

Connectionism 5

More on Representation
Autoassociation

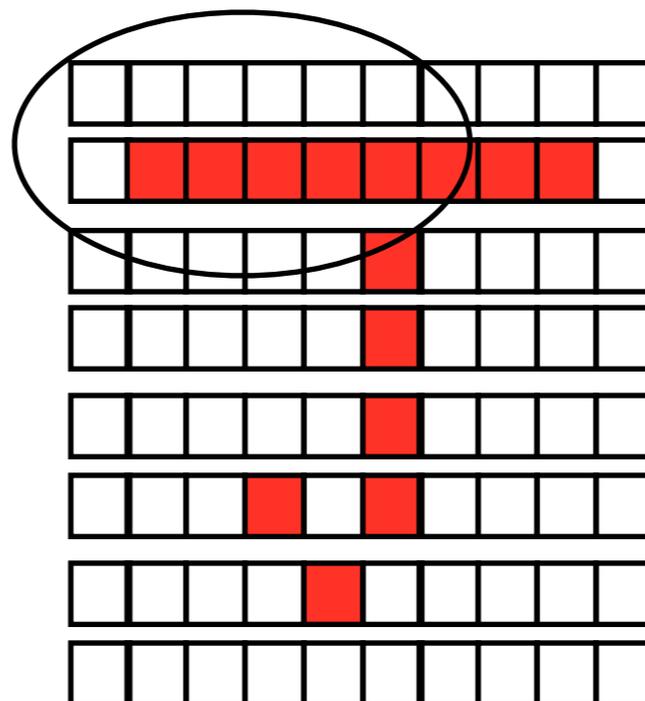
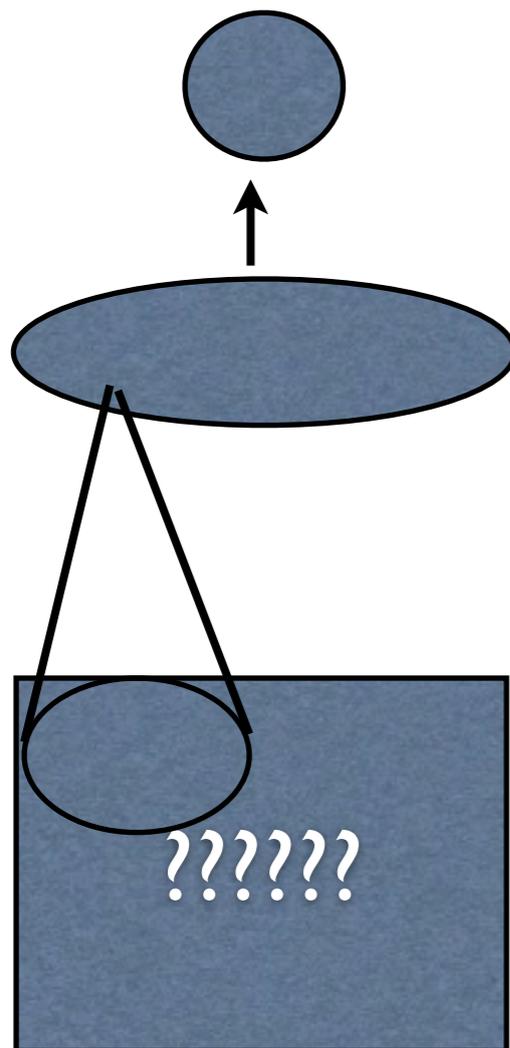
We chose to design our network to optimally 'represent' the target pattern.

Should we design representations, or discover them?

Our solution handled a limited kind of place invariance in a single dimension. How might one deal with scale invariance? rotation invariance?

Consider the task of representing letters.....

Task One: Handwriting recognition



Problem: 'j's vary a whole lot.....

8

4

4

5

4

5

5

4

4

...

4

5

4

What is in a 'j'?

Look for a crossbar

Look for a down stroke

Look for an up hook

As more features are found, the likelihood that we are seeing a 'j' increases.

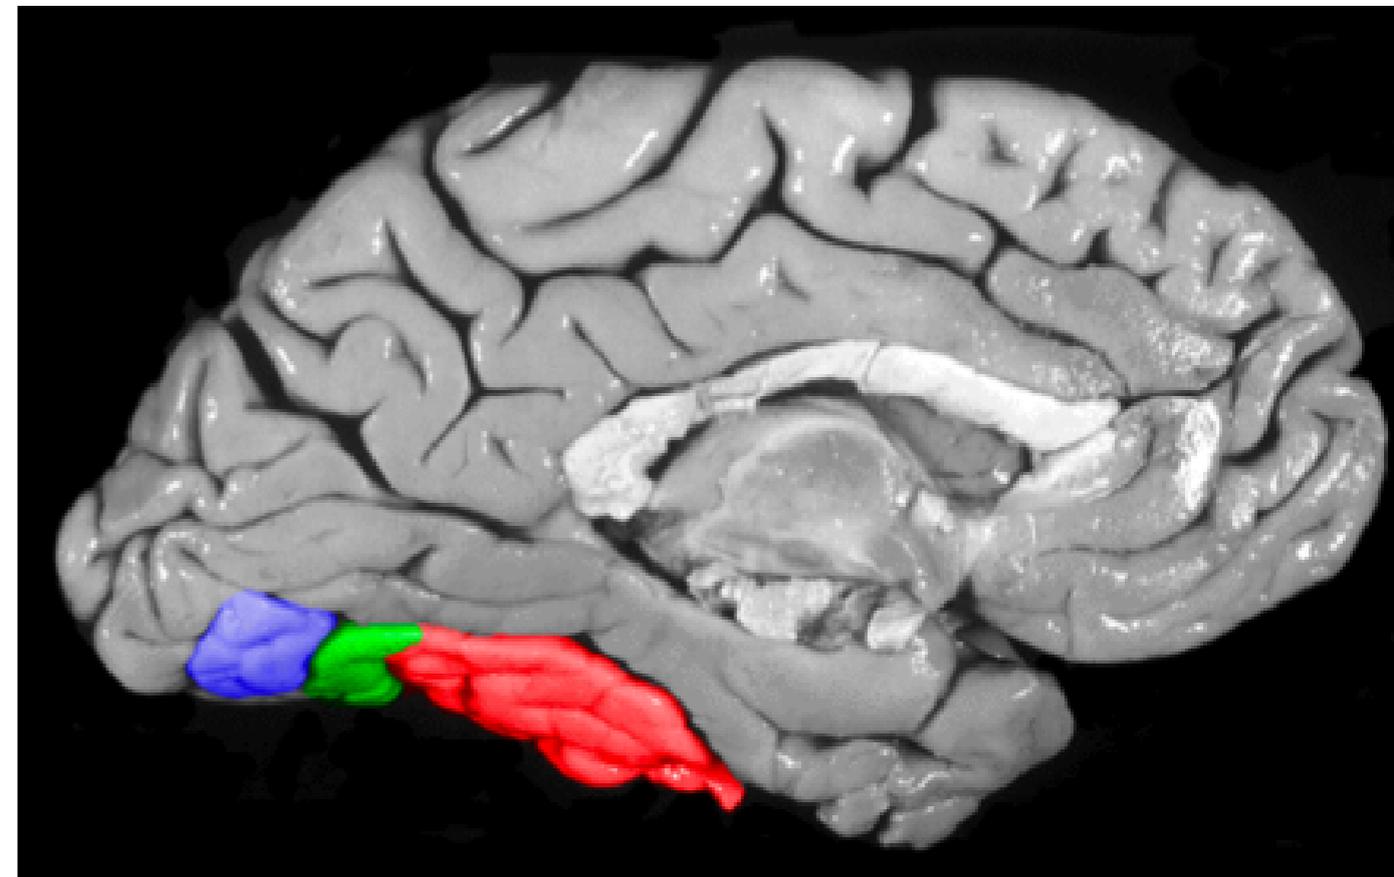
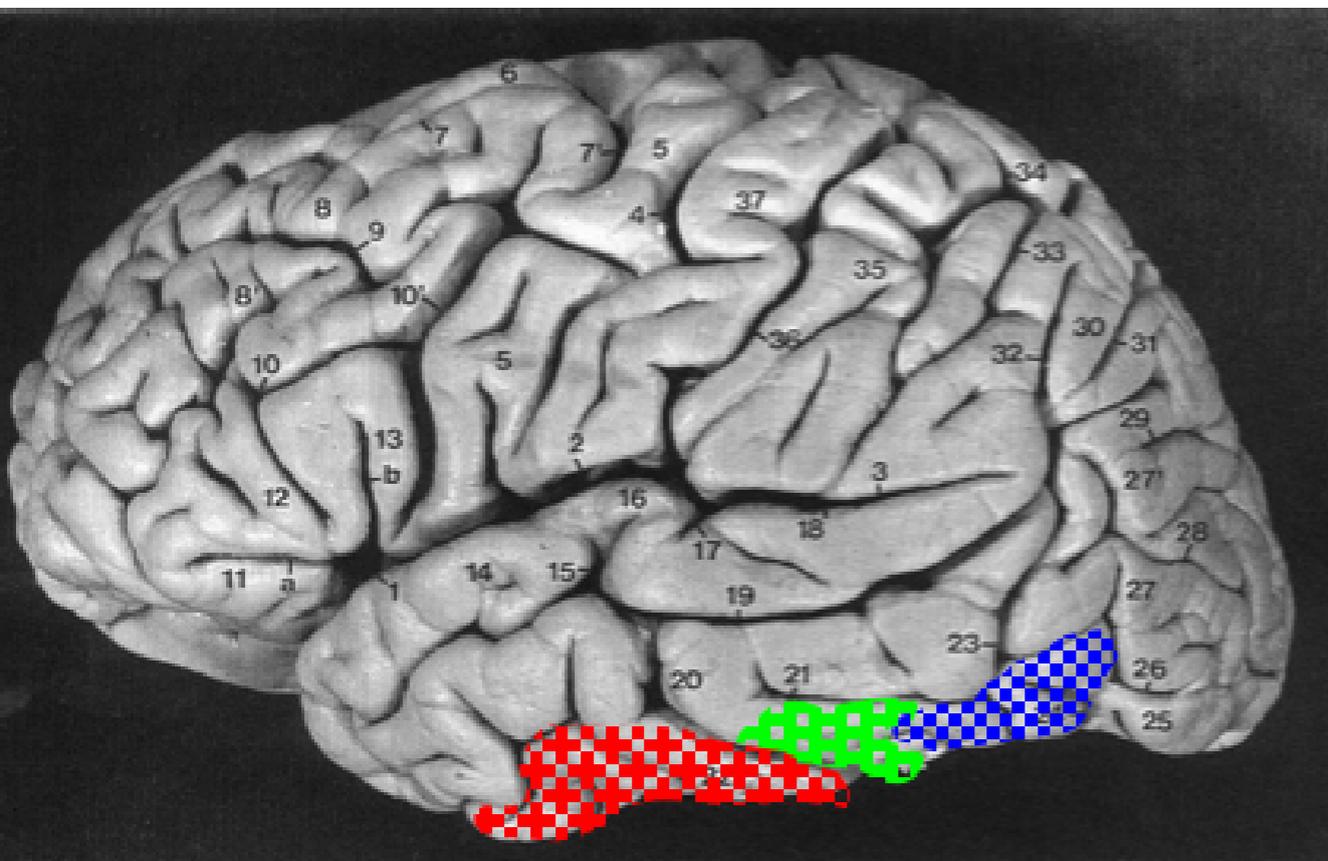
What kind of representations does the brain use?

(Does the brain use representations?)

Here, we revise our discussion of representations from the Introductory module, but with the added consideration of what the consequences are for this kind of computational model

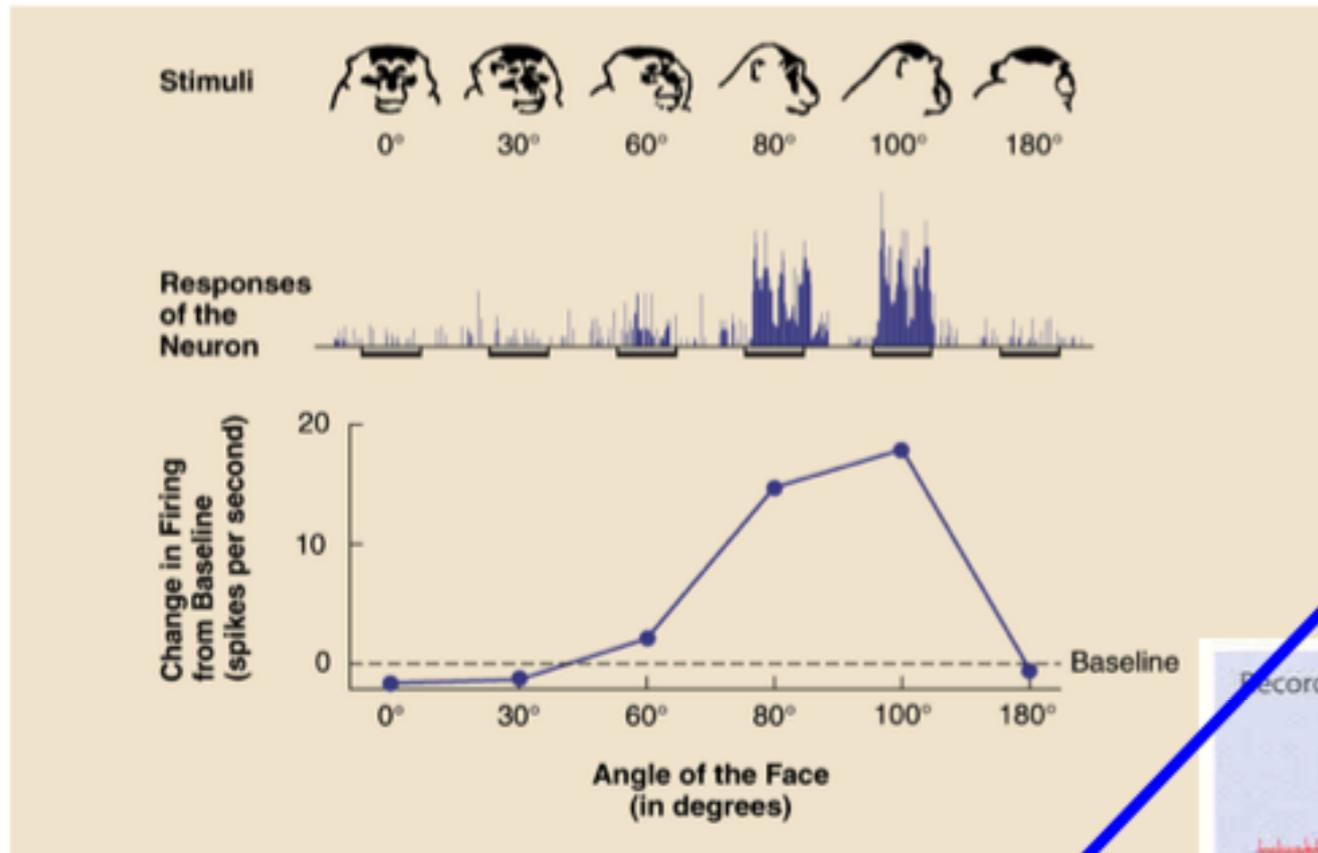
Some Neurophysiology

- The inferotemporal cortex (IT) contains neurons that respond selectively to particular objects, or complex shapes
 - like bananas or a circle with a T shape attached to it - different studies find different sensitivities.



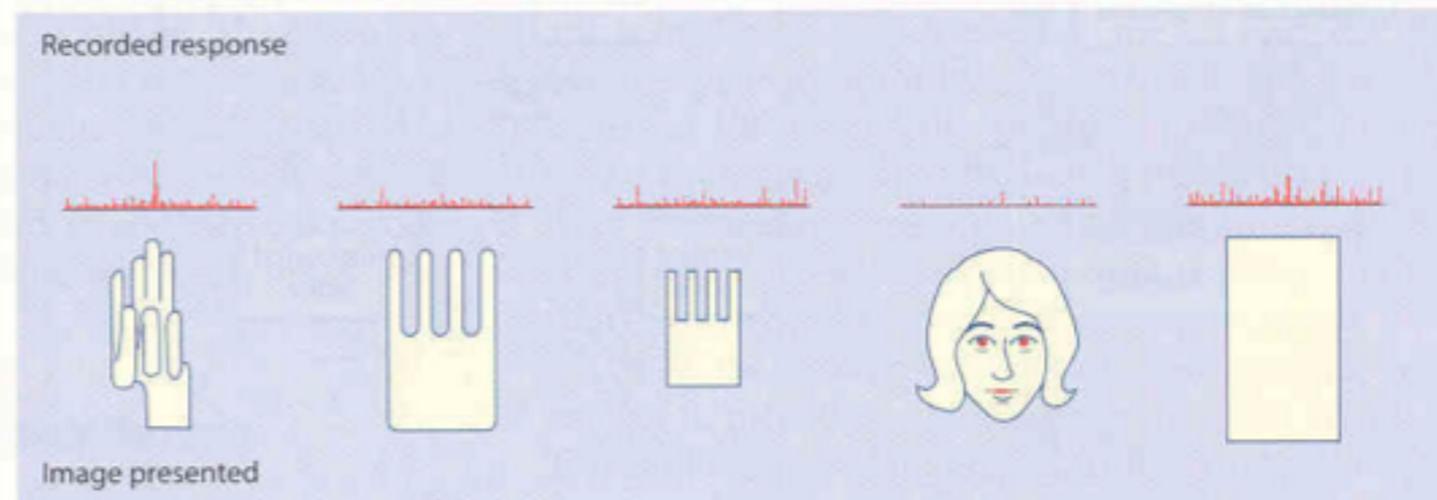
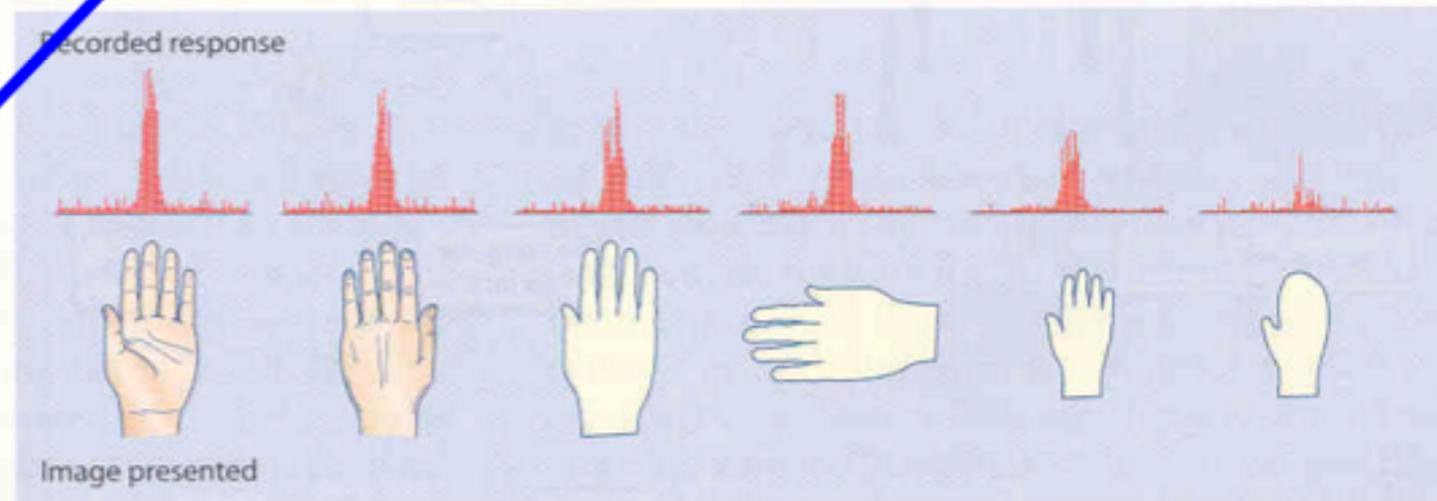
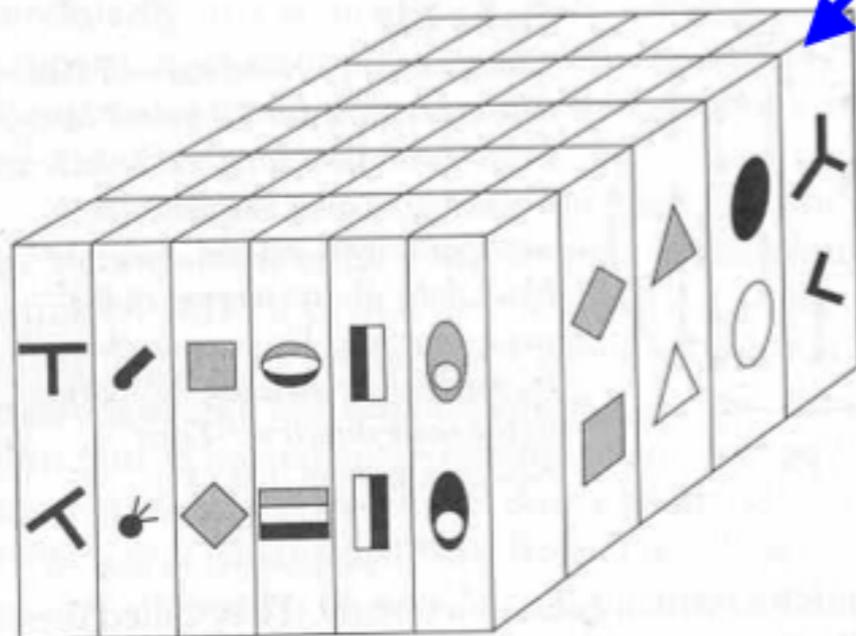
Tuning in monkey IT

► Firing Rate of a Monkey Inferotemporal Cortex Neuron



Cells respond to specific stimuli
And are clustered into columns with similar tuning

Source: Adapted from Gross et al., *Pattern Recognition Mechanisms*, p. 179-201. Berlin: Springer-Verlag 1985.

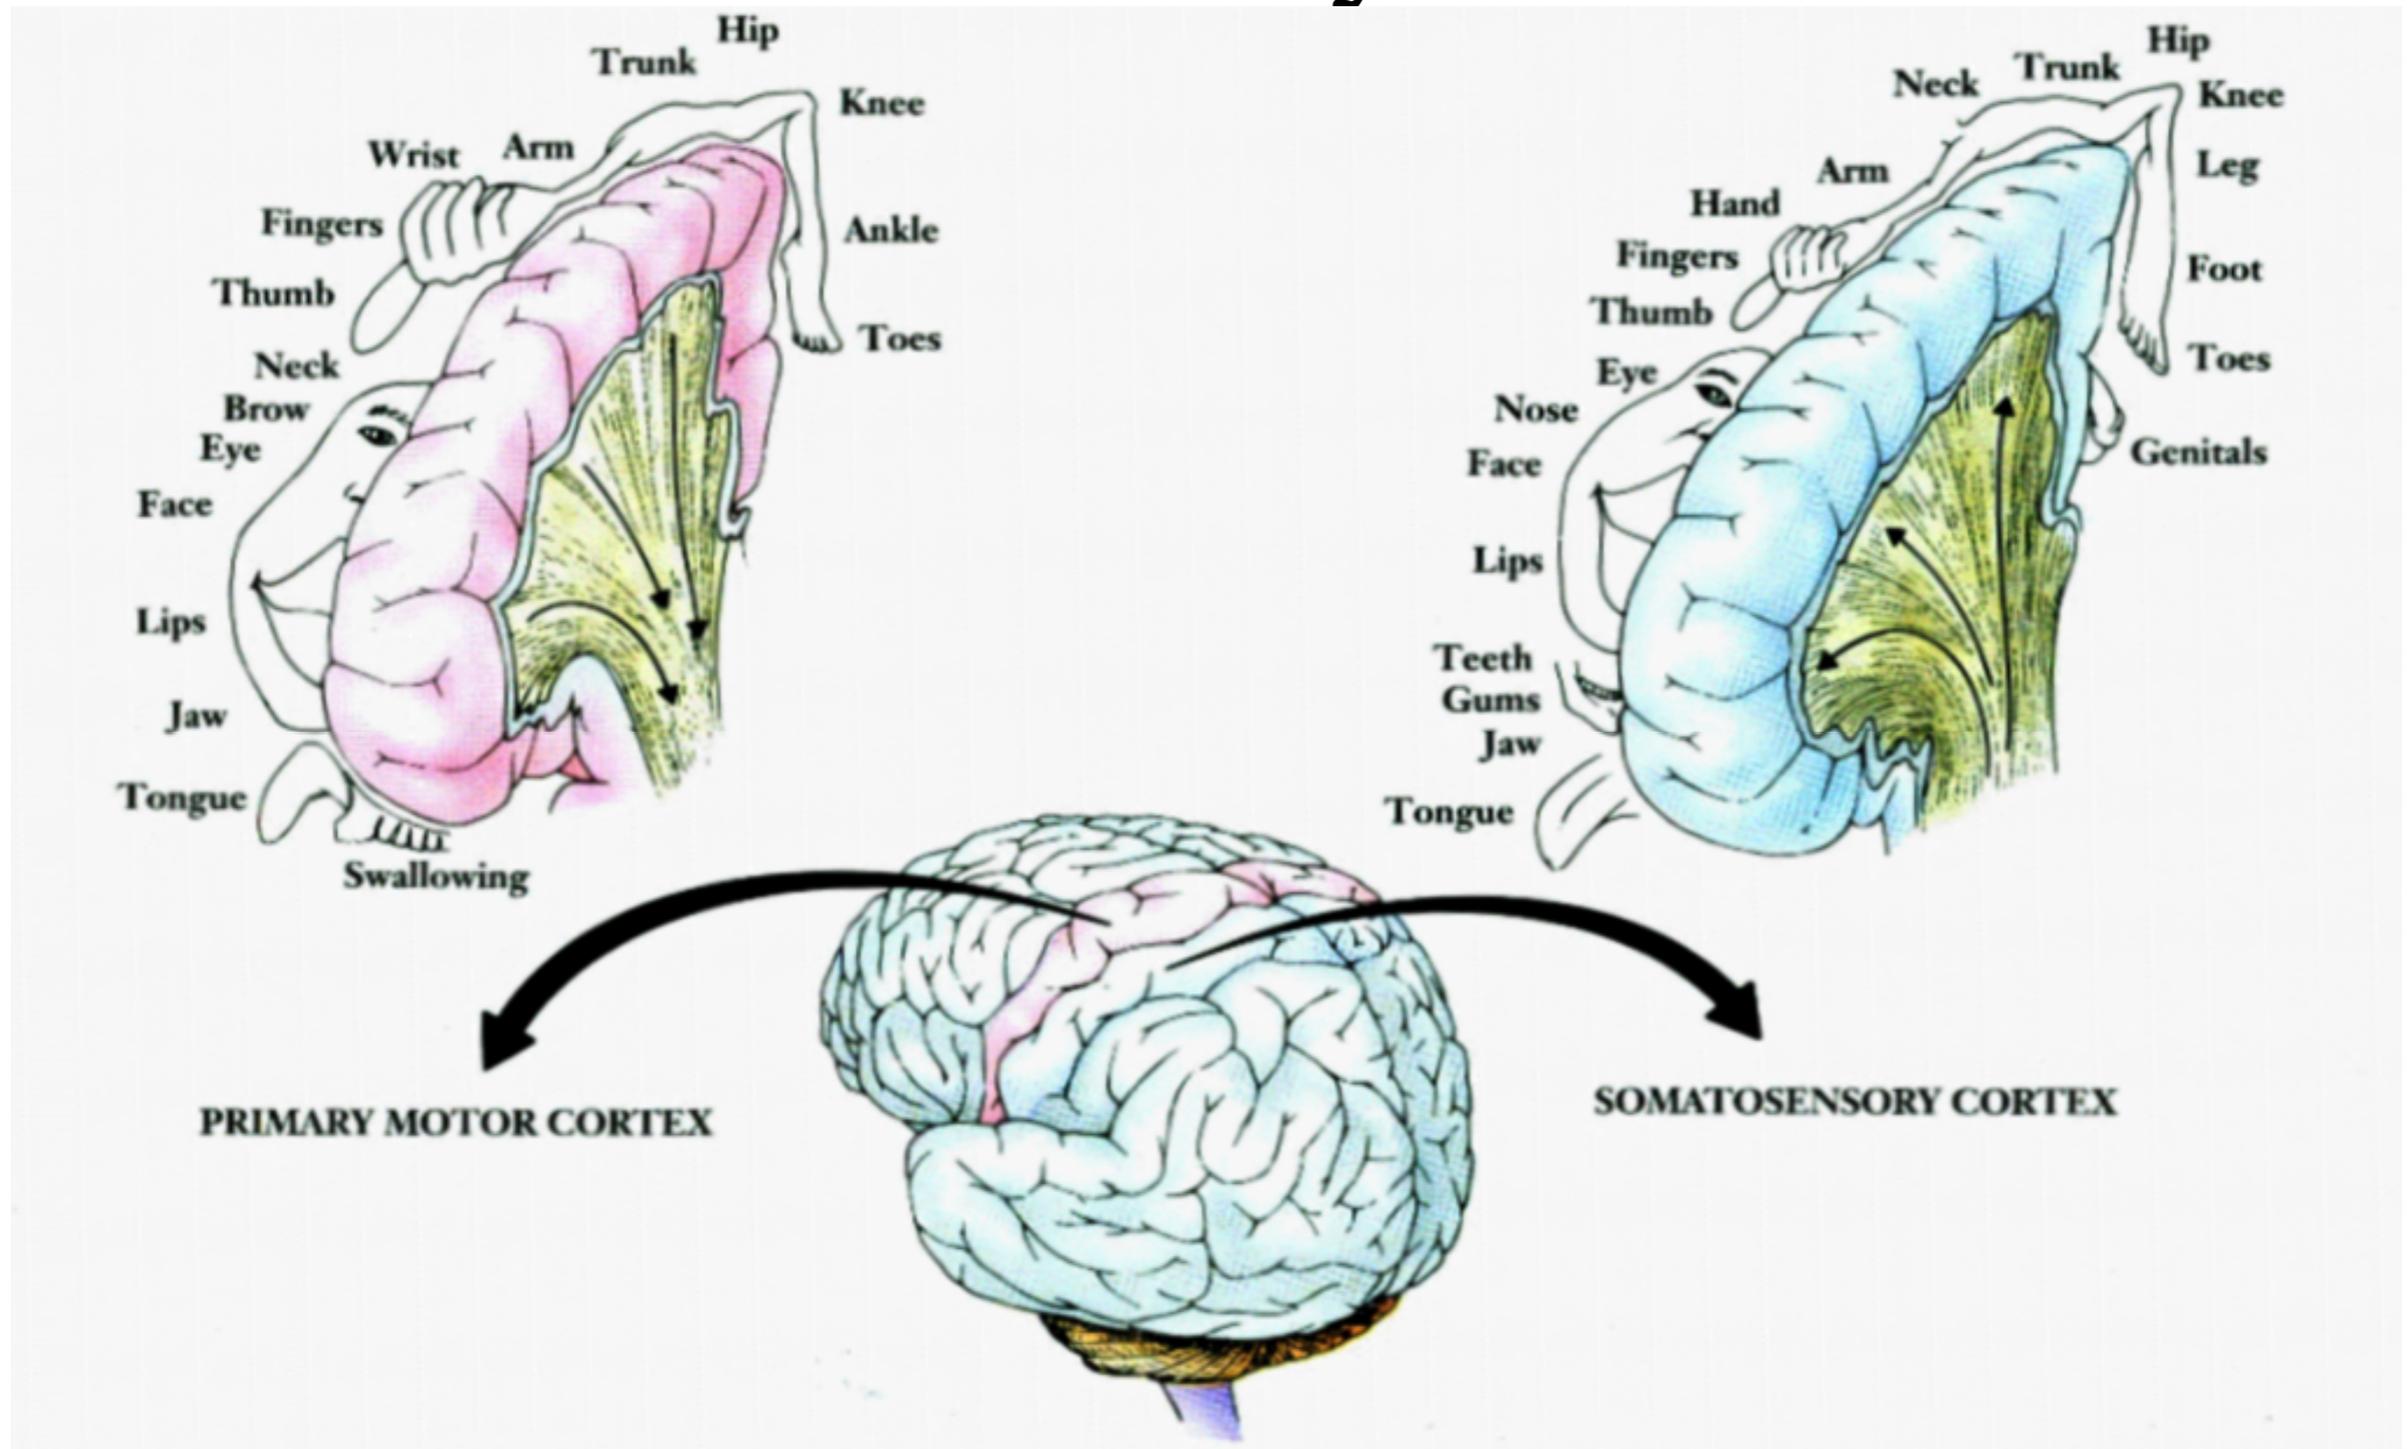


Representation: primary sensory cortex

Cells in primary sensory cortex react to first order properties of the input: e.g frequency (audition), presence of an edge at a specific orientation in a specific place.

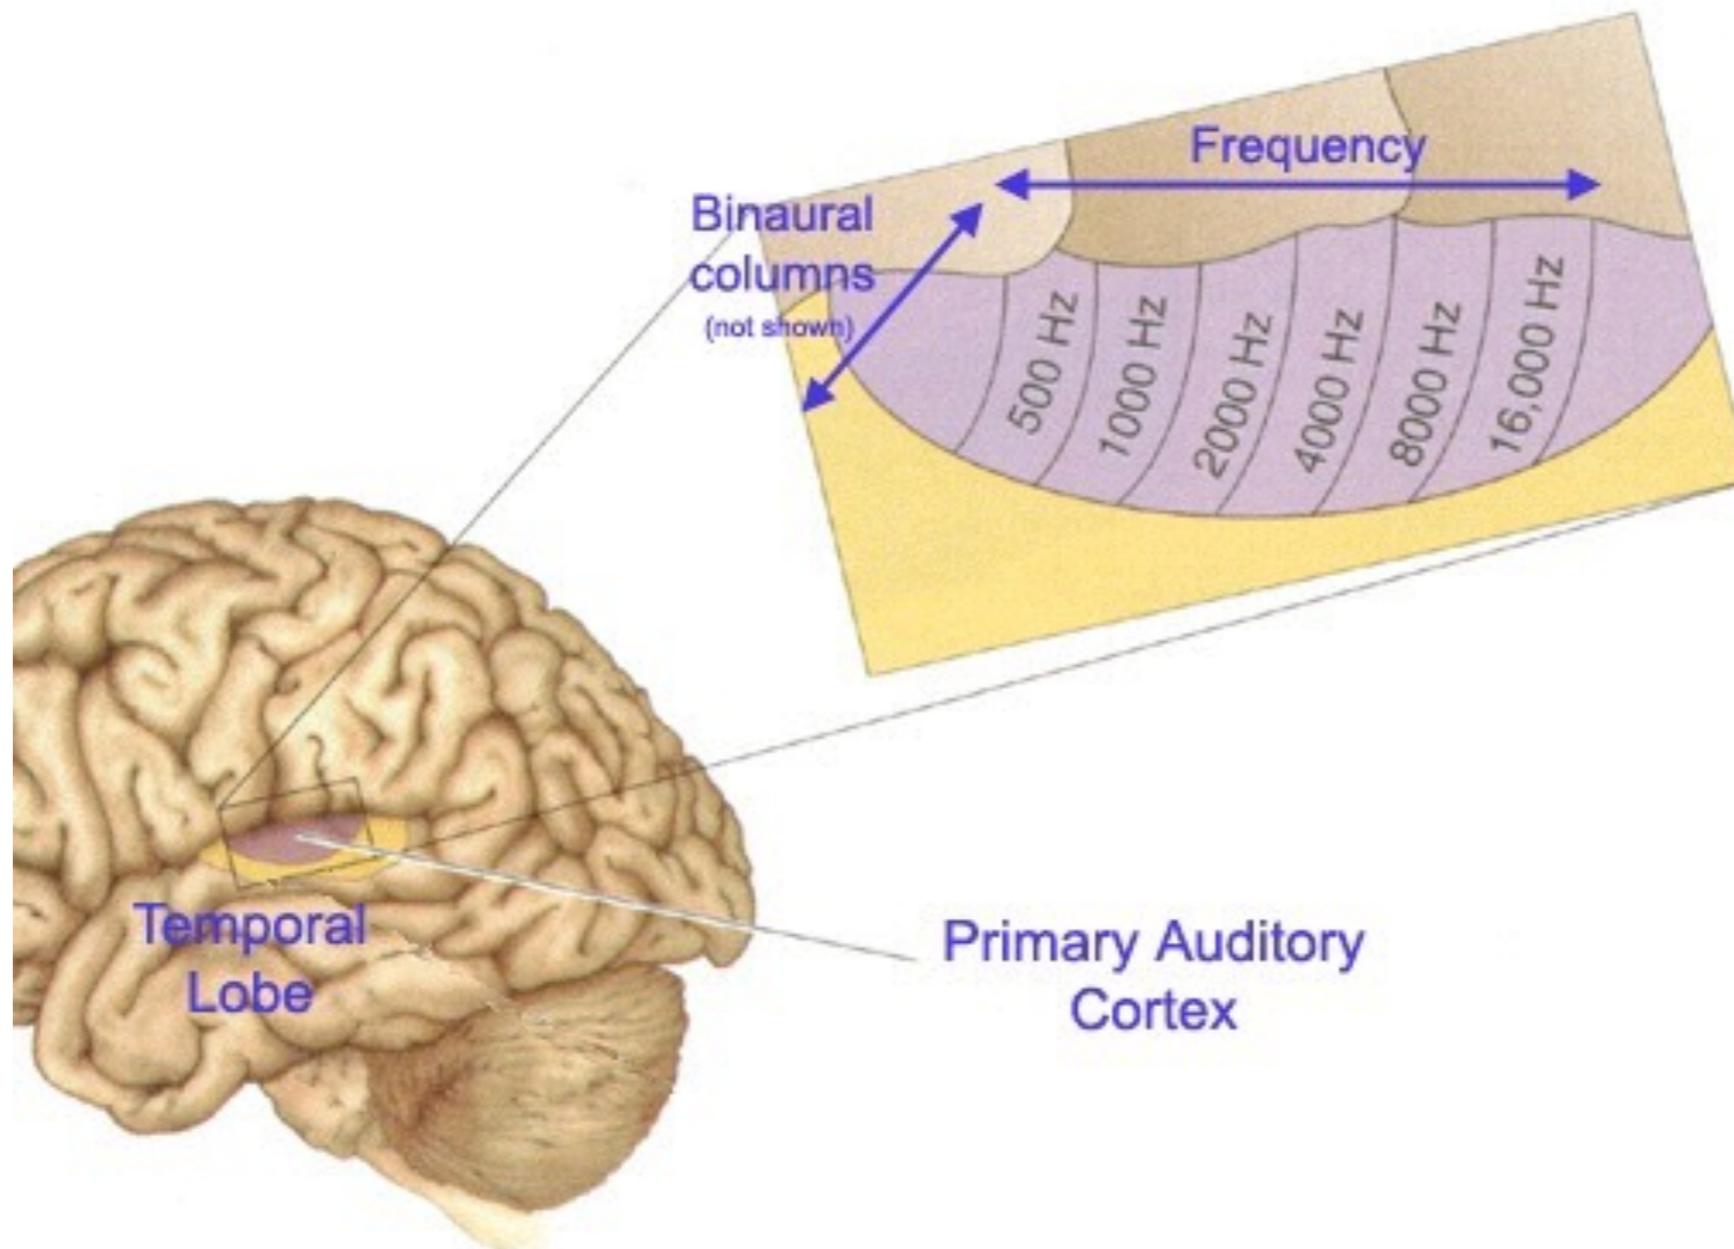
This can be implemented by applying a fixed filter to sub-sections of the input field (as we saw last week).

Cortical Pathways of Touch



Primary Auditory Cortex

Tonotopic Map Has Columnar Organization



As we ascend the processing pathways, cells are selectively responsive to more and more abstract properties

Nobody has yet found the “Grandmother Cell”

...But some cells are very specific to high-order properties of the input.. e.g. face specificity

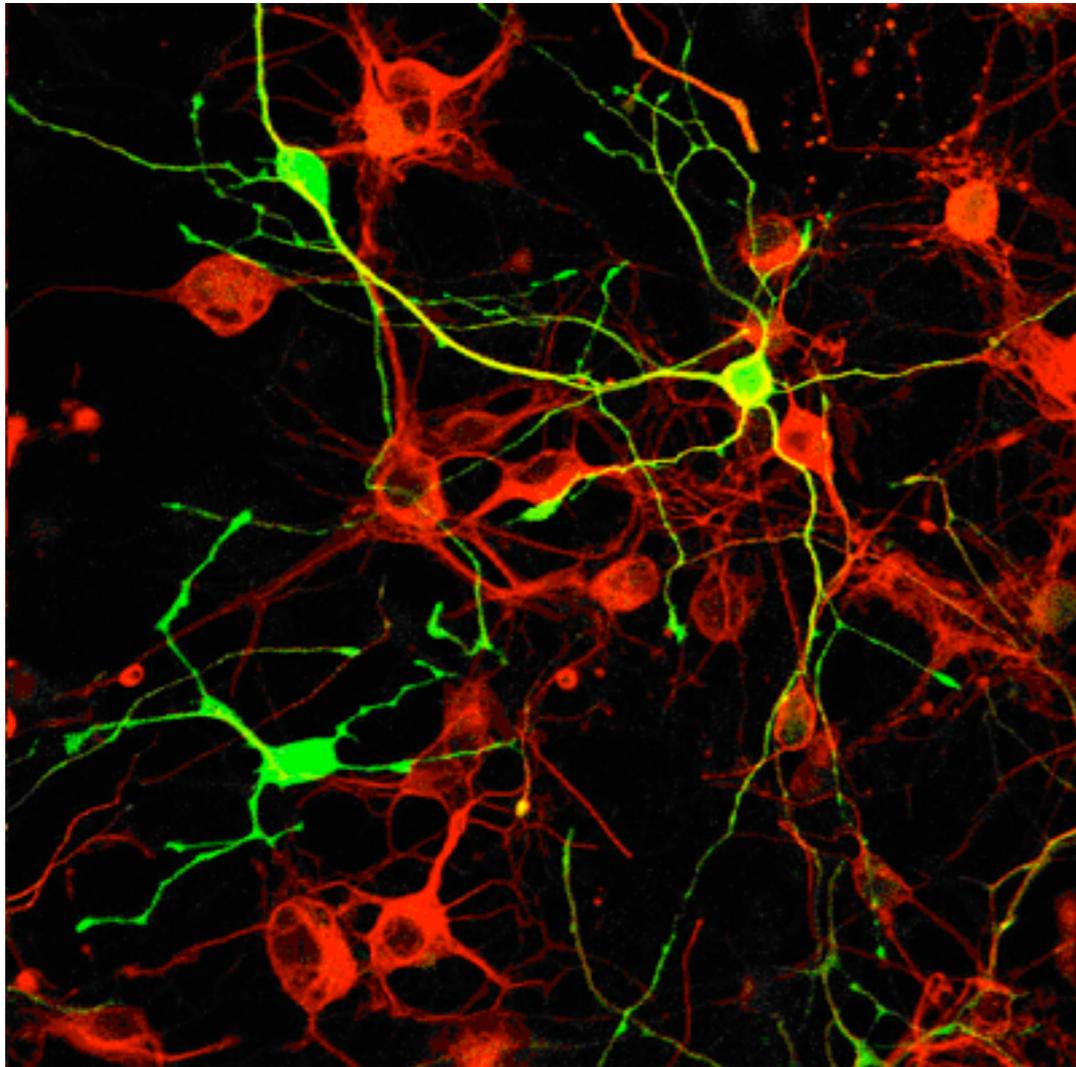






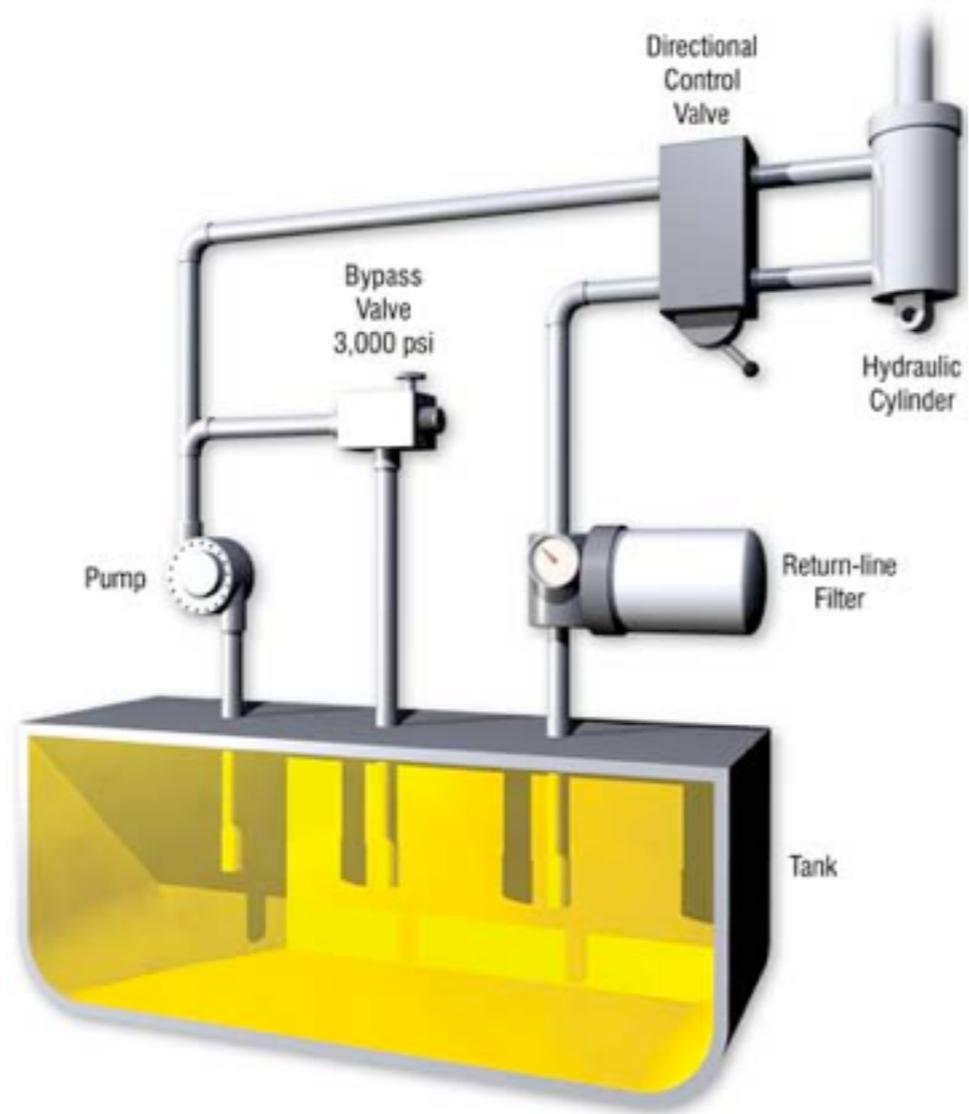


Neuroscience
studies nervous
systems....
brainzzzzzz 'n'
stuff



In what terms do we interpret findings about neurons?

What do you think brainzzzzz do?

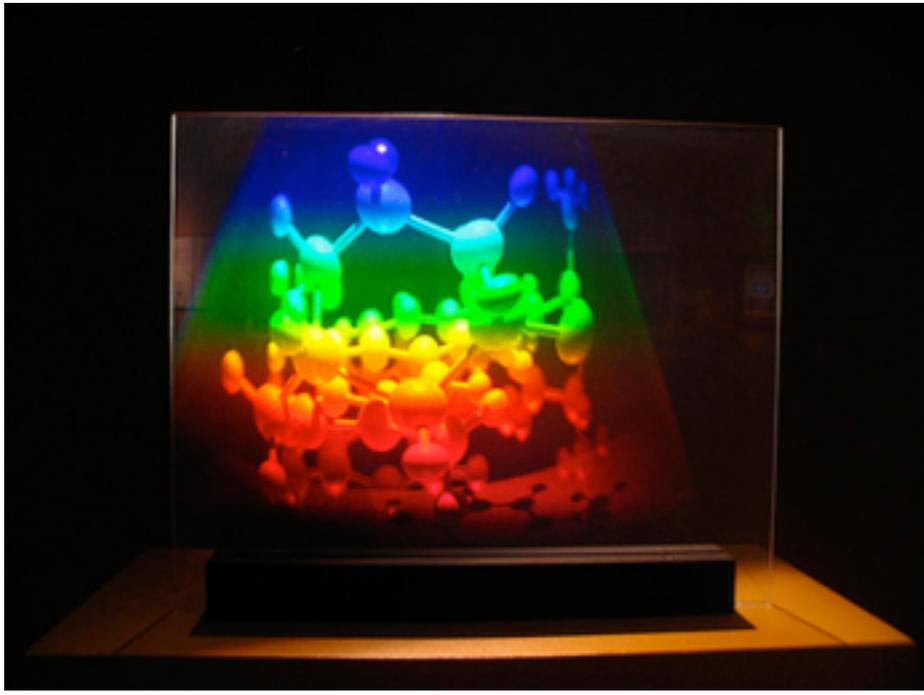


The brain is a hydraulic system
(18th Century)



The brain is a calculating engine

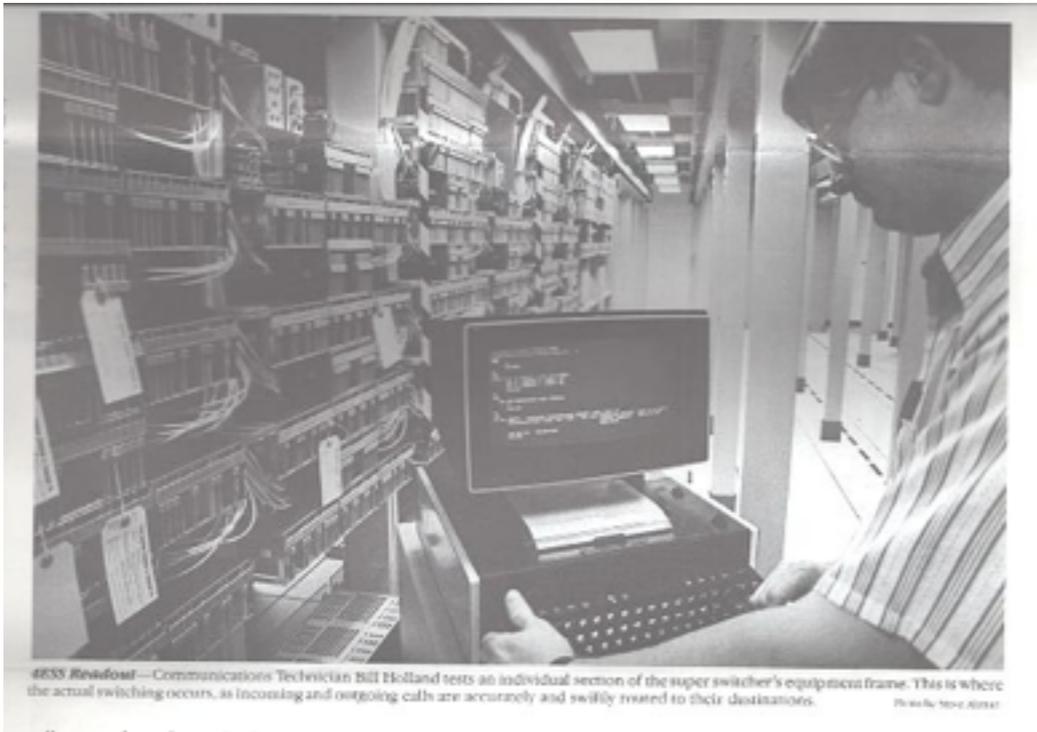
(19th Century)



The brain is a hologram

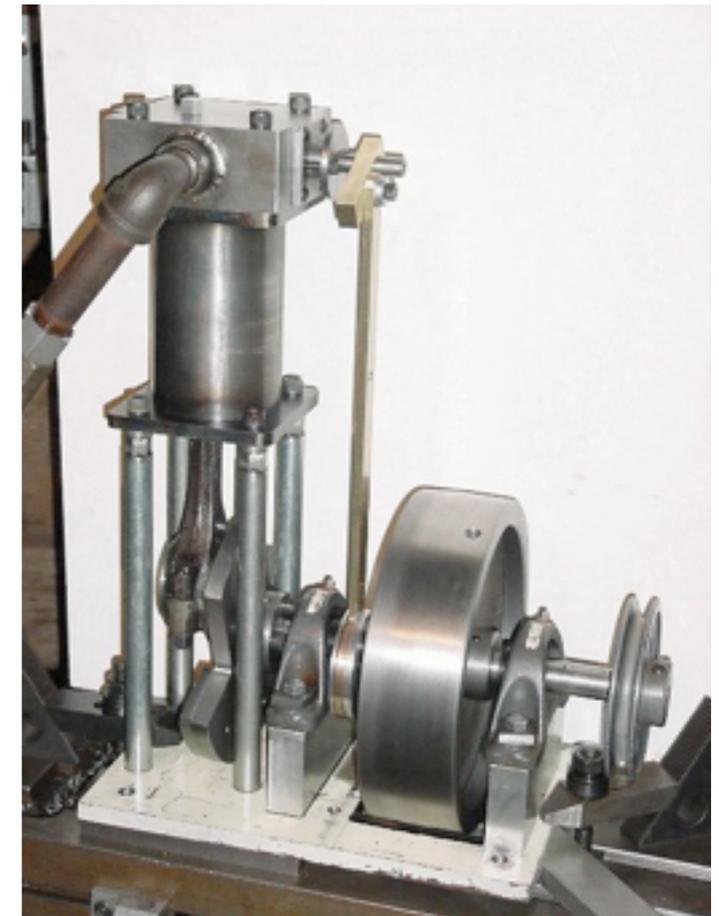
The brain is a computer

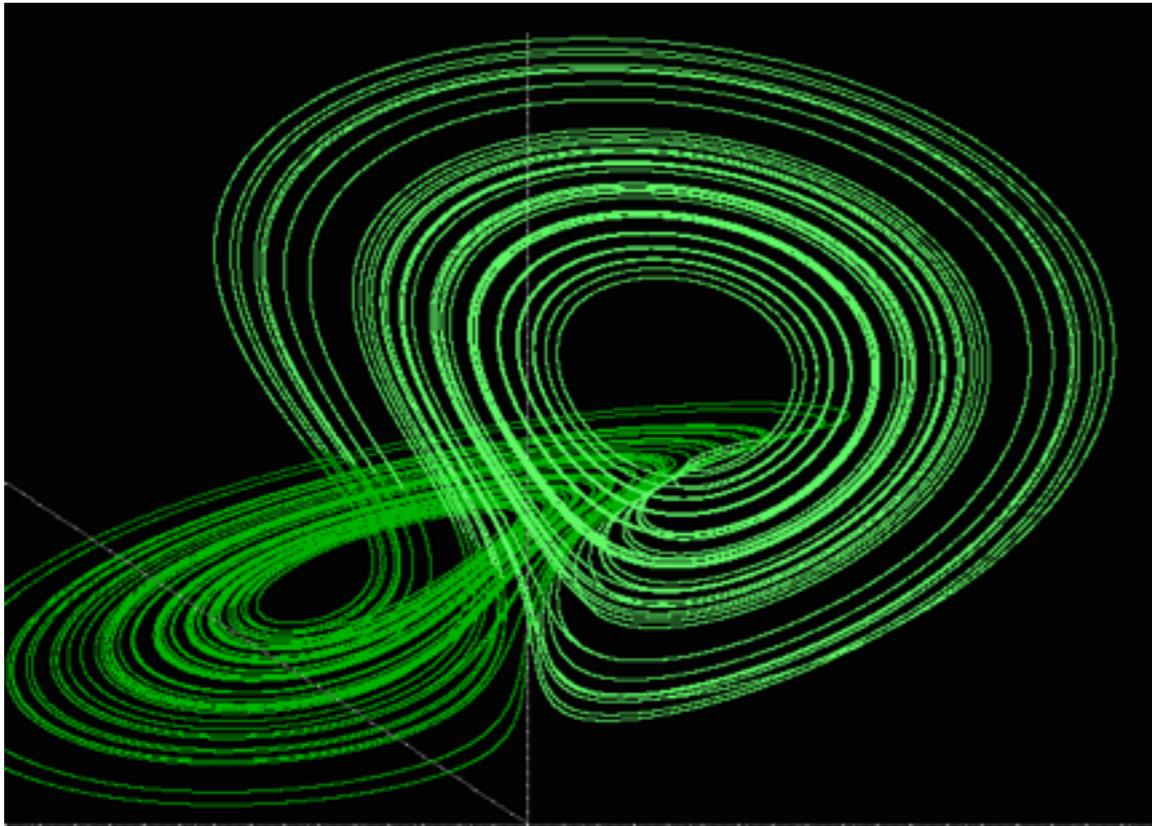




A switching network

A steam engine





The brain is a dynamical system

... and so on



Neuroscience

Psychology

cell

memory

action potential

recognition

ganglion

emotion

thalamus

intention

receptive field

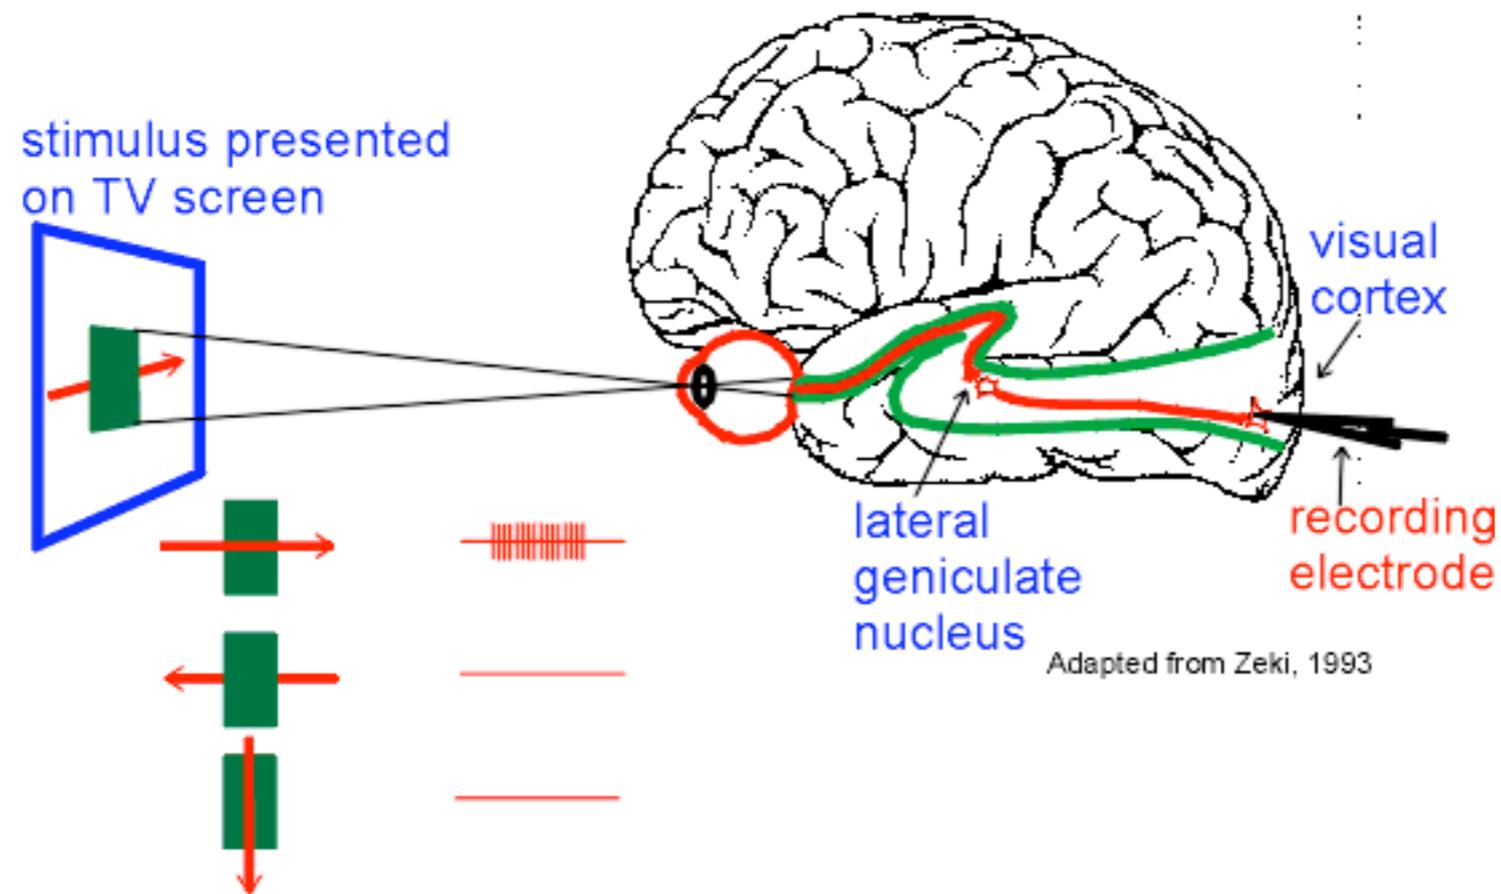
belief

There is, as yet, little agreement on
what brains do

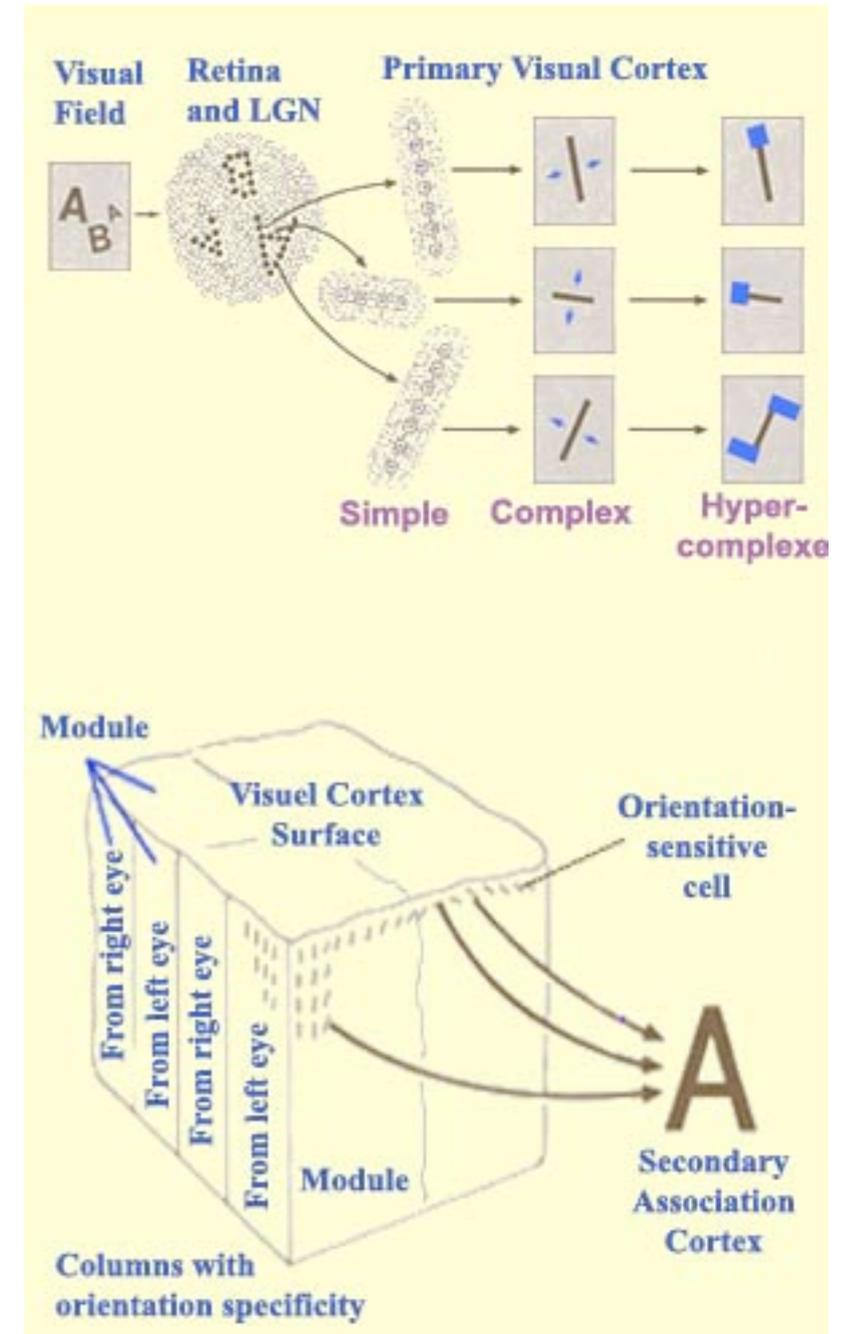
Visual representation

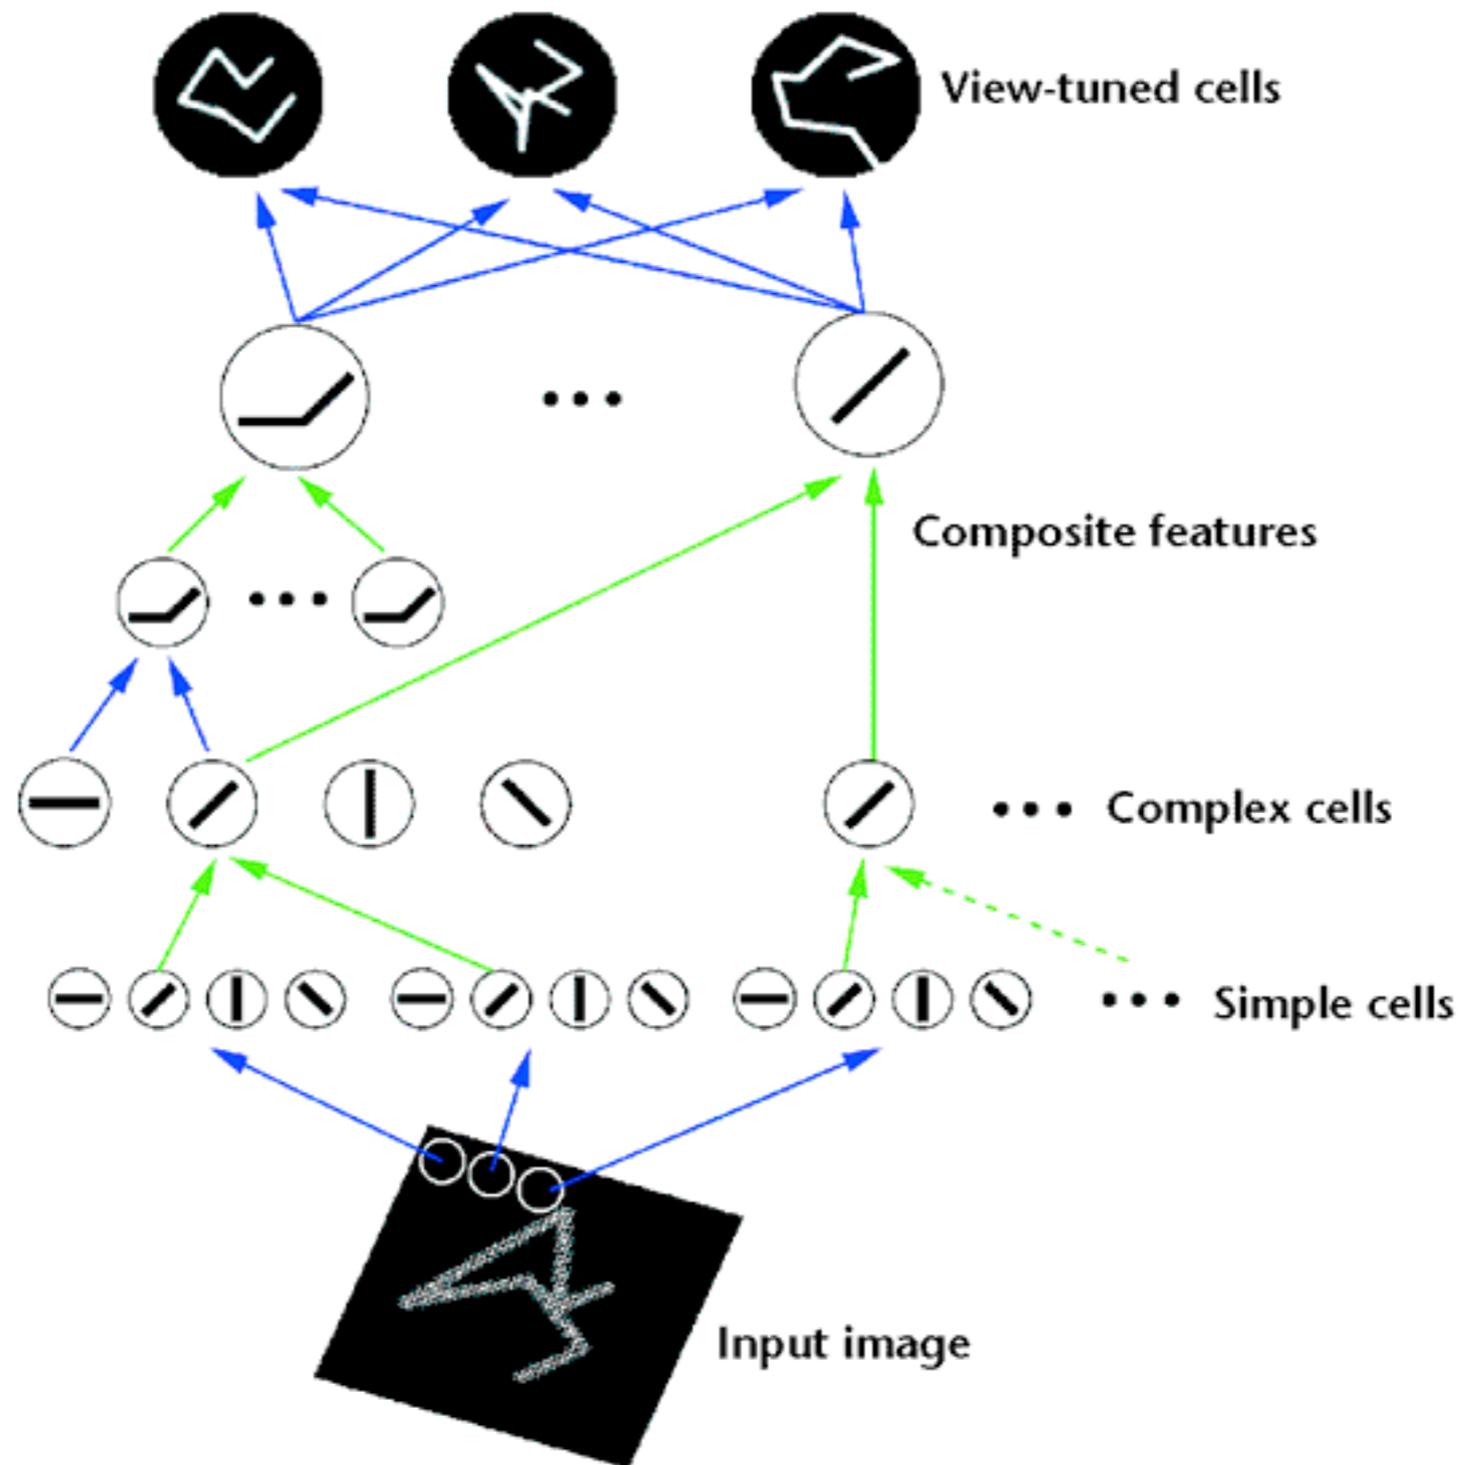


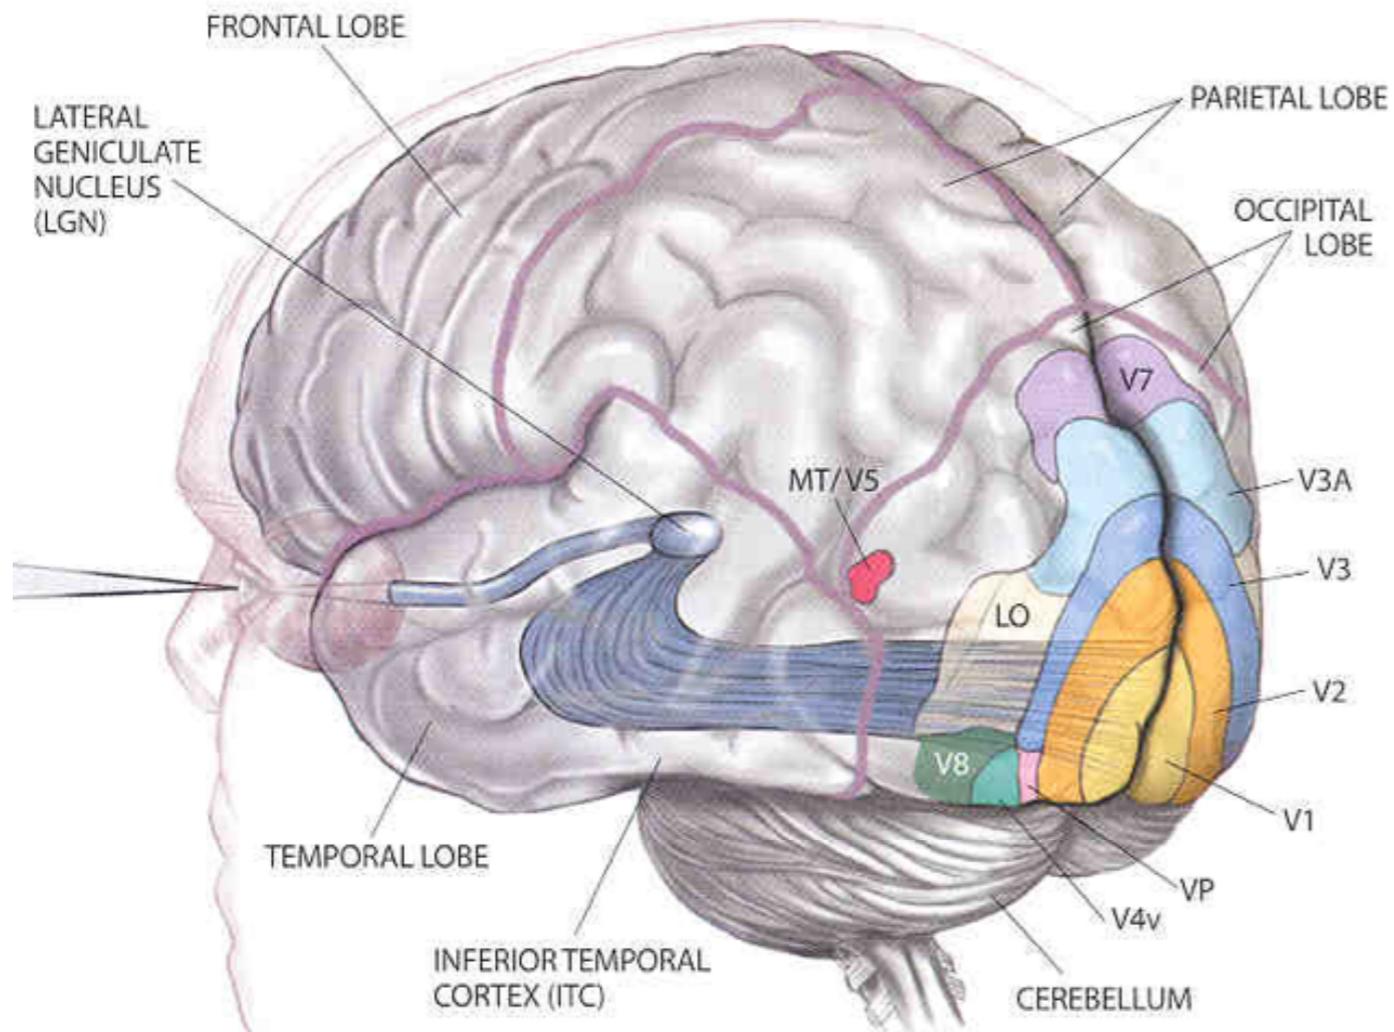
Hubel and Wiesel: Feature Detectors



Hypothetical construction of a complex representation from simple parts



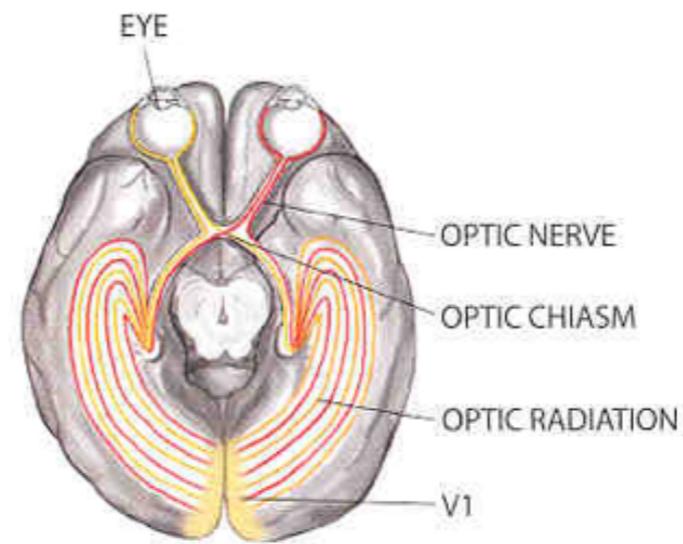
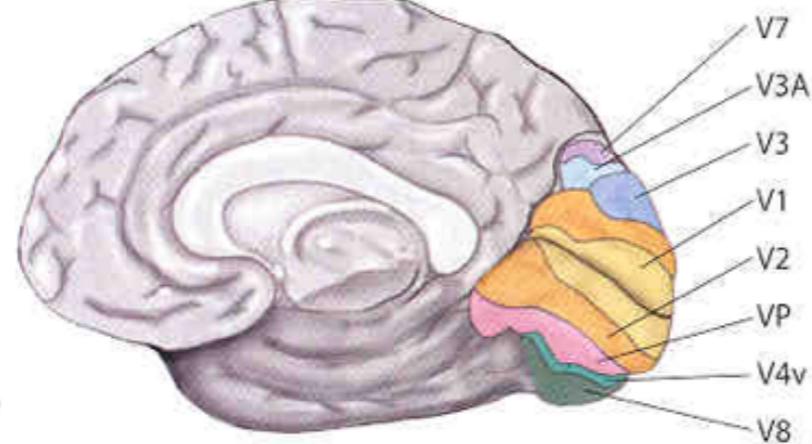




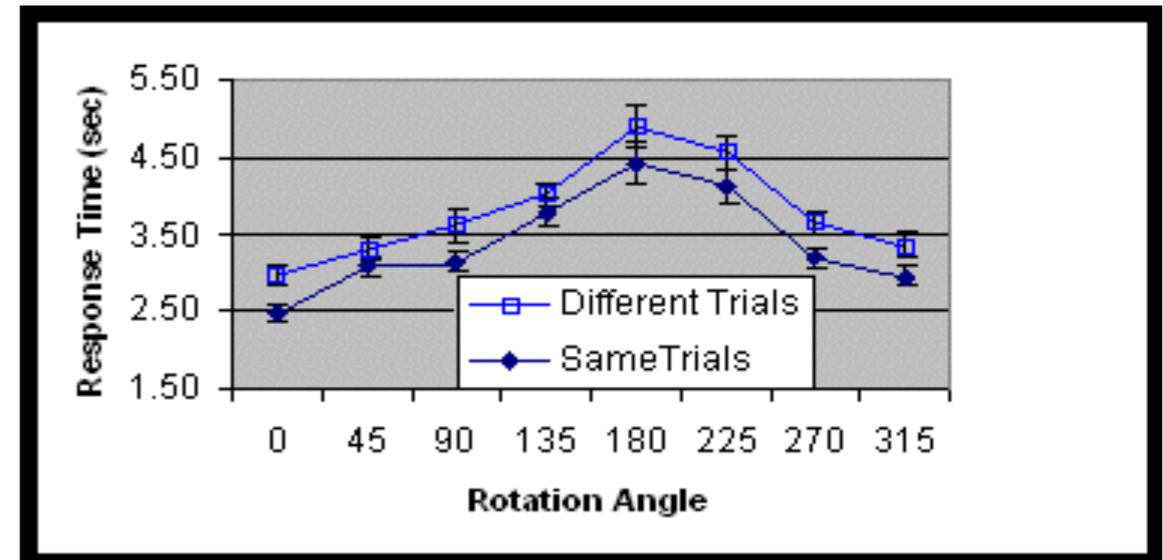
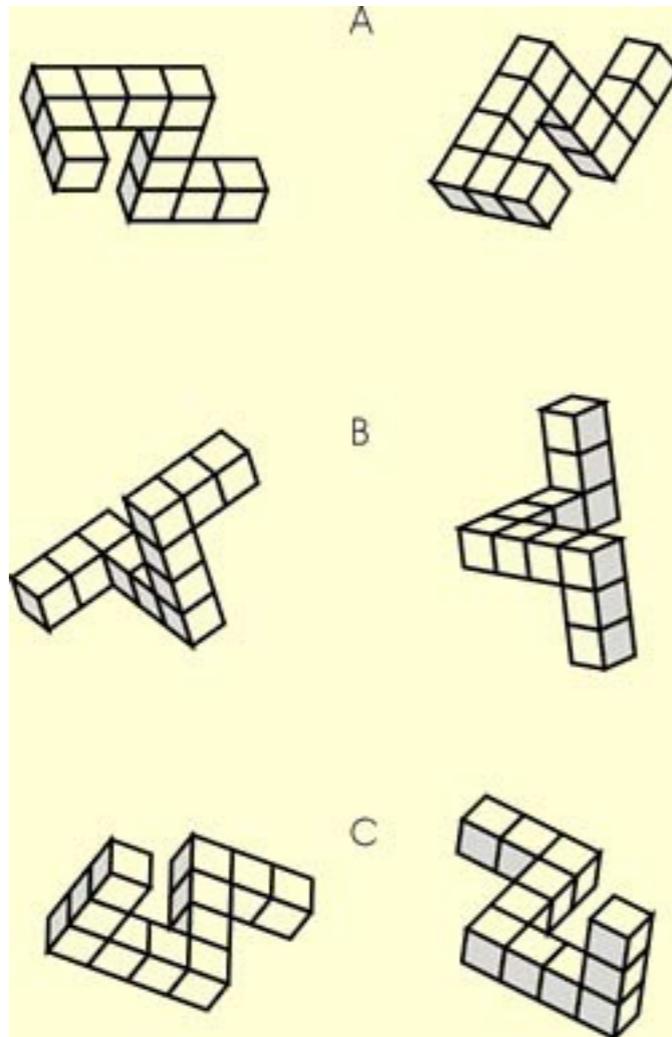
KEY TO FUNCTION

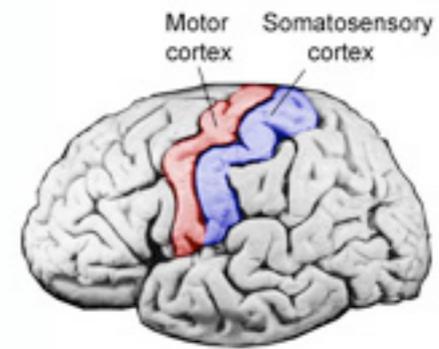
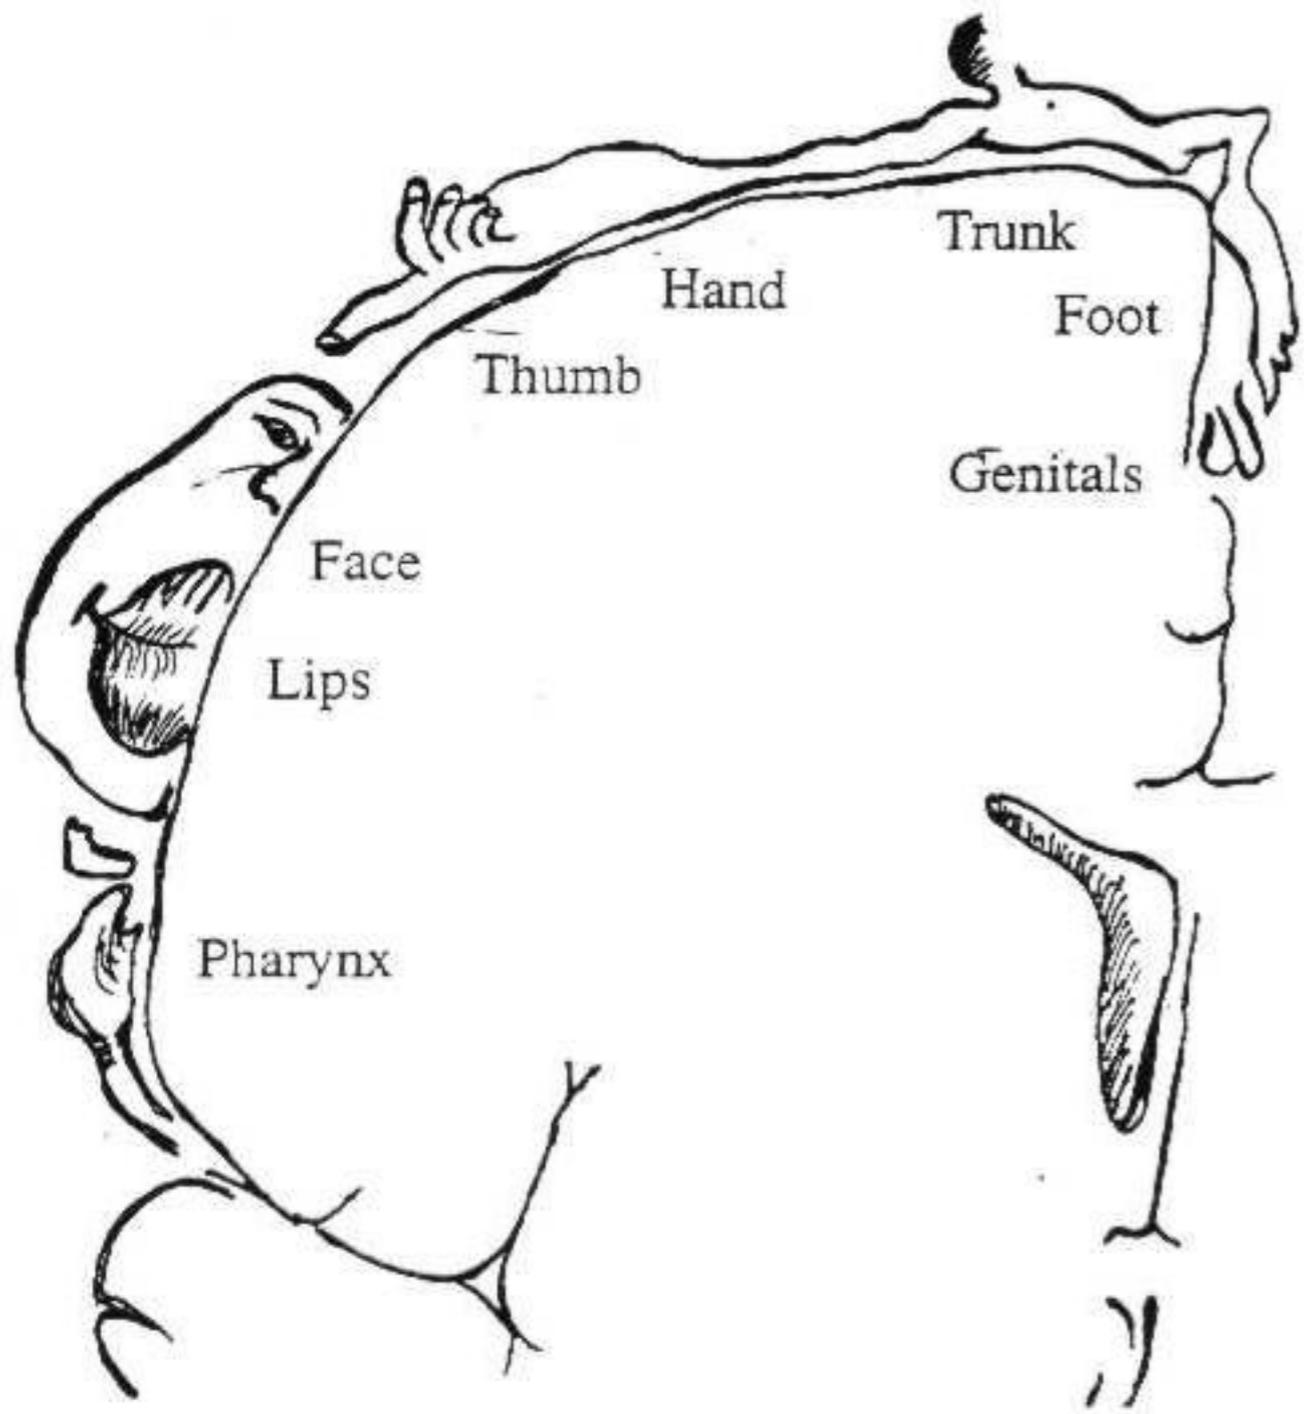
- V1:** Primary visual cortex; receives all visual input. Begins processing of color, motion and shape. Cells in this area have the smallest receptive fields.
- V2,** **V3** and **VP:** Continue processing; cells of each area have progressively larger receptive fields.
- V3A:** Biased for perceiving motion.
- V4v:** Function unknown.
- MT/V5:** Detects motion.
- V7:** Function unknown.
- V8:** Processes color vision.
- LO:** Plays a role in recognizing large-scale objects.

Note: A V6 region has been identified only in monkeys.



Representations everywhere.....





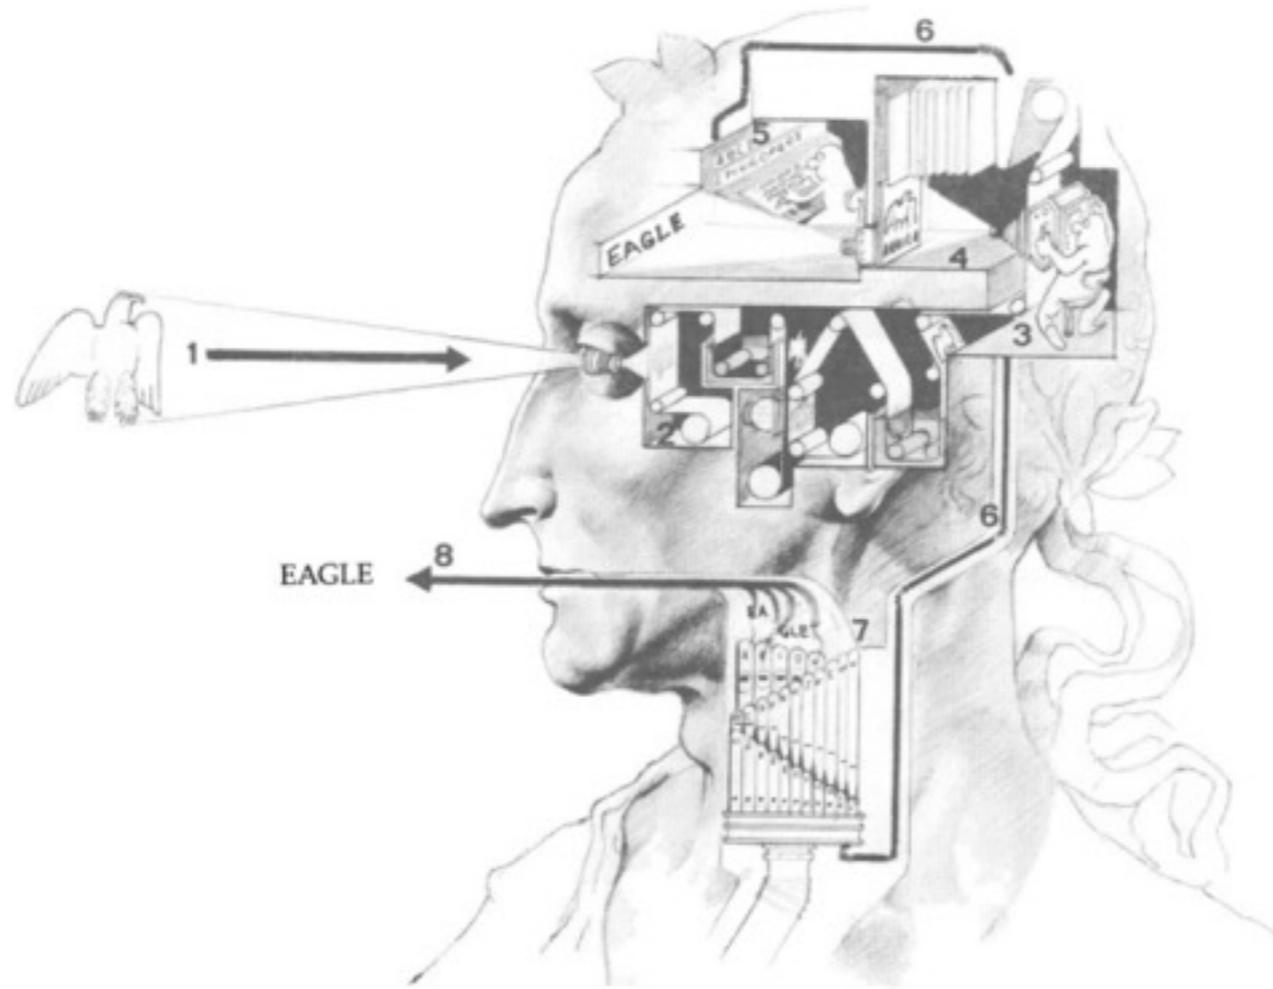


Fig. 34. Caesar according to the representationist metaphor.

Problems

Anesthetized cats are a bit special

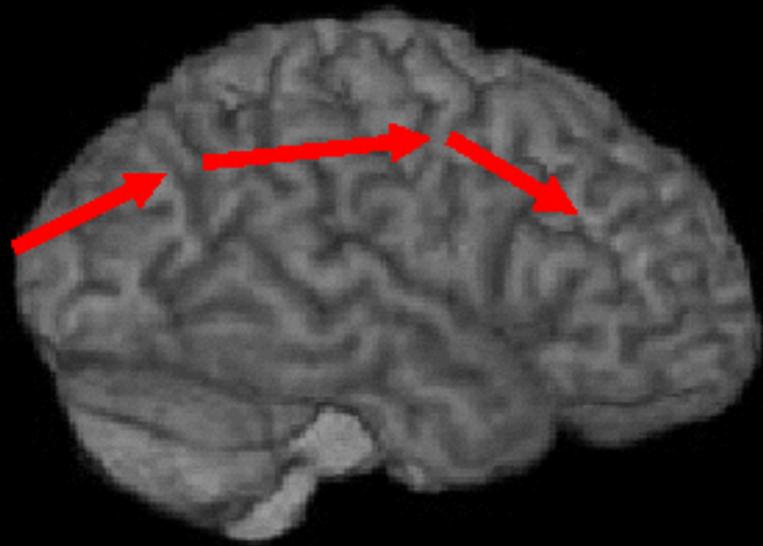


Nature **291**, 554 - 561 (18 June 1981); doi:10.1038/291554a0

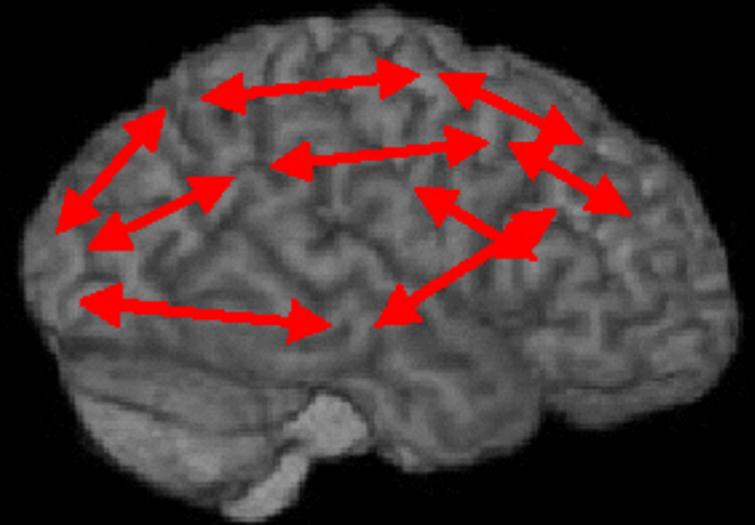
Effects of sleep and arousal on the processing of visual information in the cat

Margaret S. Livingstone* & David H. Hubel

Quick and dirty summary: Livingstone and Hubel made history simply by pinching the tail of their anesthetized cat, thereby waking it up, and observed that the same visual neurons now fired differently.



50 Years

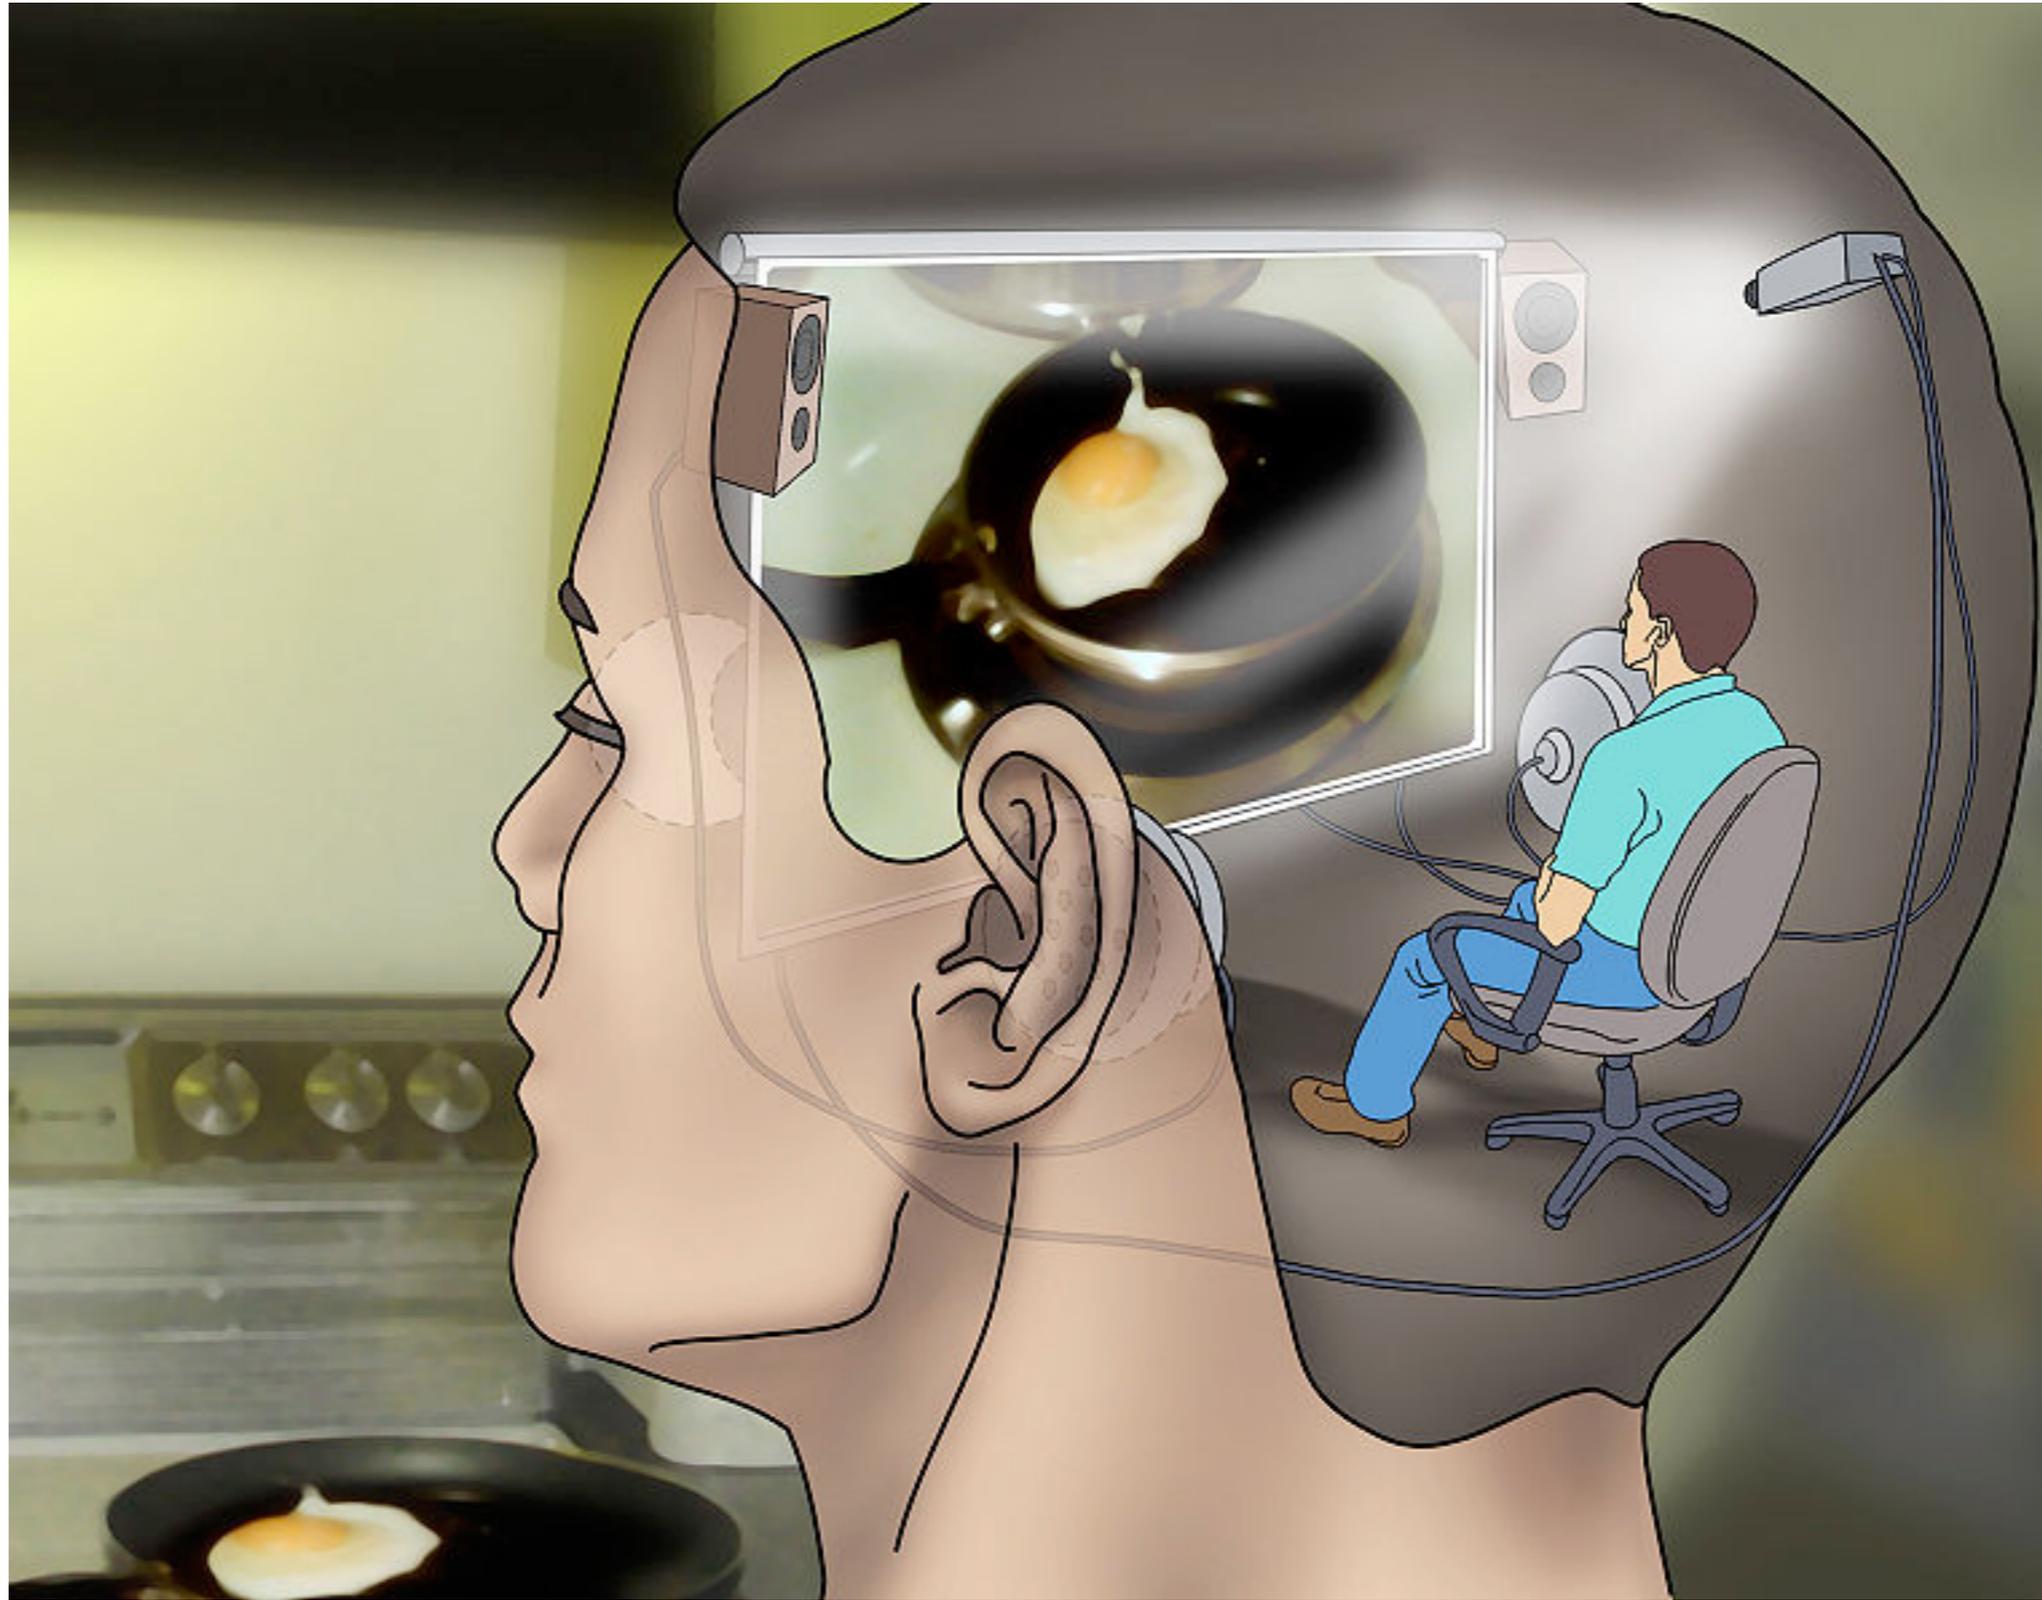


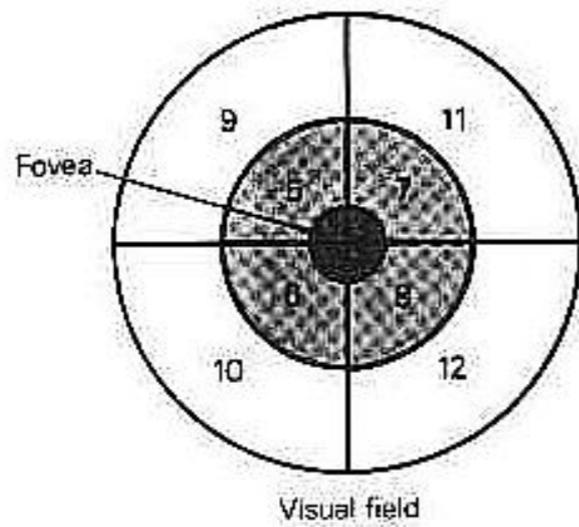


50 Years

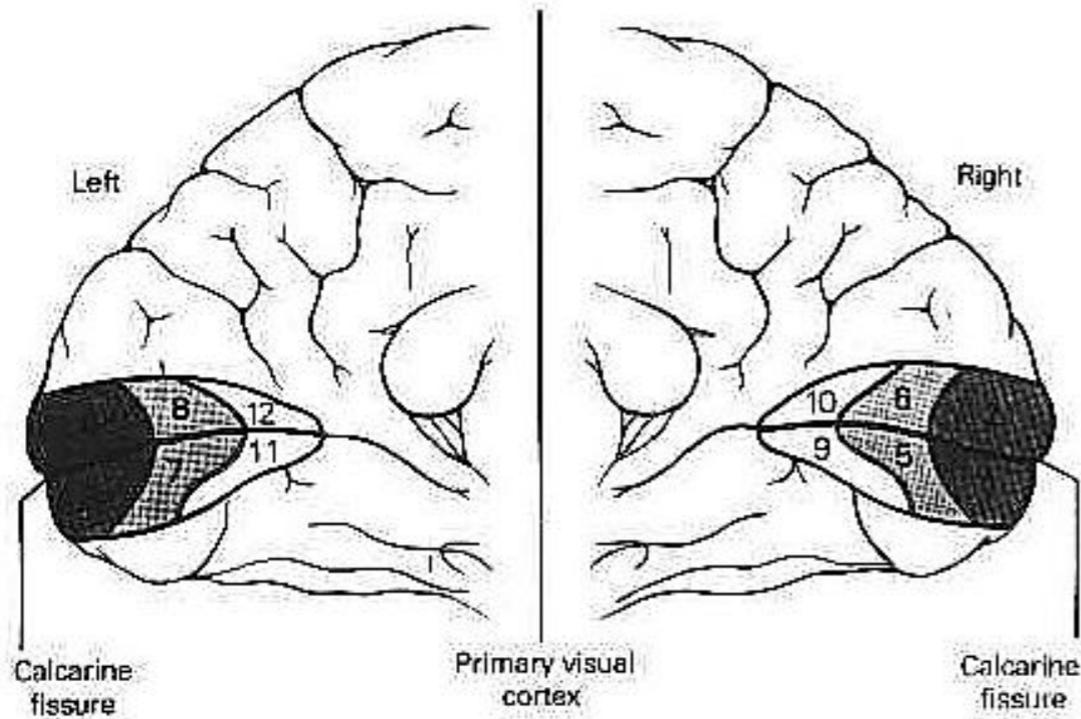


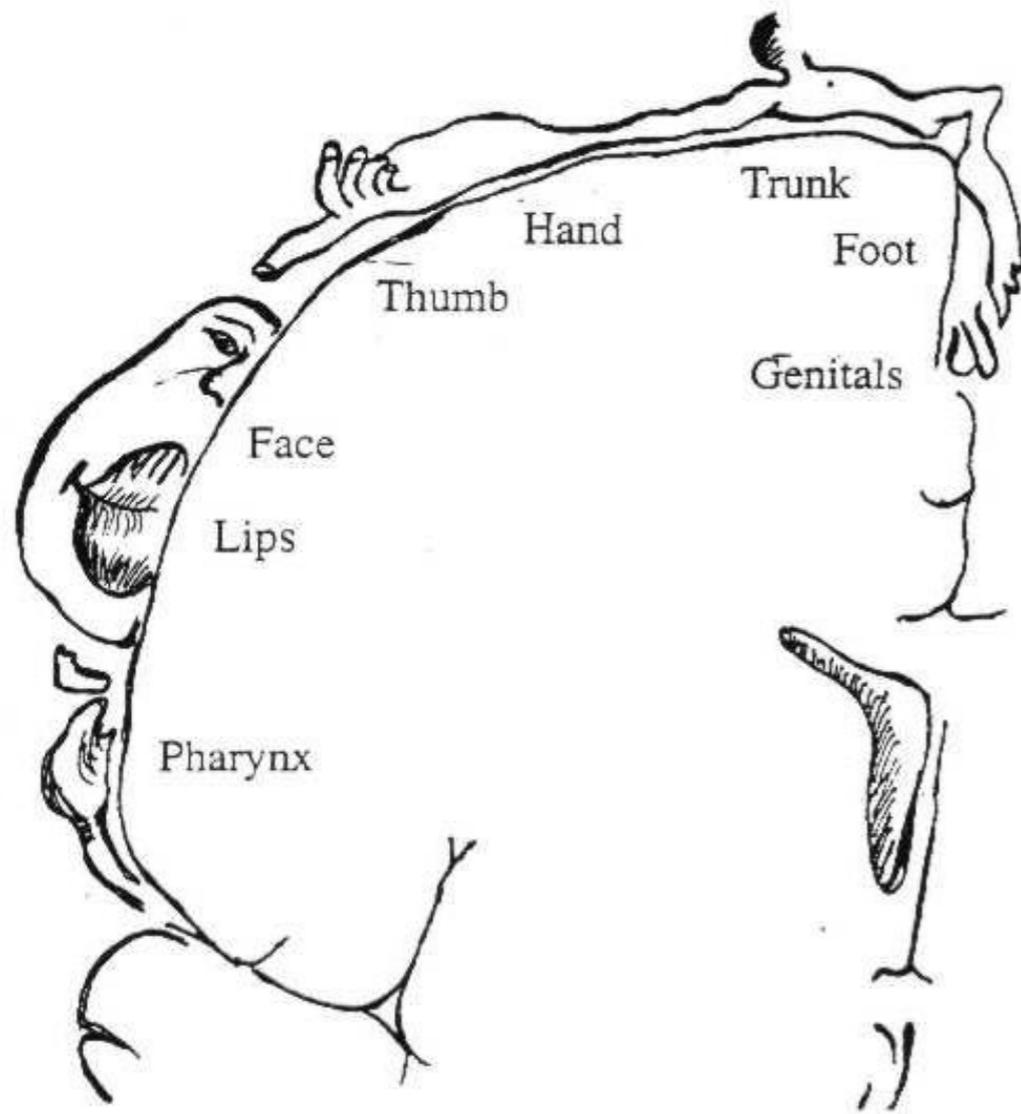
The Cartesian Theatre





Retinotopic map in monkey visual cortex





Somato-sensory map in primary sensory cortex

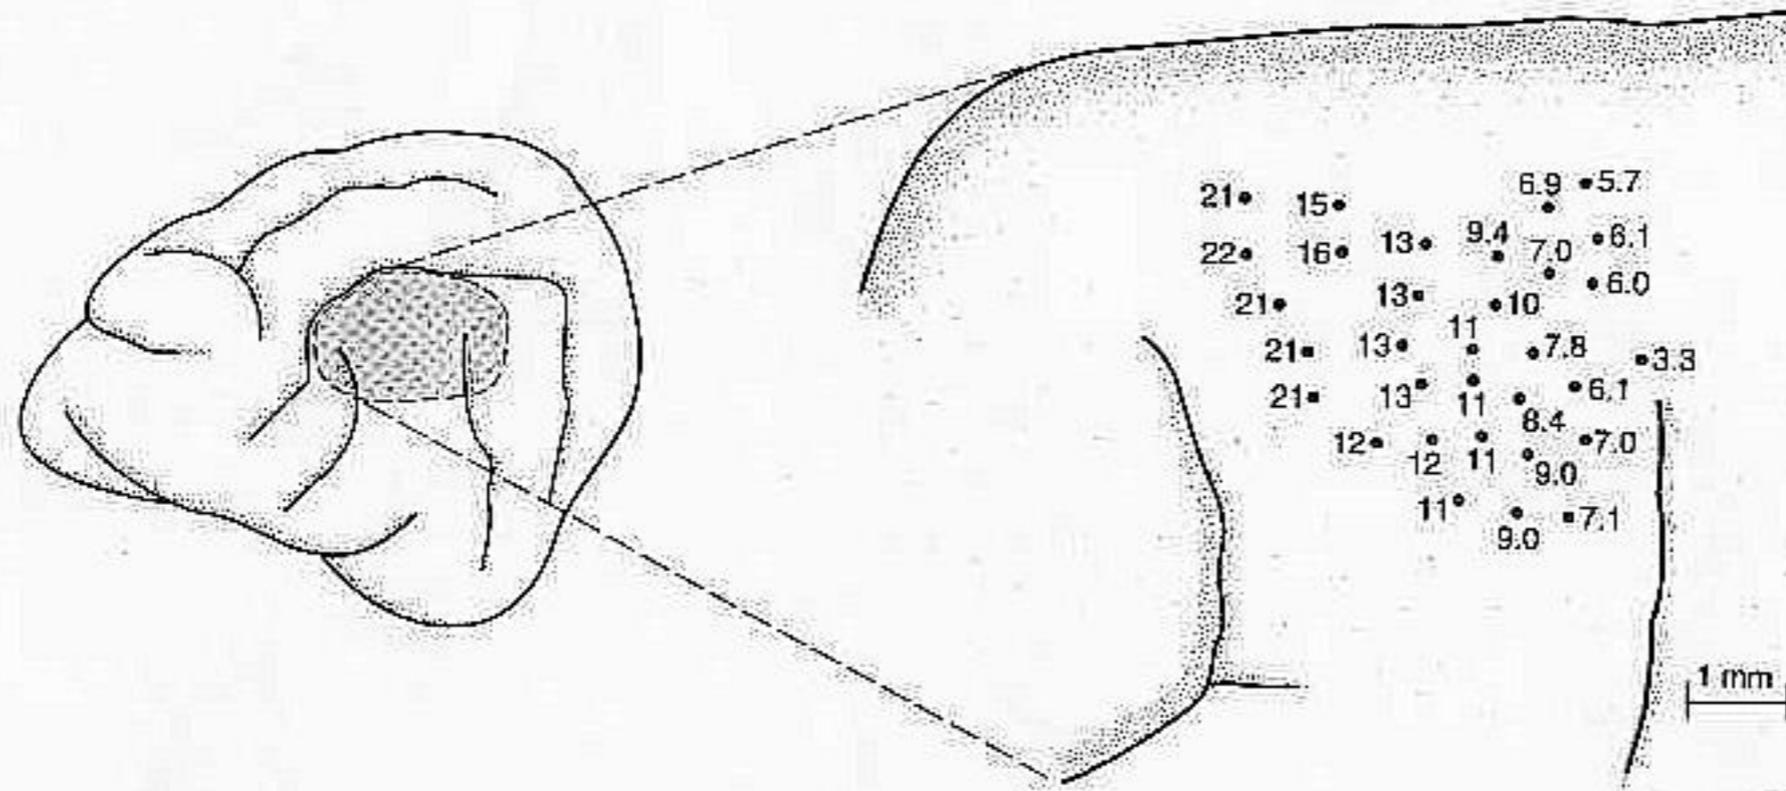


Figure 8.32

Left: The cat's brain, showing the location of the auditory receiving area (shaded). Right: A close-up of the auditory receiving area, showing the tonotopic map on the cortical surface. Each large dot represents a single neuron, and the number beside each dot represents the characteristic frequency of that neuron in thousands of Hz. (From Abeles & Goldstein, 1970, and Gulick, Gescheider, & Frisina, 1989.)

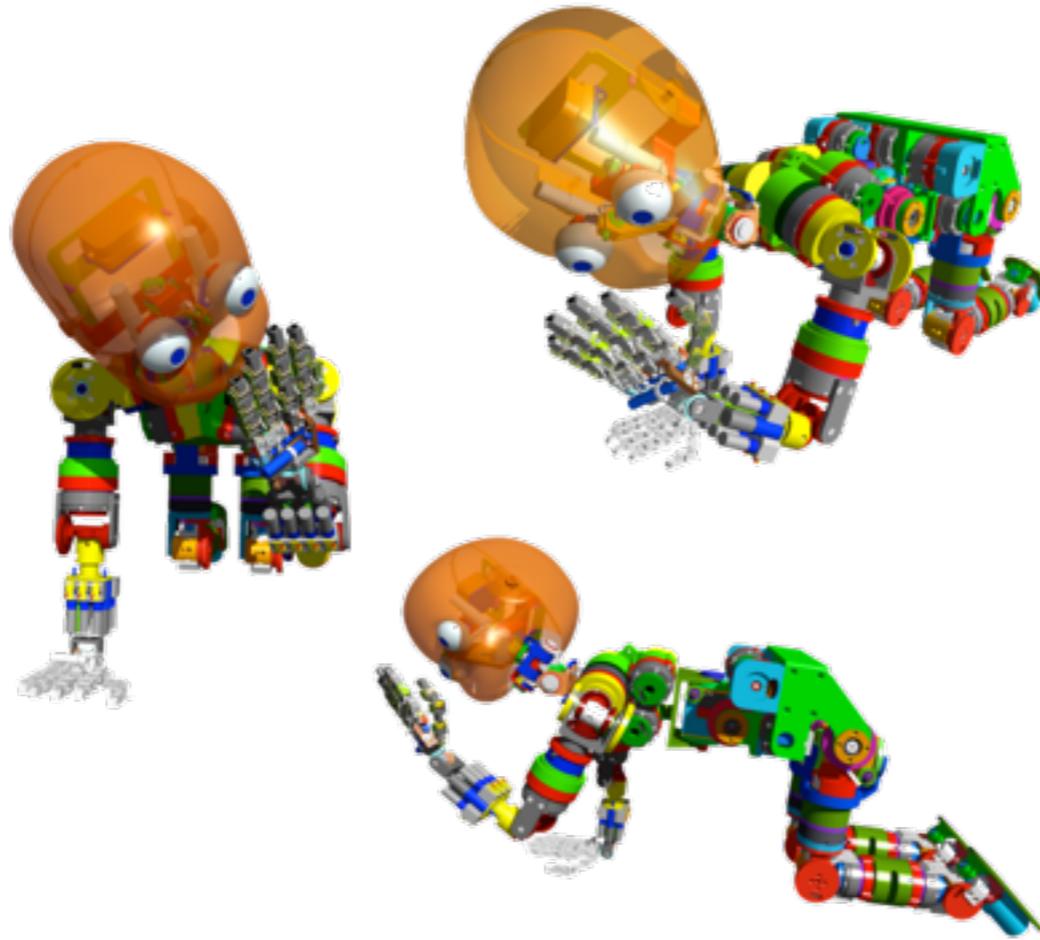
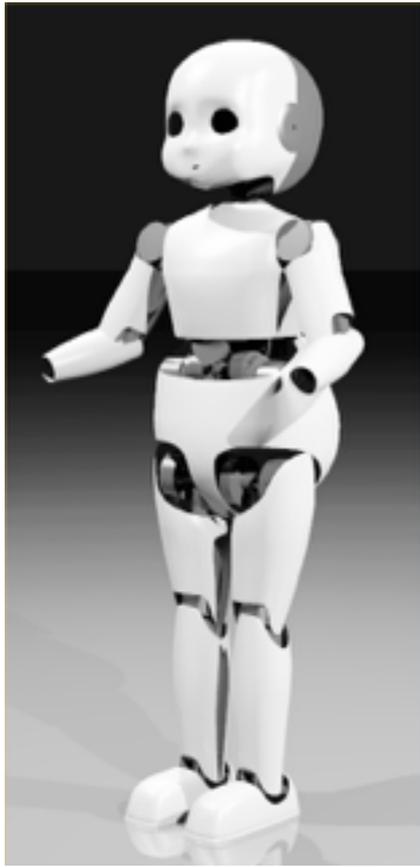
Tonotopic map in cat primary auditory cortex

No cortical maps for
.... olfaction
.... taste
.... Vestibular sense
... proprioception
... kinaesthesia
... multimodal interactions
... etc



Imprints remind us how structure in one place can be transferred





If representations work for us, they ought to be essential in artificial cognitive systems too, right?

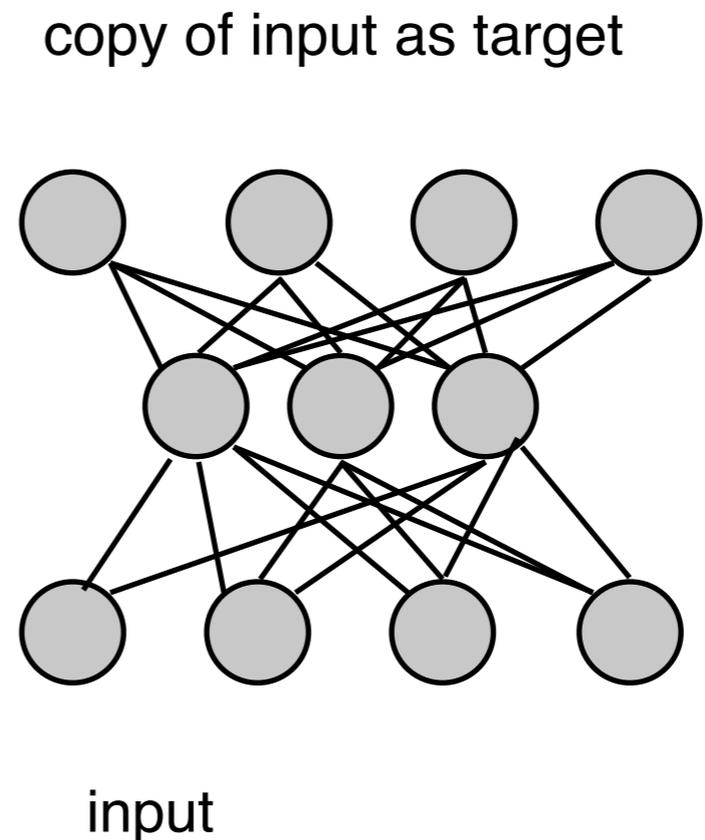
Topics

- Using autoassociators to develop principled representations
- Using Principal Component Analysis to describe high-dimensional data
- Using Self-organizing maps to reduce dimensionality

This topic is not well covered in your textbook.

Autoassociative Networks

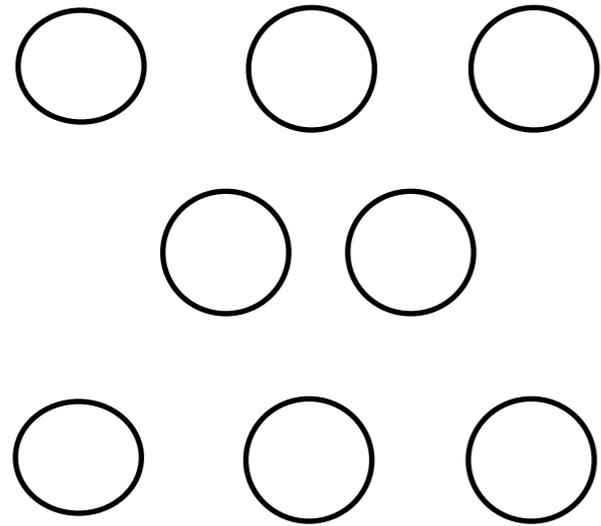
- Network trained to reproduce the input at the output layer
- Non-trivial if number of hidden units is smaller than inputs/outputs
- Forced to develop compressed representations of the patterns
- Hidden unit representations may reveal natural kinds (e.g. Vowels vs Consonants)
- Problem of explicit teacher is circumvented



Why?

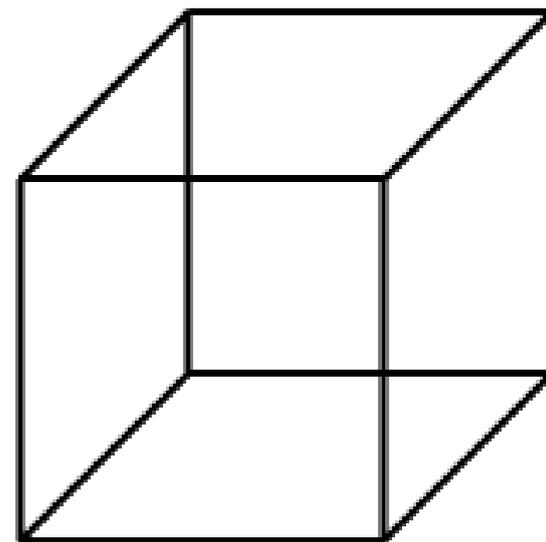
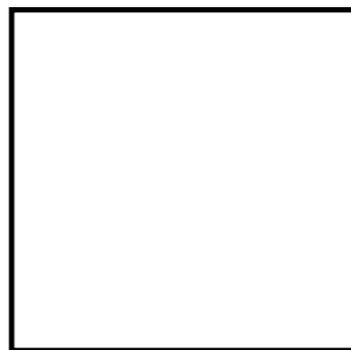
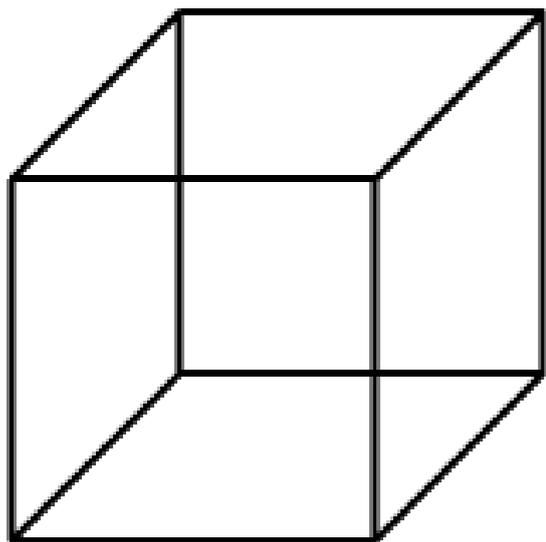
- Developing distributed representations
- Pattern recognition despite noise
- Pattern completion from partial input

Example 1:



3-2-3

**Dimension Reduction
forces similarity-based
clustering.**



Say we have a 4-2-4 network, and we train it on a set of localist patterns:

1 0 0 0

0 1 0 0

0 0 1 0

0 0 0 1

Could you design a representational scheme using only 2 bits to encode these patterns?

Now say we have a 5-2-5 network, does your solution still work?

Does any possible solution exist?

Fred's Patented Phrenological Criminal Detector™®©



Goal: train on a set of known skull images and classify them into criminal and non-criminal. Then apply the network in the field to convict criminal types before they do any harm.



From [Idebonte](#)



From [Merlene](#)



From [himbeertoni_3](#)



From [Hadjek](#)



From [Attika](#)



From [rockpaperscissors](#)



From [caito](#)



From [Koninho](#)



From [JersiBoi](#)



From [_hairplay](#)



From [XpoPen](#)



From [Hugo*](#)



From [jasonedward](#)



From [little pieces...](#)



From [Quizz](#)



From [Madman_122](#)



From [ebruli](#)



From [JourneyToNoWhere](#)



From [naunau](#)



From [Manuel Bóo](#)



From [DigitalGal](#)



From [g2bs](#)

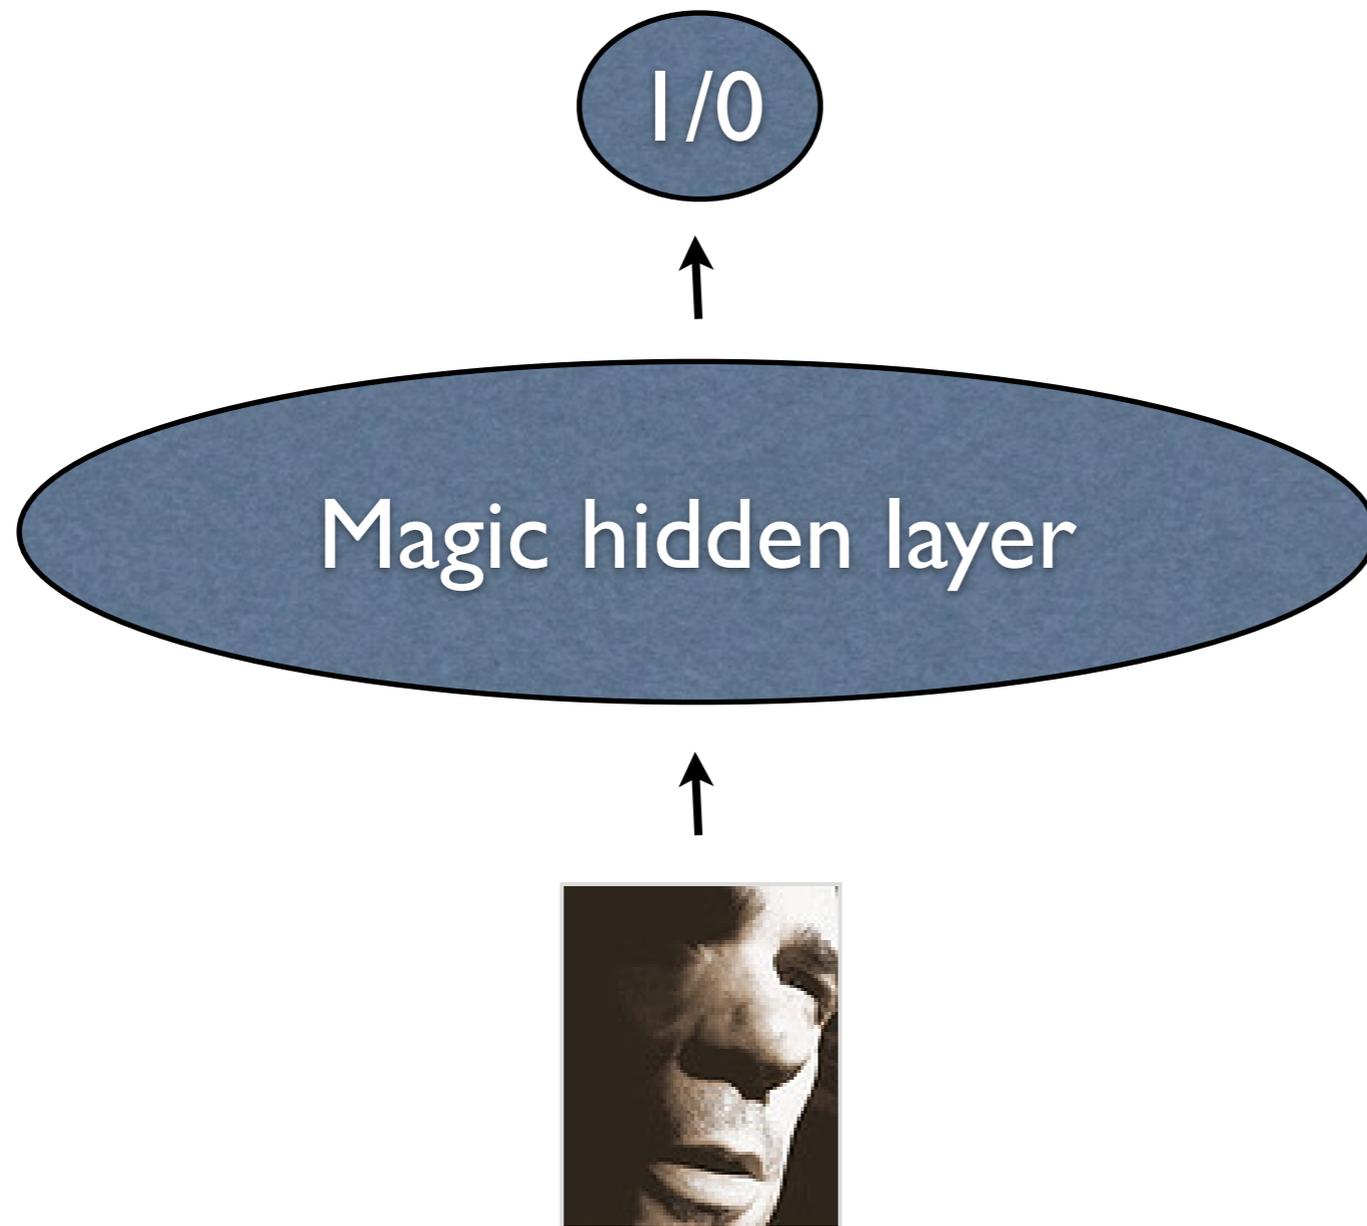


From [_Voices](#)

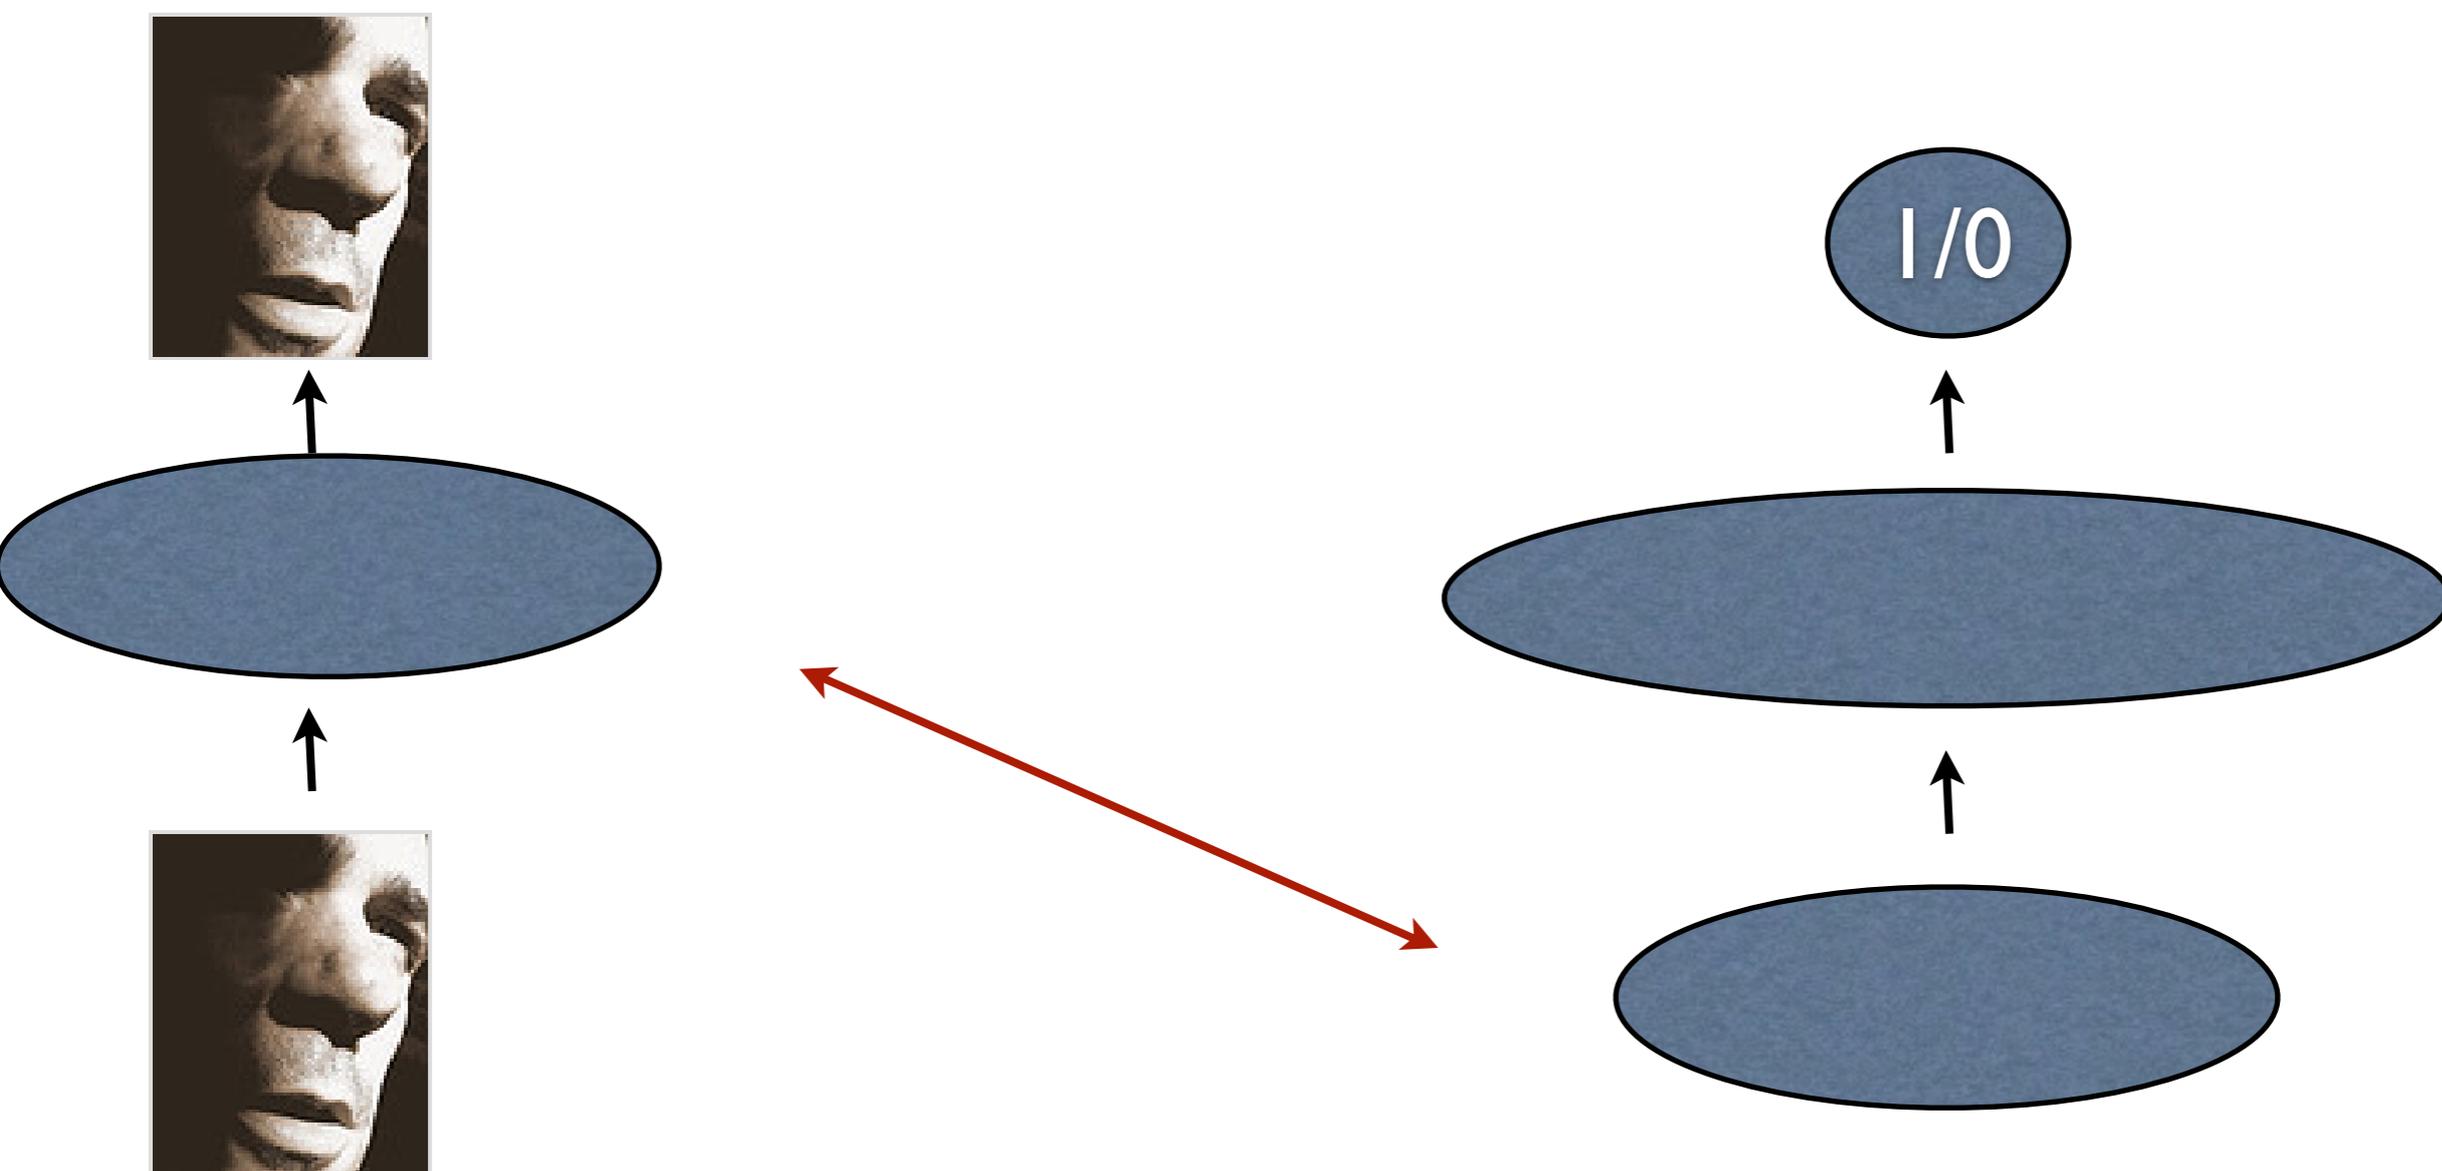
Criminals

Non-criminals

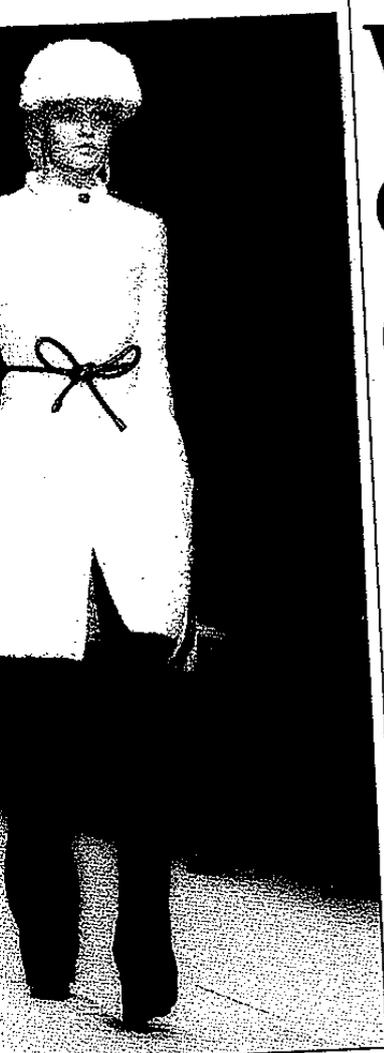
Approach 1: hope that my input representation is useful



Approach 2: Develop compressed representations which capture similarities among the faces



No problem then.



White-collar criminal has different brain, study finds

DICK AHLSTROM
in Washington

WHITE-COLLAR criminals are equipped with tremendous advantages to help them carry out their work.

Their brain structures are different and this translates into mental attributes that suit their nefarious business.

Structural brain scans showed that two areas in particular were different than in matching controls who had not become involved in credit card fraud, uttering stolen cheques and other financial misdeeds, said Prof Adrian Raine, chair of the department of criminology at the University of Pennsylvania.

He was speaking yesterday during a session entitled Nature, Nurture and Anti-Social Behaviour at the American Association for the Advancement of Science's annual meeting.

It was all part of a new science called "neurocriminology", he explained.

It involved looking for changes in normal brain structure that could have an effect on a person's behaviour.

For example, scans have shown that psychopaths tended to have an undersized amygdala, a brain structure where empathy and emotion are processed.

These individuals have a reduced ability to feel the distress experienced by their victims. It was a matter of "bad brain, bad behaviour," he suggested.

Prof Raine described the results of a preliminary study of 42 individuals - 21 white-collar criminals involved in small-time crime and 21 matched controls. He said there were no bankers involved in the study.

The group of criminals had an

altered prefrontal cortex, an area of the brain involved in planning complex cognitive behaviours, personality expression, decision-making and moderating correct social behaviour.

Another structure, the temporo-parietal junction which is the area associated with moral decision-making, was also different in the criminals.

This combination also modified the group's wider physiology, Prof Raine said.

"They have better executive function, which improves their executive skills like planning, regulation and management," he stated.

They also tend to respond more strongly to physical stimuli due to their enhanced ability to focus, and tend to be more driven, said Prof Raine.

"We are suggesting white-collar criminals have biological advantages."

Some of these are the very capabilities wanted in any business person yet "not everyone in that position becomes a white-collar criminal", he said.

There was a fine line between using these advantages for success in business or in criminality.

Neurocriminology raises intriguing ethical issues, Prof Raine suggested.

If altered brain structures can drive bad behaviour then should that be a mitigating factor when sentencing, or should a person be punished for wrongdoing regardless?

He also doubted whether we should try and find ways to reverse these brain changes given their value to politicians and business people.

"Politicians sometimes have to be ruthless, and businessmen need to take risks," Prof Raine said.

while (right), Aran knits such as

ar Kane

mid-length dresses and the general effect of light layering - tunics over skirts over narrow trousers in a melange of patterns - looked subtle and modern.

Guests arriving for the Burberry Prorsum show in Kensington Gardens created huge traffic jams; the show was streamed live at 4pm and was estimated to be seen in almost 200 countries.

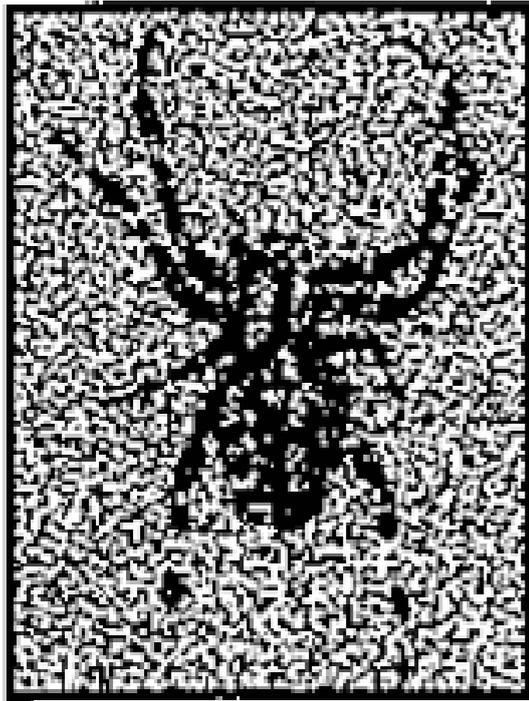
Irish viewers will have noted not only the fruity new colours for winter - tomato, pistachio and orange and the checks - but also the riffs on Aran knits like a white Aran trench and cabled insets on a white tinsel tweed coat.

Fur was used as shoulder decoration, while occasionally garish but far more desirable were the protective PVC capes paraded over their outfits by the models weathering the blizzard of snowflakes that closed the show.

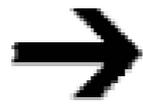
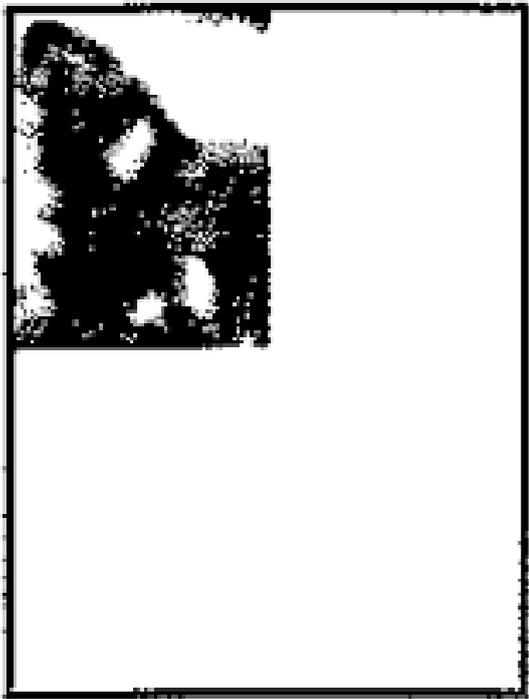
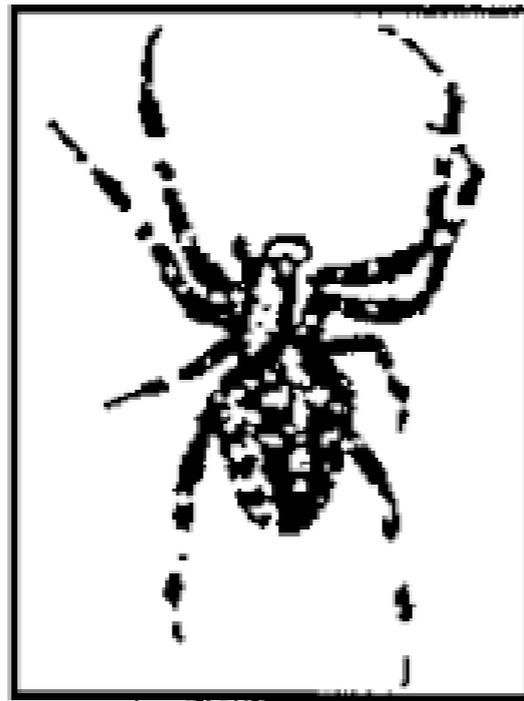
Irish Times, Feb 22, 2011

Autoassociators are often used to develop useful representations of rich data which emphasize similarity

Input



Output



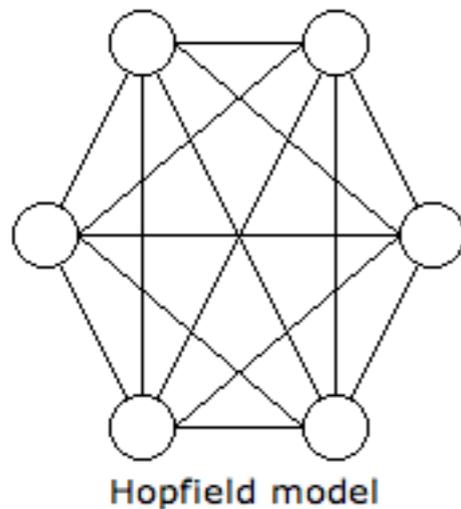
Noisy cue, partial cue
(Hertz et al, 1991)

Pattern completion

Pattern cleanup in the presence of noise

... Features of autoassociative memories

Hopfield Network:



Encodes a set of patterns, X
Weights, Y , are computed by
linear algebra

Weight matrix is square and
symmetrical

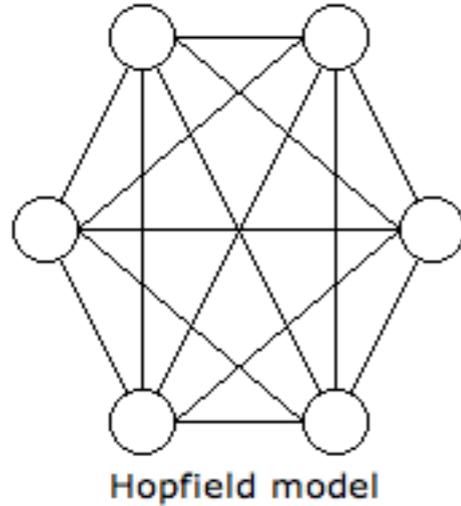
Each unit has an associated
input

Decoding:

$$X \rightarrow X' \rightarrow X'' \rightarrow X'''$$

until stabilization

Hopfield Network:



One shot learning

Correlational

Units have value -1 or 1 only

$$W_{i,j} = \frac{1}{N} \sum_m x_i^m x_j^m$$

Try out some applets....

<http://www.cbu.edu/~pong/ai/hopfield/hopfield.html>

Or Google “hopfield” and “applet”
yourself....

Content-Addressable Memory

Hopfield Network

Autoassociator

Interactive Activation Network

.....

Compare a CAM with a regular address-content link..