

Chapter 12

Large-Scale Machine Learning

Many algorithms are today classified as “machine learning.” These algorithms share, with the other algorithms studied in this book, the goal of extracting information from data. All algorithms for analysis of data are designed to produce a useful summary of the data, from which decisions are made. Among many examples, the frequent-itemset analysis that we did in Chapter 6 produces information like association rules, which can then be used for planning a sales strategy or for many other purposes.

However, algorithms called “machine learning” not only summarize our data; they are perceived as learning a model or classifier from the data, and thus discover something about data that will be seen in the future. For instance, the clustering algorithms discussed in Chapter 7 produce clusters that not only tell us something about the data being analyzed (the training set), but they allow us to classify future data into one of the clusters that result from the clustering algorithm. Thus, machine-learning enthusiasts often speak of clustering with the neologism “unsupervised learning”; the term *unsupervised* refers to the fact that the input data does not tell the clustering algorithm what the clusters should be. In *supervised* machine learning, which is the subject of this chapter, the available data includes information about the correct way to classify at least some of the data. The data classified already is called the *training set*.

In this chapter, we do not attempt to cover all the different approaches to machine learning. We concentrate on methods that are suitable for very large data and that have the potential for parallel implementation. We consider the classical “perceptron” approach to learning a data classifier, where a hyperplane that separates two classes is sought. Then, we look at more modern techniques involving support-vector machines. Similar to perceptrons, these methods look for hyperplanes that best divide the classes, so that few, if any, members of the training set lie close to the hyperplane. We end with a discussion of nearest-neighbor techniques, where data is classified according to the class(es) of their

nearest neighbors in some space.

12.1 The Machine-Learning Model

In this brief section we introduce the framework for machine-learning algorithms and give the basic definitions.

12.1.1 Training Sets

The data to which a machine-learning (often abbreviated ML) algorithm is applied is called a training set. A *training set* consists of a set of pairs (\mathbf{x}, y) , called *training examples*, where

- \mathbf{x} is a vector of values, often called a *feature vector*. Each value, or feature, can be *categorical* (values are taken from a set of discrete values, such as {red, blue, green}) or *numerical* (values are integers or real numbers).
- y is the *label*, the classification value for \mathbf{x} .

The objective of the ML process is to discover a function $y = f(\mathbf{x})$ that best predicts the value of y associated with each value of \mathbf{x} . The type of y is in principle arbitrary, but there are several common and important cases.

1. y is a real number. In this case, the ML problem is called *regression*.
2. y is a boolean value true-or-false, more commonly written as $+1$ and -1 , respectively. In this class the problem is *binary classification*.
3. y is a member of some finite set. The members of this set can be thought of as “classes,” and each member represents one class. The problem is *multiclass classification*.
4. y is a member of some potentially infinite set, for example, a parse tree for x , which is interpreted as a sentence.

12.1.2 Some Illustrative Examples

Example 12.1: Recall Fig. 7.1, repeated as Fig. 12.1, where we plotted the height and weight of dogs in three classes: Beagles, Chihuahuas, and Dachshunds. We can think of this data as a training set, provided the data includes the variety of the dog along with each height-weight pair. Each pair (\mathbf{x}, y) in the training set consists of a feature vector \mathbf{x} of the form [height, weight]. The associated label y is the variety of the dog. An example of a training-set pair would be ([5 inches, 2 pounds], Chihuahua).

An appropriate way to implement the decision function f would be to imagine two lines, shown dashed in Fig. 12.1. The horizontal line represents a height




Figure 12.1: Repeat of Fig. 7.1, indicating the heights and weights of certain dogs

of 7 inches and separates Beagles from Chihuahuas and Dachshunds. The vertical line represents a weight of 3 pounds and separates Chihuahuas from Beagles and Dachshunds. The algorithm that implements f is:

```

if (height > 7) print Beagle
else if (weight < 3) print Chihuahua
else print Dachshund;

```

Recall that the original intent of Fig. 7.1 was to cluster points without knowing which variety of dog they represented. That is, the label associated with a given height-weight vector was not available. Here, we are performing supervised learning with the same data augmented by classifications for the training data. \square

Example 12.2: As an example of supervised learning, the four points $(1, 2)$, $(2, 1)$, $(3, 4)$, and $(4, 3)$ from Fig. 11.1 (repeated here as Fig. 12.2), can be thought of as a training set, where the vectors are one-dimensional. That is, the point $(1, 2)$ can be thought of as a pair $([1], 2)$, where $[1]$ is the one-dimensional feature vector \mathbf{x} , and 2 is the associated label y ; the other points can be interpreted similarly.

Suppose we want to “learn” the linear function $f(x) = ax + b$ that best represents the points of the training set. A natural interpretation of “best” is that the RMSE of the value of $f(x)$ compared with the given value of y is




Figure 12.2: Repeat of Fig. 11.1, to be used as a training set

minimized. That is, we want to minimize

$$\sum_{x=1}^4 (ax + b - y_x)^2$$

where y_x is the y -value associated with x . This sum is

$$(a + b - 2)^2 + (2a + b - 1)^2 + (3a + b - 4)^2 + (4a + b - 3)^2$$

Simplifying, the sum is $30a^2 + 4b^2 + 20ab - 56a - 20b + 30$. If we then take the derivatives with respect to a and b and set them to 0, we get

$$\begin{aligned} 60a + 20b - 56 &= 0 \\ 20a + 8b - 20 &= 0 \end{aligned}$$

The solution to these equations is $a = 3/5$ and $b = 1$. For these values the RMSE is 3.2.

Note that the learned straight line is not the principal axis that was discovered for these points in Section 11.2.1. That axis was the line with slope 1, going through the origin, i.e., the line $y = x$. For this line, the RMSE is 4. The difference is that PCA discussed in Section 11.2.1 minimizes the sum of the squares of the lengths of the projections onto the chosen axis, which is constrained to go through the origin. Here, we are minimizing the sum of the squares of the vertical distances between the points and the line. In fact, even had we tried to learn the line through the origin with the least RMSE, we would not choose $y = x$. You can check that $y = \frac{14}{15}x$ has a lower RMSE than 4. \square

Example 12.3: A common application of machine learning involves a training set where the feature vectors \mathbf{x} are boolean-valued and of very high dimension. We shall focus on data consisting of documents, e.g., emails, Web pages, or newspaper articles. Each component represents a word in some large dictionary. We would probably eliminate stop words (very common words) from this dictionary, because these words tend not to tell us much about the subject matter of the document. Similarly, we might also restrict the dictionary to words

with high TF.IDF scores (see Section 1.3.1) so the words considered would tend to reflect the topic or substance of the document.

The training set consists of pairs, where the vector \mathbf{x} represents the presence or absence of each dictionary word in document. The label y could be +1 or -1, with +1 representing that the document (an email, e.g.) is spam. Our goal would be to train a classifier to examine future emails and decide whether or not they are spam. We shall illustrate this use of machine learning in Example 12.4.

Alternatively, y could be chosen from some finite set of topics, e.g., “sports,” “politics,” and so on. Again, \mathbf{x} could represent a document, perhaps a Web page. The goal would be to create a classifier for Web pages that assigned a topic to each. \square

12.1.3 Approaches to Machine Learning

There are many forms of ML algorithms, and we shall not cover them all here. Here are the major classes of such algorithms, each of which is distinguished by the form by which the function f is represented.

1. *Decision trees* were discussed briefly in Section 9.2.7. The form of f is a tree, and each node of the tree has a function of \mathbf{x} that determines to which child or children the search must proceed. While we saw only binary trees in Section 9.2.7, in general a decision tree can have any number of children for each node. Decision trees are suitable for binary and multiclass classification, especially when the dimension of the feature vector is not too large (large numbers of features can lead to overfitting).
2. *Perceptrons* are threshold functions applied to the components of the vector $\mathbf{x} = [x_1, x_2, \dots, x_n]$. A weight w_i is associated with the i th component, for each $i = 1, 2, \dots, n$, and there is a threshold θ . The output is +1 if

$$\sum_{i=1}^n w_i x_i \geq \theta$$

and the output is -1 otherwise. A perceptron is suitable for binary classification, even when the number of features is very large, e.g., the presence or absence of words in a document. Perceptrons are the topic of Section 12.2.

3. *Neural nets* are acyclic networks of perceptrons, with the outputs of some perceptrons used as inputs to others. These are suitable for binary or multiclass classification, since there can be several perceptrons used as output, with one or more indicating each class.
4. *Instance-based learning* uses the entire training set to represent the function f . The calculation of the label y associated with a new feature vector \mathbf{x} can involve examination of the entire training set, although usually some

preprocessing of the training set enables the computation of $f(\mathbf{x})$ to proceed efficiently. We shall consider an important kind of instance-based learning, k -nearest-neighbor, in Section 12.4. For example, 1-nearest-neighbor classifies data by giving it the same class as that of its nearest training example. There are k -nearest-neighbor algorithms that are appropriate for any kind of classification, although we shall concentrate on the case where y and the components of \mathbf{x} are real numbers.

5. *Support-vector machines* are an advance over the algorithms traditionally used to select the weights and threshold. The result is a classifier that tends to be more accurate on unseen data. We discuss support-vector machines in Section 12.3.

12.1.4 Machine-Learning Architecture

Machine-learning algorithms can be classified not only by their general algorithmic approach as we discussed in Section 12.1.3. but also by their underlying architecture – the way data is handled and the way it is used to build the model.

Training and Testing

One general issue regarding the handling of data is that there is a good reason to withhold some of the available data from the training set. The remaining data is called the *test set*. The problem addressed is that many machine-learning algorithms tend to *overfit* the data; they pick up on artifacts that occur in the training set but that are atypical of the larger population of possible data. By using the test data, and seeing how well the classifier works on that, we can tell if the classifier is overfitting the data. If so, we can restrict the machine-learning algorithm in some way. For instance, if we are constructing a decision tree, we can limit the number of levels of the tree.




Figure 12.3: The training set helps build the model, and the test set validates it

Figure 12.3 illustrates the train-and-test architecture. We assume all the data is suitable for training (i.e., the class information is attached to the data),

Generalization

We should remember that the purpose of creating a model or classifier is not to classify the training set, but to classify the data whose class we do not know. We want that data to be classified correctly, but often we have no way of knowing whether or not the model does so. If the nature of the data changes over time, for instance, if we are trying to detect spam emails, then we need to measure the performance over time, as best we can. For example, in the case of spam emails, we can note the rate of reports of spam emails that were not classified as spam.

but we separate out a small fraction of the available data as the test set. We use the remaining data to build a suitable model or classifier. Then we feed the test data to this model. Since we know the class of each element of the test data, we can tell how well the model does on the test data. If the error rate on the test data is not much worse than the error rate of the model on the training data itself, then we expect there is little, if any, overfitting, and the model can be used. On the other hand, if the classifier performs much worse on the test data than on the training data, we expect there is overfitting and need to rethink the way we construct the classifier.

There is nothing special about the selection of the test data. In fact, we can repeat the train-then-test process several times using the same data, if we divide the data into k equal-sized chunks. In turn, we let each chunk be the test data, and use the remaining $k - 1$ chunks as the training data. This training architecture is called *cross-validation*.

Batch Versus On-Line Learning

Often, as in Examples 12.1 and 12.2 we use a *batch learning* architecture. That is, the entire training set is available at the beginning of the process, and it is all used in whatever way the algorithm requires to produce a model once and for all. The alternative is *on-line learning*, where the training set arrives in a stream and, like any stream, cannot be revisited after it is processed. In on-line learning, we maintain a model at all times. As new training examples arrive, we may choose to modify the model to account for the new examples. On-line learning has the advantages that it can

1. Deal with very large training sets, because it does not access more than one training example at a time.
2. Adapt to changes in the population of training examples as time goes on. For instance, Google trains its spam-email classifier this way, adapting the classifier for spam as new kinds of spam email are sent by spammers and indicated to be spam by the recipients.

An enhancement of on-line learning, suitable in some cases, is *active learning*. Here, the classifier may receive some training examples, but it primarily receives unclassified data, which it must classify. If the classifier is unsure of the classification (e.g., the newly arrived example is very close to the boundary), then the classifier can ask for ground truth at some significant cost. For instance, it could send the example to Mechanical Turk and gather opinions of real people. In this way, examples near the boundary become training examples and can be used to modify the classifier.

Feature Selection

Sometimes, the hardest part of designing a good model or classifier is figuring out what features to use as input to the learning algorithm. Let us reconsider Example 12.3, where we suggested that we could classify emails as spam or not spam by looking at the words contained in the email. In fact, we explore in detail such a classifier in Example 12.4. As discussed in Example 12.3, it may make sense to focus on certain words and not others; e.g., we should eliminate stop words.

But we should also ask whether there is other information available that would help us make a better decision about spam. For example, spam is often generated by particular hosts, either those belonging to the spammers, or hosts that have been coopted into a “botnet” for the purpose of generating spam. Thus, including the originating host or originating email address into the feature vector describing an email might enable us to design a better classifier and lower the error rate.

Creating a Training Set

It is reasonable to ask where the label information that turns data into a training set comes from. The obvious method is to create the labels by hand, having an expert look at each feature vector and classify it properly. Recently, crowdsourcing techniques have been used to label data. For example, in many applications it is possible to use Mechanical Turk to label data. Since the “Turkers” are not necessarily reliable, it is wise to use a system that allows the question to be asked of several different people, until a clear majority is in favor of one label.

One often can find data on the Web that is implicitly labeled. For example, the Open Directory (DMOZ) has millions of pages labeled by topic. That data, used as a training set, can enable one to classify other pages or documents according to their topic, based on the frequency of word occurrence. Another approach to classifying by topic is to look at the Wikipedia page for a topic and see what pages it links to. Those pages can safely be assumed to be relevant to the given topic.

In some applications we can use the stars that people use to rate products or services on sites like Amazon or Yelp. For example, we might want to estimate the number of stars that would be assigned to reviews or tweets about a product,

even if those reviews do not have star ratings. If we use star-labeled reviews as a training set, we can deduce the words that are most commonly associated with positive and negative reviews (called *sentiment analysis*). The presence of these words in other reviews can tell us the sentiment of those reviews.

12.1.5 Exercises for Section 12.1

Exercise 12.1.1: Redo Example 12.2 for the following different forms of $f(x)$.

- (a) Require $f(x) = ax$; i.e., a straight line through the origin. Is the line $y = \frac{14}{15}x$ that we discussed in the example optimal?
- (b) Require $f(x)$ to be a quadratic, i.e., $f(x) = ax^2 + bx + c$.

12.2 Perceptrons

A *perceptron* is a linear binary classifier. Its input is a vector $\mathbf{x} = [x_1, x_2, \dots, x_d]$ with real-valued components. Associated with the perceptron is a vector of *weights* $\mathbf{w} = [w_1, w_2, \dots, w_d]$, also with real-valued components. Each perceptron has a *threshold* θ . The output of the perceptron is $+1$ if $\mathbf{w} \cdot \mathbf{x} > \theta$, and the output is -1 if $\mathbf{w} \cdot \mathbf{x} < \theta$. The special case where $\mathbf{w} \cdot \mathbf{x} = \theta$ will always be regarded as “wrong,” in the sense that we shall describe in detail when we get to Section 12.2.1.

The weight vector \mathbf{w} defines a hyperplane of dimension $d - 1$ – the set of all points \mathbf{x} such that $\mathbf{w} \cdot \mathbf{x} = \theta$, as suggested in Fig. 12.4. Points on the positive side of the hyperplane are classified $+1$ and those on the negative side are classified -1 . A perceptron classifier works only for data that is *linearly separable*, in the sense that there is some hyperplane that separates all the positive points from all the negative points. If there are many such hyperplanes, the perceptron will converge to one of them, and thus will correctly classify all the training points. If no such hyperplane exists, then the perceptron cannot converge to any particular one. In the next section, we discuss support-vector machines, which do not have this limitation; they will converge to some separator that, although not a perfect classifier, will do as well as possible under the metric to be described in Section 12.3.

12.2.1 Training a Perceptron with Zero Threshold

To train a perceptron, we examine the training set and try to find a weight vector \mathbf{w} and threshold θ such that all the feature vectors with $y = +1$ (the *positive examples*) are on the positive side of the hyperplane and all those with $y = -1$ (the *negative examples*) are on the negative side. It may or may not be possible to do so, since there is no guarantee that *any* hyperplane separates all the positive and negative examples in the training set.





Figure 12.4: A perceptron divides a space by a hyperplane into two half-spaces

We begin by assuming the threshold is 0; the simple augmentation needed to handle an unknown threshold is discussed in Section 12.2.4. The following method will converge to some hyperplane that separates the positive and negative examples, provided one exists.

1. Initialize the weight vector \mathbf{w} to all 0's.
2. Pick a *learning-rate parameter* η , which is a small, positive real number. The choice of η affects the convergence of the perceptron. If η is too small, then convergence is slow; if it is too big, then the decision boundary will “dance around” and again will converge slowly, if at all.
3. Consider each training example $t = (\mathbf{x}, y)$ in turn.
 - (a) Let $y' = \mathbf{w} \cdot \mathbf{x}$.
 - (b) If y' and y have the same sign, then do nothing; t is properly classified.
 - (c) However, if y' and y have different signs, or $y' = 0$, replace \mathbf{w} by $\mathbf{w} + \eta y \mathbf{x}$. That is, adjust \mathbf{w} slightly in the direction of \mathbf{x} .

The two-dimensional case of this transformation on \mathbf{w} is suggested in Fig. 12.5. Notice how moving \mathbf{w} in the direction of \mathbf{x} moves the hyperplane that is perpendicular to \mathbf{w} in such a direction that it makes it more likely that \mathbf{x} will be on the correct side of the hyperplane, although it does not guarantee that to be the case.

Example 12.4: Let us consider training a perceptron to recognize spam email. The training set consists of pairs (\mathbf{x}, y) where \mathbf{x} is a vector of 0's and 1's, with each component x_i corresponding to the presence ($x_i = 1$) or absence ($x_i = 0$)

Figure 12.5: A misclassified point \mathbf{x}_1 moves the vector \mathbf{w}

of a particular word in the email. The value of y is $+1$ if the email is known to be spam and -1 if it is known not to be spam. While the number of words found in the training set of emails is very large, we shall use a simplified example where there are only five words: “and,” “viagra,” “the,” “of,” and “nigeria.” Figure 12.6 gives the training set of six vectors and their corresponding classes.

| | and | viagra | the | of | nigeria | y |
|----------|-----|--------|-----|----|---------|------|
| a | 1 | 1 | 0 | 1 | 1 | $+1$ |
| b | 0 | 0 | 1 | 1 | 0 | -1 |
| c | 0 | 1 | 1 | 0 | 0 | $+1$ |
| d | 1 | 0 | 0 | 1 | 0 | -1 |
| e | 1 | 0 | 1 | 0 | 1 | $+1$ |
| f | 1 | 0 | 1 | 1 | 0 | -1 |

Figure 12.6: Training data for spam emails

In this example, we shall use learning rate $\eta = 1/2$, and we shall visit each training example once, in the order shown in Fig. 12.6. We begin with $\mathbf{w} = [0, 0, 0, 0, 0]$ and compute $\mathbf{w} \cdot \mathbf{a} = 0$. Since 0 is not positive, we move \mathbf{w} in the direction of \mathbf{a} by performing $\mathbf{w} := \mathbf{w} + (1/2)(+1)\mathbf{a}$. The new value of \mathbf{w} is thus

$$\mathbf{w} = [0, 0, 0, 0, 0] + \left[\frac{1}{2}, \frac{1}{2}, 0, \frac{1}{2}, \frac{1}{2}\right] = \left[\frac{1}{2}, \frac{1}{2}, 0, \frac{1}{2}, \frac{1}{2}\right]$$

Next, consider \mathbf{b} . $\mathbf{w} \cdot \mathbf{b} = \left[\frac{1}{2}, \frac{1}{2}, 0, \frac{1}{2}, \frac{1}{2}\right] \cdot [0, 0, 1, 1, 0] = \frac{1}{2}$. Since the associated y for \mathbf{b} is -1 , \mathbf{b} is misclassified. We thus assign

$$\mathbf{w} := \mathbf{w} + (1/2)(-1)\mathbf{b} = \left[\frac{1}{2}, \frac{1}{2}, 0, \frac{1}{2}, \frac{1}{2}\right] - \left[0, 0, \frac{1}{2}, \frac{1}{2}, 0\right] = \left[\frac{1}{2}, \frac{1}{2}, -\frac{1}{2}, 0, \frac{1}{2}\right]$$

Pragmatics of Training on Emails

When we represent emails or other large documents as training examples, we would not really want to construct the vector of 0's and 1's with a component for every word that appears even once in the collection of emails. Doing so would typically give us sparse vectors with millions of components. Rather, create a table in which all the words appearing in the emails are assigned integers $1, 2, \dots$, indicating their component. When we process an email in the training set, make a list of the components in which the vector has 1; i.e., use the standard sparse representation for the vector. If we eliminate stop words from the representation, or even eliminate words with a low TF.IDF score, then we make the vectors representing emails significantly sparser and thus compress the data even more. Only the vector \mathbf{w} needs to have all its components listed, since it will not be sparse after a small number of training examples have been processed.

Training example \mathbf{c} is next. We compute

$$\mathbf{w} \cdot \mathbf{c} = \left[\frac{1}{2}, \frac{1}{2}, -\frac{1}{2}, 0, \frac{1}{2}\right] \cdot [0, 1, 1, 0, 0] = 0$$

Since the associated y for \mathbf{c} is $+1$, \mathbf{c} is also misclassified. We thus assign

$$\mathbf{w} := \mathbf{w} + (1/2)(+1)\mathbf{c} = \left[\frac{1}{2}, \frac{1}{2}, -\frac{1}{2}, 0, \frac{1}{2}\right] + \left[0, \frac{1}{2}, \frac{1}{2}, 0, 0\right] = \left[\frac{1}{2}, 1, 0, 0, \frac{1}{2}\right]$$

Training example \mathbf{d} is next to be considered:

$$\mathbf{w} \cdot \mathbf{d} = \left[\frac{1}{2}, 1, 0, 0, \frac{1}{2}\right] \cdot [1, 0, 0, 1, 0] = 1$$

Since the associated y for \mathbf{d} is -1 , \mathbf{d} is misclassified as well. We thus assign

$$\mathbf{w} := \mathbf{w} + (1/2)(-1)\mathbf{d} = \left[\frac{1}{2}, 1, 0, 0, \frac{1}{2}\right] - \left[\frac{1}{2}, 0, 0, \frac{1}{2}, 0\right] = \left[0, 1, 0, -\frac{1}{2}, \frac{1}{2}\right]$$

For training example \mathbf{e} we compute $\mathbf{w} \cdot \mathbf{e} = \left[0, 1, 0, -\frac{1}{2}, \frac{1}{2}\right] \cdot [1, 0, 1, 0, 1] = \frac{1}{2}$. Since the associated y for \mathbf{e} is $+1$, \mathbf{e} is classified correctly, and no change to \mathbf{w} is made. Similarly, for \mathbf{f} we compute

$$\mathbf{w} \cdot \mathbf{f} = \left[0, 1, 0, -\frac{1}{2}, \frac{1}{2}\right] \cdot [1, 0, 1, 1, 0] = -\frac{1}{2}$$

so \mathbf{f} is correctly classified. If we check \mathbf{a} through \mathbf{d} , we find that this \mathbf{w} correctly classifies them as well. Thus, we have converged to a perceptron that classifies all the training set examples correctly. It also makes a certain amount of sense: it says that “viagra” and “nigeria” are indicative of spam, while “of” is indicative of nonspam. It considers “and” and “the” neutral, although we would probably prefer to give “and,” “of,” and “the” the same weight. \square

12.2.2 Convergence of Perceptrons

As we mentioned at the beginning of this section, if the data points are linearly separable, then the perceptron algorithm will converge to a separator. However, if the data is not linearly separable, then the algorithm will eventually repeat a weight vector and loop infinitely. Unfortunately, it is often hard to tell, during the running of the algorithm, which of these two cases applies. When the data is large, it is not feasible to remember all previous weight vectors to see whether we are repeating a vector, and even if we could, the period of repetition would most likely be so large that we would want to terminate the algorithm long before we repeated.

A second issue regarding termination is that even if the training data is linearly separable, the entire dataset might not be linearly separable. The consequence is that there might not be any value in running the algorithm for a very large number of rounds, in the hope of converging to a separator. We therefore need a strategy for deciding when to terminate the perceptron algorithm, assuming convergence has not occurred. Here are some common tests for termination.

1. Terminate after a fixed number of rounds.
2. Terminate when the number of misclassified training points stops changing.
3. Withhold a test set from the training data, and after each round, run the perceptron on the test data. Terminate the algorithm when the number of errors on the test set stops changing.

Another technique that will aid convergence is to lower the training rate as the number of rounds increases. For example, we could allow the training rate η to start at some initial η_0 and lower it to $\eta_0/(1 + ct)$ after the t th round, where c is some small constant.

12.2.3 The Winnow Algorithm

There are many other rules one could use to adjust weights for a perceptron. Not all possible algorithms are guaranteed to converge, even if there is a hyperplane separating positive and negative examples. One that does converge is called *Winnow*, and that rule will be described here. Winnow assumes that the feature vectors consist of 0's and 1's, and the labels are +1 or -1. Unlike the basic perceptron algorithm, which can produce positive or negative components in the weight vector \mathbf{w} , Winnow produces only positive weights.

The general Winnow Algorithm allows for a variety of parameters to be selected, and we shall only consider one simple variant. However, all variants have in common the idea that there is a positive threshold θ . If \mathbf{w} is the current weight vector, and \mathbf{x} is the feature vector in the training set that we are currently considering, we compute $\mathbf{w} \cdot \mathbf{x}$ and compare it with θ . If the dot

product is too low, and the class for \mathbf{x} is $+1$, then we have to raise the weights of \mathbf{w} in those components where \mathbf{x} has 1. We multiply these weights by a number greater than 1. The larger this number, the greater the training rate, so we want to pick a number that is not too close to 1 (or convergence will be too slow) but also not too large (or the weight vector may oscillate). Similarly, if $\mathbf{w}\cdot\mathbf{x} \geq \theta$, but the class of \mathbf{x} is -1 , then we want to lower the weights of \mathbf{w} in those components where \mathbf{x} is 1. We multiply those weights by a number greater than 0 but less than 1. Again, we want to pick a number that is not too close to 1 but also not too small, to avoid slow convergence or oscillation.

We shall give the details of the algorithm using the factors 2 and $1/2$, for the cases where we want to raise weights and lower weights, respectively. Start the Winnow Algorithm with a weight vector $\mathbf{w} = [w_1, w_2, \dots, w_d]$ all of whose components are 1, and let the threshold θ equal d , the number of dimensions of the vectors in the training examples. Let (\mathbf{x}, y) be the next training example to be considered, where $\mathbf{x} = [x_1, x_2, \dots, x_d]$.

1. If $\mathbf{w}\cdot\mathbf{x} > \theta$ and $y = +1$, or $\mathbf{w}\cdot\mathbf{x} < \theta$ and $y = -1$, then the example is correctly classified, so no change to \mathbf{w} is made.
2. If $\mathbf{w}\cdot\mathbf{x} \leq \theta$, but $y = +1$, then the weights for the components where \mathbf{x} has 1 are too low as a group. Double each of the corresponding components of \mathbf{w} . That is, if $x_i = 1$ then set $w_i := 2w_i$.
3. If $\mathbf{w}\cdot\mathbf{x} \geq \theta$, but $y = -1$, then the weights for the components where \mathbf{x} has 1 are too high as a group. Halve each of the corresponding components of \mathbf{w} . That is, if $x_i = 1$ then set $w_i := w_i/2$.

Example 12.5: Let us reconsider the training data from Fig. 12.6. Initialize $\mathbf{w} = [1, 1, 1, 1, 1]$ and let $\theta = 5$. First, consider feature vector $\mathbf{a} = [1, 1, 0, 1, 1]$. $\mathbf{w}\cdot\mathbf{a} = 4$, which is less than θ . Since the associated label for \mathbf{a} is $+1$, this example is misclassified. When a $+1$ -labeled example is misclassified, we must double all the components where the example has 1; in this case, all but the third component of \mathbf{a} is 1. Thus, the new value of \mathbf{w} is $[2, 2, 1, 2, 2]$.

Next, we consider training example $\mathbf{b} = [0, 0, 1, 1, 0]$. $\mathbf{w}\cdot\mathbf{b} = 3$, which is less than θ . However, the associated label for \mathbf{b} is -1 , so no change to \mathbf{w} is needed.

For $\mathbf{c} = [0, 1, 1, 0, 0]$ we find $\mathbf{w}\cdot\mathbf{c} = 3 < \theta$, while the associated label is $+1$. Thus, we double the components of \mathbf{w} where the corresponding components of \mathbf{c} are 1. These components are the second and third, so the new value of \mathbf{w} is $[2, 4, 2, 2, 2]$.

The next two training examples, \mathbf{d} and \mathbf{e} require no change, since they are correctly classified. However, there is a problem with $\mathbf{f} = [1, 0, 1, 1, 0]$, since $\mathbf{w}\cdot\mathbf{f} = 6 > \theta$, while the associated label for \mathbf{f} is -1 . Thus, we must divide the first, third, and fourth components of \mathbf{w} by 2, since these are the components where \mathbf{f} has 1. The new value of \mathbf{w} is $[1, 4, 1, 1, 2]$.

We still have not converged. It turns out we must consider each of the training examples \mathbf{a} through \mathbf{f} again. At the end of this process, the algorithm

| \mathbf{x} | y | $\mathbf{w} \cdot \mathbf{x}$ | OK? | and | viagra | the | of | nigeria |
|--------------|-----|-------------------------------|-----|-----|--------|-----|---------------|---------|
| | | | | 1 | 1 | 1 | 1 | 1 |
| a | +1 | 4 | no | 2 | 2 | 1 | 2 | 2 |
| b | -1 | 3 | yes | | | | | |
| c | +1 | 3 | no | 2 | 4 | 2 | 2 | 2 |
| d | -1 | 4 | yes | | | | | |
| e | +1 | 6 | yes | | | | | |
| f | -1 | 6 | no | 1 | 4 | 1 | 1 | 2 |
| a | +1 | 8 | yes | | | | | |
| b | -1 | 2 | yes | | | | | |
| c | +1 | 5 | no | 1 | 8 | 2 | 1 | 2 |
| d | -1 | 2 | yes | | | | | |
| e | +1 | 5 | no | 2 | 8 | 4 | 1 | 4 |
| f | -1 | 7 | no | 1 | 8 | 2 | $\frac{1}{2}$ | 4 |

Figure 12.7: Sequence of updates to \mathbf{w} performed by the Winnow Algorithm on the training set of Fig. 12.6

has converged to a weight vector $\mathbf{w} = [1, 8, 2, \frac{1}{2}, 4]$, which with threshold $\theta = 5$ correctly classifies all of the training examples in Fig. 12.6. The details of the twelve steps to convergence are shown in Fig. 12.7. This figure gives the associated label y and the computed dot product of \mathbf{w} and the given feature vector. The last five columns are the five components of \mathbf{w} after processing each training example. \square

12.2.4 Allowing the Threshold to Vary

Suppose now that the choice of threshold 0, as in Section 12.2.1, or threshold d , as in Section 12.2.3 is not desirable, or that we don't know the best threshold to use. At the cost of adding another dimension to the feature vectors, we can treat θ as one of the components of the weight vector \mathbf{w} . That is:

1. Replace the vector of weights $\mathbf{w} = [w_1, w_2, \dots, w_d]$ by

$$\mathbf{w}' = [w_1, w_2, \dots, w_d, \theta]$$

2. Replace every feature vector $\mathbf{x} = [x_1, x_2, \dots, x_d]$ by

$$\mathbf{x}' = [x_1, x_2, \dots, x_d, -1]$$

Then, for the new training set and weight vector, we can treat the threshold as 0 and use the algorithm of Section 12.2.1. The justification is that $\mathbf{w}' \cdot \mathbf{x}' \geq 0$ is equivalent to $\sum_{i=1}^d w_i x_i + \theta \times -1 = \mathbf{w} \cdot \mathbf{x} - \theta \geq 0$, which in turn is equivalent to $\mathbf{w} \cdot \mathbf{x} \geq \theta$. The latter is the condition for a positive response from a perceptron with threshold θ .

We can also apply the Winnow Algorithm to the modified data. Winnow requires all feature vectors to have 0's and 1's, as components. However, we can allow a -1 in the feature vector component for θ if we treat it in the manner opposite to the way we treat components that are 1. That is, if the training example is positive, and we need to increase the other weights, we instead divide the component for the threshold by 2. And if the training example is negative, and we need to decrease the other weights we multiply the threshold component by 2.

Example 12.6: Let us modify the training set of Fig. 12.6 to incorporate a sixth “word” that represents the negative $-\theta$ of the threshold. The new data is shown in Fig. 12.8.

| | and | viagra | the | of | nigeria | θ | y |
|----------|-----|--------|-----|----|---------|----------|----|
| a | 1 | 1 | 0 | 1 | 1 | -1 | +1 |
| b | 0 | 0 | 1 | 1 | 0 | -1 | -1 |
| c | 0 | 1 | 1 | 0 | 0 | -1 | +1 |
| d | 1 | 0 | 0 | 1 | 0 | -1 | -1 |
| e | 1 | 0 | 1 | 0 | 1 | -1 | +1 |
| f | 1 | 0 | 1 | 1 | 0 | -1 | -1 |

Figure 12.8: Training data for spam emails, with a sixth component representing the negative of the threshold

We begin with a weight vector \mathbf{w} with six 1's, as shown in the first line of Fig. 12.9. When we compute $\mathbf{w} \cdot \mathbf{a} = 3$, using the first feature vector \mathbf{a} , we are happy because the training example is positive, and so is the dot product. However, for the second training example, we compute $\mathbf{w} \cdot \mathbf{b} = 1$. Since the example is negative and the dot product is positive, we must adjust the weights. Since \mathbf{b} has 1's in the third and fourth components, the 1's in the corresponding components of \mathbf{w} are replaced by $1/2$. The last component, corresponding to θ , must be doubled. These adjustments give the new weight vector $[1, 1, \frac{1}{2}, \frac{1}{2}, 1, 2]$ shown in the third line of Fig. 12.9.

| x | y | $\mathbf{w} \cdot \mathbf{x}$ | OK? | and | viagra | the | of | nigeria | θ |
|----------|----|-------------------------------|-----|---------------|--------|---------------|---------------|---------|----------|
| | | | | 1 | 1 | 1 | 1 | 1 | 1 |
| a | +1 | 3 | yes | | | | | | |
| b | -1 | 1 | no | 1 | 1 | $\frac{1}{2}$ | $\frac{1}{2}$ | 1 | 2 |
| c | +1 | $-\frac{1}{2}$ | no | 1 | 2 | 1 | $\frac{1}{2}$ | 1 | 1 |
| d | -1 | $\frac{1}{2}$ | no | $\frac{1}{2}$ | 2 | 1 | $\frac{1}{4}$ | 1 | 2 |

Figure 12.9: Sequence of updates to \mathbf{w} performed by the Winnow Algorithm on the training set of Fig. 12.8

The feature vector \mathbf{c} is a positive example, but $\mathbf{w} \cdot \mathbf{c} = -\frac{1}{2}$. Thus, we must

double the second and third components of \mathbf{w} , because \mathbf{c} has 1 in the corresponding components, and we must halve the last component of \mathbf{w} , which corresponds to θ . The resulting $\mathbf{w} = [1, 2, 1, \frac{1}{2}, 1, 1]$ is shown in the fourth line of Fig. 12.9. Next, \mathbf{d} is a negative example. Since $\mathbf{w} \cdot \mathbf{d} = \frac{1}{2}$, we must again adjust weights. We halve the weights in the first and fourth components and double the last component, yielding $\mathbf{w} = [\frac{1}{2}, 2, 1, \frac{1}{4}, 1, 2]$. Now, all positive examples have a positive dot product with the weight vector, and all negative examples have a negative dot product, so there are no further changes to the weights.

The designed perceptron has a threshold of 2. It has weights 2 and 1 for “viagra” and “nigeria” and smaller weights for “and” and “of.” It also has weight 1 for “the,” which suggests that “the” is as indicative of spam as “nigeria,” something we doubt is true. Nevertheless, this perceptron does classify all examples correctly. \square

12.2.5 Multiclass Perceptrons

There are several ways in which the basic idea of the perceptron can be extended. We shall discuss transformations that enable hyperplanes to serve for more complex boundaries in the next section. Here, we look at how perceptrons can be used to classify data into many classes.

Suppose we are given a training set with labels in k different classes. Start by training a perceptron for each class; these perceptrons should each have the same threshold θ . That is, for class i treat a training example (\mathbf{x}, i) as a positive example, and all examples (\mathbf{x}, j) , where $j \neq i$, as a negative example. Suppose that the weight vector of the perceptron for class i is determined to be \mathbf{w}_i after training.

Given a new vector \mathbf{x} to classify, we compute $\mathbf{w}_i \cdot \mathbf{x}$ for all $i = 1, 2, \dots, k$. We take the class of \mathbf{x} to be the value of i for which $\mathbf{w}_i \cdot \mathbf{x}$ is the maximum, provided that value is at least θ . Otherwise, \mathbf{x} is assumed not to belong to any of the k classes.

For example, suppose we want to classify Web pages into a number of topics, such as sports, politics, medicine, and so on. We can represent Web pages by a vector with 1 for each word present in the page and 0 for words not present (of course we would only visualize the pages that way; we wouldn’t construct the vectors in reality). Each topic has certain words that tend to indicate that topic. For instance, sports pages would be full of words like “win,” “goal,” “played,” and so on. The weight vector for that topic would give higher weights to the words that characterize that topic.

A new page could be classified as belonging to the topic that gives the highest score when the dot product of the page’s vector and the weight vectors for the topics are computed. An alternative interpretation of the situation is to classify a page as belonging to all those topics for which the dot product is above some threshold (presumably a threshold higher than the θ used for training).

12.2.6 Transforming the Training Set

While a perceptron must use a linear function to separate two classes, it is always possible to transform the vectors of a training set before applying a perceptron-based algorithm to separate the classes. An example should give the basic idea.




Figure 12.10: Transforming from rectangular to polar coordinates turns this training set into one with a separating hyperplane

Example 12.7: In Fig. 12.10 we see a plot of places to visit from my home. The horizontal and vertical coordinates represent latitude and longitude of places. Some of the places have been classified into “day trips” – places close enough to visit in one day – and “excursions,” which require more than a day to visit. These are the circles and squares, respectively. Evidently, there is no straight line that separates day trips from excursions. However, if we replace the Cartesian coordinates by polar coordinates, then in the transformed space of polar coordinates, the dashed circle shown in Fig. 12.10 becomes a hyperplane. Formally, we transform the vector $\mathbf{x} = [x_1, x_2]$ into $[\sqrt{x_1^2 + x_2^2}, \arctan(x_2/x_1)]$.

In fact, we can also do dimensionality reduction of the data. The angle of the point is irrelevant, and only the radius $\sqrt{x_1^2 + x_2^2}$ matters. Thus, we can turn the point vectors into one-component vectors giving the distance of the point from the origin. Associated with the small distances will be the class label “day trip,” while the larger distances will all be associated with the label “excursion.” Training the perceptron is extremely easy. \square

12.2.7 Problems With Perceptrons

Despite the extensions discussed above, there are some limitations to the ability of perceptrons to classify some data. The biggest problem is that sometimes the data is inherently not separable by a hyperplane. An example is shown in Fig. 12.11. In this example, points of the two classes mix near the boundary so that any line through the points will have points of both classes on at least one of the sides.




Figure 12.11: A training set may not allow the existence of any separating hyperplane

One might argue that, based on the observations of Section 12.2.6 it should be possible to find some function on the points that would transform them to another space where they were linearly separable. That might be the case, but if so, it would probably be an example of *overfitting*, the situation where the classifier works very well on the training set, because it has been carefully designed to handle each training example correctly. However, because the classifier is exploiting details of the training set that do not apply to other examples that must be classified in the future, the classifier will not perform well on new data.

Another problem is illustrated in Fig. 12.12. Usually, if classes can be separated by one hyperplane, then there are many different hyperplanes that will separate the points. However, not all hyperplanes are equally good. For instance, if we choose the hyperplane that is furthest clockwise, then the point indicated by “?” will be classified as a circle, even though we intuitively see it as closer to the squares. When we meet support-vector machines in Section 12.3, we shall see that there is a way to insist that the hyperplane chosen be the one that in a sense divides the space most fairly.

Yet another problem is illustrated by Fig. 12.13. Most rules for training




Figure 12.12: Generally, more than one hyperplane can separate the classes if they can be separated at all

a perceptron stop as soon as there are no misclassified points. As a result, the chosen hyperplane will be one that just manages to classify some of the points correctly. For instance, the upper line in Fig. 12.13 has just managed to accommodate two of the squares, and the lower line has just managed to accommodate one of the circles. If either of these lines represent the final weight vector, then the weights are biased toward one of the classes. That is, they correctly classify the points in the training set, but the upper line would classify new squares that are just below it as circles, while the lower line would classify circles just above it as squares. Again, a more equitable choice of separating hyperplane will be shown in Section 12.3.

12.2.8 Parallel Implementation of Perceptrons

The training of a perceptron is an inherently sequential process. If the number of dimensions of the vectors involved is huge, then we might obtain some parallelism by computing dot products in parallel. However, as we discussed in connection with Example 12.4, high-dimensional vectors are likely to be sparse and can be represented more succinctly than would be expected from their length.

In order to get significant parallelism, we have to modify the perceptron algorithm slightly, so that many training examples are used with the same estimated weight vector \mathbf{w} . As an example, let us formulate the parallel algorithm as a MapReduce job.

The Map Function: Each Map task is given a chunk of training examples, and each Map task knows the current weight vector \mathbf{w} . The Map task computes $\mathbf{w} \cdot \mathbf{x}$ for each feature vector $\mathbf{x} = [x_1, x_2, \dots, x_k]$ in its chunk and compares that




Figure 12.13: Perceptrons converge as soon as the separating hyperplane reaches the region between classes

dot product with the label y , which is $+1$ or -1 , associated with \mathbf{x} . If the signs agree, no key-value pairs are produced for this training example. However, if the signs disagree, then for each nonzero component x_i of \mathbf{x} the key-value pair $(i, \eta y x_i)$ is produced; here, η is the learning-rate constant used to train this perceptron. Notice that $\eta y x_i$ is the increment we would like to add to the current i th component of \mathbf{w} , and if $x_i = 0$, then there is no need to produce a key-value pair. However, in the interests of parallelism, we defer that change until we can accumulate many changes in the Reduce phase.

The Reduce Function: For each key i , the Reduce task that handles key i adds all the associated increments and then adds that sum to the i th component of \mathbf{w} .

Probably, these changes will not be enough to train the perceptron. If any changes to \mathbf{w} occur, then we need to start a new MapReduce job that does the same thing, perhaps with different chunks from the training set. However, even if the entire training set was used on the first round, it can be used again, since its effect on \mathbf{w} will be different if \mathbf{w} has changed.

12.2.9 Exercises for Section 12.2

Exercise 12.2.1: Modify the training set of Fig. 12.6 so that example **b** also includes the word “nigeria” (yet remains a negative example – perhaps someone telling about their trip to Nigeria). Find a weight vector that separates the positive and negative examples, using:

- (a) The basic training method of Section 12.2.1.
- (b) The Winnow method of Section 12.2.3.

Perceptrons on Streaming Data

While we have viewed the training set as stored data, available for repeated use on any number of passes, Perceptrons can also be used in a stream setting. That is, we may suppose there is an infinite sequence of training examples, but that each may be used only once. Detecting email spam is a good example of a training stream. Users report spam emails and also report emails that were classified as spam but are not. Each email, as it arrives, is treated as a training example, and modifies the current weight vector, presumably by a very small amount.

If the training set is a stream, we never really converge, and in fact the data points may well not be linearly separable. However, at all times, we have an approximation to the best possible separator. Moreover, if the examples in the stream evolve over time, as would be the case for email spam, then we have an approximation that values recent examples more than examples from the distant past, much like the exponentially decaying windows technique from Section 4.7.

- (c) The basic method with a variable threshold, as suggested in Section 12.2.4.
- (d) The Winnow method with a variable threshold, as suggested in Section 12.2.4.

! Exercise 12.2.2: For the following training set:

$$\begin{array}{ll} ([1, 2], +1) & ([2, 1], +1) \\ ([2, 3], -1) & ([3, 2], -1) \end{array}$$

describe all the vectors \mathbf{w} and thresholds θ such that the hyperplane (really a line) defined by $\mathbf{w} \cdot \mathbf{x} - \theta = 0$ separates the points correctly.

! Exercise 12.2.3: Suppose the following four examples constitute a training set:

$$\begin{array}{ll} ([1, 2], -1) & ([2, 3], +1) \\ ([2, 1], +1) & ([3, 2], -1) \end{array}$$

- (a) What happens when you attempt to train a perceptron to classify these points using 0 as the threshold?

!! (b) Is it possible to change the threshold and obtain a perceptron that correctly classifies these points?

- (c) Suggest a transformation using quadratic polynomials that will transform these points so they become linearly separable.

12.3 Support-Vector Machines

We can view a *support-vector machine*, or SVM, as an improvement on the perceptron that is designed to address the problems mentioned in Section 12.2.7. An SVM selects one particular hyperplane that not only separates the points in the two classes, but does so in a way that maximizes the *margin* – the distance between the hyperplane and the closest points of the training set.

12.3.1 The Mechanics of an SVM

The goal of an SVM is to select a hyperplane $\mathbf{w} \cdot \mathbf{x} + b = 0$ ¹ that maximizes the distance γ between the hyperplane and any point of the training set. The idea is suggested by Fig. 12.14. There, we see the points of two classes and a hyperplane dividing them.




Figure 12.14: An SVM selects the hyperplane with the greatest possible margin γ between the hyperplane and the training points

Intuitively, we are more certain of the class of points that are far from the separating hyperplane than we are of points near to that hyperplane. Thus, it is desirable that all the training points be as far from the hyperplane as possible (but on the correct side of that hyperplane, of course). An added advantage of choosing the separating hyperplane to have as large a margin as possible is that there may be points closer to the hyperplane in the full data set but not in the training set. If so, we have a better chance that these points will be classified properly than if we chose a hyperplane that separated the training points but allowed some points to be very close to the hyperplane itself. In that case, there is a fair chance that a new point that was near a training point that

¹Constant b in this formulation of a hyperplane is the same as the negative of the threshold θ in our treatment of perceptrons in Section 12.2.

was also near the hyperplane would be misclassified. This issue was discussed in Section 12.2.7 in connection with Fig. 12.13.

We also see in Fig. 12.14 two parallel hyperplanes at distance γ from the central hyperplane $\mathbf{w} \cdot \mathbf{x} + b = 0$, and these each touch one or more of the *support vectors*. The latter are the points that actually constrain the dividing hyperplane, in the sense that they are all at distance γ from the hyperplane. In most cases, a d -dimensional set of points has $d + 1$ support vectors, as is the case in Fig. 12.14. However, there can be more support vectors if too many points happen to lie on the parallel hyperplanes. We shall see an example based on the points of Fig. 11.1, where it turns out that all four points are support vectors, even though two-dimensional data normally has three.

A tentative statement of our goal is:

- Given a training set $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)$, maximize γ (by varying \mathbf{w} and b) subject to the constraint that for all $i = 1, 2, \dots, n$,

$$y_i(\mathbf{w} \cdot \mathbf{x}_i + b) \geq \gamma$$

Notice that y_i , which must be $+1$ or -1 , determines which side of the hyperplane the point \mathbf{x}_i must be on, so the \geq relationship to γ is always correct. However, it may be easier to express this condition as two cases: if $y = +1$, then $\mathbf{w} \cdot \mathbf{x} + b \geq \gamma$, and if $y = -1$, then $\mathbf{w} \cdot \mathbf{x} + b \leq -\gamma$.

Unfortunately, this formulation doesn't really work properly. The problem is that by increasing \mathbf{w} and b , we can always allow a larger value of γ . For example, suppose that \mathbf{w} and b satisfy the constraint above. If we replace \mathbf{w} by $2\mathbf{w}$ and b by $2b$, we observe that for all i , $y_i((2\mathbf{w}) \cdot \mathbf{x}_i + 2b) \geq 2\gamma$. Thus, $2\mathbf{w}$ and $2b$ is always a better choice than \mathbf{w} and b , so there is no best choice and no maximum γ .

12.3.2 Normalizing the Hyperplane

The solution to the problem that we described intuitively above is to normalize the weight vector \mathbf{w} . That is, the unit of measure perpendicular to the separating hyperplane is the unit vector $\mathbf{w}/\|\mathbf{w}\|$. Recall that $\|\mathbf{w}\|$ is the Frobenius norm, or the square root of the sum of the squares of the components of \mathbf{w} . We shall require that \mathbf{w} be such that the parallel hyperplanes that just touch the support vectors are described by the equations $\mathbf{w} \cdot \mathbf{x} + b = +1$ and $\mathbf{w} \cdot \mathbf{x} + b = -1$, as suggested by Fig. 12.15.

Our goal becomes to maximize γ , which is now the multiple of the unit vector $\mathbf{w}/\|\mathbf{w}\|$ between the separating hyperplane and the parallel hyperplanes through the support vectors. Consider one of the support vectors, say \mathbf{x}_2 shown in Fig. 12.15. Let \mathbf{x}_1 be the projection of \mathbf{x}_2 onto the far hyperplane, also as suggested by Fig. 12.15. Note that \mathbf{x}_1 need not be a support vector or even a point of the training set. The distance from \mathbf{x}_2 to \mathbf{x}_1 in units of $\mathbf{w}/\|\mathbf{w}\|$ is 2γ . That is,




Figure 12.15: Normalizing the weight vector for an SVM

$$\mathbf{x}_1 = \mathbf{x}_2 + 2\gamma \frac{\mathbf{w}}{\|\mathbf{w}\|} \quad (12.1)$$

Since \mathbf{x}_1 is on the hyperplane defined by $\mathbf{w} \cdot \mathbf{x} + b = +1$, we know that $\mathbf{w} \cdot \mathbf{x}_1 + b = 1$. If we substitute for \mathbf{x}_1 using Equation 12.1, we get

$$\mathbf{w} \cdot \left(\mathbf{x}_2 + 2\gamma \frac{\mathbf{w}}{\|\mathbf{w}\|} \right) + b = 1$$

Regrouping terms, we see

$$\mathbf{w} \cdot \mathbf{x}_2 + b + 2\gamma \frac{\mathbf{w} \cdot \mathbf{w}}{\|\mathbf{w}\|} = 1 \quad (12.2)$$

But the first two terms of Equation 12.2, $\mathbf{w} \cdot \mathbf{x}_2 + b$, sum to -1 , since we know that \mathbf{x}_2 is on the hyperplane $\mathbf{w} \cdot \mathbf{x} + b = -1$. If we move this -1 from left to right in Equation 12.2 and then divide through by 2, we conclude that

$$\gamma \frac{\mathbf{w} \cdot \mathbf{w}}{\|\mathbf{w}\|} = 1 \quad (12.3)$$

Notice also that $\mathbf{w} \cdot \mathbf{w}$ is the sum of the squares of the components of \mathbf{w} . That is, $\mathbf{w} \cdot \mathbf{w} = \|\mathbf{w}\|^2$. We conclude from Equation 12.3 that $\gamma = 1/\|\mathbf{w}\|$.

This equivalence gives us a way to reformulate the optimization problem originally stated in Section 12.3.1. Instead of maximizing γ , we want to minimize $\|\mathbf{w}\|$, which is the inverse of γ if we insist on normalizing the scale of \mathbf{w} . That is:

- Given a training set $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)$, minimize $\|\mathbf{w}\|$ (by varying \mathbf{w} and b) subject to the constraint that for all $i = 1, 2, \dots, n$,

$$y_i(\mathbf{w} \cdot \mathbf{x}_i + b) \geq 1$$

Example 12.8: Let us consider the four points of Fig. 11.1, supposing that they alternate as positive and negative examples. That is, the training set consists of

$$\begin{array}{ll} ([1, 2], +1) & ([2, 1], -1) \\ ([3, 4], +1) & ([4, 3], -1) \end{array}$$

Let $\mathbf{w} = [u, v]$. Our goal is to minimize $\sqrt{u^2 + v^2}$ subject to the constraints we derive from the four training examples. For the first, where $\mathbf{x}_1 = [1, 2]$ and $y_1 = +1$, the constraint is $(+1)(u + 2v + b) = u + 2v + b \geq 1$. For the second, where $\mathbf{x}_2 = [2, 1]$ and $y_2 = -1$, the constraint is $(-1)(2u + v + b) \geq 1$, or $2u + v + b \leq -1$. The last two points are analogously handled, and the four constraints we derive are:

$$\begin{array}{ll} u + 2v + b \geq 1 & 2u + v + b \leq -1 \\ 3u + 4v + b \geq 1 & 4u + 3v + b \leq -1 \end{array}$$

We shall cover in detail the subject of how one optimizes with constraints; the subject is broad and many packages are available for you to use. Section 12.3.4 discusses one method – gradient descent – in connection with a more general application of SVM, where there is no separating hyperplane. An illustration of how this method works will appear in Example 12.9.

In this simple example, the solution is easy to see: $b = 0$ and $\mathbf{w} = [u, v] = [-1, +1]$. It happens that all four constraints are satisfied exactly; i.e., each of the four points is a support vector. That case is unusual, since when the data is two-dimensional, we expect only three support vectors. However, the fact that the positive and negative examples lie on parallel lines allows all four constraints to be satisfied exactly. \square

12.3.3 Finding Optimal Approximate Separators

We shall now consider finding an optimal hyperplane in the more general case, where no matter which hyperplane we chose, there will be some points on the wrong side, and perhaps some points that are on the correct side, but too close to the separating hyperplane itself, so the margin requirement is not met. A typical situation is shown in Fig. 12.16. We see two points that are misclassified; they are on the wrong side of the separating hyperplane $\mathbf{w} \cdot \mathbf{x} + b = 0$. We also see two points that, while they are classified correctly, are too close to the separating hyperplane. We shall call all these points *bad* points.

Each bad point incurs a penalty when we evaluate a possible hyperplane. The amount of the penalty, in units to be determined as part of the optimization process, is shown by the arrow leading to the bad point from the hyperplane on the wrong side of which the bad point lies. That is, the arrows measure the distance from the hyperplane $\mathbf{w} \cdot \mathbf{x} + b = 1$ or $\mathbf{w} \cdot \mathbf{x} + b = -1$. The former is the baseline for training examples that are supposed to be above the separating hyperplane (because the label y is $+1$), and the latter is the baseline for points that are supposed to be below (because $y = -1$).

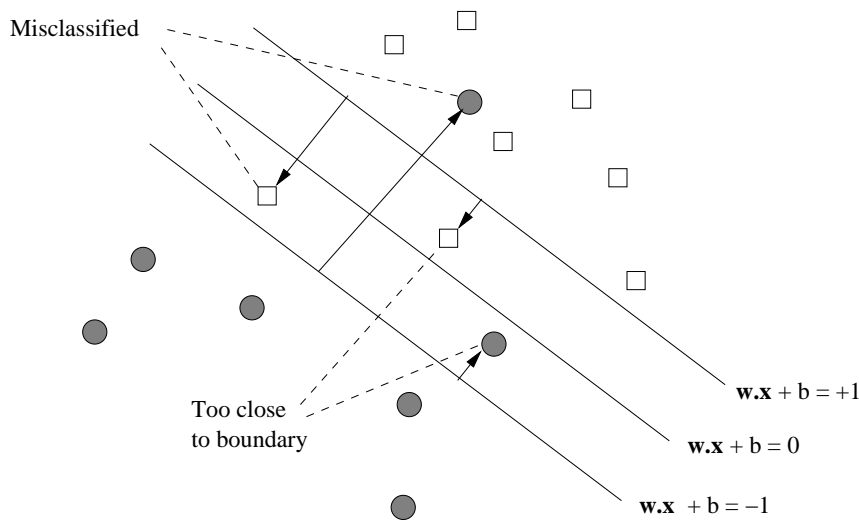


Figure 12.16: Points that are misclassified or are too close to the separating hyperplane incur a penalty; the amount of the penalty is proportional to the length of the arrow leading to that point

We have many options regarding the exact formula that we wish to minimize. Intuitively, we want $\|\mathbf{w}\|$ to be as small as possible, as we discussed in Section 12.3.2. But we also want the penalties associated with the bad points to be as small as possible. The most common form of a tradeoff is expressed by a formula that involves the term $\|\mathbf{w}\|^2/2$ and another term that involves a constant times the sum of the penalties.

To see why minimizing the term $\|\mathbf{w}\|^2/2$ makes sense, note that minimizing $\|\mathbf{w}\|$ is the same as minimizing any monotone function of $\|\mathbf{w}\|$, so it is at least an option to choose a formula in which we try to minimize $\|\mathbf{w}\|^2/2$. It turns out to be desirable because its derivative with respect to any component of \mathbf{w} is that component. That is, if $\mathbf{w} = [w_1, w_2, \dots, w_d]$, then $\|\mathbf{w}\|^2/2$ is $\frac{1}{2} \sum_{i=1}^n w_i^2$, so its partial derivative $\partial/\partial w_i$ is w_i . This situation makes sense because, as we shall see, the derivative of the penalty term with respect to w_i is a constant times each x_i , the corresponding component of each feature vector whose training example incurs a penalty. That in turn means that the vector \mathbf{w} and the vectors of the training set are commensurate in the units of their components.

Thus, we shall consider how to minimize the particular function

$$f(\mathbf{w}, b) = \frac{1}{2} \sum_{j=1}^d w_j^2 + C \sum_{i=1}^n \max\left\{0, 1 - y_i \left(\sum_{j=1}^d w_j x_{ij} + b \right)\right\} \quad (12.4)$$

The first term encourages small \mathbf{w} , while the second term, involving the constant C that must be chosen properly, represents the penalty for bad points

in a manner to be explained below. We assume there are n training examples (\mathbf{x}_i, y_i) for $i = 1, 2, \dots, n$, and $\mathbf{x}_i = [x_{i1}, x_{i2}, \dots, x_{id}]$. Also, as before, $\mathbf{w} = [w_1, w_2, \dots, w_d]$. Note that the two summations $\sum_{j=1}^d$ express the dot product of vectors.

The constant C , called the *regularization parameter*, reflects how important misclassification is. Pick a large C if you really do not want to misclassify points, but you would accept a narrow margin. Pick a small C if you are OK with some misclassified points, but want most of the points to be far away from the boundary (i.e., the margin is large).

We must explain the penalty function (second term) in Equation 12.4. The summation over i has one term

$$L(\mathbf{x}_i, y_i) = \max\left\{0, 1 - y_i \left(\sum_{j=1}^d w_j x_{ij} + b\right)\right\}$$

for each training example \mathbf{x}_i . The quantity L is a *hinge function*, suggested in Fig. 12.17, and we call its value the *hinge loss*. Let $z_i = y_i(\sum_{j=1}^d w_j x_{ij} + b)$. When z_i is 1 or more, the value of L is 0. But for smaller values of z_i , L rises linearly as z_i decreases.




Figure 12.17: The hinge function decreases linearly for $z \leq 1$ and then remains 0

Since we shall have need to take the derivative with respect to each w_j of $L(\mathbf{x}_i, y_i)$, note that the derivative of the hinge function is discontinuous. It is $-y_i x_{ij}$ for $z_i < 1$ and 0 for $z_i > 1$. That is, if $y_i = +1$ (i.e., the i th training example is positive), then

$$\frac{\partial L}{\partial w_j} = \text{if } \sum_{j=1}^d w_j x_{ij} + b \geq 1 \text{ then } 0 \text{ else } -x_{ij}$$

Moreover, if $y_i = -1$ (i.e., the i th training example is negative), then

$$\frac{\partial L}{\partial w_j} = \mathbf{if} \sum_{j=1}^d w_j x_{ij} + b \leq -1 \mathbf{then} 0 \mathbf{else} x_{ij}$$

The two cases can be summarized as one, if we include the value of y_i , as:

$$\frac{\partial L}{\partial w_j} = \mathbf{if} y_i \left(\sum_{j=1}^d w_j x_{ij} + b \right) \geq 1 \mathbf{then} 0 \mathbf{else} -y_i x_{ij} \quad (12.5)$$

12.3.4 SVM Solutions by Gradient Descent

A common approach to solving Equation 12.4 is to use quadratic programming. For large-scale data, another approach, *gradient descent* has an advantage. We can allow the data to reside on disk, rather than keeping it all in memory, which is normally required for quadratic solvers. To implement gradient descent, we compute the derivative of the equation with respect to b and each component w_j of the vector \mathbf{w} . Since we want to minimize $f(\mathbf{w}, b)$, we move b and the components w_j in the direction opposite to the direction of the gradient. The amount we move each component is proportional to the derivative with respect to that component.

Our first step is to use the trick of Section 12.2.4 to make b part of the weight vector \mathbf{w} . Notice that b is really the negative of a threshold on the dot product $\mathbf{w} \cdot \mathbf{x}$, so we can append a $(d+1)$ st component b to \mathbf{w} and append an extra component with value $+1$ to every feature vector in the training set (not -1 as we did in Section 12.2.4).

We must choose a constant η to be the fraction of the gradient that we move \mathbf{w} in each round. That is, we assign

$$w_j := w_j - \eta \frac{\partial f}{\partial w_j}$$

for all $j = 1, 2, \dots, d+1$.

The derivative $\frac{\partial f}{\partial w_j}$ of the first term in Equation 12.4, $\frac{1}{2} \sum_{j=1}^d w_j^2$, is easy; it is w_j . However, the second term involves the hinge function, so it is harder to express. We shall use an if-then expression to describe these derivatives, as in Equation 12.5. That is:

$$\frac{\partial f}{\partial w_j} = w_j + C \sum_{i=1}^n \left(\mathbf{if} y_i \left(\sum_{j=1}^d w_j x_{ij} + b \right) \geq 1 \mathbf{then} 0 \mathbf{else} -y_i x_{ij} \right) \quad (12.6)$$

Note that this formula gives us a partial derivative with respect to each component of \mathbf{w} , including w_{d+1} , which is b , as well as to the weights w_1, w_2, \dots, w_d . We continue to use b instead of the equivalent w_{d+1} in the if-then condition to remind us of the form in which the desired hyperplane is described.

To execute the gradient-descent algorithm on a training set, we pick:

1. Values for the parameters C and η .
2. Initial values for \mathbf{w} , including the $(d + 1)$ st component b .

Then, we repeatedly:

- (a) Compute the partial derivatives of $f(\mathbf{w}, b)$ with respect to the w_j 's.
- (b) Adjust the values of \mathbf{w} by subtracting $\eta \frac{\partial f}{\partial w_j}$ from each w_j .

Example 12.9: Figure 12.18 shows six points, three positive and three negative. We expect that the best separating line will be horizontal, and the only question is whether or not the separating hyperplane and the scale of \mathbf{w} allows the point $(2, 2)$ to be misclassified or to lie too close to the boundary. Initially, we shall choose $\mathbf{w} = [0, 1]$, a vertical vector with a scale of 1, and we shall choose $b = -2$. As a result, we see in Fig. 12.18 that the point $(2, 2)$ lies on the initial hyperplane and the three negative points are right at the margin. The parameter values we shall choose for gradient descent are $C = 0.1$, and $\eta = 0.2$.




Figure 12.18: Six points for a gradient-descent example

We begin by incorporating b as the third component of \mathbf{w} , and for notational convenience, we shall use u and v as the first two components, rather than the customary w_1 and w_2 . That is, we take $\mathbf{w} = [u, v, b]$. We also expand the two-dimensional points of the training set with a third component that is always 1. That is, the training set becomes

$$\begin{array}{lll} ([1, 4, 1], +1) & ([2, 2, 1], +1) & ([3, 4, 1], +1) \\ ([1, 1, 1], -1) & ([2, 1, 1], -1) & ([3, 1, 1], -1) \end{array}$$

In Fig. 12.19 we tabulate the if-then conditions and the resulting contributions to the summations over i in Equation 12.6. The summation must be

| | | | for u | for v | for b |
|-----------|-----------------------|-------------|---------|-------------|---------|
| if | $u + 4v + b \geq +1$ | then | 0 | else | -1 |
| if | $2u + 2v + b \geq +1$ | then | 0 | else | -2 |
| if | $3u + 4v + b \geq +1$ | then | 0 | else | -3 |
| if | $u + v + b \leq -1$ | then | 0 | else | +1 |
| if | $2u + v + b \leq -1$ | then | 0 | else | +2 |
| if | $3u + v + b \leq -1$ | then | 0 | else | +3 |

Figure 12.19: Sum each of these terms and multiply by C to get the contribution of bad points to the derivatives of f with respect to u , v , and b

multiplied by C and added to u , v , or b , as appropriate, to implement Equation 12.6.

The truth or falsehood of each of the six conditions in Fig. 12.19 determines the contribution of the terms in the summations over i in Equation 12.6. We shall represent the status of each condition by a sequence of x 's and o 's, with x representing a condition that does not hold and o representing one that does. The first few iterations of gradient descent are shown in Fig. 12.20.

| | $\mathbf{w} = [u, v]$ | b | Bad | $\partial/\partial u$ | $\partial/\partial v$ | $\partial/\partial b$ |
|-----|-----------------------|--------|---------|-----------------------|-----------------------|-----------------------|
| (1) | [0.000, 1.000] | -2.000 | oxoooo | -0.200 | 0.800 | -2.100 |
| (2) | [0.040, 0.840] | -1.580 | oxoxxx | 0.440 | 0.940 | -1.380 |
| (3) | [-0.048, 0.652] | -1.304 | oxoxxx | 0.352 | 0.752 | -1.104 |
| (4) | [-0.118, 0.502] | -1.083 | xxxxxxx | -0.118 | -0.198 | -1.083 |
| (5) | [-0.094, 0.542] | -0.866 | oxoxxx | 0.306 | 0.642 | -0.666 |
| (6) | [-0.155, 0.414] | -0.733 | xxxxxxx | | | |

Figure 12.20: Beginning of the process of gradient descent

Consider line (1). It shows the initial value of $\mathbf{w} = [0, 1]$. Recall that we use u and v for the components of \mathbf{w} , so $u = 0$ and $v = 1$. We also see the initial value of $b = -2$. We must use these values of u and v to evaluate the conditions in Fig. 12.19. The first of the conditions in Fig. 12.19 is $u + 4v + b \geq +1$. The left side is $0 + 4 + (-2) = 2$, so the condition is satisfied. However, the second condition, $2u + 2v + b \geq +1$ fails. The left side is $0 + 2 + (-2) = 0$. The fact that the sum is 0 means the second point $(2, 2)$ is exactly on the separating hyperplane, and not outside the margin. The third condition is satisfied, since $0 + 4 + (-2) = 2 \geq +1$. The last three conditions are also satisfied, and in fact are satisfied exactly. For instance, the fourth condition is $u + v + b \leq -1$. The left side is $0 + 1 + (-2) = -1$. Thus, the pattern oxoooo represents the outcome of these six conditions, as we see in the first line of Fig. 12.20.

We use these conditions to compute the partial derivatives. For $\partial f/\partial u$, we

use u in place of w_j in Equation 12.6. This expression thus becomes

$$u + C(0 + (-2) + 0 + 0 + 0 + 0) = 0 + \frac{1}{10}(-2) = -0.2$$

The sum multiplying C can be explained this way. For each of the six conditions of Fig. 12.19, take 0 if the condition is satisfied, and take the value in the column labeled “for u ” if it is not satisfied. Similarly, for v in place of w_j we get $\partial f/\partial v = 1 + \frac{1}{10}(0 + (-2) + 0 + 0 + 0 + 0) = 0.8$. Finally, for b we get $\partial f/\partial b = -2 + \frac{1}{10}(0 + (-1) + 0 + 0 + 0 + 0) = -2.1$.

We can now compute the new \mathbf{w} and b that appear on line (2) of Fig. 12.20. Since we chose $\eta = 1/5$, the new value of u is $0 - \frac{1}{5}(-0.2) = -0.04$, the new value of v is $1 - \frac{1}{5}(0.8) = 0.84$, and the new value of b is $-2 - \frac{1}{5}(-2.1) = -1.58$.

To compute the derivatives shown in line (2) of Fig. 12.20 we must first check the conditions of Fig. 12.19. While the outcomes of the first three conditions have not changed, the last three are no longer satisfied. For example, the fourth condition is $u + v + b \leq -1$, but $0.04 + 0.84 + (-1.58) = -0.7$, which is not less than -1 . Thus, the pattern of bad points becomes *oxxxxx*. We now have more nonzero terms in the expressions for the derivatives. For example $\partial f/\partial u = 0.04 + \frac{1}{10}(0 + (-2) + 0 + 1 + 2 + 3) = 0.44$.

The values of \mathbf{w} and b in line (3) are computed from the derivatives of line (2) in the same way as they were computed in line (2). The new values do not change the pattern of bad points; it is still *oxxxxx*. However, when we repeat the process for line (4), we find that all six conditions are unsatisfied. For instance, the first condition, $u + 4v + b \geq +1$ is not satisfied, because $(-0.118 + 4 \times 0.502 + (-1.083)) = 0.807$, which is less than 1. In effect, the first point has become too close to the separating hyperplane, even though it is properly classified.

We can see that in line (5) of Fig. 12.20, the problems with the first and third points are corrected, and we go back to pattern *oxxxxx* of bad points. However, at line (6), the points have again become too close to the separating hyperplane, so we revert to the *xxxxxx* pattern of bad points. You are invited to continue the sequence of updates to \mathbf{w} and b for several more iterations.

One might wonder why the gradient-descent process seems to be converging on a solution where at least some of the points are inside the margin, when there is an obvious hyperplane (horizontal, at height 1.5) with a margin of $1/2$, that separates the positive and negative points. The reason is that when we picked $C = 0.1$ we were saying that we really don’t care too much whether there are points inside the margins, or even if points are misclassified. We were saying also that what was important was a large margin (which corresponds to a small $\|\mathbf{w}\|$), even if some points violated that same margin. \square

12.3.5 Stochastic Gradient Descent

The gradient-descent algorithm described in Section 12.3.4 is often called *batch* gradient descent, because at each round, all the training examples are consid-

ered as a “batch.” While it is effective on small datasets, it can be too time-consuming to execute on a large dataset, where we must visit every training example, often many times before convergence.

An alternative, called *stochastic* gradient descent, considers one training example, or a few training examples at a time and adjusts the current estimate of the error function (\mathbf{w} in the SVM example) in the direction indicated by only the small set of training examples considered. Additional rounds are possible, using other sets of training examples; these can be selected randomly or according to some fixed strategy. Note that it is normal that some members of the training set are *never* used in a stochastic gradient descent algorithm.

Example 12.10: Recall the UV-decomposition algorithm discussed in Section 9.4.3. This algorithm was described as an example of batch gradient descent. We can regard each of the nonblank entries in the matrix M we are trying to approximate by the product UV as a training example, and the error function is the root-mean-square error between the product of the current matrices U and V and the matrix M , considering only those elements where M is nonblank.

However, if M has a very large number of nonblank entries, as would be the case if M represented, say, purchases of items by Amazon customers or movies that Netflix customers had rated, then it is not practical to make repeated passes over the entire set of nonblank entries of M when adjusting the entries in U and V . A stochastic gradient descent would look at a single nonblank entry of M and compute the change to each element of U and V that would make the product UV agree with that element of M . We would not make that change to the elements of U and V completely, but rather choose some learning rate η less than 1 and change each element of U and V by the fraction η of the amount that would be necessary to make UV equal M in the chosen entry. \square

12.3.6 Parallel Implementation of SVM

One approach to parallelism for SVM is analogous to what we suggested for perceptrons in Section 12.2.8. You can start with the current \mathbf{w} and b , and in parallel do several iterations based on each training example. Then average the changes for each of the examples to create a new \mathbf{w} and b . If we distribute \mathbf{w} and b to each mapper, then the Map tasks can do as many iterations as we wish to do in one round, and we need use the Reduce tasks only to average the results. One iteration of MapReduce is needed for each round.

A second approach is to follow the prescription given here, but implement the computation of the second term in Equation 12.4 in parallel. The contribution from each training example can then be summed. This approach requires one round of MapReduce for each iteration of gradient descent.

12.3.7 Exercises for Section 12.3

Exercise 12.3.1: Continue the iterations of Fig. 12.20 for three more iterations.

Exercise 12.3.2: The following training set obeys the rule that the positive examples all have vectors whose components sum to 10 or more, while the sum is less than 10 for the negative examples.

$$\begin{array}{lll} ([3, 4, 5], +1) & ([2, 7, 2], +1) & ([5, 5, 5], +1) \\ ([1, 2, 3], -1) & ([3, 3, 2], -1) & ([2, 4, 1], -1) \end{array}$$

(a) Which of these six vectors are the support vectors?

! (b) Suggest a vector \mathbf{w} and constant b such that the hyperplane defined by $\mathbf{w} \cdot \mathbf{x} + b = 0$ is a good separator for the positive and negative examples. Make sure that the scale of \mathbf{w} is such that all points are outside the margin; that is, for each training example (\mathbf{x}, y) , you have $y(\mathbf{w} \cdot \mathbf{x} + b) \geq +1$.

! (c) Starting with your answer to part (b), use gradient descent to find the optimum \mathbf{w} and b . Note that if you start with a separating hyperplane, and you scale \mathbf{w} properly, then the second term of Equation 12.4 will always be 0, which simplifies your work considerably.

! **Exercise 12.3.3:** The following training set obeys the rule that the positive examples all have vectors whose components have an odd sum, while the sum is even for the negative examples.

$$\begin{array}{lll} ([1, 2], +1) & ([3, 4], +1) & ([5, 2], +1) \\ ([2, 4], -1) & ([3, 1], -1) & ([7, 3], -1) \end{array}$$

(a) Suggest a starting vector \mathbf{w} and constant b that classifies at least three of the points correctly.

!! (b) Starting with your answer to (a), use gradient descent to find the optimum \mathbf{w} and b .

12.4 Learning from Nearest Neighbors

In this section we consider several examples of “learning,” where the entire training set is stored, perhaps, preprocessed in some useful way, and then used to classify future examples or to compute the value of the label that is most likely associated with the example. The feature vector of each training example is treated as a data point in some space. When a new point arrives and must be classified, we find the training example or examples that are closest to the new point, according to the distance measure for that space. The estimated label is then computed by combining the closest examples in some way.

12.4.1 The Framework for Nearest-Neighbor Calculations

The training set is first preprocessed and stored. The decisions take place when a new example, called the *query example* arrives and must be classified.

There are several decisions we must make in order to design a nearest-neighbor-based algorithm that will classify query examples. We enumerate them here.

1. What distance measure do we use?
2. How many of the nearest neighbors do we look at?
3. How do we weight the nearest neighbors? Normally, we provide a function (the *kernel function*) of the distance between the query example and its nearest neighbors in the training set, and use this function to weight the neighbors.
4. How do we define the label to associate with the query? This label is some function of the labels of the nearest neighbors, perhaps weighted by the kernel function, or perhaps not. If there is no weighting, then the kernel function need not be specified.

12.4.2 Learning with One Nearest Neighbor

The simplest cases of nearest-neighbor learning are when we choose only the one neighbor that is nearest the query example. In that case, there is no use for weighting the neighbors, so the kernel function is omitted. There is also typically only one possible choice for the labeling function: take the label of the query to be the same as the label of the nearest neighbor.

Example 12.11: Figure 12.21 shows some of the examples of dogs that last appeared in Fig. 12.1. We have dropped most of the examples for simplicity, leaving only three Chihuahuas, two Dachshunds, and two Beagles. Since the height-weight vectors describing the dogs are two-dimensional, there is a simple and efficient way to construct a *Voronoi diagram* for the points, in which the perpendicular bisectors of the lines between each pair of points is constructed. Each point gets a region around it, containing all the points to which it is the nearest. These regions are always convex, although they may be open to infinity in one direction.² It is also a surprising fact that, even though there are $O(n^2)$ perpendicular bisectors for n points, the Voronoi diagram can be found in $O(n \log n)$ time.

In Fig. 12.21 we see the Voronoi diagram for the seven points. The boundaries that separate dogs of different breeds are shown solid, while the boundaries

²While the region belonging to any one point is convex, the union of the regions for two or more points might not be convex. Thus, in Fig. 12.21 we see that the region for all Dachshunds and the region for all Beagles are not convex. That is, there are points p_1 and p_2 that are both classified Dachshunds, but the midpoint of the line between p_1 and p_2 is classified as a Beagle, and vice versa.




Figure 12.21: Voronoi diagram for the three breeds of dogs

between dogs of the same breed are shown dashed. Suppose a query example q is provided. Note that q is a point in the space of Fig. 12.21. We find the region into which q falls, and give q the label of the training example to which that region belongs. Note that it is not too hard to find the region of q . We have to determine to which side of certain lines q falls. This process is the same as we used in Sections 12.2 and 12.3 to compare a vector \mathbf{x} with a hyperplane perpendicular to a vector \mathbf{w} . In fact, if the lines that actually form parts of the Voronoi diagram are preprocessed properly, we can make the determination in $O(\log n)$ comparisons; it is not necessary to compare q with all of the $O(n \log n)$ lines that form part of the diagram. \square

12.4.3 Learning One-Dimensional Functions

Another simple and useful case of nearest-neighbor learning has one-dimensional data. In this situation, the training examples are of the form $([x], y)$, and we shall write them as (x, y) , identifying a one-dimensional vector with its lone component. In effect, the training set is a collection of samples of the value of a function $y = f(x)$ for certain values of x , and we must interpolate the function f at all points. There are many rules that could be used, and we shall only outline some of the popular approaches. As discussed in Section 12.4.1, the approaches vary in the number of neighbors they use, whether or not the

neighbors are weighted, and if so, how the weight varies with distance.

Suppose we use a method with k nearest neighbors, and x is the query point. Let x_1, x_2, \dots, x_k be the k nearest neighbors of x , and let the weight associated with training point (x_i, y_i) be w_i . Then the estimate of the label y for x is $\sum_{i=1}^k w_i y_i / \sum_{i=1}^k w_i$. Note that this expression gives the weighted average of the labels of the k nearest neighbors.

Example 12.12: We shall illustrate four simple rules, using the training set $(1, 1)$, $(2, 2)$, $(3, 4)$, $(4, 8)$, $(5, 4)$, $(6, 2)$, and $(7, 1)$. These points represent a function that has a peak at $x = 4$ and decays exponentially on both sides. Note that this training set has values of x that are evenly spaced. There is no requirement that the points have any regular pattern. Some possible ways to interpolate values are:




Figure 12.22: Results of applying the first two rules in Example 12.12

1. *Nearest Neighbor.* Use only the one nearest neighbor. There is no need for a weighting. Just take the value of any $f(x)$ to be the label y associated with the training-set point nearest to query point x . The result of using this rule on the example training set described above is shown in Fig. 12.22(a).
2. *Average of the Two Nearest Neighbors.* Choose 2 as the number of nearest neighbors to use. The weights of these two are each $1/2$, regardless of how far they are from the query point x . The result of this rule on the example training set is in Fig. 12.22(b).
3. *Weighted Average of the Two Nearest Neighbors.* We again choose two nearest neighbors, but we weight them in inverse proportion to their distance from the query point. Suppose the two neighbors nearest to query

point x are x_1 and x_2 . Suppose first that $x_1 < x < x_2$. Then the weight of x_1 , the inverse of its distance from x , is $1/(x - x_1)$, and the weight of x_2 is $1/(x_2 - x)$. The weighted average of the labels is

$$\left(\frac{y_1}{x - x_1} + \frac{y_2}{x_2 - x}\right) / \left(\frac{1}{x - x_1} + \frac{1}{x_2 - x}\right)$$

which, when we multiply numerator and denominator by $(x - x_1)(x_2 - x)$, simplifies to

$$\frac{y_1(x_2 - x) + y_2(x - x_1)}{x_2 - x_1}$$

This expression is the linear interpolation of the two nearest neighbors, as shown in Fig. 12.23(a). When both nearest neighbors are on the same side of the query x , the same weights make sense, and the resulting estimate is an *extrapolation*. We see extrapolation in Fig. 12.23(a) in the range $x = 0$ to $x = 1$. In general, when points are unevenly spaced, we can find query points in the interior where both neighbors are on one side.

4. *Average of Three Nearest Neighbors.* We can average any number of the nearest neighbors to estimate the label of a query point. Figure 12.23(b) shows what happens on our example training set when the three nearest neighbors are used.

□




Figure 12.23: Results of applying the last two rules in Example 12.12

12.4.4 Kernel Regression

A way to construct a continuous function that represents the data of a training set well is to consider all points in the training set, but weight the points using a kernel function that decays with distance. A popular choice is to use a normal distribution (or “bell curve”), so the weight of a training point x when the query is q is $e^{-(x-q)^2/\sigma^2}$. Here σ is the standard deviation of the distribution and the query q is the mean. Roughly, points within distance σ of q are heavily weighted, and those further away have little weight. The advantage of using a kernel function that is itself continuous and that is defined for all points in the training set is to be sure that the resulting function learned from the data is itself continuous (see Exercise 12.4.6 for a discussion of the problem when a simpler weighting is used).

Example 12.13: Let us use the seven training examples of Example 12.12. To make calculation simpler, we shall not use the normal distribution as the kernel function, but rather another continuous function of distance, namely $w = 1/(x - q)^2$. That is, weights decay as the square of the distance. Suppose the query q is 3.5. The weights w_1, w_2, \dots, w_7 of the seven training examples $(x_i, y_i) = (i, 8/2^{|i-4|})$ for $i = 1, 2, \dots, 7$ are shown in Fig. 12.24.

| | | | | | | | | |
|-----|-----------|------|-----|----|----|------|------|------|
| (1) | x_i | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| (2) | y_i | 1 | 2 | 4 | 8 | 4 | 2 | 1 |
| (3) | w_i | 4/25 | 4/9 | 4 | 4 | 4/9 | 4/25 | 4/49 |
| (4) | $w_i y_i$ | 4/25 | 8/9 | 16 | 32 | 16/9 | 8/25 | 4/49 |

Figure 12.24: Weights of points when the query is $q = 3.5$

Lines (1) and (2) of Fig. 12.24 give the seven training points. The weight of each when the query is $q = 3.5$ is given in line (3). For instance, for $x_1 = 1$, the weight $w_1 = 1/(1 - 3.5)^2 = 1/(-2.5)^2 = 4/25$. Then, line (4) shows each y_i weighted by the weight from line (3). For instance, the column for x_2 has value $8/9$ because $w_2 y_2 = 2 \times (4/9)$.

To compute the label for the query $q = 3.5$ we sum the weighted values of the labels in the training set, as given by line (4) of Fig. 12.24; this sum is 51.23. We then divide by the sum of the weights in line (3). This sum is 9.29, so the ratio is $51.23/9.29 = 5.51$. That estimate of the value of the label for $q = 3.5$ seems intuitively reasonable, since q lies midway between two points with labels 4 and 8. \square

12.4.5 Dealing with High-Dimensional Euclidean Data

We saw in Section 12.4.2 that the two-dimensional case of Euclidean data is fairly easy. There are several large-scale data structures that have been developed for finding near neighbors when the number of dimensions grows, and

Problems in the Limit for Example 12.13

Suppose q is exactly equal to one of the training examples x . If we use the normal distribution as the kernel function, there is no problem with the weight of x ; it is 1. However, with the kernel function discussed in Example 12.13, the weight of x is $1/(x-q)^2 = \infty$. Fortunately, this weight appears in both the numerator and denominator of the expression that estimates the label of q . It can be shown that in the limit as q approaches x , the label of x dominates all the other terms in both numerator and denominator, so the estimated label of q is the same as the label of x . That makes excellent sense, since $q = x$ in the limit.

the training set is large. We shall not cover these structures here, because the subject could fill a book by itself, and there are many places available to learn about these techniques, collectively called *multidimensional index structures*. The references for this chapter give some of these sources for information about such structures as *kd*-Trees, R-Trees, and Quad Trees.

Unfortunately, for high-dimensional data, there is little that can be done to avoid searching a large portion of the data. This fact is another manifestation of the “curse of dimensionality” from Section 7.1.3. Two ways to deal with the “curse” are the following:

1. *VA Files*. Since we must look at a large fraction of the data anyway in order to find the nearest neighbors of a query point, we could avoid a complex data structure altogether. Accept that we must scan the entire file, but do so in a two-stage manner. First, a summary of the file is created, using only a small number of bits that approximate the values of each component of each training vector. For example, if we use only the high-order (1/4)th of the bits in numerical components, then we can create a file that is (1/4)th the size of the full dataset. However, by scanning this file we can construct a list of candidates that *might* be among the k nearest neighbors of the query q , and this list may be a small fraction of the entire dataset. We then look up only these candidates in the complete file, in order to determine which k are nearest to q .
2. *Dimensionality Reduction*. We may treat the vectors of the training set as a matrix, where the rows are the vectors of the training example, and the columns correspond to the components of these vectors. Apply one of the dimensionality-reduction techniques of Chapter 11, to compress the vectors to a small number of dimensions, small enough that the techniques for multidimensional indexing can be used. Of course, when processing a query vector q , the same transformation must be applied to q before searching for q 's nearest neighbors.

12.4.6 Dealing with Non-Euclidean Distances

To this point, we have assumed that the distance measure is Euclidean. However, most of the techniques can be adapted naturally to an arbitrary distance function d . For instance, in Section 12.4.4 we talked about using a normal distribution as a kernel function. Since we were thinking about a one-dimensional training set in a Euclidean space, we wrote the exponent as $-(x-q)^2$. However, for any distance function d , we can use as the weight of a point x at distance $d(x, q)$ from the query point q the value of

$$e^{-(d(x-q))^2/\sigma^2}$$

Note that this expression makes sense if the data is in some high-dimensional Euclidean space and d is the usual Euclidean distance or Manhattan distance or any other distance discussed in Section 3.5.2. It also makes sense if d is Jaccard distance or any other distance measure.

However, for Jaccard distance and the other distance measures we considered in Section 3.5 we also have the option to use locality-sensitive hashing, the subject of Chapter 3. Recall these methods are only approximate, and they could yield false negatives – training examples that were near neighbors to a query but that do not show up in a search.

If we are willing to accept such errors occasionally, we can build the buckets for the training set and keep them as the representation of the training set. These buckets are designed so we can retrieve all (or almost all, since there can be false negatives) training-set points that have a minimum similarity to a given query q . Equivalently, one of the buckets to which the query hashes will contain all those points within some maximum distance of q . We hope that as many nearest neighbors of q as our method requires will be found among those buckets.

Yet if different queries have radically different distances to their nearest neighbors, all is not lost. We can pick several distances $d_1 < d_2 < d_3 < \dots$. Build the buckets for locality-sensitive hashing using each of these distances. For a query q , start with the buckets for distance d_1 . If we find enough near neighbors, we are done. Otherwise, repeat the search using the buckets for d_2 , and so on, until enough nearest neighbors are found.

12.4.7 Exercises for Section 12.4

Exercise 12.4.1: Suppose we modified Example 12.11 to look at the two nearest neighbors of a query point q . Classify q with the common label if those two neighbors have the same label, and leave q unclassified if the labels of the neighbors are different.

- (a) Sketch the boundaries of the regions for the three dog breeds on Fig. 12.21.
 ! (b) Would the boundaries always consist of straight line segments for any training data?

Exercise 12.4.2: Suppose we have the following training set

$$\begin{array}{ll} ([1, 2], +1) & ([2, 1], -1) \\ ([3, 4], -1) & ([4, 3], +1) \end{array}$$

which is the training set used in Example 12.9. If we use nearest-neighbor learning with the single nearest neighbor as the estimate of the label of a query point, which query points are labeled +1?

Exercise 12.4.3: Consider the one-dimensional training set

$$(1, 1), (2, 2), (4, 3), (8, 4), (16, 5), (32, 6)$$

Describe the function $f(q)$, the label that is returned in response to the query q , when the interpolation used is:

- (a) The label of the nearest neighbor.
- (b) The average of the labels of the two nearest neighbors.
- ! (c) The average, weighted by distance, of the two nearest neighbors.
- (d) The (unweighted) average of the three nearest neighbors.

! **Exercise 12.4.4:** Apply the kernel function of Example 12.13 to the data of Exercise 12.4.3. For queries q in the range $2 < q < 4$, what is the label of q ?

Exercise 12.4.5: What is the function that estimates the label of query points using the data of Example 12.12 and the average of the four nearest neighbors?

!! **Exercise 12.4.6:** Simple weighting functions such as those in Example 12.12 need not define a continuous function. We can see that the constructed functions in Fig. 12.22 and Fig. 12.23(b) are not continuous, but Fig. 12.23(a) is. Does the weighted average of two nearest neighbors always give a continuous function?

12.5 Comparison of Learning Methods

Each of the methods discussed in this chapter and elsewhere has its advantages. In this closing section, we shall consider:

- Does the method deal with categorical features or only with numerical features?
- Does the method deal effectively with high-dimensional feature vectors?
- Is the model that the method constructs intuitively understandable?

Perceptrons and Support-Vector Machines: These methods can handle millions of features, but they only make sense if the features are numerical. They only are effective if there is a linear separator, or at least a hyperplane that approximately separates the classes. However, we can separate points by a nonlinear boundary if we first transform the points to make the separator be linear. The model is expressed by a vector, the normal to the separating hyperplane. Since this vector is often of very high dimension, it can be very hard to interpret the model.

Nearest-Neighbor Classification and Regression: Here, the model is the training set itself, so we expect it to be intuitively understandable. The approach can deal with multidimensional data, although the larger the number of dimensions, the sparser the training set will be, and therefore the less likely it is that we shall find a training point very close to the point we need to classify. That is, the “curse of dimensionality” makes nearest-neighbor methods questionable in high dimensions. These methods are really only useful for numerical features, although one could allow categorical features with a small number of values. For instance, a binary categorical feature like {male, female} could have the values replaced by 0 and 1, so there was no distance in this dimension between individuals of the same gender and distance 1 between other pairs of individuals. However, three or more values cannot be assigned numbers that are equidistant. Finally, nearest-neighbor methods have many parameters to set, including the distance measure we use (e.g., cosine or Euclidean), the number of neighbors to choose, and the kernel function to use. Different choices result in different classification, and in many cases it is not obvious which choices yield the best results.

Decision Trees: We have not discussed this commonly used method in this chapter, although it was introduced briefly in Section 9.2.7. Unlike the methods of this chapter, decision trees are useful for both categorical and numerical features. The models produced are generally quite understandable, since each decision is represented by one node of the tree. However, this approach is only useful for low-dimension feature vectors. The reason is that building decision trees with many levels leads to overfitting, where below the top levels, the decisions are based on peculiarities of small fractions of the training set, rather than fundamental properties of the data. But if a decision tree has few levels, then it cannot even mention more than a small number of features. As a result, the best use of decision trees is often to create an *ensemble* of many, low-depth trees and combine their decision in some way.

12.6 Summary of Chapter 12

- ◆ *Training Sets:* A training set consists of a feature vector, each component of which is a feature, and a label indicating the class to which the object represented by the feature vector belongs. Features can be categorical – belonging to an enumerated list of values – or numerical.

- ◆ *Test Sets and Overfitting:* When training some classifier on a training set, it is useful to remove some of the training set and use the removed data as a test set. After producing a model or classifier without using the test set, we can run the classifier on the test set to see how well it does. If the classifier does not perform as well on the test set as on the training set used, then we have overfit the training set by conforming to peculiarities of the training-set data which is not present in the data as a whole.
- ◆ *Batch Versus On-Line Learning:* In batch learning, the training set is available at any time and can be used in repeated passes. On-line learning uses a stream of training examples, each of which can be used only once.
- ◆ *Perceptrons:* This machine-learning method assumes the training set has only two class labels, positive and negative. Perceptrons work when there is a hyperplane that separates the feature vectors of the positive examples from those of the negative examples. We converge to that hyperplane by adjusting our estimate of the hyperplane by a fraction – the learning rate – of the direction that is the average of the currently misclassified points.
- ◆ *The Winnow Algorithm:* This algorithm is a variant of the perceptron algorithm that requires components of the feature vectors to be 0 or 1. Training examples are examined in a round-robin fashion, and if the current classification of a training example is incorrect, the components of the estimated separator where the feature vector has 1 are adjusted up or down, in the direction that will make it more likely this training example is correctly classified in the next round.
- ◆ *Nonlinear Separators:* When the training points do not have a linear function that separates two classes, it may still be possible to use a perceptron to classify them. We must find a function we can use to transform the points so that in the transformed space, the separator is a hyperplane.
- ◆ *Support-Vector Machines:* The SVM improves upon perceptrons by finding a separating hyperplane that not only separates the positive and negative points, but does so in a way that maximizes the margin – the distance perpendicular to the hyperplane to the nearest points. The points that lie exactly at this minimum distance are the support vectors. Alternatively, the SVM can be designed to allow points that are too close to the hyperplane, or even on the wrong side of the hyperplane, but minimize the error due to such misplaced points.
- ◆ *Solving the SVM Equations:* We can set up a function of the vector that is normal to the hyperplane, the length of the vector (which determines the margin), and the penalty for points on the wrong side of the margins. The regularization parameter determines the relative importance of a wide margin and a small penalty. The equations can be solved by several methods, including gradient descent and quadratic programming.

- ◆ *Nearest-Neighbor Learning*: In this approach to machine learning, the entire training set is used as the model. For each (“query”) point to be classified, we search for its k nearest neighbors in the training set. The classification of the query point is some function of the labels of these k neighbors. The simplest case is when $k = 1$, in which case we can take the label of the query point to be the label of the nearest neighbor.
- ◆ *Regression*: A common case of nearest-neighbor learning, called regression, occurs when there is only one feature vector, and it, as well as the label, are real numbers; i.e., the data defines a real-valued function of one variable. To estimate the label, i.e., the value of the function, for an unlabeled data point, we can perform some computation involving the k nearest neighbors. Examples include averaging the neighbors or taking a weighted average, where the weight of a neighbor is some decreasing function of its distance from the point whose label we are trying to determine.

12.7 References for Chapter 12

The perceptron was introduced in [11]. [7] introduces the idea of maximizing the margin around the separating hyperplane. A well-known book on the subject is [10].

The Winnow algorithm is from [9]. Also see the analysis in [1].

Support-vector machines appeared in [6]. [5] and [4] are useful surveys. [8] talks about a more efficient algorithm for the case of sparse features (most components of the feature vectors are zero). The use of gradient-descent methods is found in [2, 3].

1. A. Blum, “Empirical support for winnow and weighted-majority algorithms: results on a calendar scheduling domain,” *Machine Learning* **26** (1997), pp. 5–23.
2. L. Bottou, “Large-scale machine learning with stochastic gradient descent,” *Proc. 19th Intl. Conf. on Computational Statistics* (2010), pp. 177–187, Springer.
3. L. Bottou, “Stochastic gradient tricks, neural networks,” in *Tricks of the Trade, Reloaded*, pp. 430–445, Edited by G. Montavon, G.B. Orr and K.-R. Mueller, Lecture Notes in Computer Science (LNCS 7700), Springer, 2012.
4. C.J.C. Burges, “A tutorial on support vector machines for pattern recognition,” *Data Mining and Knowledge Discovery* **2** (1998), pp. 121–167.
5. N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods*, Cambridge University Press, 2000.

6. C. Cortes and V.N. Vapnik, "Support-vector networks," *Machine Learning* **20** (1995), pp. 273–297.
7. Y. Freund and R.E. Schapire, "Large margin classification using the perceptron algorithm," *Machine Learning* **37** (1999), pp. 277–296.
8. T. Joachims, "Training linear SVMs in linear time." *Proc. 12th ACM SIGKDD* (2006), pp. 217–226.
9. N. Littlestone, "Learning quickly when irrelevant attributes abound: a new linear-threshold algorithm," *Machine Learning* **2** (1988), pp. 285–318.
10. M. Minsky and S. Papert, *Perceptrons: An Introduction to Computational Geometry* (2nd edition), MIT Press, Cambridge MA, 1972.
11. F. Rosenblatt, "The perceptron: a probabilistic model for information storage and organization in the brain," *Psychological Review* **65**:6 (1958), pp. 386–408.