

CS 374 Assignment #1

Introduction to Machine Learning

Linear Models, Perceptrons

Due the week of February 4, 2008

Our first week of readings and exercises will cover two topics. The first goal is to give you a general introduction to machine learning. The second goal is to begin to consider a simple yet powerful class of learning models: linear models. In particular, we will consider the perceptron learning algorithm, an online algorithm for learning a linear classification model from data. We will also read an article that describes one application of machine learning – user profiling for detection of fraudulent cellular calls – that makes use of linear models, among other techniques.

Readings below are from the Alpaydin text, unless indicated otherwise. Complete references to textbooks can be found by following the “Resources” link from the course homepage.

1 Introduction to Machine Learning

Your first objectives for the week should be to get a sense of the breadth of the field of machine learning and to become familiar with machine learning terminology.

1.1 Reading

Please read

- Chapter 1: Introduction
- Chapter 2: Supervised Learning

The first chapter in Alpaydin begins with a survey of some applications that can be addressed by machine learning. Because the types of applications we are concerned with are so diverse, no single algorithm can be expected to address them all. There are many factors that affect the design of machine learning algorithms including: the type of information available from which the algorithm can learn, the desired representation of the learned knowledge, expert theory about the domain of application, etc.

In reading Chapter 1, make note of the definitions of important terms such as *supervised learning*, *clustering*, *reinforcement learning*. These, as well as other terms highlighted in the chapter, will be part of our common vocabulary throughout the course.

The second chapter in Alpaydin introduces supervised learning in more detail. It discusses both classifier learning and regression, with a focus on the former.

A classifier considers examples, which it then designates as belonging to one class or another. There are many applications where classification is important. In medicine, for example, the process of diagnosis is one of classification. That is, the goal is to consider a patient (or, more specifically, a description of a patient, including history, test results, etc.) and to assign a “label” to the patient that gives the name of their disease. In banking, the process of loan approval is a classification problem. Here the goal is to consider an applicant and to assign a “label” indicating whether that applicant is low-risk (i.e., a good candidate for a loan) or high-risk (a bad candidate).

There are many applications for which the rules of classification are either unknown or are difficult to articulate. In these situations, we can use machine learning algorithms to automatically induce the knowledge that will allow us to perform classification.

Again, as you read this chapter, please note all new terms, such as *positive examples*, *negative examples*, *generalization*, *VC dimension*, *PAC learning*, *training set*, *test set*, *validation set*, *cross-validation*. Also note that Alpaydin tends to use the word “any” to mean “all”. At times he also uses it to mean “at least one”. The meanings should be clear from context, but if not, please ask me.

If you have any interest in reading other introductions to the topic of machine learning, you might consider the first chapters in Mitchell or Duda, Hart, and Stork. The

1.2 Exercises

1. (Modified from Exercise 1.3 in Alpaydin) Assume you are given the task of building a system that can distinguish junk e-mail from non-junk. What attributes of e-mail signal to us that it is junk? How might junk be identified through a syntactic analysis?
2. (Modified from Exercise 1.5 in Alpaydin) In basket analysis, our goal is to find dependencies between pairs of items X and Y. Given a database of customer transactions, how might you find these dependencies? How would you generalize this to more than two items? What are the problems (if any) with your approach?
3. Exercise 2.9 in Alpaydin.
4. On page 24, Alpaydin takes us through an analysis to find the number of training examples needed for PAC learnability when our hypothesis is taken to be the tightest rectangle bounding the positive examples. Carefully explain his analysis, filling in details where he has omitted them.

2 Linear Models and the Perceptron Learning Algorithm

Our study of machine learning begins with algorithms for supervised learning of classifiers. A very simple type of classifier is the linear discriminant.

2.1 Reading

You were introduced to the notion of supervised learning of classifiers in Chapter 2. However, there the emphasis was on the use of rectangular regions to separate examples from different classes. Please read the following sections of Chapter 10 on Linear Discrimination:

- Section 10.1: Introduction (You can basically ignore up to the middle of page 198. Start really paying attention with the paragraph that begins, “In the discriminant-based approach...”)
- Section 10.3: Geometry of the Linear Discriminant
- Section 10.4: Pairwise Separation

It is important that you understand the basic geometry of the linear discriminant, and a number of the exercises below will guide you through this process. You might want to do the exercises that refer to this material before going on to the next set of readings.

Once you understand the geometry of the linear discriminant, you will explore one algorithm for learning a linear separator from data: the perceptron learning algorithm. While there are many algorithms for learning linear separators, the perceptron learning algorithm is interesting in that (1) it makes no assumptions about the underlying distribution of the data, and (2) it is an online algorithm

(i.e., it does not require that you provide it with all the data at once; instead, it updates a linear model incrementally, with each new example that it sees).

Please read:

- Sections 11.1-11.4 of Chapter 11.

You might find it useful to read parts of Chapter 4 in Mitchell or Section 20.5 in Russell and Norvig. In particular, pages 90-91 in Mitchell give a nice visualization of the hypothesis space, as well a clear intuition for the gradient descent rule. You should certainly be sure you understand what the sigmoid (or logistic) function looks like and have some intuition for why it makes sense to use this as the output of a linear model, rather than the simple linear combination of weights and inputs.

While linear discriminants are simple and, in theory, not very powerful, they often prove to be useful in practice. Please read

- “Combining Data Mining and Machine Learning for Effective User Profiling” by Fawcett and Provost.

It discusses the application of user profiling for detection of cellular phone fraud. Fawcett and Provost’s method combines several learning algorithms, but simple linear thresholding plays an important role in their system.

2.2 Exercises

1. In explaining the geometry of the linear discriminant, Alpaydin re-writes \mathbf{x} as $\mathbf{x}_p + r \frac{\mathbf{w}}{\|\mathbf{w}\|}$. (See bottom of page 201.) Briefly explain why \mathbf{x} can be expressed in this way.
2. Starting with the alternate way to express \mathbf{x} , follow Alpaydin’s explanation to derive r , the distance of \mathbf{x} to the hyperplane, given by Equation 10.4. Be sure to fill in the details that Alpaydin omits.
3. Now show that the distance of the hyperplane to the origin is $\frac{w_0}{\|\mathbf{w}\|}$, as given by Equation 10.5. (This should be very easy.)
4. Equation 11.8 tells us how to update the weights of a perceptron for each new training example. For this question, you will derive this update rule.

As Mitchell (pages 90-91) tells us, it is helpful to visualize the entire hypothesis space of weight vectors and their associated error values. The error values make up an error surface. As our goal is to minimize error, we want to find the point on the error surface that is minimal so that we can find the best (corresponding) weights.

Gradient descent search determines a weight vector that minimizes error by starting with an arbitrary initial weight vector, then repeatedly modifying it in small steps. At each step, the weight vector is altered in the direction that produces the steepest descent along the error surface. This process continues until a minimum is reached.

We calculate the direction of steepest descent along the error surface by computing the derivative of the error function, E , with respect to each component of \mathbf{w} . This is the *gradient* of the error function with respect to \mathbf{w} . When interpreted as a vector in weight space, the gradient specifies the direction that produces the steepest increase in E . The negative of this vector therefore gives the direction of steepest decrease.

Use the information above to derive the update rule given in Equation 11.8, assuming that error is measured by cross-entropy (see top of page 238 in Alpaydin).

5. Now use the update rule to learn appropriate weights for the following classifier learning problem:

You have been asked to learn a linear model that will help distinguish normal cell phone usage from fraudulent usage. You have decided to use a perceptron as your linear discriminator. The data you have been given are described by two boolean features, x_1 and x_2 , in addition to the class label. x_1 represents *Call time = midnight* and x_2 represents *Call origin = Bronx*. The training data are as follows:

Call time = midnight	Call origin = Bronx	Fraud
0	0	0
0	1	0
1	0	0
1	1	1

If the initial weights are: $w_0 = -0.5$, $w_1 = 1$, $w_2 = 1$, and if $\eta = 0.1$, trace through the updates for each of the training examples. You should assume that the data are augmented by an input, $x_0 = 1$. To make the calculations easier, use the threshold function – Equation 11.3 on page 234 of Alpaydin – rather than the sigmoid as output of the perceptron.

The function to be modeled in this exercise is the boolean AND function. Why are the final weights different from those given on page 240 in Alpaydin?