folder to test your algorithm. At the end of each section save the images as mentioned in the comments.

- Texton Maps are saved to `../Images/TextonMap/TextonMap_ImageName.png` (ImageName is the file name of the image)
- Texture Gradients to `../Images/tg/tg_ImageName.png`
- Brightness Maps to `../Images/BrightnessMap/BrightnessMap_ImageName.png`
- Brightness Gradients to `../Images/bg/bg_ImageName.png`
- Sobel Pb outputs to `../Images/SobelPb/SobelPb_ImageName.png`
- Canny Pb outputs to `../Images/CannyPb/CannyPb_ImageName.png`
- Pb Lite outputs to `../Images/PbLite/PbLite_ImageName.png`
- Plot PR curves and save it as `../PR_Curve.png` (code will be given soon)

You can refer to some sample reports (which might help you understand the pipeline better) at <http://cs.brown.edu/courses/cs143/2011/results/proj2/>.

7 Submission Guidelines

7.1 Milestone 1 (Due Thursday, Feb 15th)

Your submission on Canvas must be a zip file, following the naming convention “**YourDirectoryID_proj1_milestone1.zip**”. For example, jrasiel_proj1_milestone1.zip. The file must have the following directory structure, based on the starter files:

```
YourDirectoryID_proj1_milestone1.zip
├── Code/
│   ├── Starter.m
│   └── (Any dependencies of Starter.m)
├── Images/
└── report.pdf
```

Note: can also include the other files and folders provided in the project starter files.

When run, **Starter.m** must generate the filter banks and half-disc masks, and save them as the files `Images/GaussianFB_ImageName.png` and `Images/HDMasks_ImageName.png` (as described in the comments in `Starter.m`). **Your code must be able to run on our machines, or you won't receive credit for it.**

report.pdf should include example output of your code, along with any implementation details and observations. **We will deduct 2.5 points for not including these images.**

7.2 Milestone 2 (Due Thursday, Feb 22nd)

Your submission on Canvas must be a zip file, following the naming convention “**YourDirectoryID_proj1_milestone2.zip**”. It must have the same directory structure as the first milestone. When run, **Starter.m** must generate and save all of the images described in Section 6. Again, **your code must be able to run on our machines, or you won't receive credit for it.**

report.pdf should include example output of your code, along with any implementation details and observations. **We will deduct 2.5 points for not including these images.**

If your code does not comply with the above guidelines, you'll be given **ZERO** credit.

Also make a presentation if you want to present it in the class **Note: Every student should present AT LEAST once in the class (you can volunteer to present for more than one project) and can choose to do for any of the Projects from 1 to 4. Good presentations will receive cookie points.**

8 Allowed Matlab functions

`imfilter`, `conv2`, `imrotate`, `im2double`, `rgb2gray`, `rgb2lab`, `rgb2ycbcr`, `rgb2hsv`, `kmeans` and all other plotting and matrix operation/manipulation functions are allowed. `fspecial`, `imgaussfilt` are NOT ALLOWED!

9 Collaboration Policy

You are restricted to discuss the ideas with at most two other people. But the code you turn-in should be your own and if you **DO USE** (try not to and it is not permitted) other external codes/codes from other students - do cite them. For other honor codes refer to the CMSC426 Spring 2018 website.

DON'T FORGET TO HAVE FUN AND PLAY AROUND WITH IMAGES!.

Acknowledgements

This fun project was inspired from 'Introduction to Computer Vision' (CS 143) course of Brown University (<http://cs.brown.edu/courses/cs143/2011/proj2/>).

References

- [1] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2011.