1 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

2 Networks

Layers and layers of detectors...

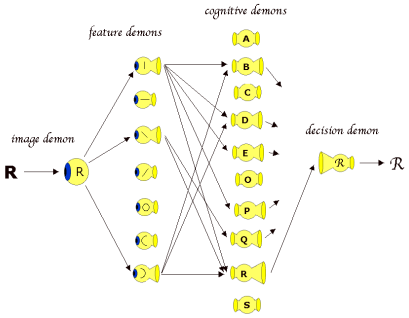

3 Networks

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.

4 Cortex: Neurons

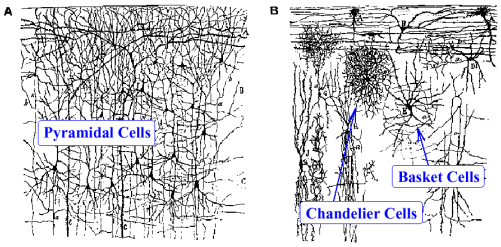

Two **separate** populations:

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

5 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

6 The Neuron and its Ions

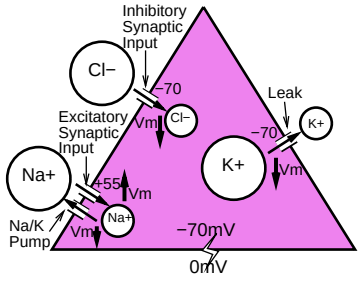

Glutamate →

7

The Neuron and its Ions

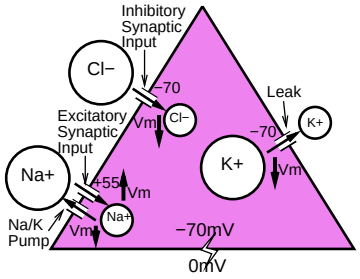

Glutamate → opens Na+ channels →

8

The Neuron and its Ions

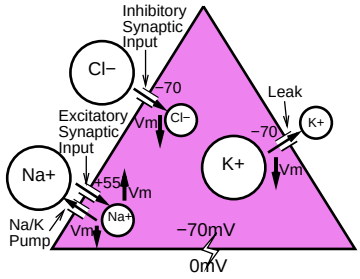

Glutamate → opens Na+ channels → Na+ enters (excitatory)

9

The Neuron and its Ions

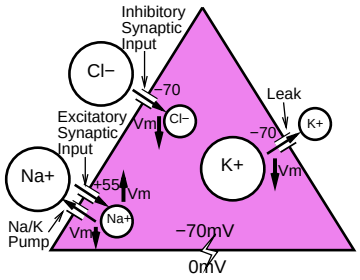

Glutamate → opens Na+ channels → Na+ enters (excitatory)

GABA →

10

The Neuron and its Ions

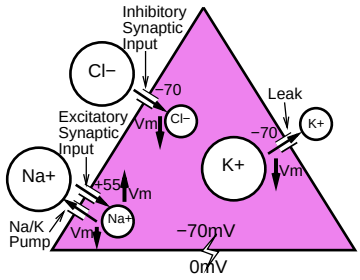

Glutamate → opens Na+ channels → Na+ enters (excitatory)

GABA → opens Cl- channels →

11

The Neuron and its Ions

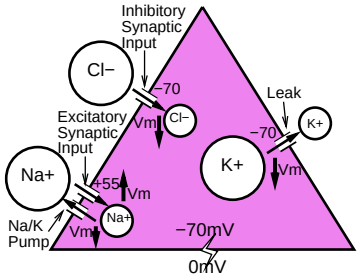

Glutamate → opens Na+ channels → Na+ enters (excitatory)

GABA → opens Cl- channels → Cl- enters if $V_m \uparrow$ (inhibitory)

12

Laminar Structure of Cortex

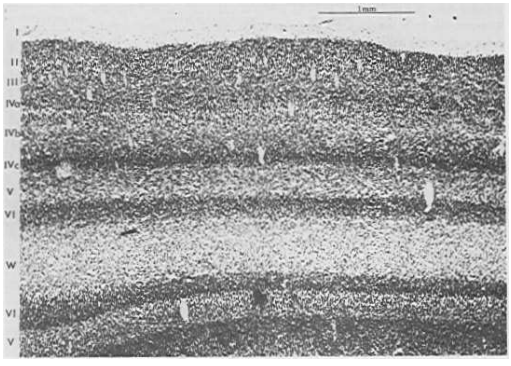

13

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)

14

Laminar Structure of Cortex
Three *Functional* Layers

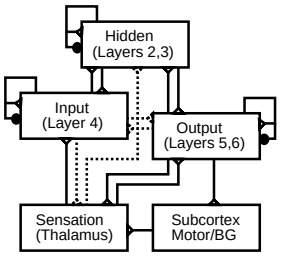

Hidden layer *transforms* input-output mappings

More hidden layers → richer transformations → less reflex-like...

smarter? more "free"?

1

Area Structure of Cortex

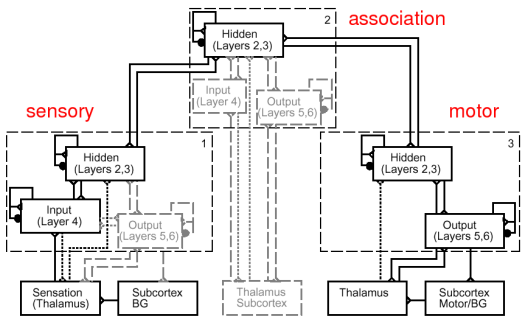

16

Area Structure of Cortex

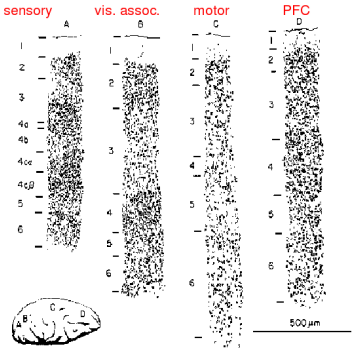

17

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.

18

Excitation (Unidirectional): Transformations

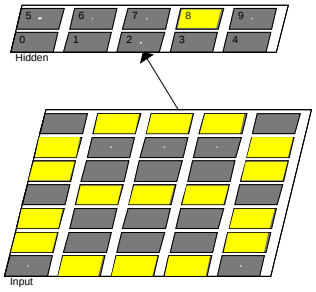

- Detectors work in parallel to *transform* input activity pattern to hidden activity pattern.

19 Excitation (Unidirectional): Transformations

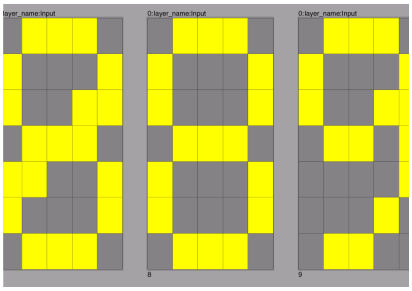

- Emphasizes some distinctions, collapses across others.
- Function of what detectors detect (and what they ignore).

20 Emphasizing/Collapsing Distinctions

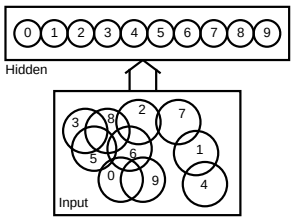

Other (more interesting) examples?...

21 [transform.proj]

digit detectors:

- tested with noisy digits
- tested with letters

22 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

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

- Example...

23 Making Friends With Cluster Plots

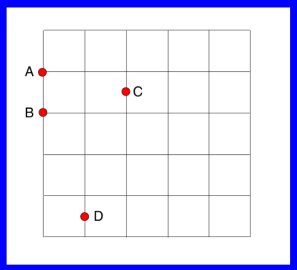

24 Making Friends With Cluster Plots

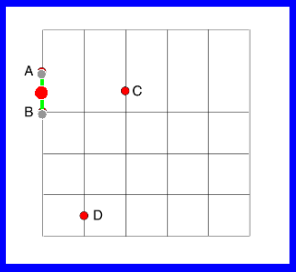

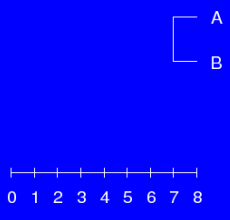

25 Making Friends With Cluster Plots

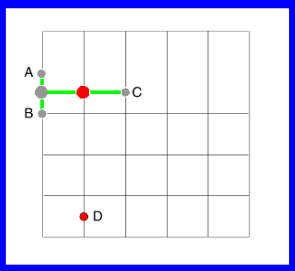

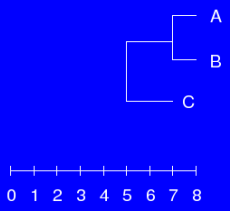

26 Making Friends With Cluster Plots

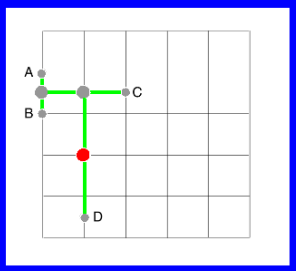

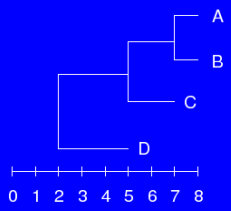

27 [transform.proj]

cluster plots (digits, noisy digits, hidden).

28 Emphasizing/Collapsing Distinctions: Categorization

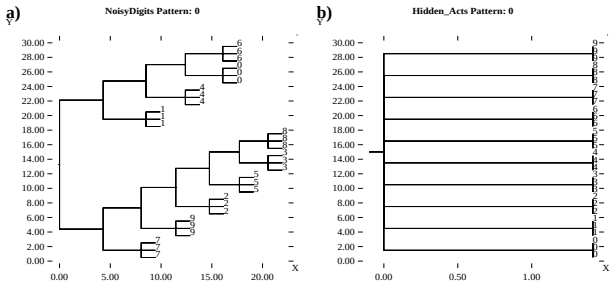

Emphasize distinctions: Different digits separated, even though they have perceptual overlap.

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

29 Detectors are Dedicated, Content-Specific

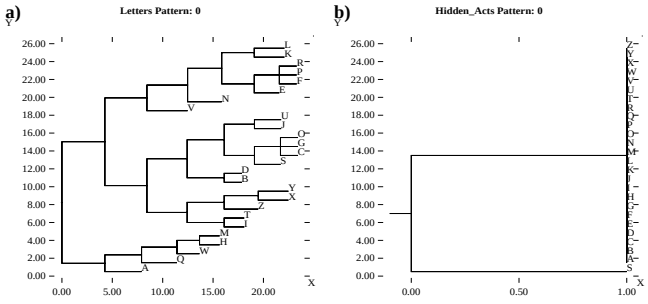

30 Localist vs. Distributed Representations

- 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

31 Digits With Distributed Representations

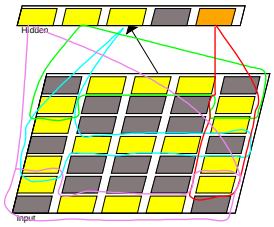

32 Advantages of Distributed Representations

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

33 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

34

35

36

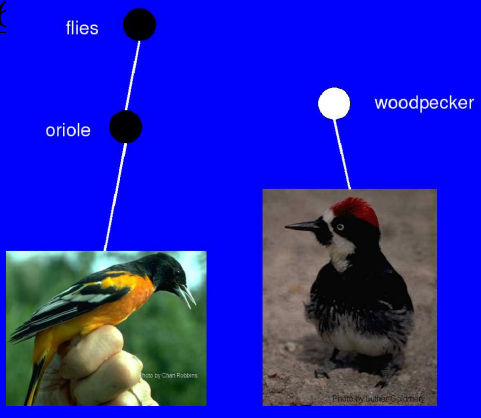

37

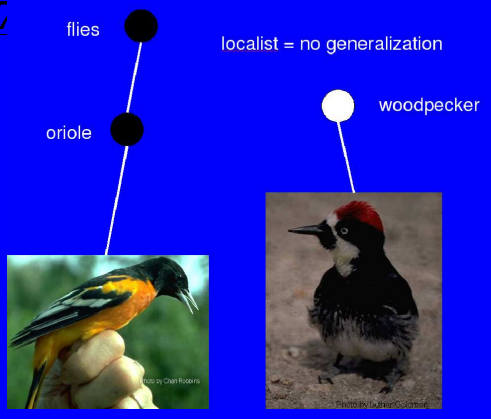

38

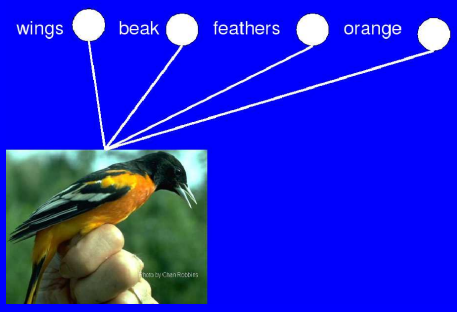

39

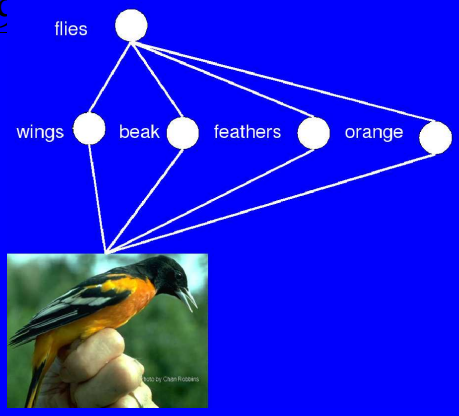

40

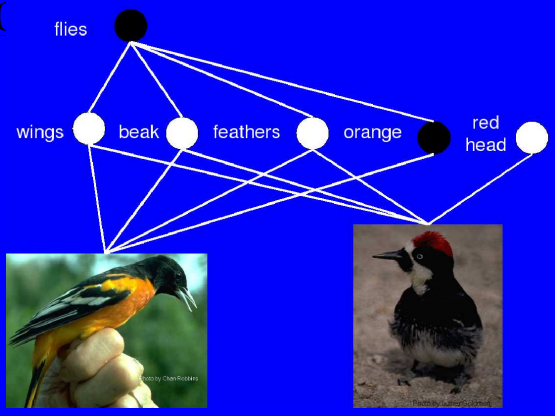

41

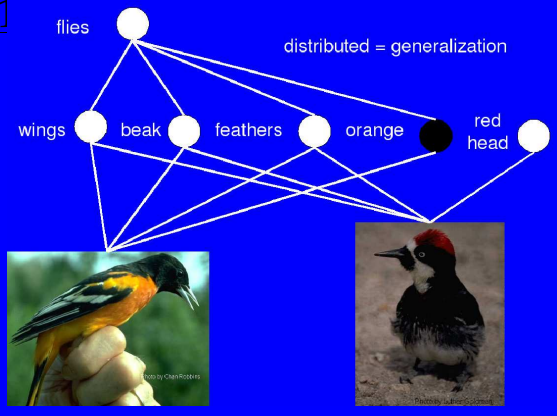

42 Advantages of Distributed Representations

Robustness (Graceful Degradation):

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

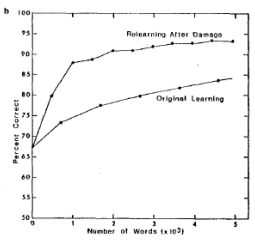

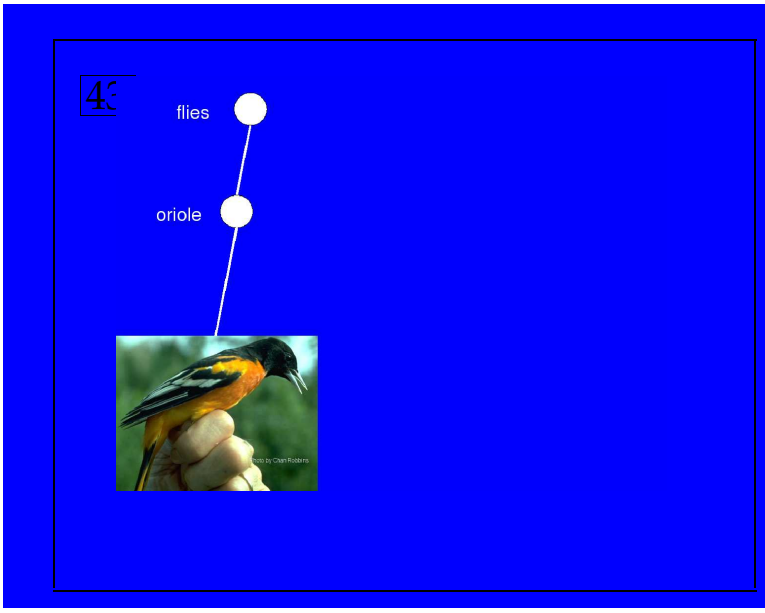

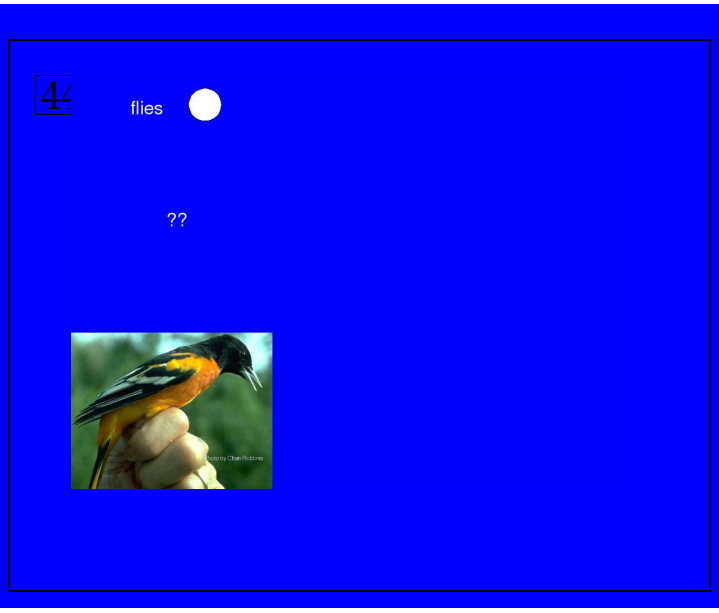

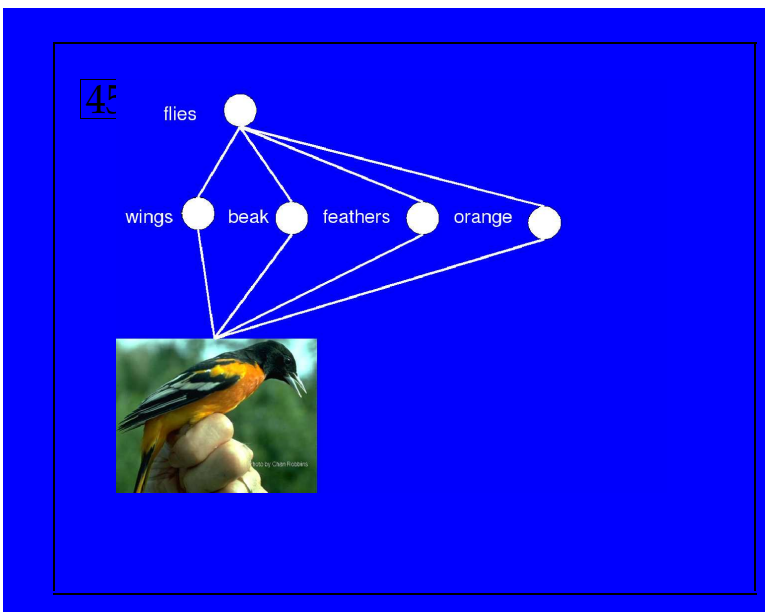

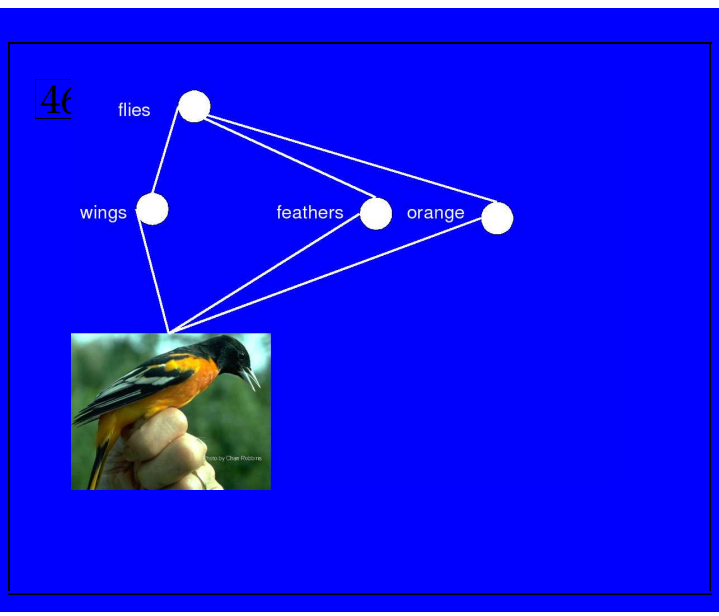

47 Advantages of Distributed Representations

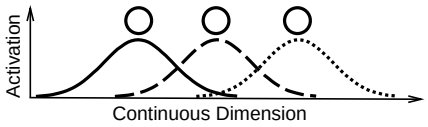

Accuracy: By coarse-coding. (e.g, color, position)

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

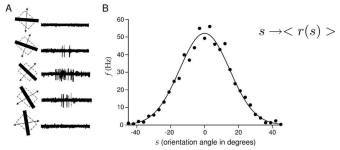

48 Advantages of Distributed Representations

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.

49

Networks: Bidirectional Excitation

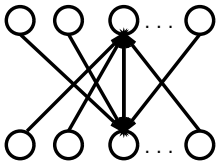

1. Top-down expectations about low-level features.
2. Pattern completion.

50

Top-Down Processing (Imagery)

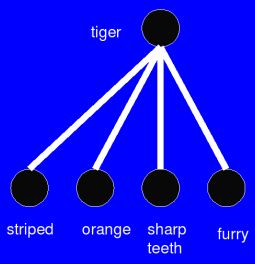

51

Top-Down Processing (Imagery)

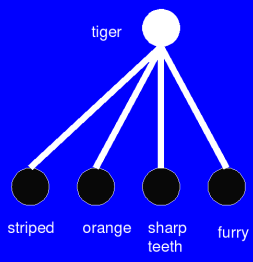

52

Top-Down Processing (Imagery)

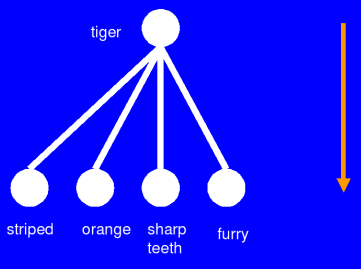

53

Pattern Completion

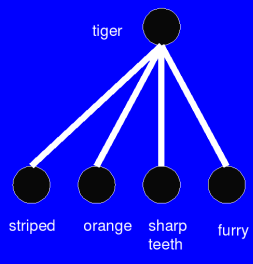

54

Pattern Completion

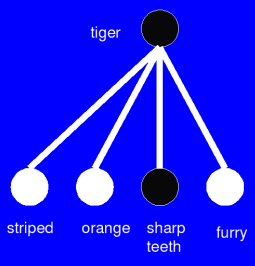

55

Pattern Completion

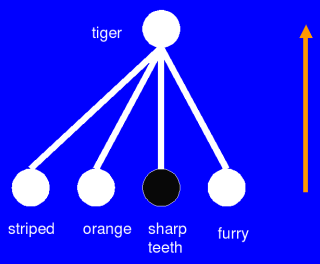

56

Pattern Completion

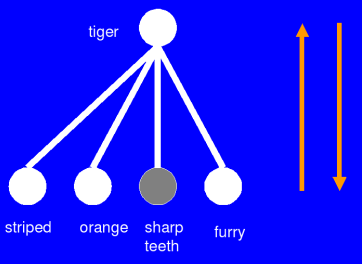

57 [pat_complete.proj]

58 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?

59 [amp_topdown.proj]

Amplification

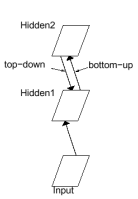

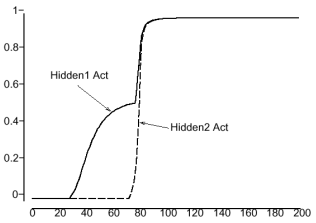

60 Application to Word Superiority Effect

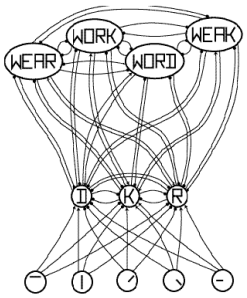

Figure 1. Interactive Activation Network Model (after McClelland and Rumelhart, 1981).

61 Networks: Bidirectional Excitation

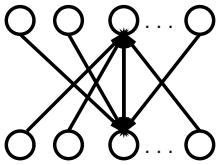

1. Top-down processing (“imagery”).
2. Pattern completion.
3. Amplification/bootstrapping.
4. Need inhibition! [amp_top_down_dist.proj]

62 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.