Chapter NLP:VI

VI. Semantics

- Intro
- Lexical and Word-level semantics
- Compositional semantics
- Distributional semantics

Semantics

What is semantics?

- The meaning of single words and compositions of words.

“The man sighed.
It’s raining cats and dogs, he felt.”

Meaning

What is meaning?

- Propositional content in terms of validity or truth conditions.

  Inference How to tell if one statement/sentence follows from another? Can this be automatically computed?

"All men are mortal."	$\forall x : man(x) \rightarrow mortal(x)$
"Socrates is a man."	$man(Socrates)$
Socrates is mortal.	$mortal(Socrates)$

Meaning

What is meaning?

- Propositional content in terms of validity or truth conditions.

  Inference How to tell if one statement/sentence follows from another? Can this be automatically computed?

"All men are mortal."	$\forall x : man(x) \rightarrow mortal(x)$
"Socrates is a man."	$man(Socrates)$
Socrates is mortal.	$mortal(Socrates)$

- Often requires common-sense reasoning based on world knowledge.

  "Max can open Tim's safe.
  He knows the combination."

  "Max can open Tim's safe.
  He should change the combination."

Meaning

What is meaning?

- Propositional content in terms of validity or truth conditions.

  Inference How to tell if one statement/sentence follows from another? Can this be automatically computed?

"All men are mortal."	$\forall x : man(x) \rightarrow mortal(x)$
"Socrates is a man."	$man(Socrates)$
Socrates is mortal.	$mortal(Socrates)$

- Often requires common-sense reasoning based on world knowledge.

  "Max can open Tim's safe. He knows the combination."

  "Max can open Tim's safe. He should change the combination."

- Includes expressed emotional content.

  "That poor cat!"

  "Fortunately, Max can open Tim's safe."

Meaning

What is meaning?

- Propositional content in terms of validity or truth conditions.

  Inference How to tell if one statement/sentence follows from another? Can this be automatically computed?

"All men are mortal."	$\forall x : man(x) \rightarrow mortal(x)$
"Socrates is a man."	$man(Socrates)$
Socrates is mortal.	$mortal(Socrates)$

- Often requires common-sense reasoning based on world knowledge.

  "Max can open Tim's safe. He knows the combination."

  "Max can open Tim's safe. He should change the combination."

- Includes expressed emotional content.

  "That poor cat!"

  "Fortunately, Max can open Tim's safe."

Construction of meaning

- Linguistic form vs. context of use **or** Lexical vs. Compositional

Meaning

Linguistic Form

Meaning that can often be derived from linguistic form

- Constant meaning of language across different contexts of use.

Max: “Linda, it is raining.”

(the meaning of “raining” is clear without context)

Meaning

Linguistic Form

Meaning that can often be derived from linguistic form

- Constant meaning of language across different contexts of use.

  Max: "Linda, it is raining."

  (the meaning of "raining" is clear without context)

- Meaning a speaker publicly commits to by using a certain form.

  Linda: "It's perfectly dry outside. You're mocking me?"

  (since Max said it's raining, he committed to the outside being wet)

Meaning

Linguistic Form

Meaning that can often be derived from linguistic form

- Constant meaning of language across different contexts of use.

  Max: "Linda, it is raining."

  (the meaning of "raining" is clear without context)

- Meaning a speaker publicly commits to by using a certain form.

  Linda: "It's perfectly dry outside. You're mocking me?"

  (since Max said it's raining, he committed to the outside being wet)

- Inferences about a speaker's private cognitive states.

  Linda: "Or do you want me to take an umbrella?" Max: "No."

  (there must be a reason why Max said that)

Meaning

Linguistic Form

Meaning that can often be derived from linguistic form

- Constant meaning of language across different contexts of use.

  Max: "Linda, it is raining."

  (the meaning of "raining" is clear without context)

- Meaning a speaker publicly commits to by using a certain form.

  Linda: "It's perfectly dry outside. You're mocking me?"

  (since Max said it's raining, he committed to the outside being wet)

- Inferences about a speaker's private cognitive states.

  Linda: "Or do you want me to take an umbrella?" Max: "No."

  (there must be a reason why Max said that)

- Social meaning, such as politeness, formality, peer-group style, ...

  Linda: "Could you be serious, please?" Max: "Sorry, I was just mocking you."

  (Linda indicates that she wants to avoid unnecessary discussions)

Meaning

Context of Use

Meaning that can often only be derived from context of use

- Scope of quantifiers, such as “Every student reads some book”.
- Word sense ambiguities, such as “I’m making it.”.
- Semantic relations between nouns in compounds, such as “play book”.

... and many others...

Meaning

Context of Use

Meaning that can often only be derived from context of use

- Scope of quantifiers, such as “Every student reads some book”.
- Word sense ambiguities, such as “I’m making it.”.
- Semantic relations between nouns in compounds, such as “play book”.

... and many others...

Interpretation interacts with non-linguistic perception

- Time, such as “now”, “tomorrow”, ...
- Location, such as “here”, ”there”, “That’s a beautiful city.”
- Speaker and hearer, such as “I”, “you”, ...

Chapter NLP:VI

VI. Semantics

- Intro
- Lexical and Word-level semantics
- Compositional semantics
- Discourse-level semantics
- Distributional semantics

Lexical Semantics

What is lexical semantics?

- The meaning of words and multi-word expressions and representing knowledge about word meaning.
- Covers word senses, semantic roles, and connotations.

Lexical Semantics

What is lexical semantics?

- The meaning of words and multi-word expressions and representing knowledge about word meaning.
- Covers word senses, semantic roles, and connotations.

Word senses

- Distinctions in meaning between different uses of the same form.
- Shared meanings between different forms.

Semantic roles

- Number of arguments of a predicate.
- Specific relationship the arguments bear to the predicate.

Connotation

- What word choice conveys beyond truth-conditional semantics.

Lexical Semantics

Word senses

What is a word sense?

- The meaning of a word.
- Words can have multiple senses, due to *polysemy* and *homonymy*.
- **Polysemy:** a word has different meanings that are related to each other (e.g. “university”)
- **Homonyny:** a word has different meanings that have no obvious relation to each other (e.g. “bank”)
- **Word sense disambiguation:** the task of identifying word meanings, i.e. tagging each word token with its sense
- **Synsets:** groups words into sets of synonyms called synsets and describes semantic relationships between synsets – As used in [WordNet]

Lexical Semantics

Word senses

String similarity – Simple Assumption: Similar wordforms share semantic properties

- Matchings (sub-)parts of strings
  - **begin:** begin, beginner, beginning . . . beggar . . . beer . . . bike
- share of character-trigrams
  - begin : end $\rightarrow$ 0 trigrams shared
  - ##street## : ##straight## $\rightarrow$ 4 trigrams shared (##s, #st, str, t##)
  - ##street## : ##streets## $\rightarrow$ 6 trigrams shared (##s, #st, str, tre, ree, eet)
- edit distance (Levenshtein distance) – see reference for algorithm [Wikipedia]
  - number of edit operations (insert, substitute, delete) on single characters
  - begin : end $\rightarrow$ 4
  - street : straight $\rightarrow$ 3

Lexical Semantics

Word senses

String patterns

- String = concatenated symbols $\sim$ representative of semantic unit
  - String matching in texts: retrieval of semantic unit
- dictionaries = list of equivalent strings $\sim$ representatives of semantic unit
  - s_1 = `United States, US, USA, America, ...`
  - s_2 = `United States, Germany, Israel, Ghana, ...`
  - s_3 = `because of, instead of, thus, ...`
- equivalency relation (w.r.t. to tertium comparationis c)
  - $a \equiv b \iff a \in s_i \wedge b \in s_i \wedge c(a) = c(b)$
- Regular expressions for patterns (see part Regular Grammars for detailed explanation)
  - high expressiveness by complex string patterns
  - composition, dictionaries, combinations

Lexical Semantics

Word senses

Example: “ride” has 16 senses, here is a selection:

- ride over, along, or through
- sit and travel on the back of animal, usually while controlling its motions
- be carried or travel on or in a vehicle
- be contingent on
- harass with persistent criticism or carping
- keep partially engaged by slightly depressing a pedal with the foot
- continue undisturbed and without interference
- move like a floating object

Lexical Semantics

Polysemy vs. Homonymy

Constructional polysemy

- Related senses that have the same lexical entry.

"newspaper" (physical object vs. abstract content)

Sense extension polysemy

- Regular ways of deriving new word senses given a member of a class.

"chicken" (animal vs. meat of the animal)

Lexical Semantics

Polysemy vs. Homonymy

Constructional polysemy

- Related senses that have the same lexical entry.

"newspaper" (physical object vs. abstract content)

Sense extension polysemy

- Regular ways of deriving new word senses given a member of a class.

"chicken" (animal vs. meat of the animal)

Homonoymy

- Unrelated word senses that have the same lexical entry.

"bank" (river bank vs. money bank)

Lexical Semantics

Word Sense Goes Wild *

Word senses may change over time

(Hamilton et al., ACL 2016)

Lexical Semantics

Word Sense Goes Wild *

Word senses may change over time

(Hamilton et al., ACL 2016)

Metaphoric word senses

- Metaphors add senses to words in (theoretically) unbounded ways.

"I have always despised politics.
But I have climbed to the top of that greasy pole."

Lexical Semantics

Semantic Roles

What are semantic roles?

- The roles the arguments of a predicate have in the state or activity captured by the predicate.
- Not to be confused with syntactic roles, such as subject or object.
- Different predicates have different semantic roles.

"She saw Max." vs. "She kissed Max." vs. "She ressembled Max."

Lexical Semantics

Semantic Roles

What are semantic roles?

- The roles the arguments of a predicate have in the state or activity captured by the predicate.
- Not to be confused with syntactic roles, such as subject or object.
- Different predicates have different semantic roles.

"She saw Max." vs. "She kissed Max." vs. "She ressembled Max."

Why is this *lexical* semantics?

Lexical Semantics

Semantic Roles

What are semantic roles?

- The roles the arguments of a predicate have in the state or activity captured by the predicate.
- Not to be confused with syntactic roles, such as subject or object.
- Different predicates have different semantic roles.

"She saw Max." vs. "She kissed Max." vs. "She ressembled Max."

Why is this *lexical* semantics?

- Syntax is important for identifying what roles an argument plays.
- But the predicate defines the semantic roles.

Lexical Semantics

Semantic Roles

What are semantic roles?

- The roles the arguments of a predicate have in the state or activity captured by the predicate.
- Not to be confused with syntactic roles, such as subject or object.
- Different predicates have different semantic roles.

"She saw Max." vs. "She kissed Max." vs. "She ressembled Max."

Why is this *lexical* semantics?

- Syntax is important for identifying what roles an argument plays.
- But the predicate defines the semantic roles.

Semantic role labeling

- The text analysis that finds the arguments taking on the semantic roles in a predicate.
- Used in text mining when deeper language understanding is required.

Chapter NLP:VI

VI. Semantics

- Intro
- Lexical and Word-level semantics
- Compositional semantics
- Discourse-level semantics
- Distributional semantics

Distributional Semantics

What is Distributional semantics?

- Describe word meaning in terms of the word's distribution
- Semantically "similar" words tend to occur in the context of the same words; here: "similar" means roughly "synonymous"
- The most important commonplace in NLP: **You shall know a word by the company it keeps** [Harris 1951, Firth 1957]
- Can by modeled by the cooccurrence of words in semantical meaningful units (sentences, paragraphs, words)
- Main ideas:
  - Meaning of words can be inferred from observations in large-scale, empirical language data
  - Meaning of words can be encoded in numerical vectors
  - Composition of word meanings can be modeled by compositional functions of their numerical representations
- Five hypothesis of distributional semantics: Statistical semantics hypothesis, Bag of words hypothesis, Distributional hypothesis, Extended distributional hypothesis, Latent relation hypothesis

Distributional Semantics

What is Distributional semantics?

Statistical semantics hypothesis

- “Statistical patterns of human word usage can be used to figure out what people mean (Weaver, 1955; Furnas et al., 1983)”
- word frequencies, collocation of words

Bag of words hypothesis:

- “The frequencies of words in a document tend to indicate the relevance of the document to a query (Salton et al., 1975)”
- Frequency is important, order can be neglected

Distributional Semantics

Distributional Hypothesis

"Words that occur in similar contexts tend to have similar meanings (Harris, 1954; Firth, 1957; Deerwester et al., 1990)"

- The symptoms of the **virus** can include dry **cough**, **fever** and fatigue.

  If you're making choices to spread the **virus** then more healthcare workers will have to treat more patients and more healthcare workers will get **sick**.

  So if you have symptoms, respiratory symptoms of **cough**, **fever**, sore throat, runny nose, headache, aches and pains, it is most likely that you have **Covid**, not flu.

  The union is compiling a list of contractors who will not grant full **sick** pay if employees get **Covid** or need to go into isolation.

  Software can be used to develop apps that detect when a user has spent time near another user who later tests positive for the **virus**.

- **Example:** covid, virus

	covid	virus	sick	fever	cough	software
covid	0	433	54	30	50	0
virus	433	0	230	99	354	11
sick	54	230	0	19	23	0
fever	30	99	19	0	780	0
cough	50	354	23	780	0	0
software	0	11	0	0	0	0

Distributional Semantics

Usage in Vector Space Model

	covid	virus	sick	fever	cough	software
covid	0	433	54	30	50	0
virus	433	0	230	99	354	11
sick	54	230	0	19	23	0
fever	30	99	19	0	780	0
cough	50	354	23	780	0	0
software	0	11	0	0	0	0

- Word context can be modeled as Word-Context-Matrix (also: Term-Term-Matrix)
- Event: cooccurrence of type A and type B within certain window
- Symmetry $|V| \times |V|$
- Possible windows: right / left neighbor (asymmetric), n tokens, sentence, paragraph, document
- Vector similarity is interpreted as semantic similarity. (e.g. Cosine Distance, see VSM in Text Model Section)

sim_{cosine}(virus, covid) = 0.83
sim_{cosine}(covid, death) = 0.73
sim_{cosine}(remote, death) = 0.48

sim_{cosine}(remote, work) = 0.71
sim_{cosine}(player, football) = 0.84
sim_{cosine}(goal, death) = 0.33

Distributional Semantics

Cooccurrences Graph

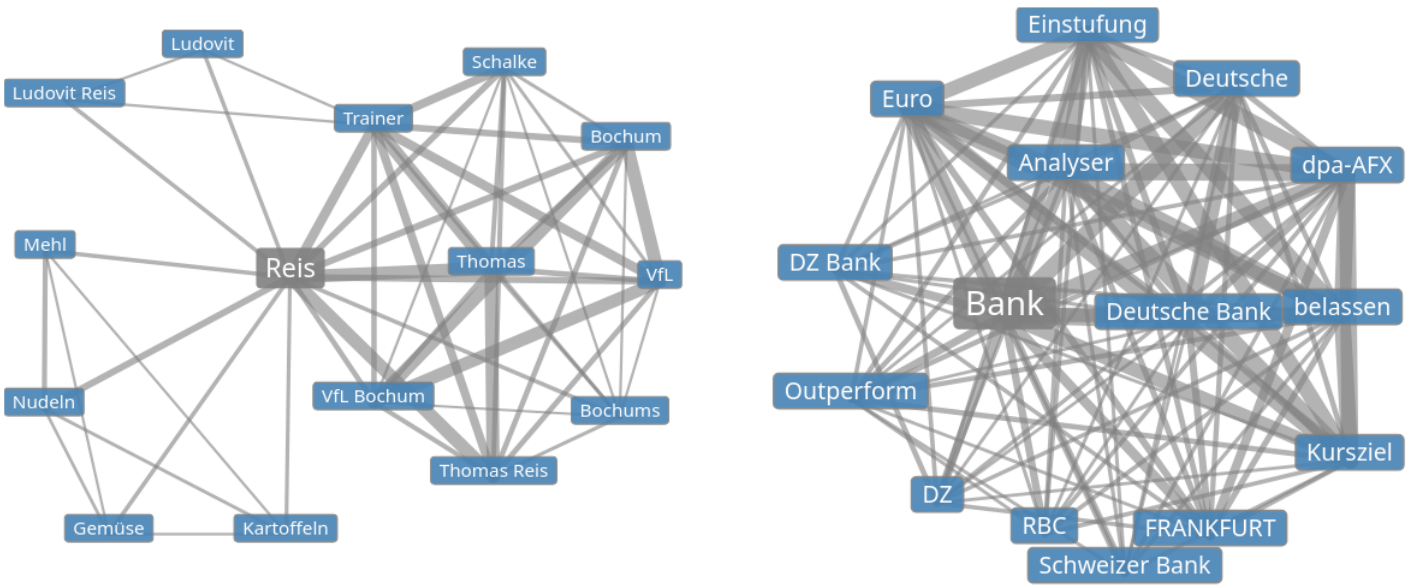

Source: Wortschatz Leipzig - Corpus *deu_news_2022 - German news corpus based on material from 2022* (**Reis** / **Bank**)

Distributional Semantics

Dimensional Reduction, Topic Models, Language Model Embedding

Problem: co-occurrence vectors have very high dimension; each context word is represented (i.e. one for each word)

Typical approach: dimensionality reduction algorithms

- Latent Semantic Analysis: dimensionality reduction via singular value decomposition
- Principal Component Analysis
- Topic Modelling
- Neural Network Based Language Models, Word2Vec, FastText

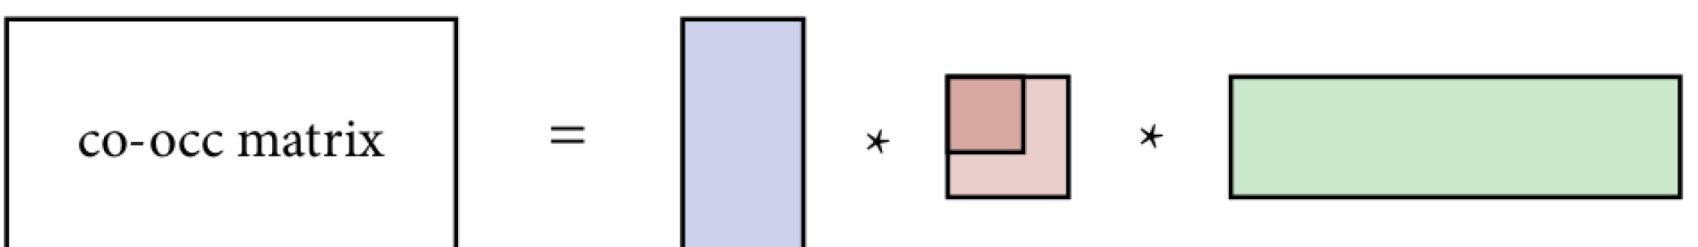

Remarks:

- Vector space distributional semantics can be used in order to mine a corpus for synsets. (Paradigmatic)
- If used with on left/right neighbor windows we can see common collocations. (Syntagmatic)
- Using significance tests for the word association strongly supports the interpretation and usefulness.
- See section cooccurrence analysis in part applications