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?

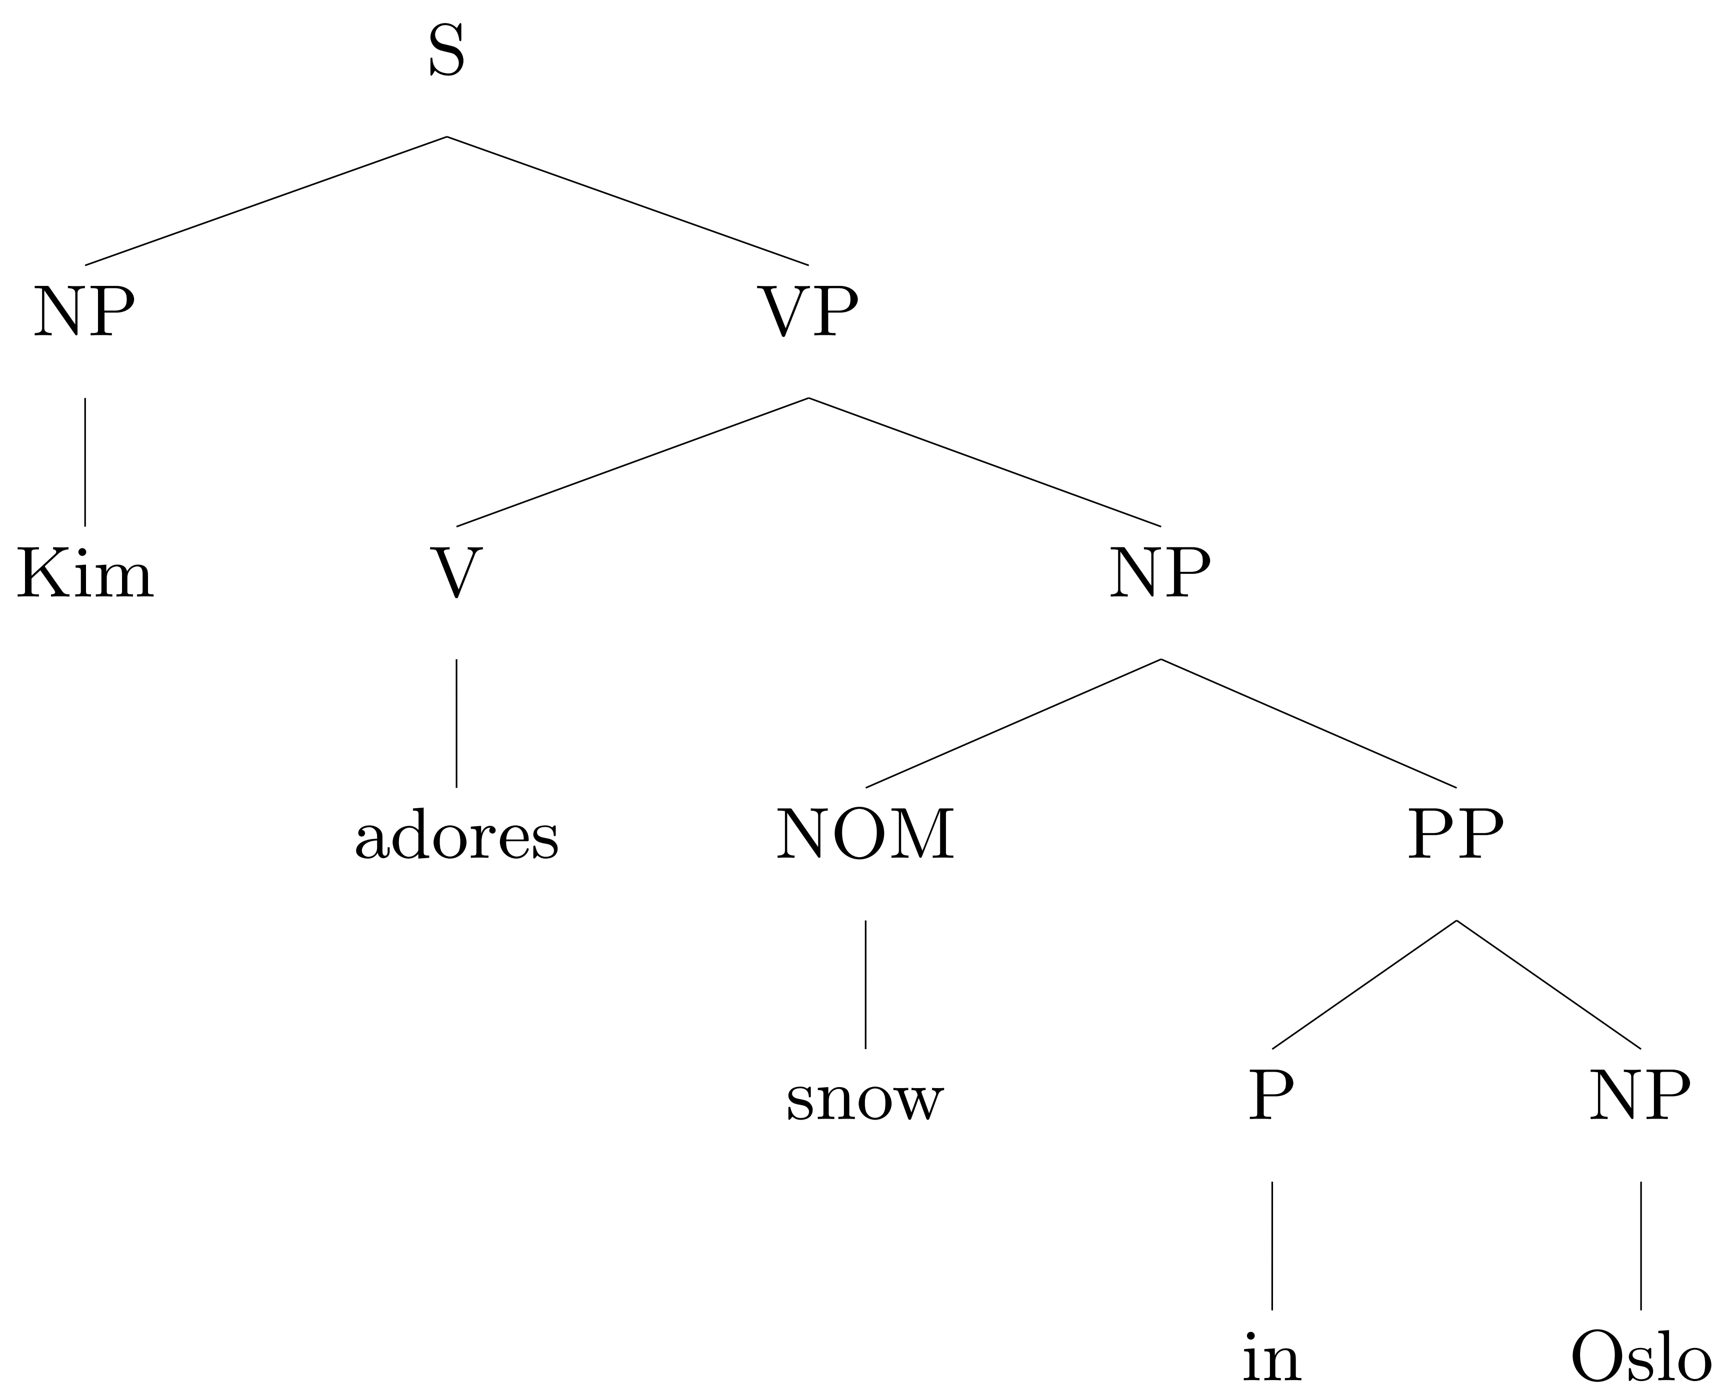

Why represent semantics?

- 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?

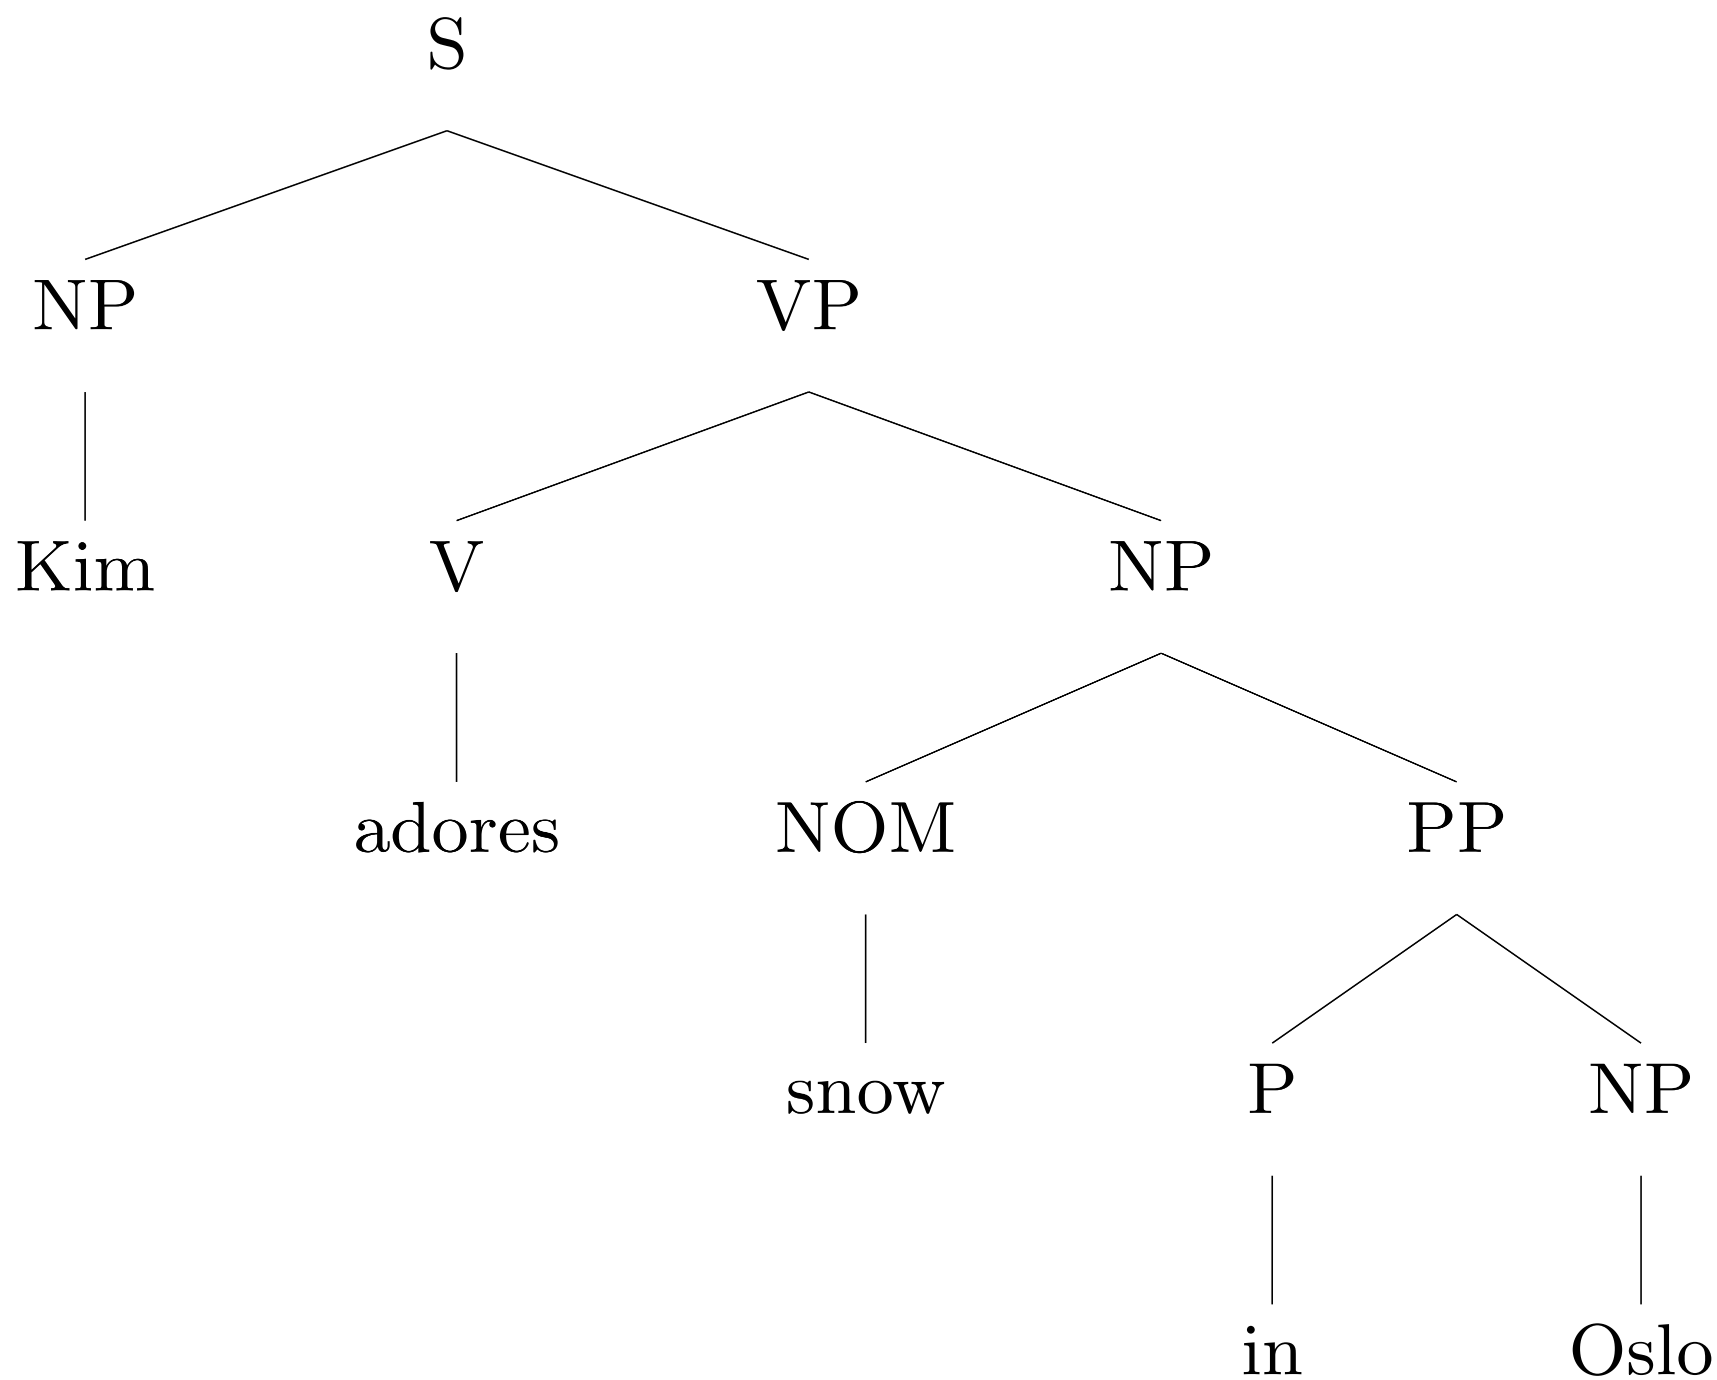

Semantics

- 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]] $\{ \langle P,Q \rangle \mid P \subseteq Q \}$

  - [[every cat is angry]] is True *iff* $\{ x \mid x \text{ is a cat} \} \subseteq \{ y \mid y \text{ is angry} \}$

  - [[some]] $\{ \langle P,Q \rangle \mid P \cap Q \neq \emptyset \}$

  - [[some cat is angry]] is True *iff* $\{ x \mid x \text{ is a cat} \} \cap \{ y \mid y \text{ is angry} \} \neq \emptyset$

- 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?

Overview

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

Scheduling term project presentations

- June 2
  - 1
  - 2
  - 3
  - 4
  - 5
- June 4
  - 1
  - 2
  - 3
  - 4
  - 5

Reading questions

- How do systems take the FOL and check it against the state of the actual world? Just information online?
- What does a knowledge base look like? Is it manually created? How is it ordered for the models to access and scan for information?
- Does a meaning representation evaluate to true if it's in the knowledge base?
- How should the domain of a model be determined? Should it be as narrow as possible? How does this relate to the expressiveness from JM Ch 16.1, where it says that a meaning representation scheme/ model should be expressive enough to handle a wide range of subject matter?

Reading questions

- The textbook says that Italian food is still vague. Then what exactly is a representation that is not vague? Is asking for pasta still a vague representation since there are multiple kinds of pasta?

- Can subsumption make use of inference or is it limited to its knowledge base? Using the chapter's example where IlFornaio is not linked to VegetarianRestaurant, what should be the expected outcome of asking "If all vegetarians can eat at IlFornaio, is it a vegetarian restaurant?"

Reading questions

- Does this apply for all languages? I'm assuming yes because regardless of the language, there would be a way to represent it using first-order logic since there has to be some sort of relation in semantics?
- How would we adjust logical representations of sentence meaning to account for slang or idioms that are units of knowledge rather than sentences built word-by-word?

Reading questions

- Why does non-logical vocabulary of meaning representation seem so complex in comparison to logical vocabulary? Like having things have denotation, domain, etc., that are absent from logical vocabulary (or just not explained)?

Reading questions

- The chapter briefly discusses lambda notation and says that it “provides a way to incrementally gather arguments to a predicate when they do not all appear together as daughters of the predicate in a parse tree.” What does this mean, exactly?

- Why do we need the lambda notation? Can the described reduction not be applied to the expressions with quantifiers?

- What are lambdas capable of representing that FOL on its own can't?

Reading questions

- It seems like the concepts and practices described in this chapter are almost exclusively from the computer science side of things, as opposed to linguistics. Is that really representative of categorical semantic representations in use today?

- The vegetarian sentence "I'd like to find a restaurant where I can get vegetarian food." was represented as Serves(X, Vegetarian food). Does this mean most of computational semantics is hard coded? Or similar to syntax, does the computer simply select the most likely candidate?

Reading questions

- If description logics are a subset of first order logic, what exactly is not included in description logics.
- How do we move from some natural language input to a semantic structure? The text mentions dependency parses, so do we typically just take the root and its dependents as a function and its variables? And then, what is a knowledge base and how does it come into play?
- Have there been any attempts to make standard semantic representation frameworks, like UD is for dependency parsing?

Reading questions

- How is everything then implemented in application i.e. how are first-order logic representations implemented in large NLP systems and how well do these representations perform in said systems? Do N-grams play a role in this as well and/or are first-order logic representations applied to N-gram models?

- Now that we are stepping away from syntax and more into a semantic light, which has more input and impact in the field of NLP? Which is developing more/more quickly?

- There doesn't seem to be an efficient way of encoding aspects in FOL. Is this a tradeoff that is usually accepted?

Reading questions

- I feel like getting data for this is a much harder thing to deal with as opposed to something such as Treebank, which can use a set of annotated trees as a training set, and extrapolate information from those. Here however, you can't train any system on the data--instead you have to continually make sure you have enough data for this to work efficiently and also constantly update outdated data. In practice, how do we efficiently collect data sets large enough to efficiently use this type of representation? Is it mainly through manual entry, a computerized scan of the Internet for information, or something else?

Reading questions

- This chapter uses pairs of questions and answers to demonstrate the usefulness of meaning representations. Is this the only application of meaning representations in NLP?

- Program verification commonly works on the First Order Logic level and has shown popularity among Programming Languages research. Part of the reason why this is possible is that PL is strictly defined compared to natural language. I'd like to know more about if there are attempts to use formal verification of natural language as part of a pipeline for QA tasks or "fake news detection"?

Reading questions

- Although I'm somewhat familiar with first-order-logic, I didn't fully understand forward chaining and backward chaining. Why is backward chaining sound, but reasoning backward is invalid? What is plausible reasoning/abduction?