Due Date 25/Oct. Fall 2016. Homework 2

Prof. Alan Yuille

October 12, 2016

Due on 25/Oct/2016. Email pdf file to reisinger@cogsci.jhu.edu by 11:59:59 PM on the due date. Format file name as *firstname*-*lastname*-`hw2.pdf`.

Question 1. Decision Theory

1. Briefly describe how to make decisions with uncertain observations by using Bayes Decision Theory. What are priors, likelihood functions, and loss functions? What are the formulas for Bayes Risk and the Bayes Decision Rule? What are maximum likelihood (ML) estimation or maximum a posteriori (MAP) estimation? When does the Bayes Decision Rule reduce to them?

2. What are false positives and false negatives? Give a formula for these when the likelihood functions are Gaussians of one variable with the same variance σ^2, but different means μ_T, μ_D. Express the false positives and false negatives in terms of the error function (integrals of Gaussians). How do they vary with the threshold? Give a formula for how the threshold depends on the prior and the loss function?

3. How can Bayes Decision Theory be applied to edge detection? Why is a first order derivative filter good for edge detection? And why is the second order

derivative filter less good? Given the hierarchical nature of visual processing, and the difficulty of edge detection, what is a good loss function for edge detection? What should be the trade-off between false positives and false negatives?

Question 2.

1. Recall that Fourier theory represents an image by a linear combination of basis functions, namely sinusoids, where the coefficients of the basis functions can be obtained by the Fourier inverse transform. Write down the formulas for the Fourier transform and inverse Fourier transform. What property of the basis functions ensures that there is a simple expression for the coefficients?

2. What is sparse coding? How do the basis functions learnt using sparse coding differ from those learnt using methods like Fourier theory? The sparsity penalty encourages many of the coefficients of basis functions to be zero. How does the degree of sparsity depend on the parameter λ which penalizes the sum of the magnitude of the coefficients. Motivate this by studying the one-dimensional case with the function $f(x;a) = (x-a)^2 + \lambda|x|$ and the rule $\hat{x} = \arg\min f(x;a)$. Show that, for some values of a, the answer is sparse (i.e. $\hat{x} = 0$).

Question 3. Experimental Section: Edge Detection and Sparse Coding

This consists of two separate projects as described in this IPython notebook: `http://nbviewer.jupyter.org/github/drew-reisinger/AS.50.375HWFall16/blob/master/HW2Intro.ipynb`

Project 1: Statistical Edge Detection. Apply Bayes Decision Theory to edge detection.

Project 2: Learn a Sparse Code for Natural Images. Use an unsupervised sparse coding technique to learn receptive fields from naturally occurring image statistics.

The Sparse Coding project strongly encourages people to team up in groups of up to three people with at least person being technically proficient, and teams should have a mix of engineering and non-engineering students as much as possible. The code for this project is unstable on Windows, so it is best that each team uses OS X or Linux.