

# Introduction to Neural Networks

## U. Minn. Psy 5038

Daniel Kersten

---

## Introduction

### Last time

Modeled aspects of spatial lateral inhibition in the visual system:

- Simple linear network: intensity  $\rightarrow$  spatial filter  $\rightarrow$  response

- intensity and response were represented as vectors

- the spatial filter was represented as a matrix

If the spatial filter is a neural network, then the elements of the matrix can be interpreted as strengths of synaptic connections. The idea is to represent the synaptic weights for the  $i^{\text{th}}$  output neuron by the values in the  $i^{\text{th}}$  row of the matrix. The matrix of weights is sometimes called a connection matrix just because non-zero weights are between connected neurons, and neurons that aren't connected get fixed zero weights.

This model is what we called a linear feedforward network, and used the generic neuron model.

We generalized the linear discrete model to a linear continuous one with feedback. Although more complex, its steady-state solution is identical to the linear feedforward model.

We studied the continuous system by approximating it as a discrete time system with  $\epsilon = \Delta t$  as a free parameter.

## Today

There is a large body of mathematical results on linear algebra and matrices, and it is worth our while to spend some time going over some of the basics of matrix manipulation.

We will first review matrix arithmetic (addition and multiplication). We will review the analog of division, namely finding the inverse of a matrix--something that was used in Lecture 5 to show how the steady-state solution of the feed-back model of lateral inhibition was equivalent to a feedforward model with the appropriate weights.

You may wonder at times how all this is used in neural modeling. But as this course goes on, we will see how otherwise obscure notions of things like an "outer product" between two vectors, or the "eigenvectors" of a matrix are meaningful for neural networks. For example, the "outer product" between two vectors can be used in modeling learning, and the eigenvectors corresponding to a "matrix of memories" can represent stored prototypes, and in dimensionality reduction.

---

## Basic matrix arithmetic

### Definition of a matrix: a list of scalar lists

Traditional and Standard form output format defaults.

#### ■ Defining arrays or lists of function outputs using indices

As we've already seen, `Table[ ]` can be used to generate a matrix, or list of lists. E.g.

```
In[17]:= H = Table[i^2 + j^2, {i, 1, 3}, {j, 1, 3}]
```

Before going on, use the Cell menu to set: **Cell->Default Output FormatType->TraditionalForm**. This will cause vectors and matrices to be printed out in traditional mathematics format. Then we don't have to explicitly type `MatrixForm` every time to see the output in traditional matrix form. But if you want to see the output in standard *Mathematica* form, use `//StandardForm` or `StandardForm[ ]`.

```
In[18]:= StandardForm[H]
```

## ■ Defining arrays or lists of symbols

Although most of the time, we'll be working with numerical matrices, *Mathematica* also allows one to specify arrays or lists of variables.

An  $m \times n$  matrix has  $m$  rows, and  $n$  columns. Here is a  $3 \times 4$  matrix of symbolic elements  $w[i,j]$ :

```
In[19]:= Array[w, {3, 4}]
```

You can produce the same array with: `Table[w[i,j],{i,1,3},{j,1,4}]`. You can also have *Mathematica* display a subscripted list

```
In[20]:= L = Table[wi,j, {i, 1, 3}, {j, 1, 4}]
```

*Mathematica* lists can have elements of almost any type:

```
In[21]:= w2,3 = 3.0;
w3,4 = Pi;
w2,4 = "hi";
w1,1 = 3;
```

```
In[25]:= L
```

## Adding, subtracting and multiplying by a scalar

As with vectors, matrices are added, subtracted, and multiplied by a scalar component by component:

```
In[26]:= A = {{a,b},{c,d}};
B = {{x,y},{u,v}};
```

Add **A** to **B**:

```
In[28]:= A+B
```

Subtract **B** from **A**:

```
In[29]:= A-B
```

Multiply **A** by 3:

```
In[30]:= 3 A
```

## Multiplying two matrices

We have already seen how to multiply a vector by a matrix: we replace the  $i^{\text{th}}$  row (i.e. element) of the output vector by the inner product of the  $i^{\text{th}}$  row of the matrix with the vector.

In order to multiply a matrix  $\mathbf{A}$ , by another matrix  $\mathbf{B}$  to get  $\mathbf{C} = \mathbf{AB}$ , we calculate the  $ij^{\text{th}}$  component of the output matrix by taking the inner product of the  $i^{\text{th}}$  row of  $\mathbf{A}$  with the  $j^{\text{th}}$  column of  $\mathbf{B}$ :

```
In[33]:= A . B
```

Note that  $\mathbf{AB}$  is not necessarily equal to  $\mathbf{BA}$ :

```
In[34]:= B . A
```

## Laws of commutation, association and distribution

In particular, look at the element in the upper left of the matrix  $\mathbf{BA}$  above--there is no reason, in general, for  $\mathbf{ax+bu}$  to equal  $\mathbf{ax + cy}$ . That is, matrix multiplication does not *commute*.

Apart from commutation for matrix multiplication, the usual laws of commutation, association, and distribution that hold for scalars hold for matrices. Matrix addition and subtraction do commute. Matrix multiplication is associative, so  $(\mathbf{AB})\mathbf{C} = \mathbf{A}(\mathbf{BC})$ . The distributive law works too:  $\mathbf{A}(\mathbf{B+C}) = \mathbf{AB} + \mathbf{AC}$

## Non-square matrices

It is not necessary for  $\mathbf{A}$  and  $\mathbf{B}$  to be square matrices (i.e. have the same number of rows as columns) to multiply them. But if  $\mathbf{A}$  is an  $m \times n$  matrix, then  $\mathbf{B}$  has to be an  $n \times p$  matrix in order for  $\mathbf{AB}$  to make sense. In other words, the left matrix ( $\mathbf{A}$ ) has to have the same number of columns as the right matrix ( $\mathbf{B}$ ) has rows.

For example, here  $\mathbf{F}$  is a  $3 \times 2$  matrix, and  $\mathbf{G}$  is a  $2 \times 4$  matrix.

```
In[82]:= F = {{a,b},{c,d},{e,f}}
          G = {{p,q,r,s},{t,u,v,w}}
```

Check the dimensions with

```
In[37]:= Dimensions[F]
          Dimensions[G]
```

Because **F** has 2 columns, and **G** has 2 rows, it makes sense to multiply **G** by **F**:

```
In[85]:= Print[F, G, " =", F.G]
```

However, because the number of columns of **G** (4) do not match the number of rows of **F** (3), **G.F** is not well-defined:

Try this:

```
In[40]:= Print[F, G, " =", G.F]
```

---

## Inverse of a Matrix

### "Dividing" a matrix by a matrix: the identity matrix & matrix inverses

The matrix corresponding to 1 or unity is the **identity matrix**. Like 1, the identity matrix is fundamental enough, that *Mathematica* provides a special function to generate n-dimensional identity matrices. Here is a 2x2:

```
In[41]:= IdentityMatrix[2]
```

It is easy to show that the identity matrix plays the role for matrix arithmetic that the scalar 1 plays for scalar arithmetic.

How can one divide one matrix, say **B**, by another, say **A**? We can divide numbers, x by y, by multiplying x times the inverse of y, i.e. 1/y. So to do the equivalent of dividing **B** by **A**, we need to find a matrix **Q** such that when **A** is multiplied by **Q**, we get the matrix equivalent of unity, i.e. the identity matrix. Then "**B/A**" can be achieved by calculating the matrix product: **B.Q**.

```
In[43]:= A = {{a,b},{c,d}};
```

*Mathematica* provides a built-in function to compute matrix inverses:

```
In[44]:= Q = Inverse[A]
```

We can test to see whether the product of a **A** and **Q** is the identity matrix, but *Mathematica* won't go through the work of simplifying the algebra in this case, unless we specifically ask it to.

```
In[45]:= Q.A
```

```
In[46]:= Simplify[Q.A]
```

Here is a simple numerical example. Verify that **R** is the inverse of **B** by seeing whether the product is the Identity matrix.

```
In[47]:= B = {{1,-1},{3,2}};
R = Inverse[B]
```

### Singular and badly conditioned matrices

You might have noticed that elements of the inverse matrix could have zero denominators if we happened to have unfortunate values of a,b,c,d. In general, the inverse of a matrix is not guaranteed to exist. Let's look at the conditions for which the inverse is not defined.

What if one row is a scaled version of another row?

```
In[50]:= B1 = {{1.5,1},{3,2.0}}
```

Then the rows of the matrix are not linearly independent. In this case, the inverse is not defined and is said to be *singular*.

### Ask for the inverse of B1

```
In[51]:= Inverse[B1]
```

Sometimes the rows are almost, but not quite, linearly dependent (because the elements are represented as approximate floating point approximations to the actual values). *Mathematica* may warn you that the matrix is badly conditioned. *Mathematica* may try to find a solution to the inverse, but you should be suspicious of the solution. In general, one has to be careful of badly conditioned matrices.

Try finding the inverse of the following matrix:

```
In[52]:= B2 = {{-2,-1},{4.00000000000001,2.0}};
Inverse[B2]
```

### Were we lucky or not?

```
In[54]:= B2.Inverse[B2]
```

## Determinant of a matrix

There is a scalar function of a matrix called the determinant. If a matrix has an inverse, then its determinant is non-zero. Does **B2** have an inverse?

```
In[55]:= Det [ B2 ]
```

Yes, but it is suspiciously small.

```
In[56]:= Det [ B ]
```

---

## Matrix transpose

### Transpose definition

We will use the transpose operation quite a bit in this course. It interchanges the rows and columns of a matrix:

```
In[57]:= Y = Table [ yi,j, { i, 1, 3 }, { j, 1, 4 } ]
```

In StandardForm on the input line, the transpose is written:

```
In[58]:= Transpose [ Y ]
```

The output default is TraditionalForm. On the input line, in TraditionalForm, the transpose is written:

```
In[62]:= YT
StandardForm [% ]
```

### Getting parts of a matrix

We can pull out the  $i^{\text{th}}$  row of a matrix in *Mathematica* by simply writing  $W[[i]]$ . For example, the 2nd row of **Y** is:

```
In[65]:= Y [ [ 2 ] ]
```

What about  $i^{\text{th}}$  column of a matrix? There is no equally simple way of getting the column of a matrix in *Mathematica*, but we can use the transpose operation to do it. `Transpose[W][[i]]` produces the  $i$ th column of matrix  $\mathbf{W}$ . For example, the 3rd column of  $\mathbf{Y}$  is:

```
In[66]:= Transpose[Y][[3]]
```

Try putting `//ColumnForm` after the above expression.

**Pull out the element  $y_{2,3}$  from  $\mathbf{Y}$**

---

## Symmetric matrices

For a square matrix, the diagonal elements remain the same under transpose.

If the transpose of a matrix equals itself,  $\mathbf{H}^T = \mathbf{H}$ ,  $\mathbf{H}$  is said to be a **symmetric matrix**. Symmetric matrices occur quite often in physical systems (e.g. the force on particle  $i$  by particle  $j$  is equal to the force of  $j$  on  $i$ ). This means that the elements of a symmetric matrix  $\mathbf{H}$  look like they are reflected about the diagonal. We constructed a symmetric matrix in Lecture 5 when we modeled the exponential drop-off in lateral inhibition:

```
In[67]:= H = Table[N[Exp[-Abs[i-j]]], 1], {i, 5}, {j, 5}]
```

The notion of a Hermitian matrix is a generalization of symmetric, where  $\mathbf{H}^* = \mathbf{H}$ , and  $\mathbf{H}^*$  is the transpose of the complex conjugate of  $\mathbf{H}$ .

## Neural networks and symmetric connections

Do neural systems have symmetric connections? Most real neural networks probably do not. Although, when the nature of the processing would not be expected to favor a particular asymmetry, we might expect that there should be symmetric connections on average. We made this assumption when setting up our lateral inhibition weight matrix, as seen above for  $\mathbf{H}$ . Symmetric matrices have so many nice properties, that they are appealing to neural modelers on theoretical grounds. We'll see this later when we study the Hopfield networks. Lack of symmetry can have profound effects on the dynamics of non-linear networks and can produce chaotic trajectories of the state vector.



## Outer product of two vectors

We've already seen that the inner product of two vectors produces a scalar. In the next lecture, we will begin our discussion of Hebbian learning in neural networks. In this context, the **outer product** of an input and output vector will be used to model synaptic modification. Consider two vectors:

```
In[68]:= fv = Table[fi, {i, 1, 3}]
gv = Table[gi, {i, 1, 3}]
```

The outer product is just all of the pairwise products of the elements of **f** and **g** arranged in a nice (and special) order:

```
In[70]:= h = Outer[Times, fv, gv]
```

To get a preview, the idea in Hebbian learning is that an adjustment to the strength of a connection between say the *i*th and *j*th neurons is proportional to the product of the strength of pre- and post-synaptic activities, represented by  $g_i$  and  $f_j$ .

The outer product is also written in traditional row and column format as:  $\mathbf{f} \mathbf{g}^T$

$$\mathbf{f} \mathbf{g}^T = \begin{pmatrix} f_1 \\ f_2 \\ f_3 \end{pmatrix} (g_1 \ g_2 \ g_3) \quad (1)$$

We'll also use outer product to calculate autocorrelation matrices.

## Eigenvectors and eigenvalues

### Basic idea

Eigenvalues and eigenvectors crop up in many ways. For example, the solution to a set of linear first-order (i.e. no derivative orders bigger than 1) differential equations can be found in terms of eigenvalues and eigenvectors. For a simplified version of the limulus equations

$$\frac{d\mathbf{f}}{dt} = \mathbf{A} \cdot \mathbf{f} \quad (2)$$

the solution is determined by the eigenvalues and eigenvectors of **A** and the initial conditions.

Eigenvectors also crop up in statistics and related neural network algorithms. For example, principal components analysis (PCA) is used in dimensionality reduction--also important in neural networks and "natural computation".

So what are eigenvalues, eigenvectors? An eigenvector,  $\mathbf{x}$ , of a matrix,  $\mathbf{A}$ , is a vector that when you multiply it by  $\mathbf{A}$ , you get an output vector that points in the same (or opposite, depending on the sign of  $\lambda$ ) direction as  $\mathbf{x}$ :

$$\mathbf{Ax} = \lambda\mathbf{x}$$

where  $\lambda$  is an eigenvalue--a scalar that adjusts the length change of  $\mathbf{x}$ . The eigenvalues are found by solving:

$$\mathbf{Det}[\mathbf{A} - \lambda\mathbf{Id}] = 0, \text{ for } \lambda \text{ where } \mathbf{Id} \text{ is the identity matrix.}$$

Let's find the eigenvalues of:

```
In[71]:= A = {{1, 2}, {3, 4}};
```

```
In[72]:= Id = IdentityMatrix[2]
```

```
In[73]:= Solve[Det[A - λ * Id] == 0, λ]
```

Now pick one of the solutions for  $\lambda$ , find the vector  $\mathbf{x} = \{x_1, x_2\}$  that satisfies the basic definition,  $\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$

```
In[74]:= Solve[{A . {x1, x2} == 1/2 (5 - sqrt(33)) * {{1, 0}, {0, 1}} . {x1, x2}, {x1, x2}]
```

So any multiple of

$$\mathbf{x} = \left\{ 1, -\frac{1}{6} (\sqrt{33} + 3) \right\}$$

is a valid eigenvector.

From algebra, because the  $\mathbf{Det}[\ ] = 0$  is a polynomial of order  $n$ , there can't be any more than  $n$  distinct eigenvectors for an  $n \times n$  matrix--and there may be less.

We can similarly find the second distinct eigenvector for  $\lambda \rightarrow \frac{1}{2} (5 + \sqrt{33})$ , but *Mathematica* provides a simpler way of doing things.

## Built-in functions for finding eigenvalues and eigenvectors

### ■ Eigenvalues[ ]

The eigenvalues are given by:

```
In[75]:= Eigenvalues[A]
```

Eigenvalues and eigenvector elements do not have to be real numbers. They can be complex, that is an element can be the sum of a real and imaginary number. Imaginary numbers are represented by multiples (or fractions) of  $\mathbf{I}$ , the square root of  $-1$ :

```
In[76]:= Sqrt[-1]
          Sqrt[-1]//OutputForm
```

```
Out[76]= i
```

```
Out[77]//OutputForm=
```

```
I
```

```
In[78]:= B = {{1,2},{-3,4}};
          Eigenvalues[B]
```

### ■ Eigenvectors[ ]

The built-in *Mathematica* function **Eigenvectors[A]** returns the eigenvectors of matrix **A** as the rows of a matrix, lets call **eig**:

```
In[80]:= eig = Eigenvectors[A]
```

Verify that **eig[[1]]** and **A.eig[[1]]** lie along the line (i.e. in the same or opposite directions) by taking the dot product of the unit vectors pointing in the directions of each. Use this function:

```
In[81]:= normalize[x_] := N[x/Sqrt[x.x]];
```

## Preview of eigenvectors and dimensionality reduction

A measure of the structure of an ensemble of signals (e.g. of vectors) is the degree of correlation between their elements. Signals, such as sounds and visual images, have correlational structure that is taken advantage of in sound and image compression algorithms. One simple kind of statistical structure is characterized by the degree to which one can predict one element of a vector from a nearby element. For images, the color of a pixel at location  $i$  is a pretty good predictor of the color at location  $j = i+1$ . As  $j$  gets far from  $i$ , however, the prediction gets worse. It is possible to characterize the degree of predictability by a matrix whose  $ij^{\text{th}}$  element is big if the  $i^{\text{th}}$  pixel of an image is a good predictor of the  $j^{\text{th}}$ .

Natural sounds also have a similar correlational structure. One measure of predictability is the autocorrelation matrix where, for example, the  $i,j^{\text{th}}$  element represents the degree of correlation between the  $i^{\text{th}}$  and  $j^{\text{th}}$  pixels, i.e. that are  $|j - i|$  apart. An autocorrelation matrix has real elements and is symmetric. The eigenvectors of the matrix capture the subspace with most of the variance, i.e. where the "action" is. The eigenvalues correspond to the amount of "action"--i.e. the variance in the eigenvector directions.

**Verify that the eigenvectors of the symmetric matrix above (H) are orthogonal.**

---

---

## Next time

### Linear systems analysis & introduction to learning/memory

In the next lecture, we'll apply what we've learned about matrices and vectors to two problems. We'll first take an in-depth look at the properties of the linear neural network to see how it relates to general linear systems. We'll also show the advantages of using the eigenvectors of the linear transformation matrix as a useful coordinate system to represent input and output pattern vectors.

We will introduce the topic of learning and memory, and see how the outer product rule can be used to model associative learning.

## Preview of linear systems analysis

Linear systems are an important, and tractable class of input/output models. Matrix operations on vectors provide the prototype linear system.

Consider the generic 2-layer network. It consists of a weighted average of the inputs (stage 1), followed by a point-nonlinearity (the squash function of stage 2), and added noise (stage 3). Although later we will see how the non-linearity enables computations that are not possible without it, useful functions can be realized with just the linear or stage 1 part of the network. We've seen one application already with the model of the limulus eye. In the next lecture, we will see how linear networks can be used to model associative memory. But first, let us take what we've learned so far about modeling linear networks and look at in the general context of linear systems theory. Our 2-layer net is a matrix of weights that operates on a vector of input activities by computing a weighted sum. One property of such a system is that it satisfies the fundamental definition of a "linear system", which we will define shortly.

The world of input/output systems can be divided up into linear and non-linear systems. Linear systems are nice because the mathematics that describes them is not only well-known, but also has a mature elegance. On the other hand, it is a fair statement to say that most real-world systems are not linear, and thus hard to analyze...but fascinating if for that reason alone. Scientists were lucky with the limulus eye. That nature is usually non-linear doesn't mean one shouldn't familiarize oneself with the basics of linear system theory. Many times non-linear systems can be approximated by a linear one over some restricted range of parameter values.

So what is a "linear system"?

The notion of a "linear system" is a generalization of the input/output properties of a straight line passing through zero. The matrix equation  $\mathbf{W}\cdot\mathbf{x} = \mathbf{y}$  is a linear system. This means that if  $\mathbf{W}$  is a matrix,  $\mathbf{x1}$  and  $\mathbf{x2}$  are vectors, and  $a$  and  $b$  are scalars:

$$\mathbf{W}\cdot(a \mathbf{x1} + b \mathbf{x2}) = a \mathbf{W}\cdot\mathbf{x1} + b \mathbf{W}\cdot\mathbf{x2}$$

This is a consequence of the laws of matrix algebra. The idea of a linear system has been generalized beyond matrix algebra. Imagine we have a box that takes inputs such as  $f$ , and outputs  $g = T[f]$ .

The abstract definition of a linear system is that it satisfies:

$$T[a f + b g] = a T[f] + b T[g]$$

where  $T$  is the transformation that takes the sum of scaled inputs  $f$ ,  $g$  (which can be functions or vectors) to the sum of the scaled transformation of  $f$  and  $g$ . The property, that the output of a sum is the sum of the outputs, is sometimes known as the *superposition principle* for linear systems. The fact that linear systems show superposition is good for doing theory, but as we will see later, it limits the kind of computations that can be done with linear systems, and thus with linear neural network models.

## References

Linear Algebra and Its Applications by Gilbert Strang, 3rd Edition , Hardcover, 505 pages, Published by Hbj College & School Div. Publication date: February 1988

© 1998, 2001, 2003, 2005 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.