Semantic hierarchy

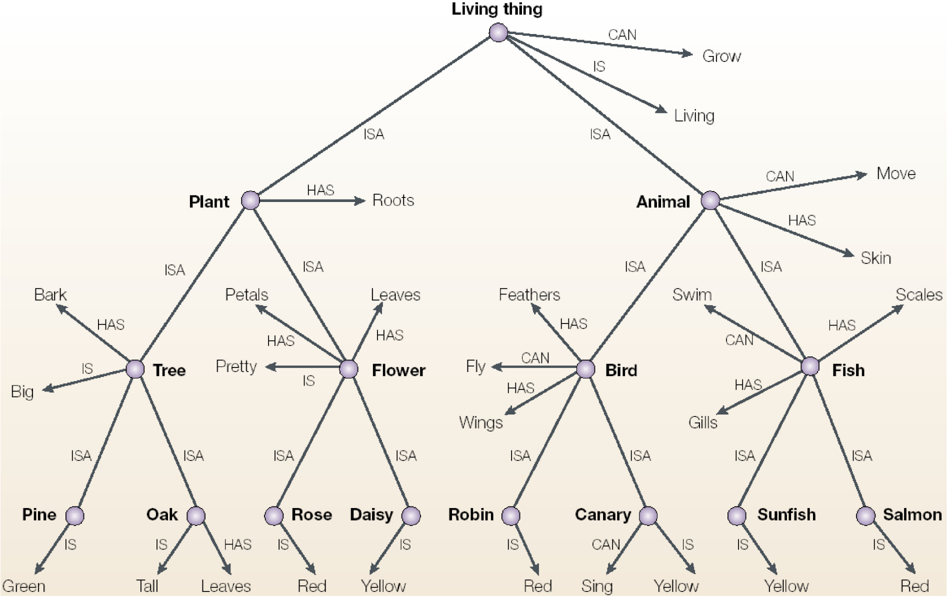

Progressive differentiation in development

(Keil, 1981, *Psych. Rev.*)

"Is it ok or silly to say that a pig is an hour long?"

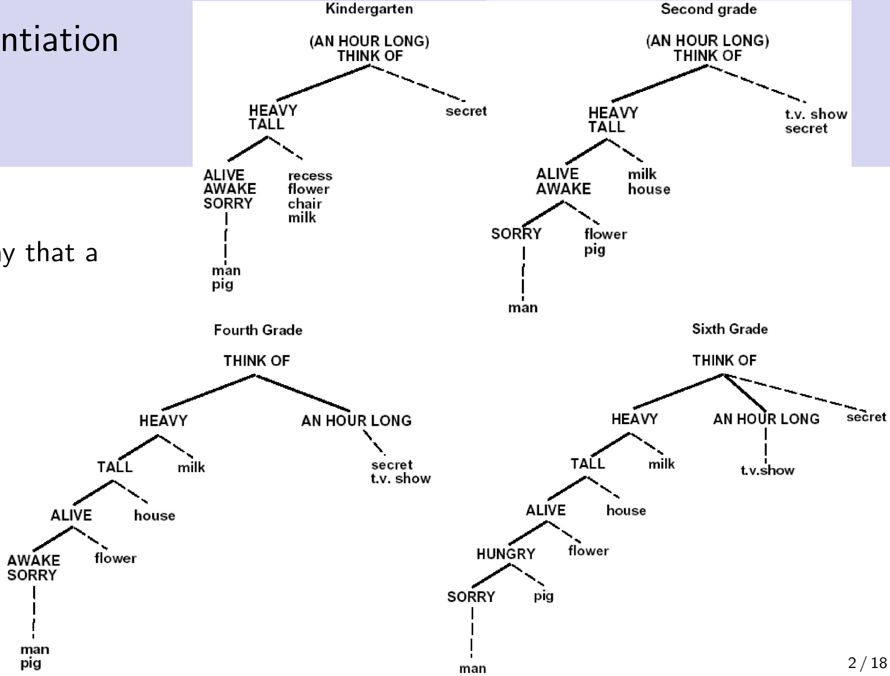

Progressive deterioration in semantic dementia

(McClelland & Rogers, 2003, *Nat. Neuro. Rev.*)

Patients with temporal variant of fronto-temporal dementia ("semantic dementia")

- more localized than Alzheimer-type dementia
- seems to specifically affect object semantics

a

Picture naming responses for JL

Item	Sept. 91	March 92	March 93
Bird	+	+	Animal
Chicken	+	+	Animal
Duck	+	Bird	Dog
Swan	+	Bird	Animal
Eagle	Duck	Bird	Horse
Ostrich	Swan	Bird	Animal
Peacock	Duck	Bird	Vehicle
Penguin	Duck	Bird	Part of animal
Rooster	Chicken	Chicken	Dog

b

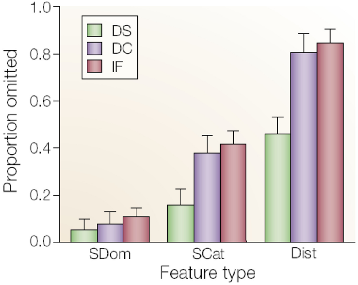

c IF's delayed copy of a camel

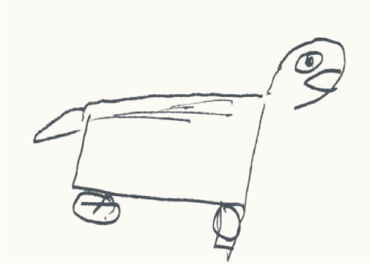

d DC's delayed copy of a swan

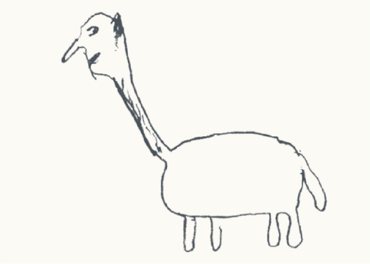

Rumelhart and Todd (1993)

- Feedforward network trained on all the “facts” from the semantic hierarchy
  - Canary CAN Grow Move Fly Sing
- Localist inputs for objects and relations
- Network learns to recode objects over **Representation** hidden layer in a way that captures the similarities among their properties
  - These representations correspond to “semantics” in the model

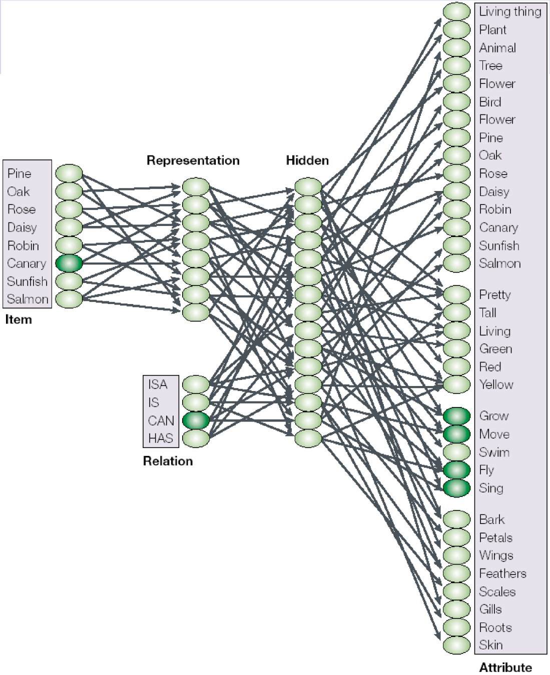

Progressive differentiation of concepts

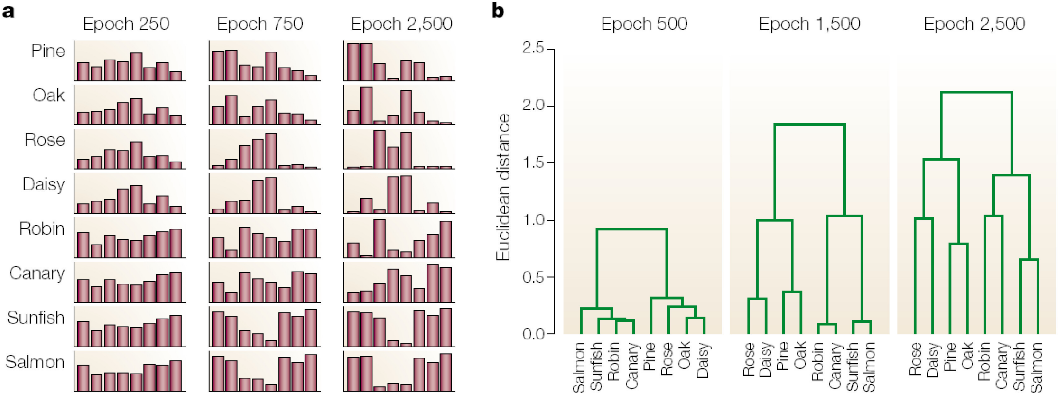

a) Activations of “representation” units progressively differentiate over training

b) Hierarchical clustering analysis of same data

Progressive differentiation of concepts

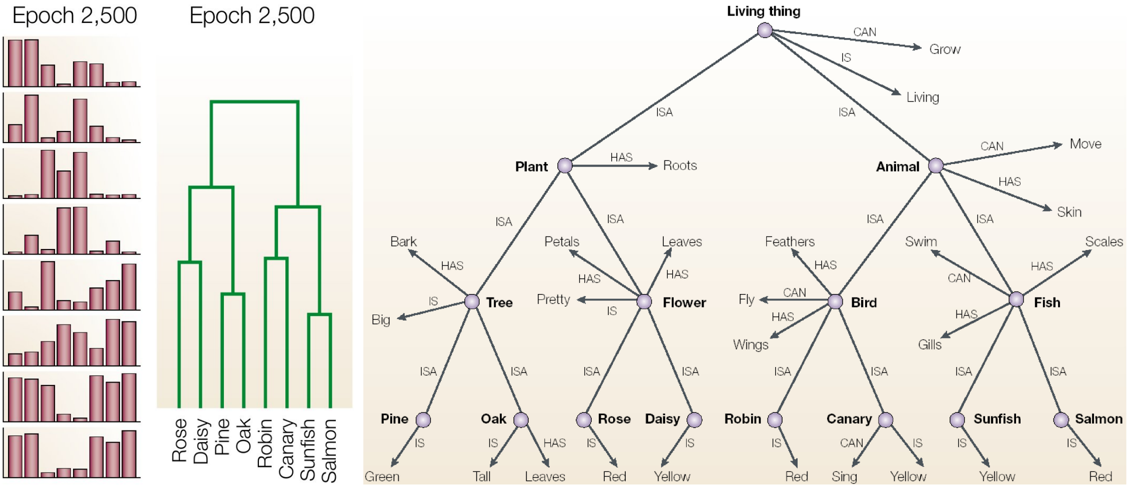

- Similarities among learned internal representations recapitulates semantic hierarchy
  - but in their *functional* relationships, not in the structure of the model

Time course of learning

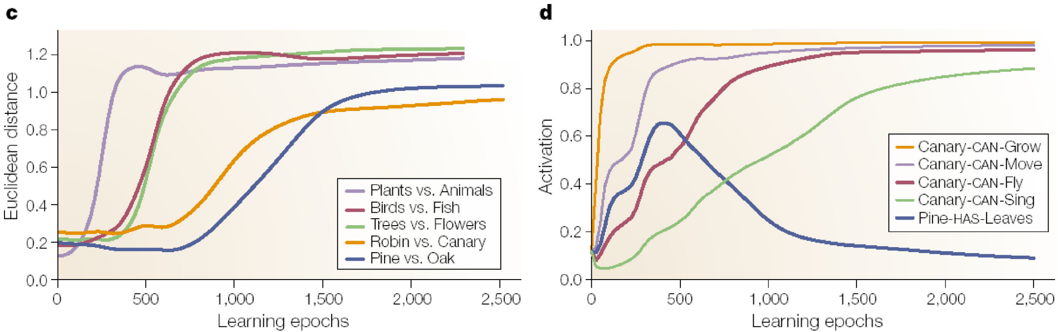

c) Differentiation from general to specific concepts

d) For a given object, general facts are learned before specific ones;

Temporary misatribution of common facts to exceptions (U-shaped performance)

- e.g., Pine HAS Leaves (which is false; analogous to over-regularizations “goed”, “taked”)

Internal representations

- Spatial organization of represenations in "state space"
  - reduced to two dimensions using Principal Components Analysis (PCA)

- Degradation by noise compromises specific before general information

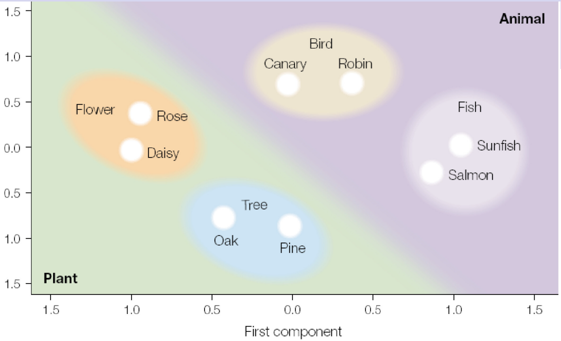

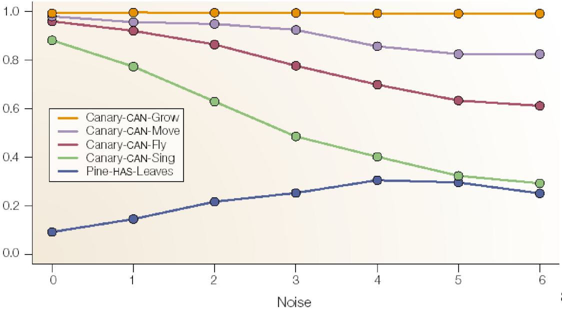

Internal representations

- Spatial organization of represenations in "state space"
  - reduced to two dimensions using Principal Components Analysis (PCA)

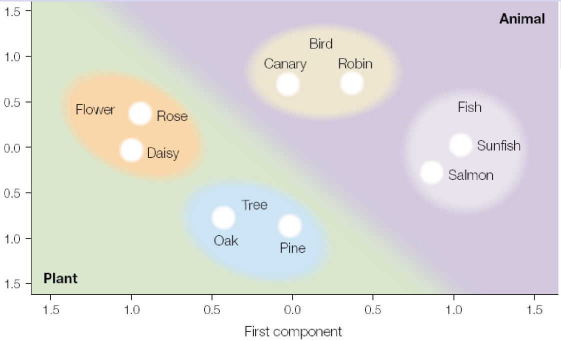

- Degradation by noise compromises specific before general information

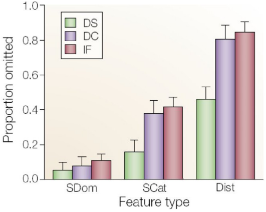

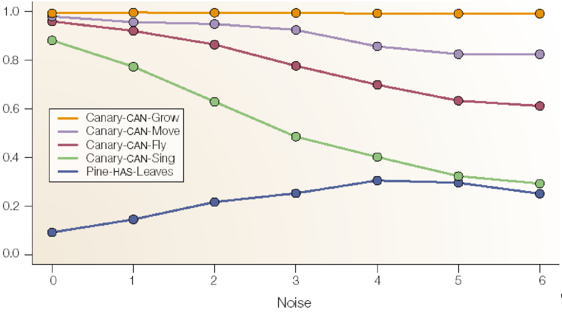

“Basic” level (Rosch et al., 1976, *Cog. Psych.*)

- Train version of network with more objects
  - Much higher frequency of **Dog**
- (b) Intermediate (“basic”) level learned before General and Specific levels
- (c) Tendency to over-attributed Dog label (Goat ISA Dog > Goat)

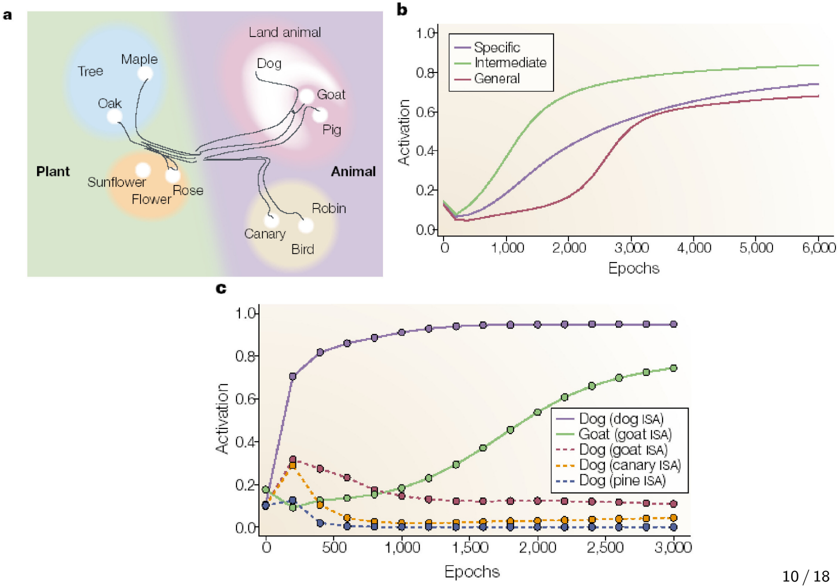

Perceptual to conceptual shift

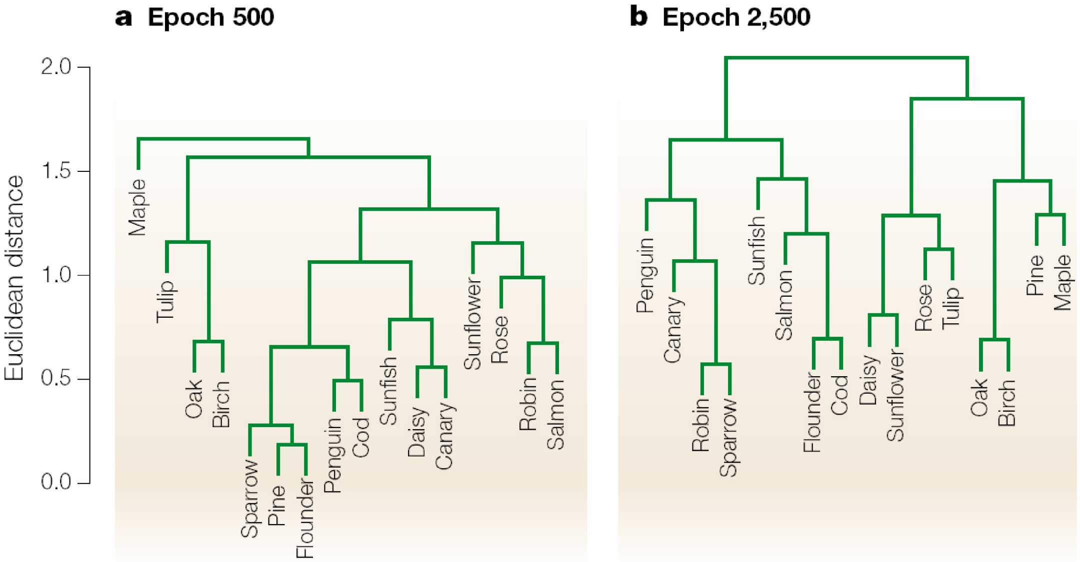

Word embeddings (Derive representations of word meanings from large text corpora)

Assumption: The meaning of a word can be inferred from how it is used—that is, from all of the "contexts" in which it occurs

- e.g., sentences, paragraphs, or even whole "documents" (e.g., Wikipedia entries)

Word embeddings (Derive representations of word meanings from large text corpora)

Assumption: The meaning of a word can be inferred from how it is used—that is, from all of the "contexts" in which it occurs

- e.g., sentences, paragraphs, or even whole "documents" (e.g., Wikipedia entries)

Latent Semantic Analysis (LSA) (Landauer & Dumais, 1997, *Psych. Rev.*)

- Count how many times any two words both occur in the same document
  - word order is irrelevant ("bag-of-words" or BOW)

Word embeddings (Derive representations of word meanings from large text corpora)

Assumption: The meaning of a word can be inferred from how it is used—that is, from all of the "contexts" in which it occurs

- e.g., sentences, paragraphs, or even whole "documents" (e.g., Wikipedia entries)

Latent Semantic Analysis (LSA) (Landauer & Dumais, 1997, *Psych. Rev.*)

- Count how many times any two words both occur in the same document
  - word order is irrelevant ("bag-of-words" or BOW)
- Words that co-occur often are associated (e.g., DOG, BONE)

Word embeddings (Derive representations of word meanings from large text corpora)

Assumption: The meaning of a word can be inferred from how it is used—that is, from all of the "contexts" in which it occurs

- e.g., sentences, paragraphs, or even whole "documents" (e.g., Wikipedia entries)

Latent Semantic Analysis (LSA) (Landauer & Dumais, 1997, *Psych. Rev.*)

- Count how many times any two words both occur in the same document
  - word order is irrelevant ("bag-of-words" or BOW)
- Words that co-occur often are associated (e.g., DOG, BONE)
- Words that have <u>similar co-occurrence statistics</u> (i.e., co-occur with the same other words, even if not with each other) are semantically related (e.g., ROAD, STREET)

Word embeddings (Derive representations of word meanings from large text corpora)

Assumption: The meaning of a word can be inferred from how it is used—that is, from all of the "contexts" in which it occurs

- e.g., sentences, paragraphs, or even whole "documents" (e.g., Wikipedia entries)

Latent Semantic Analysis (LSA) (Landauer & Dumais, 1997, *Psych. Rev.*)

- Count how many times any two words both occur in the same document
  - word order is irrelevant ("bag-of-words" or BOW)
- Words that co-occur often are associated (e.g., DOG, BONE)
- Words that have <u>similar co-occurrence statistics</u> (i.e., co-occur with the same other words, even if not with each other) are semantically related (e.g., ROAD, STREET)
- Apply dimensionality-reduction techniques (here: Singular Value Decomposition) to produce a lower-dimensional vector representation of each word

Word embeddings (Derive representations of word meanings from large text corpora)

Assumption: The meaning of a word can be inferred from how it is used—that is, from all of the "contexts" in which it occurs

- e.g., sentences, paragraphs, or even whole "documents" (e.g., Wikipedia entries)

Latent Semantic Analysis (LSA) (Landauer & Dumais, 1997, *Psych. Rev.*)

- Count how many times any two words both occur in the same document
  - word order is irrelevant ("bag-of-words" or BOW)
- Words that co-occur often are associated (e.g., DOG, BONE)
- Words that have <u>similar co-occurrence statistics</u> (i.e., co-occur with the same other words, even if not with each other) are semantically related (e.g., ROAD, STREET)
- Apply dimensionality-reduction techniques (here: Singular Value Decomposition) to produce a lower-dimensional vector representation of each word

Hyperspace Analogue to Language (HAL) (Burgess & Lund, 1997, *Lang. & Cog. Proc.*)

- Restrict co-occurence statistics to neighboring words (e.g., 10 on either side, weighted by distance)
  - also used different dimensionality reduction

word2vec (Mikolov et al., 2013)

Train feedforward networks to

a) predict word from surrounding words   b) predict surrounding words from the word itself

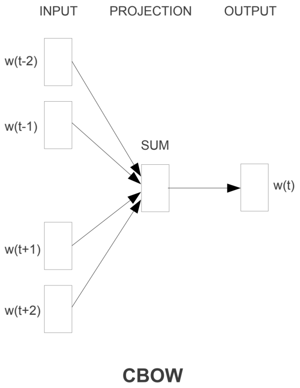

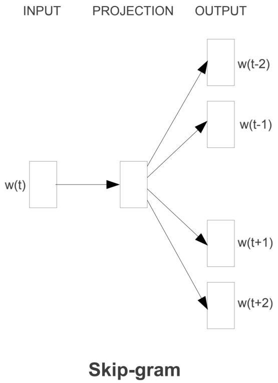

Can't be completely correct, but network will develop internal (hidden) representations (word embeddings) that minimize error