CS691A: Written problem set

Erik G. Learned-Miller
Department of Computer Science
University of Massachusetts, Amherst
Amherst, MA 01003

March 11, 2010

Abstract

READ THIS!!!!!!

For the exam, you will be tested on all of the material found on the course web site in the form of handouts (there are 4 of them), and on concepts from the first 3 homeworks.

This problem set exercises some, but not all, of the skills you will need on the exam. The best strategy is to read the material in the 4 handouts, and review the problem sets, and then see if you can do this problem set without referring to those materials. This will give you a sense of how prepared you are for the exam.

Problem 1: Classification of 2-pixel images

Consider the following supervised learning problem. You are given two classes of images, with training data from each class. Each image has just two pixels, each of which is either black or white.

Training images for class A.

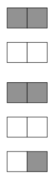

Training images for class B.

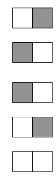

You are given one test image, and your job is to classify it. Here is the test image:

Part A.1. Pixels are dependent. (5 points total) *Assuming the pixels in each image are NOT independent random variables*, estimate the values in the following probability table, using the frequencies of the training data. Each entry in the table represents the probability of one of the possible images, given the class, i.e. $Prob(im|class)$.

Class A

	Pixel 1 black	Pixel 1 white
Pixel 2 black		
Pixel 2 white		

Class B

	Pixel 1 black	Pixel 1 white
Pixel 2 black		
Pixel 2 white		

A.2. (10 points) Assume that the *prior probability* of class A is 0.6 and that the prior probability of class B is 0.4. Use Bayes' Rule to compute an estimated posterior probability of class A and class B given your estimated likelihoods (from the table above) and the given prior probabilities. That is, you should compute $Prob(class = A|testImage)$ and $Prob(class = B|testImage)$ for the test image given on the previous page.

What class do you choose based on Bayes' Rule with no assumption of dependence among the pixels?

Part B.1. Pixels are independent. (5 points total) Now, *assuming the pixels in each image ARE independent random variables*, estimate the values in the following probability table, using the frequencies of the training data. Each entry in the table represents the probability of one of the possible images, given the class, i.e. $Prob(im|class)$.

Class A

	Pixel 1 black	Pixel 1 white
Pixel 2 black		
Pixel 2 white		

Class B

	Pixel 1 black	Pixel 1 white
Pixel 2 black		
Pixel 2 white		

B.2. (10 points) Again, assuming that the *prior probability* of class A is 0.6 and that the prior probability of class B is 0.4. Use Bayes' Rule to compute an estimated posterior probability of class A and class B given your estimated likelihoods (from the table above) and the given prior probabilities. That is, you should compute $Prob(class = A|testImage)$ and $Prob(class = B|testImage)$ for the test image given on the previous page.

What class do you choose based on Bayes' Rule when the pixels are assumed to be independent?

Part C.1. (2 points) Billy-Bob wants to try to decide if the pixels in the training images in Part A are really independent or not. He knows that there is a function, which we shall call FOO, which is zero when two variables are independent, and non-zero when the variables are not independent. What is the real name of this function of two random variables?

C.2. (10 points) Now Billy-Bob wants to estimate the function $FOO(pixel1, pixel2)$ using all of the data in the training images. As you may recall, the formula for $FOO(X, Y)$ is:

$$FOO(X, Y) = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} Prob(x, y) \log_2 \frac{Prob(x, y)}{Prob(x)Prob(y)},$$

where x and y here will represent pixel 1 and pixel 2 in the 2-pixel images.

Your job is to use the training data and the estimates of probabilities you've made so far to estimate this function. For the purposes of this exercise, assume that the frequencies you estimate from the training data are exact probabilities. In other words, assume you're estimated probabilities are exactly correct.

NOTE: Beware that the probabilities you estimated in Part A.1, for example, are conditioned on class, and that you want probabilities that are NOT conditioned on class. In other words, in Part A.1, you computed probabilities of the form $Prob(x, y|class)$, but what you need for this problem is probabilities of the form $Prob(x, y)$, so you will have to derive these from the data and from the class priors given before.

You may not use a calculator, but you may leave your final expression in an unsimplified form, include the presence of logs. Please show work clearly, as there will be partial credit on this problem.

Part D.1 (2 points each) For this question, the answer can be either "unknown", "class A", or "class B".

What would the nearest neighbor classifier give for the test image in Part A?

What would a k-nearest neighbor classifier give with k = 3?

With k= 5?

Problem 2. Alignment

Discuss alignment in computer vision. Try to address these questions. What is alignment? Why is alignment an important subproblem in computer vision? Why is it difficult? What are some techniques for doing alignment? What are their pros and cons. Two paragraphs should suffice.

Problem 3: Pinhole cameras

Part A (10 points) You are going to a party at Tiger Woods' house, and you have been hired by the National Enquirer to get a picture of him, for which you will be paid a million dollars. You know that you will be checked for cameras at the door, and hence will not be able to take a regular camera to the party. However, you can hide a piece of photographic paper taped to your back. The photographic paper, when the protective coating is removed, will record a picture if exposed to light. Also, you know that Tiger keeps lots of Nike shoe boxes in his closet, one of which you can obtain when no one is looking. Your scheme is to build a pinhole camera with a pinhole of radius one millimeter and take a single picture of Woods.

Let the dimensions of the shoe box be $30cm$ by $30cm$ by $60cm$. You will put the pinhole in the middle of one of the ends of the shoe box and put the piece of photographic paper on the other end of the shoe box.

If Tiger is exactly 2 meters tall (a bit more than in reality), and he is standing 10 meters away when you manage to take your pinhole picture, how tall will the picture of Tiger be on your film?

To answer this question, draw a picture of the shoe box, of Tiger, and of the image of tiger on the end of the shoe box.

Part B (5 points) After successfully building the pinhole camera, you take a single picture of Tiger while he is showing off with a golf club. You expose the film for exactly 1 second. While most of the picture looks pretty good, the parts of Tiger that were moving (his arm and the golf club and the golf ball) are very blurry. Fortunately, you watched MacGyver a lot as a kid, and have an idea for improving your photo.

You know that, to a first approximation, the amount of time you need to expose film to light is inversely proportional to the amount of light your camera is capturing. Thus, by capturing twice as much light, you could take a picture of the same brightness in half the time, and hence reduce the blurring. You also know that, to a first approximation, a good lens or a pinhole captures an amount of light that is proportional to its cross sectional area.

Telling the other guests you are going to get a beer, you go to the kitchen and find the lid of a pot which is circular with radius 10cm. You fill the lid with water and put it in the freezer. Of course, Tiger has a high-end freezer that will freeze water in minutes. You take the lid out of the freezer in a few minutes, and you have created an ice lens! Being very good with duct tape and cardboard, you manage to build a camera using a large box (which was used to ship golf clubs) and your ice lens. You flip over your original piece of film paper, which, fortunately, was two-sided, so you have one more chance to take a picture!

Assuming your ice lens is perfect (it focuses all of the light coming into it to the proper place), and you are able to surreptitiously take another picture of Tiger with your giant and highly improbable device, how long do you need to expose the film to obtain an image of the same brightness as with your pinhole camera? Recall that the pinhole camera had a hole with radius of 1mm.

Problem 4: Radiation (10 points)

Suppose you have just mounted K square meters of solar panels on your house, and that they are oriented so that they are perfectly perpendicular to the direction of the sun. (Assume the sun is a point light source.) Discounting the absorption and scattering of the atmosphere, compute the total number of watts of power being radiated upon your solar panels, assuming the sun is exactly M kilometers (NOT METERS!) from your solar panels. The formula for the surface area of a sphere is:

$$A = 4 \times \pi \times r^2,$$

where r is the radius of the sphere. Assume the power output of the sun is S watts.

Problem 5. Matlab

For each of the 4 occurrences of the word "COMMAND" below, write the shortest Matlab command or sequence of commands that you can which would produce the output immediately below it.

```
>> a=[1:5]
a =
1 2 3 4 5


COMMAND:


ans =
1 2 3 4 5
1 2 3 4 5
1 2 3 4 5


>> b=rand(3,3)
b =
0.8147 0.9134 0.2785
0.9058 0.6324 0.5469
0.1270 0.0975 0.9575


COMMAND:


ans =
0.8147
0.9058
0.1270
0.9134
0.6324
0.0975
0.2785
0.5469
0.9575
```

```
COMMAND:

b =
0.8147 0.9134 0.2785
0.9058 0 0
0.1270 0 0

>> a=eye(3)
a =
1 0 0
0 1 0
0 0 1

>> b=rand(3)
b =
0.9649 0.9572 0.1419
0.1576 0.4854 0.4218
0.9706 0.8003 0.9157

COMMAND:

ans =
0.9649 0 0
0 0.4854 0
0 0 0.9157
>>
```