

CS 374 Assignment #3

Estimating Probabilities

Regression

Due the week of September 26, 2005

This week you will follow up on some ideas from last week. In particular, you will explore the theory behind last week's techniques for estimating prior probabilities and likelihoods from a data set. After doing this, you will move on to the topic of regression, where the result to be predicted for each example is a real value rather than a nominal class value.

1 Estimating Probabilities

Last week you implemented Naive Bayes learning, which requires that you compute priors and likelihoods from a training set of data. Witten and Frank do a nice job of describing how to calculate priors and likelihoods, and the methods they outline certainly make a great deal of intuitive sense. What wasn't clear in your readings last week, however, is that there is sound theory to back up these techniques for estimating probabilities. This week you will explore that theory.

1.1 Reading

In many areas of machine learning, we make assumptions that our data are drawn from a distribution that obeys a particular probability model. The techniques that make this assumption are called *parametric*. Once we assume a model, our task is to estimate the parameters of the model, based on the data that we have.

Please read Sections 4.1, 4.2, and 4.5 of Alpaydin, which trace through the theory behind the estimators for a number of different probability models.

1.2 Exercises

1. On page 63, Alpaydin sketches the derivation of the rule for estimating the parameter, p , of a Bernoulli distribution. Fill in the details of the derivation. In particular, show dL/dp and solve for 0 to get Equation 4.4.
2. Now fill in the details for the derivation of Equation 4.8, the maximum likelihood estimators for a Gaussian distribution.

2 Regression

Up to this point, we have devoted our attention to the general task of classification. Given a labeled set of training examples (labeled with correct class names, that is), our goal has been to induce a classifier to predict the labels of previously unseen examples. Now we turn our attention to a related, but different, task – regression. Here, given an example, our goal will be to predict a real-valued output, rather than a discrete class.

2.1 Reading

Read Sections 4.6 and 5.8 in Alpaydin, which introduce regression.

2.2 Exercise

In the simplest case, as discussed in Section 4.6, we assume that the real-valued output is a function of a single input attribute plus some Gaussian noise. In linear regression, we assume that the function to be learned is linear. We can also assume more generally that the function to be learned is a polynomial of order k . While Alpaydin introduces the more general version, we will concentrate on the linear case.

This exercise will take you through the process of implementing linear regression, so that you can gain a better understanding of how it works and what sorts of results it can produce.

Your program should begin by generating a training set.¹ To help you do this, I will provide for you a method (implemented in Java) to generate Gaussian noise. The method can be found in

```
~andrea/shared/cs374/noise.java
```

Though it is implemented in Java, you should have no trouble converting it into the language of your choice.

There are several different ways to generate Gaussian noise. The method I have implemented for you uses a technique called the Box-Muller Method. It starts with two uniformly distributed random numbers and returns a pair of numbers from a normal distribution with a given a mean and variance. (The mean and variance are passed in as parameters.)

When you set up your data generator, be sure that you can generate data sets of different sizes. Then, since we are assuming a linear model, you should specify values for w_1 and w_0 . Next you need to generate values for the input attribute x , and, for each x , calculate the corresponding output value r by plugging x into your linear function and then adding some Gaussian noise.

You can use the noise generator to help you generate x values. (Perhaps a better name for the method would be “gaussian” rather than “noise”.) Just specify a mean and variance of your choice. The added noise should be 0-mean Gaussian noise. Again, you select the variance. (I like to use 1.0.)

Now you are ready to implement linear regression as described on pages 73-75 of Alpaydin. Print the resulting w_1 and w_0 values. How do they compare to those you defined earlier? How does the linear function that you output perform on previously unseen data? To determine this, generate a test set. (This time don't add Gaussian noise, so that you can see how your results compare to the ideal.) Compute the sum of squared errors. How well did the algorithm work?

3 Using Naive Bayes for Regression

Linear regression as described in Alpaydin is not the only technique we have for inducing a regression model (i.e., a model that predicts real values, rather than discrete classes). In fact, we can employ Naive Bayes for this purpose.

3.1 Reading

Please read the paper “Naive Bayes for Regression”, by Frank et al. (A link can be found on the Assignments page of the course website.) This paper appeared in the journal Machine Learning in 2000. You need not follow all of the details in Section 2, but you should be able to summarize the approach described there at a high level.

¹It might seem silly to generate your own data, but it is often very useful. There are times when we want a data set that will serve as a “control” in an experimental evaluation of a technique. Therefore, it is important to learn how to generate data sets that satisfy given sets of constraints.

3.2 Exercise

Summarize (and be prepared to discuss) the algorithm and results described in the paper. Does the authors' conclusion match your intuition about the performance of Naive Bayes for regression? Do the authors make a convincing case for their conclusion?