Ling/CSE 472: Introduction to Computational Linguistics

May 12: Meaning representation

Overview

- Semantics
- Semantics in NLP
- Scheduling term project presentations
- Reading questions

Parsing makes explicit inherent structure.
So, does this tree represent meaning?

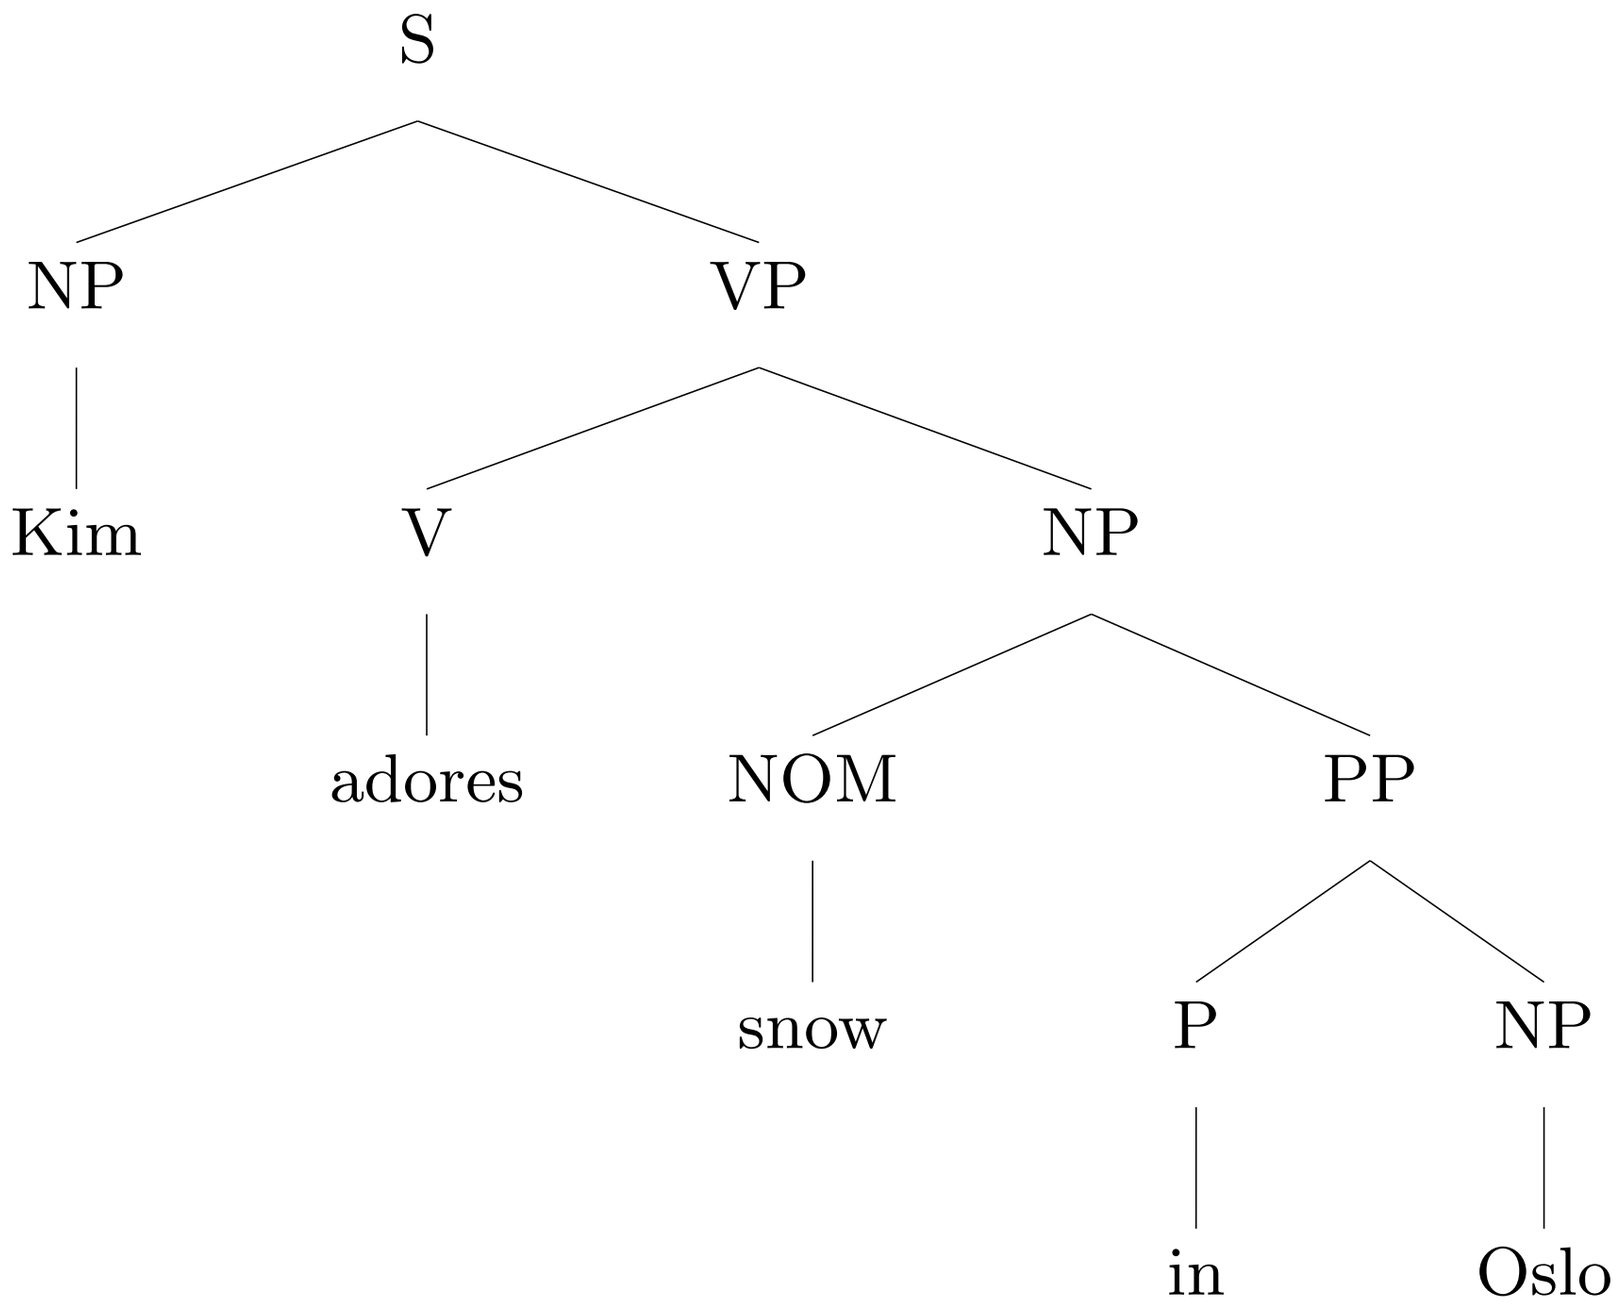

Why represent semantics?

- When “earlier” levels aren’t enough
- Bridge between linguistics and real world items/models

How could we put this tree in correspondence to a model of the world?

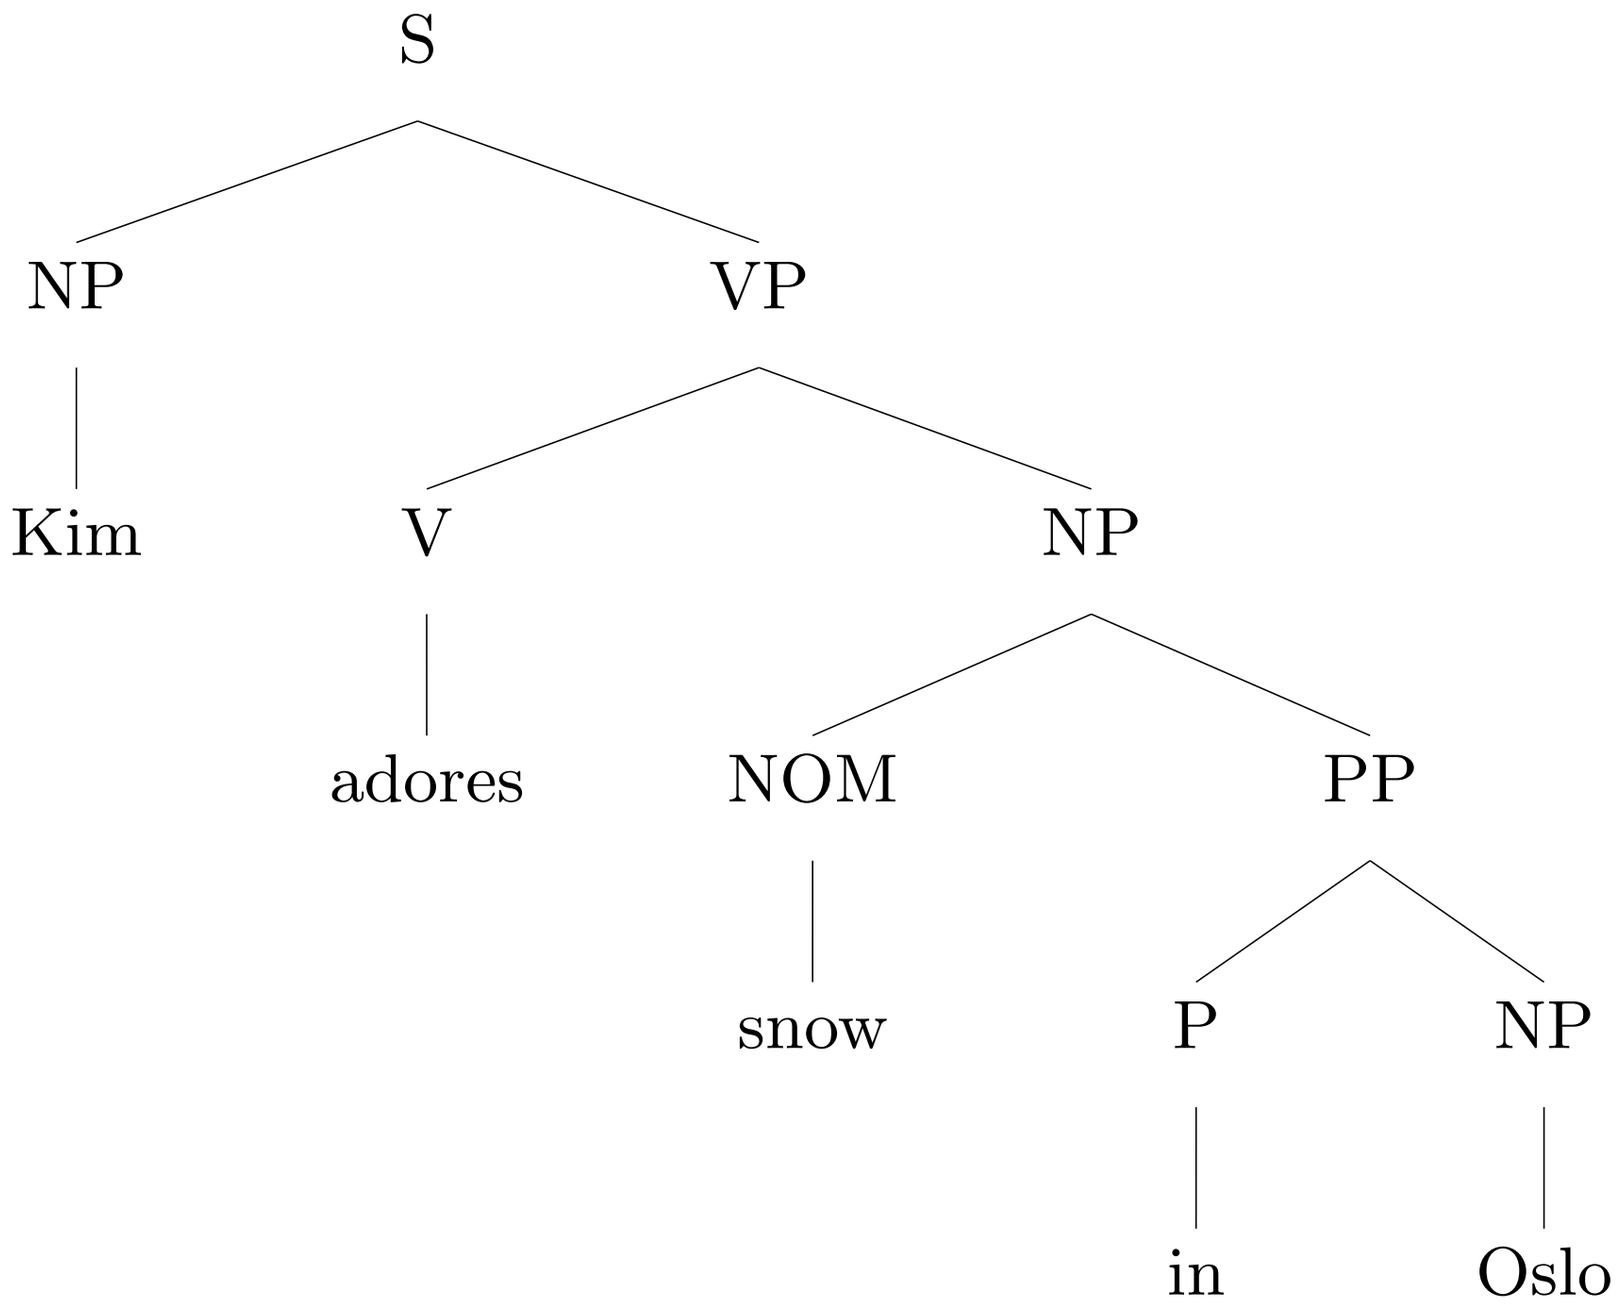

Semantics

- Create representations which can be put in correspondence with models of the world
- ... and which can be built compositionally via parsing

Basic model-theoretic semantics

- Create a model of the world, consisting of elements, sets of elements and relations

- Create an interpretation function which maps linguistic elements (parts of the semantic structure) to parts of the model

- Simple propositions are interpreted by checking their truth in the model

- Define semantics for “logical vocabulary”: *and, or, not, if, every, some*, ...

Model theoretic semantics example

- Entities: Joey:  Fluffy:  Tiger: 
- Properties: calm: {   }; angry: {  }
- Relations: knows: { <  ,  > , <  ,  > }

Model theoretic semantics example: denotations

- [[Fluffy]] = 
- [[angry]] = { x | x is angry } = {  }
- [[Fluffy is angry]] = True *iff* the entity denoted by *Fluffy* is in the set denoted by *angry*
- Compositionality: The process of determining the truth conditions of *Fluffy is angry* based on the denotations of its parts and its syntactic structure

Logical vocabulary gets special treatment

- *Fluffy is angry and Joey is not angry.*
  - What does *and* mean? (How does it affect the truth conditions of the whole?)
  - What does *not* mean?
- *Every cat is angry.*
  - What does *cat* mean? (Is this a logical operator?)
  - What does *every* mean?
- Is the division into logical and non-logical vocabulary an inherent property of language or an artifact of the system of meaning representation?

More on quantifiers

- The semantic type of a quantifier is a relation between sets, called the *restriction* and *body* (or *scope*) of the quantifier

  - [[every]] { <P,Q> | P ⊆ Q}

  - [[every cat is angry]] is True *iff* { x | x is a cat } ⊆ { y | y is angry }

  - [[some]] { <P,Q> | P ∩ Q ≠ ∅}

  - [[some cat is angry]] is True *iff* { x | x is a cat } ∩ { y | y is angry } ≠ ∅

- Where do those sets come from?

Why represent semantics?

- When “earlier” levels aren’t enough
- Bridge between linguistics and real world items/models

Semantics in NLP

- Construct knowledge base or model of the world
- Extract meaning representations from linguistic input
- Match input to world knowledge
- Produce replies/take action on the basis of the results
- In what other cases might semantic representations be useful?

Semantics in NLP

- In what other cases might semantic representations be useful?
  - Transfer-based MT
  - Building a knowledge base by “reading” the web (or wikipedia or...)
  - Generation

Semantic representations: Desiderata (Jurafsky & Martin)

- Verifiability: We must be able to compare the representation to a knowledge base
- Lack of ambiguity: A semantic representation should have just one interpretation
- Canonical form: A given interpretation should have just one representation
  - Does Maharani have vegetarian dishes?
  - Do they have vegetarian food at Maharani?
  - Are vegetarian dishes served at Maharani?
  - Does Maharani have vegetarian fare?
  - But not: Can vegetarians eat at Maharani?
- Expressiveness: Must be able to adequately represent a wide range of expressions

Semantic Representations: Desiderata (Copestake et al 2005)

- Expressive Adequacy: The framework must allow linguistic meanings to be expressed correctly

- Grammatical Compatibility: Semantic representations must be linked clearly to other kinds of grammatical information (most notably syntax)

- Computational Tractability: It must be possible to process meanings and to check semantic equivalence and to express relationships between semantic representations straightforwardly

- Underspecifiability: Semantic representations should allow underspecification (leaving semantic distinctions unresolved), in such a way as to allow flexible, monotonic resolution of such partial semantic representations

Evaluation slide

- How would we evaluate a system of semantic representations?
- How would we evaluate a parsing system which produces semantic representations from input?
  - What's the gold standard?
  - What's the baseline?
  - What are the metrics?
  - What else might we need?

Reading questions

- Are the symbols they use for existential and universal quantifiers the same that are used within syntax classes?

- Is there a basic/standard logical vocabulary for natural languages or do the contents of the logical vocabulary depend on the task or language?

Reading questions

- Would basic semantic parsers be able to deal with things like lingo or words that possibly mean different things than their, say, dictionary meanings, like words that have a definition both in a normal dictionary and another in, say, the urban dictionary? Also thinking about words who's meanings depend on their context.

- Could/how could semantic parsers deal with ambiguity involving words or phrases that have multiple and relatively uncommon, informal meanings, like slang? Where do the abilities of semantic parsers tend to fall short?

Reading questions

- How might semantic parsers deal with words that has multiple meanings under different contexts? Moreover, I wonder if they could identify special structures, for example double negatives expressing negative/positive in English. ("I ain't got no money", and "The name Bob is not uncommon")

- How effective are semantic parsers with non-straightforward sentences like ones with slang, a lot of semantic ambiguity, informal/unusual syntax, etc?

- How effective are semantic parsers in understanding and interpreting figurative language, such as metaphors, similes, and analogies? Humans often rely on shared cultural knowledge/context to understand implicit information within text, but it seems like these forms of information would be difficult to encode and recognize.

Reading questions

- Where's the limit on how much a semantic parser can infer or interpret vagueness? The chapter gives us the example of different ways to ask whether a restaurant has vegetarian options, but human speech goes far beyond that. For example, saying "it's cold in here" could mean "can you close that window?" or "can I borrow your jacket?" or "can we go to a different room?" And a human being can almost always correctly infer what the speaker wants to convey, but how does a semantic parser compare in that regard?

Reading questions

- For the First-Order Logic representation of any sentence, can we think of it as we are breaking any sentence to the smallest logical sub-sentence and construct a statement, either function or variable, out of one or two objects with a relation between them? From here, we use logical connectives to combine multiple of the "sub-sentence" into a full sentence?

- If so, it looks really like a relational database where each function and variable are just tables. Parsing a sentence is like adding rows to corresponding tables of function and variable. When we want to generate a response to a question, we just parse the question and do a query on all possible objects?

- Is first order logic the most popular or are there other forms that are more suitable and better in practice?

Reading questions

- How exactly does semantic parsing interact with syntactic parsing? Is semantic parsing a sort of layer that is 'added on' after we have parsed a text for syntactic information?

- Are some types of meaning representations or models better for certain kinds of semantic parsing tasks? Or, what other factors influence which meaning representation and model is best to use?

Reading questions

- Wow the idea that a simple True of False of some role can tell us interesting things about some statement is really fascinating to me. In the reading it mentioned that we don't have simple concepts such as differentiating between eater and eaten and adding detail about location and time. How useful is the semantic parsing as shown in the book without these important details? Of course it won't be completely useless by any means, but is having the information presented in Chapters 22 and 24 important for effective semantic parsing?

Dagstuhl seminar reportback

- How does dependency grammar handle ambiguity?
- What we’re doing in the working groups here…