6.869 Advances in Computer Vision

Exam 2

Submission: Submission deadline is Wednesday, April 28, before midnight.

Requirements: The exam has 2 questions (each is worth the same points). The exam is open book and open web, but you need to cite all the materials (books, web pages, etc) that you referred to. Take-home exams may not be discussed. No collaboration is allowed. You should include the following two statements on the front page of your exam write-up, which you need to sign physically or electronically:

1. “Below is a listing of the sources I used for this open book, open web test:”

2. “I affirm that I did all the work for this exam myself, and did not receive help from anyone else. I also did not give help on this exam to anyone else in the class.

Signature:

Date:

1. Image deblurring.

Consider this 1-d blind deconvolution problem: we observe the "image": **y** = [1 1 2 1 1]. We seek to explain the observation as the product of a latent image **x** convolved with a blur kernel, **b**.

Listed below are 3 candidate solutions for **x** and **b**, which all explain the data **y** equally well (apart from convolution edge effects, which we'll ignore for this problem). For this problem, we also assume that these are the only 3 solutions for **y = x** ** **b** (where ** denotes convolution); the probability of all other solutions is zero.

solution 1: blur kernel: **b** = [1 1 1] / 3 latent image: **x** = [0 3 0 3 0]
solution 2: blur kernel: **b** = [1 0 1] / 2 latent image: **x** = [2 6 4 6 2] / 3
solution 3: blur kernel: **b** = [0 1 0] latent image: **x** = [1 1 2 1 1]

We assume the following prior probabilities for latent images and blur kernels:

- All 3 blur kernels b are equally likely.
- The log prior probability of a latent images x is minus the sum of the absolute value of its derivatives:

$$p(x) = \exp\left(-\sum_{n=1}^{4}\left|x(n+1) - x(n)\right|\right)$$

(a) What is the Bayesian MAP solution? (write the form of the posterior)

(b) Does this Bayesian solution favor the explanation of sharp kernel **b** applied to a blurry latent image **x** or a blurry kernel b applied to a sharp latent image **x**?

2. Boosted classifier

The goal of this question is to train a boosted classifier given a set of training examples. The training set consists of 6 examples. For each example we have computed 3 binary features. The features are in **x** and the labels are in **y**.

Features:

x = [+1 +1 -1 +1 -1 +1;
-1 +1 -1 +1 +1 -1;
+1 -1 -1 +1 -1 -1]

Labels:

y = [-1 +1 +1 +1 +1 -1]

We are going to train a classifier F(x) with the form:

F(x) = f1(x) + f2(x)

Note that there we will use no weights to do the combination. The final classification function G(x) will be obtained by thresholding the output of the function F(x).

G(x) = 2*(F(x)>0)-1

The functions f1(x) and f2(x) will be chosen from a predefined dictionary of weak classifiers:

h1(x) = 1
h2(x) = -1
h3(x) = x(1)
h4(x) = x(2)
h5(x) = x(2)*x(3)

We will use boosting to select the functions f1(x) and f2(x). There are many different versions of boosting. Here we will use Gentle boosting (described in class). [Note: You do not need to write general boosting code. Here we will only focus on the first two iterations of boosting. If you do the iterations by hand it is fine.]

(a) Write the cost function that boosting is optimizing, and the learning equations for each boosting iteration that you will need to solve this problem.

(b) What is the function f1(x) that boosting will select in the first iteration?

(c) What is the function f2(x) selected in the second iteration?

(d) Do both functions f1(x) and f2(x) contribute to the final classification?