Left over from units..

- Hodgkin-Huxley model

$$I_{net} = g_{Na} m^3 h (V_m - E_{Na}) + g_k n^4 (V_m - E_k) + (V_m - E_l)$$

  m, h, n: voltage gating variables with their own dynamics that determine when channels open and close

- Bias weight

Left over from units..

- Hodgkin-Huxley model

  $$I_{net} = g_{Na} m^3 h (V_m - E_{Na}) + g_k n^4 (V_m - E_k) + (V_m - E_l)$$

  m, h, n: voltage gating variables with their own dynamics that determine when channels open and close

- Bias weight

  General cell excitability is plastic, regardless of particular synaptic inputs (weights). See review by Mozzachido & Byrne, 2010 *TINS* on non-synaptic plasticity.

Vol 466 | 1 July 2010 | doi:10.1038/nature09086 nature

LETTERS

Sensitivity to perturbations *in vivo* implies high noise and suggests rate coding in cortex

Michael London[1], Arnd Roth[1], Lisa Beeren[1], Michael Häusser[1] & Peter E. Latham[2]

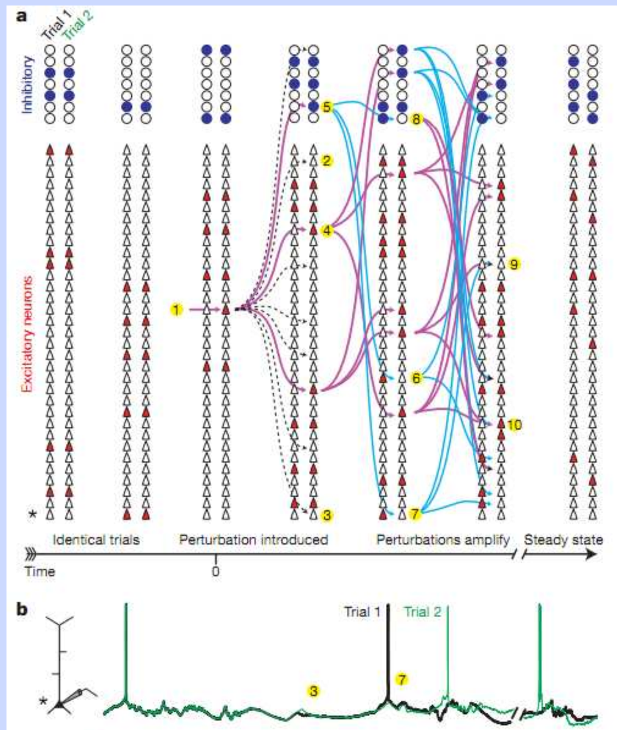

Figure 1 | The effect of an extra spike on a neuronal network. a, The

see also Shadlen & Newsome '98

Networks

Layers and layers of detectors...

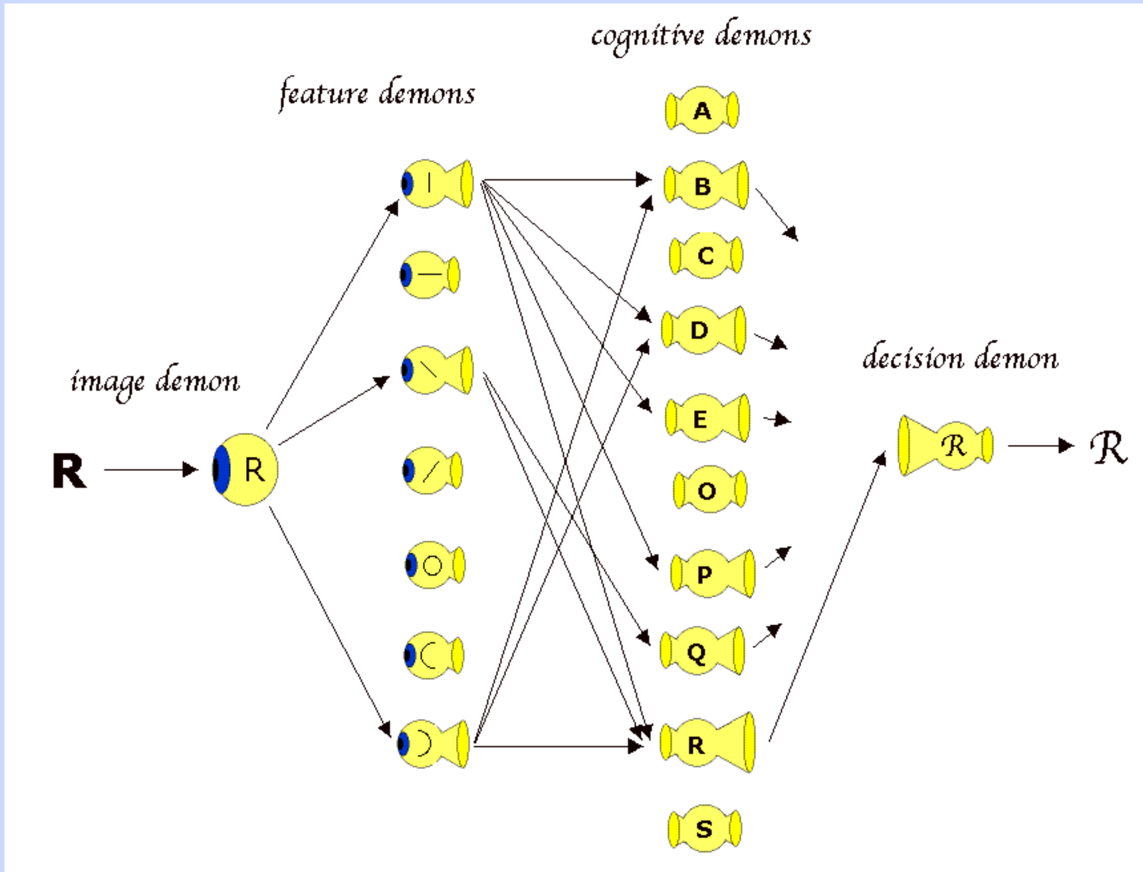

Networks

Layers and layers of detectors...

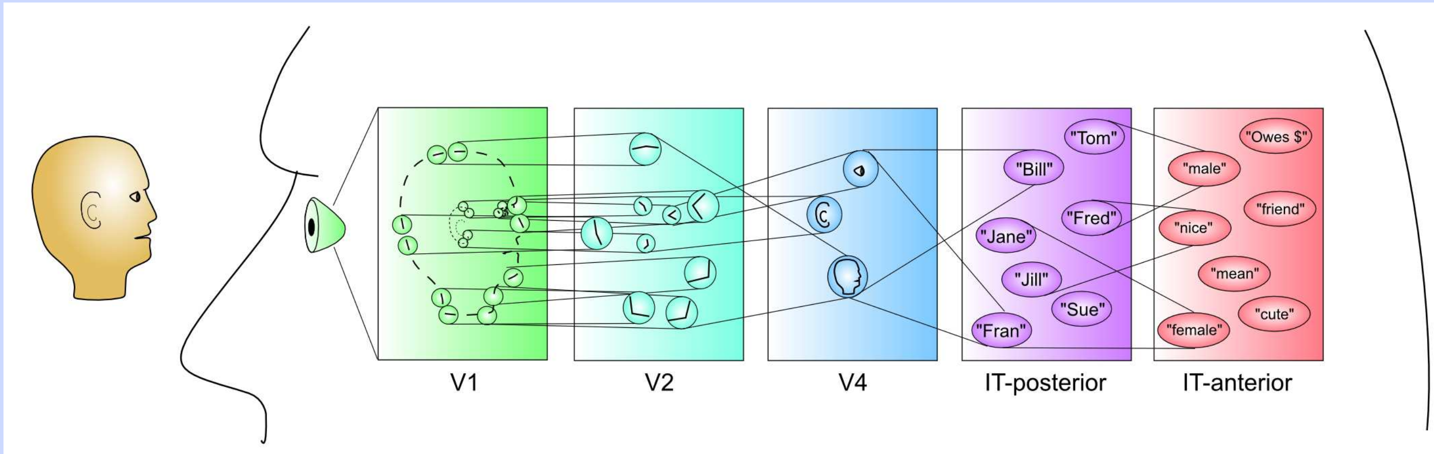

Networks

1. Biology of networks: the cortex
2. Excitation:
   - Unidirectional (transformations)
   - Bidirectional (pattern completion, amplification)
3. Inhibition: Controlling bidirectional excitation.
4. Constraint Satisfaction: Putting it all together.

Cortex: Neurons

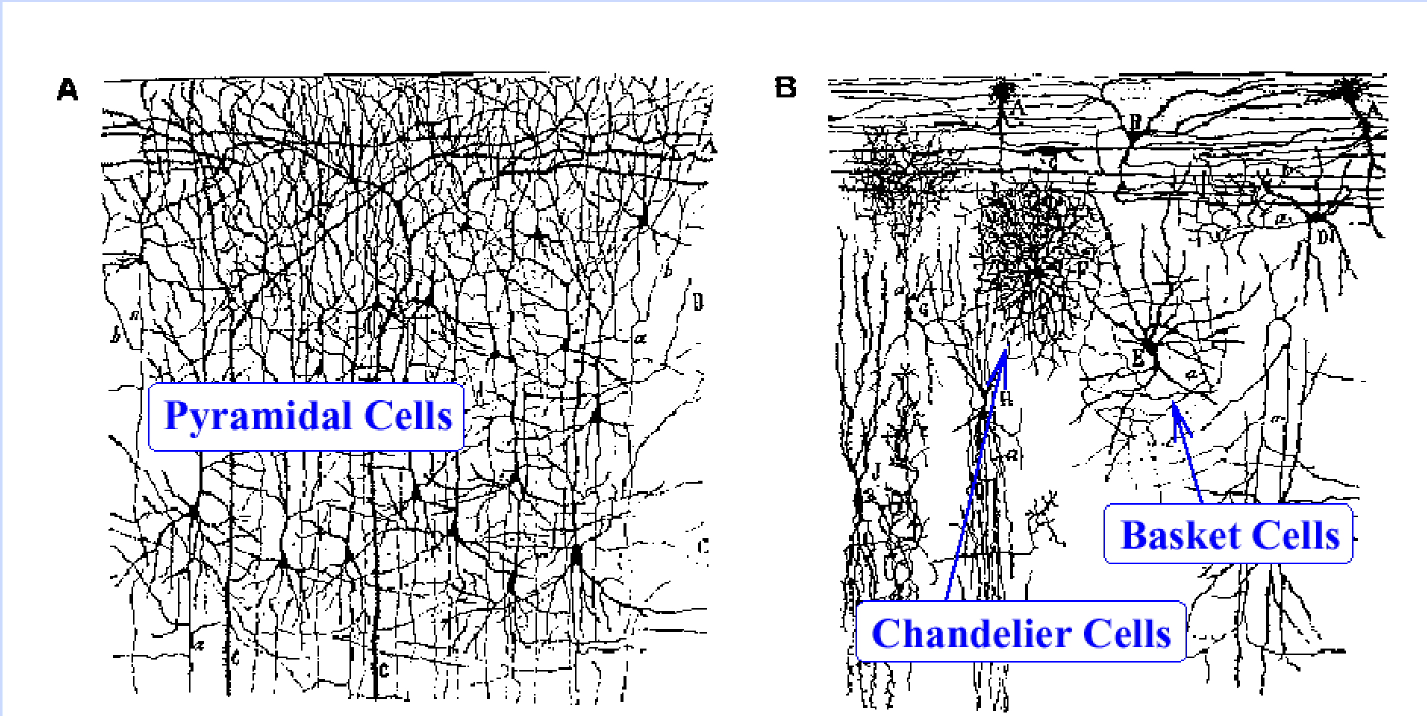

Two **separate** populations:

- Excitatory (glutamate): Pyramidal, Spiny stellate.
- Inhibitory (GABA): Chandelier, Basket.

More recent images..

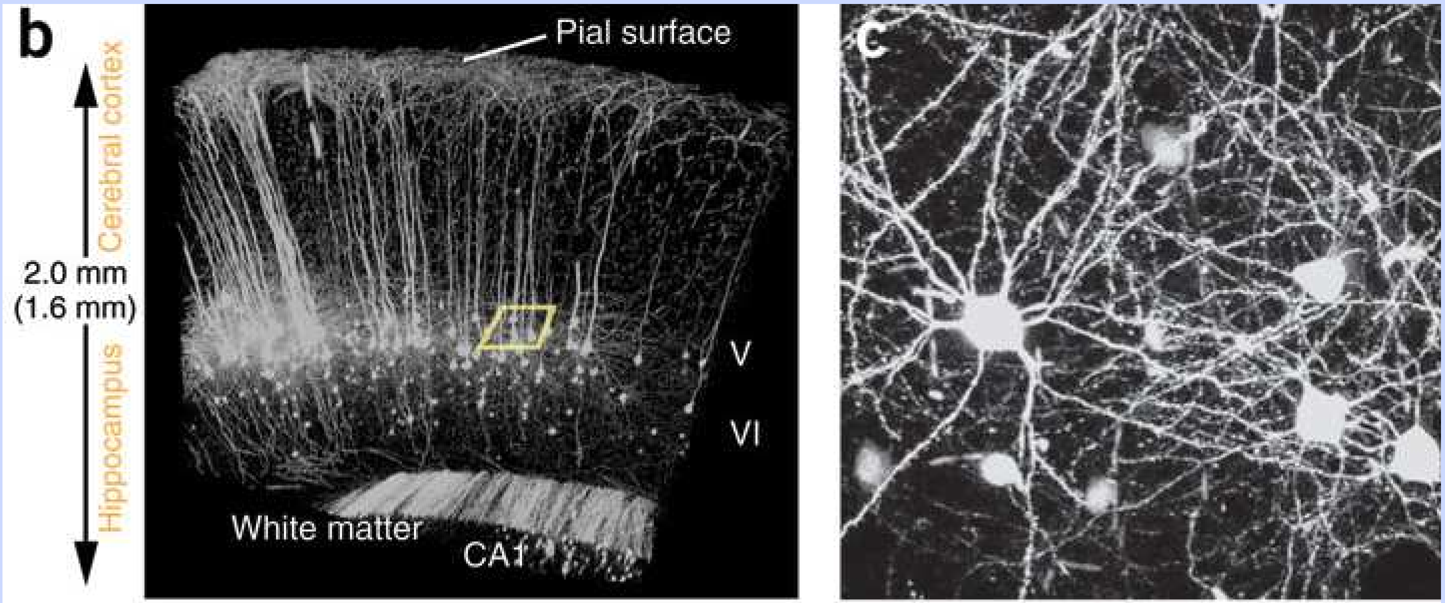

Animation: file:///Users/frankmj/teach/cogsim/ctxhippo.mpg

Excitatory vs Inhibitory Neurons

- Excitatory neurons both project locally and make long-range projections between different cortical areas
- Inhibitory neurons primarily project within small, localized regions of cortex
- Excitatory neurons carry the information flow (long range projections)
- Inhibitory neurons are responsible for (locally) regulating the activation of excitatory neurons

Laminar Structure of Cortex

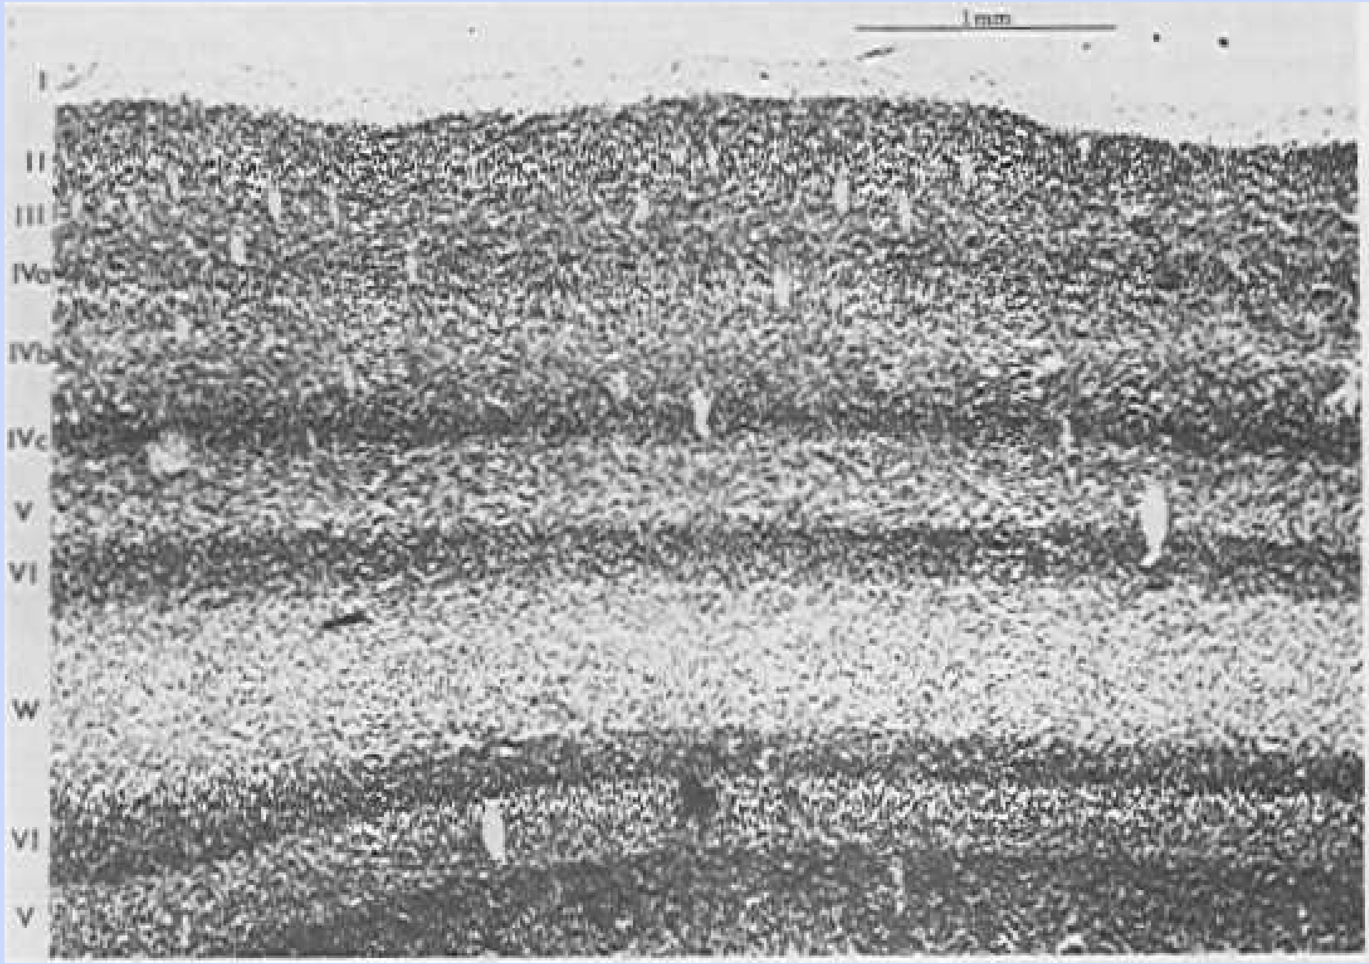

Layers

- Layer = a bunch of neurons with similar connectivity
- Localized to a particular region (physically contiguous)
- All cells within a layer receive input from approximately the same places (i.e,. from a common collection of layers)
- All cells within a layer send their outputs to approximately the same places (i.e., to a common collection of layers)

Laminar Structure of Cortex
Three *Functional* Layers

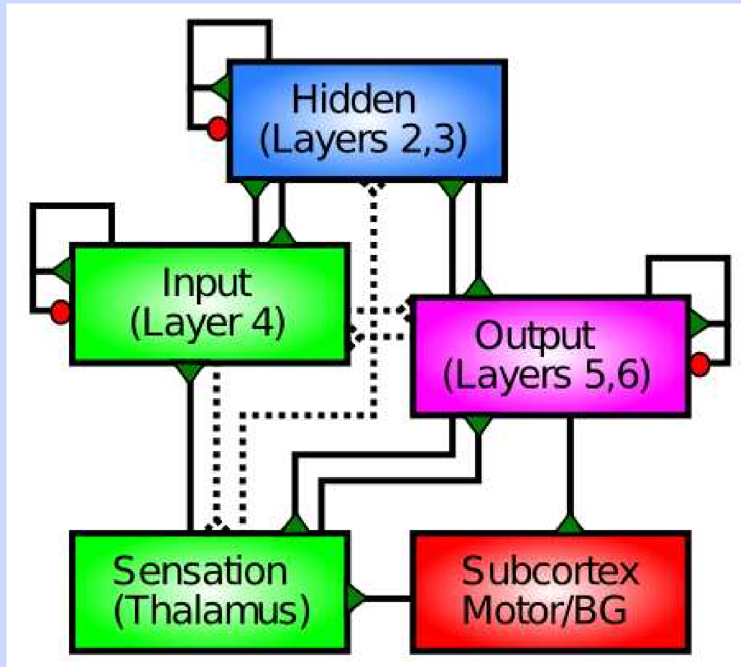

Laminar Structure of Cortex
Three *Functional* Layers

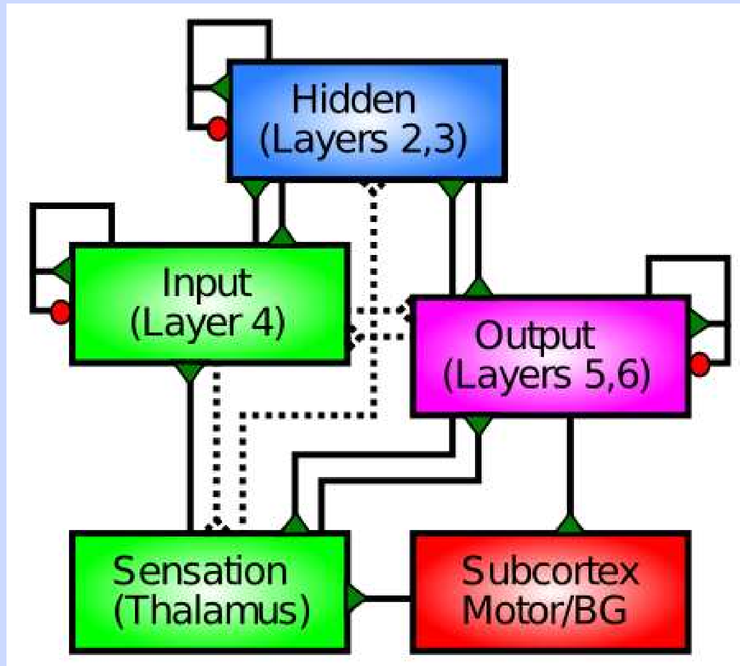

Hidden layer *transforms* input-output mappings

Laminar Structure of Cortex
Three *Functional* Layers

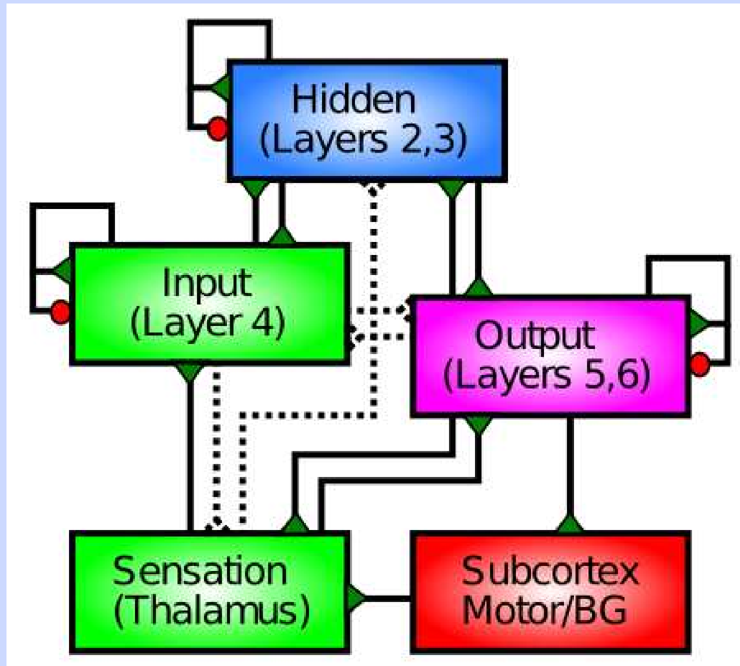

Hidden layer *transforms* input-output mappings

More hidden layers → richer transformations

Laminar Structure of Cortex
Three *Functional* Layers

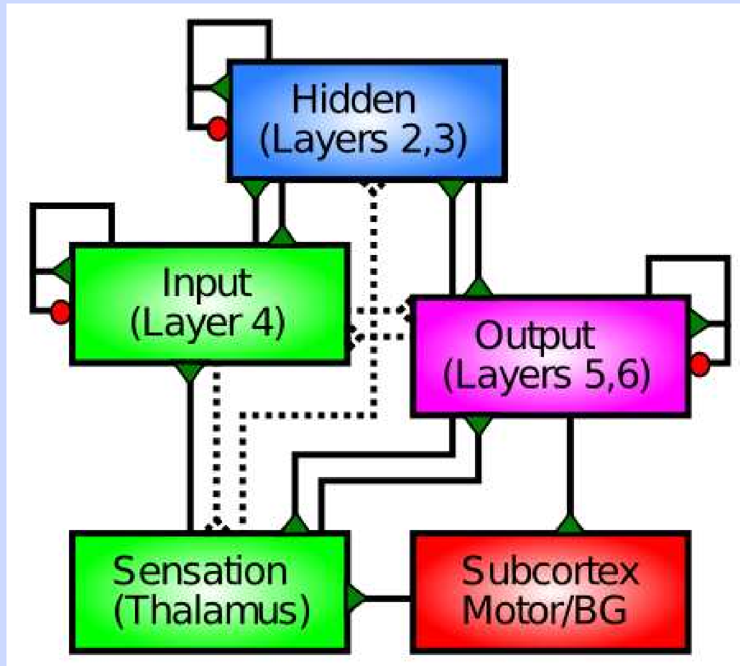

Hidden layer *transforms* input-output mappings

More hidden layers $\rightarrow$ *richer transformations*

$\rightarrow$ *less reflex-like... smarter? more "free"?*

Area Structure of Cortex

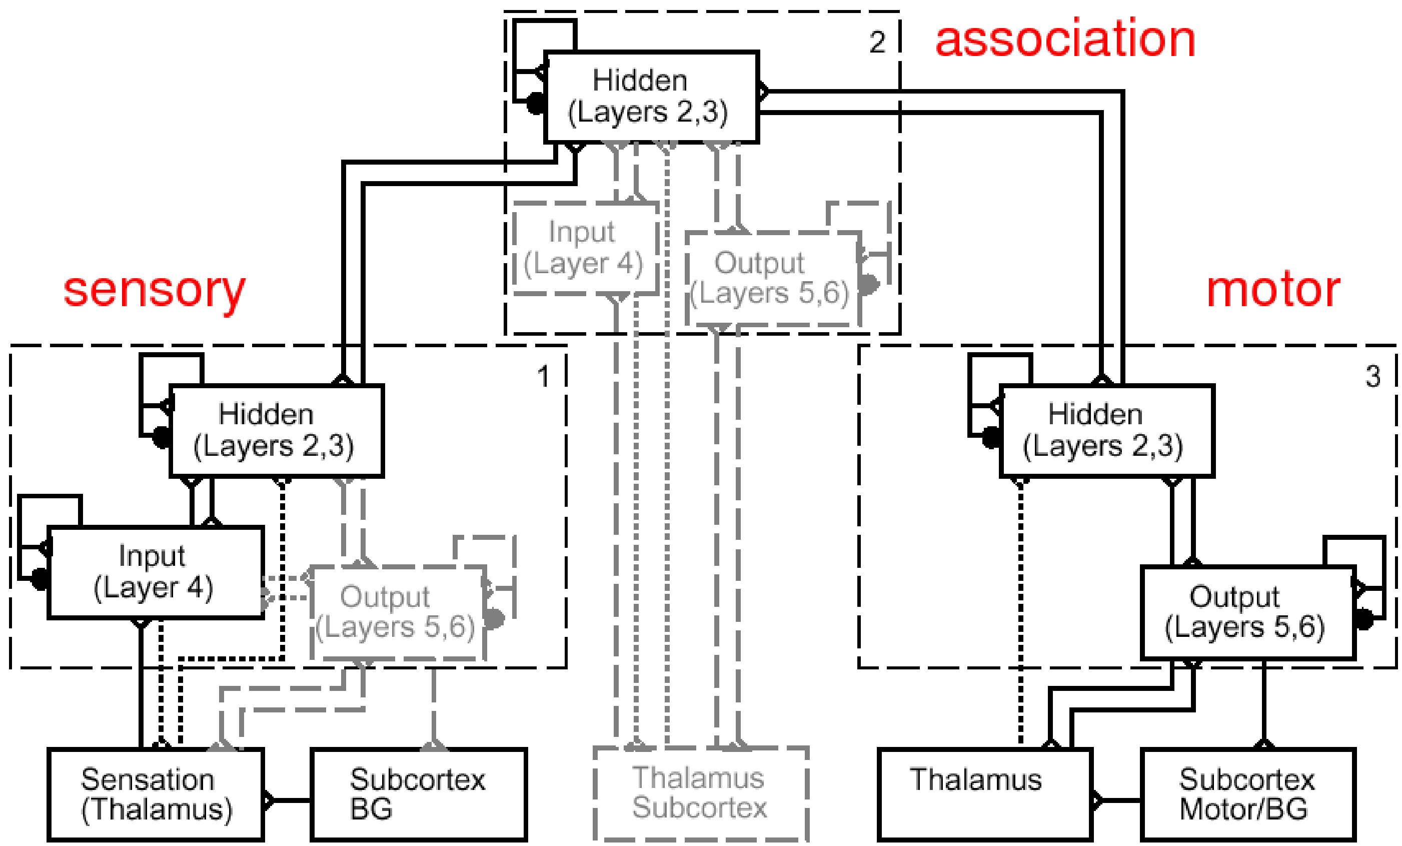

Area Structure of Cortex

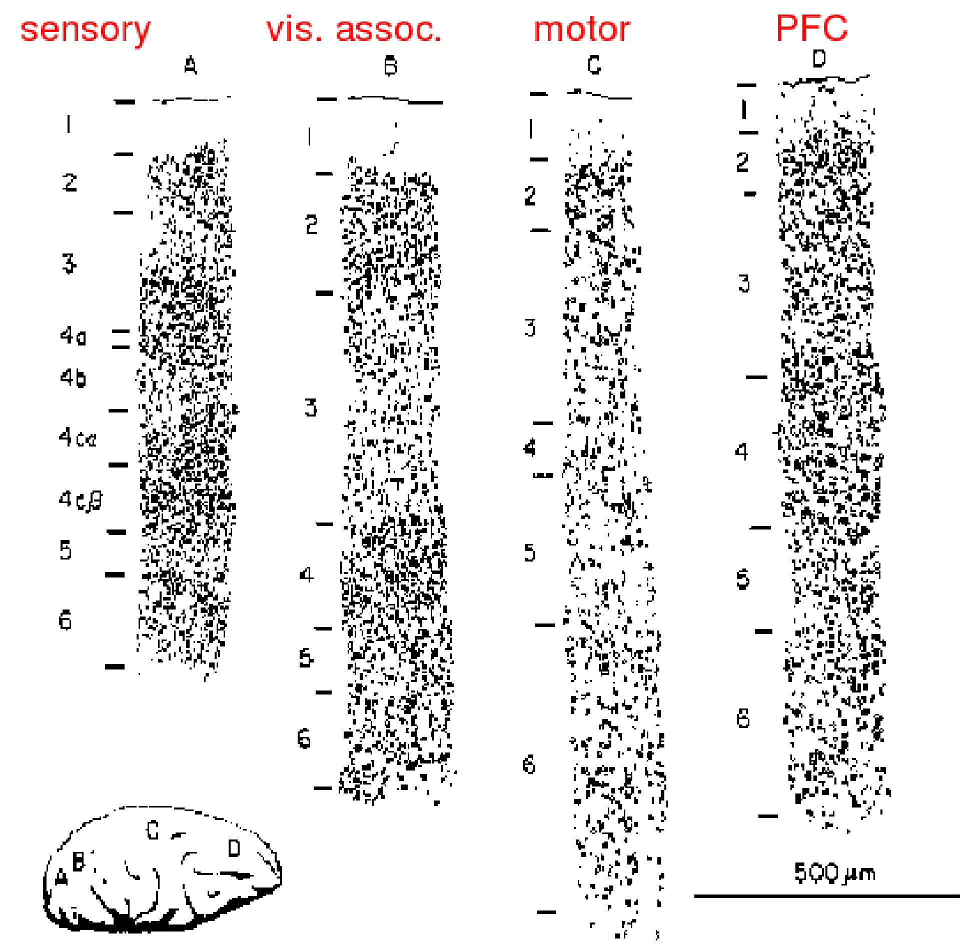

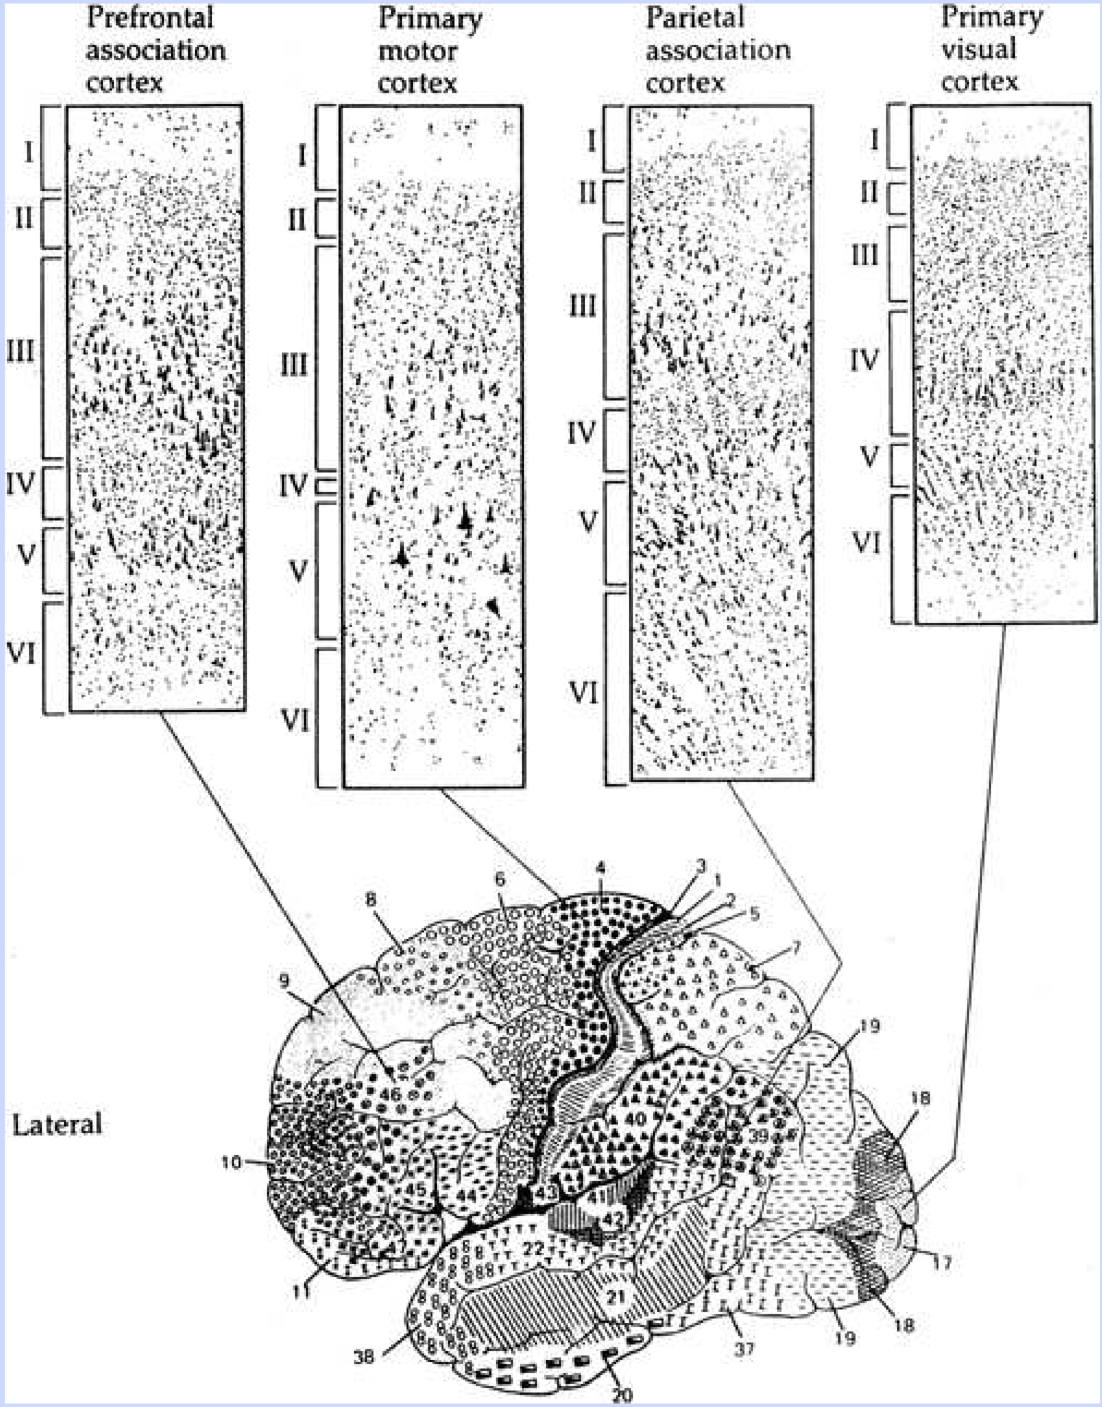

Prefrontal association cortex
Primary motor cortex
Parietal association cortex
Primary visual cortex
I
II
III
IV
V
VI
Lateral
1
2
3
4
5
6
7
8
9
10
11
17
18
19
20
21
22
37
38
39
40
41
42
43
44
45
46

Networks

1. Biology: Cortical layers and neurons
2. Excitation:
   - Unidirectional (transformations)
   - Bidirectional (pattern completion, amplification)
3. Inhibition: Controlling bidirectional excitation.
4. Constraint Satisfaction: Putting it all together.

Excitation (Unidirectional): Transformations

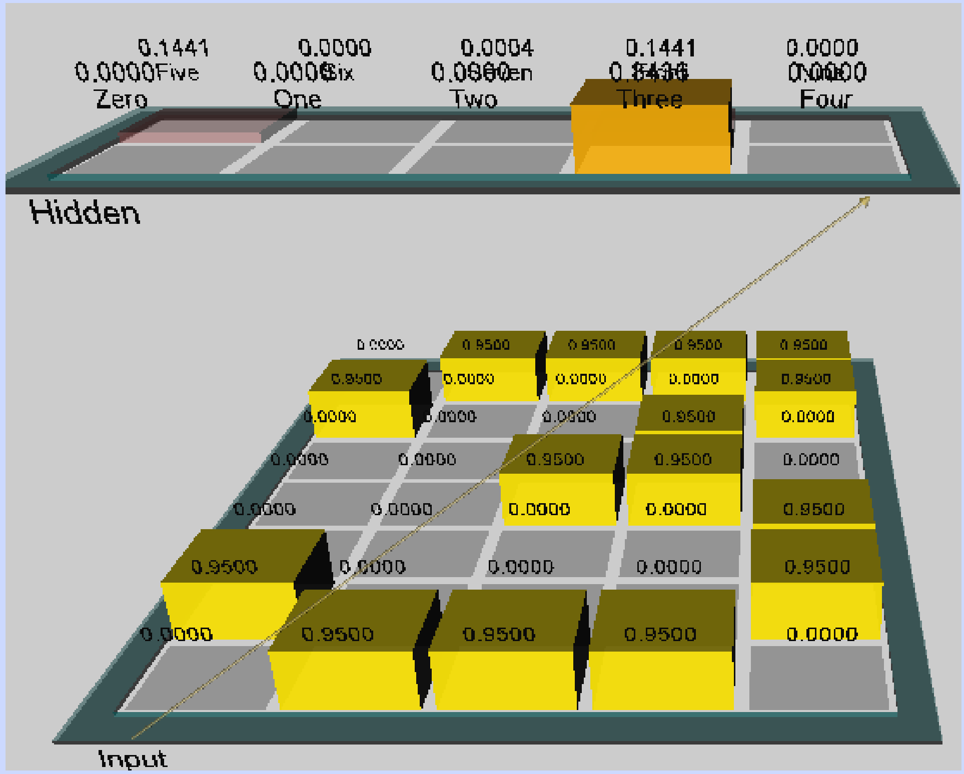

Excitation (Unidirectional): Transformations

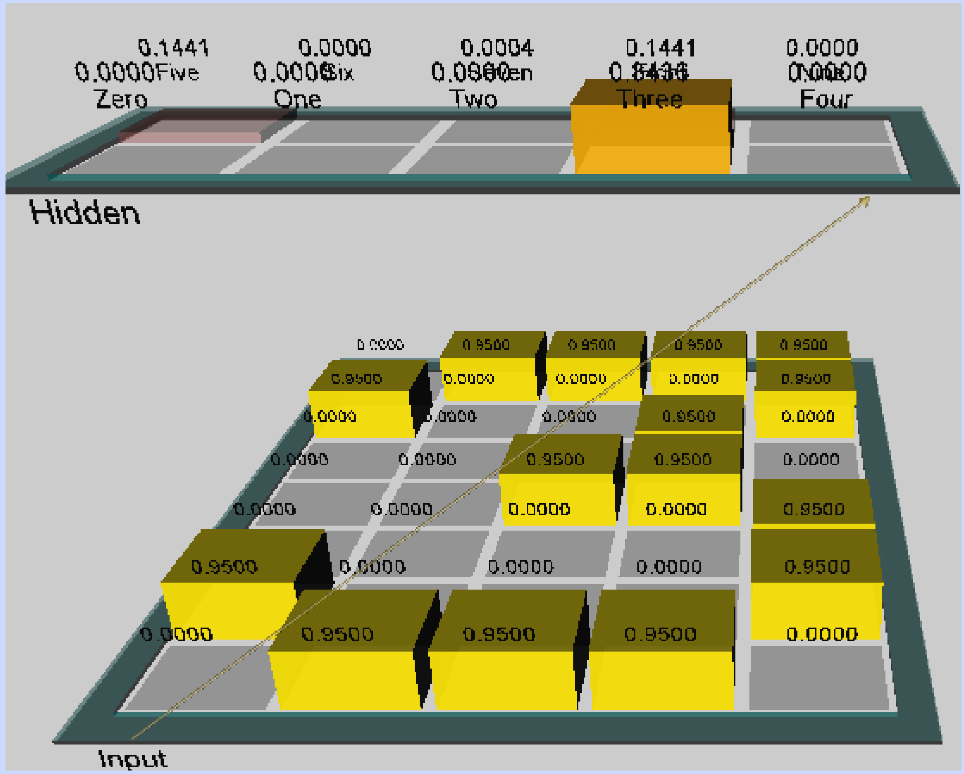

- Detectors work in parallel to *transform* input activity pattern to hidden activity pattern.

Excitation (Unidirectional): Transformations

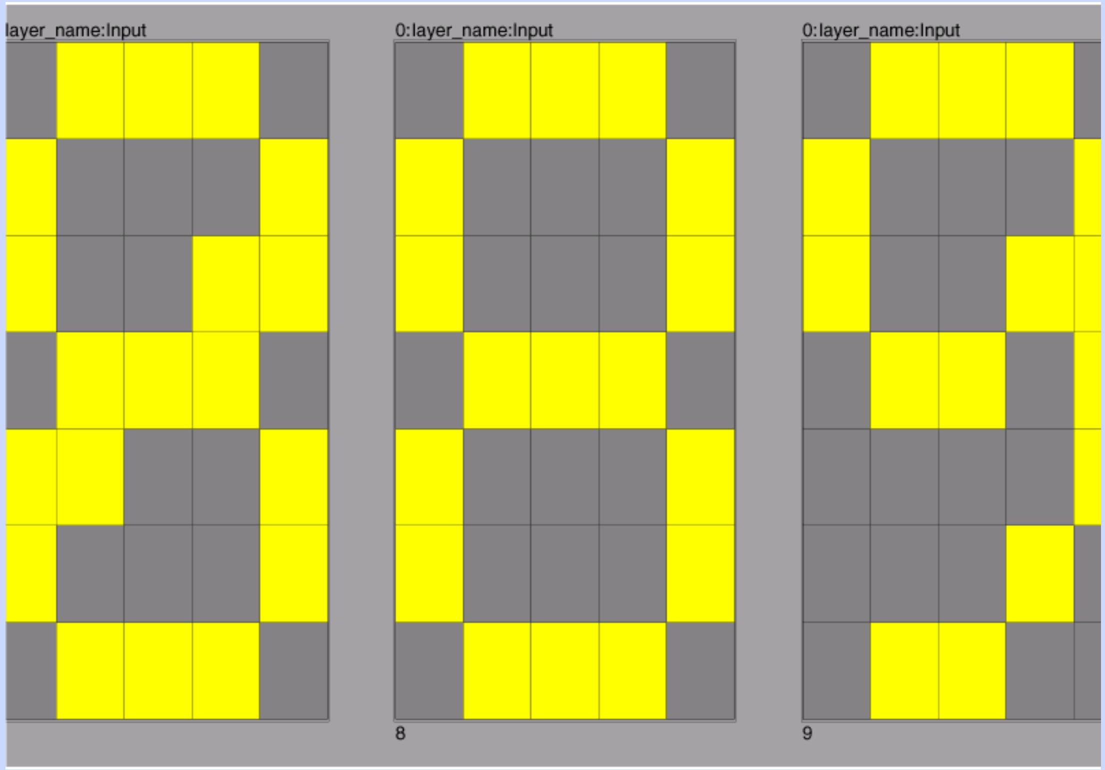

- Emphasizes some distinctions, collapses across others.

Excitation (Unidirectional): Transformations

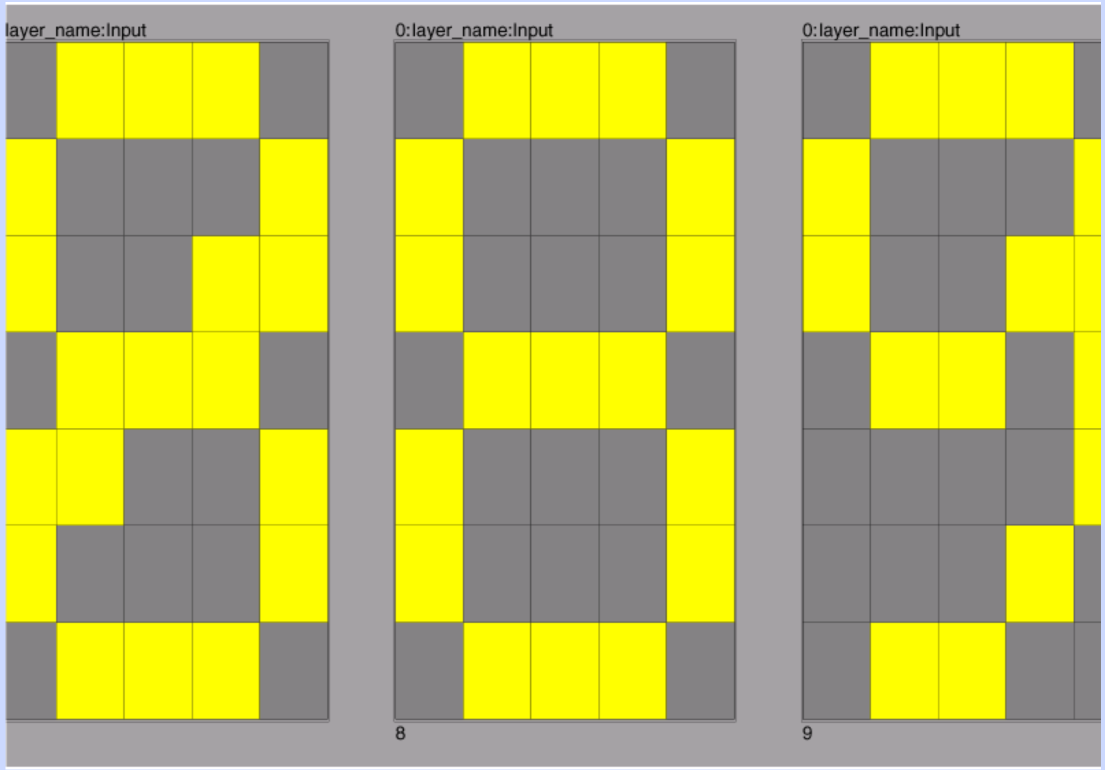

- Emphasizes some distinctions, collapses across others.
- Function of what detectors detect (and what they ignore).

Emphasizing/Collapsing Distinctions

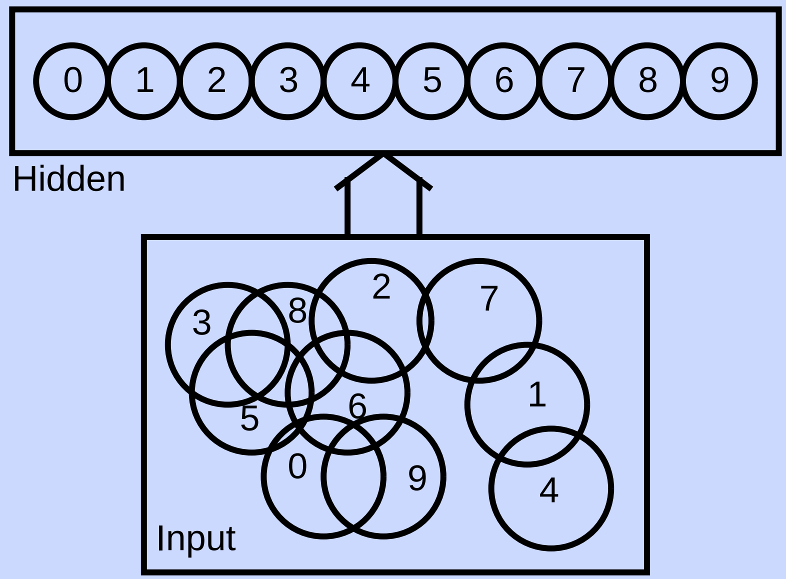

Emphasizing/Collapsing Distinctions

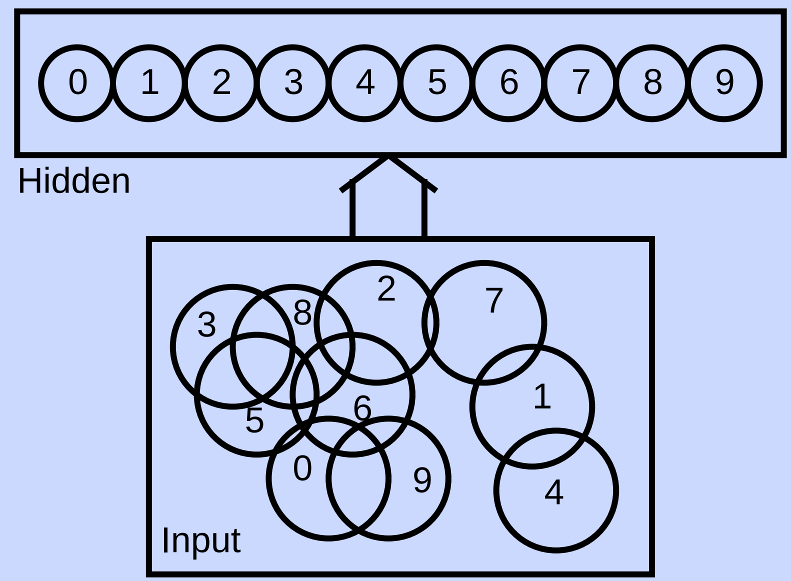

Other (more interesting) examples?...

[transform.proj]

digit detectors:

- tested with noisy digits
- tested with letters

Cluster Plots

- Cluster plots provide a means of visualizing similarity relationships between patterns of activity in a network
- Cluster plots are constructed based on the **distances** between patterns of activity
- **Euclidean distance** = sum (across all units) of the squared difference in activation

$$d = \sqrt{\sum_i (x_i - y_i)^2} \tag{1}$$

- Example...

Making Friends With Cluster Plots

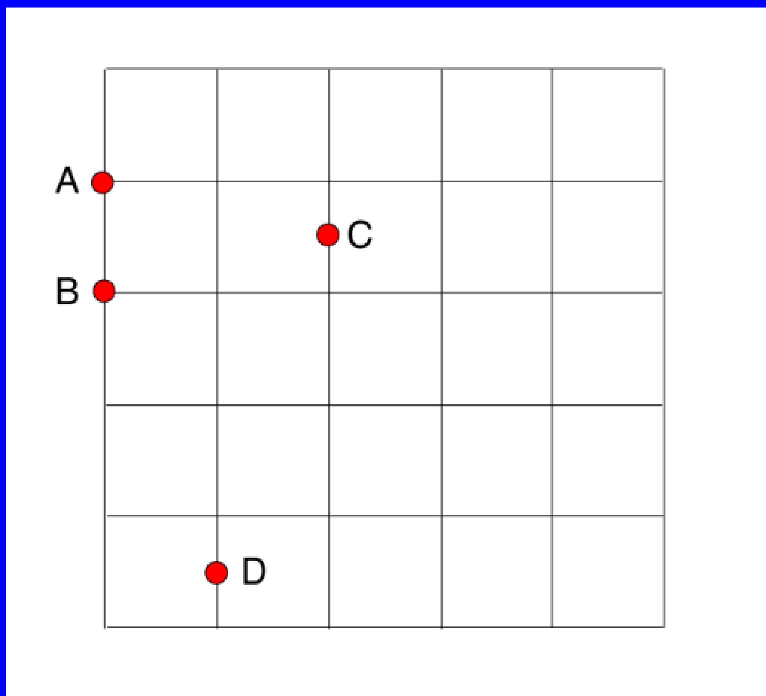

Making Friends With Cluster Plots

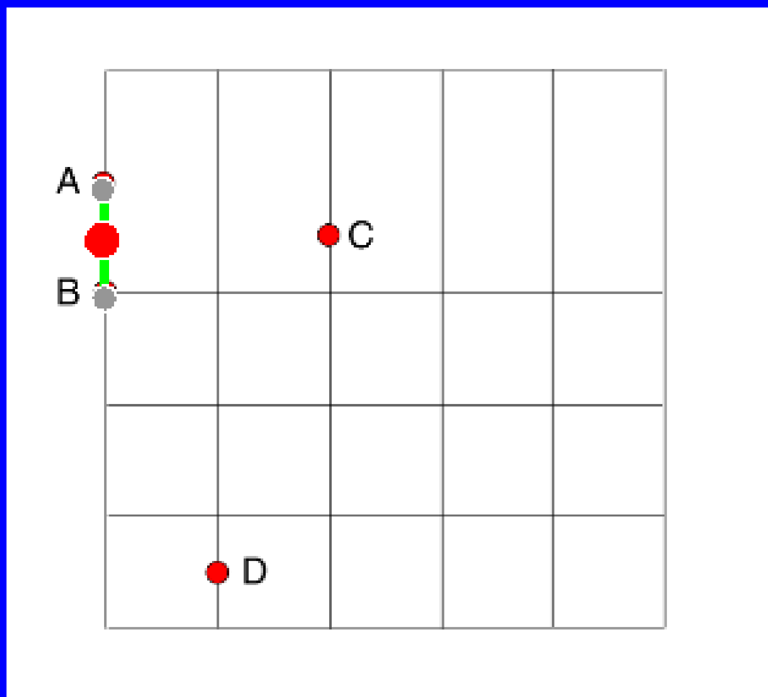

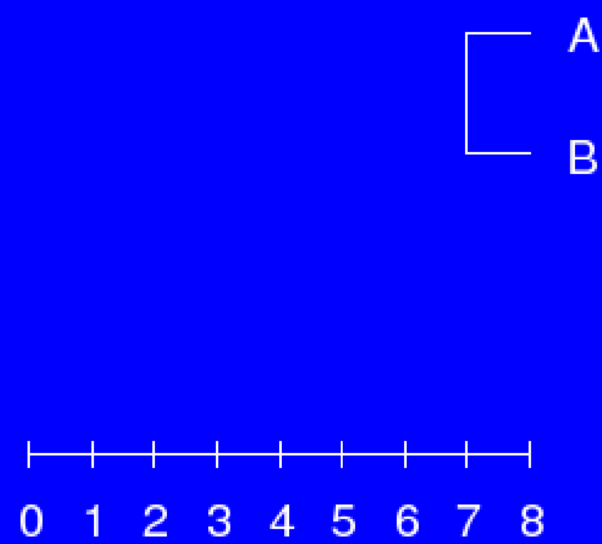

Making Friends With Cluster Plots

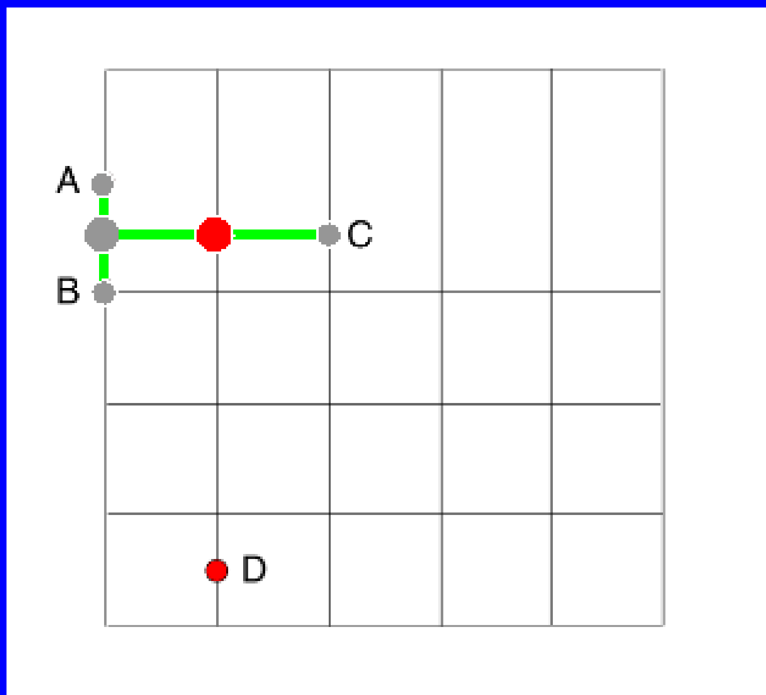

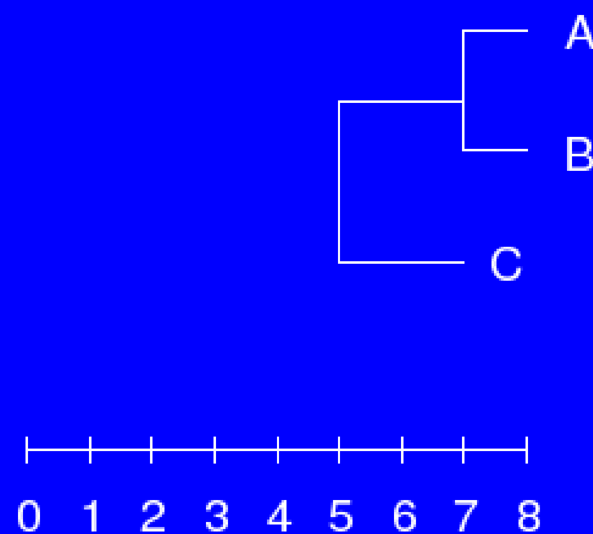

Making Friends With Cluster Plots

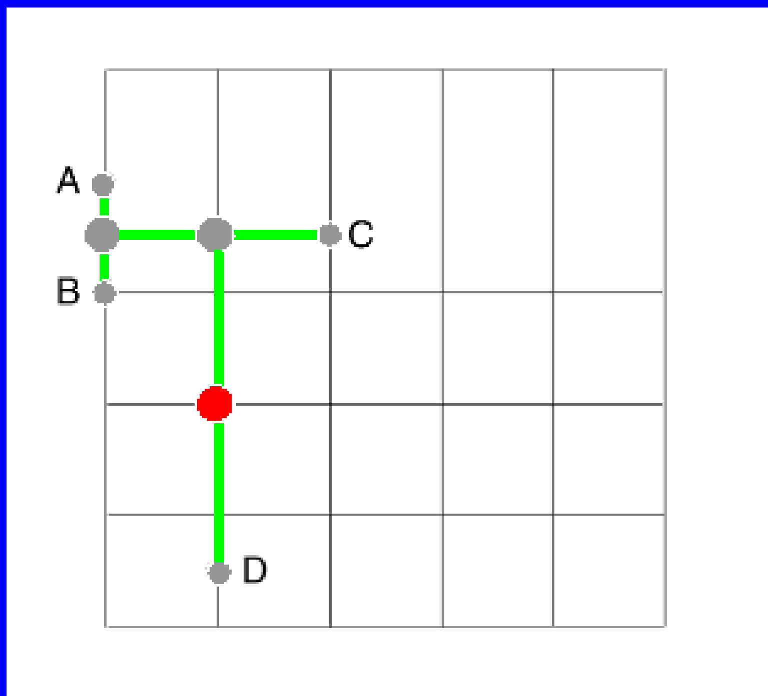

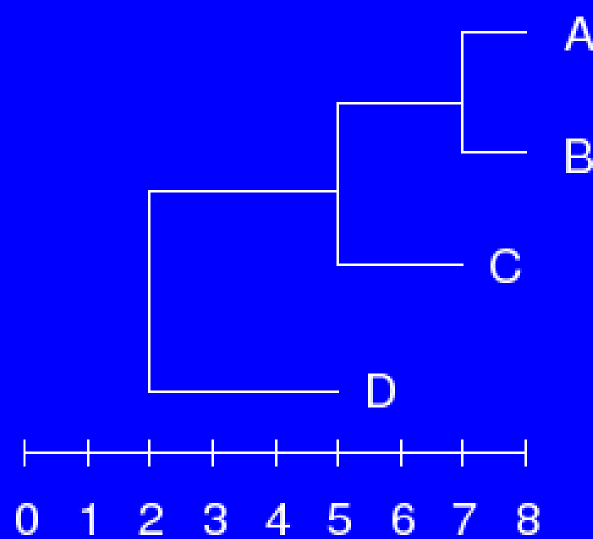

[transform.proj]

cluster plots (digits, noisy digits, hidden).

Emphasizing/Collapsing Distinctions: Categorization

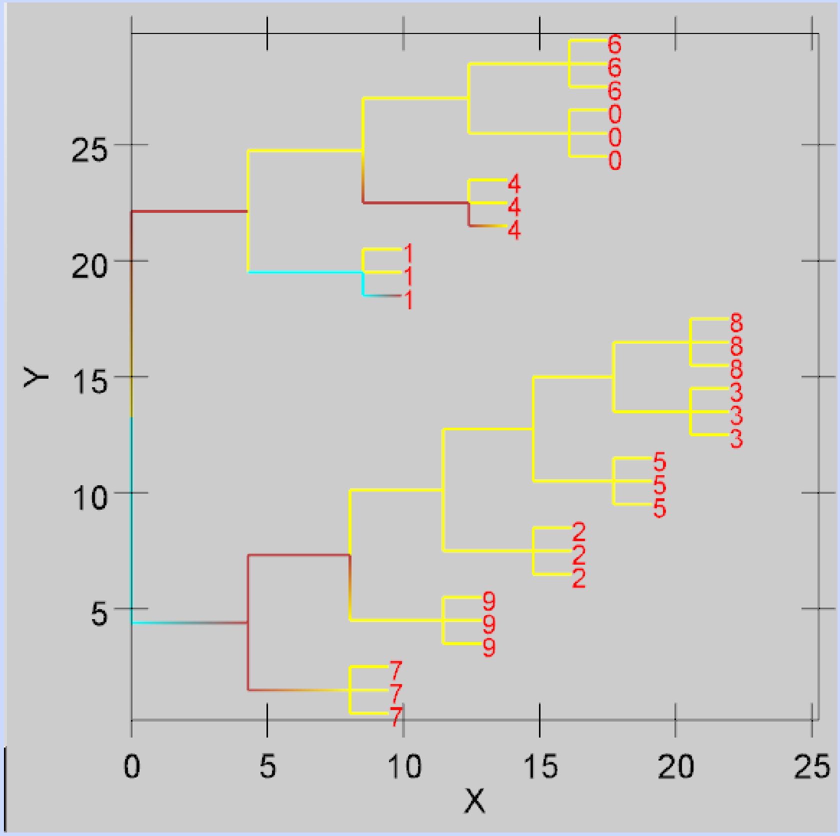

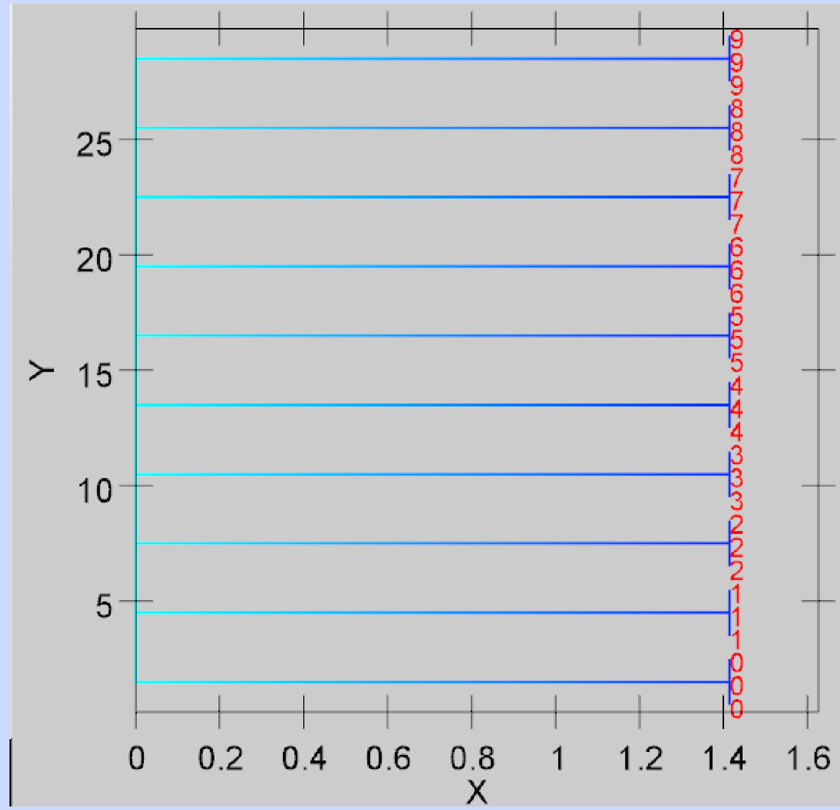

Emphasizing/Collapsing Distinctions: Categorization

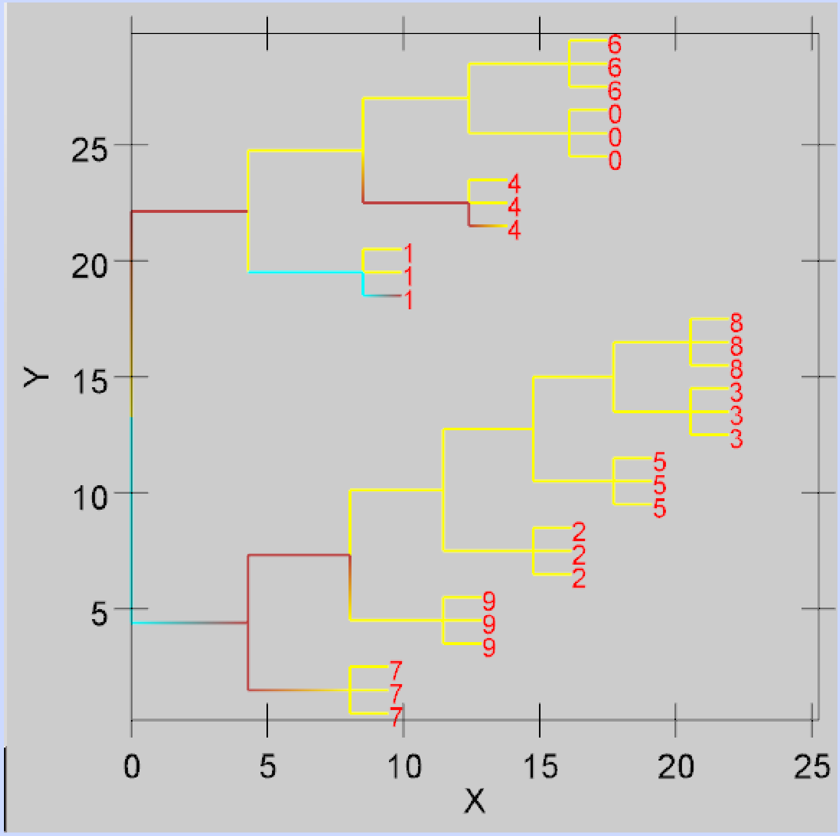

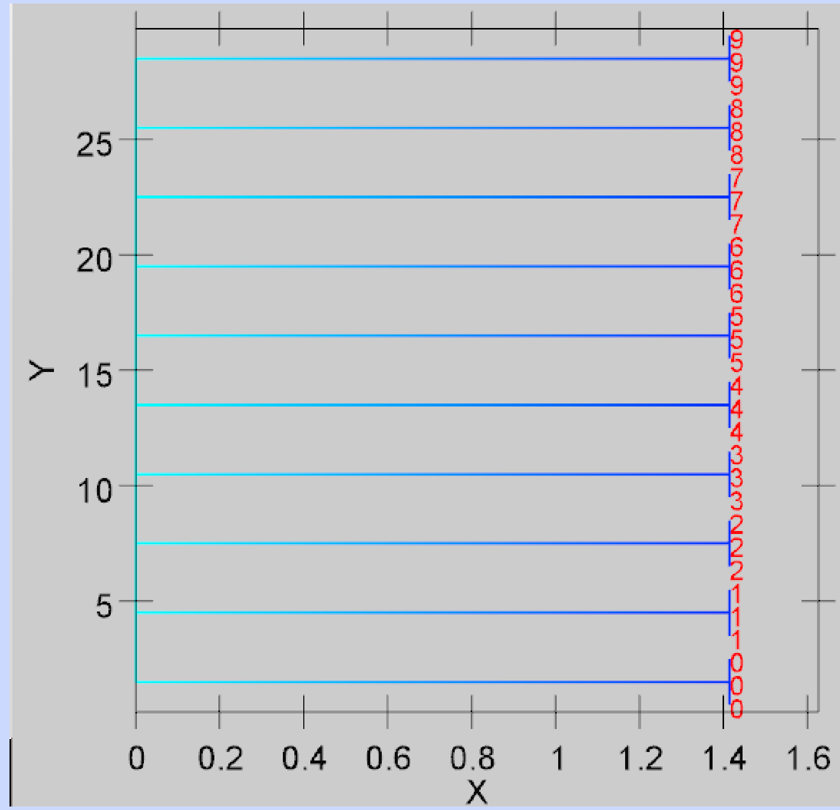

Emphasize distinctions: digits separated, even though they have perceptual overlap.

Emphasizing/Collapsing Distinctions: Categorization

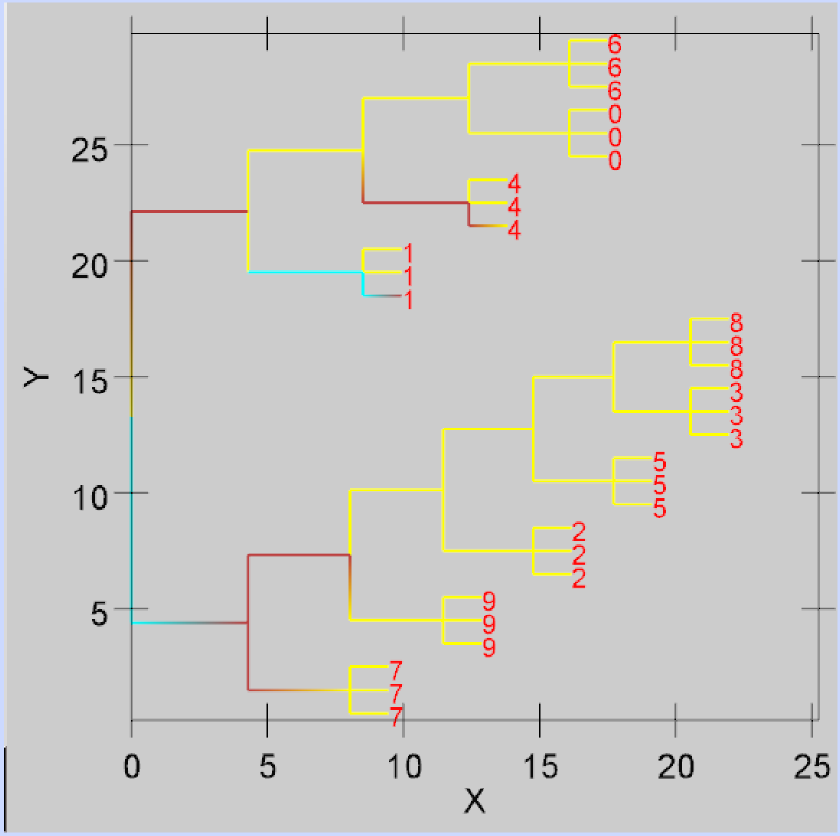

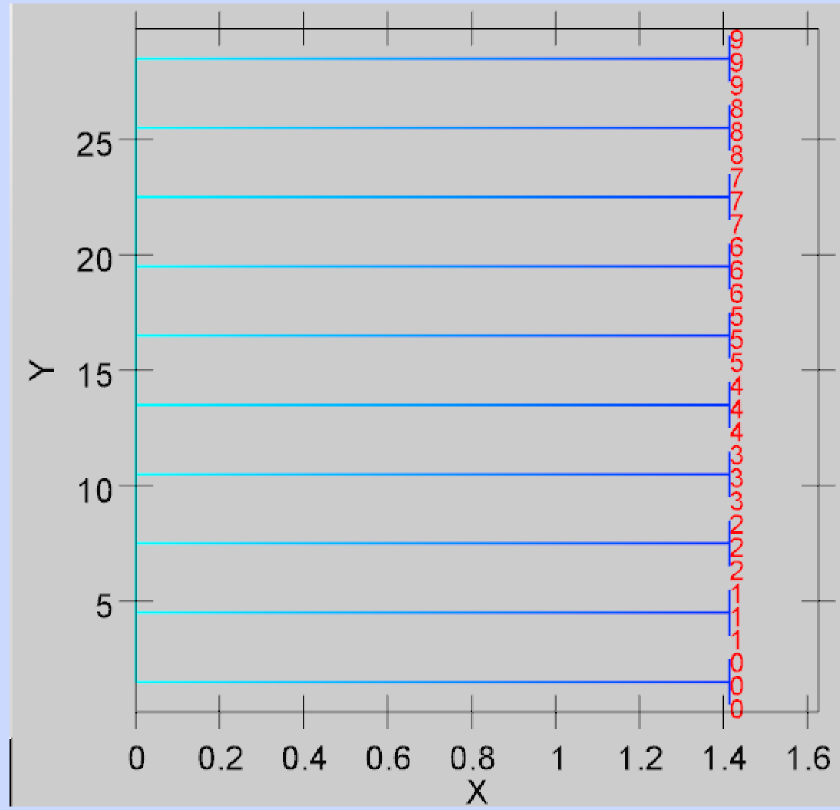

Emphasize distinctions: digits separated, even though they have perceptual overlap.

Collapse distinctions: Noisy digits categorized as same, even though they have perceptual differences.

Detectors are Dedicated, Content-Specific

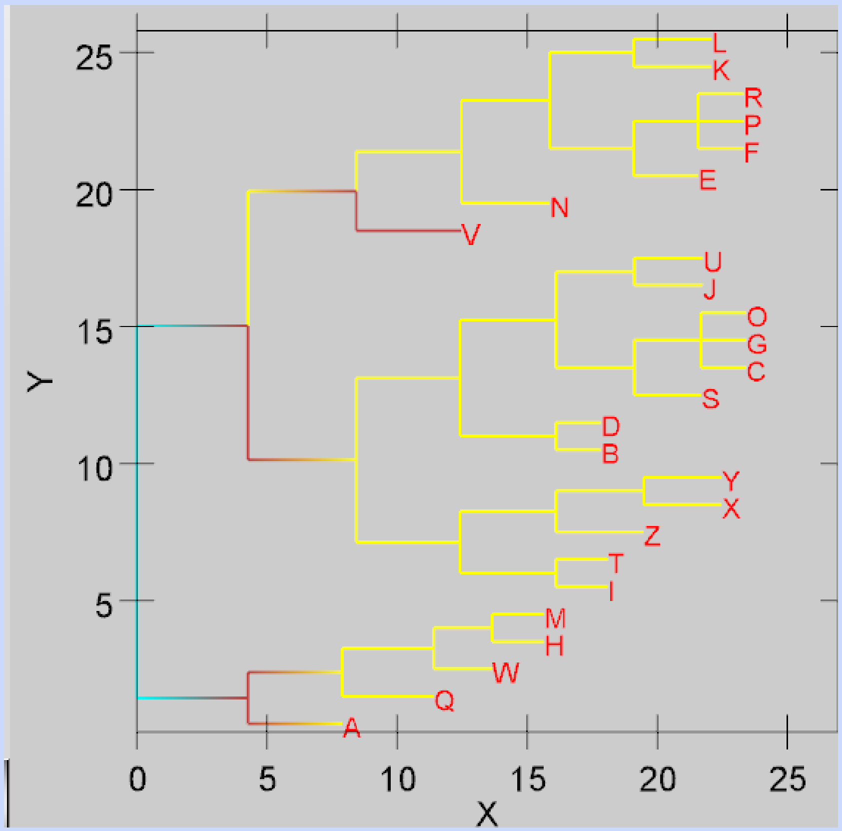

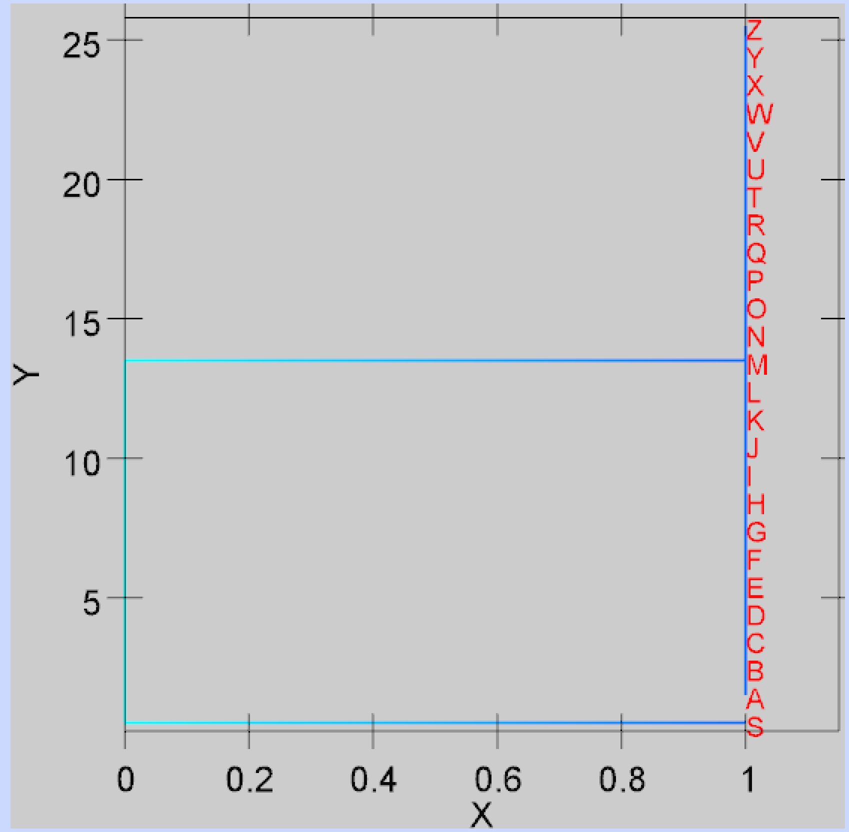

What are these??

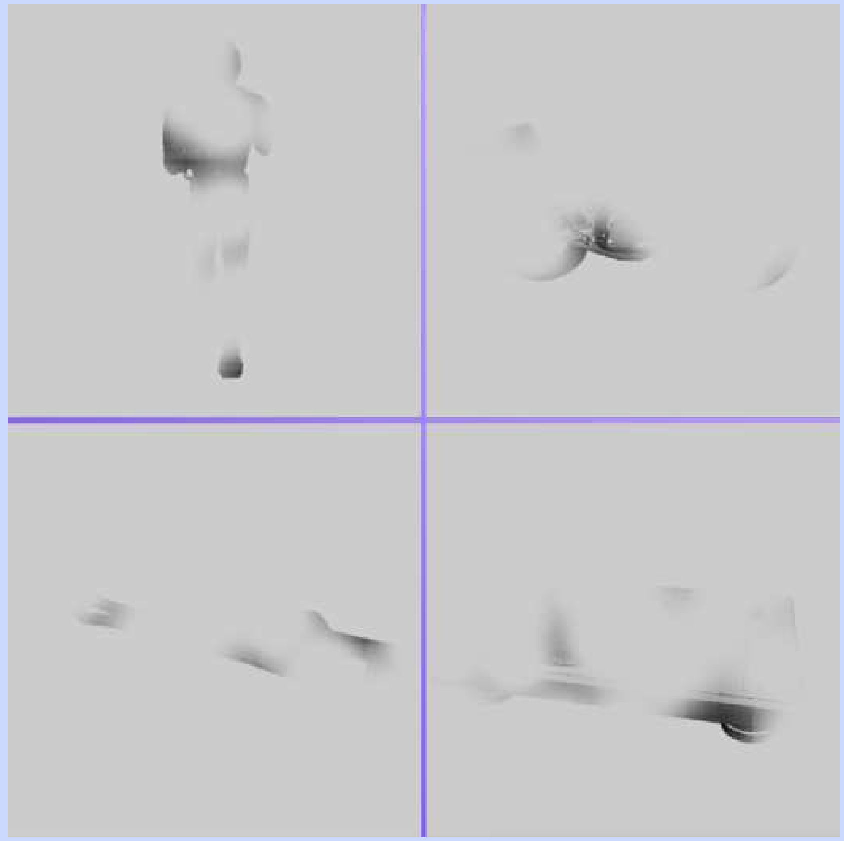

Networks: Bidirectional Connectivity

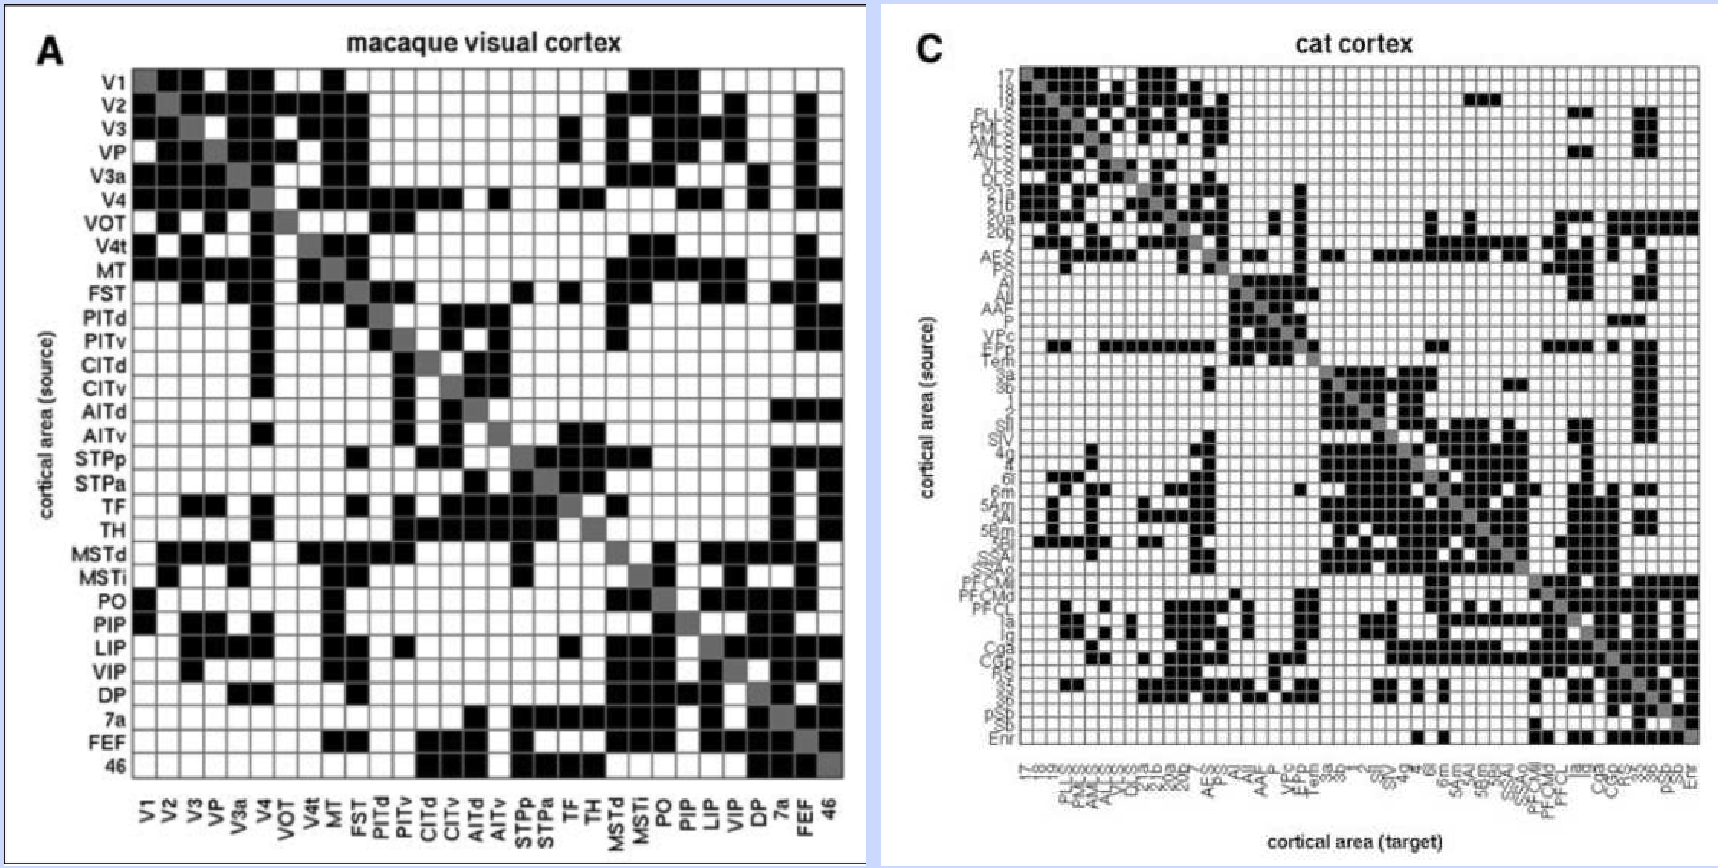

connectivity matrices

Sporns & Zwi 04

Networks: Bidirectional Excitation

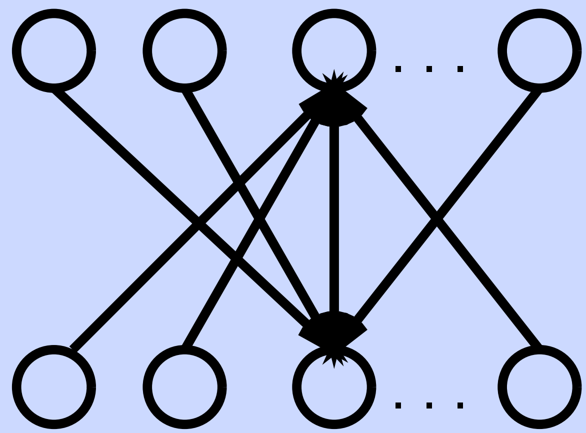

Networks: Bidirectional Excitation

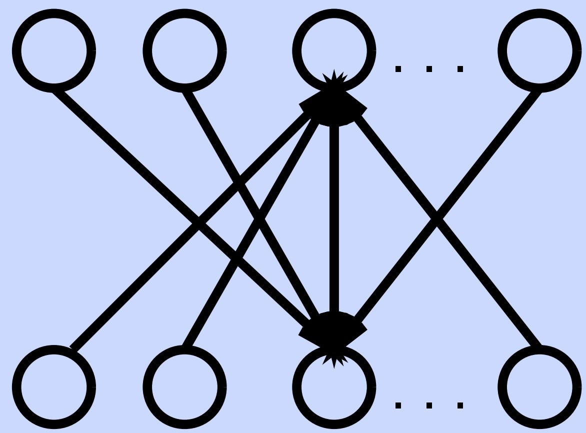

1. Top-down expectations about low-level features.

Networks: Bidirectional Excitation

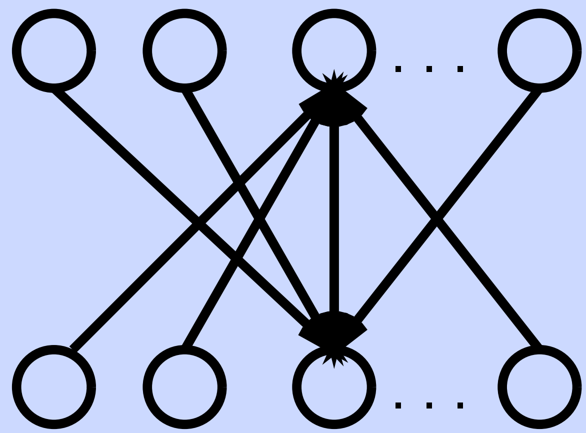

1. Top-down expectations about low-level features.

2. Pattern completion.

Top-Down Processing (Imagery)

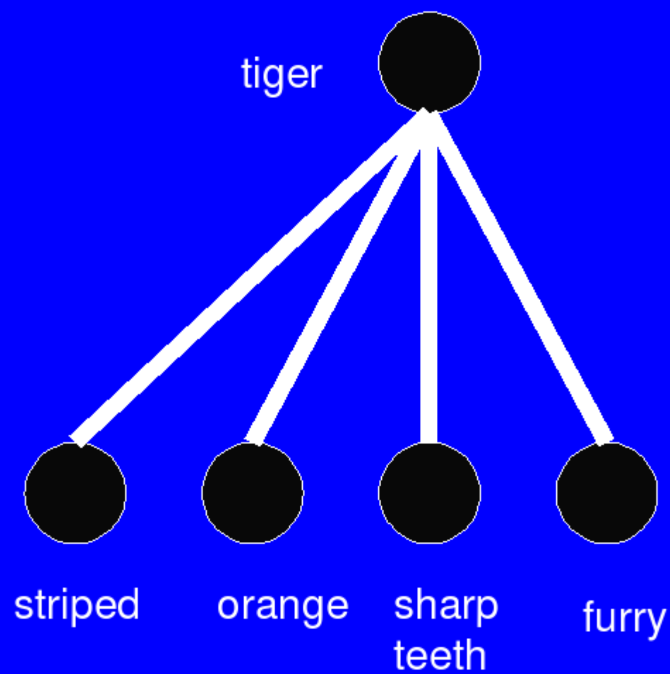

Top-Down Processing (Imagery)

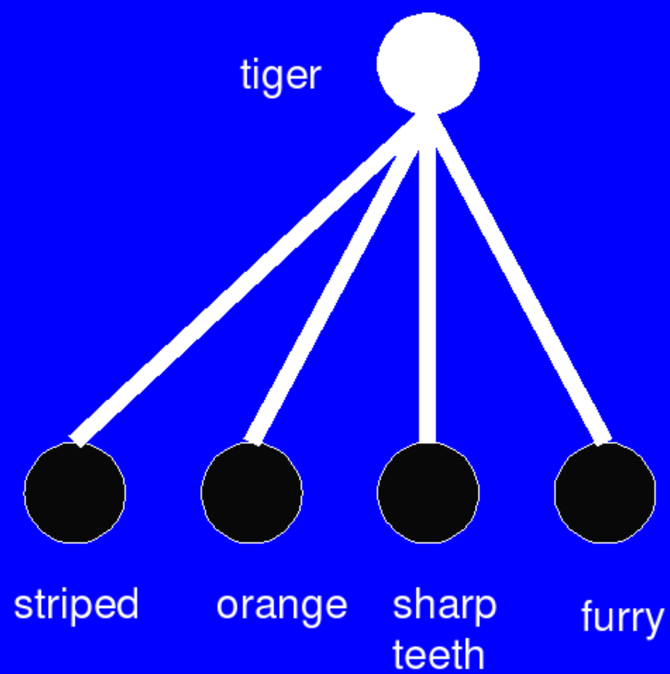

Top-Down Processing (Imagery)

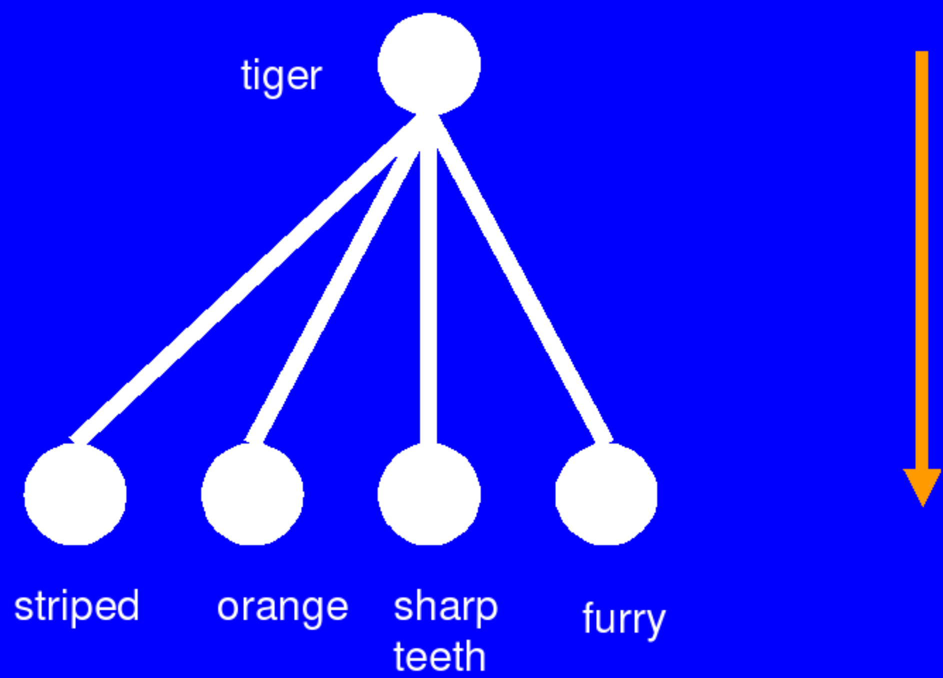

Pattern Completion

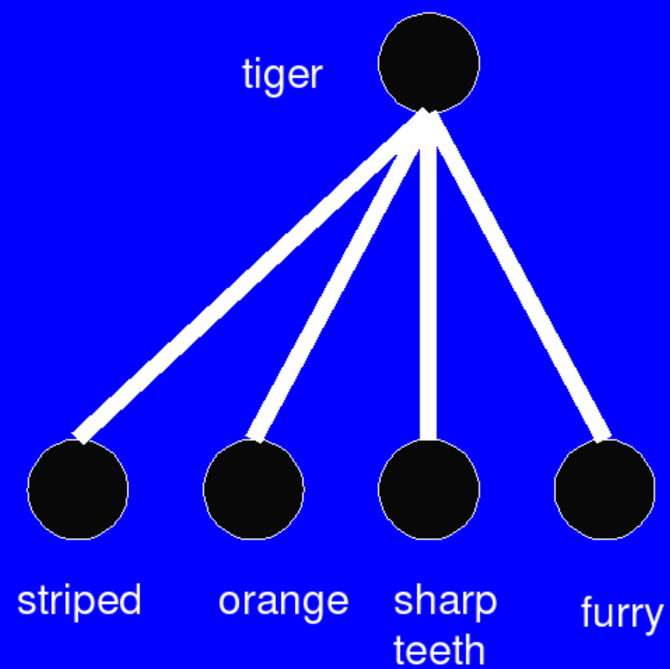

Pattern Completion

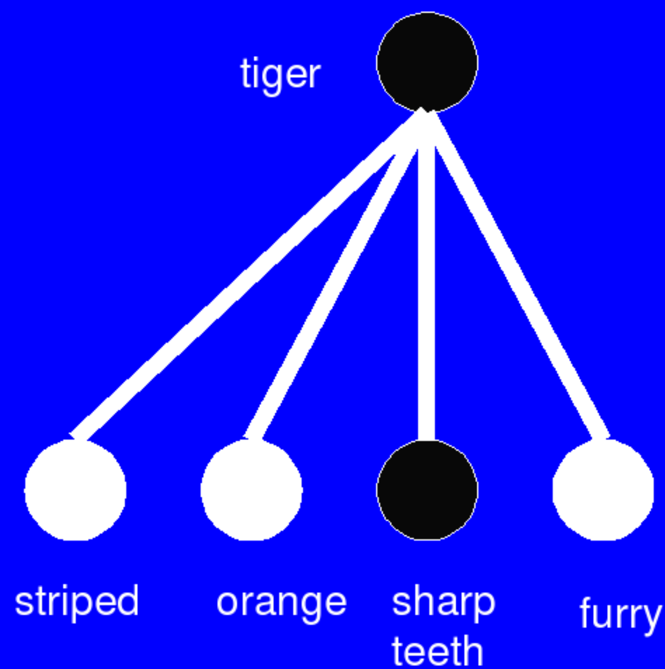

Pattern Completion

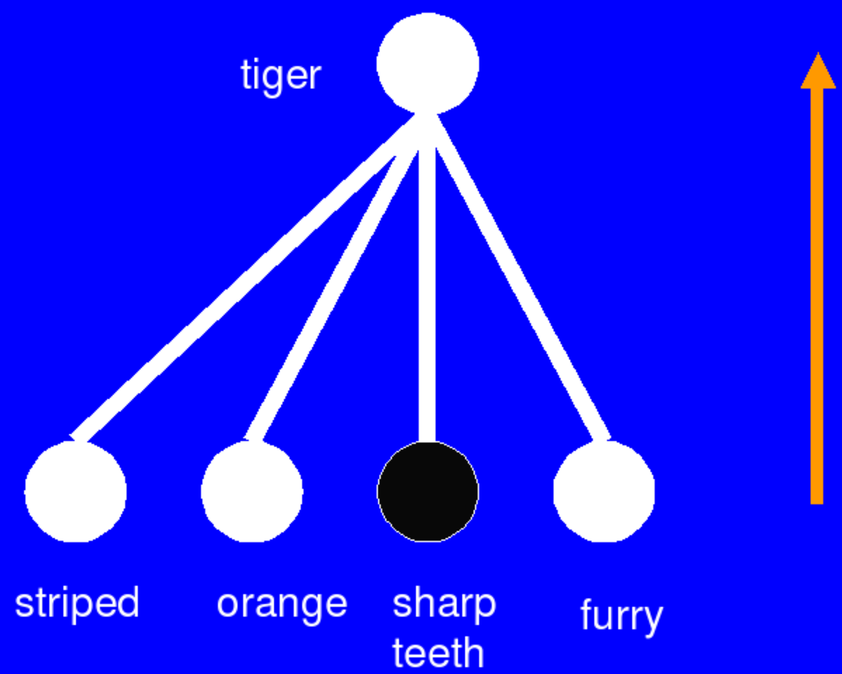

Pattern Completion

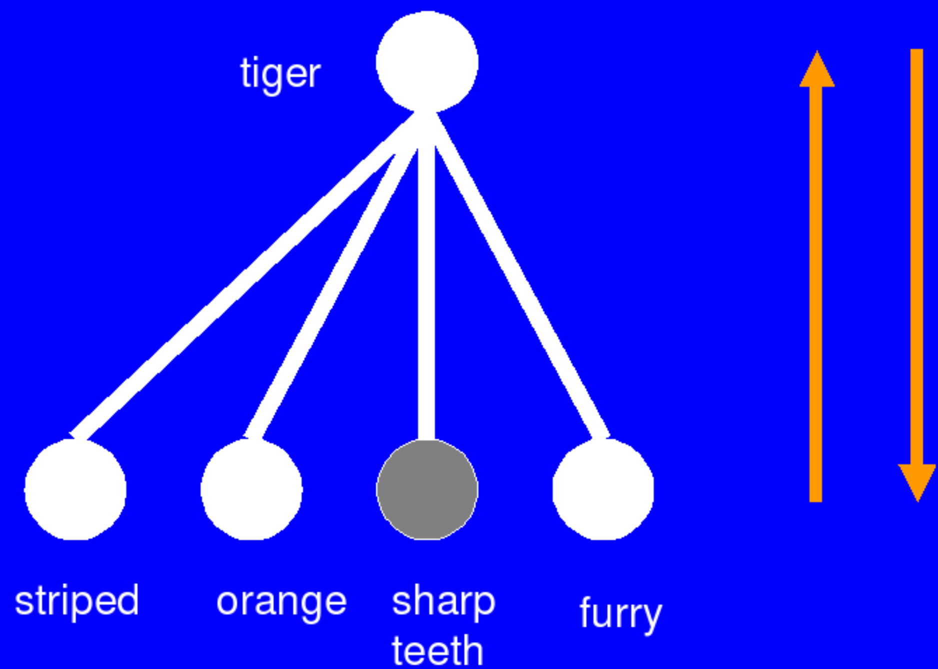

[face_categ.proj]

Word Superiority Effect: Top-Down Amplification

Word Superiority Effect: Top-Down Amplification

Identify second letter in:

NEST (faster)

DEST (slower)

Word Superiority Effect: Top-Down Amplification

Identify second letter in:

NEST (faster)

DEST (slower)

Weird! You have to recognize the letter before you can recognize the word, so how can the word help letter recognition?

[amp_topdown.proj]

[amp_topdown.proj]

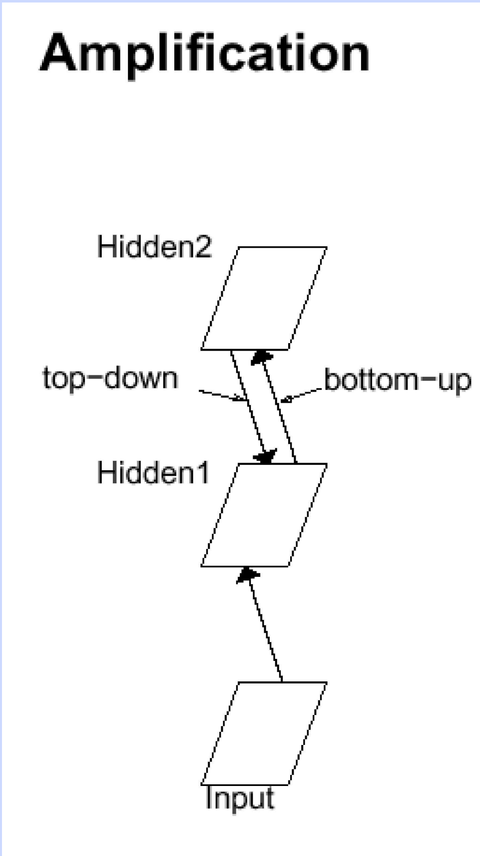

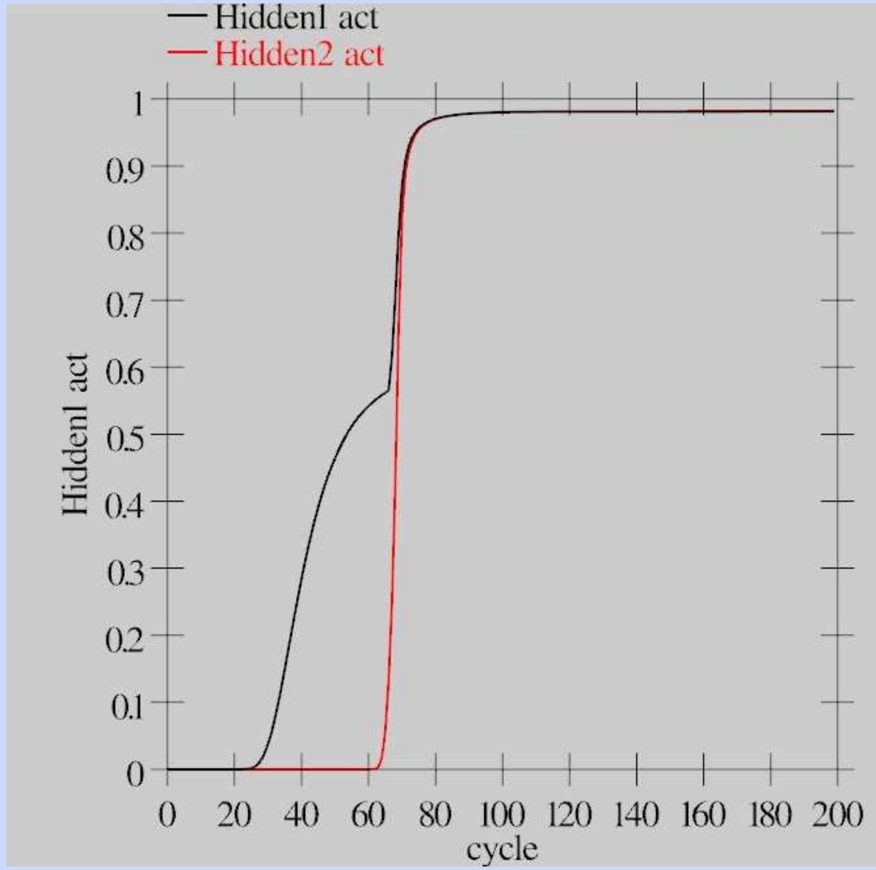

Application to Word Superiority Effect

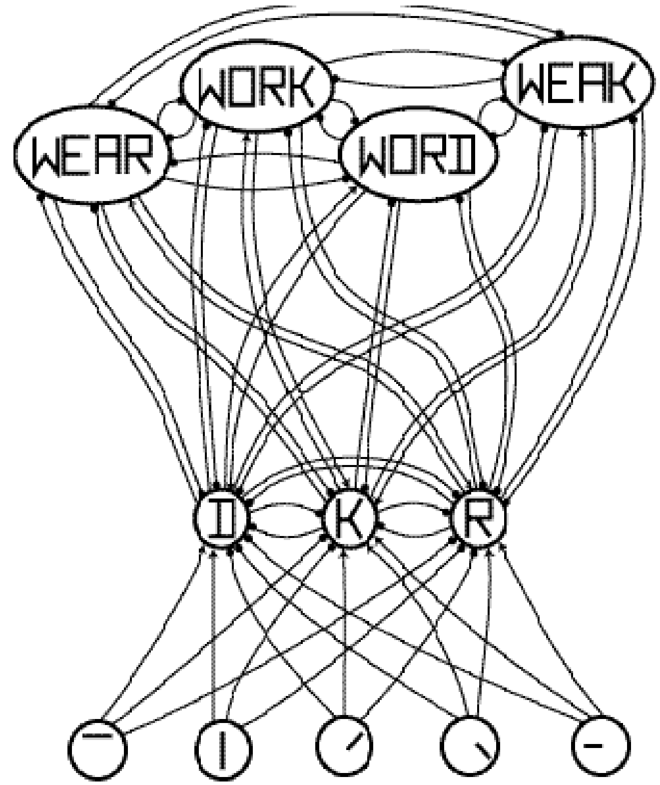

Figure 1. Interactive Activation Network Model
(after McClelland and Rumelhart, 1981).

Bigger network model (details later)

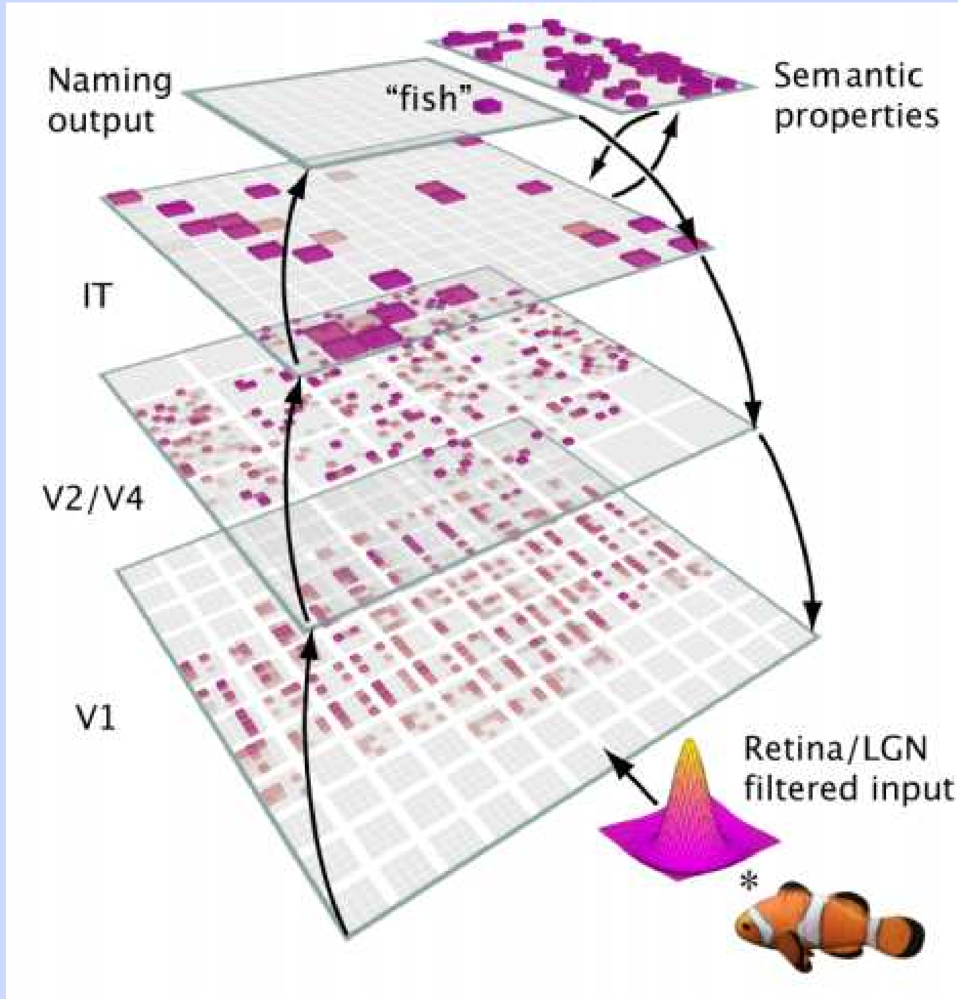

Bidirectional Dynamics

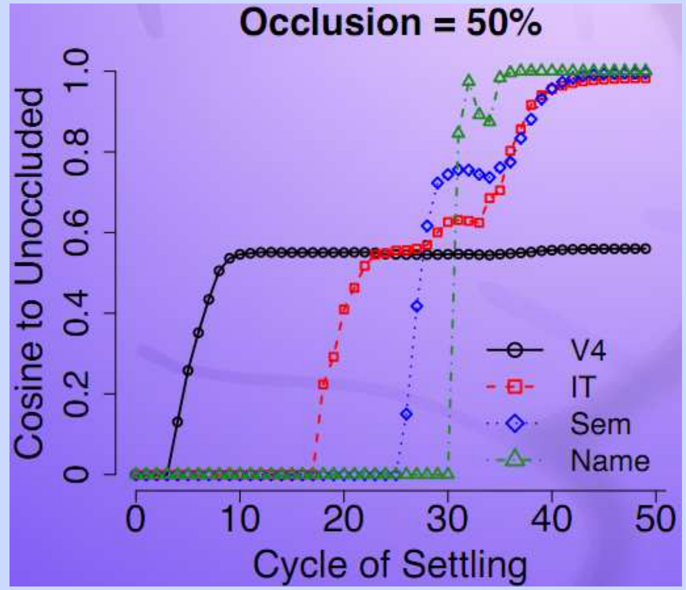

Networks: Bidirectional Excitation

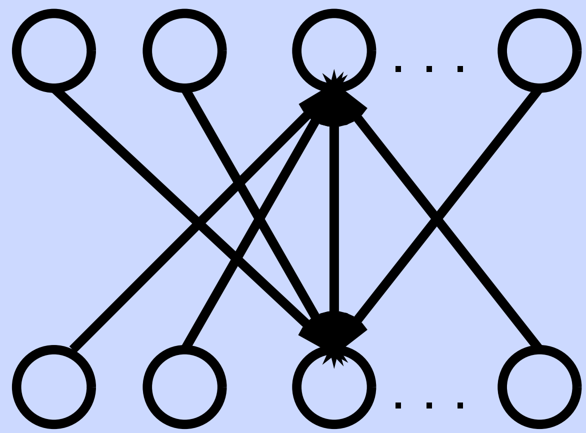

1. Top-down processing ("imagery").

2. Pattern completion.

3. Amplification/bootstrapping.

Networks: Bidirectional Excitation

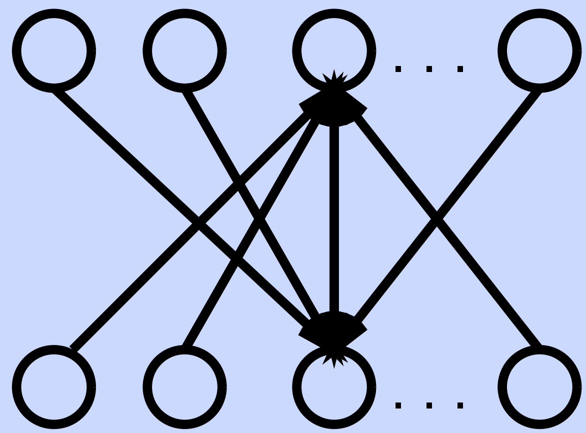

1. Top-down processing ("imagery").

2. Pattern completion.

3. Amplification/bootstrapping.

4. Need inhibition!

Localist vs. Distributed Representations

Localist vs. Distributed Representations

- Localist = 1 unit responds to 1 thing (e.g., digits, grandmother cell).

Localist vs. Distributed Representations

- Localist = 1 unit responds to 1 thing (e.g., digits, grandmother cell).
- Distributed = Many units respond to 1 thing, one unit responds to many things.

Localist vs. Distributed Representations

- Localist = 1 unit responds to 1 thing (e.g., digits, grandmother cell).
- Distributed = Many units respond to 1 thing, one unit responds to many things.
- With distributed representations, units correspond to stimulus *features* as opposed to complete stimuli

Digits With Distributed Representations

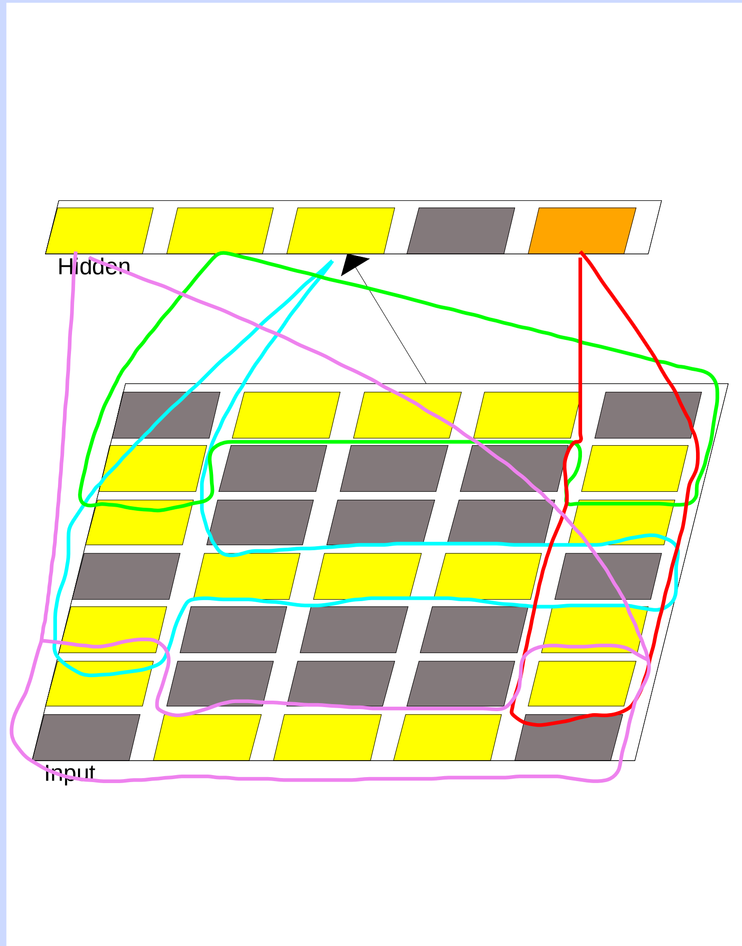

Advantages of Distributed Representations

Efficiency: Fewer Units Required

Advantages of Distributed Representations

Efficiency: Fewer Units Required

The digits network can represent 10 digits using 5 “feature” units

Each digit is a unique combination of the 5 features, e.g.,

“0” = feature 3
“1” = features 1, 4
“2” = features 1, 2
“3” = features 1, 2, 5
“4” = features 3, 4
“5” = features 1, 2, 3

There are 32 unique ways to combine 5 features
There are $>$ 1 million unique ways to combine 20 features

Advantages of Distributed Representations

Similarity and **Generalization**:

Advantages of Distributed Representations

Similarity and **Generalization**:

If you represent stimuli in terms of their constituent features, stimuli with similar features get assigned similar representations

This allows you to generalize to novel stimuli based on their similarity to previously encountered stimuli oriole


flies
oriole
Photo by Chan Robbins

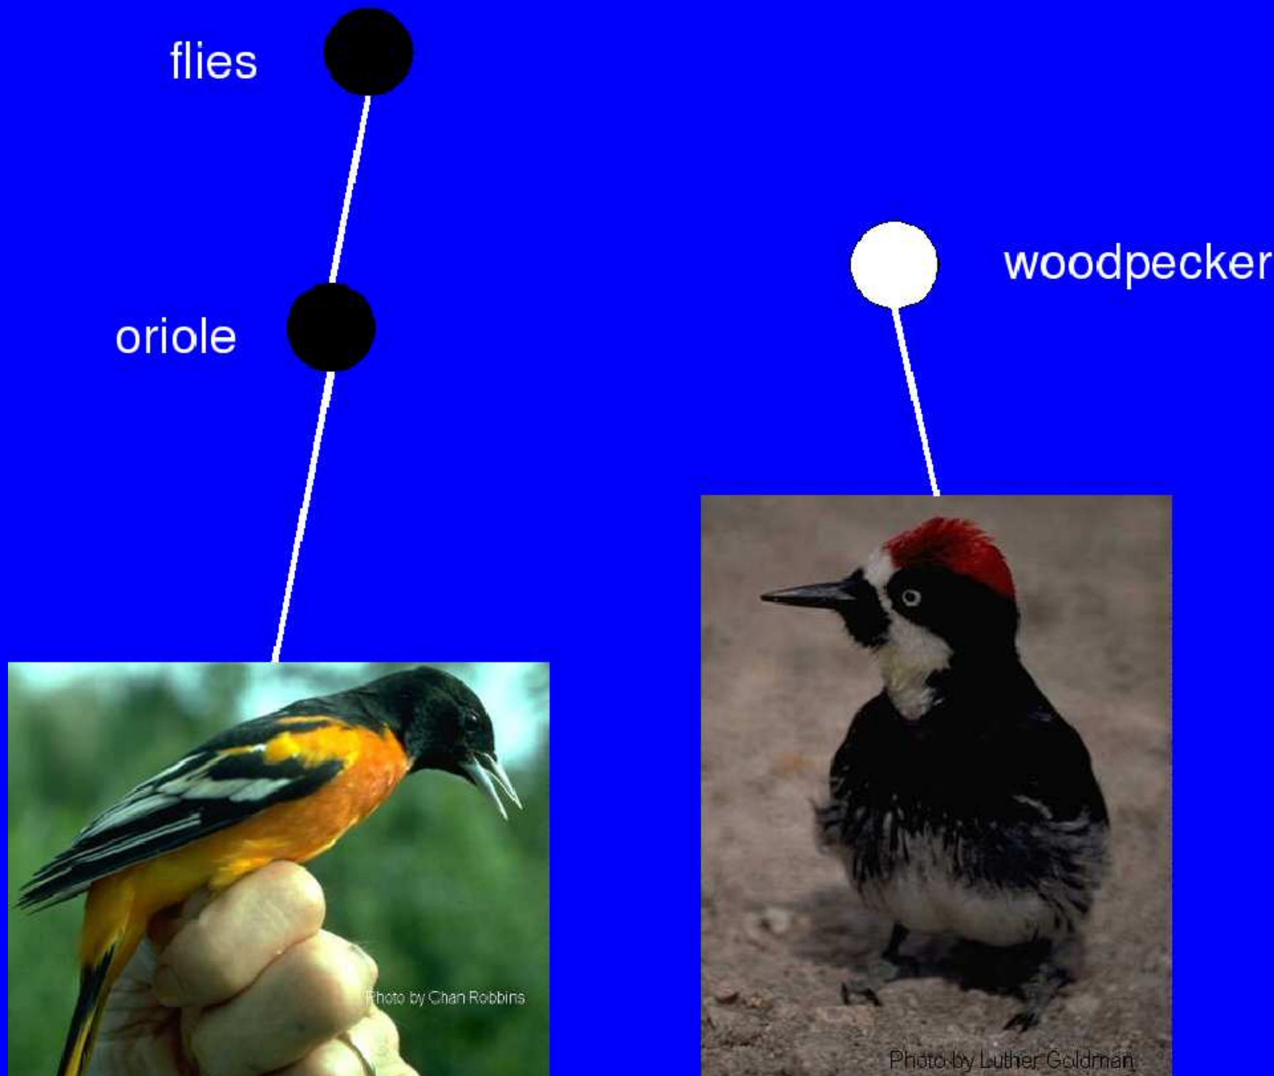
flies
oriole
woodpecker
Photo by Chan Robbins
Photo by Luther Goldman

localist = no generalization

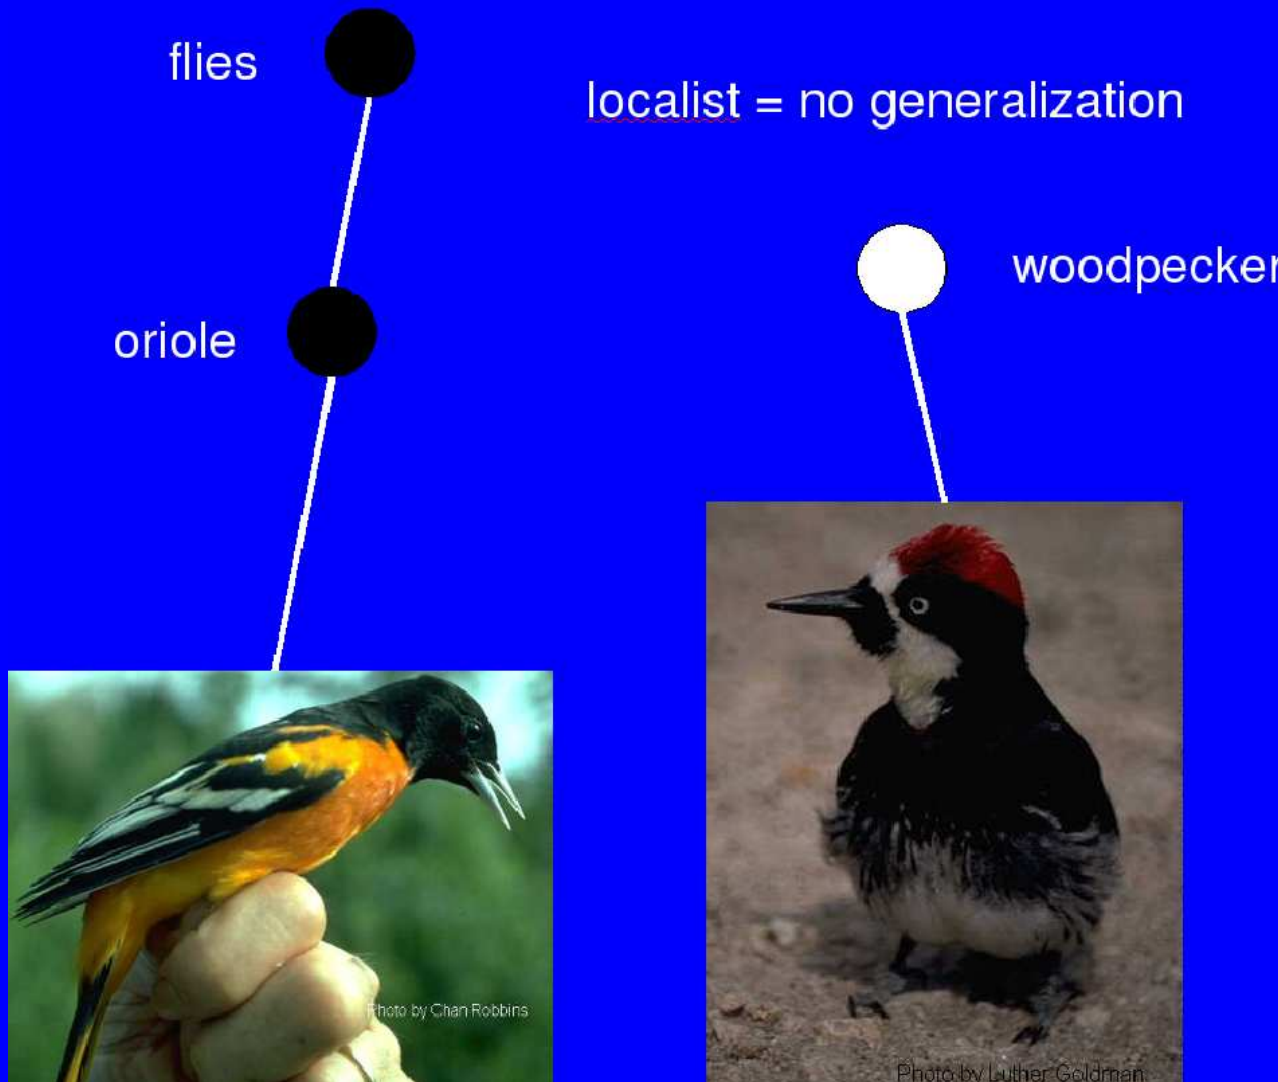

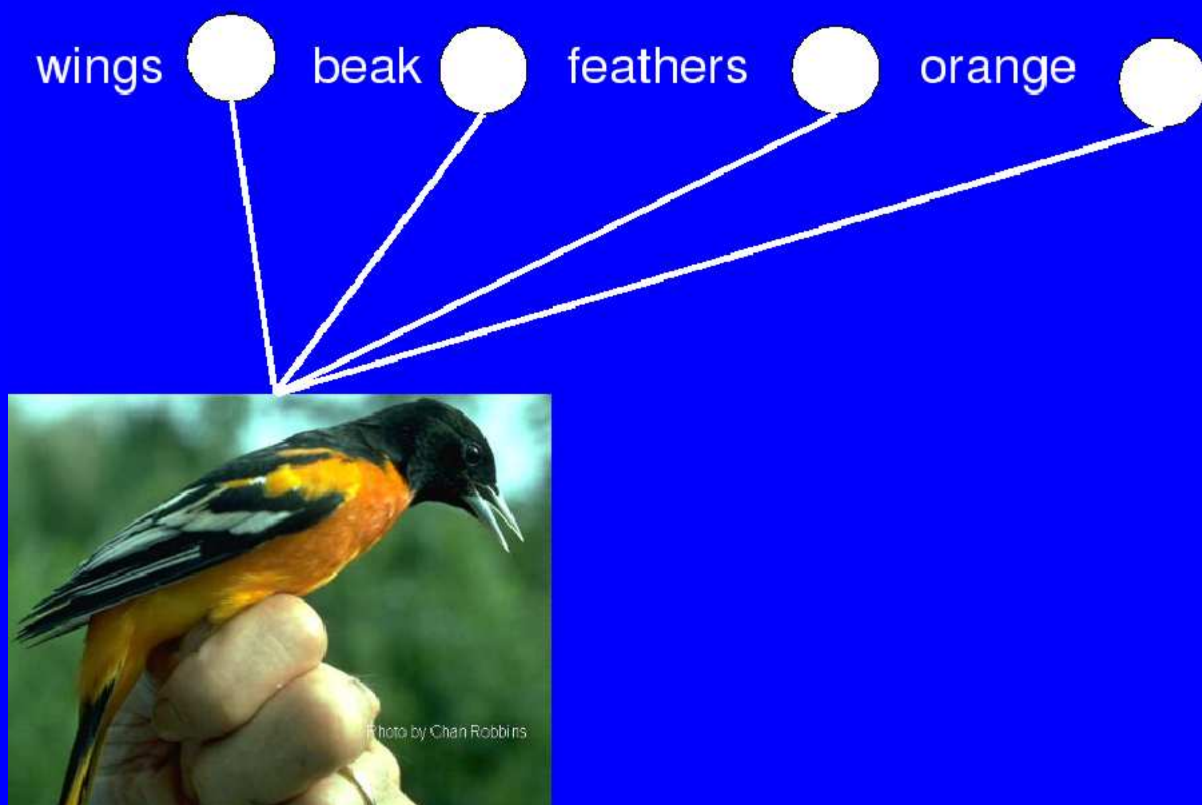
wings
beak
feathers
orange
Photo by Chan Robbins

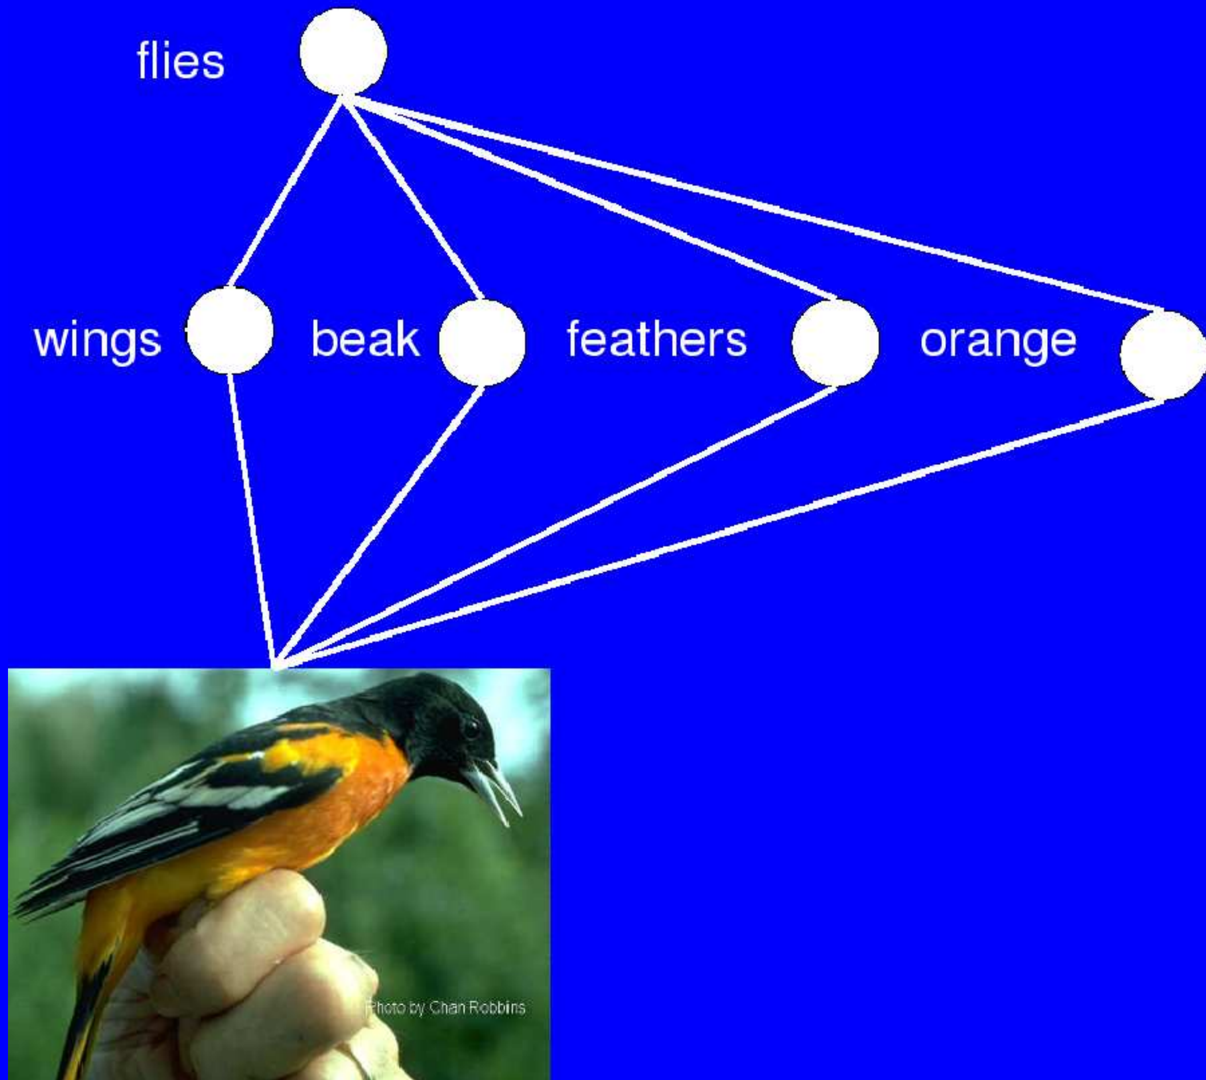
flies
wings
beak
feathers
orange
Photo by Chan Robbins

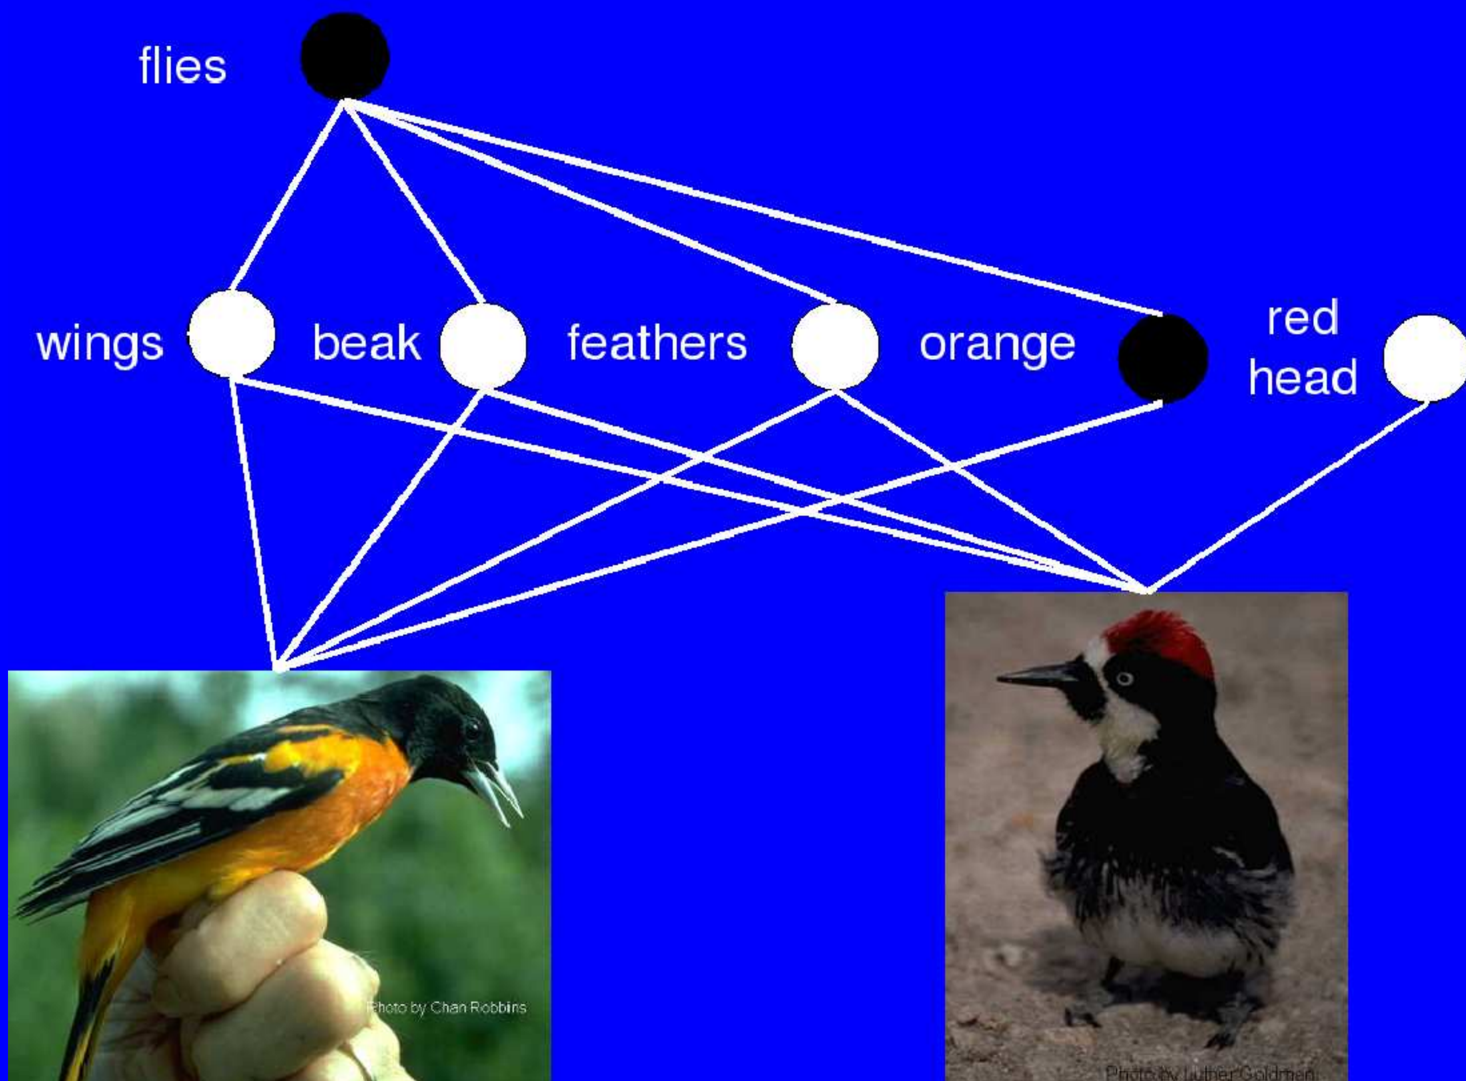
flies
wings
beak
feathers
orange
red
head
Photo by Chan Robbins
Photo by Luther Goldman

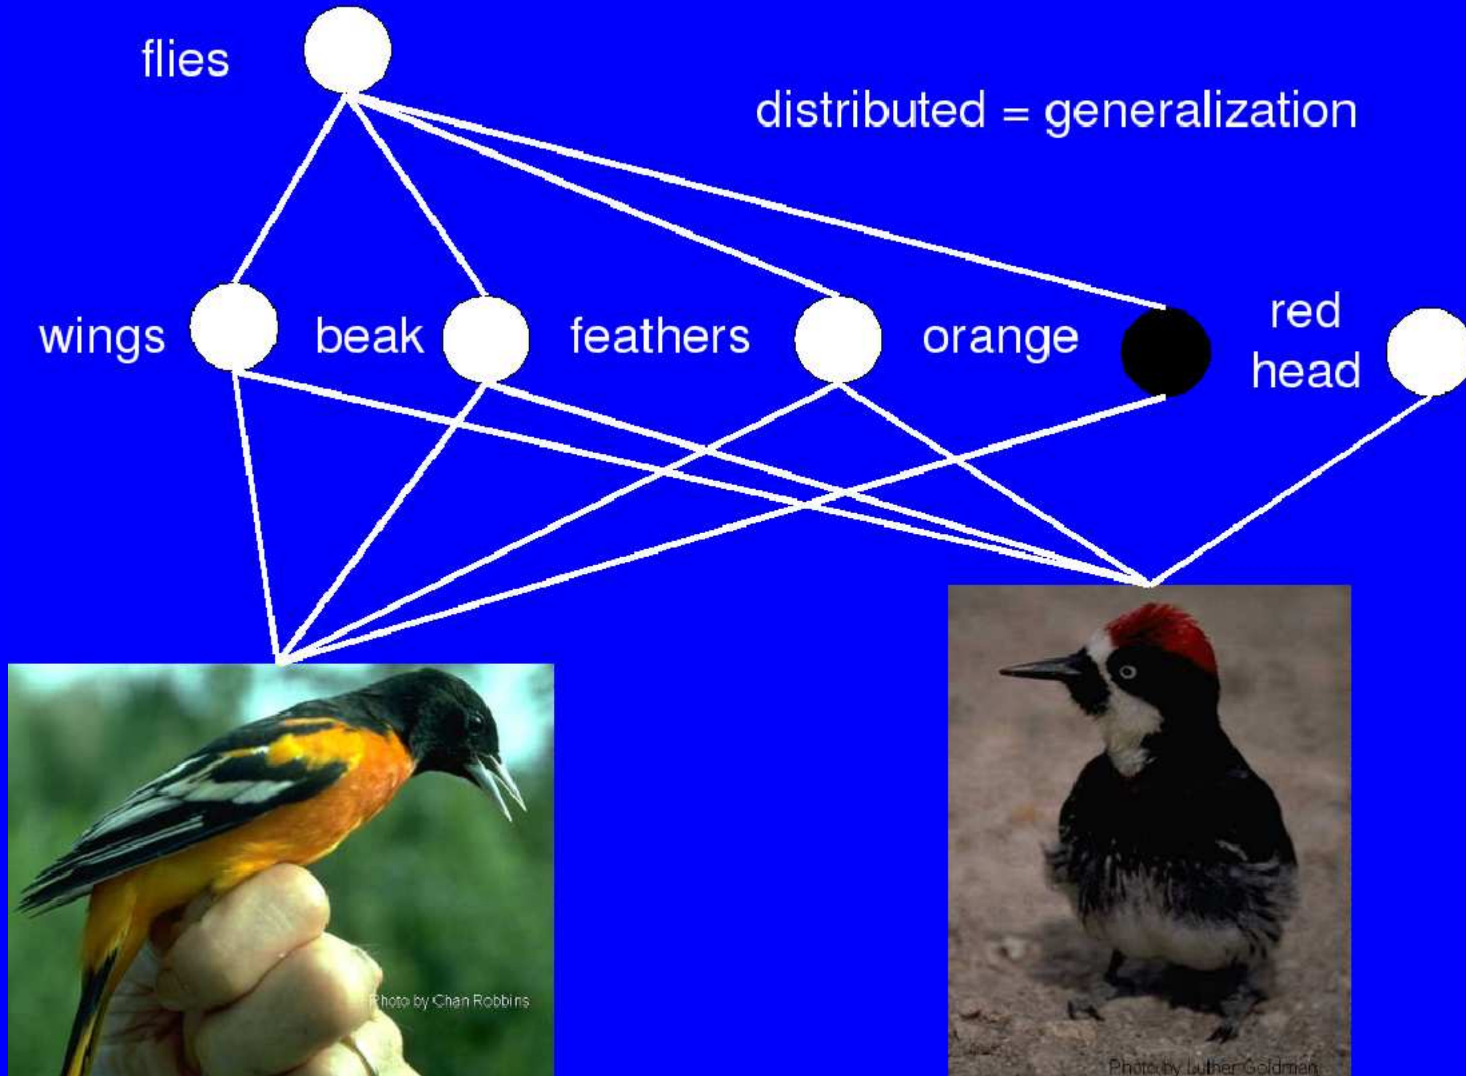
flies
distributed = generalization
wings
beak
feathers
orange
red
head
Photo by Chan Robbins
Photo by Luther Goldman

Advantages of Distributed Representations

Robustness (Graceful Degradation):

Advantages of Distributed Representations

Robustness (Graceful Degradation):

Damage has less of an effect on networks with distributed (vs. localist) representations

Advantages of Distributed Representations

Robustness (Graceful Degradation):

Damage has less of an effect on networks with distributed (vs. localist) representations

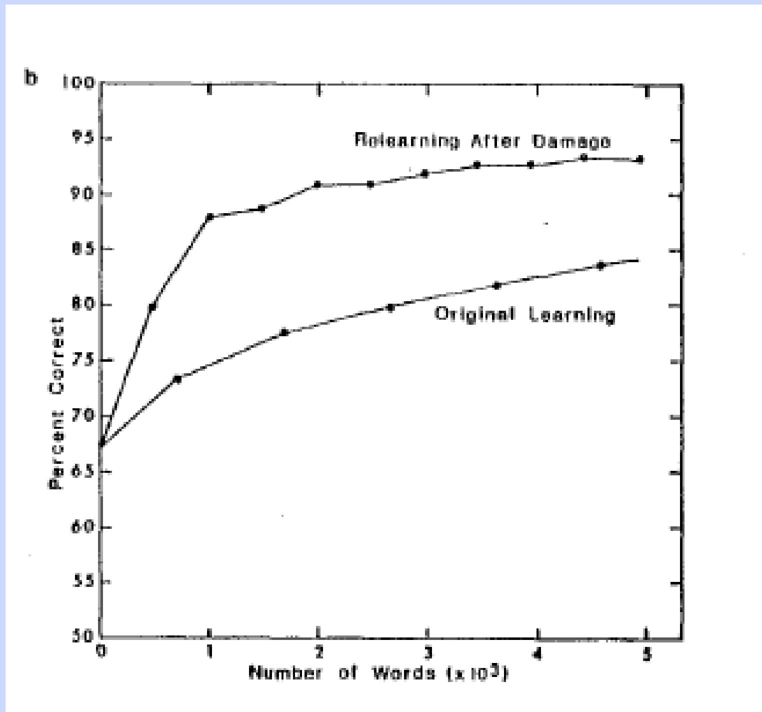


flies
oriole
Photo by Chan Robbins

flies

??

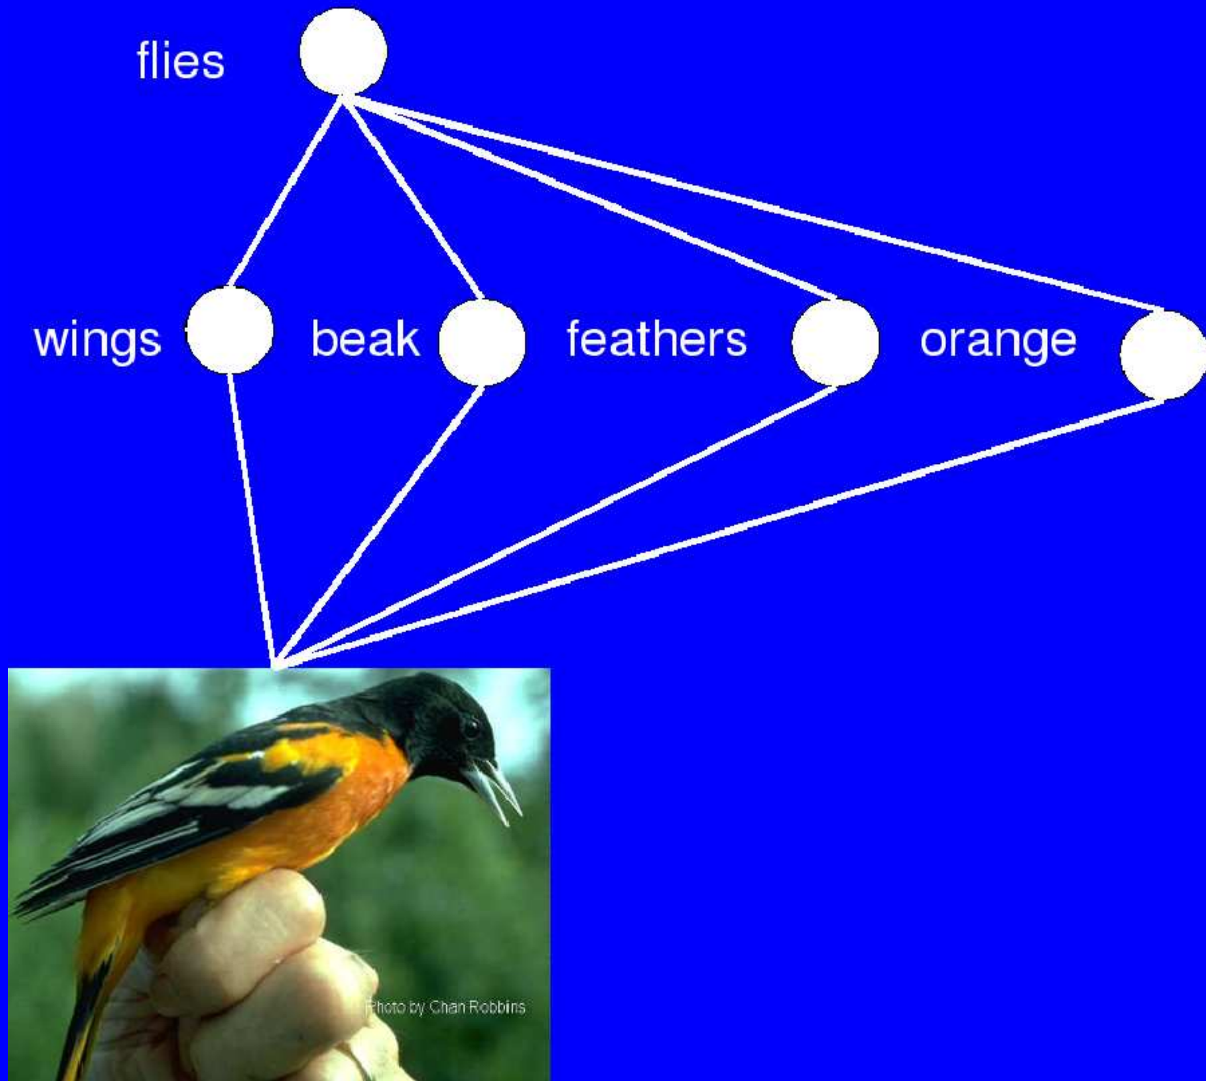
flies
wings
beak
feathers
orange
Photo by Chan Robbins

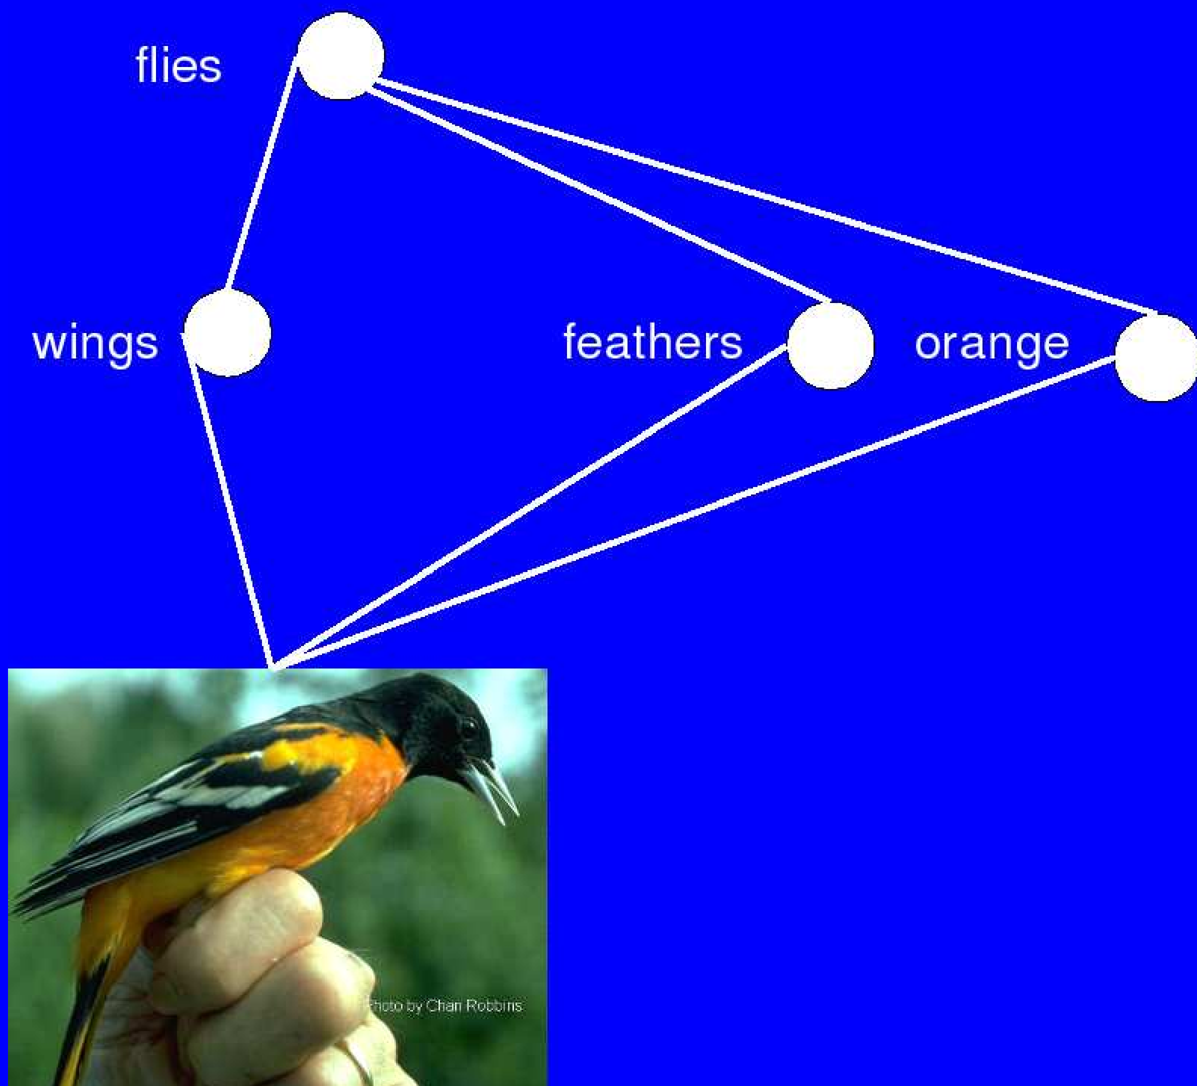
flies
wings
feathers
orange
Photo by Chan Robbins

Advantages of Distributed Representations

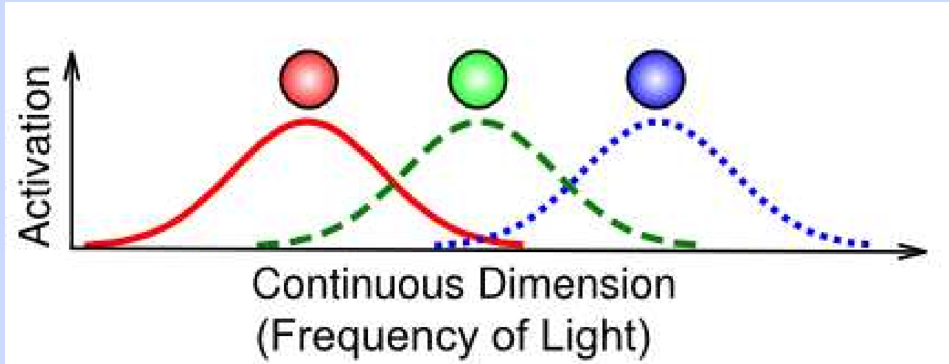

Accuracy: By coarse-coding. (e.g, color, position)

Advantages of Distributed Representations

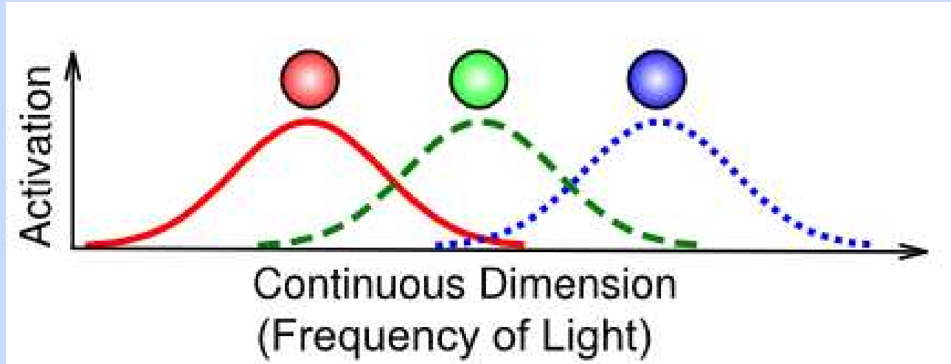

Accuracy: By coarse-coding. (e.g, color, position)

These tuning curves are commonly seen, e.g. in V1:

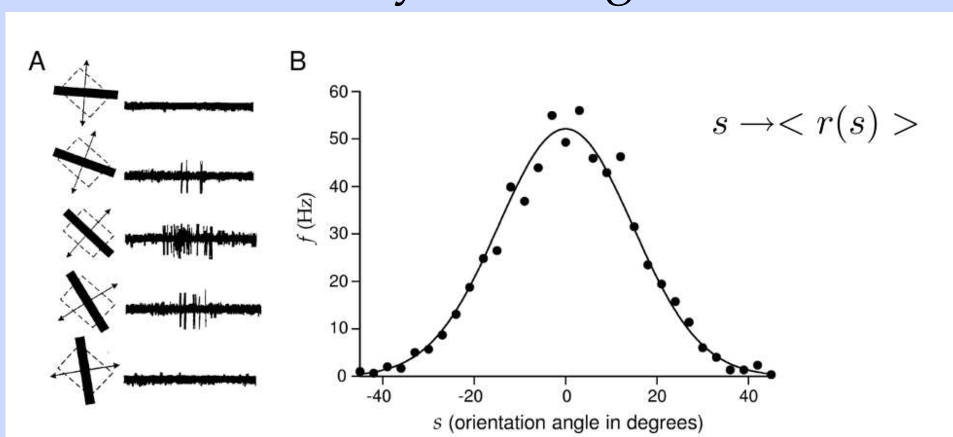

Advantages of Distributed Representations

Advantages of Distributed Representations

Efficiency: Fewer total units required.

Similarity: As a function of overlap.

Generalization: Can use novel combinations.

Robustness: Redundancy: damage has less of an effect

Accuracy: By coarse-coding.

Networks

1. Biology: The cortex
2. Excitation:
   - Unidirectional (transformations)
   - Bidirectional (top-down processing, pattern completion, amplification)
3. Inhibition: Controlling bidirectional excitation.
4. Constraint Satisfaction: Putting it all together.