

Introduction to Neural Networks

U. Minn. Psy 5038

Daniel Kersten

Lecture 1-Introduction

Goal

Understand the functioning of the brain as a computational device. Tools to explain brain & behavior. Relation to connectionist, neuromorphic, computational neuroscience research, and cognitive science.

Relation to Cognitive Science

Cognitive Science: The interdisciplinary study of the acquisition, storage, retrieval and utilization of knowledge.

Problems: perception, learning, memory, planning, action

Often, we don't know how to solve a problem even in principle. For others, we have solutions, but they don't resemble how a biological system might solve the problem.

What kinds of problems can large interconnected systems of model neurons solve? What are the limitations? What are the strengths?

How do neural networks relate to the larger field of statistical pattern recognition?

Understanding the relation between brain and behavior requires...

A multidisciplinary approach

Multiple levels of explanation.

Multidisciplinary approach

Three primary areas or disciplines influence current neural network research:

Neuroscience, computational neuroscience

Understand the basic building blocks or "hardware" of the nervous system

these are: nerve cells or neurons, and their connections, the synapses

Our emphasis is on: large scale neural networks. Requires great simplification in the model of the

neuron...in order to compute and theorize about what large numbers of them can do.

Compare with other areas of Computational Neuroscience that emphasize the biology. Here we emphasize "brain-style"

computation. Often wrong in detail, but driven by a curiosity about how the complex processes of

perception, and memory work.

What can these large scale neural systems do? That is, what can they compute? And how?

■ Computational theory, mathematics, statistical pattern recognition

Statistical inference, engineering (information and communication theory), statistical physics and computer science.

Provide the tools and analogs to abstract and formalize for analysis and simulation.

One of the characteristics of this course is to try to relate the neural models to statistical methods of inference and regression in order to understand the computational principles and power behind a neural implementation.

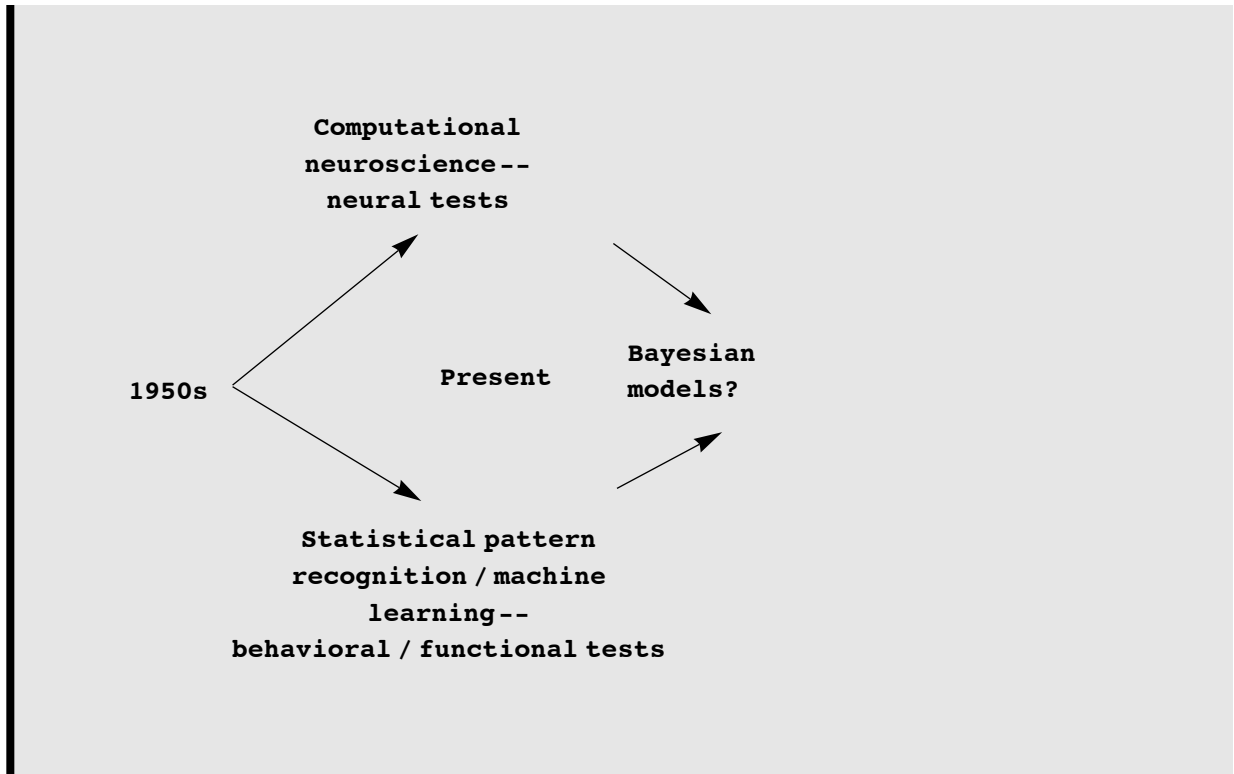
What should these large scale neural systems compute? What are the ways in which information is represented? How can a system be designed to get from input to output representations?

■ Behavioral sciences, psychology, cognitive science and ethology

Understand what subsystems are supposed to do as a functioning organism in the environment.

Psychology & Computational theory =>The brain is NOT a general purpose computer.

■ How do different theoretical neural network approaches relate? A little history and future...



Multiple levels of explanation

■ **Functional/Behavioral level**

Psychology/Cognitive Science/Ethology tells us what is actually solved by functioning behaving organisms. Descriptions of behavior.

■ Statistical Inference level

Theories of pattern recognition, inference & estimation.

Functionalities supported by neural network computing provide a useful way of categorizing models in terms of the computational tasks required:

1. Learning input/output mappings from examples (learning as regression, classification boundaries)

associative memory (->neural networks: Hopfield net, back-prop, local minima can be useful.)

2. Inferring outputs from inputs (continuous estimation, discrete classification)

memory recall, perceptual inference

optimization or constraint satisfaction (->neural networks: Hopfield net, Boltzmann machine, global minimum is desired, local minima are problems)

3. Modeling data (learning as probability density estimation)

self-organization of sensory data into useful representations or classes (e.g principal components analysis, clustering)

(Can view 1. Learning input/output mappings as a special case)

■ Neural network level: Algorithms, implementation

Algorithms: Mathematics of computation tells us what is computable and how. Practical limits. Parallel vs. serial.

Input and output representation, and algorithms for getting from input to output. Programming rules, data structures.


Implementation: Wetware, hardware.

Neuroscience, neurophysiology and anatomy tell us the adequacies and inadequacies of our modeling assumptions.

■ Emphasis in this course

Understand high-level functions such as vision, pattern recognition, learning, memory, inference and control. In the brain these functions involve large-scale systems each with many "modules" and 10s to 100s of thousands of neurons in each. Appropriate level tends towards more abstract where we can manage the complexity through mathematical models.

The interaction between levels of analysis considers a function (e.g. pattern recognition), the theory to understand the function (e.g. through statistical inference), and how the function may be realized in a neural system (neural networks):



The Big Picture: Overview of the Brain

Before we look at models of neurons and their interactions, let us get an overview of the large scale context.

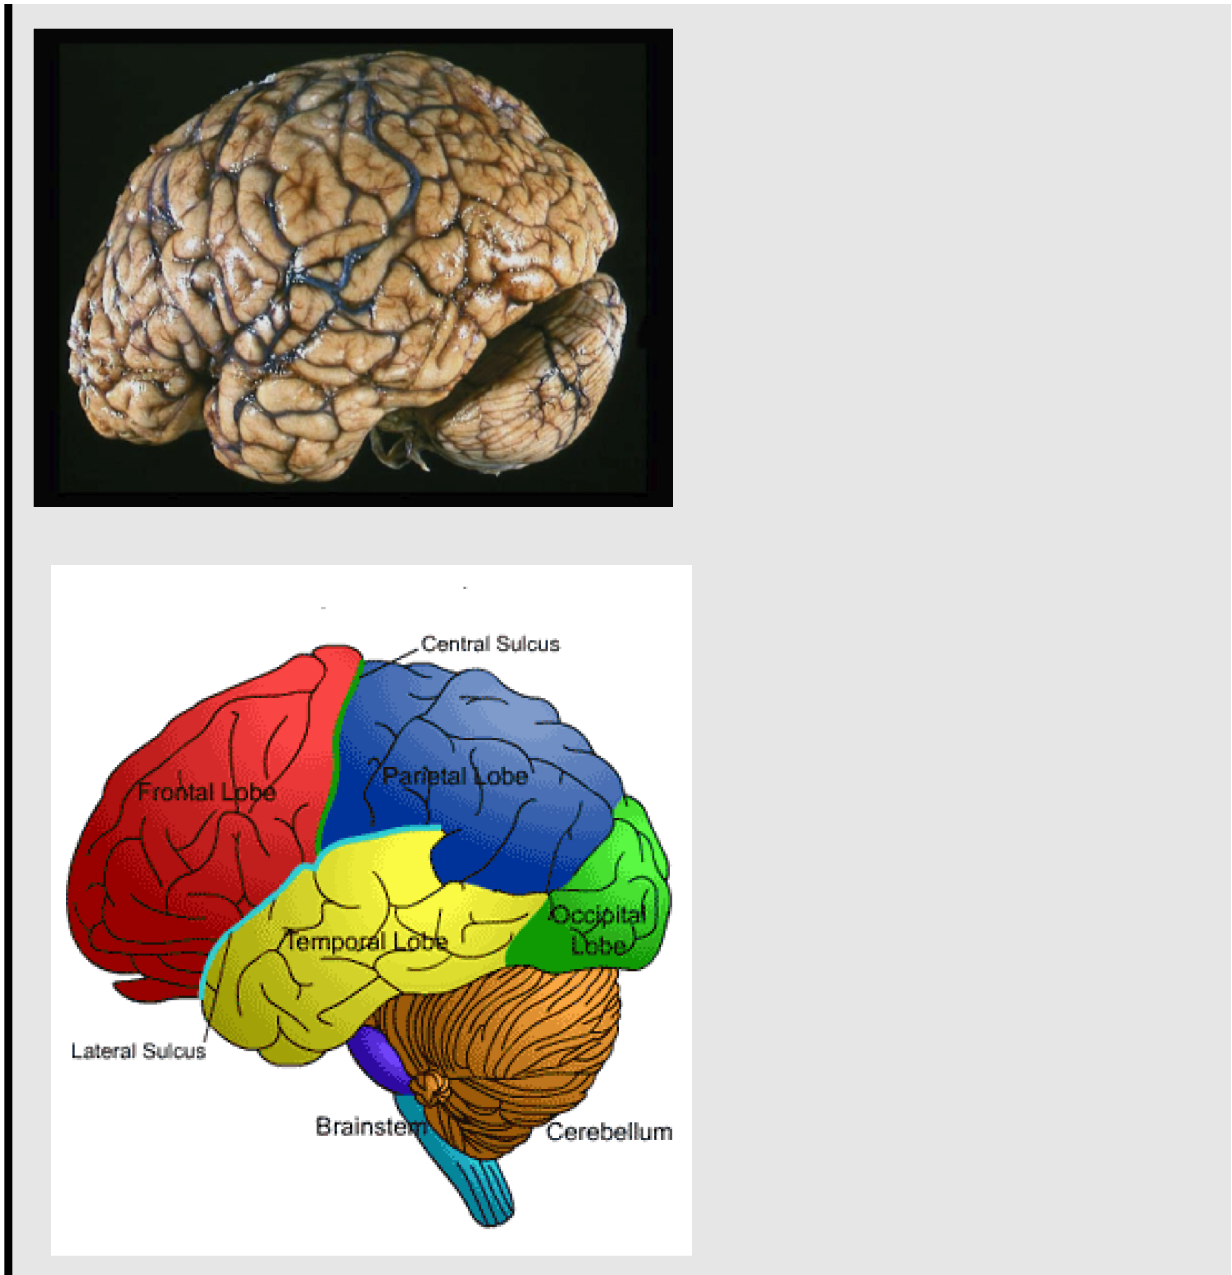
Understanding function means understanding how an organism's information processing is determined by the structure of its environmental inputs (e.g. natural images, objects to be avoided, places to go), and the nature of its outputs (e.g. estimates of shapes of objects to be manipulated, terrain to walk on, movements, and decisions). The brain doesn't operate in isolation, these inputs and outputs are intimately tied to the sensory and motor neurons that make up the peripheral nervous system, as well as the physical make-up of the body itself.

The brain has both surface and interior structures. Surface structures that are visible in the side view below are: frontal, temporal, parietal and occipital lobes, and the cerebellum. But apart from the cerebellum, it isn't totally obvious where one part stops and another begins. Landmarks are the sulci (valleys) and gyri (bumps). E.g. the lateral fissure (sulcus) is perhaps the easiest to spot. It separates the temporal lobe from the frontal and parietal lobes.

There are internal components too: Thalamus (sensory and motor relays), hypothalamus (control of endocrine activity, temperature, food intake, etc.), basal ganglia (relay for motor behavior), limbic system (expression of emotion), medulla (part of lower brain stem, breathing, heart rate).

We believe that much of what makes us interesting as humans, our thoughts, imaginations, words and actions, depends on having a large and complex cortex.

A common view shows the surface or cortex--the gray matter.




From: <http://www.utdallas.edu/~kilgard/brain.jpg>

Check out: <http://psychology.uwo.ca/fmri4newbies/PrimeronCorticalSulci.html>

The anatomy provides an important, but static view of the brain. Recent exciting progress in functional imaging provides dynamic pictures that illustrate relationships between human functions (seeing, imagining, etc.) are related to various cortical areas. (From our own lab, see: <http://gandalf.psych.umn.edu/~kersten/kersten-lab/Perceptual.html>)

Levels of organization, scales



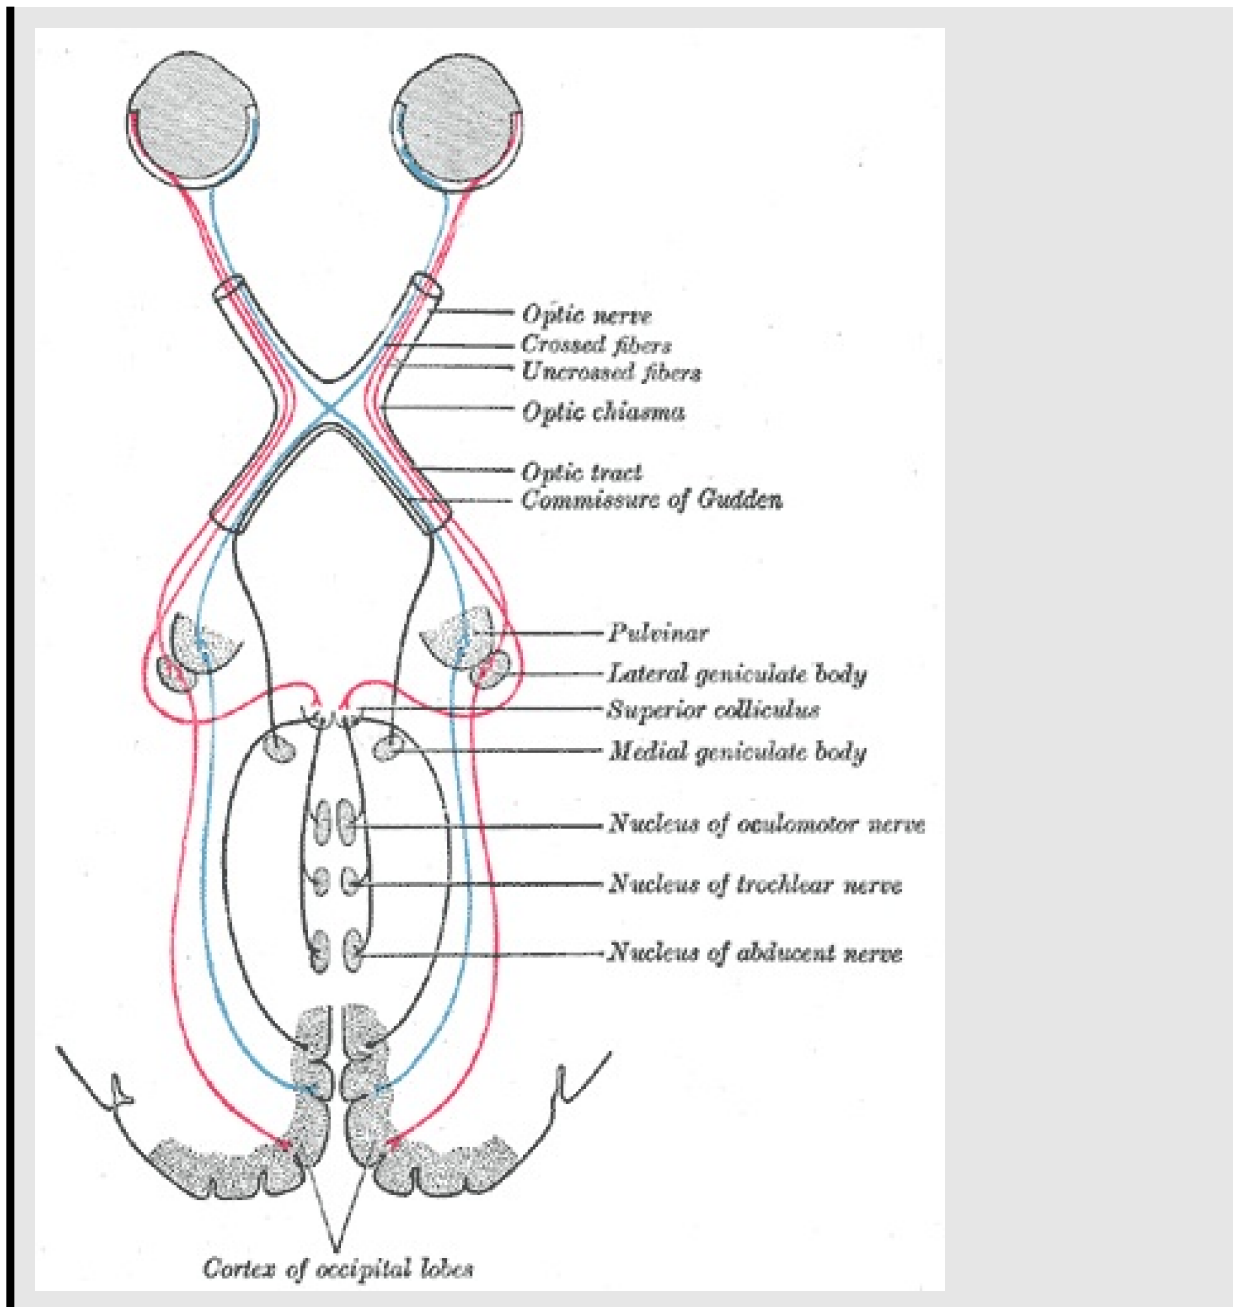
From: Churchland & Sejnowski. (10,000 \AA to a micron)

Example: Visual system

Let's take a look from two points of view: 1) information flow through a specific system--the visual system; 2) levels of organization at successive stages.

At a very coarse spatial scale, we know that we have eyes and a portion of the brain that processes the incoming images. This system enables us to recognize and manipulate objects, and that help us to navigate.

■ Anatomical view of a system

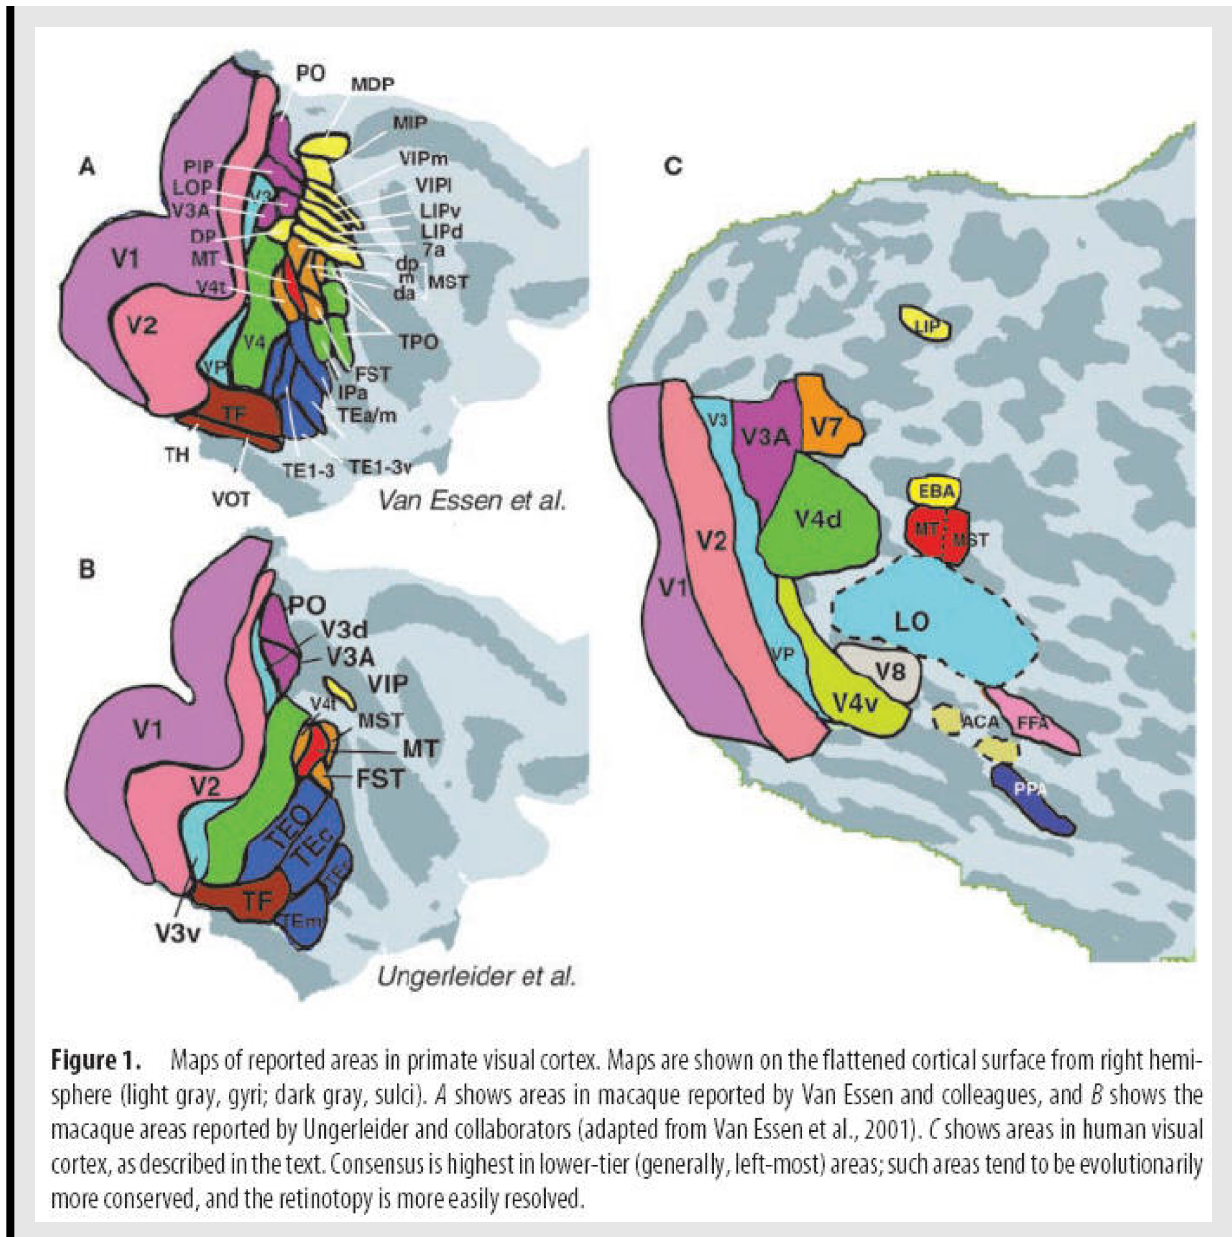


(From wikipedia entry for "Visual System")

The above picture ends at primary visual cortex (V1, area 17), but there are more than 30 other visual areas after that.

■ Visual cortex & maps (~ 1 cm scale)

Higher visual areas in cortex. Maps.



From: Tootell, R. B., Tsao, D., & Vanduffel, W. (2003). Neuroimaging weighs in: humans meet macaques in "primate" visual cortex. *J Neurosci*, 23(10), 3981-3989.


■ Cortical layers & circuits (~ 1 mm scale)

Structure within maps.

Casual inspection shows that the brain has gross structure. What is not immediately apparent is that structures do not consist of randomly connected nerve cells. There is a medium-level organization into multiple functional groupings.


The neocortex has 6 more or less distinguishable layers, there is a microorganization into vertical columns. In the primary visual cortical area (V1, see above figure), there are ocular dominance and orientation selectivity columns which are believed to form a functional unit called a hypercolumn (1 to 2 mm). Each hypercolumn takes into account local image intensities and colors to represent information or features for a single point of the visual field.

Organization of the primary visual cortex: the hypercolumn



<http://www.ualr.edu/~klwennstrom/hypercolumn.gif>

■ Retina, circuits (~100 microns)

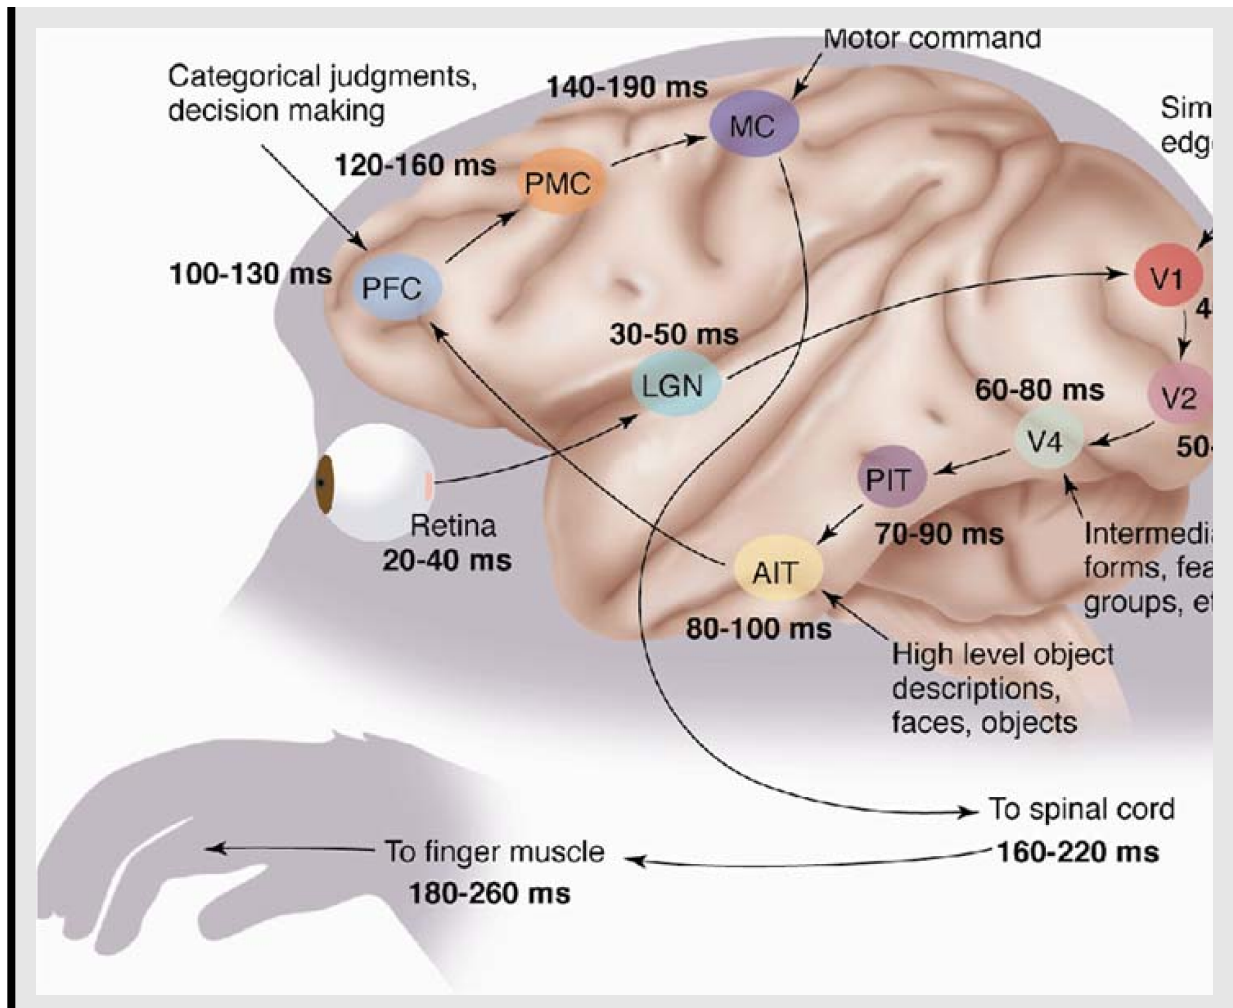


Above Adapted from Rodieck 1998

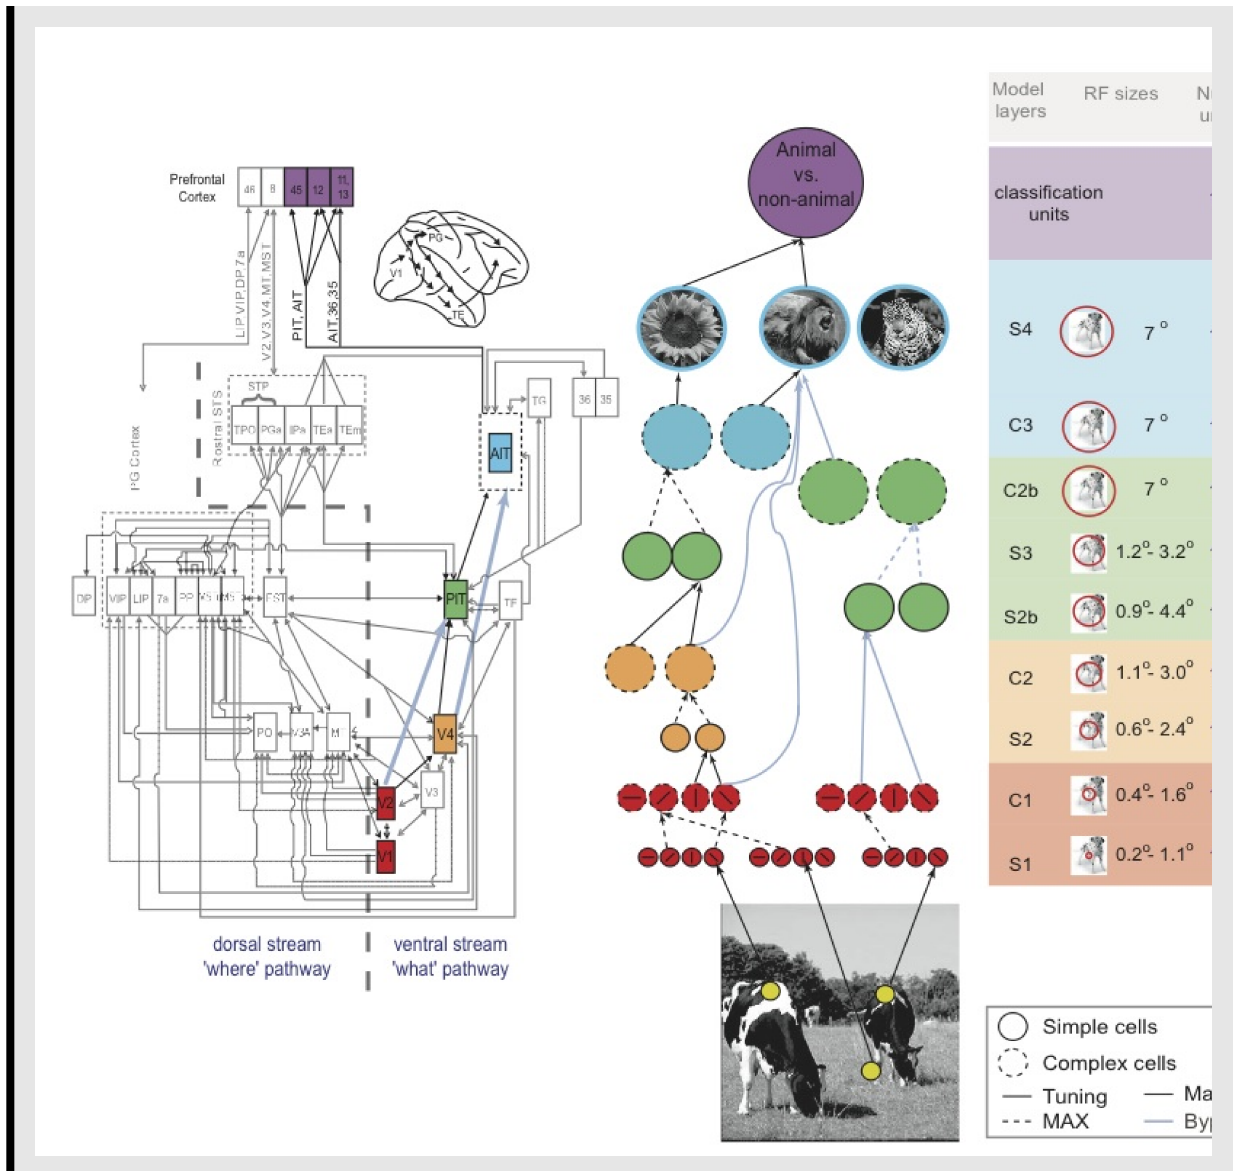
Rods, cones. Rods & rhodopsin molecules. Synapses. Different types of neurons: Spike-generating neurons, the ganglion cells.

Networks of neurons that behave as "image filters". M & P pathways.

■ Functional view of a process



(example from Simon Thorpe's lab: http://www.cerco.ups-tlse.fr/fr_vers/cerco_eng/index.php?url=objet.htm)



From Serre, T., Oliva, A., & Poggio, T. (2007). A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci U S A*, 104(15), 6424-6429.

■ Some brain specs

"cortico-centric" cognitive neuroscience

The human brain is:: volume - 1.4 liters, Cortex 2 mm, volume 0.32 liters

Cortex: 1.6 x 10¹⁰ neurons, with about 4000 synapses/neuron, about 6x10¹³ connections.

Area of cortex	1.60 E + 05	mm ^ 2		
Thickness of cortex	2.00 E + 00	mm		
Volume of cortex	3.20 E + 05	mm ^ 3	3.20 E - 01	liters

Cortex synapse density	4.00 E + 03	synapse / neuron
Cortex connectivity	2.00 E + 08	synapses / mm ³
connectivity / neuron	5.00 E + 00	mm
connection length / mm ³	2.50 E + 04	mm
neuron density in cortex	5.00 E + 04	neurons / mm³
Total brain volume	1.40 E + 00	liters
Total neurons in cortex	1.60 E + 10	
Total visual neurons	8.00 E + 09	
(50 % visual neurons is often quoted and used here, but is difficult to pin down)		
Total visual connection lengths	4.00 E + 09	mm
	4.00 E + 07	m or 24 874 miles of connections
(about the distance around the earth at the equator, 24, 901 miles)		

Some reference numbers taken or inferred from those published by :
Cherniak, J. of Cog. Neurosc., 1990, vol 2., pp 58 - 68

Getting started with *Mathematica*

Mathematica vs. Matlab vs. neural network simulation packages.

- **Front-end and Notebooks: Organize, outline, document. Program, evaluations, data all in one place**

Kernel: Separate program does the calculations

To get an overview of *Mathematica* and its help resources:

- **In *Mathematica 7*, go to *Welcome Screen* under the *Help* menu of *Mathematica* and check out the links under, **GETTING STARTED** and **LEARN AND EXPLORE**. In *Mathematica 6*, go to the **Startup palette**, and check out the "**First Five Minutes with *Mathematica***". In either version, the **Documentation Center** is a good place to go and return to, to find information.**
- **Then check out the Virtual Book under COMPLETE DOCUMENTATION**

Let's practice

- **Numerical Calculations.** You can do arithmetic. For example, type $5+7$ as shown in the cell below, and then hit the "enter" key. Note that if you try division, e.g. $2/3$, you get the exact answer back. To get a decimal approximation, type $N[2/3]$.

```
5+7
```

```
12
```

```
N[2/3]
```

```
0.666667
```

You can go back and select an expression by clicking on the brackets on the far right. These brackets serve to organize text and calculations into a Notebook with outlining features. You can group or ungroup cells for text, graphs, and expressions in various ways to present your calculations. Explore these options under Cell in the menu. You can see the possible cell types under the **Style** menu.

Try some other operations, 5^3 , $4*3$ (note that $4\ 3$, where a space separates the digits is also interpreted as multiplication).

Evaluate $4*3$

Compare with $4\ 3$ (i.e. 4 followed by a space, and then 3). Note that $4\ 3$, where a space separates the digits is also interpreted as multiplication.

Seems pretty dumb so far...but you can see that *Mathematica's* default handling of arithmetic is special:

Compare the square roots of: 12345678987654321 and 12345678987654321.0

Compare $(2^{.000000000001})^{1000000000000}$ with $(2^{(1/1000000000000)})^{1000000000000}$

If you don't explicitly tell *Mathematica* that your expressions involving floating point numbers, it will treat the numbers as true integers. And in general, *Mathematica* will do symbolic processing as a default, and assume you want as general an answer as possible. Because symbolic processing demands considerably more computer resources than numerical processing, this can be annoying. So remember to include a decimal point somewhere, or use `N[]` to force a floating point representation, and consequently numerical computations.

■ Front-end stuff

There's a lot to explore stylistically in Notebooks, but one of the most common things you will do is to select and manipulate the brackets on the far right. You can go back and select an expression by clicking on the brackets on the far right. These brackets are features of the user interface and serve to organize text and calculations into a Notebook with outlining features. You can group or ungroup cells for text, graphs, and expressions in various ways to present your calculations. Explore these options under **Cell** in the menu. You can see the possible cell types under the **Style** menu.

■ By ending an expression with `;` you can suppress the output--this is useful later when the output might be a list of a 10,000 neural activity levels!

```
(3/4)/6;  
(3 4)/6;
```

Before *Mathematica 6*, you had to be particularly careful... without a semi-colon before evaluating the next cell, you could end up with a Notebook full of 50,000 "hi"s:


```
a = Table["hi", {50 000}]
```

A very large output was generated. Here is a sample of it:

```
{hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi,
  hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi,
  hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi,
  hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi, hi}
```

[Show Less](#)
[Show More](#)
[Show Full Output](#)
[Set Size Limit...](#)

The most recent result of a calculation is given by %, the one before by %%, and so forth. Try it on the previous two outputs

Built-in functions

Mathematica has a very large library of built-in functions. They all begin with an uppercase letter and the arguments are enclosed by square brackets. Knowing that, you can often guess the form of a function.

Try taking the logarithm of 8.0

Did it return log to the base 10 or e? Check the definition by typing "?Log" or by typing "Log[E]" and "Log[10]"

You can get information more about a function, by clicking on the resulting link [>>](#)

Try Log[8,2]

You can get information about a function, e.g. for the exponential of a function, or for plotting graphs by selecting the function (e.g. Exp) and going to "Find Selected Function" in the Help menu. Or you can enter:

```
?Exp
```

```
Exp[z] is the exponential function. >>
```

and then click on >> to take you to the documentation page.

?Plot

`Plot[f, {x, xmin, xmax}]` generates a plot of f as a function of x from x_{min} to x_{max} .
`Plot[{f1, f2, ...}, {x, xmin, xmax}]` plots several functions f_i . >>

If you type two question marks before a function, **??Plot**, you'll get more information. Try it. What does the **Random** function do?

Defining your own functions

Let's illustrate function definition by building a really simple model of a neuron. We'll see the justification later. Suppose a neuron takes a set of inputs, say **x1**, **x2**, **x3**, (e.g. from the outputs of three other neurons in the network) and produces an output signal, call it **y**. For a so-called linear model, the output is the weighted sum of the inputs. We'll assume the weights are fixed and given by w_1, w_2, w_3 .

```
w1 = -1; w2 = 2; w3 = -1;

y[x1_, x2_, x3_] := w1 * x1 + w2 * x2 + w3 * x3;
```

(This is pretty primitive and not very general--don't worry, we'll get more sophisticated later).

The underscore in **x1_** is important because it tells Mathematica that **x1** represents a slot, not an expression.

■ := vs. =

Note that when defining a function, we used a colon followed by equals (`:=`) instead of just an equals sign (`=`). When you use an equals sign, the value is calculated and assigned immediately. When there is a colon in front of the equals, the value is calculated only when called on later. So we use `:=` for function definition because we need to define the function for later use and evaluation, when we may have new values for its arguments.

A double equals (`==`) has yet another meaning and is used to represent a symbolic equation which evaluates to `True` if left and right hand sides are identical.

Let's define **xv** using `:=`, and **xf** using `=`

```
xv := RandomInteger[10];
xf = RandomInteger[10];
```

Now evaluate `r1` and `r2` three times each. What is the difference between the two definitions?

```
{xv, xf}
```

```
{5, 7}
```

..but if we evaluate `xf = RandomInteger[10]` again, we re-initialize it:

```
xf = RandomInteger[10]
```

```
1
```

```
xf = RandomInteger[10];  
y[xf, xf, xf]
```

```
0
```

```
y[xv, xv, xv]
```

```
-15
```

Why does our neuron model always output a zero when the inputs all have the same value? What other family of inputs will all produce zero output?

Hint: The weights we defined can be thought of as a discrete approximation to a 2nd derivative operator in differential calculus.

Define a new function `squashedExp[]` that applies `squash` to an exponentiated value (i.e. takes `Exp[x]` as the argument of `squash[]`)

(There's a hidden cell just above this text. In the menu, go to Cell -> Cell Properties to open it if curious.)

Graphics & more function definitions


Later on we'll require defining a function that suppresses small outputs and "squashes" or clamps large outputs to a maximum level. Here is an example:

```
squash[x_] := N[1/(1 + Exp[-x+4])];
```

Also note that our squashing function was defined with **N[]**. Again, remember that *Mathematica* tries to keep everything exact as long as possible and thus will try to do symbol manipulation if we don't explicitly tell it that we want numerical representations and calculations.

Let's plot a graph of the squash function using the syntax we discovered above for $-5 < x < 10$


```
Plot[squash[x], {x, -5, 10}]
```



We'll use this and similar functions to model the small-signal compression and large signal saturation characteristics of neural output.

Plot squashedExp[] for x going from -5 to 5

```
Plot[{e^x, squashedExp[x]}, {x, -5, 5}]
```




Even though Exp grows exponentially fast with x (by definition!), squash keeps a lid on it.

Ask *Mathematica* for the definition of $\text{squashedExp}[]$

It can be important to check your definitions like this. One reason is that, as we will see later, *Mathematica* definitions can be built up with multiple constraints. And sometimes you might add to a function unwittingly and it appears to misbehave. You can check your definition by asking *Mathematica* for it.

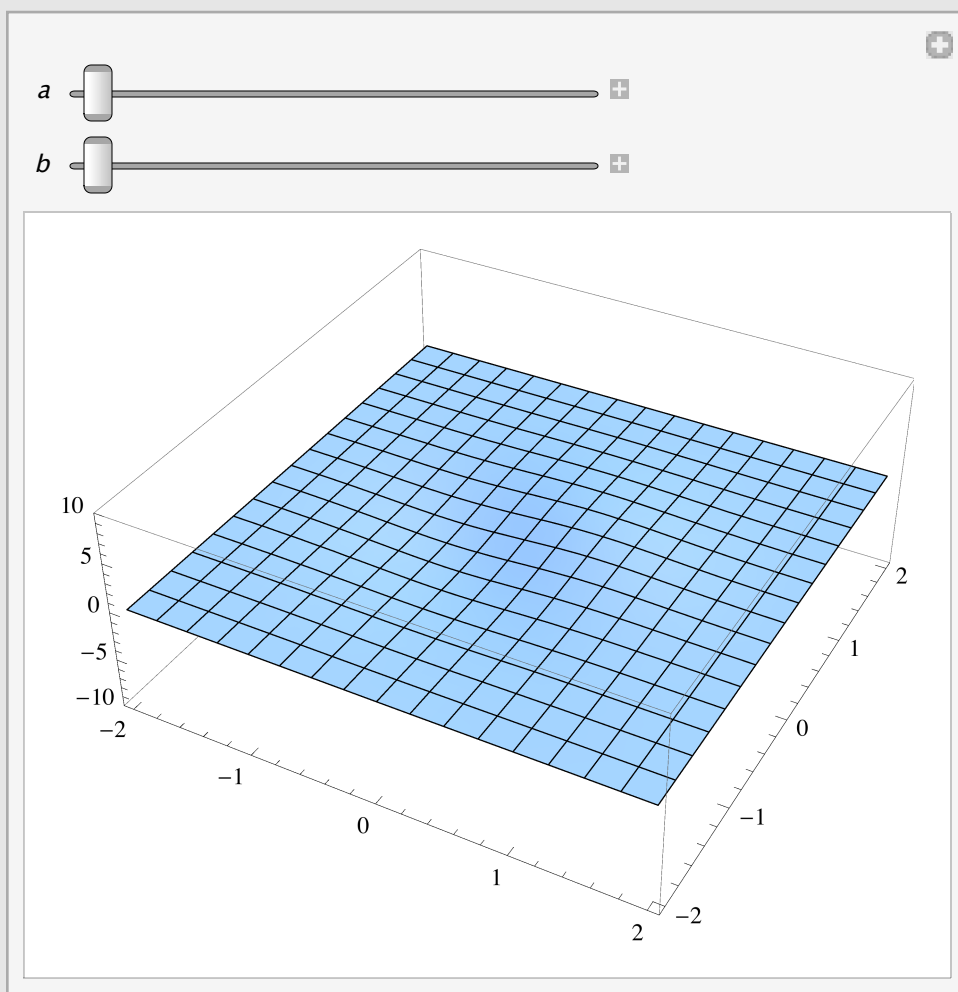
Redefine `squash` to include a steepness term λ , and use `Manipulate` to control λ

```
dsquash[x_,λ_] := N[1/(1 + Exp[(-x+4)/λ])];  
Manipulate[Plot[dsquash[x,λ], {x, -5, 10}],{{λ,1},0.1,2}]
```



Mathematica 7 has a very useful and easy-to-program GUI tools. One of the most common is `Manipulate[]`

```
Manipulate[Plot3D[Exp[-x^2 - y^2] * Sin[a * x + b * y], {x, -2, 2},  
  {y, -2, 2}, PlotRange -> {-10, 10}], {a, 1, 6}, {b, 1, 6}]
```



Try using your mouse and mouse button to rotate the above plot. What does the Option or Alt key do?

Mathematica Demonstrations project

Check out the Demonstrations center at the Wolfram *Mathematica* site for some cool examples, such as:

```

Manipulate[Module[{bluecol, colfunc},
  bluecol[n_] := Blend[{Black, Blue, White}, n];
  If[usecol, colfunc = bluecol, colfunc = GrayLevel];

  Graphics[{colfunc[left], Rectangle[{0, 0}, {3, 3}], colfunc[right],
    Rectangle[{3, 0}, {6, 3}], colfunc[mid], Rectangle[{1, 1}, {2, 2}],
    Rectangle[{4, 1}, {5, 2}]}], {{left, .35, "left square"}, .2, .8},
  {{right, .75, "right square"}, .2, .8},
  {{mid, .5, "middle squares"}, .2, .8},
  {{usecol, True, ""}, {True → "blues", False → "grays"}},
  ControlPlacement → Bottom]

```

"The Simultaneous Contrast Effect" from The Wolfram Demonstrations Project <http://demonstrations.wolfram.com/TheSimultaneousContrastEffect/>

References

Bear, M. F., Connors, B. W., & Paradiso, M. A. (2006). *Neuroscience : exploring the brain* (3rd ed.). Baltimore, MD: Lippincott Williams & Wilkins.

Churchland, P. S., & Sejnowski, T. J. (1992). *The Computational Brain* . Cambridge, MA: MIT Press.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex*, 1(1), 1-47.

Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2000). *Principles of neural science* (4th ed.). New York: McGraw-Hill, Health Professions Division.

Levy, I., Hasson, U., & Malach, R. (2004). One picture is worth at least a million neurons. *Curr Biol*, 14(11), 996-1001.

Milner, D., & Goodale, M. (1995). *The Visual Brain in Action* . Oxford: Oxford University Press.

Purves, D. (2007). *Neuroscience* (4th ed.). Sunderland, Mass.: Sinauer Associates, Publishers.

Rolls, E.T., & Treves, A. (1998). *Neural networks and brain function*. Oxford ; New York: Oxford University Press.

Zeki, S. (1993). *A Vision of the Brain*. Oxford: Blackwell Scientific Publications.

Links

<http://www.med.harvard.edu/AANLIB/cases/caseM/case.html>

<http://thalamus.wustl.edu/course/>

© 1998,2001,2003,2005,2007, 2009 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.