CS 182
Sections 101 - 102

Leon Barrett

bad puns alert!

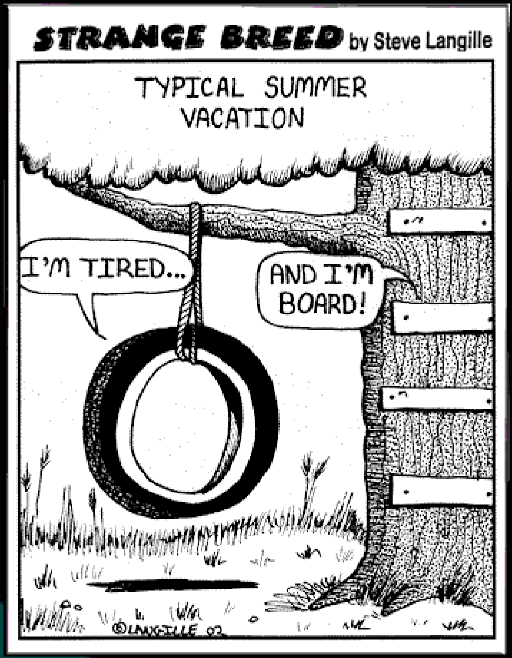

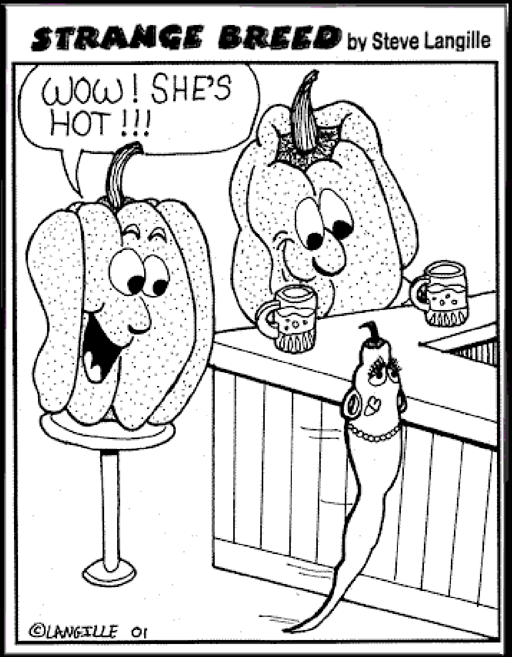

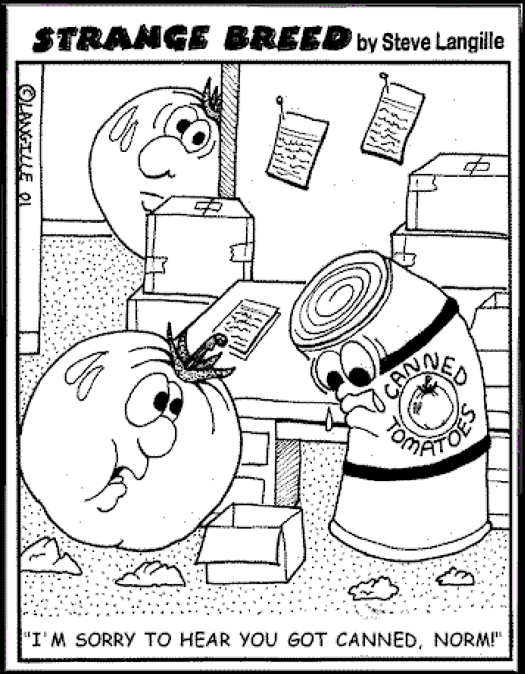

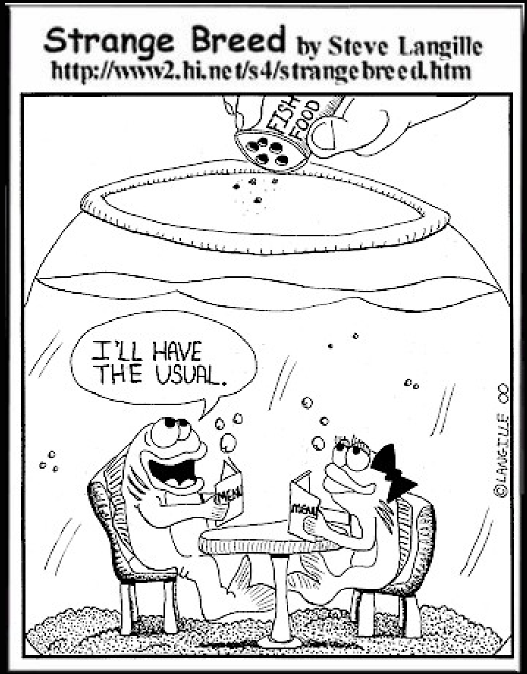

(http://www2.hi.net/s4/strangebreed.htm)

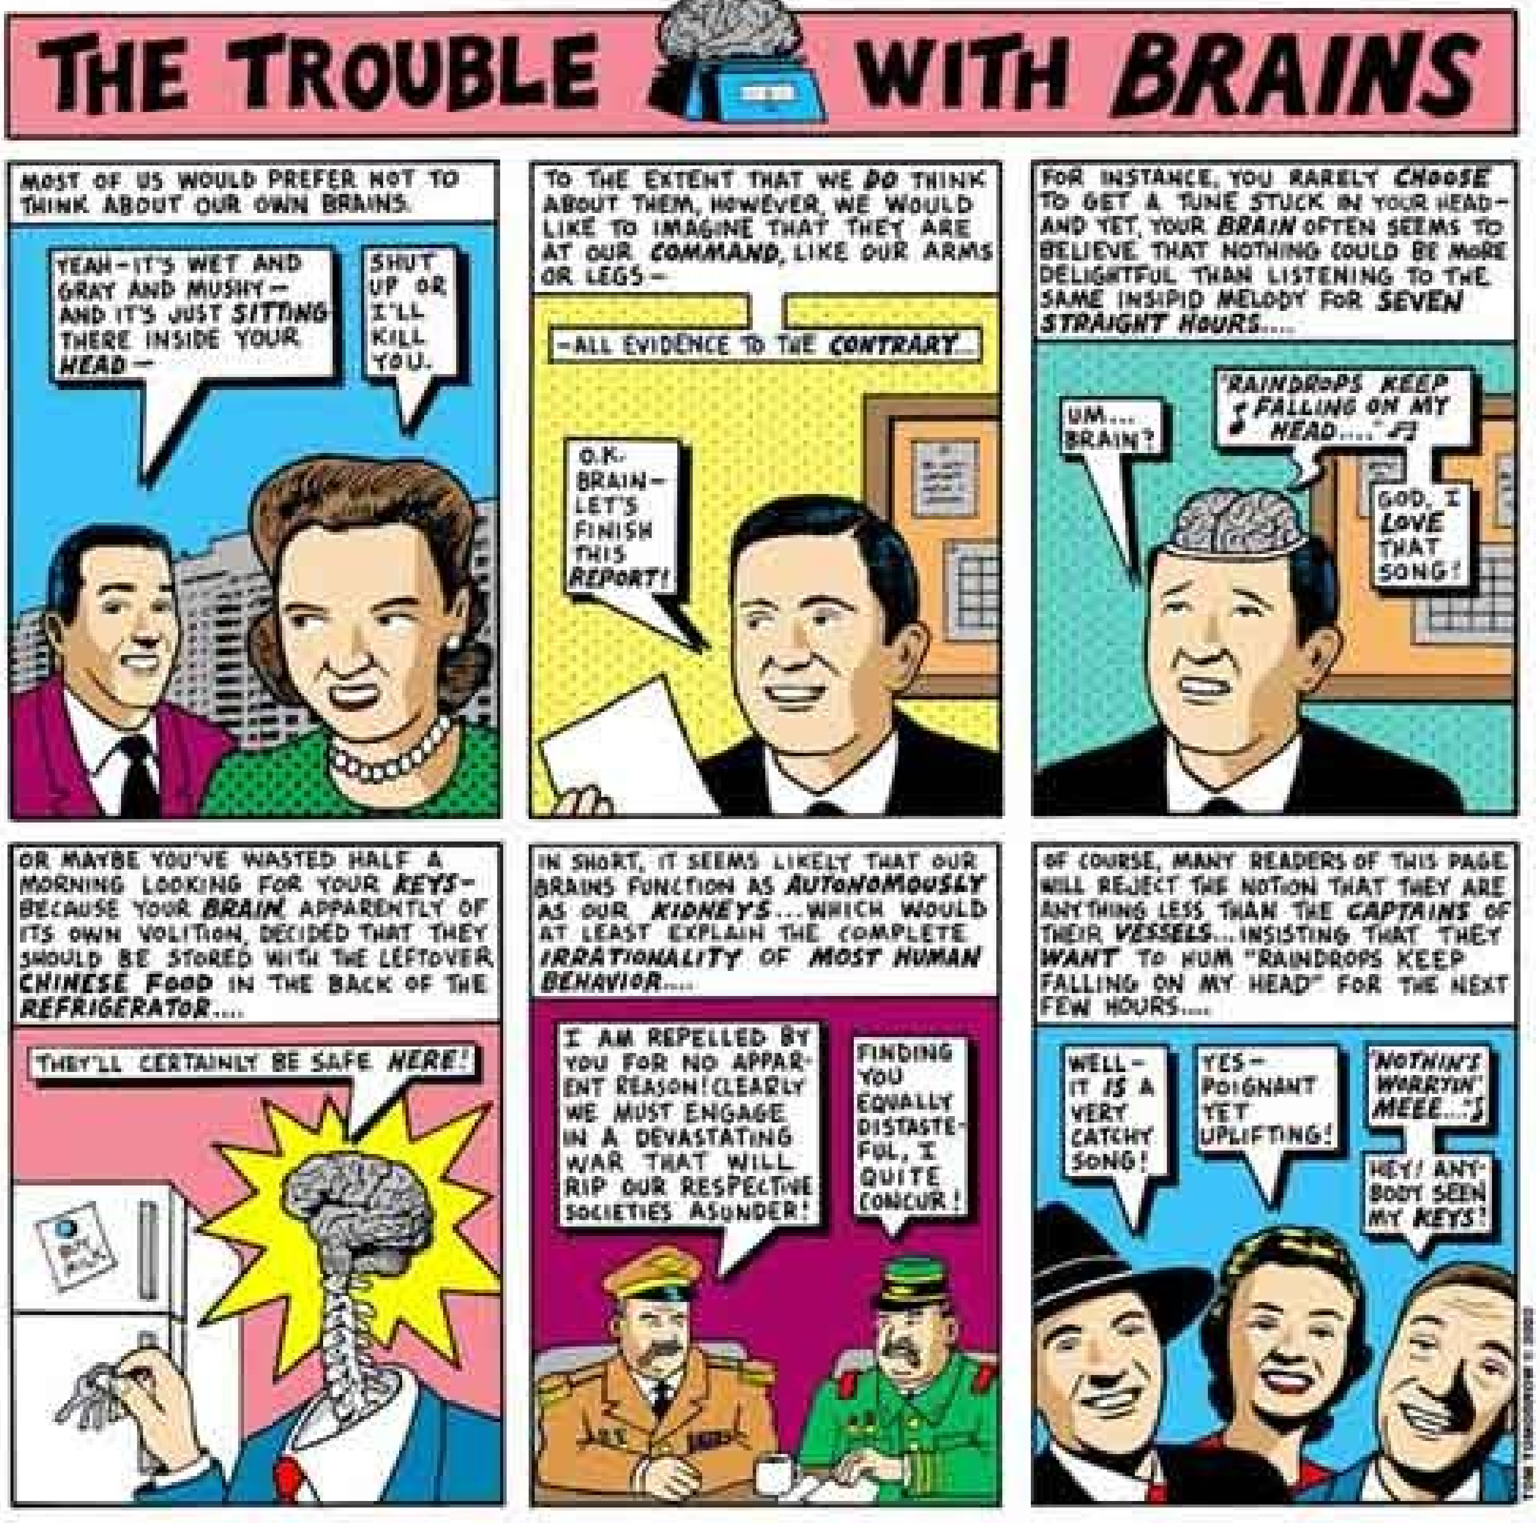
THE TROUBLE WITH BRAINS
MOST OF US WOULD PREFER NOT TO THINK ABOUT OUR OWN BRAINS.
YEAH—IT'S WET AND GRAY AND MUSHY—AND IT'S JUST SITTING THERE INSIDE YOUR HEAD—
SHUT UP OR I'LL KILL YOU.
TO THE EXTENT THAT WE DO THINK ABOUT THEM, HOWEVER, WE WOULD LIKE TO IMAGINE THAT THEY ARE AT OUR COMMAND, LIKE OUR ARMS OR LEGS—
—ALL EVIDENCE TO THE CONTRARY...
O.K. BRAIN—LET'S FINISH THIS REPORT!
FOR INSTANCE, YOU RARELY CHOOSE TO GET A TUNE STUCK IN YOUR HEAD—AND YET, YOUR BRAIN OFTEN SEEMS TO BELIEVE THAT NOTHING COULD BE MORE DELIGHTFUL THAN LISTENING TO THE SAME INSIPID MELODY FOR SEVEN STRAIGHT HOURS....
UM... BRAIN?
"RAINDROPS KEEP FALLING ON MY HEAD...."
GOD, I LOVE THAT SONG!
OR MAYBE YOU'VE WASTED HALF A MORNING LOOKING FOR YOUR KEYS—BECAUSE YOUR BRAIN, APPARENTLY OF ITS OWN VOLITION, DECIDED THAT THEY SHOULD BE STORED WITH THE LEFTOVER CHINESE FOOD IN THE BACK OF THE REFRIGERATOR....
THEY'LL CERTAINLY BE SAFE HERE!
IN SHORT, IT SEEMS LIKELY THAT OUR BRAINS FUNCTION AS AUTONOMOUSLY AS OUR KIDNEYS... WHICH WOULD AT LEAST EXPLAIN THE COMPLETE IRRATIONALITY OF MOST HUMAN BEHAVIOR....
I AM REPELLED BY YOU FOR NO APPARENT REASON! CLEARLY WE MUST ENGAGE IN A DEVASTATING WAR THAT WILL RIP OUR RESPECTIVE SOCIETIES ASUNDER!
FINDING YOU EQUALLY DISTASTEFUL, I QUITE CONCUR!
OF COURSE, MANY READERS OF THIS PAGE WILL REJECT THE NOTION THAT THEY ARE ANYTHING LESS THAN THE CAPTAINS OF THEIR VESSELS... INSISTING THAT THEY WANT TO HUM "RAINDROPS KEEP FALLING ON MY HEAD" FOR THE NEXT FEW HOURS....
WELL—IT IS A VERY CATCHY SONG!
YES—POIGNANT YET UPLIFTING!
"NOTHIN'S WORRYIN' MEEE..."
HEY! ANYBODY SEEN MY KEYS?

Announcements

- a3 part 1 is due tomorrow night (submit as a3-1)
- The second tester file is up, so please start part 2.
- If you don't like your solution to Part 1, you can get our solution on Sunday morning.

- The quiz is graded (get it after class).

Where we stand

- Last Week
  - Learning
  - backprop
  - color
- This Week
  - cognitive linguistics

Back-Propagation Algorithm

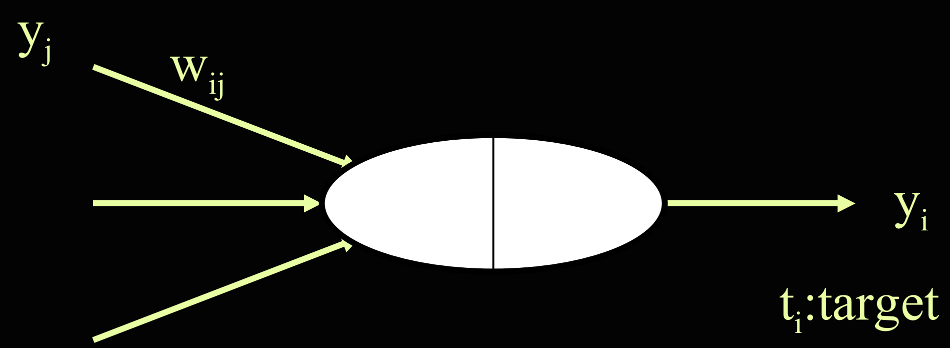

$x_i = \sum_j w_{ij} y_j$

$y_i = f(x_i)$

Sigmoid: $$y_i = f(x_i) = \frac{1}{1+e^{-x_i}}$$

We define the error term for a single node to be $t_i - y_i$

Gradient Descent

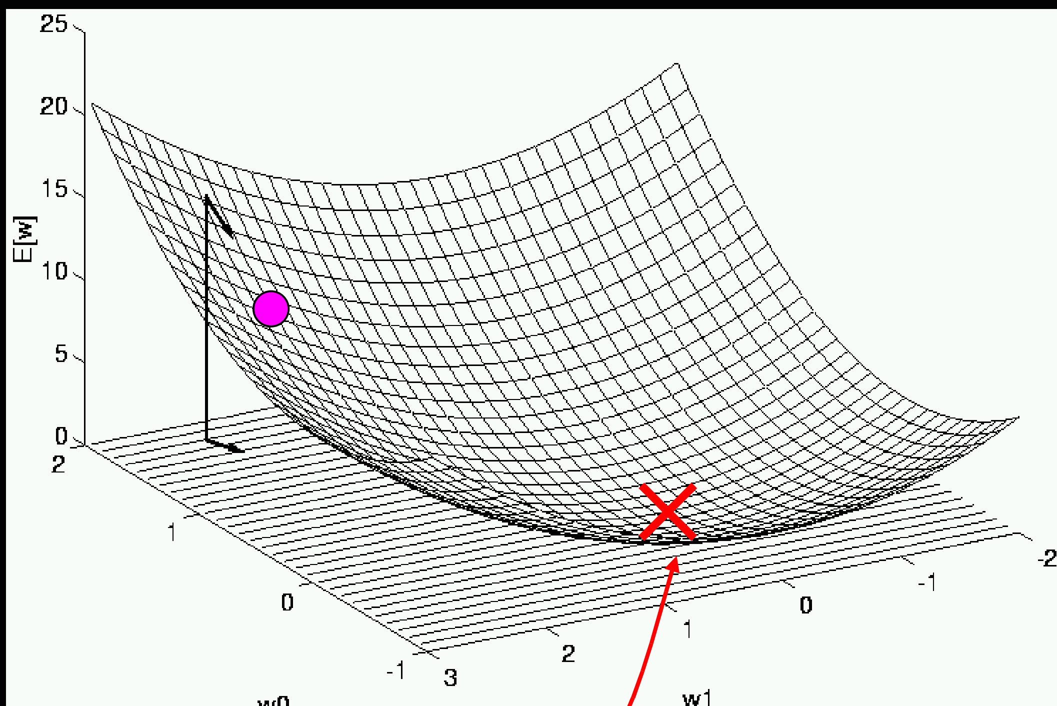

global mimimum:
this is your goal

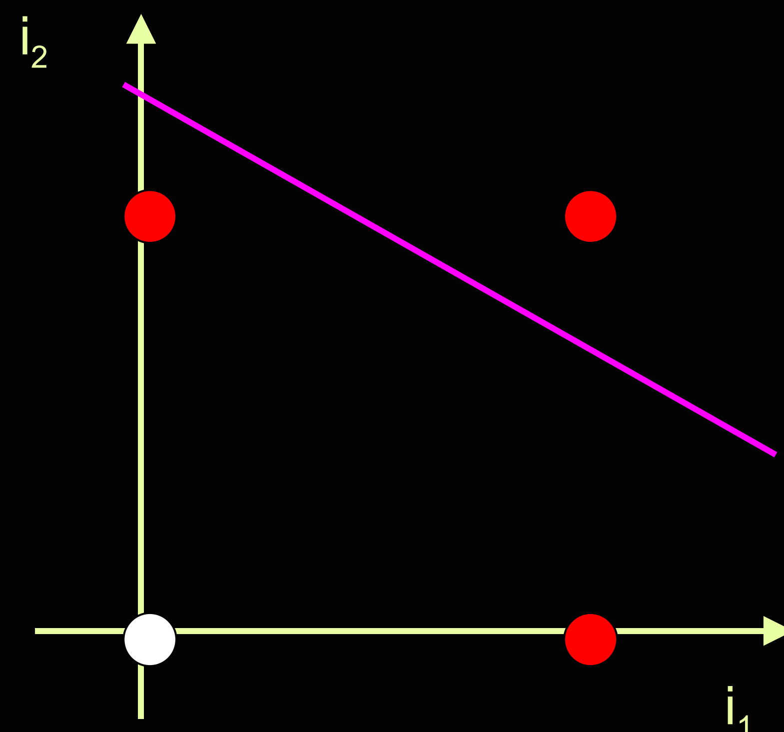

it should be 4-D (3 weights) but you get the idea

Equations of Backprop

- Weight update shown on following slides; important equations highlighted in green
- Note momentum equation:
  - dW(t) = change in weight at time t
  - dW(t-1) = change in weight at time t-1
  - so using momentum:
  - dW(t) = -learning_rate * -input * delta(i) + momentum * dW(t-1)
  - the first part of that comes from last slides below
  - the second part is the momentum term

The output layer

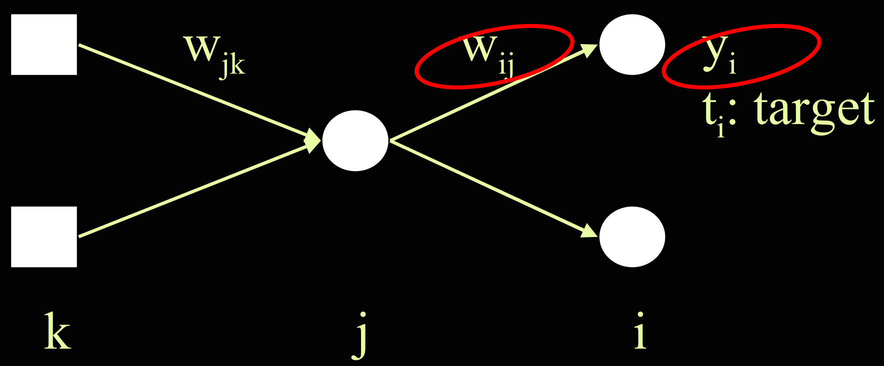

learning rate

$$W_{ij} \leftarrow W_{ij} - \alpha \cdot \frac{\partial E}{\partial W_{ij}}$$

$$\Delta W_{ij} = -\alpha \cdot \frac{\partial E}{\partial W_{ij}}$$

E = Error = $\frac{1}{2} \sum_i (t_i - y_i)^2$

$$\frac{\partial E}{\partial W_{ij}} = \frac{\partial E}{\partial y_i} \cdot \frac{\partial y_i}{\partial x_i} \cdot \frac{\partial x_i}{\partial W_{ij}} = -(t_i - y_i) \cdot f'(x_i) \cdot y_j$$

The derivative of the sigmoid is just $y_i(1 - y_i)$

$$\Delta W_{ij} = -\alpha \cdot -(t_i - y_i) \cdot y_i(1 - y_i) \cdot y_j$$

$$\Delta W_{ij} = -\alpha \cdot -y_j \cdot \delta_i \qquad \delta_i = (t_i - y_i) \cdot y_i(1 - y_i)$$

The hidden layer

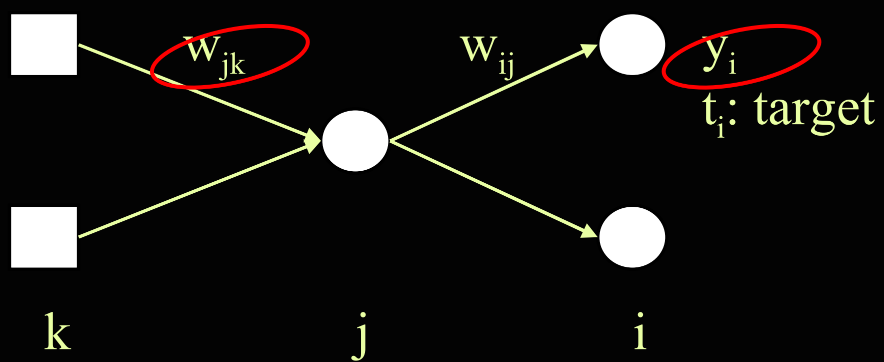

$$\boxed{\Delta W_{jk} = -\alpha \cdot \frac{\partial E}{\partial W_{jk}}}$$

$$\frac{\partial E}{\partial W_{jk}} = \frac{\partial E}{\partial y_j} \cdot \frac{\partial y_j}{\partial x_j} \cdot \frac{\partial x_j}{\partial W_{jk}}$$

$$\frac{\partial E}{\partial y_j} = \sum_i \frac{\partial E}{\partial y_i} \cdot \frac{\partial y_i}{\partial x_i} \cdot \frac{\partial x_i}{\partial y_j} = \sum_i -(t_i - y_i) \cdot f'(x_i) \cdot W_{ij}$$

$$\frac{\partial E}{\partial W_{jk}} = \left(-\sum_i (t_i - y_i) \cdot f'(x_i) \cdot W_{ij} \right) \cdot f'(x_j) \cdot y_k$$

E = Error = $\frac{1}{2} \sum_i (t_i - y_i)^2$

$$\Delta W_{jk} = -\alpha \cdot \left(-\sum_i (t_i - y_i) \cdot y_i (1 - y_i) \cdot W_{ij} \right) \cdot y_j (1 - y_j) \cdot y_k$$

$$\boxed{\Delta W_{jk} = -\alpha \cdot -y_k \cdot \delta_j}$$

$$\delta_j = \left(\sum_i (t_i - y_i) \cdot y_i (1 - y_i) \cdot W_{ij} \right) \cdot y_j (1 - y_j)$$

$$\boxed{\delta_j = \left(\sum_i W_{ij} \cdot \delta_i \right) \cdot y_j (1 - y_j)}$$

Let's just do an example

i_1	i_2	y_0
0	0	0
0	1	1
1	0	1
1	1	1

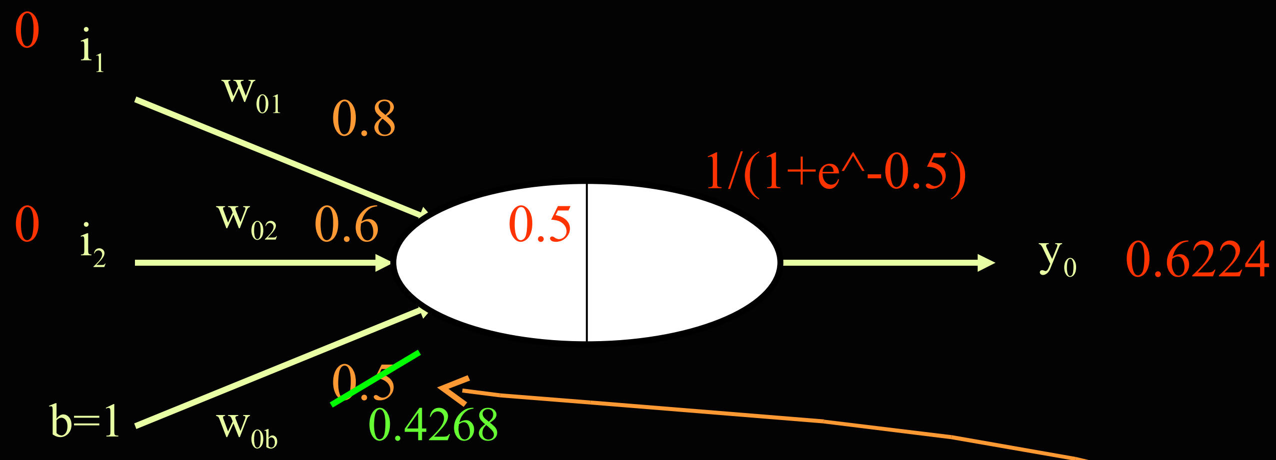

E = Error = $\frac{1}{2}\sum_i (t_i - y_i)^2$

$E = \frac{1}{2}(t_0 - y_0)^2$

$E = \frac{1}{2}(0 - 0.6224)^2 = 0.1937$

$$\Delta W_{ij} = -\alpha \cdot -y_j \cdot \delta_i$$

$$\Delta W_{01} = -\alpha \cdot -y_1 \cdot \delta_0 \quad = -\alpha \cdot -\overset{0}{i_1} \delta_0$$

$$\Delta W_{02} = -\alpha \cdot -y_2 \cdot \delta_0 \quad = -\alpha \cdot -\overset{0}{i_2} \cdot \delta_0$$

$$\Delta W_{0b} = -\alpha \cdot -y_b \cdot \delta_0 \quad = -\alpha \cdot -b \cdot \delta_0$$

$$= \alpha \cdot -0.1463$$

learning rate

$$\delta_i = (t_i - y_i) \cdot y_i (1 - y_i)$$

$$\delta_0 = (t_0 - y_0) \cdot y_0 (1 - y_0)$$

$$\delta_0 = (0 - 0.6224) \cdot 0.6224 (1 - 0.6224)$$

$$\delta_0 = -0.1463$$

suppose $\alpha = 0.5$

$$\Delta W_{0b} = 0.5 \cdot -0.1463 = -0.0731$$

Biological learning

1. What is Hebbian learning?
2. Where has it been observed?
3. What is wrong with Hebbian learning as a story of how animals learn?
   - hint – it's the opposite of what's wrong with backprop

LTP and Hebb's Rule

- Hebb's Rule:
  neurons that fire together wire together

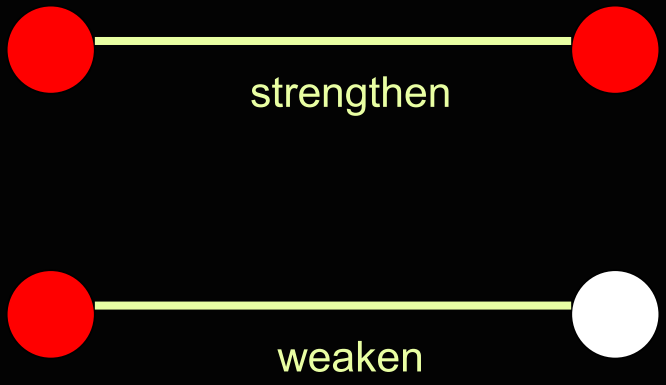

- Long Term Potentiation (LTP) is the biological basis of Hebb's Rule
- Calcium channels are the key mechanism

Why is Hebb's rule incomplete?

- here's a contrived example:

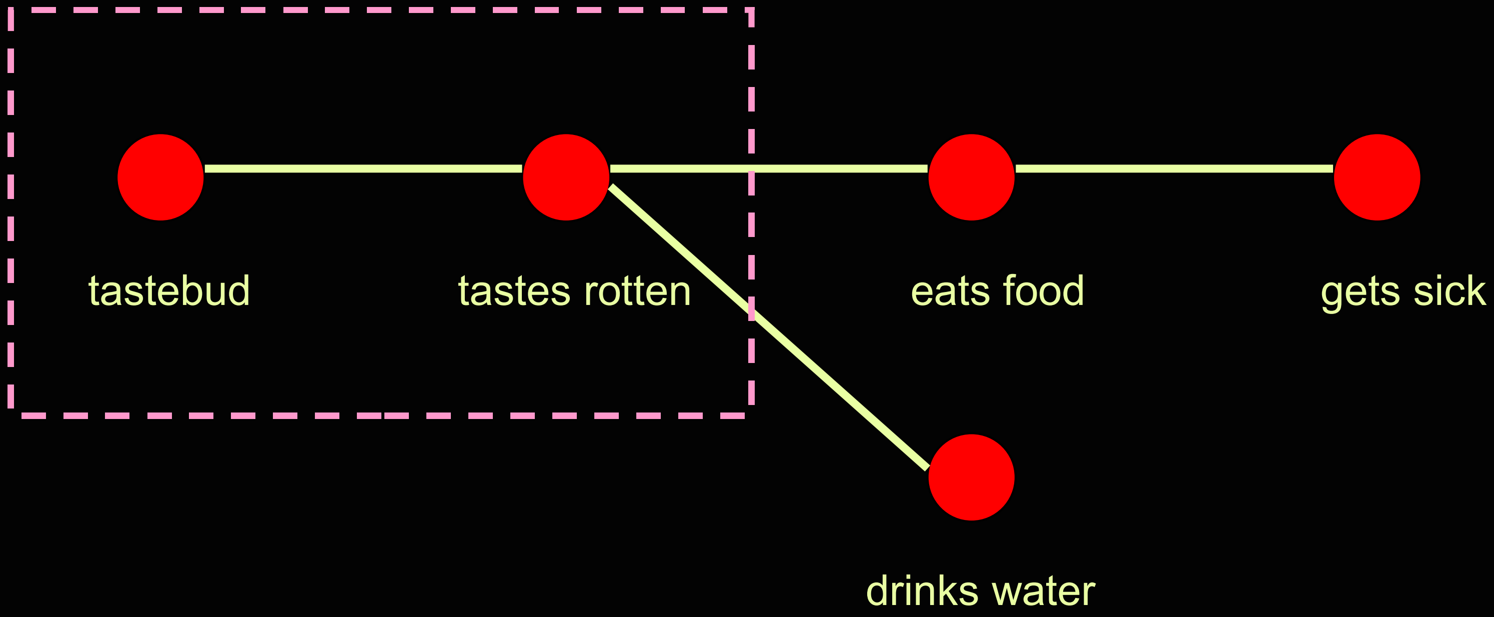

- should you "punish" all the connections?

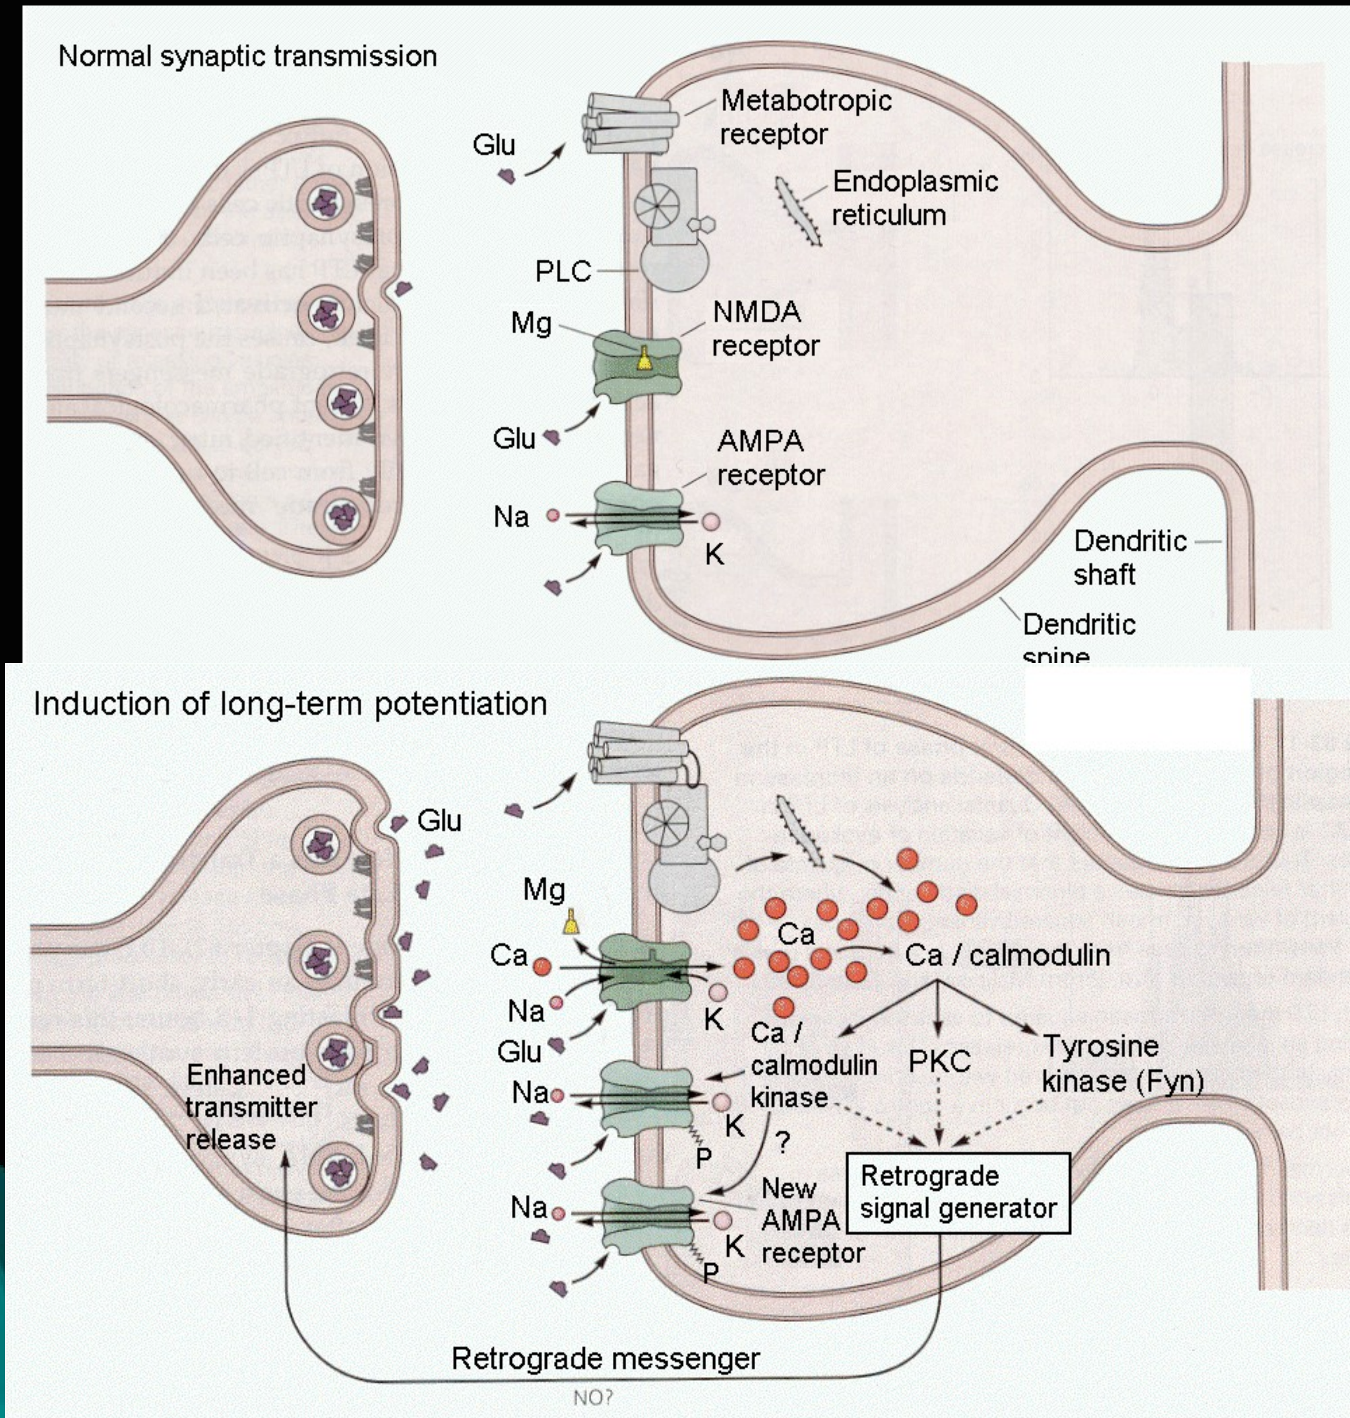

During normal low-frequency trans-mission, glutamate interacts with NMDA and non-NMDA (AMPA) and metabotropic receptors.

With high-frequency stimulation, Calcium comes in

Recruitment learning

- What is recruitment learning?
- Why do we need it in our story?
- How does it relate to triangle nodes?

Models of Learning

- Hebbian ~ coincidence
- Recruitment ~ one trial
- Supervised ~ correction (backprop)
- Reinforcement ~ delayed reward
- Unsupervised ~ similarity

Questions!

1. How do humans detect color biologically?
2. Are color names arbitrary? What are the findings surrounding this?

Questions!

- How do humans detect color biologically?
- Are color names arbitrary? What are the findings surrounding this?

A Tour of the Visual System

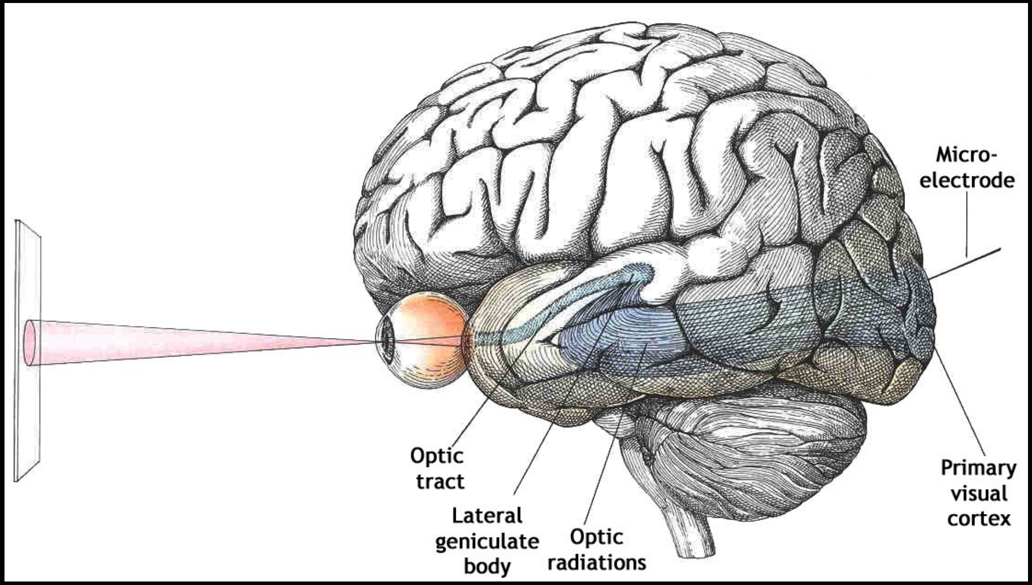

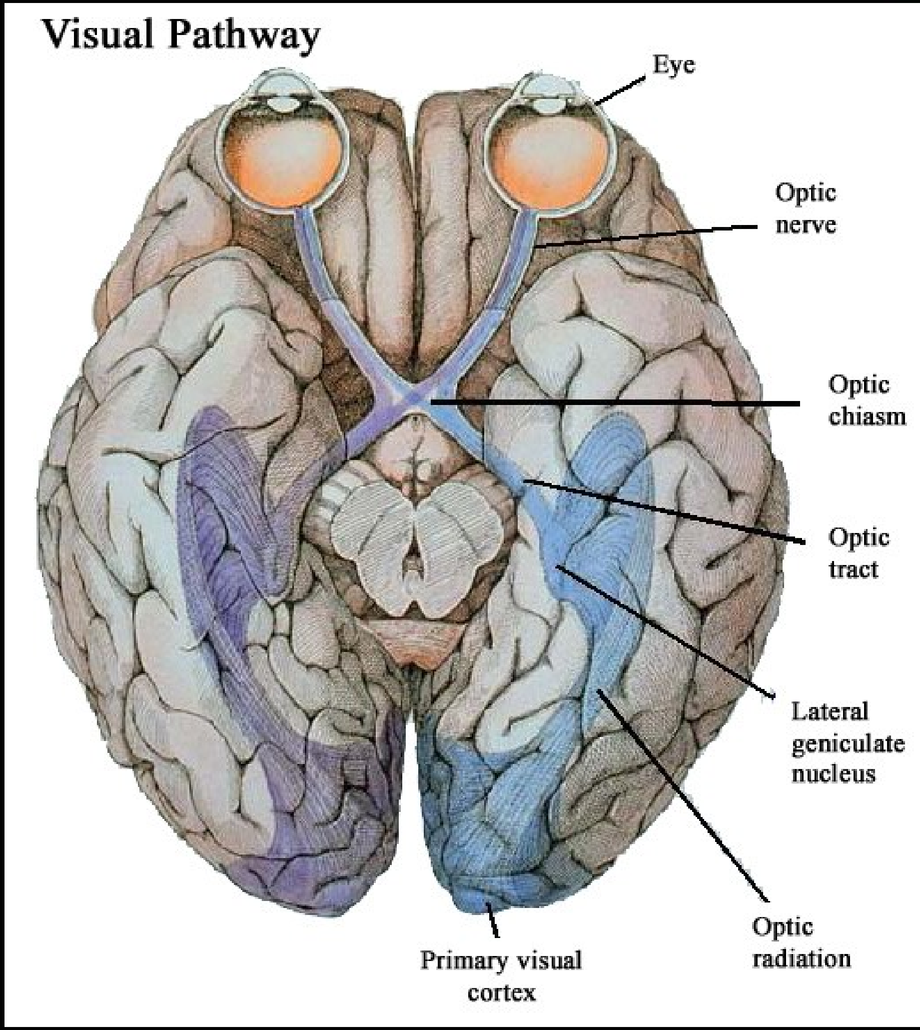

- two regions of interest:
  - retina
  - LGN

Rods and Cones in the Retina

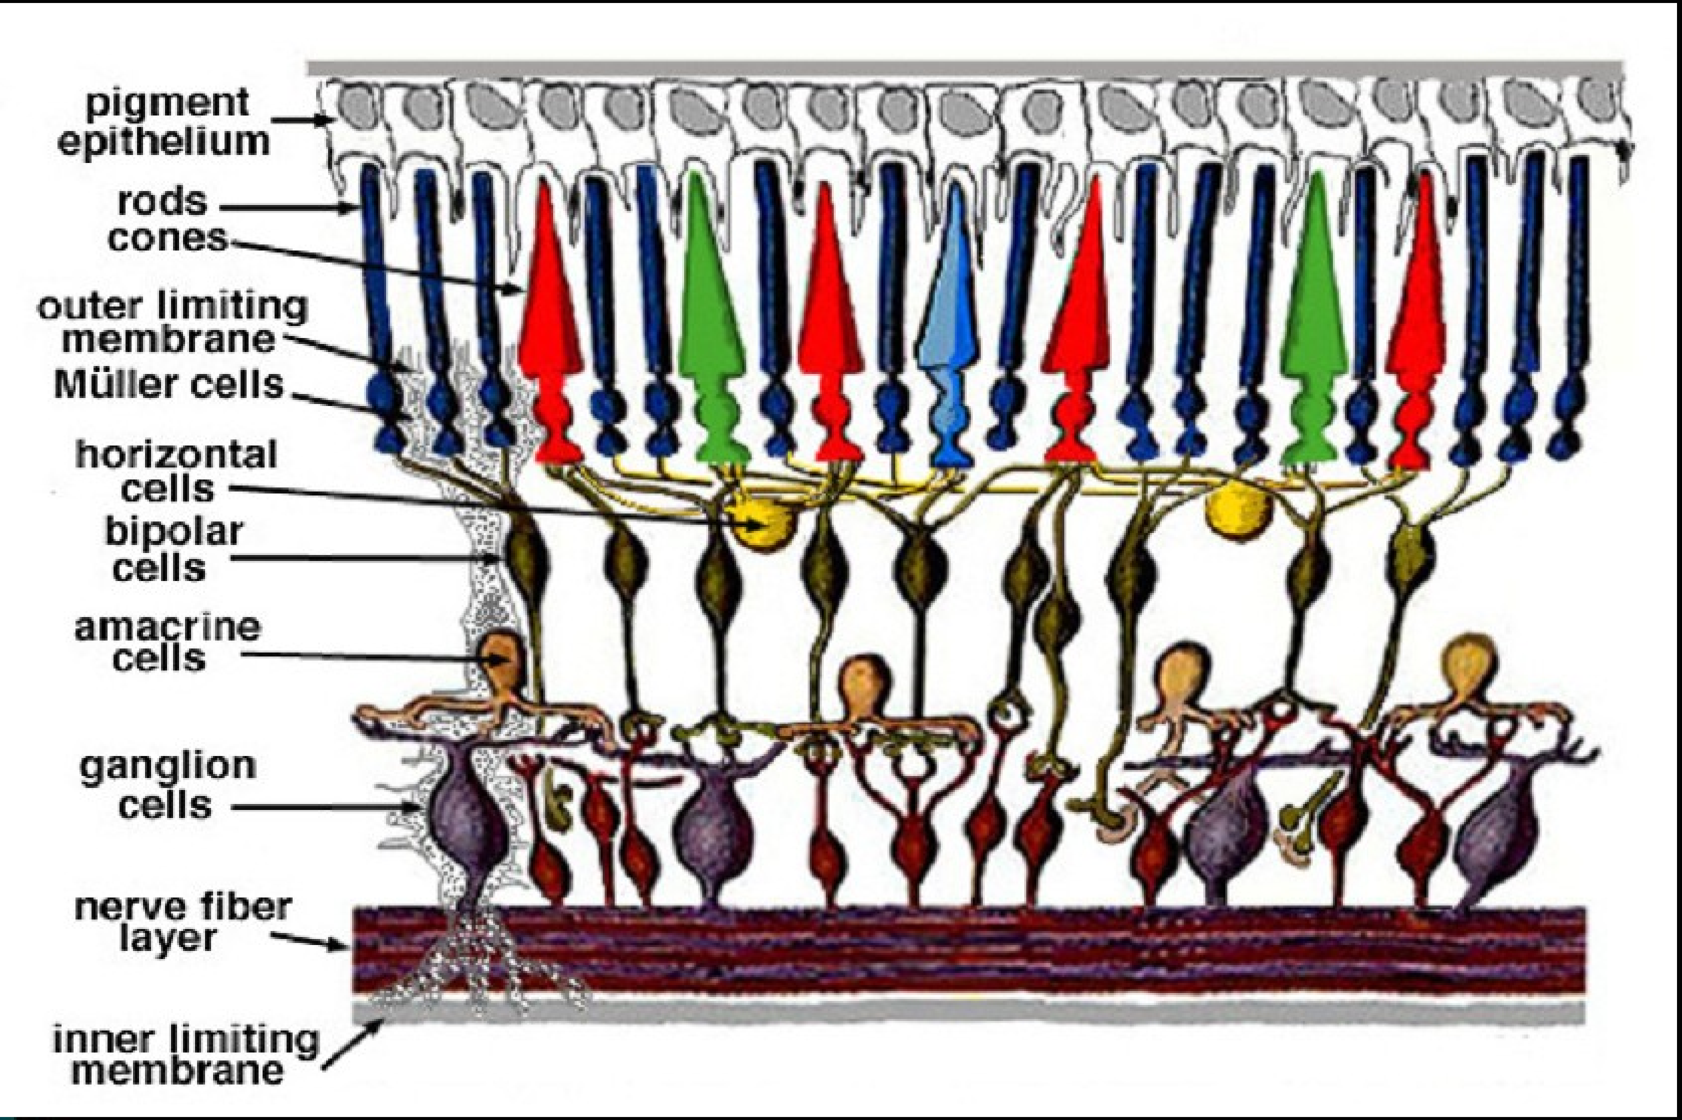

http://www.iit.edu/~npr/DrJennifer/visual/retina.html

The Microscopic View

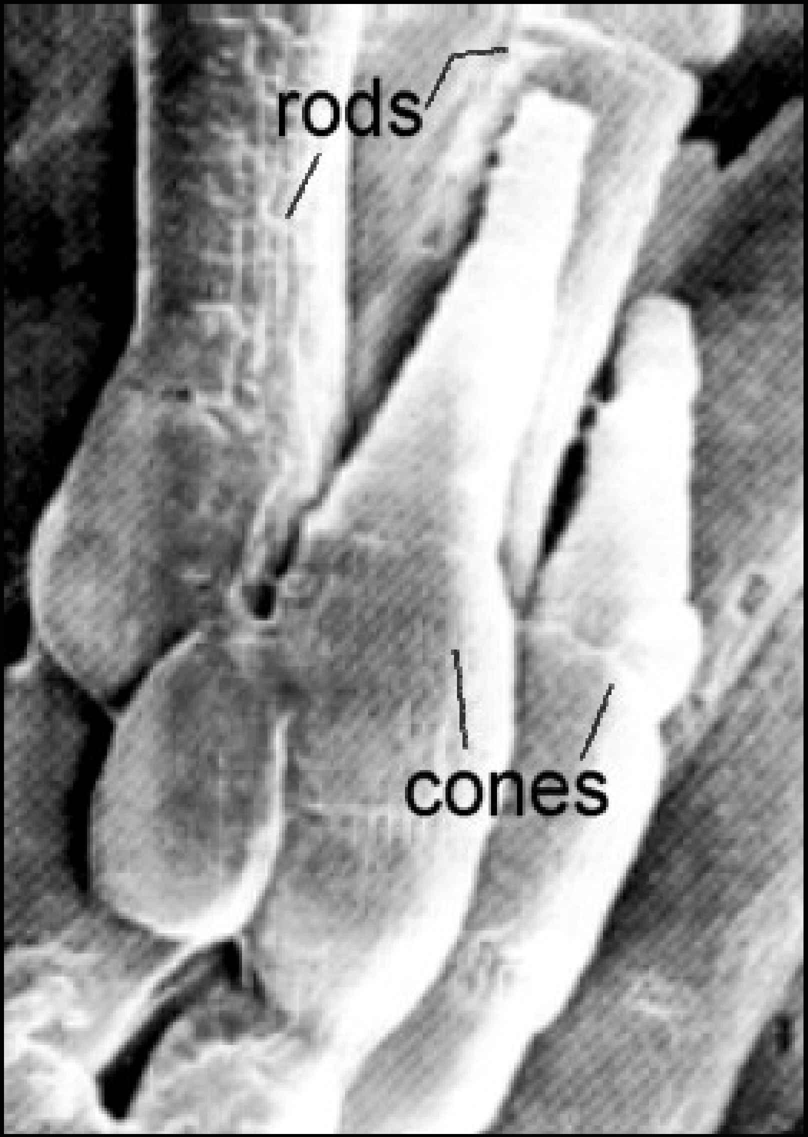

What Rods and Cones Detect

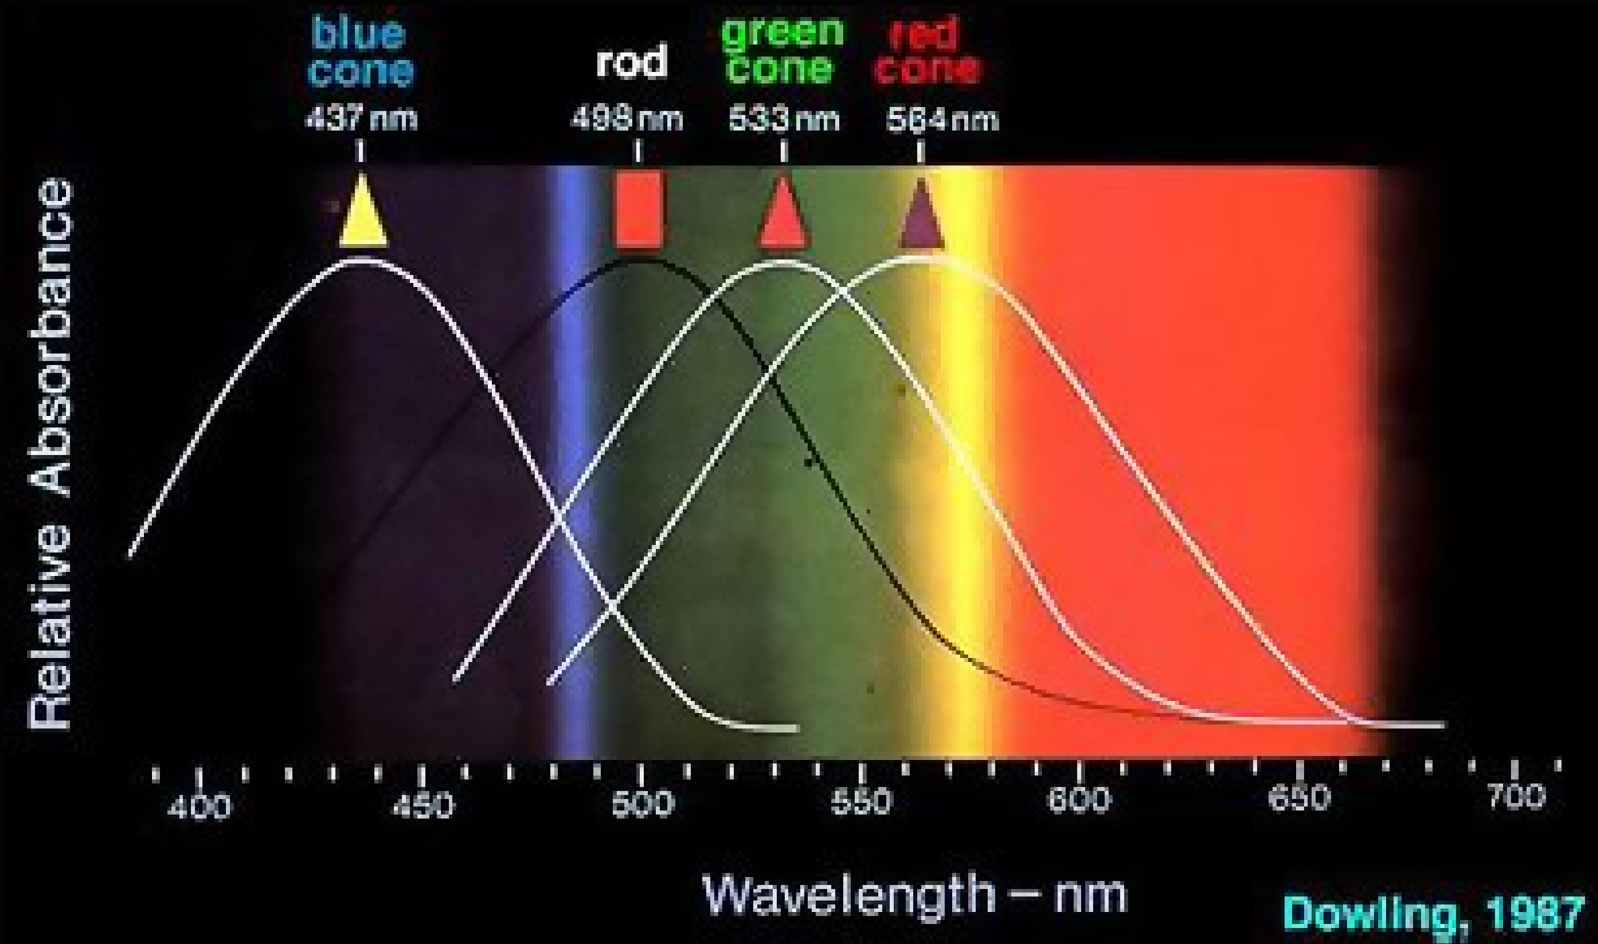

Notice how they aren't distributed evenly, and the rod is more sensitive to shorter wavelengths

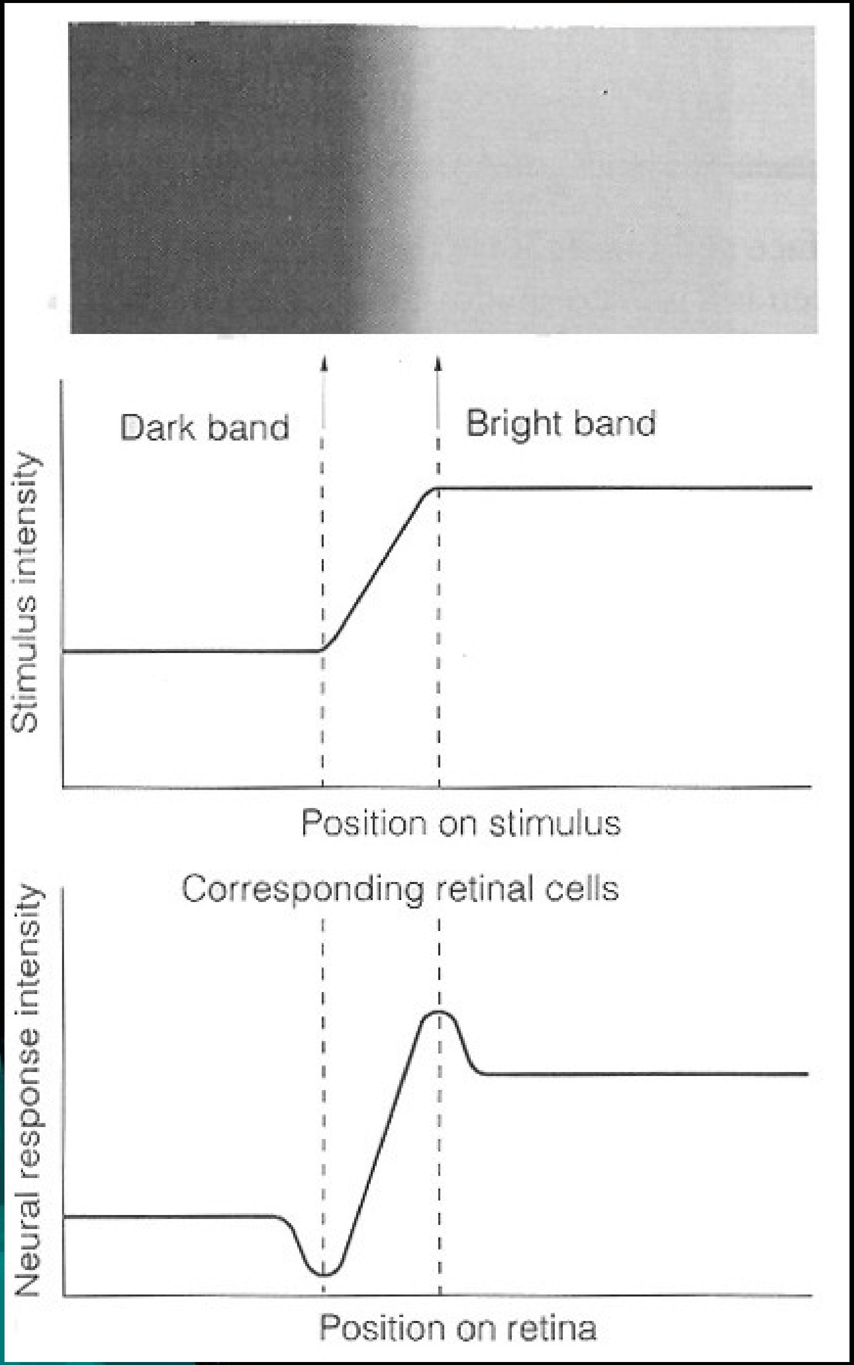

Center / Surround

- Strong activation in center, inhibition on surround
- The effect you get using these center / surround cells is enhanced edges

top: the stimuli itself

middle: brightness of the stimuli

bottom: response of the retina

- You'll see this idea get used in Regier's model

http://www-psych.stanford.edu/~lera/psych115s/notes/lecture3/figures1.html

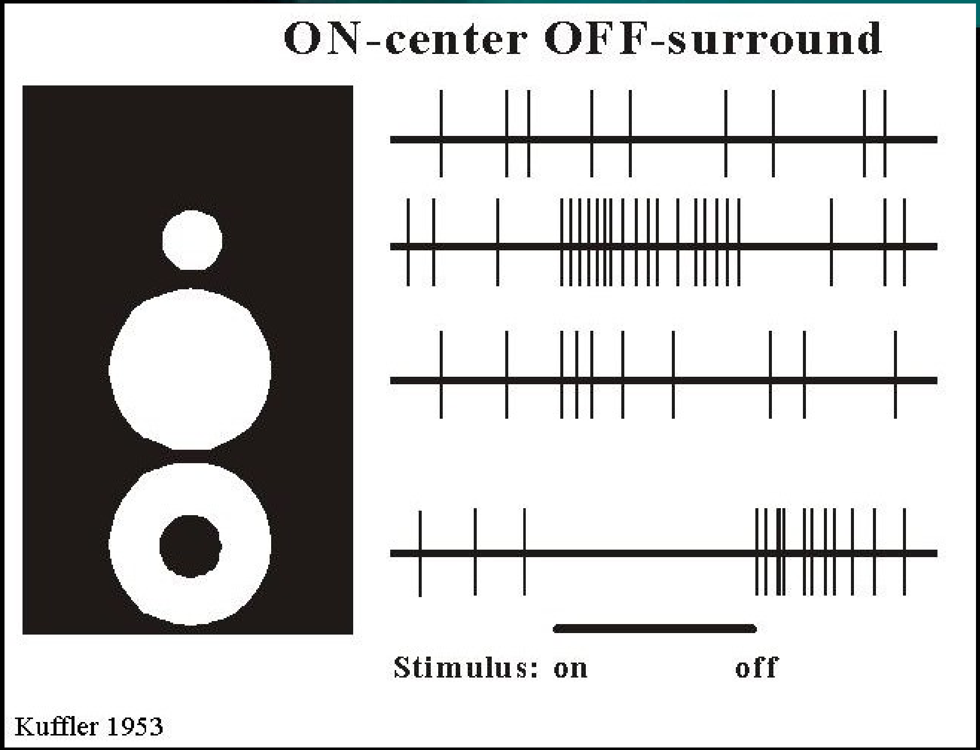

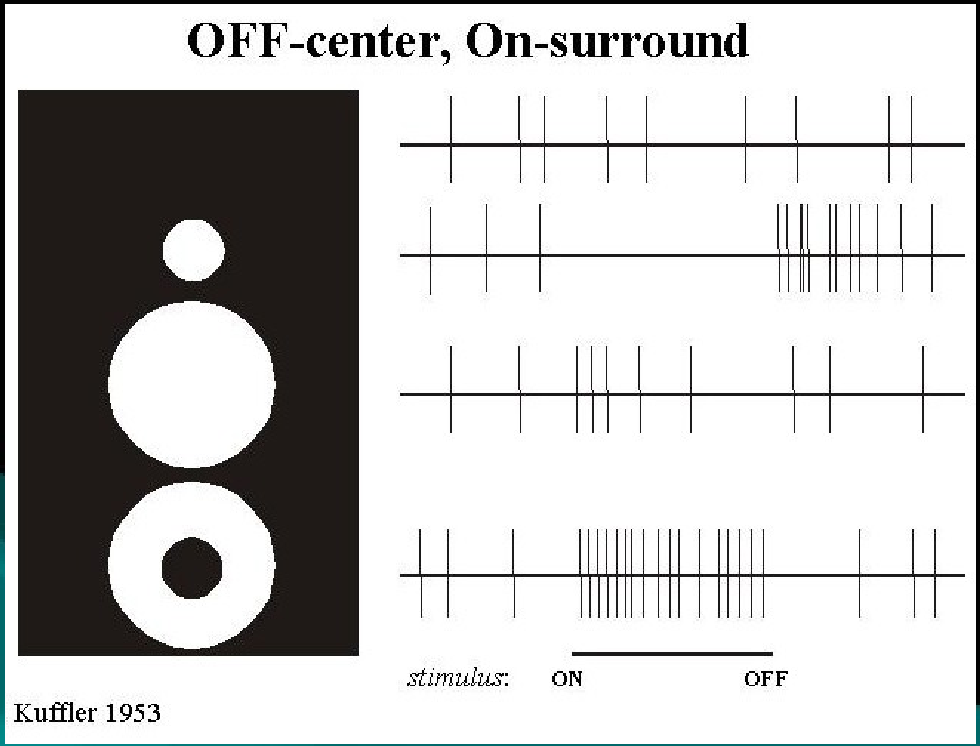

How They Fire

- No stimuli:
  - both fire at base rate
- Stimuli in center:
  - ON-center-OFF-surround fires rapidly
  - OFF-center-ON-surround doesn't fire
- Stimuli in surround:
  - OFF-center-ON-surround fires rapidly
  - ON-center-OFF-surround doesn't fire
- Stimuli in both regions:
  - both fire slowly

Color Opponent Cells

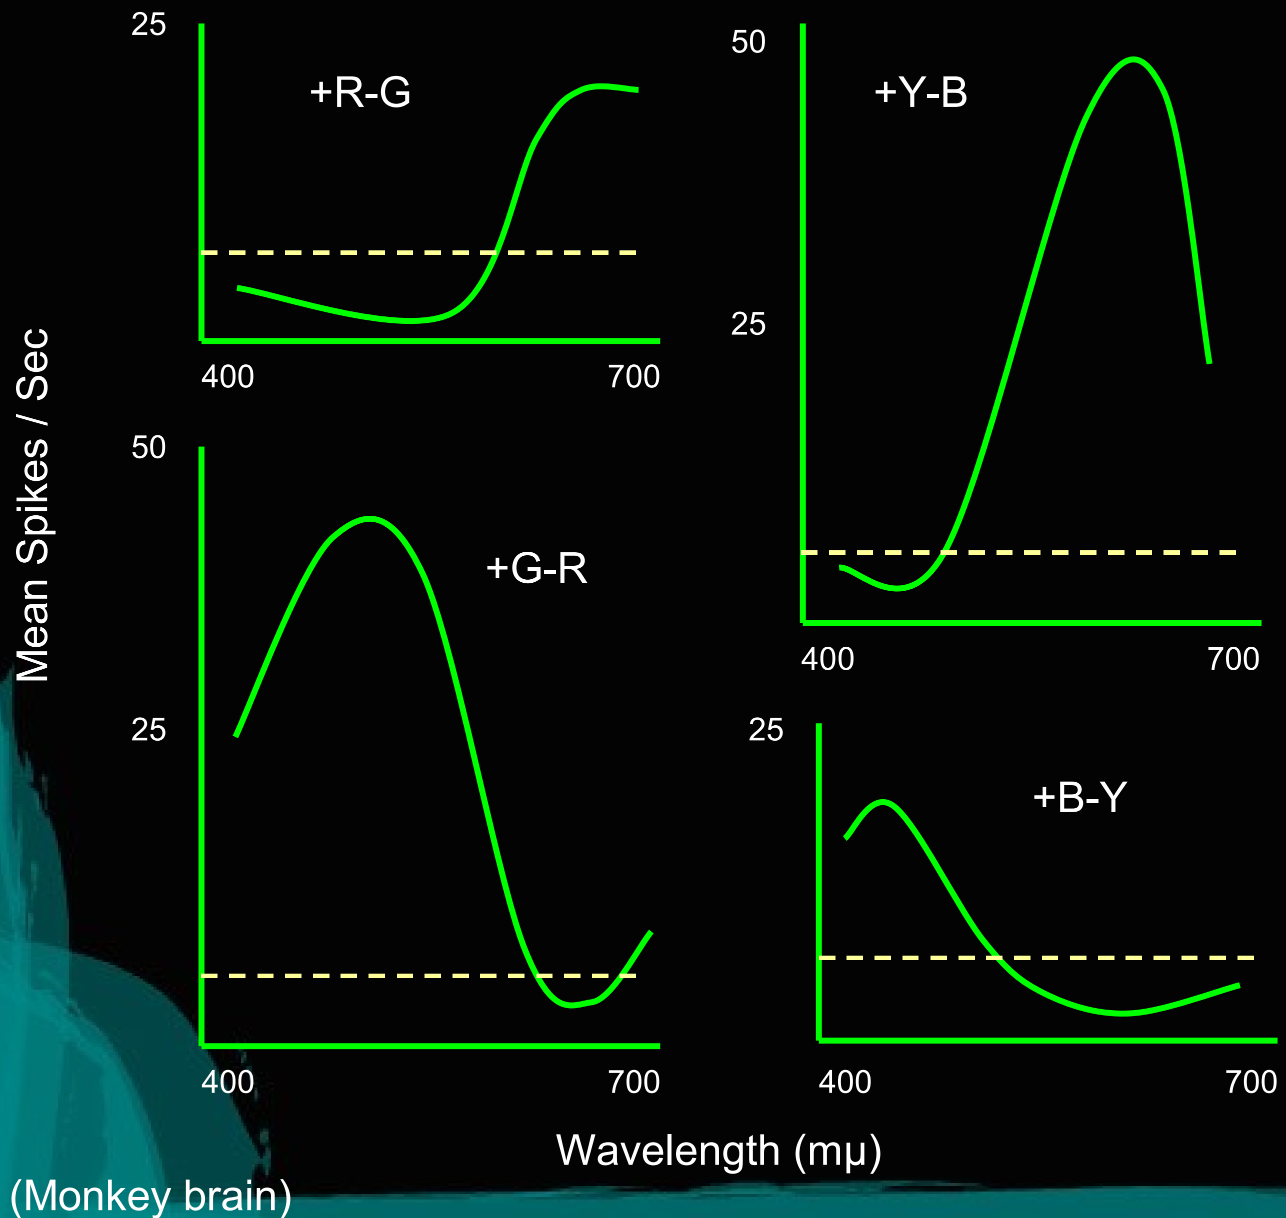

- These cells are found in the LGN
- Four color channels: Red, Green, Blue, Yellow
- R/G , B/Y pairs
- much like center/surround cells
- We can use these to determine the visual system's fundamental hue responses

The WCS Color Chips

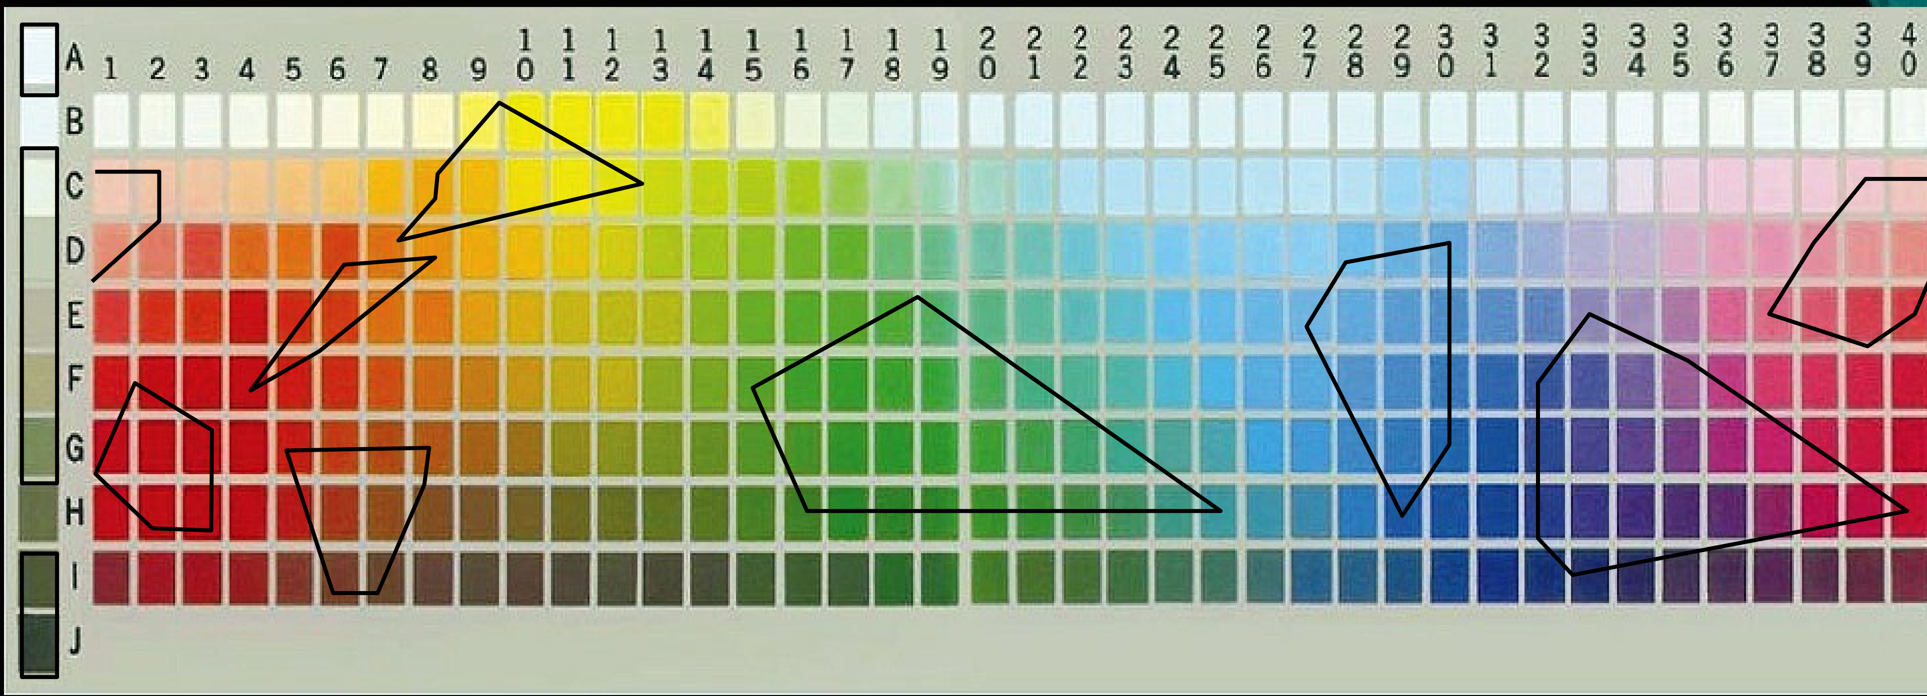

- Basic color terms:
  - Single word (not *blue-green*)
  - Frequently used (not *mauve*)
  - Refers primarily to colors (not *lime*)
  - Applies to any object (not *blonde*)

FYI:
English has 11
basic color terms

Results of Kay's Color Study

Stage I	II	IIIa / IIIb	IV	V	VI	VII
W or R or Y	W	W	W	W	W	W
Bk or G or Bu	R or Y	R or Y	R	R	R	R
	Bk or G or Bu	G or Bu	Y	Y	Y	Y
		Bk	G or Bu	G	G	G
			Bk	Bu	Bu	Bu
		W		Bk	Bk	Bk
		R			Y+Bk (Brown)	Y+Bk (Brown)
		Y				R+W (Pink)
		Bk or G or Bu				R + Bu (Purple)
						R+Y (Orange)
						B+W (Grey)

If you group languages into the number of basic color terms they have, as the number of color terms increases, additional terms specify focal colors