

Introduction to Neural Networks
U. Minn. Psy 5038

Discriminant functions
Gradient descent
Regression, least squares, and Widrow-Hoff

Initialization

```
Off[SetDelayed::write]  
Off[General::spell1]
```

Introduction

Last time

- **Turning linear networks into classifiers with a hard threshold**
- **Perceptrons, perceptron learning rule.**
Linear part of a Perceptron-unit is a **linear discriminant**. Will return to discriminant functions in **Lecture 27**.

Today

- **Linear regression & brain-style learning -- the Widrow-Hoff learning rule**

Regression, Pseudoinverse, and Widrow-Hoff learning

Introduction

We have been studying a linear matrix model of memory based on the storage of connection weights that follow a particular Hebbian rule. We have studied the "psychology" of some operations of linear algebra and have seen some interesting parallels to human memory, such as interference and pattern reconstruction.

We've learned that linear networks can be configured to be supervised or unsupervised learning devices. But linear mappings are severely limited in what they can compute. The simple perceptron unit added a threshold. The network then makes discrete decisions.

The kinds of classifications, however, are still limited by the linearity of the decision surface, i.e. by the fact that it is a hyperplane. More complex networks can be built by adding layers, but then we have another problem: how can we learn the weights in this more complicated setting?

In this section, we return to the study of linear models of memory. However, we are going to view memory from the point of view of statistical regression. The idea is to treat memory as an attempt to fit past input/output associations into a model that can be used both for recall and for generalization.

The problem of regression in statistics is: given a set of vector inputs $\{\mathbf{x}_i\}$, and a set of corresponding vector outputs $\{\mathbf{y}_i\}$, one tries to find a transformation \mathbf{W} that will map $\mathbf{x} \rightarrow \mathbf{y}$ as closely as possible over the data available.

For this we need a model for \mathbf{W} , and a measure of *goodness of fit*. We will assume below a linear model for \mathbf{W} , so for the discrete case, \mathbf{W} is a matrix. Our measure of goodness of fit will be the *sum of the squared differences between predicted output and actual output*. In this linear case, this is *least squares regression*.

■ Overview of our strategy

Our return to strictly linear networks (no threshold) will allow us to introduce techniques (gradient descent learning in the context of finding the weights of a linear matrix transformation) and concepts that will generalize to non-linear networks with multiple layers of weights. Gradient descent will lead us to the Widrow-Hoff learning rule. By treating the linear case first, we will be able to see how the Widrow-Hoff learning rule relates to classic problems of statistical regression. This, in turn, will provide the introduction to the generalization of this rule to multiple layer networks with smooth non-linear squashing functions--the error back-propagation algorithm applied to the generic "connectionist" network model.

More terminology regarding supervised learning

Consider **supervised learning**. We have a "training set" $\{\mathbf{f}_i, \mathbf{t}_i\}$ representing inputs \mathbf{f}_i , and target outputs \mathbf{t}_i . The training set in some sense "samples" the larger space of possible input/output pairs $\{\mathbf{f}, \mathbf{t}\}$. We would like to learn a general mapping: $T: \mathbf{f} \rightarrow \mathbf{g}$ in such a way that T is a good fit to the training data (i.e. \mathbf{g} is close to \mathbf{t}), and generalizes well to novel inputs. The set of target data is the feedback for the "teacher".

In general, feedback can vary in the degree of precision it provides for learning. It can specify whether the mapping is correct or not. Or the feedback can provide information as to how far off the map T 's prediction of \mathbf{f} (i.e. How far is $T[\mathbf{f}_i]$ is from \mathbf{t}_i ?).

After training, one can require that T always maps members of the training set to exactly the target members, and generalizes appropriately for other inputs. This means that the learning should be *consistent*. **Interpolation** is between data points on a graph is an example of consistent learning. An example would be drawing lines connecting data points on a graph.

Or, we may require that the T maps the original members of the training set to outputs \mathbf{g} , that are close to the original targets \mathbf{t} . This is called **approximation** learning. Linear regression is an example of *approximation* learning. The linear associator was our first example of supervised learning. Approximation doesn't necessarily exactly fit the data points, but the goal is that it should come close. We are going to study linear regression, a fundamental case of approximation learning.

Least squares regression - linear models

The idea behind least squares regression is given a set of N training pairs $\{\mathbf{x}_i; y_i\}$, where i runs from 1 to N , we would like to find a function that given an input \mathbf{x} , the function approximates well the output y . If the function reproduces the association between input and output that it has seen before, this is like "remembering". But in addition, regression generalizes. So novel input values get mapped to predicted outputs based on past "experience". We've already seen how linear heteroassociative learning does this.

An example of linear regression

A fundamental assumption behind any learning system is that there is an underlying structure to the data--the relationship between associative pairs is not arbitrary. When trying to understand how a relationship can be learned between a set of two patterns, it is important to have some understanding of the structure of the relationship. For example, one shouldn't try to fit a straight line to data when there is evidence to indicate that the underlying process is quadratic. We will study this more when we learn about the "bias/variance dilemma". So let us assume that the data have an underlying structure that we are going to try to discover or approximate using \mathbf{W} .

We will study a simple "toy" problem that has the following very specific generative structure. The inputs are randomly located points on a 2D plane, and the outputs are heights above these points. The outputs lie approximately on a planar surface that runs through the origin and whose orientation is characterized by two parameters, \mathbf{a} and \mathbf{b} .

It may seem like overkill, but we are going to estimate \mathbf{W} for the same set of data in 4 (yes, 4!) different ways. The first two are drawn from standard linear algebra (a least squares solution using transpose and inverse, and the second using the pseudoinverse). The third introduces the method of "gradient descent" a common numerical technique which is used in many applications. We will use it later in several contexts. The fourth method is the most relevant to neural network theory--we estimate \mathbf{W} using a biologically plausible learning rule (the Widrow-Hoff rule).

■ Generative model: Synthetic training pairs

Although we will illustrate our examples with small dimensions, everything we do generalizes to higher dimensional inputs and outputs, and in fact the demonstration code below will work with higher dimensional input/output vectors.

Let x_1 and x_2 be the (scalar) inputs, and z be the (scalar) output:

$$\{x_1, x_2\} \rightarrow z,$$

and assuming the mapping is a plane through the origin and additive noise (1)

$$z = a x_1 + b x_2 + \text{noise}$$


We will want to learn a and b from the training pairs: $\{\{x_1, x_2\}, z\}$.

```
In[1]:= rsurface[a_, b_] :=  
  N[Table[{x1 = 1 RandomReal[], x2 = 1 RandomReal[],  
    a x1 + b x2 + 0.5 RandomReal[] - 0.25}, {120}], 2];
```

```
In[2]:= data = rsurface[2,3];
```

```
In[3]:= Graphics3D[Point /@ data, AxesLabel → {"x1", "x2", "z"}, Axes → True,  
  Ticks → None, BoxRatios → {1, 1, 1}, ImageSize → Small]
```

```
Out[3]=
```



```
In[6]:= Dimensions[data]
```

```
Out[6]= {120, 3}
```

You can also use: **ListPointPlot3D**[**data**]

Outdata represents the 1 dimensional z values, and **Indata** the 2-dimensional input values:

```
In[4]:= Outdata = data[[All, 3]];  
  Indata = data[[All, 1 ;; 2]];
```

■ 1. Least squares regression to find **W**

So let's assume we want to find a matrix **W** that will come close to reproducing values **y**, given inputs **x**. Of course, because we generated the data, we know the underlying structure and what the matrix **W** should be. It should be a 1x2 matrix = **{{2,3}}**. But let's assume we don't know the answer, and want to discover the weights from **Outdata** and **Indata**.

Imagine you have knobs that you can twiddle that let you adjust every element of the matrix **W**. For particular settings of these knobs, you can calculate **W.x_i** for any input/output pair. In general you won't be so lucky as to have **y_i=W.x_i**. So you take the opportunity to adjust the knobs to produce **W'** so that **W'.x_i** is closer to **y_i**. One way to measure "how close" is **|y_i - W'.x_i|²**. Of course, we usually have lots of data, so we could add up all the discrepancies between what **W** predicts, and what the "teacher" has provided in the training pair (i.e. **y_i** for all of the **i**'s).

In least squares regression, we try to find the values of the matrix that will minimize **e(W)**.

$$e(\mathbf{W}) = \sum_{i=1}^N |y_i - \mathbf{W}\mathbf{x}_i|^2$$


In many cases our error function will be over a very high dimensional space determined by the number of elements in **W**. HOWEVER, it is useful to get an intuition for the problem in a low dimensional space. We can actually get a picture of the error function **e(W)** for our 2-dimensional case as a function of the weights **w₁** and **w₂**:

```

e(W):=
  sum[{outdata[[i]] - {w1\ w2} . indata[[i]]} .
      {outdata[[i]] - {w1\ w2} . indata[[i]]} \
      {1\ 1\ pseudoinv[indata]\ 1}]
  
```

```
In[8]:= g = Plot3D[eW[w1, w2], {w1, 0, 4}, {w2, 0, 6},
ViewPoint -> {1.78, -2.861, 0.312}, AxesLabel -> {"w1", "w2", "e"},
ImageSize -> Small]
```

Out[8]=




Note that because our input data are 2-element vectors, and our output or target values are 1-D, the matrix W is not square--it has 1 row and 2 columns. So we represent W above as a vector.

The minimum appears to be near $\{2,3\}$. It is easier to see whether there is a minimum or not using `ContourPlot[]`.

```
In[9]:= gc = ContourPlot[eW[w1, w2], {w1, 0, 4}, {w2, 0, 6}, Contours -> 16,
ImageSize -> Small]
```

Out[9]=



Note: You can get a rough idea of the coordinates near the middle by moving your computer mouse over the above graphic after initializing the cell below:

```
In[10]:= Dynamic[MousePosition["Graphics"]]
```

```
Out[10]= None
```

Of course, in high dimensional spaces, \mathbf{W} is huge, so we can't visualize it, and we can't hope to actually "adjust knobs". But we can find the exact location of the minimum by finding \mathbf{W} at "the bottom of the bowl", which from calculus should be where the gradient of the error function, e is zero. Although there are a lot of indices to worry about, the familiar analogy from introductory calculus is to find the point at which the derivative of a function, say $f(x)$ is zero. $e(\mathbf{W})$ plays the role of $f(x)$, but \mathbf{W} is a matrix of variables. The solution can be written very concisely in terms of vector outerproducts, inversion, and matrix multiplication:

$$e(\mathbf{W}) = \sum_{i=1}^N |y_i - \mathbf{W}\mathbf{x}_i|^2$$

$$\frac{\partial e}{\partial \mathbf{W}} = -2 \sum_{i=1}^N (y_i - \mathbf{W}\mathbf{x}_i)\mathbf{x}_i^T = 0$$

$$\sum_{i=1}^N y_i \mathbf{x}_i^T - \mathbf{W} \sum_{i=1}^N \mathbf{x}_i \mathbf{x}_i^T$$

$$\mathbf{W} = \sum_{i=1}^N y_i \mathbf{x}_i^T \left(\sum_{i=1}^N \mathbf{x}_i \mathbf{x}_i^T \right)^{-1}$$

\mathbf{W} can be calculated by adding up two sets of outer products. Let's translate the last line above into *Mathematica* instructions:

```
In[13]:= sumX = Sum[Outer[Times, Indata[[i]], Indata[[i]]], {i, Length[Indata]}];
sumYX = Sum[Outer[Times, {Outdata[[i]]}, Indata[[i]]], {i, Length[Indata]}];
W = sumYX.Inverse[sumX]
```

```
Out[15]= (2.00417 3.03145)
```

The values for \mathbf{W} come close to what we would expect from the structure of our data, a plane with parameters [{2,3}](#).

■ 2. Pseudoinverse to find \mathbf{W}

For linear algebra aficionados, There is another way of solving the linear least squares regression that uses the "pseudoinverse" of matrix to find the least-squares solution.

For a square matrix \mathbf{X} , the inverse of \mathbf{X} is chosen so that $\mathbf{X}\mathbf{X}^{-1}$ is equal to the identity matrix, \mathbf{I} . The pseudoinverse, \mathbf{X}^* , of a rectangular matrix \mathbf{X} is chosen so that $\mathbf{X}\mathbf{X}^*$ is close to the identity matrix in the sense that the sum of the squares of all of the entries of $\mathbf{X}\mathbf{X}^*\mathbf{I}$ is least.

$$\mathbf{X}^* = \left(\mathbf{X}^T \mathbf{X}\right)^{-1} \mathbf{X}^T$$

or as:

$$\mathbf{X}^* = \mathbf{X}^T \left(\mathbf{X}\mathbf{X}^T\right)^{-1}$$

For our simple generative model, let's take all of the input vectors \mathbf{x} , and arrange them as columns in a matrix \mathbf{X} :

$$\mathbf{X} = \begin{pmatrix} x_1^1 & x_1^2 & \cdots & x_1^N \\ x_2^1 & x_2^2 & \cdots & x_2^N \end{pmatrix}$$

Now do the same for the \mathbf{y} 's:

$$\mathbf{Y} = \left(y_1 \quad y_2 \quad y_3 \cdots y_N\right)$$

And what is the matrix that maps the x 's to the y 's with least squared error? It is $\mathbf{X}^*\mathbf{Y}$:

```
In[17]:= Inverse[Transpose[Indata].Indata].Transpose[Indata].Outdata
```

```
Out[17]= {2.00417, 3.03145}
```

\mathbf{X}^* is the *pseudoinverse* (sometimes called the *generalized inverse*) of the matrix \mathbf{X} . The **PseudoInverse[]** function is built into *Mathematica*, so once we know about it, we can calculate the least-squares solution for our problem in one simple expression:

```
In[18]:= PseudoInverse[Indata].Outdata
```

```
Out[18]= {2.00417, 3.03145}
```

It can be shown that **PseudoInverse[X].X** is the identity matrix. We won't prove it here, but we can verify that it is the case with our data:

```
In[19]:= Chop[PseudoInverse[Indata].Indata]//MatrixForm
```

```
Out[19]//MatrixForm=
```

$$\begin{pmatrix} 1. & 0 \\ 0 & 1. \end{pmatrix}$$

Question: What are the dimensions of the above PseudoInverse?

■ 3. Gradient descent

Let's go back to the original global error function that we used above for standard linear least squares regression. There we found the weights that gave the minimum total error by setting the gradient of the error function to zero, and then solving for the weights.

Gradient descent, is a more general way of finding the minimum of an error function. It can be used when the error function is much more complicated (e.g. it isn't quadratic), and there is no linear solution. This is the kind of situation we will run into when we try to learn weights via supervised learning in multiple layer neural networks.

The idea is to start off at some location, \mathbf{W}_0 in weight space (which is just a guess), and iteratively move towards the minimum by always taking a step downhill. Again we take the elements of \mathbf{W} and think of them as defining a point in weight space, i.e. a vector in weight space. We want to find a vector, say \mathbf{W}_{best} , that tells us which direction to move in weight space so that that the error measured at location \mathbf{W}_0 goes decreases by as much as possible. We don't want to take too big of step or too small, so we need to control the length of \mathbf{W}_{best} , with a scalar η . In general, if we are at the i th point of an iteration, we would like:

$$\mathbf{W}^{i+1} = \mathbf{W}^i + \eta (\mathbf{W}_{\text{best}} \text{ at } \mathbf{W}^i)$$


To derive a rule for \mathbf{W}_{best} , we assume that the error function is a smooth function of \mathbf{W} . The downhill direction is given by the negative gradient of the error function.

$$\frac{d\mathbf{W}}{dt} = - \frac{\partial e}{\partial \mathbf{W}} = - \nabla e$$

$$\mathbf{W}^{i+1} = \mathbf{W}^i - \eta \left. \frac{\partial e}{\partial \mathbf{W}} \right|_{\mathbf{W}^i}$$

and then approximate the differential equation by a discrete update rule. $\mathbf{W}_{\text{best}} = - \nabla e$ evaluated at the i th weight space point in the iteration.

The figure below gives a graphical view in the small dimensional case of our example problem.




From the expression for the gradient which we wrote earlier in terms of outer products, we can obtain an expression for $\Delta \mathbf{W}$, with η taking the place of Δt :

$$\frac{\partial e}{\partial \mathbf{W}} = -2 \left(\sum_{i=1}^N \mathbf{y}_i \mathbf{x}_i^T - \mathbf{W} \sum_{i=1}^N \mathbf{x}_i \mathbf{x}_i^T \right)$$

```
In[21]:= eWgradient[wg_] :=
  Sum[Outer[Times, {Outdata[[i]]}, Indata[[i]]], {i, Length[Indata]}
  wg.Sum[Outer[Times, Indata[[i]], Indata[[i]]],
  {i, Length[Indata]}];
```

Now define a function T that specifies an updated weight matrix $\mathbf{W}=\mathbf{wg}$. You may have to adjust eta ($=\eta$) to make sure the steps are sufficiently small to get convergence, but not so small as to take long to converge.

```
In[22]:= i=0; wg={{RandomReal[5.0],RandomReal[5.0]}}; eta = .005; wglist = {};  
T[wg_] := wg + eta eWgradient[wg];  
w1 = Nest[T,wg,40];  
Clear[w1List]; w1List = NestList[T, wg, 20];  
g1 = ListPlot[Flatten[w1List, 1], Joined -> True,  
PlotStyle -> {Orange, Dashed, Thick}];  
  
Show[gc, g1]  
w1
```



Out[29]= (1.99777 3.03785)

■ 4. "Brain-style" learning: Iterative Widrow-Hoff learning to estimate \mathbf{W}

So far so good. But there are a several problems. First, we are interested in brain-style computation. Based on what we think we know about neurons, how could the brain compute transposes, do matrix inversion and multiplication?

Second, when we learn we don't seem to gather information on a whole set of training pairs, and then suddenly build a memory matrix. We typically learn *incrementally, trial by trial*.

Another problem is purely computational. What if the dimensionality of the vectors is really big? It is computationally expensive to invert large matrices. The above gradient descent procedure avoids the problem of inverting large matrices, but it involved computing a global error term over all the training pairs. We would like a method which would learn the regression map without having to store all the training pairs with the accompanying computation of a global error term. Instead, we'd like *to compute an error term incrementally, trial-by-trial*.

Can we discover the mapping \mathbf{W} in such a way so as to be biologically plausible, *and* avoid having to invert a large matrix? The *Widrow-Hoff* rule provides an answer. The basic idea behind *Widrow-Hoff* learning is to update \mathbf{W} iteratively with each new training pair. Let's start off with an arbitrary \mathbf{W} , find out which direction we would have to go in weight space to reduce the discrepancy between what \mathbf{W} tells us \mathbf{x} should map to and what it actually is, namely \mathbf{y} . We take our clue from gradient descent, but apply it each time a new training pair comes along. We recompute the error term each time a new training pair comes along.

$$e(\mathbf{W}) = \left| \mathbf{y}_i - \mathbf{W} \mathbf{x}_i \right|^2$$

$$\frac{\partial e}{\partial \mathbf{W}} = -2(\mathbf{y}_i - \mathbf{W} \mathbf{x}_i) \mathbf{x}_i^T$$

$$\frac{\partial e}{\partial \mathbf{W}} = -\frac{d\mathbf{W}}{dt}$$

$$\mathbf{W}^{i+1} = \mathbf{W}^i + \eta_i (\mathbf{y}_i - \mathbf{W}^i \mathbf{x}_i) \mathbf{x}_i^T$$

Let's try out this update rule on our synthetic training pairs.

```
In[30]:= ww1 = {{0,0}}; ww1list = {}; ww2list = {};
```

To visual how the learning progresses, we are using `ww1list` and `ww2list` to store the first and second weights respectively, for each learning iteration.

```

In[34]:= i=0;eta = 5;
While[i<Length[data],
  ++i;
  in = {data[[i,1]],data[[i,2]]} ; out = {data[[i,3]]};
  ww1 = ww1 + (eta/i) Outer[Times,(out - ww1.in),in];
  ww1list = Append[ww1list,ww1[[1,1]]];
  ww2list = Append[ww2list,ww1[[1,2]]];
];
ww1

```

```
Out[36]= (2.02028 3.02859)
```


Keep executing the above input cell. You may have to run through the above loop several times before reaching stable convergence. We can plot up the two weights as a function of the iteration number to see how the Widrow-Hoff rule for weight modification eventually leads to two stable weights:

```

In[37]:= ListPlot[{ww1list, ww2list}, PlotRange -> {-1, 8}, AxesOrigin -> {0, 0},
  Joined -> True]

```

```
Out[37]=
```



■ Memory recall

We've seen several ways of finding the weights of a matrix that will approximately reproduce an output, given an input it has seen before. Let's try it out.

So in order to "recall" a response, from an input `Indata[[6]]`, we run it through the "network" memory matrix `w1`:

```
In[38]:= w1.Indata[[22]]
```

```
Out[38]= {2.0259}
```

And we can check to see how well it recalls:

```
In[39]:= Outdata[[22]]
```

```
Out[39]= 2.12715
```

One property of the regression model of memory is that it also generalizes. We saw a 3-dimensional example of interpolation in a graphical demonstration in the last lecture. But we can also extrapolate.

For example, $\{11,15\}$ wasn't in the training set, but the expected output is:

```
In[41]:= w1.{11,15}
```

```
Out[41]= {67.5432}
```

The network has "learned" a surface, (e.g. for one particular training set, $\mathbf{z} = 1.9\mathbf{x} + 3.07\mathbf{y}$) through the points specified in the training set.

The main point is that the linear associator will try to fit a plane (or hyperplane) through the data. If the data do not fit that model, then the memory and generalization will not be good.

An obvious generalization is to fit hypersurfaces, rather than hyperplanes. And that is the direction we will head. But first let us look at linear regression from a point of view that you may not have thought of before.

Underconstrained problems and redundancy


This example is very similar to the preceding example, except that we are going to try to learn **2D responses from 1D stimuli**. At first this doesn't seem to make sense, because the mapping from 1D to 2D in general would be expected to be underconstrained by the data. But this is exactly what is needed in certain kinds of inverse problems that are said to be "ill-posed" (Poggio et al., 1985). They are made "well-posed" when there is some underlying regularity (sometimes called "smoothness") in the high-dimensional output space (e.g. Kersten et al., 1987).

The "shape-from-shading" problem in vision can be formulated as a problem of mapping N pixel intensities to N surface normals, where each surface normal is specified by two numbers--thus the mapping goes from N to $2N$. E.g. given a 15×15 grid of intensities, estimate the 15×15 grid of 2D vectors that are normal to each point of the underlying surface. Thus one might start off with 225 numbers as input, but outputs 550 numbers.

```
In[42]:= GradientFieldPlot[f_, {x_, xmin_, xmax_}, {y_, ymin_, ymax_},
  opts : OptionsPattern[]] :=
  VectorPlot[Evaluate[D[f, {{x, y}}]], {x, xmin, xmax}, {y, ymin, ymax},
  opts]
```

```
In[43]:= ssurface[x_, y_] := Sin[x y];
ggss = Plot3D[ssurface[x, y], {x, -4, 4}, {y, -4, 4},
  MeshFunctions -> {#1 &, #3 &}];
gsss = Plot3D[ssurface[x, y], {x, -4, 4}, {y, -4, 4}, ViewPoint -> {0, 0, 30},
  Axes -> False,
  Lighting -> {{"Directional", White, {{1, 0, 1}, {1, 1, 0}}}},
  Mesh -> 15, MeshStyle -> Directive[Thick, Blue]];
gf = GradientFieldPlot[ssurface[x, y], {x, -4, 4}, {y, -4, 4}];
Show[GraphicsRow[{ggss, gsss, gf}]]
```

Out[47]=



The left panel illustrates an underlying surface that generates the image shown in the middle. The shape-from-shading problem tries to take the graylevel values at points (approximated by grid boxes) in the middle panel, and output a 2-D vector for that grid, as shown in the right-hand panel.

Let's try a simple version of learning a mapping from a low to high-dimensional case. We'll use low-dimensional synthetic data whose generation process we'll keep hidden for now.

■ Synthetic data -- try not to look closely...

Here is the generative model:

$$\text{output} = \{x_2, z\} = \{a x_1, b x_1 + \text{noise}\}$$

We want to learn the parameters $\{a, b\}$ from training pairs of inputs and outputs: $\{x_1, \{x_2, z\}\}$

```
In[48]:= r3Dline[a_, b_] :=
  N[Table[{x1 = 1 RandomReal[], a x1, b x1 + 0.5 RandomReal[] - 0.25}, {30}],
  2];
```



```
In[49]:= data = r3Dline[2,3];
```

```
In[60]:= Indata = data[[All, 1]];
Outdata =data[[All, 2 ;; 3]];
```

So the input data is a list of 1D scalars, and the output data a list of 2D vectors. Let's apply the Widrow-Hoff algorithm to learn the relationship between the input stimuli, **Indata**, and the output responses, **Outdata**.

■ Iterative Widrow-Hoff learning


```
In[52]:= i=0; w1 = {{0},{0}}; w1list = {}; w2list = {}; eta = 0.4;
Dimensions[w1]
```

```
Out[53]= {2, 1}
```

```
In[54]:= While[i<Length[data],
  ++i;
  out = {data[[i,2]],data[[i,3]]} ; in = {data[[i,1]]};
  w1 = w1 + eta Outer[Times,(out - w1.in),in];
  w1list = Append[w1list,w1[[1]][[1]]];
  w2list = Append[w2list,w1[[2]][[1]]];
];
w1
```

```
Out[55]=  $\begin{pmatrix} 1.9832 \\ 2.89941 \end{pmatrix}$ 
```

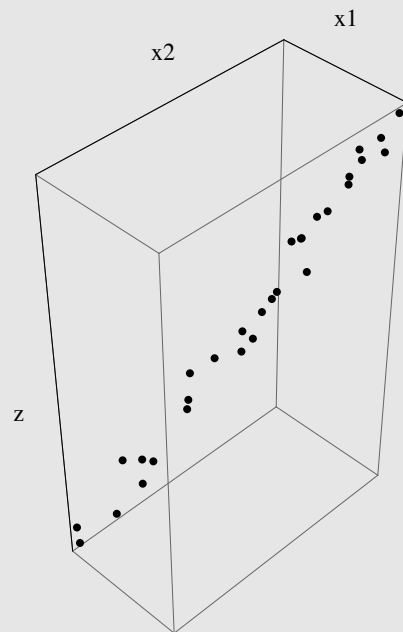
```
In[56]:= ListPlot[{w1list, w2list}, PlotRange -> {0, 6}, AxesOrigin -> {0, 0},  
Joined -> True]
```



■ Visualizing the underlying structure

So why does it work to learn to predict a 2D value from a 1D input? The reason, of course, is that the underlying structure of the output data is highly constrained, and in fact lies close to a straight line in 3-space.

```
Show[Graphics3D[Point /@ data], Axes → True, AxesLabel → {"x1", "x2", "z"},
  Ticks → None]
```



Here is the generative model for our synthetic data:

$x1 \rightarrow \{x2, z\}$,

and given the form is a straightline through the origin with some additive noise, the output $\{x2, z\}$ is given by:

$$\{x2, z\} = \{a x1, b x1 + \text{noise}\}$$

An input, a scalar quantity $x1$ specifies a plane. We've found a line that intersects the plane and for which points in our data set our close to that line. The intersection point gives us the coordinates $\{x2, z\}$ for our regression fit, or "recall".

We learned the parameters $\{a, b\}$, corresponding to the weights, from training pairs of inputs and outputs: $\{x1, \{\{x2, z\}\}\}$

```
r3Dline[a_, b_] :=
  N[Table[{x1 = 1 RandomReal[], a x1, b x1 + 0.5 RandomReal[] - 0.25}, {30}],
    2];
```

```
data = r3Dline[2,3];
```

It produced the coordinates of a noisy line in 3-space.

Making estimates of high-dimensional functions from lower dimensional inputs happens alot, and as mentioned above is common to many so called "early vision" problems (Poggio et al., 1985).

Appendix

■ Pseudoinverse solution

We can calculate the memory matrix using the **PseudoInverse** in this case too:

```
In[107]:= Dimensions[Indata]
          Dimensions[Outdata]
```

```
Out[107]= {30}
```

```
Out[108]= {30, 2}
```

```
In[109]:= matmem = Transpose[PseudoInverse[Transpose[{Indata}]]].Outdata]
```

```
Out[109]=  $\begin{pmatrix} 2. \\ 2.96766 \end{pmatrix}$ 
```

■ Recall

Let's check out a few values to see how well the memory matrix can recall a 2 dimensional output, given a one dimensional input.

```
In[110]:= matmem.Transpose[{Indata}][[5]]  
matmem.Transpose[{Indata}][[13]]  
matmem.Transpose[{Indata}][[28]]
```

```
Out[110]= {0.157826, 0.234186}
```

```
Out[111]= {0.149802, 0.222281}
```

```
Out[112]= {0.422584, 0.627043}
```

```
In[113]:= Outdata[[5]]  
Outdata[[13]]  
Outdata[[28]]
```

```
Out[113]= {0.157826, 0.270793}
```

```
Out[114]= {0.149802, 0.311916}
```

```
Out[115]= {0.422584, 0.769959}
```

References

- Bishop, C. M. (1995). Neural Networks for Pattern Recognition. Oxford: Oxford University Press.
- 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.
- Kersten, D., O'Toole, A. J., Sereno, M. E., Knill, D. C., & Anderson, J. A. (1987). Associative learning of scene parameters from images. *Appl. Opt.*, 26, 4999-5006.
- Knill, D. C., & Kersten, D. (1990). Learning a near-optimal estimator for surface shape from shading. *CVGIP*, 50(1), 75-100.
- Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature*, 317, 314-319.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1), 1-58.