Word Meaning

Informatics 1 CG: Lecture 23

Mirella Lapata

School of Informatics
University of Edinburgh
mlap@inf.ed.ac.uk

March 13, 2014

Reading:

Trevor Harley's The Psychology of Language, Chapter 10

Recap: Categorization

Categorization is one of the classical problems in the field of cognitive science, one with a histiry dating back to Aristotle.

- Ability to generalize from experience underlies a variety of common mental tasks
- Perception, learning, and the use of language.
- Definitional, prototype, exemplar, and theories theory.
- Basic-level categories, prototype, family resemblance.

How do we Represent the Meaning of Words?

Semantic knowledge can be thought of as knowledge about relations among several types of elements, including words, concepts, and percepts.

Word-concept relations

Knowledge that the word *dog* refers to the concept DOG the word *animal* refers to the concept ANIMAL or the word *toaster* refers to the concept TOASTER.

Word-concept relations

Knowledge that the word *dog* refers to the concept DOG the word *animal* refers to the concept ANIMAL or the word *toaster* refers to the concept TOASTER.

dog

Word-concept relations

Knowledge that the word *dog* refers to the concept DOG the word *animal* refers to the concept ANIMAL or the word *toaster* refers to the concept TOASTER.

animal

Word-concept relations

Knowledge that the word *dog* refers to the concept DOG the word *animal* refers to the concept ANIMAL or the word *toaster* refers to the concept TOASTER.

toaster

Concept-concept relations

Knowledge that *dogs* are a kind of animal, that *dogs* have tails and can bark, or that *animals* have bodies and can move.

animal
|
IS-A
|
vertebrate
|
IS-A
|
mammal
|
IS-A
|
canine
|
dog

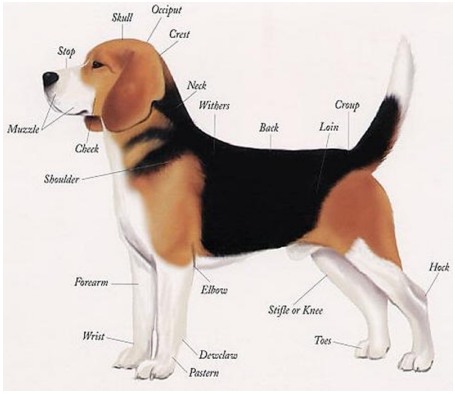

Concept-percept, Concept-action Relations

Knowledge about what dogs look like, how a dog can be distinguished from a cat, or how to pet a dog or operate a toaster.

Word-word relations

Knowledge that the word *dog* tends to be associated with or co-occur with words such as *tail*, *bone*, and *cat* or that the word *toaster* tends to be associated with *kitchen*, *oven*, or *bread*.

Word-word relations

Knowledge that the word *dog* tends to be associated with or co-occur with words such as *tail*, *bone*, and *cat* or that the word *toaster* tends to be associated with *kitchen*, *oven*, or *bread*.

What are the associates of *apple*?

Word-word relations

Knowledge that the word *dog* tends to be associated with or co-occur with words such as *tail*, *bone*, and *cat* or that the word *toaster* tends to be associated with *kitchen*, *oven*, or *bread*.

What are the associates of *apple*?

red

orange

pie

worm

Semantic Networks

- Emphasizes abstract conceptual structure, focusing on relations among concepts and relations between concepts and percepts or actions.
- This knowledge is represented in terms of systems of abstract propositions, such as *canary is-a bird*, *canary has bird wings*.
- Concepts represented in network of interconnecting nodes
- Distance between nodes represents similarity between them.
- Concept defined in terms of the connections with other concepts

Collins and Quillian (1969)

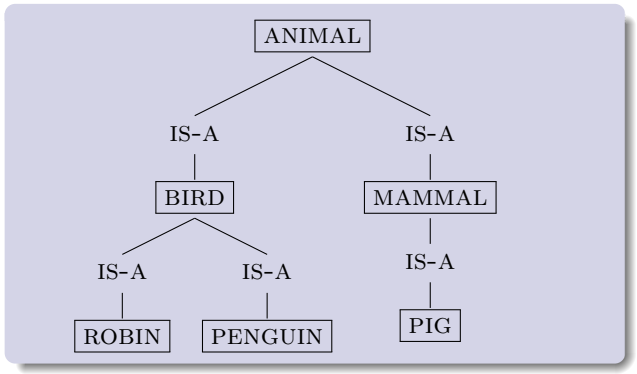

- Useful for representing natural kind terms.
- Economical method for storing information.
- Most common links are IS-A links
- Attributes stored at lowest possible node at which they are true of all lower nodes in network

Collins and Quillian (1969)

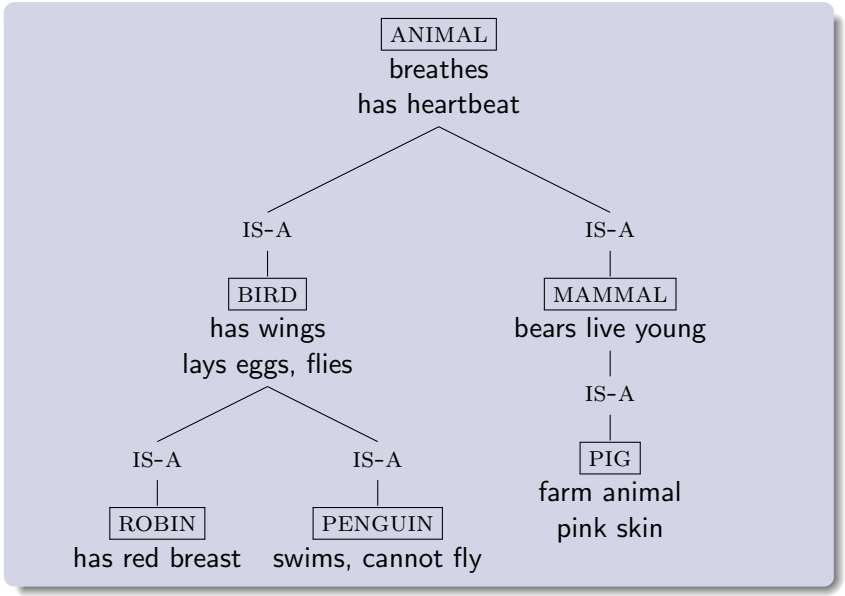

Sentence Verification Task

Participants are presented with simple "facts" and have to press one button if the sentence is true, another if it is false. The reaction time is an index of how difficult the decision was.

(1) A robin is a robin.
(2) A robin is a bird.
(3) A robin is an animal.
(4) A robin is a fish.

- Response time to (1) < (2) < (3) < (4).
- Participants start off from robin and travel through the network until they find the necessary infromation.
- The farther away the information the slower the response time.

Sentence Verification Task

Participants are presented with simple "facts" and have to press one button if the sentence is true, another if it is false. The reaction time is an index of how difficult the decision was.

(1) A robin is a robin. (baseline measure) Yes!
(2) A robin is a bird.
(3) A robin is an animal.
(4) A robin is a fish.

- Response time to (1) < (2) < (3) < (4).
- Participants start off from robin and travel through the network until they find the necessary infromation.
- The farther away the information the slower the response time.

Informatics 1 CG: Lecture 23 Word Meaning 12

Sentence Verification Task

Participants are presented with simple "facts" and have to press one button if the sentence is true, another if it is false. The reaction time is an index of how difficult the decision was.

(1) A robin is a robin. (baseline measure) Yes!
(2) A robin is a bird. Yes!
(3) A robin is an animal.
(4) A robin is a fish.

- Response time to (1) < (2) < (3) < (4).
- Participants start off from robin and travel through the network until they find the necessary infromation.
- The farther away the information the slower the response time.

Informatics 1 CG: Lecture 23 Word Meaning 12

Sentence Verification Task

Participants are presented with simple "facts" and have to press one button if the sentence is true, another if it is false. The reaction time is an index of how difficult the decision was.

(1) A robin is a robin. (baseline measure) Yes!
(2) A robin is a bird. Yes!
(3) A robin is an animal. Yes!
(4) A robin is a fish.

- Response time to (1) < (2) < (3) < (4).
- Participants start off from robin and travel through the network until they find the necessary infromation.
- The farther away the information the slower the response time.

Informatics 1 CG: Lecture 23 Word Meaning 12

Sentence Verification Task

Participants are presented with simple "facts" and have to press one button if the sentence is true, another if it is false. The reaction time is an index of how difficult the decision was.

(1)	A robin is a robin.	(baseline measure) Yes!
(2)	A robin is a bird.	Yes!
(3)	A robin is an animal.	Yes!
(4)	A robin is a fish.	No!

- Response time to (1) < (2) < (3) < (4).
- Participants start off from robin and travel through the network until they find the necessary infromation.
- The farther away the information the slower the response time.

Sentence Verification Task

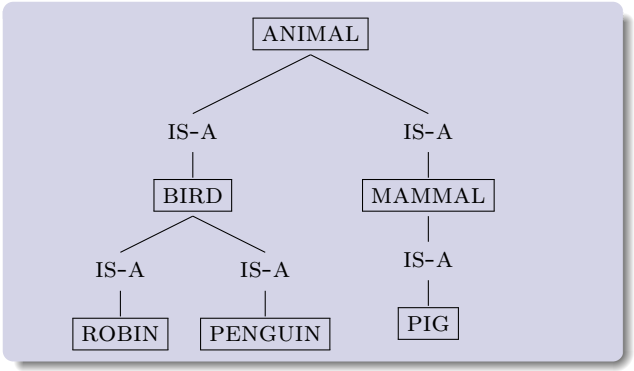

http://www.penscreen.com/Sentence%20Verification%20Test.htm

Sentence Verification Task

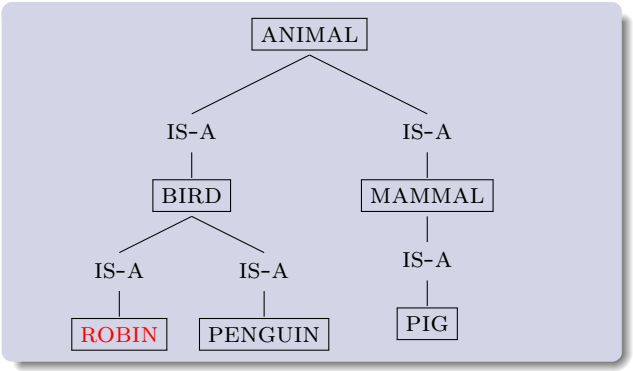

A robin is a robin.

http://www.penscreen.com/Sentence%20Verification%20Test.htm

Sentence Verification Task

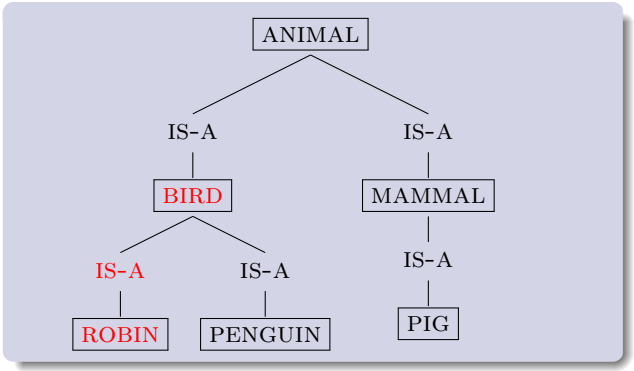

A robin is a bird.

http://www.penscreen.com/Sentence%20Verification%20Test.htm

Sentence Verification Task

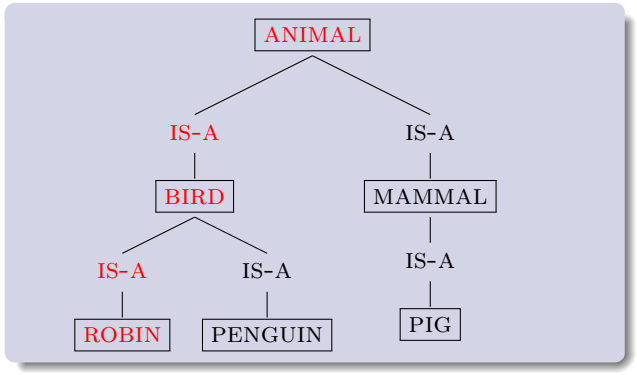

A robin is an animal.

http://www.penscreen.com/Sentence%20Verification%20Test.htm

Sentence Verification Task

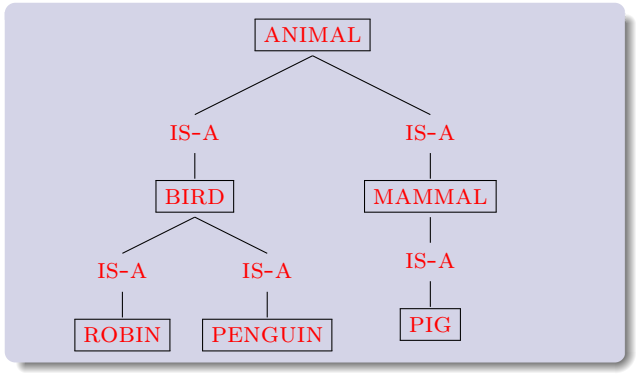

A robin is a fish.

http://www.penscreen.com/Sentence%20Verification%20Test.htm

Problems with the Collins and Quillian Model

- Not all information is easily represented in hierarchical form (what is the relation between *truth*, *justice*, and *law*?)
- Experiments confound distance in network with conjoint frequency (*robin* and *bird* often co-occur).
- The model makes some incorrect predictions: (5) < (6), (7) < (8), (9) < (10).

(5) A cow is an animal.
(6) A cow is a mammal.
(7) A pine is a church.
(8) A pine is a flower.
(9) A robin is a bird.
(10) A penguin is a bird.

Problems with the Collins and Quillian Model

- Not all information is easily represented in hierarchical form (what is the relation between *truth*, *justice*, and *law*?)
- Experiments confound distance in network with conjoint frequency (*robin* and *bird* often co-occur).
- The model makes some incorrect predictions: (5) < (6), (7) < (8), (9) < (10).

(5) A cow is an animal.
(6) A cow is a mammal.
(7) A pine is a church.
(8) A pine is a flower.
(9) A robin is a bird.
(10) A penguin is a bird.

Problems with the Collins and Quillian Model

- Not all information is easily represented in hierarchical form (what is the relation between *truth*, *justice*, and *law*?)
- Experiments confound distance in network with conjoint frequency (*robin* and *bird* often co-occur).
- The model makes some incorrect predictions: (5) < (6), (7) < (8), (9) < (10).

(5) A cow is an animal.
(6) A cow is a mammal.
(7) A pine is a church.
(8) A pine is a flower.
(9) A robin is a bird.
(10) A penguin is a bird.

Problems with the Collins and Quillian Model

- Not all information is easily represented in hierarchical form (what is the relation between *truth*, *justice*, and *law*?)
- Experiments confound distance in network with conjoint frequency (*robin* and *bird* often co-occur).
- The model makes some incorrect predictions: (5) < (6), (7) < (8), (9) < (10).

(5) A cow is an animal.
(6) A cow is a mammal.
(7) A pine is a church.
(8) A pine is a flower.
(9) A robin is a bird.
(10) A penguin is a bird.

Informatics 1 CG: Lecture 23 Word Meaning 14

Collins and Loftus (1975)

- Model is based on idea of spreading activation.
- More complex network structure with links varying in strength or distance; structure is no longer hierarchical.
- Connections represent: categorical relations, degree of association, typicality.
- When you think about a concept, that concept will become activated, and that activation will spread to other concepts that are linked to it.
- Verification times depend on closeness concepts in network.
- Hard to see what sort of experiments could falsify this model.

Collins and Loftus (1975)

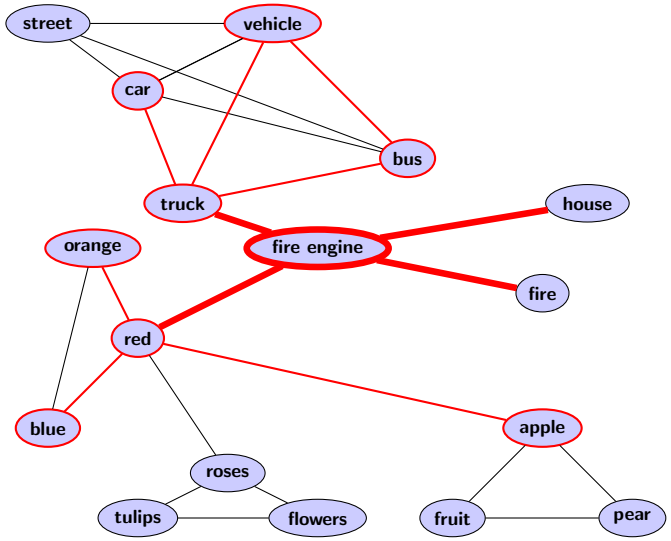

Feature Comparison Model

	animate	feathered	has_a_beak	flies	sings
BIRD	+	+	+	+	+
ROBIN	+	+	+	+	+
PENGUIN	+	+	+	−	−
PIG	+	−	−	−	−

- Concepts are represented as a set of features.
- Features are ordered in terms of definingness.[1]
- Defining features: essential to meaning of word and relate to preperties that things must have to be members of category.
- Characteristic features: are usually but not necessarily true (most birds can fly but penguins and ostriches cannot).

[1] Distinction between defining and characteristic features is arbitrary.

Feature Comparison Model

defining ⟶ characteristic

	animate	feathered	has_a_beak	flies	sings
BIRD	+	+	+	+	+
ROBIN	+	+	+	+	+
PENGUIN	+	+	+	−	−
PIG	+	−	−	−	−

- Concepts are represented as a set of features.
- Features are ordered in terms of definingness.[1]
- Defining features: essential to meaning of word and relate to preperties that things must have to be members of category.
- Characteristic features: are usually but not necessarily true (most birds can fly but penguins and ostriches cannot).

[1] Distinction between defining and characteristic features is arbitrary.

Feature Comparison Model

One way to test defining and characteristic features is to carry out a similarity test (Smith and colleagues, 1978).

Question	Feature	Decision	RT
A cat is a pencil (low similarity)	defining	false	fast
A cat is a mammal (high similarity)	defining	true	fast
A cat is fluffy (intermediate similarity)	characteristic	true or false	slow

Sentence Verification (Again)

A robin is a bird.

Sentence Verification (Again)

A robin is a bird.

Stage 1
Compare all features of robin and bird to dermine featural similarity.

Sentence Verification (Again)

A robin is a bird.

Stage 1
Compare all features of robin and bird to dermine featural similarity.

High Overlap

Fast yes

Sentence Verification (Again)

A robin is a pig.

Stage 1
Compare all features of robin and pig to dermine featural similarity.

Low Overlap

fast no

Sentence Verification (Again)

A penguin is a bird.

Stage 1
Compare all features of penguin and bird to dermine featural similarity.

Sentence Verification (Again)

A penguin is a bird.

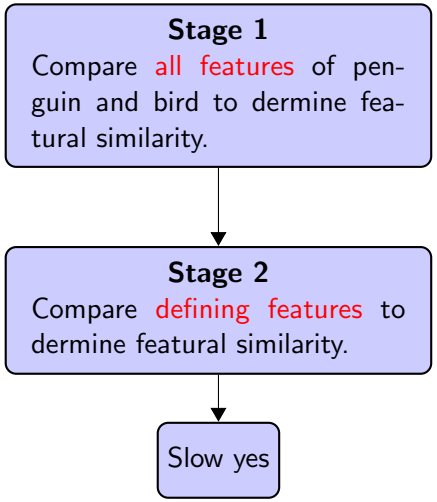

Problems with Feature Comparison Model

- Many words do not have obvious defining features!
- Model is tied to sentence verification paradigm.
- Probabilistic feature model (Smith and Medin, 1981).
- Distinguishes between essential defining features of concepts and aspects of meaning for identifying instances of concept.
- Features are weighted based on salience and probability of being true for category (has four limbs vs bears live young)
- Candidate instance is accepted if exceeds some critical weighted sum of features.
- Categories now have fuzzy boundaries.
- How is this model different from prototype model?

Where Do the Features Come from?

- Participants are presented with set of concept names
- Asked to write down up to n features they think are important for each concept
- McRae et al (2005) collected feature norms for 541 living and nonliving concepts
- Largest set in existence (2,526 features), collected over several years
- Reveal psychologically salient dimensions of meaning

Where Do the Features Come from?

Where Do the Features Come from?

MOOSE/ELK

Where Do the Features Come from?

MOOSE/ELK

Feature	Freq	Classification
is_large	27	visual
has_antlers	23	visual
has_legs	14	visual
is_brown	10	visual
has_fur	7	visual
has_hooves	5	visual
eaten_as_meat	5	function
lives_in_woods	14	encyclopedic
an_animal	17	taxonomic
a_mammal	9	taxonomic

Feature norms from McRae et al. (2005).

Representing Word Meaning

	eats_seads	has_beak	has_claws	has_handlebar	has_wheels	has_wings	has_feathers	made_of_metal	made_of_wood
TROLLEY	.00	.00	.00	.30	.32	.00	.00	.06	.25
ROBIN	.05	.24	.15	.00	.00	.19	.34	.00	.00

- McRae spent 10 years collecting his feature norms! (541 basic-level nouns).
- What about verbs or abstract concepts (e.g., *move*, *peace*)?
- But, humans naturally express word meaning using features.

Representing Word Meaning

	eats_seads	has_beak	has_claws	has_handlebar	has_wheels	has_wings	has_feathers	made_of_metal	made_of_wood
TROLLEY	.00	.00	.00	.30	.32	.00	.00	.06	.25
ROBIN	.05	.24	.15	.00	.00	.19	.34	.00	.00

- McRae spent 10 years collecting his feature norms! (541 basic-level nouns).
- What about verbs or abstract concepts (e.g., *move*, *peace*)?
- But, humans naturally express word meaning using features.

Representing Word Meaning

	eats_seads	has_beak	has_claws	has_handlebar	has_wheels	has_wings	has_feathers	made_of_metal	made_of_wood
TROLLEY	.00	.00	.00	.30	.32	.00	.00	.06	.25
ROBIN	.05	.24	.15	.00	.00	.19	.34	.00	.00

- McRae spent 10 years collecting his feature norms! (541 basic-level nouns).
- What about verbs or abstract concepts (e.g., *move, peace*)?
- But, humans naturally express word meaning using features.

Summary

How do we represent the meaning of words? How is semantic knowledge organized?

- Semantic information in encoded in networks of linked nodes.
- Collins and Quillian network emphasizes hierarchical relations and cognitive economy; sentence verification times.
- Does not explain similarity and relatedness effects.
- Spreading activation model does but is difficult to falsify.
- Word meaning can be decomposed into semantic features.
- Feature-list theories account for sentence verification times by postulating that we compare lists of defining and characteristic features.

Next lecture: associationist view of meaning, vector space model.