

Perceptron Classifiers

Charles Elkan
elkan@cs.ucsd.edu

October 2, 2007

Suppose we have N training examples. The training data are a matrix with N rows and p columns, where each example is represented by values for p different features. Assume that each feature value is a real number. Let feature value j for example number i be written x_{ij} . The label of example i is y_i . For example, $y_i = 1$ if message i is spam and $y_i = 0$ if it is not spam.

We have separate training and test sets of examples. Each test example is also represented as a row vector of length p . The label y for a test example is unknown. The output of a classifier is a guess at y .

The simplest way to distinguish between two classes in p -dimensional space is a hyperplane. In Euclidean space, this is a linear subspace of dimension $p - 1$. The parameters defining a hyperplane are a vector w in R^p and a scalar b . The former gives the orientation of the hyperplane, which is at right angles (also called perpendicular, also called orthogonal) to w .

The projection of a vector (i.e. a point) u onto another vector v is the scalar $u \cdot v = \sum_p u_p v_p$, where \cdot is the symbol for scalar product. Think of $u \cdot v$ as the length of the shadow of u falling onto v . The hyperplane consists of all points x whose projection onto w equals b , i.e. $x \cot w = w \cdot x = b$. The points on one side have projection greater than b , while the points on the other side have projection less than b . The distance along w from the origin to the hyperplane is $b/\|w\|$ where $\|w\|$ is the length of w so $\|w\|^2 = w \cdot w = \sum_p w_p^2 = 1$. Note that the pair $\langle w, b \rangle$ defines the same hyperplane as $\langle \lambda w, \lambda b \rangle$ for any $\lambda \neq 0$.

The hyperplane classifies $\langle x, y \rangle$ correctly if and only if $y(w \cdot x - b) > 0$, that is if and only if y and $w \cdot x - b$ have the same sign. A hyperplane is a convenient classifier because it is fast to apply to a test example z . It only takes $O(p)$ time to compute $w \cdot z$.




Figure 1: A hyperplane defined by a vector w and a scalar ρ . All points x indicated by circles have $x \cdot w > \rho$. Picture from an unknown web source.

Without loss of generality, we will only look at learning hyperplanes that go through the origin, that is with $b = 0$. “Without loss of generality” means that if we have an algorithm for learning hyperplanes through the origin, we can use it to learn hyperplanes with shift b by simply rewriting the training examples. Specifically, we extend each original example x with an extra coefficient $x_{p+1} = 1$. Let x' be the extended version of x and suppose $y(w' \cdot x') > 0$ for every training example $\langle x, y \rangle$. Then $w' \cdot x' = w \cdot x - b$ where $b = -w'_{p+1}$ so w' includes a vector w and a scalar b defining a hyperplane shifted away from the origin. (Illustration of changing 2D data on a circle to data on a circle on a sphere.)

The task of learning a hyperplane through the origin is:

Task: Given training data $\{\langle x_i, y_i \rangle\}$ find w such that $y_i(w \cdot x_i) > 0$ for all i . Geometrically, the constraint for each training example specifies a half-space that w must lie inside: $w \cdot x_i > 0$ if $y_i > 0$ and $w \cdot x_i < 0$ if $y_i < 0$. Some obvious questions arise: (1) What if there is no solution? (2) If more than one solution w exists, then which one should we choose? That is, what is the objective function of w to maximize?

If we fix a linear objective function, then the learning task becomes what is called a “linear programming” problem: maximizing an objective function that is linear subject to constraints that are also linear. This raises a third issue: (3) Even if a solution exists, algorithms for linear programming are complicated and slow and/or biologically not plausible.

In 1958 Frank Rosenblatt created much enthusiasm when he published a method called the “perceptron algorithm” that answers the three concerns above.¹ For convenience, let the training labels y_i be -1 or $+1$. The basic algorithm is very simple:

- $w := 0$
- while any training example $\langle x, y \rangle$ is misclassified, set $w := w + yx$

Consider an example of the algorithm in operation. Suppose all training examples are points on a circle, with positive points in one hemicircle, and negative points in the opposite hemicircle. Remember that the hyperplane goes through the origin and is perpendicular to the vector w .

Notice that as the algorithm proceeds the coefficients of the vector w may keep growing in magnitude, and/or in digits of precision.

¹For interesting and important historical background, see the Wikipedia biographies of Rosenblatt and of Marvin Minsky.

The perceptron algorithm has an “online” version also. Given an infinite stream of training examples $\langle x_t, y_t \rangle$ where time is indexed $t = 1, 2, 3, \dots$, this version is as follows:

- $w := 0$
- for $t = 1, 2, 3, \dots$
if $y_t(w \cdot x_t) \leq 0$ then set $w := w + y_t x_t$

This online learning algorithm is biologically plausible, because (i) it uses only simple arithmetic, (ii) the learner does not need to memorize examples, (iii) examples can arrive in any order selected by nature, and (iv) the learner can run the algorithm throughout its life.

The algorithm also has aspects that are not biologically plausible. The operation $w := w + y_t x_t$ needs more and more digits of precision to be performed correctly. It is plausible that neurons can do basic arithmetic easily, and that they may have evolved to implement an algorithm like the online perceptron. However, it is not plausible that they can do high-precision arithmetic.

The perceptron algorithm is mathematically important because of a theorem about its convergence, due to Novikoff in 1962.

Assumptions: Let the training set be finite or infinite. Let $R = \max_t \|x_t\|$. If the training data are all on the unit sphere, then $R = 1$ for example. Suppose the learning task is solvable, i.e. there exists some vector w^* of unit length and some $\delta > 0$ such that $y_t(w^* \cdot x_t) \geq \delta$ for all t .

Theorem: Under these assumptions, the perceptron algorithm converges after at most $(R/\delta)^2$ updates.

Proof: Let w_n be the w vector after n updates and let $w_0 = 0$. We will argue that whenever w is updated it becomes closer to w^* .

Suppose w_{n+1} is an update, i.e. w_n misclassifies x and hence $w_{n+1} = w_n + yx$. Consider

$$w_{n+1} \cdot w^* = (w_n + yx) \cdot w^* = w_n \cdot w^* + yx \cdot w^* \geq w_n \cdot w^* + \delta.$$

This says that the projection of w_{n+1} onto w^* has increased. We would like this to mean that w_{n+1} is closer to w^* . However, what it really means is that w_{n+1} is closer to w^* and/or w_{n+1} has grown bigger. So, consider the Euclidean length of w_{n+1} :

$$\|w_{n+1}\|^2 = \|w_n + yx\|^2 = \|w_n\|^2 + 2y(w_n \cdot x) + \|x\|^2 \leq \|w_n\|^2 + R^2$$

since $y(w_n \cdot x) \leq 0$. So, after N actual updates we know two facts: $\|w_n\|^2 \leq NR^2$ and $w_n \cdot w^* \geq N\delta$. Putting these together gives a contradiction if N is too large: $w_N \cdot w^* \leq \|w_N\| \|w^*\| = \|w_N\|$ so $N\delta \leq \|w_N\| \leq R\sqrt{N}$ so $\sqrt{N} \leq R/\delta$. **End of proof.**

Can the perceptron learn to distinguish between any two classes? If yes, it is a general-purpose biologically plausible learning algorithm! The answer is obvious in retrospect: No. The reason is simple. “You cannot learn what you cannot represent.” Many concepts (i.e. distinctions between classes) cannot be represented by a hyperplane. The simplest example of a non-representable, and hence non-learnable, concept is exclusive-or. $x_1 \text{ xor } x_2$ means $x_1 = x_2 = 0$ or $x_1 = x_2 = 1$

The insight that the perceptron algorithm is incapable of learning many very simple distinctions killed most interest in it for many years. However, there are good reasons to investigate the situation further.

So, why is the perceptron algorithm important? The online version of the algorithm is a lifelong learning method for an intelligent agent such as an animal or robot. The convergence theorem says that the final classifier is a $+1/-1$ weighted sum of a finite number of training points, even if the input is an infinite stream of data. The proof says that that the algorithm can generalize from a *finite* training set to a concept that is valid for an *infinite* set.

Of course, the guarantee is only true under some conditions. First, some such valid concept must actually exist. Specifically, learning converges after at most $(R/\delta)^2$ updates, where $R = \max \|x_t\|$ and δ is a level of precision: we need $y_t(w^* \cdot x_t) \geq \delta$ for all t .

A modern observation is that the number of updates until convergence does not depend on the dimensionality of the data. This suggests that perceptron learning will be useful for very high-dimensional data such as images. Indeed, this is true.

The resurgence of interest in perceptron-type classifiers came in the 1980s because backpropagation was invented as a training algorithm for multiple perceptrons connected sequentially.

It turns out that a multilayer perceptron with a single hidden layer is a universal approximator. This means that the network can mimic any continuous function $\mathbb{R}^n \rightarrow \mathbb{R}^m$ with any specified level of accuracy. However, as the desired accuracy increases, and as the function to be represented becomes less smooth, the number of nodes needed increases, and/or the number of digits of numerical precision needed. One of the first proofs of this fact was given by Hal White (professor of Economics at UCSD) [HSW89].

The universal-approximator property depends on having at least one hidden layer, and on the nodes being nonlinear, which are biologically plausible. How-

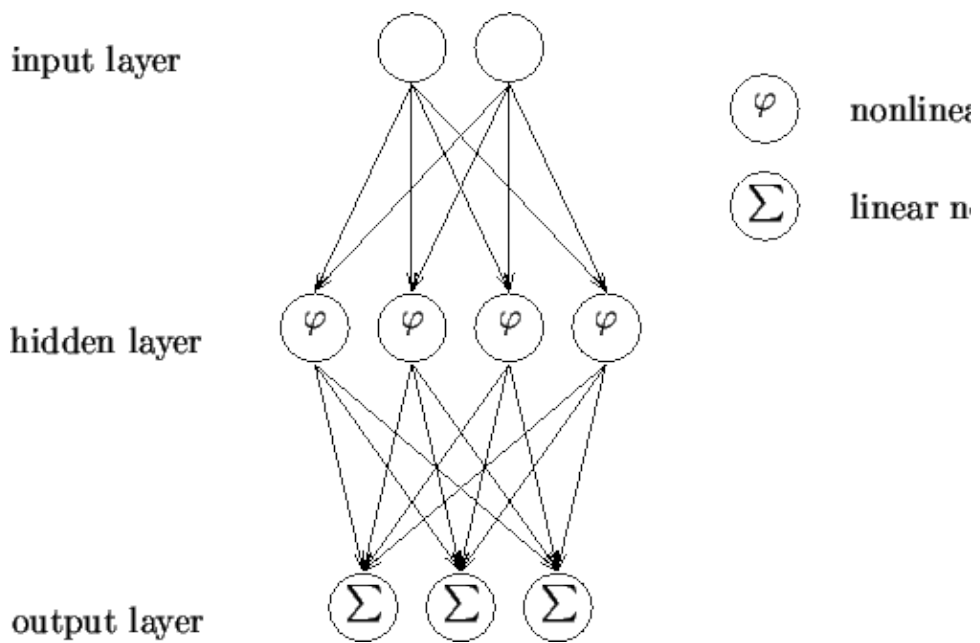


Figure 2: A multilayer perceptron with a single hidden layer. Picture from www.cis.hut.fi/ahonkela/dippa/.

ever it also depends on high-precision arithmetic, which may not biologically plausible. Moreover, plausibility for the representation does not necessarily imply plausibility for the training algorithm.

Overall, biological plausibility remains an open question. On the one hand, Terry Sejnowski (professor at UCSD) has written “No biological significance is claimed for the algorithm (backpropagation) by which the network developed” [LS88, page 454]. On the other hand, current findings in neuroscience suggest that forward and backward waves of activity, until quiescence, may be widespread in the brain. Also, some modern training algorithms benefit from injecting random noise, so high-precision arithmetic may not be needed.

The basic perceptron algorithm will never converge on nonseparable data. An idea for dealing with nonseparable data is to just iterate through the training data for a fixed number K of epochs. The disadvantage of this procedure is that just before finishing, it might do an update on an outlier, and hence terminate with a very bad concept w . A better idea is, for every w ever considered, to remember for how many training points it was correct.

This is the idea behind a method called the “voted perceptron” due to Yoav Freund (professor at UCSD) and Robert Schapire [FS99]:

```

n := 1
w1 := 0
c1 := 0
repeat for T epochs:
  for i = 1 to i = m (this is one epoch)
    if ⟨xi, yi⟩ is classified correctly then increment cn
    otherwise:
      increment n
      wn := wn-1 + yixi
      cn := 1

```

When the algorithm terminates we have a set of classifiers w_n each with a weight c_n that is its survival time. The survival times add up to Tm . We know that each w_n was correct on either c_n or $c_n - 1$ training examples, so c_n is a reasonable measure of the reliability of w_n .

The final classifier is nonlinear. It is

$$f(x) = \text{sign}\left(\sum_n c_n \text{sign}(w_n \cdot x)\right).$$

Practically, computing $f(x)$ requires storing all the intermediate w_n in memory. The classifier can be linearized as

$$f'(x) = \text{sign}\left(\sum_n c_n(w_n \cdot x)\right) = \text{sign}\left(x \cdot \sum_n c_n w_n\right).$$

In experiments this averaging method works slightly better than the voting method.

Why is the voted perceptron important? First, there is a theorem about generalization accuracy with m training examples, k training mistakes, and bounds R and δ . The result of the theorem is an upper bound on the error probability for an iid (independent identically distributed) test example. Second, the algorithm works well in practice on high-dimensional separable as well as nonseparable data. For example, with a kernel function 28 by 28 images of digits are separable.

References

- [FS99] Yoav Freund and Robert E. Schapire. Large margin classification using the perceptron algorithm. *Machine Learning*, 37(3):277–296, 1999.
- [HSW89] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2:359–366, 1989.
- [LS88] S. R. Lehky and T. J. Sejnowski. Network model of shape-from-shading: neural function arises from both receptive and projective fields. *Nature*, 333(6172):452–454, 1988.