

## Introduction to Neural Networks U. Minn. Psy 5038

### Linear discriminant, Bias/Variance & Model Selection, Pointers to support vector approaches

#### Initialize

##### ■ Read in Statistical Add-in packages:

```
Off[General::spell1];
<< Statistics`DescriptiveStatistics`
<< Statistics`DataManipulation`
<< Statistics`NormalDistribution`
<< Statistics`MultiDescriptiveStatistics`
<< Statistics`MultinormalDistribution`
```

#### Review Discriminant functions

Let's review earlier material on discriminant functions. Perceptron learning is an example of nonparametric statistical learning, because it doesn't require knowledge of the underlying probability distributions generating the data (such distributions are characterized by a relatively small number of "parameters", such as the mean and variance of a Gaussian distribution). Of course, how well it does will depend on the generative structure of the data. Much of the material below is covered in Duda and Hart (1978).

#### Linear discriminant functions: Two category case

A discriminant function,  $g(x)$  divides input space into two category regions depending on whether  $g(x) > 0$  or  $g(x) < 0$ . (We've switched notation,  $x=f$ ). The linear case corresponds to the simple perceptron unit we studied earlier:

$$g(x) = w \cdot x + w_0 \quad (1)$$

where  $w$  is the weight vector and  $w_0$  is the threshold (sometimes called bias, although this "bias" has nothing to do with statistical "bias").

Discriminant functions can be generalized, for example to quadratic decision surfaces:

$$g(x) = w_0 + \sum_{i=1} w_i x_i + \sum_{i=1} \sum_{j=1} w_{ij} x_i x_j \quad (2)$$

We've seen how  $g(x)=0$  defines a decision surface which in the linear case is a hyperplane. Suppose  $x_1$  and  $x_2$  are points sitting on the hyperplane, then their difference is a vector lying in the hyperplane

$$\begin{aligned} w \cdot x_1 + w_0 &= w \cdot x_2 + w_0 \\ w \cdot (x_1 - x_2) &= 0 \end{aligned} \quad (3)$$

so the weight vector  $w$  is normal to any vector lying in the hyperplane. Thus  $w$  determines how the plane is oriented. The normal vector  $w$  points into the region for which  $g(x) > 0$ , and  $-w$  points into the region for which  $g(x) < 0$ .

Let  $x$  be a point on the hyperplane. If we project  $x$  onto the normalized weight vector  $x \cdot w/|w|$ , we have the normal distance of the hyperplane from the origin equal to:


$$w \cdot x / |w| = -w_0 / |w| \quad (4)$$

Thus, the threshold determines the position of the hyperplane.

One can also show that the normal distance of  $x$  to the hyperplane is given by:

$$|g(x)| / |w| \quad (5)$$

So we've seen that: 1) discriminant function divides the input space by a hyperplane decision surface; 2) The orientation of the surface is determined by the weight vector  $w$ ; 3) the location is determined by the threshold  $w_0$ ; 4) the discriminant function gives a measure of how far in input vector is from the hyperplane.




## Multiple classes

Suppose there are  $c$  classes. There are a number of ways to define multiple class discriminant rules. One way that avoids undefined regions is:

$$g_i(\mathbf{x}) = \mathbf{w}_i \cdot \mathbf{x} + w_{i0}, \quad i = 1, \dots, c \quad (6)$$

$$\text{Assign } \mathbf{x} \text{ to the } i\text{th class if: } g_i(\mathbf{x}) > g_j(\mathbf{x}) \text{ for all } j \neq i. \quad (7)$$



It can be shown that this classifier partitions the input space into simply connected convex regions. This means that if you connect any two feature vectors belonging to the same class by a line, all points on the line are in the same class. Thus this linear classifier won't be able to handle problems for which there are disconnected clusters of features that all belong to the same class. Also, from a probabilistic perspective, if the underlying generative probability model for a given class has multiple modes, this linear classifier won't do a good job either.

## Task-dependent Dimensionality reduction

### Fisher's linear "discriminant"

The idea is that the original input space may be impractically huge, but if we can find a subspace (hyperplane) that preserves the distinctions between categories as well as possible, we can make our decisions in smaller space. We will derive the Fisher linear "discriminant".

This is closely related to the psychology idea of finding "distinctive" features. E.g. consider bird identification. If I want to discriminate cardinals from other birds in my backyard, I can make use of the fact that (males) cardinals may be the only birds that are red. So even tho' the image of a bird can have lots of dimensions, if I project the image on to the "red" axis, I can do fairly well with just one number. How about male vs. female human faces?

### ■ Generative model: two nearby gaussian classes

Pick one of Case 1 or Case 2:

Case 1:

```
(ar = {{1, .9}, {1, 1}};
ndista = MultinormalDistribution[{0, 0}, ar];
(br = {{1, .9}, {1, 2}};
ndistb = MultinormalDistribution[{0, 0}, br];
```

Case 2:

```
(ar = {{1, .1}, {1, 1}};
ndista = MultinormalDistribution[{0, 0}, ar];
(br = {{1, .1}, {1, 2}};
ndistb = MultinormalDistribution[{0, 0}, br];
```


```
pdf = PDF[ndista, {x1, x2}]
```

```
0.365126 e1/2 (-x1 (5.26316 x1 - 4.73684 x2) - x2 (-4.73684 x1 + 5.26316 x2))
```

```

nsamples = 500;
a = Table[Random[ndista], {nsamples}];
ga = ListPlot[a, AspectRatio -> 1,
  PlotStyle -> Hue[0.2], DisplayFunction -> Identity];
b = Table[Random[ndistb] + {0, 3}, {nsamples}];
gb = ListPlot[b, AspectRatio -> 1,
  PlotStyle -> Hue[0.6], DisplayFunction -> Identity];
Show[ga, gb, DisplayFunction -> $DisplayFunction];

```



- Try out different projections of the data by varying the slope (m) of the discriminant line

```

m = 1 / 2.0;
wnvec = {1, m} / Sqrt[1 + m^2];

```


```

aproj = Map[#1 * wnvec &, a.wnvec];
gaproj = ListPlot[aproj, AspectRatio -> 1,
  PlotStyle -> Hue[0.3], DisplayFunction -> Identity];

bproj = Map[#1 * wnvec &, b.wnvec];
gbproj = ListPlot[bproj, AspectRatio -> 1,
  PlotStyle -> Hue[0.7], DisplayFunction -> Identity];

Show[ga, gb, gaproj, gbproj,
  DisplayFunction -> $DisplayFunction, AspectRatio -> Automatic];

```



Find a value of m that separates the classes well along the projection line

---

Calculate the "signal-to-noise" ratio along the projection line: difference between the means divided by the square root of the product of the standard deviations along the line

---

#### ■ Theory for simple 2-class case

(see Duda and Hart for general case)

A measure of the separation between the projections is the difference between the means:

$$\begin{aligned}
 & | \mathbf{w} \cdot (\mu_1 - \mu_2) | \\
 & \text{and} \\
 \mu_1 &= \frac{1}{N} \sum_{i=1}^N \mathbf{x}, \text{ summed over the } N \text{ } \mathbf{x}' \text{ s from class 1} \\
 \mu_2 &= \frac{1}{M} \sum \mathbf{x}, \text{ summed over the } M \text{ } \mathbf{x}' \text{ s from class 2}
 \end{aligned} \tag{8}$$

where  $w$  is the unknown weight, or slope of the discriminant line.

To improve separation, we can't just scale  $w$ , because the noise scales too.

Measure of within-class scatter, the within-class scatter matrix

$$\mathbf{S}_W = \sum_{i=1}^2 \sum_{\mathbf{x} \in \text{class } i} (\mathbf{x} - \mu_i) (\mathbf{x} - \mu_i)^T \tag{9}$$

Measure of between class scatter

$$\mathbf{S}_B = (\mu_1 - \mu_2) \cdot (\mu_1 - \mu_2)^T \tag{10}$$

Find weights (e.g. slope) to maximize the criterion function

$$\mathcal{J}(\mathbf{w}) = \frac{\mathbf{w}^T \cdot \mathbf{S}_B \cdot \mathbf{w}}{\mathbf{w}^T \cdot \mathbf{S}_W \cdot \mathbf{w}} \tag{11}$$

Answer:

$$\mathbf{w} = \mathbf{S}_W^{-1} \cdot (\mu_1 - \mu_2) \tag{12}$$

#### ■ Demo: Finding Fisher's linear discriminant

```
normalize[x_] := x / Sqrt[x.x];
```

```
ma = Mean[a];
mb = Mean[b];
```


```
Sa = Sum[Outer[Times, a[[i]] - ma, a[[i]] - ma], {i, 1, nsamples}];
Sb = Sum[Outer[Times, b[[i]] - mb, b[[i]] - mb], {i, 1, nsamples}];
Sw = Sa + Sb;
wldf = normalize[Inverse[Sw].(ma - mb)]
```

```
{0.6524, -0.757875}
```

```
aproj = Map[#1*wldf &, a.wldf];
gaproj = ListPlot[aproj, AspectRatio -> 1,
  PlotStyle -> Hue[0.3], DisplayFunction -> Identity];

bproj = Map[#1*wldf &, b.wldf];
gbproj = ListPlot[bproj, AspectRatio -> 1,
  PlotStyle -> Hue[0.7], DisplayFunction -> Identity];

Show[ga, gb, gaproj, gbproj, DisplayFunction -> $DisplayFunction];
```



We started off with a 2-dimensional input problem and turned it into a 1-D problem. For the n-dimensional case, see Duda and Hart.

## Model selection & The bias/variance dilemma

Above we assumed linear separability. If the data are not linearly separable, then we resort to more complicated decision surfaces. That is what originally led us to non-linear multi-layer networks with hidden units. At first thought, one might think that we could just add lots of hidden units and solve any input-output problem. But there is a problem--the bias/variance dilemma.

Consider the regression problem, fitting data that may be a complex function of the input.

The problem in general is how to choose the function that both remembers the relationship between  $x$  and  $y$ , and generalizes with new values of  $x$ . The function has to be parameterized in some way that allows easy computation. What form should the function take? At first one might think that it should be as general as possible to allow all kinds of maps. For example, if one is fitting a curve, you might wish to use a very high-order polynomial, or a back-prop network with lots of hidden units. There is a drawback to the flexibility afforded by extra degrees of freedom in fitting the data. We can get drastically different fits for different sets of data that are randomly drawn from the same underlying process. The fact that we get different fit parameters (e.g. slope of a regression line) each time means that although we may exactly fit the data each time, we introduce variation between the average fit (over all data sets) and the fit for a single data set. We could get around this problem with a huge amount of data, but the problem is that the amount of required data typically grows exponentially with the order of the fit--an example of the so-called "curse of dimensionality".

On the other hand, if the function is restrictive, (e.g. straight lines through the origin), then we will get similar fits for different data sets, because all we have to adjust is one parameter--the slope. The problem here, is that the fit is only good if the underlying process is in fact a straight line through the origin. If it isn't a straight line for instance, there will be a fixed error or bias that will never go away, no matter how much data we collect. Statisticians refer to this problem as the *bias/variance* dilemma.

To sum up, lots of parameter flexibility (or lots of hidden units) has the benefit of fitting anything, but at the cost of sensitivity to variability in the data set--there is *variance* introduced by the fits found over multiple training sets (e.g. of a small fixed size). A fit with very few parameters is not as sensitive to the inevitable variability in the training set, but can give large constant errors or *bias* if the data do not match the underlying model. There is no general way of getting around this problem, and neural networks are no exception. We generalized linear regression to non-linear fits using error back-propagation. Because back-propagation models can have lots of hidden layers with many units and weights, they form a class of very flexible approximators and can fit almost any function. But these models can show high variability in their fits from one data set to the next, even when the data comes from the same underlying process. The linear associator is very restrictive, so would show high bias. One has to study the problem one is trying to model (independently of the neural network parameters if possible) and choose the appropriate network model to do the fit. Because the straight line, or in higher dimensions, a linear model is, in a mathematical sense, the "simplest", it does make good sense to try it first. Later on, one can complicate matters if it turns out that the data one is trying to fit is non-linear.

For a formal definition of the bias/variance trade-off, see pdf notes in class outline.

The bias/variance dilemma has led to techniques for model selection. This is the problem of how to pick the "simplest" model that is still consistent with the data. Much research in data modeling and pattern learning is devoted to the problem of how to achieve good generalization after learning from finite example sets. Good models combine some faithfulness to the data seen, while minimizing the complexity of the model used to fit the data. There are Bayesian approaches (e.g. MacKay, 1995) and related minimum description length techniques (cf. Bishop). But another rather different approach is support vector machines.

## Beyond linear separation, beyond sigmoidal kernels

Map the data (through some non-linear mapping, e.g. polynomial) to a higher-dimensional space to find the optimal hyperplane separating the data. But what is "optimal". Require good generalization, small VC dimension. Construct the hyperplane on a small number of support vectors, then the generalization ability will be high.

Main source: Vapnik (1995)

Demo links:

<http://svm.dcs.rhnc.ac.uk/pagesnew/GPat.shtml>

<http://svm.research.bell-labs.com/SVT/SVMsvt.html>

## References

- Duda, R. O., & Hart, P. E. (1973). *Pattern classification and scene analysis*. New York: John Wiley & Sons.
- Duda, R. O., Hart, P. E., & Stork, D. G. (2001). *Pattern classification* (2nd ed.). New York: Wiley. (Amazon.com)
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1), 1-58.
- MacKay, D. J. C. (1992). Bayesian interpolation. *Neural Computation*, 4(3), 415-447.
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge, UK: Cambridge University Press.
- Vapnik, V. N. (1995). *The nature of statistical learning*. New York: Springer-Verlag.
- <http://neuron.eng.wayne.edu/software.html>

© 1998, 2001 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.