

Introduction to Neural Networks

U. Minn. Psy 5038

Lecture 8

Examples of heteroassociation and autoassociation

Initialization

```
In[1]:= Off[SetDelayed::write]
Off[General::spell1]
SetOptions[ListPlot, Joined -> True];
SetOptions[ArrayPlot, ColorFunction -> "GrayTones", ImageSize->Tiny];
```

Introduction

Last time

■ Outer product

Outer product, $\mathbf{g}\mathbf{f}^T$ models the increase in weight strength when input \mathbf{f} is associated with an output \mathbf{g} :

```
In[5]:= Outer[Times, Array[g, 5], Array[f, 4]] // MatrixForm
```

Out[5]/MatrixForm=

```
( f(1) g(1) f(2) g(1) f(3) g(1) f(4) g(1)
  f(1) g(2) f(2) g(2) f(3) g(2) f(4) g(2)
  f(1) g(3) f(2) g(3) f(3) g(3) f(4) g(3)
  f(1) g(4) f(2) g(4) f(3) g(4) f(4) g(4)
  f(1) g(5) f(2) g(5) f(3) g(5) f(4) g(5) )
```

■ Learning & recall

1. Learning

Let $\{\mathbf{f}_n, \mathbf{g}_n\}$ be a set of input/output activity pairs. Memories are stored by superimposing new weight changes on old ones. Information from many associations is present in *each* connection strength.

$$\mathbf{W}_{n+1} = \mathbf{W}_n + \mathbf{g}_n \mathbf{f}_n^T \quad (1)$$

2. Recall

Let \mathbf{f} be an input possibly associated with output pattern \mathbf{g} . For recall, the neuron acts as a linear summer:

$$\mathbf{g} = \mathbf{W}\mathbf{f} \quad (2)$$

$$g_i = \sum_j w_{ij} f_j \quad (3)$$

3. Condition for perfect recall

If $\{\mathbf{f}_n\}$ are orthonormal, the system shows perfect recall:

$$\begin{aligned} \mathbf{W}_n \mathbf{f}_m &= (\mathbf{g}_1 \mathbf{f}_1^T + \mathbf{g}_2 \mathbf{f}_2^T + \dots + \mathbf{g}_n \mathbf{f}_n^T) \mathbf{f}_m \\ &= \mathbf{g}_1 \mathbf{f}_1^T \mathbf{f}_m + \mathbf{g}_2 \mathbf{f}_2^T \mathbf{f}_m + \dots + \mathbf{g}_m \mathbf{f}_m^T \mathbf{f}_m + \dots + \mathbf{g}_n \mathbf{f}_n^T \mathbf{f}_m \\ &= \mathbf{g}_m \end{aligned} \quad (4)$$

since,

$$\mathbf{f}_n^T \mathbf{f}_m = \begin{cases} 1, & n = m \\ 0, & n \neq m \end{cases} \quad (5)$$

Today

A common distinction in neural networks is between supervised and unsupervised learning ("self-organization"). The heteroassociative network is supervised, in the sense that a "teacher" supplies the proper output to associate with the input. Learning in autoassociative networks is unsupervised in the sense that they just take in inputs, and try to organize a useful internal representation based the inputs. What "useful" means depends on the application. We explore the idea that if memories are stored with autoassociative weights, it is possible to later "recall" the whole pattern after seeing only part of the whole. Later we'll see how heteroassociative learning can be treated as a special case of autoassociative learning.

■ Simulations

Heterassociation

Autoassociation

Superposition and interference


```
In[8]:= Pmatrix = {
  {0, 0, 0, 0, 0, 0, 0, 0, 0, 0},
  {0, 1, 1, 1, 1, 1, 0, 0, 0, 0},
  {0, 1, 0, 0, 0, 0, 1, 0, 0, 0},
  {0, 1, 0, 0, 0, 0, 1, 0, 0, 0},
  {0, 1, 0, 0, 0, 0, 1, 0, 0, 0},
  {0, 1, 1, 1, 1, 1, 0, 0, 0, 0},
  {0, 1, 0, 0, 0, 0, 0, 0, 0, 0},
  {0, 1, 0, 0, 0, 0, 0, 0, 0, 0},
  {0, 1, 0, 0, 0, 0, 0, 0, 0, 0},
  {0, 1, 0, 0, 0, 0, 0, 0, 0, 0},
  {0, 0, 0, 0, 0, 0, 0, 0, 0, 0}};
```

```
In[9]:= Tv = N[Normalize[Flatten[Tmatrix]]];
Iv = N[Normalize[Flatten[Imatrix]]];
Pv = N[Normalize[Flatten[Pmatrix]]];
size = Dimensions[Imatrix][[1]];
maxv = Max[Flatten[Tmatrix]];
maxi = Max[Tv];
```

```
In[15]:= maxi
```

```
Out[15]= 0.267261
```

■ Sidenote: Making images into vectors: Flatten[] and Partition[]

Flatten[] takes a list of lists and turns it into a list of elements, that is, it removes all of the inner braces:

```
In[16]:= Flatten[{{a,b},{c,d}}]
```

```
Out[16]= {a, b, c, d}
```

Partition[] is sort of like an inverse for **Flatten[]** and takes a list of elements and structures it back into a list of lists:


```
In[17]:= Partition[%,2]
```

```
Out[17]=  $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ 
```

For our purposes, **Flatten[]** turns a matrix representing a 2D picture (e.g. the letter I, or T) into a vector that we can store in a weight matrix memory. Later, we use **Partition[]** to turn whatever the matrix remembers into a 2D picture for comparison with the input picture originally learned.


```
In[18]:= ArrayPlot[Tmatrix, PlotRange -> {0, maxv}]
```

```
Out[18]=
```



```
In[19]:= ArrayPlot[Partition[Tv, size], PlotRange -> {0, maxi}, Mesh -> False]
```

```
Out[19]=
```



Learning an association I --> T

■ Learning

Let's use the outer product to represent the change in the synaptic weights caused by the simultaneous activity of **T** and **I** which assuming a Hebb-type rule, is proportional to the product of the activities:

```
In[40]:= weights = Outer[Times, Tv, Iv];
```

```
In[41]:= Dimensions[weights]
```


```
Out[41]= {100, 100}
```

Now if sometime later, the weights matrix is "stimulated" with the letter **I**, it produces as a response the letter **T**:

■ Recall: Remembering T from I

```
In[42]:= response = weights.Iv;
ArrayPlot[Partition[response, size], Mesh -> False]
```

Out[43]=




Exercise

What if Tv was the input? What if a random vector was the input? What response would you expect?

```
In[31]:= response = weights.Tv;
ArrayPlot[Partition[response, size]]
```

Out[32]=



What if you used stimulated the Transpose of weights with Tv ?

Add the association $T \leftrightarrow P$ to the mix

■ Learning P from T, storing it with the association between I and T

```
In[48]:= weights = weights + Outer[Times, Pv, Tv];
```

■ Recall: Stimulate with T: Response? P, I, or a mixture?

```
In[49]:= response = weights.Tv; ArrayPlot[Partition[response, size]]
```

Out[49]=



Exercise

If you look carefully, you can see some evidence for interference. Why might you expect this based on the two inputs **Iv** and **Tv**? Try comparing the dot products of the various inputs.

Can you think of a network modification for recall that might help to reduce the interference?

Exercise

The **ArrayPlot** functions don't always give the best way of seeing variations in a function or list. Try **ListPlot[response]**. What do you notice?

Exercise

The plot function automatically scales the plot range so that white corresponds to the maximum value in the list. Use **Max[Tv]** to find the peak value in a list.

Add the association I <--> P to the mix


■ Learning P & I, store it with the associations between I & T, P & T

```
In[50]:= weights = weights + Outer[Times, Iv, Pv];
```

■ **Recall: Stimulate with P: Response? T, I, or a mixture?**

```
In[51]:= response = weights.Iv; ArrayPlot[Partition[response, size]]
```

Out[51]=



Autoassociation

If $\mathbf{f}=\mathbf{g}$, then we have an autoassociative system. There is only one set of units, and each element potentially connects to each other element. Later we will see how this architecture is used in non-linear networks. Autoassociation stores information about the relationships between the elements or features of a stimulus pattern (vector). We can show how this kind of knowledge can be used to predict or reconstruct missing information. Neural networks of this sort build internal models of the statistical structure of the ensemble they are exposed to.

Reconstructive property

■ **Autoassociation can reconstruct missing parts of a stimulus.**

$$\mathbf{x} = \begin{pmatrix} f_1 \\ f_2 \\ \cdot \\ \cdot \\ f_m \\ g_1 \\ \cdot \\ \cdot \\ \cdot \\ g_n \end{pmatrix} \quad (6)$$

Suppose a whole pattern \mathbf{x} , consists of two parts $\{f_1, f_2, f_3\}$, and $\{g_1, g_2, g_3, g_4\}$, and the association of vector \mathbf{x} with itself is represented by the outer product in matrix \mathbf{W} :


```
In[52]:= f={f1,f2,f3,0,0,0,0};
g={0,0,0,g1,g2,g3,g4};
x=f+g;
W=Outer[Times,x,x];
```

```
In[56]:= x = f + g
```

```
Out[56]= {f1, f2, f3, g1, g2, g3, g4}
```

Sometime later, we input a version of \mathbf{x} , but with "missing" elements--i.e. \mathbf{x} with some elements set to zero--namely, vector \mathbf{f} . What do we get in response?

```
In[57]:= Simplify[W.f]
```

```
Out[57]= {f1 (f1^2 + f2^2 + f3^2), f2 (f1^2 + f2^2 + f3^2), f3 (f1^2 + f2^2 + f3^2),
g1 (f1^2 + f2^2 + f3^2), g2 (f1^2 + f2^2 + f3^2), g3 (f1^2 + f2^2 + f3^2), g4 (f1^2 + f2^2 + f3^2)}
```

So if \mathbf{f} is normalized, each sum of squared elements of \mathbf{f} is equal to 1, and $\mathbf{W.f}$ is equal to \mathbf{x} . In general, the matrix \mathbf{W} restores \mathbf{f} to the pattern \mathbf{x} up to a scale factor α :

```
In[58]:= % /. (f1^2 + f2^2 + f3^2) -> alpha
```

```
Out[58]= {f1 alpha, f2 alpha, f3 alpha, g1 alpha, g2 alpha, g3 alpha, g4 alpha}
```

Exercise

What does it mean to have a "missing" part of a pattern?

■ Autoassociation includes heteroassociation

At first it may seem that an autoassociative system is a more restrictive type of association than heteroassociation. But if we form a new vector \mathbf{f}' in which we stack \mathbf{f} on top of \mathbf{g} , then autoassociation outer product matrix contains within it, the heteroassociation between \mathbf{f} and \mathbf{g} .

Question:

What can you say about the eigenvectors of the weight matrix from autoassociative learning? Take a look at : Outer[Times,Array[ff,5],Array[ff,4]]//MatrixForm

Autoassociative example with TIP pictures

- Learn about T, learn about I, and store the associations together by superimposing their weight matrices


```
In[78]:= Clear[weights];
weights = Outer[Times, Tv, Tv] + Outer[Times, Iv, Iv];
```

- Sometime later, stimulate the network with an impoverished T, missing some bits

Let's delete the 2nd row of T:

```
In[80]:= forgettingTmatrix = Partition[Tv, size];
forgettingTmatrix[[2]] = Table[0, {10}];
forgettingT = Flatten[forgettingTmatrix];
ArrayPlot[Partition[forgettingT, size]]
```


Out[83]=



- Recall of the original T, from the missing bits

```
In[84]:= rememberingT = weights.forgettingT;
ArrayPlot[Partition[rememberingT, size]]
```

Out[84]=




Interference: Corrupt T again, this time with some other random bits missing

Let's do something a little more drastic to **T**. We'll randomly "delete" pixels of the picture:


```
In[85]:= pepper = RandomInteger[1, Dimensions[Tv][[1]]];  
peppermatrix = DiagonalMatrix[pepper]; forgettingT = peppermatrix.Tv;  
ArrayPlot[Partition[forgettingT, size]]
```

Out[85]=



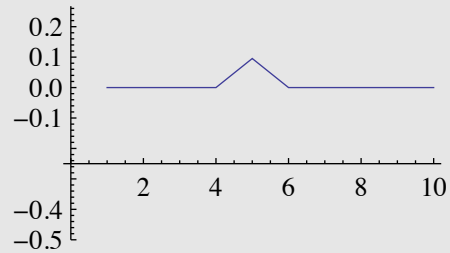
```
In[86]:= rememberingT = weights.forgettingT;  
ArrayPlot[Partition[rememberingT, size]]
```

Out[86]=



```
In[87]:= ListPlot[Partition[rememberingT, size][[5]], AxesOrigin -> {0, -0.25},
  PlotRange -> {-0.5, maxi}, Joined -> True]
```

Out[87]=



- Recall that you can get information about the options as well as the functions with a ?? query:

```
??AxesOrigin
```


AxesOrigin is an option for graphics functions which specifies where any axes drawn should cross. >

```
Attributes[AxesOrigin] = {Protected}
```

- **Interference: Corrupt T, with added noise**

```
In[88]:= forgettingT = Tv; forgettingT[[59]] = 0.27;
  ArrayPlot[Partition[forgettingT, size]]
```

Out[88]=



```
In[89]:= rememberingT = weights.forgettingT;
ArrayPlot[Partition[rememberingT, size]]
```

Out[89]=



Although the memory looks pretty good, it is not perfect because although **T_v** and **I_v** were almost orthogonal, with a cosine of about .09, they were not perfectly orthogonal. In fact, we can get a measure of how close **rememberingT** is to **T_v** in terms of the cosine of the angle between them:

```
In[90]:= Tv.Normalize[rememberingT]
```

Out[90]= 0.995682


■ Interference with more autoassociations: I, T, and now P too

If we have the connection matrix, weights store another letter, **P**, then we will begin to get even more interference when we try to recall **T** from a fragment of **T**:

```
In[91]:= weights = weights +
          Outer[Times, Pv, Pv];
```

```
In[92]:= rememberingT = weights.forgettingT;
ArrayPlot[Partition[rememberingT, size]]
```

Out[92]=



This is because the patterns we've stored are not mutually orthogonal, and in particular, **P** is too close to **I** and **T**:

Map the function over the list

This means that to apply `step[]` to `rememberingT`, you would have to use the `Map[]` function:

The `Map[]` function is used often enough, that *Mathematica* has a short-hand:

```
In[98]:= Clear[f, a, b, c];  
Map[f, {a, b, c}]
```

```
Out[99]= {f(a), f(b), f(c)}
```

```
In[100]:= f /@ {a, b, c}
```

```
Out[100]= {f(a), f(b), f(c)}
```

If the function has multiple arguments, then use `#` to identify which function slots get the variable, with the function end defined by `&`:

```
In[101]:= Map[step[#, 0] &, rememberingT];
```

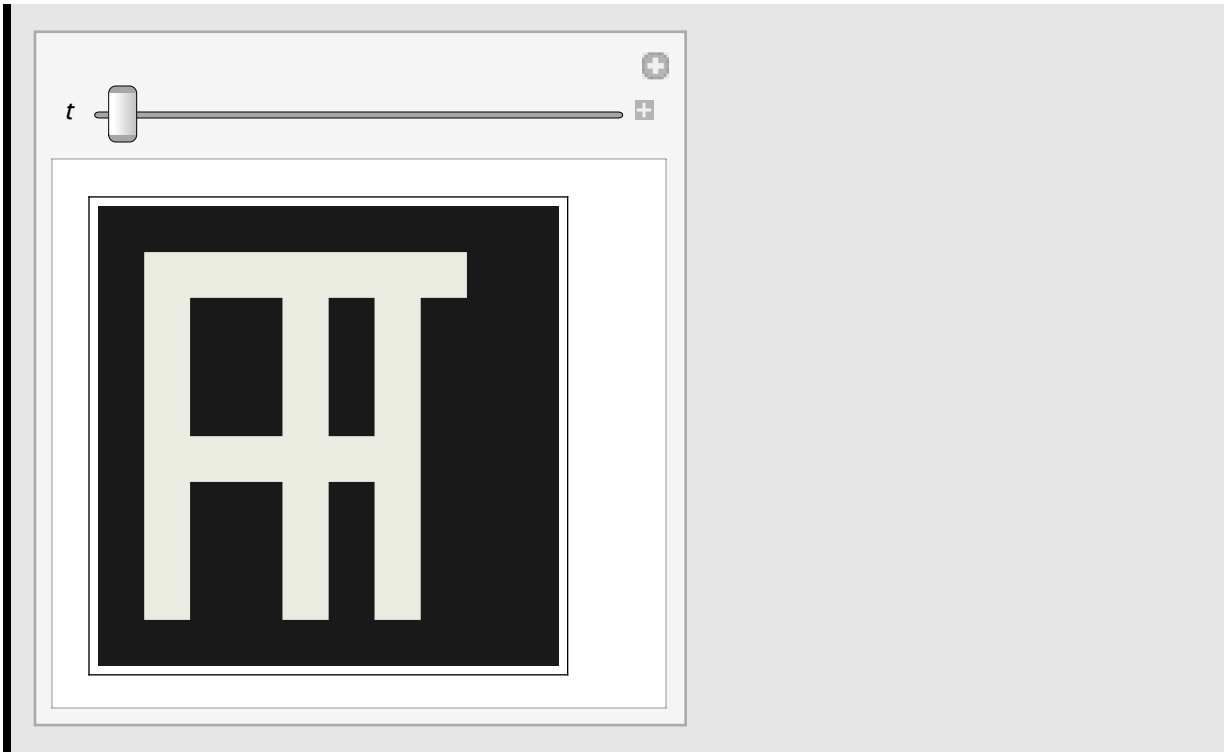
Using the short-hand version:

```
In[102]:= step[#, 0] & /@ rememberingT;
```

We'll put it all in one line, together with a slider to adjust the threshold:

```
In[103]:= Manipulate[ArrayPlot[Partition[step[#, t] & /@ rememberingT, size],
  ImageSize -> Small], {t, 0, 1}]
```

```
Out[103]=
```



The threshold choice is clearly important. How could you make the threshold automatic?

Autoassociation with pictures

Autoassociation with some other patterns: Einstein or Shannon


Get einstein64x64.jpg:

```
ieinstein = Import[SystemDialogInput["FileOpen"], "Data"];
```

```
size2 = Dimensions[ieinstein][[1]];
```



```
geinstein = ArrayPlot[ieinstein, ColorFunction -> GrayLevel]
einstein = N[Normalize[Flatten[ieinstein]]];
```



Get shannon64x64.jpg:

```
ishannon = Import[SystemDialogInput["FileOpen"], "Data"];
```

```
gshannon = ArrayPlot[ishannon, ColorFunction -> GrayLevel]
shannon = N[Normalize[Flatten[ishannon]]];
```



Autoassociatively store einstein

```
weights = Outer[Times, einstein, einstein];
```

```
Dimensions[einstein]
```

```
{4096}
```

How big is weights?

Make an einstein picture with some "salt and pepper noise"--random intensities at random locations, **einstein64x64missing.jpg**:

```

forgettingeinstein = einstein;
Table[forgettingeinstein[[RandomInteger[{1, 4096}]]] =
  RandomReal[{0.0, Max[einstein]}], {1000}];
gforgettingeinstein = ArrayPlot[Partition[forgettingeinstein, size2],
  ColorFunction -> GrayLevel]

```



Recall einstein, with only a part of the original pixels in einstein, i.e. with forgettingeinstein as input:

```

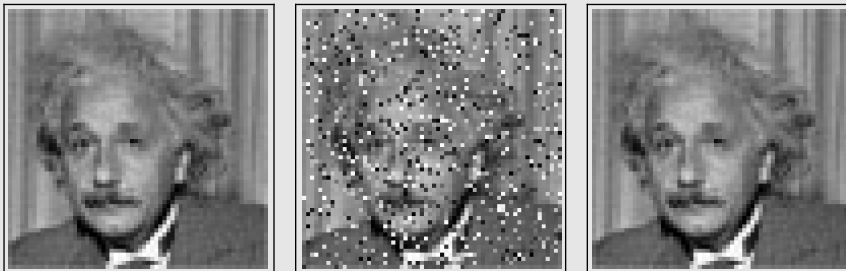
rememberingeinstein = weights.forgettingeinstein;
grememberingeinstein = ArrayPlot[Partition[rememberingeinstein,size2], Color

```

```

Show[GraphicsRow[{geinstein, gforgettingeinstein, grememberingeinstein}]]

```



Now superimpose the outerproduct weight matrices for einstein and shannon:

```

deltaweights = Outer[Times,shannon,shannon];

```

```

weights = weights + deltaweights;

```

```

rememberingeinstein = weights.forgettingeinstein;
grememberingeinstein = ArrayPlot[Partition[rememberingeinstein,size2], Color

```

```
Show[GraphicsRow[{geinstein, gshannon, gforgettingeinstein,  
grememberingeinstein}]]
```



Next time

- Explore a simple application, summed vector memories.
- Neural networks as statistical pattern processing

Then we are going to step back in a sense, and ask ourselves what these networks are doing in the sense of information or pattern processing. This will lead naturally to a statistical framework for understanding both heteroassociative and autoassociative networks.

Heteroassociation will lead to regression.

Summed vector memories will lead to a discussion of simple ("first-order") statistical learning.

Autoassociation will lead to second-order statistical learning.

A deeper understanding of the information processing roles of 1) our learning rule; and 2) our recall mechanism, helps in two ways. First, we will have a better idea of what we need to do to get the network to solve a given problem. Second, we may discover that there is a different problem for which it is better suited. So linear heteroassociation recall may be a poor classification model, but it might be a good interpolation model. Linear autoassociation learning may be a poor way to store information about specific memories, but may be a good way of learning about the statistics of ensembles.

As a preview of the latter, consider another application of autoassociation learning, where the goal is not to remember a particular previous input, but rather to learn something about an ensemble of inputs that all belong to the same class. Practical examples are collections of networks that can learn about their own special environments. For example, you want to build an automated driving system. The problem is complex, in part, due to different kinds of demands placed by the driving environment. So you decide you need three "experts": one for single-lane country roads, another for two-lane highways, and another for four-lane divided highways. Now you put them all in one vehicle and let them all monitor their sensory inputs as the vehicle is being driven (by one of them). When given a new environment, these experts "compare notes" to see how well this new environment fits their internal models or domain of expertise. Which ever expert "knows" the new environment the best, gets to drive the car. This is related to the idea of mental modules. The part of the driving expert that validates the environment could be realized by an autoassociative network. It can do this by testing how well its prediction of the environment's input to its sensors fits the actual sensor measurements.

Improve learning or recall?

We've seen that the linear associator has problems of interference. How do we improve the network? We have two general strategies: 1) improve the learning, so that linear recall will do better; 2) improve the recall, so that a simple Hebbian outer product rule can still be used.

Later we will derive a new learning rule that builds better association matrices for certain problems like interpolation, regression, and generalization. And we will look at non-linear recall mechanisms that make cleaner classifications for memory problems.

For specific example, the outerproduct learning rule can be seen as part of a computation that computes autocovariance matrices. These are useful because the input ensemble in some sense "lives" in the space defined by the eigenvectors with the largest eigenvalues. The problem is that to project input vectors into this space using simple matrix multiplication requires one to have the eigenvector matrix, NOT the autocovariance matrix. So we have two possibilities. First, we can find an alternative learning rule that will give us the eigenvector matrix. Or we can look for a different recall rule that will give us the projection we want.

© 1998, 2001, 2005, 2007, 2009 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.