Outline of today's lecture

Lecture 7: Lexical semantics

Lexical semantics: semantic relations

Polysemy

Word sense disambiguation

Lexical semantics

- Limited domain: mapping to some knowledge base term(s). Knowledge base constrains possible meanings.
- Issues for broad coverage systems:
  - Boundary between lexical meaning and world knowledge.
  - Representing lexical meaning.
  - Acquiring representations.
  - Polysemy and multiword expressions.

Approaches to lexical meaning

- Formal semantics: extension — what words denote (e.g., cat′: the set of all cats). But ...
- Semantic primitives: e.g., *kill* means CAUSE (NOT (ALIVE)). But ...
- Meaning postulates:

$$\forall e, x, y[\text{kill}'(e, x, y) \rightarrow \exists e'[\text{cause}'(e, x, e') \wedge \text{die}'(e', y)]]$$

  But ...
- Ontological relationships: informal or formal (description logics): this lecture (informal approaches).
- Distributional approaches (lecture 8).

Approaches to lexical meaning

- Formal semantics: extension — what words denote (e.g., cat′: the set of all cats). But ...
- Semantic primitives: e.g., *kill* means CAUSE (NOT (ALIVE)). But ...
- Meaning postulates:

$$\forall e, x, y[\text{kill}'(e, x, y) \rightarrow \exists e'[\text{cause}'(e, x, e') \wedge \text{die}'(e', y)]]$$

  But ...
- Ontological relationships: informal or formal (description logics): this lecture (informal approaches).
- Distributional approaches (lecture 8).

Approaches to lexical meaning

- Formal semantics: extension — what words denote (e.g., cat′: the set of all cats). But . . .
- Semantic primitives: e.g., *kill* means CAUSE (NOT (ALIVE)). But . . .
- Meaning postulates:

$$\forall e, x, y[\text{kill}'(e, x, y) \rightarrow \exists e'[\text{cause}'(e, x, e') \wedge \text{die}'(e', y)]]$$

  But . . .
- Ontological relationships: informal or formal (description logics): this lecture (informal approaches).
- Distributional approaches (lecture 8).

Approaches to lexical meaning

- Formal semantics: extension — what words denote (e.g., cat′: the set of all cats). But . . .
- Semantic primitives: e.g., *kill* means CAUSE (NOT (ALIVE)). But . . .
- Meaning postulates:

$$\forall e, x, y[\text{kill}'(e, x, y) \rightarrow \exists e'[\text{cause}'(e, x, e') \wedge \text{die}'(e', y)]]$$

  But . . .
- Ontological relationships: informal or formal (description logics): this lecture (informal approaches).
- Distributional approaches (lecture 8).

Approaches to lexical meaning

- Formal semantics: extension — what words denote (e.g., cat′: the set of all cats). But . . .
- Semantic primitives: e.g., *kill* means CAUSE (NOT (ALIVE)). But . . .
- Meaning postulates:

$$\forall e, x, y[\text{kill}'(e, x, y) \rightarrow \exists e'[\text{cause}'(e, x, e') \wedge \text{die}'(e', y)]]$$

But . . .

- Ontological relationships: informal or formal (description logics): this lecture (informal approaches).
- Distributional approaches (lecture 8).

Approaches to lexical meaning

- Formal semantics: extension — what words denote (e.g., cat′: the set of all cats). But . . .
- Semantic primitives: e.g., *kill* means CAUSE (NOT (ALIVE)). But . . .
- Meaning postulates:

$$\forall e, x, y[\text{kill}'(e, x, y) \rightarrow \exists e'[\text{cause}'(e, x, e') \wedge \text{die}'(e', y)]]$$

  But . . .
- Ontological relationships: informal or formal (description logics): this lecture (informal approaches).
- Distributional approaches (lecture 8).

Approaches to lexical meaning

- Formal semantics: extension — what words denote (e.g., cat′: the set of all cats). But ...
- Semantic primitives: e.g., *kill* means CAUSE (NOT (ALIVE)). But ...
- Meaning postulates:

$$\forall e, x, y[\text{kill}'(e, x, y) \rightarrow \exists e'[\text{cause}'(e, x, e') \wedge \text{die}'(e', y)]]$$

  But ...
- Ontological relationships: informal or formal (description logics): this lecture (informal approaches).
- Distributional approaches (lecture 8).

Examples to think about

- tomato
- table
- thought
- democracy
- push
- sticky

Hyponymy: IS-A

- (a sense of) *dog* is a hyponym of (a sense of) *animal*
- *animal* is a hypernym of *dog*
- hyponymy relationships form a taxonomy
- works best for concrete nouns

Some issues concerning hyponymy

- not useful for all words: *thought*, *democracy*, *push*, *sticky*?
- individuation differences: is *table* a hyponym of *furniture*?
- multiple inheritance: e.g., is *coin* a hyponym of both *metal* and *money*?
- what does the top of the hierarchy look like?

Other semantic relations

Classical relations:

Meronomy: PART-OF e.g., *arm* is a meronym of *body*, *steering wheel* is a meronym of *car* (piece vs part)

Synonymy e.g., *aubergine/eggplant*.

Antonymy e.g., *big/little*

Also:

Near-synonymy/similarity e.g., *exciting/thrilling*
e.g., *slim/slender/thin/skinny*

WordNet

- large scale, open source resource for English
- hand-constructed
- wordnets being built for other languages
- organized into synsets: synonym sets (near-synonyms)

Overview of adj red:

```
1. (43) red, reddish, ruddy, blood-red, carmine,
cerise, cherry, cherry-red, crimson, ruby,
ruby-red, scarlet - (having any of numerous
bright or strong colors reminiscent of the color
of blood or cherries or tomatoes or rubies)
2. (8) red, reddish - ((used of hair or fur)
of a reddish brown color; "red deer";
reddish hair")
```

Hyponymy in WordNet

```
Sense 6
big cat, cat
       => leopard, Panthera pardus
           => leopardess
           => panther
       => snow leopard, ounce, Panthera uncia
       => jaguar, panther, Panthera onca,
                                   Felis onca
       => lion, king of beasts, Panthera leo
           => lioness
           => lionet
       => tiger, Panthera tigris
           => Bengal tiger
           => tigress
```

Polysemy

- homonymy: unrelated word senses. *bank* (raised land) vs *bank* (financial institution)
- *bank* (financial institution) vs *bank* (in a casino): related but distinct senses.
- *bank* (N) (raised land) vs *bank* (V) (to create some raised land): regular polysemy. Compare *pile*, *heap* etc
- vagueness: *bank* (river vs snow vs cloud)?

No clearcut distinctions.
Dictionaries are not consistent.

Word sense disambiguation

Needed for many applications, problematic for large domains. Assumes that we have a standard set of word senses (e.g., WordNet)

- frequency: e.g., *diet*: the food sense (or senses) is much more frequent than the parliament sense (Diet of Wurms)
- collocations: e.g. *striped bass* (the fish) vs *bass guitar*: syntactically related or in a window of words (latter sometimes called ‘cooccurrence’). Generally ‘one sense per collocation’.
- selectional restrictions/preferences (e.g., *Kim eats bass*, must refer to fish

WSD techniques

- supervised learning: cf. POS tagging from lecture 3. But sense-tagged corpora are difficult to construct, algorithms need far more data than POS tagging
- unsupervised learning (see below)
- Machine readable dictionaries (MRDs): e.g., look at overlap with words in definitions and example sentences
- selectional preferences: don't work very well by themselves, useful in combination with other techniques

WSD by (almost) unsupervised learning

Disambiguating *plant* (factory vs vegetation senses):
1. Find contexts in training corpus:

sense	training example
?	company said that the *plant* is still operating
?	although thousands of *plant* and animal species
?	zonal distribution of *plant* life
?	company manufacturing *plant* is in Orlando
	etc

Yarowsky (1995): schematically

Initial state

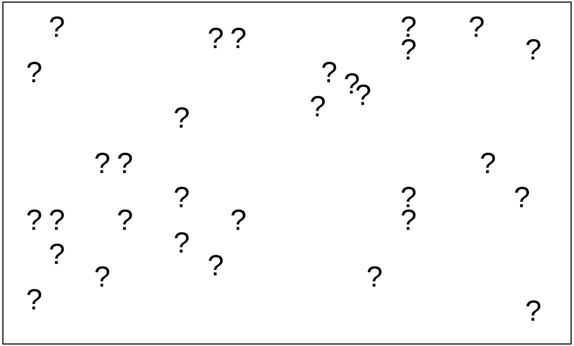

2. Identify some seeds to disambiguate a few uses. e.g., ‘plant life’ for vegetation use (A) ‘manufacturing plant’ for factory use (B):

sense	training example
?	company said that the *plant* is still operating
?	although thousands of *plant* and animal species
A	zonal distribution of *plant* life
B	company manufacturing *plant* is in Orlando
	etc

Seeds

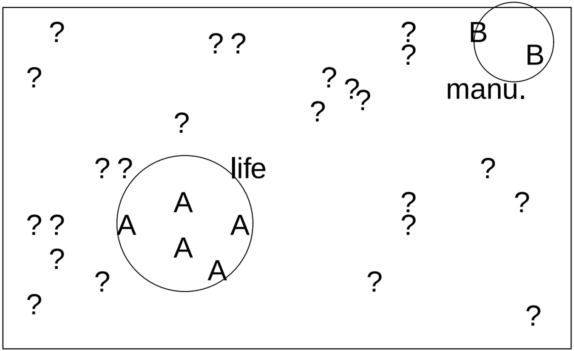

3. Train a decision list classifier on the Sense A/Sense B examples.

reliability	criterion	sense
8.10	*plant* life	A
7.58	manufacturing *plant*	B
6.27	*animal* within 10 words of *plant*	A
	etc	

Decision list classifier: automatically trained if/then statements. Experimenter decides on classes of test by providing definitions of features of interest: system builds specific tests and provides reliability metrics.

4. Apply the classifier to the training set and add reliable examples to A and B sets.

sense	training example
?	company said that the *plant* is still operating
A	although thousands of *plant* and animal species
A	zonal distribution of *plant* life
B	company manufacturing *plant* is in Orlando
	etc

5. Iterate the previous steps 3 and 4 until convergence

Iterating:

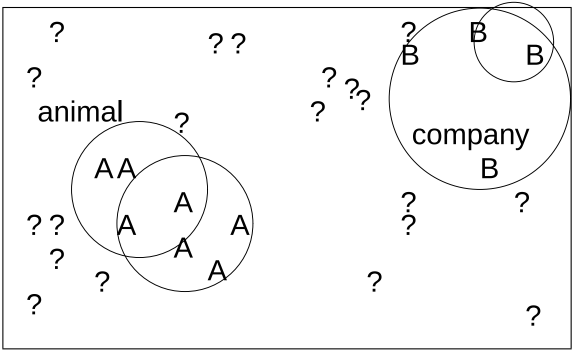

Final:

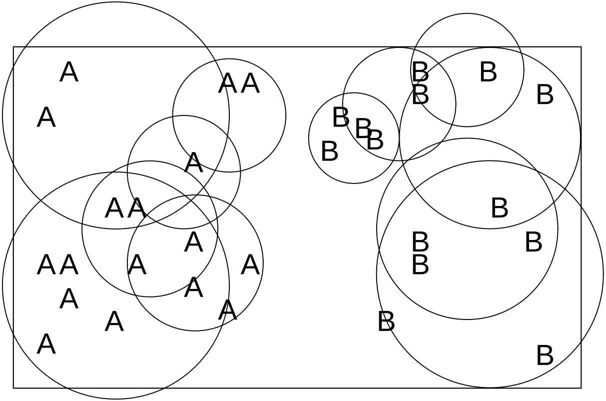

6. Apply the classifier to the unseen test data

‘one sense per discourse’: can be used as an additional refinement
e.g., once you’ve disambiguated *plant* one way in a particular text/section of text, you can assign all the instances of *plant* to that sense

Evaluation of WSD

- SENSEVAL competitions
- evaluate against WordNet
- baseline: pick most frequent sense — hard to beat (but don't always know most frequent sense)
- human ceiling varies with words
- MT task: more objective but sometimes doesn't correspond to polysemy in source language