6.034 Notes: Section 9.2

Slide 9.2.1

In this chapter, we take a quick survey of some aspects of natural language understanding. Our goal will be to capture the **meaning** of sentences in some detail. This will involve finding representations for the sentences that can be connected to more general knowledge about the world. This is in contrast to approaches to dealing with language that simply try to match textual patterns, for example, web search engines.

We will briefly provide an overview of the various levels and stages of natural language processing and then begin a more in-depth exploration of language syntax.

6.034 Artificial Intelligence

- Natural Language Understanding
  - Getting at the meaning of text and speech
  - Not just pattern matching
- Overview
- Syntax

tlp • Spring 02 • 1

Applications of NLU

- Interfaces to databases (weather, financial,…)
- Automated customer service (banking, travel,…)
- Voice control of machines (PCs, VCRs, cars,…)
- Grammar and style checking
- Summarization (news, manuals, …)
- Email routing
- Smarter Web Search
- Translating documents
- Etc.

tlp • Spring 02 • 2

Slide 9.2.2

The motivation for the study of natural language understanding is twofold. One is, of course, that language understanding is one of the quintessentially human abilities and an understanding of human language is one of key steps in the understanding of human intelligence.

In addition to this fundamental long-term scientific goal, there is a pragmatic shorter-term engineering goal. The potential applications of in-depth natural language understanding by computers are endless. Many of the applications listed here are already available in some limited forms and there is a great deal of research aimed at extending these capabilities.
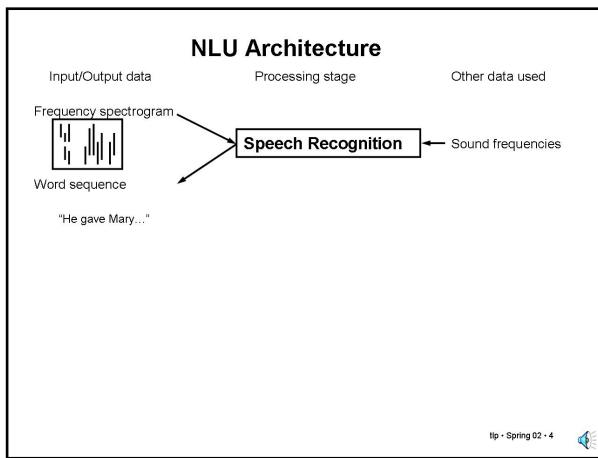

Slide 9.2.3

Language is an enormously complex process, which has been studied in great detail for a long time. The study of language is usually partitioned into a set of separate sub-disciplines, each with a different focus. For example, phonetics concerns the rules by which sounds (phonemes) combine to produce words. Morphology studies the structure of words: how tense, number, etc is captured in the form of the word. Syntax studies how words are combined to produce sentences. Semantics studies how the meaning of words are combined with the structure of a sentence to produce a meaning for the sentence, usually a meaning independent of context. Pragmatics concerns how context factors into the meaning (e.g. "it's cold in here") and finally there's the study of how background knowledge is used to actually understand the meaning the utterances.

We will consider the process of understanding language as one of progressing through various "stages" or processing that break up along the lines of these various subfields. In practice, the processing may not be separated as cleanly as that, but the division into stages allows us to focus on one type of problem at a time.

Levels of language analysis

- Phonetics: sounds → words
- Morphology: morphemes → words (jump+ed=jumped)
- Syntax: word sequence → sentence structure
- Semantics: sentence structure + word meaning → sentence meaning
- Pragmatics: sentence meaning + context → deeper meaning
- Discourse and World Knowledge: connecting sentences and background knowledge to utterances.

tlp • Spring 02 • 3

NLU Architecture

Slide 9.2.4

If one considers the problem of understanding speech, the first stage of processing is, conceptually, that of converting the spoken utterance into a string of words. This process is extremely complex and quite error prone and, today, cannot be solved without a great deal of knowledge about what the words are likely to be. But, in limited domains, fairly reliable transcription is possible. Even more reliability can be achieved if we think of this stage as producing a few alternative interpretations of the speech signal, one of which is very likely to be the correct interpretation.

Slide 9.2.5

The next step is **syntax**, that is, computing the structure of the sentence, usually in terms of phrases, such as noun phrases, verb phrases and prepositional phrases. These nested phrases will be the basis of all subsequent processing. Syntactic analysis is probably the best developed area in computational linguistics but, nevertheless, there is no universally reliable "grammar of English" that one can use to parse sentences as well as trained people can. There are, however, a number of wide-coverage grammars available.

We will see later that, in general, there will not be a unique syntactic structure that can be derived from a sequence of words.

NLU Architecture

NLU Architecture
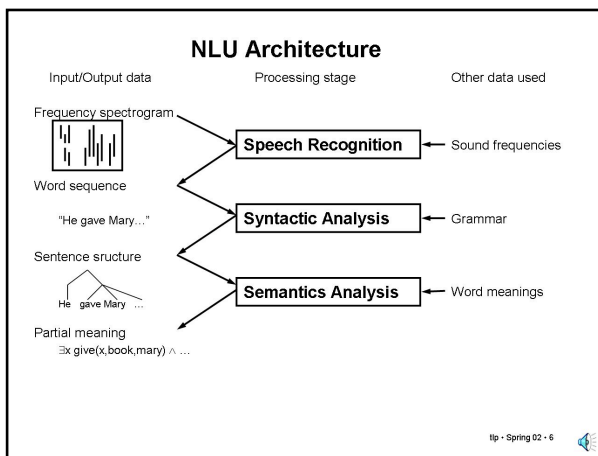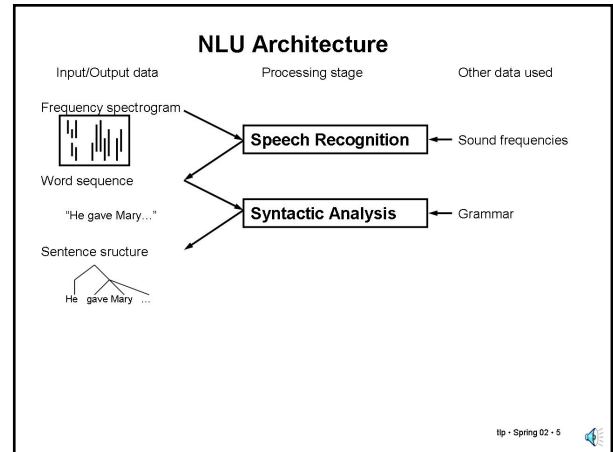
Slide 9.2.6

Given the sentence structure, we can begin trying to attach meaning to the sentence. The first such phase is known as **semantics**. The usual intent here is to translate the syntactic structure into some form of logical representation of the meaning - but without the benefit of context. For example, who is being referred to by a pronoun may not be determined at this point.
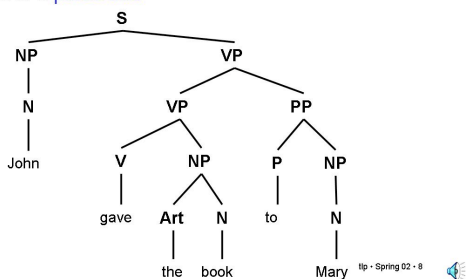
Slide 9.2.7

We will focus in this chapter on syntax and semantics, but clearly there is a great deal more work to be done before a sentence could be understood. One such step, sometimes known as **pragmatics**, involves among other things disambiguating the various possible senses of words, possible syntactic structures, etc. Also, trying to identify the referent of pronouns and descriptive phrases. Ultimately, we have to connect the meaning of the sentence with general knowledge in order to be able to act on it. This is by far the least developed aspect of the whole enterprise. In practice, this phase tends to be very application specific.

NLU Architecture
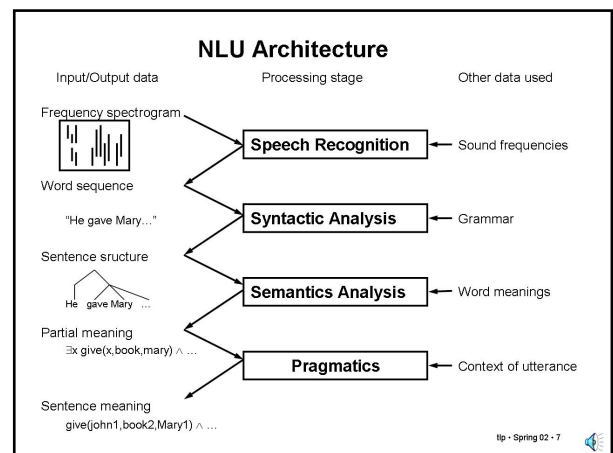
Syntax

- Grammar captures legal structures in the language
- Parsing involves finding the legal structure(s) for a sentence
- The result is a parse tree

```
                    S
          ┌─────────┴─────────┐
         NP                   VP
          │            ┌───────┴───────┐
          N           VP               PP
          │        ┌───┴───┐        ┌───┴───┐
        John       V      NP        P      NP
                   │    ┌──┴──┐      │       │
                 gave  Art    N     to       N
                        │     │              │
                       the   book          Mary
```
tlp · Spring 02 · 8

Slide 9.2.8

In the rest of this section, we will focus on syntax. The description of the legal structures in a language is called a **grammar**. We'll see examples of these later. Given a sentence, we use the grammar to find the legal structures for a sentence. This process is called **parsing** the sentence. The result is one or more **parse trees**, such as the one shown here, which indicates that the sentence can be broken down into two **constituents**, a noun phrase and a verb phrase. The verb phrase, in turn, is composed of another verb phrase followed by a prepositional phrase, etc.

Our attempt to understand sentences will be based on assigning meaning to the individual constituents and then combining them to construct the meaning of the sentence. So, in this sense, the constituent phrases are the atoms of meaning.

Slide 9.2.9

A grammar is typically written as a set of **rewrite rules** such as the ones shown here in blue. Bold-face symbols, such as S, NP and VP, are known as non-terminal symbols, in that they can be further re-written. The non-bold-face symbols, such as John, the and boy, are the words of the language - also known as the terminal symbols.

Grammars

- We can think of grammars as a set of rules for rewriting strings of symbols:
  - **S** → **NP VP**
  - **NP** → **Name**
  - **NP** → **Art N**
  - **Name** → John
  - **Art** → the
  - **N** → boy

tlp · Spring 02 · 9

Grammars

- We can think of grammars as a set of rules for rewriting strings of symbols:
  - **S** → **NP VP**
  - **NP** → **Name**
  - **NP** → **Art N**
  - **Name** → John
  - **Art** → the
  - **N** → boy
- The string S can be rewritten as NP followed by VP

tlp · Spring 02 · 10

Slide 9.2.10

The first rule, S -> NP VP, indicates that the symbol S (standing for sentence) can be rewritten as NP (standing for noun phrase) followed by VP (standing for verb phrase).

Slide 9.2.11

The symbol NP, can be rewritten either as a Name or as an Art(icle), such as the, followed by a N(oun), such as boy.

Grammars

- We can think of grammars as a set of rules for rewriting strings of symbols:
  - **S** → **NP VP**
  - **NP** → **Name**
  - **NP** → **Art N**
  - **Name** → John
  - **Art** → the
  - **N** → boy
- The string S can be rewritten as NP followed by VP
- The string NP can be rewritten either as Name (which can be rewritten as John) or as an Art (such as the) followed by an N (such as boy).

tlp · Spring 02 · 11

Grammars

- We can think of grammars as a set of rules for rewriting strings of symbols:
  - $S \rightarrow NP\ VP$
  - $NP \rightarrow Name$
  - $NP \rightarrow Art\ N$
  - $Name \rightarrow John$
  - $Art \rightarrow the$
  - $N \rightarrow boy$
- The string S can be rewritten as NP followed by VP
- The string NP can be rewritten either as Name (which can be rewritten as John) or as an Art (such as the) followed by an N (such as boy).
- A sentence is legal if we can find a sequence of rewrite rules that, starting from the symbol S, generate the sentence. This is called parsing the sentence.

tlp • Spring 02 • 12

Slide 9.2.12

If we can find a sequence of rewrite rules that will rewrite the initial S into the input sentence, the we have successfully parsed the sentence and it is legal.

Note that this is a search process like the ones we have studied before. We have an initial state, S, at any point in time, we have to decide which grammar rule to apply (there will generally be multiple choices) and the result of the application is some sequence of symbols and words. We end the search when the words in the sentence have been obtained or when we have no more rules to try.

Slide 9.2.13

Note that the successful sequence of rules applied to achieve the rewriting give us the parse tree. Note that this excludes any "wrong turns" we might have taken during the search.

Grammars

- We can think of grammars as a set of rules for rewriting strings of symbols:
  - $S \rightarrow NP\ VP$
  - $NP \rightarrow Name$
  - $NP \rightarrow Art\ N$
  - $Name \rightarrow John$
  - $Art \rightarrow the$
  - $N \rightarrow boy$
- The string S can be rewritten as NP followed by VP
- The string NP can be rewritten either as Name (which can be rewritten as John) or as an Art (such as the) followed by an N (such as boy).
- A sentence is legal if we can find a sequence of rewrite rules that, starting from the symbol S, generate the sentence. This is called parsing the sentence.
- The sequence of rules applied also give us the parse tree.

tlp • Spring 02 • 13

Good Grammars

- Differentiates between "correct" and "incorrect" sentences
  - The boy hit the ball   correct
  - The hit boy the ball (*)   incorrect

tlp • Spring 02 • 14

Slide 9.2.14

What makes a good grammar?

The primary criterion is that it differentiates correct sentences from incorrect ones. (By convention an asterisk next to a sentence indicates that it is not grammatical).

Slide 9.2.15

The other principal criterion is that it assigns "meaningful" structures to sentences. In our case, this literally means that it should be possible to assign meaning to the sub-structures. For example, a noun phrase will denote an object while a verb phrase will denote an event or an action, etc.

Good Grammars

- Differentiates between "correct" and "incorrect" sentences
  - The boy hit the ball
  - The hit boy the ball (*)
- Assigns meaningful structure to the sentences
  - (The boy) (hit the ball)
  - (The) (boy hit) (the ball)

tlp • Spring 02 • 15

Good Grammars

- Differentiates between "correct" and "incorrect" sentences
  - The boy hit the ball
  - The hit boy the ball (*)
- Assigns meaningful structure to the sentences
  - (The boy) (hit the ball)
  - (The) (boy hit) (the ball)
- Compact and modular, e.g. all these NPs can be used in any context NPs are allowed:
  - **NP → Name**   John
  - **NP → Art N**   the boy
  - **NP → Art Adj N**   the tall girl
  - **NP → Art N that VP**   the dog that barked

tlp • Spring 02 • 16

Slide 9.2.16

Among the grammars that meet our principal criteria we prefer grammars that are compact, that is, have fewer rules and are modular, that is, define structures that can be re-used in different contexts - such as noun-phrase in this example. This is partly for efficiency reasons in parsing, but is partly because of Occam's Razor - the simplest interpretation is best.

Slide 9.2.17

There are many possible types of grammars. The three types that are most common in computational linguistics are regular grammars, context-free grammars and context-sensitive grammars. These grammars can be arranged in a hierarchy (the Chomsky hierarchy) according to their generality. In this hierarchy, the grammars in higher levels fully contain those below and there are languages in the more general grammars not expressible in the less general grammars.

The least general grammar of some interest in computational linguistics are the **regular grammars**. These grammars are composed of rewrite rules of the form A -> x or A -> x B. That is, a non-terminal symbol can be rewritten as a string of terminal symbols or by a string of terminal symbols followed by a non-terminal symbol.

Types of Grammars

- There's a hierarchy of grammar types that can be classified by their generality. Some common types in wide use (from less general to more general):
  - Regular grammars – Rules are of the form:
    - $A \rightarrow x$ or $A \rightarrow x\,B$

  - A, B are single non-terminal symbols
  - x is a string of terminal symbols
  - α, β, γ are strings of terminal & non-terminal symbols.

tlp • Spring 02 • 17

Types of Grammars

- There's a hierarchy of grammar types that can be classified by their generality. Some common types in wide use (from less general to more general):
  - Regular grammars – Rules are of the form:
    - $A \rightarrow x$ or $A \rightarrow x\,B$
  - Context Free grammars – Rules are of the form:
    - $A \rightarrow \gamma$

  - A, B are single non-terminal symbols
  - x is a string of terminal symbols
  - α, β, γ are strings of terminal & non-terminal symbols.

tlp • Spring 02 • 18

Slide 9.2.18

At the next level are the **context-free grammars**. In these grammars, a non-terminal symbol can be rewritten into any combination of terminal and non-terminal symbols. Note that since the non-terminal appears alone in the left-hand side (lhs) of the rule, it is re-written independent of the context in which it appears - and thus the name.

Slide 9.2.19

Finally, in **context-sensitive** grammars, we are allowed to specify a context for the rewriting operation.

There are even more general grammars (known as Type 0) which we will not deal with at all.

Types of Grammars

- There's a hierarchy of grammar types that can be classified by their generality. Some common types in wide use (from less general to more general):
  - Regular grammars – Rules are of the form:
    - $A \rightarrow x$ or $A \rightarrow x\,B$
  - Context Free grammars – Rules are of the form:
    - $A \rightarrow \gamma$
  - Context Sensitive grammars – Rules are of the form:
    - $\alpha\,A\,\beta \rightarrow \alpha\,\gamma\,\beta$

  - A, B are single non-terminal symbols
  - x is a string of terminal symbols
  - α, β, γ are strings of terminal & non-terminal symbols.

tlp • Spring 02 • 19

Types of Grammars

- Regular grammars cannot capture some of the nested structures of natural language. The language $a^n b^n$ is not a regular language and there are legal sentences with that type of structure.

tlp • Spring 02 • 20

Slide 9.2.20

The language of parenthesized expressions, that is, n left parens followed by n right parens is the classic example of a non-regular language that requires us to move to context-free grammars. There are legal sentences in natural languages whose structure is isomorphic to that of parenthesized expressions (the cat likes tuna; the cat the dog chased likes tuna; the cat the dog the rat bit chased likes tuna). Therefore, we need at least a context-free grammar to capture the structure of natural languages.

Slide 9.2.21

There have been several empirical proofs that there exist natural languages that have non-context-free structure.

Types of Grammars

- Regular grammars cannot capture some of the nested structures of natural language. The language $a^n b^n$ is not a regular language and there are legal sentences with that type of structure.
- Some constructions in some natural languages have also been shown not to be context free.

tlp • Spring 02 • 21

Types of Grammars

- Regular grammars cannot capture some of the nested structures of natural language. The language $a^n b^n$ is not a regular language and there are legal sentences with that type of structure.
- Some constructions in some natural languages have also been shown not to be context free.
- But, much of the structure of natural languages can be captured in a context free language and we will restrict ourselves to context free grammars.

tlp • Spring 02 • 22

Slide 9.2.22

However, much of natural language can be expressed in context-free grammars extended in various ways. We will limit ourselves to this class.

Slide 9.2.23

Here's an example of a context free grammar for a small subset of English. Note that the vertical band is a short hand which can be read as "or"; it is a notation for combining multiple rules with identical left hand sides. Many variations on this grammar are possible but this illustrates the style of grammar that we will be considering.

A Simple Context-Free Grammar

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

- Article → the | a | an | this | that ...
- Preposition → to | in | on | near ...
- Conjunction → and | or | but ...
- Pronoun → I | you | he | me | him ...
- Noun → book | flight | meal ...
- Name → John | Mary | Boston ...
- Verb → book | include | prefer ...
- Adjective → first | earliest | cheap ...

tlp • Spring 02 • 23

Grammar Rules

- We can use our rule language to write grammar rules (we've seen this before).
- Assume words in sentence is represented as a set of facts:
  - (John 0 1)
  - (ran 1 2)
- Then a rule would be represented:
  - S → NP VP
  - `((S ?s1 ?s3) :- (NP ?s1 ?s2)(VP ?s2 ?s3))`
  - With, for example, ?s1=0, ?s2=1 and ?s3=2, r1 would match the sentence "John ran"

tlp • Spring 02 • 24

Slide 9.2.24

At this point, we should point out that there is a strong connection between these grammar rules that we have been discussing and the logic programming rules that we have already studied. In particular, we can write context-free grammar rules in our simple Prolog-like rule language.

We will assume that a set of facts are available that indicate where the particular words in a sentence start and end (as shown here). Then, we can write a rule such as S -> NP VP as a similar Prolog-like rule, where each non-terminal is represented by a fact that indicates the type of the constituent and the start and end indices of the words.

Slide 9.2.25

In the rest of this Chapter, we will write the rules in a simpler shorthand that leaves out the word indices. However, we will understand that we can readily convert that notation into the rules that our rule-interpreters can deal with.

Grammar Rules

- We can use our rule language to write grammar rules (we've seen this before).
- Assume words in sentence is represented as a set of facts:
  - (John 0 1)
  - (ran 1 2)
- Then a rule would be represented:
  - S → NP VP
  - `((S ?s1 ?s3) :- (NP ?s1 ?s2)(VP ?s2 ?s3))`
  - With, for example, ?s1=0, ?s2=1 and ?s3=2, r1 would match the sentence "John ran"
- We will write these grammar rules in the following shorthand:
  - ((S) :- (NP) (VP))
  which would just generate the rule above.

tlp • Spring 02 • 25

Grammar Rules

- We can use our rule language to write grammar rules (we've seen this before).
- Assume words in sentence is represented as a set of facts:
  - (John 0 1)
  - (ran 1 2)
- Then a rule would be represented:
  - S → NP VP
  - `((S ?s1 ?s3) :- (NP ?s1 ?s2)(VP ?s2 ?s3))`
  - With, for example, ?s1=0, ?s2=1 and ?s3=3, r1 would match the sentence "John ran"
- We will write these grammar rules in the following shorthand:
  - ((S) :- (NP) (VP))
  which would just generate the rule above.
- Rules indicating the category for particular words will be written:
  - ((NP) :- John)

tlp • Spring 02 • 26

Slide 9.2.26

We can also use the same syntax to specify the word category of individual words and also turn these into rules.

Slide 9.2.27

We can make a small modification to the generated rule to keep track of the parse tree as the rules are being applied. The basic idea is to introduce a new argument into each of the facts which keeps track of the parse tree rooted at that component. So, the parse tree for the sentence is simply a list, starting with the symbol S, and whose other components are the trees rooted at the NP and VP constituents.

Parsing

- We can construct a parse tree by modifying the generated rules slightly:
  - `((S (s ?np ?vp) ?s1 ?s3)) :-`
    `(NP ?np ?s1 ?s2)(VP ?vp ?s2 ?s3))`
- This rule now shows how to build the parse tree for the whole sentence (s ?np ?vp) from the parse trees for the constituent NP tree (bound to ?np) and VP tree (bound to ?vp).

tlp • Spring 02 • 27

Parsing

- We can construct a parse tree by modifying the generated rules slightly:
  - `((S (s ?np ?vp) ?s1 ?s3)) :- (NP ?np ?s1 ?s2)(VP ?vp ?s2 ?s3))`
  - This rule now shows how to build the parse tree for the whole sentence (s ?np ?vp) from the parse trees for the constituent NP tree (bound to ?np) and VP tree (bound to ?vp).
- This can all be generated automatically from the same shorthand notation:
  - `((S) :- (NP) (VP))`

tlp • Spring 02 • 28

Slide 9.2.28

This additional bit of bookkeeping can also be generated automatically from the shorthand notation for the rule.

Slide 9.2.29

Note that given the logic rules from the grammar and the facts encoding a sentence, we can use chaining (either forward or backward) to parse the sentence. Let's look at this in more detail.

Parsing

- We can construct a parse tree by modifying the generated rules slightly:
  - `((S (s ?np ?vp) ?s1 ?s3)) :- (NP ?np ?s1 ?s2)(VP ?vp ?s2 ?s3))`
  - This rule now shows how to build the parse tree for the whole sentence (s ?np ?vp) from the parse trees for the constituent NP tree (bound to ?np) and VP tree (bound to ?vp).
- This can all be generated automatically from the same shorthand notation:
  - `((S) :- (NP) (VP))`
- Given a set of rules (and word facts), we can use forward or backward chaining to parse a sentence.

tlp • Spring 02 • 29

Top-Down vs Bottom-Up

- Simple parsers are often classified as top-down or bottom-up where top and bottom refer to the parse tree.
- Backwards chaining on the grammar rules we have seen is a top-down approach to parsing (starts with S and works towards the words).
- Forward chaining on the grammar rules is a bottom up approach (starts with the words and works towards S).

tlp • Spring 02 • 30

Slide 9.2.30

A word on terminology. Parsers are often classified into **top-down** and **bottom-up** depending whether they work from the top of the parse tree down towards the words or vice-versa. Therefore, backward-chaining on the rules leads to a top-down parser, while forward-chaining, which we will see later, leads to a bottom-up parser. There are more sophisticated parsers that are neither purely top-down nor bottom-up, but we will not pursue them here.
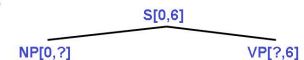
Slide 9.2.31

Let us look at how the sample grammar can be used in a top-down manner (backward-chaining) to parse the sentence "John gave the book to Mary". We start backchaining with the goal S[0,6]. The first relevant rule is the first one and so we generate two subgoals: NP[0,?] and VP[?,6].

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]

NP[0,?]   VP[?,6]

tlp • Spring 02 • 31

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]

NP[0,?]   VP[?,6]

Pronoun[0,?]

John[0,1]

tlp · Spring 02 · 32

Slide 9.2.32

Assuming we examine the rules in order, we first attempt to apply the NP -> Pronoun rule. But that will fail when we actually try to find a pronoun at location 0.

Slide 9.2.33

Then we try to see if NP -> Name will work, which it does, since the first word is John and we have the rule that tells us that John is a Name. Note that this will also bind the end of the VP phrase and the start of the VP to be at position 1.

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]

NP[0,1]   VP[1,6]

Name[0,1]

John[0,1]

tlp · Spring 02 · 33

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]

NP[0,1]   VP[1,6]

Name[0,1]   Verb[1,6]

John[0,1]   gave[1,2] ...

tlp · Spring 02 · 34

Slide 9.2.34

So, we move on to the pending VP. Our first relevant rule is VP -> Verb, which will fail. Note, however, that there is a verb starting at location 1, but at this point we are looking for a verb phrase from positions 1 to 6, while the verb only goes from 1 to 2.

Slide 9.2.35

So, we try the next VP rule, which will look for a verb followed by a noun phrase, spanning from words 1 to 6. The Verb succeeds when we find "gave" in the input.

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]

NP[0,1]   VP[1,6]

Name[0,1]   Verb[1,2]   NP[2,6]

John[0,1]   gave[1,2]

tlp · Spring 02 · 35

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   Verb[1,2]   NP[2,6]
John[0,1]   gave[1,2]   Pronoun[2,6]
✗
the[2,3]

tlp • Spring 02 • 36

Slide 9.2.36

Now we try to find an NP starting at position 2. First we try the pronoun rule, which fails.

Slide 9.2.37

Then we try the name rule, which also fails.

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   Verb[1,2]   NP[2,6]
John[0,1]   gave[1,2]   Name[2,6]
✗
the[2,3]

tlp • Spring 02 • 37

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   Verb[1,2]   NP[2,6]
John[0,1]   gave[1,2]   Art[2,?]   N[?,6]
the[2,3]   book[3,4]

tlp • Spring 02 • 38

Slide 9.2.38

Then we try the article followed by a noun.

Slide 9.2.39

The article succeeds when we find "the" in the input. Now we try to find a noun spanning words 3 to 6. We have a noun in the input but it only spans one word, so we fail.

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   Verb[1,2]   NP[2,6]
John[0,1]   gave[1,2]   Art[2,3]   N[3,6]
✗
the[2,3]   book[3,4]

tlp • Spring 02 • 39

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   Verb[1,2]   Adj[2,6] ✗
John[0,1]   gave[1,2]   the[2,3]

tlp • Spring 02 • 40

Slide 9.2.40

We eventually fail back to our choice of the VP rule and so we try the next VP rule candidate, involving a Verb followed by an adjective, which also fails.

Slide 9.2.41

The next VP rule, looks for a VP followed by prepositional phrase.

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   VP[1,?]   PP[?,6]
John[0,1]

tlp • Spring 02 • 41

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   VP[1,2]   PP[2,6]
John[0,1]   Verb[1,2]
            gave[1,2]

tlp • Spring 02 • 42

Slide 9.2.42

The first VP succeeds by finding the verb "gave", which now requires us to find a prepositional phrase starting at position 2.

Slide 9.2.43

We proceed to try to find a preposition at position 2 and fail.

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S[0,6]
NP[0,1]   VP[1,6]
Name[0,1]   VP[1,2]   PP[2,6]
John[0,1]   Verb[1,2]   P[2,?] ✗   NP[?,6]
            gave[1,2]   the[2,3]

tlp • Spring 02 • 43

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

tlp • Spring 02 • 44

Slide 9.2.44

We fail back to trying an alternative rule (verb followed by NP) for the embedded VP, which now successfully parses "gave the book" and we proceed to look for a prepositional phrase in the range 4 to 6.

Slide 9.2.45

Which successfully parses, "to Mary", and the complete parse succeeds.

Top Down Parsing
John gave the book to Mary

- S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

tlp • Spring 02 • 45

Problems with Top Down Parsing

- Generates sub-trees without checking the input
  - **NP → Pronoun** is tested when input is John, etc.
- Left-recursive rules lead to infinite loops
  - **NP → NP PP**
  - When looking for an **NP** where there isn't one, this rule will loop forever, generating new **NP** sub-goals.
  - Grammar needs to be rewritten to avoid these rules.
- Repeated parsing of sub-trees (after failure and backup)
  - In our simple example, **VP → Verb → gave** is parsed 3 times.
  - If we store intermediate results in fact database, can save some of this work.

tlp • Spring 02 • 46

Slide 9.2.46

There are a number of problems with this top-down parsing strategy. One that substantially impacts efficiency is that rules are chosen without checking whether the next word in the input can possibly be compatible with that rule. There are simple extensions to the top-down strategy to overcome this difficulty (by keeping a table of constituent types and the lexical categories that can begin them).

A more substantial problem, is that rules such as NP -> NP PP (left-branching rules) will cause an infinite loop for this simple top-down parsing strategy. It is possible to modify the grammar to turn such rules into right-branching rules - but that may not be the natural interpretation.

Note that the top-down strategy is carrying out a search for a correct parse and it ends up doing wasted work, repeatedly parsing parts of the sentence during its attempts. This can be avoided by building a table of parses that have been previously discovered (stored in the fact database) so they can be reused rather than re-discovered.

Slide 9.2.47

So far we have been using our rules together with our backchaining algorithm for logic programming to do top-down parsing. But, that's not the only way we can use the rules.

An alternative strategy starts by identifying any rules for which all the literals in their right hand side can be unified (with a single unifier) to the known facts. These rules are said to be **triggered**. For each of those triggered rules, we can add a new fact for the left hand side (with the appropriate variable substitution). Then, we repeat the process. This is known as **forward chaining** and corresponds to bottom-up parsing, as we will see next.

Forward Chaining

- Identify those rules whose antecedents (rhs) can be unified with the ground facts in the database.
- These rules are said to be triggered.
- Don't trigger rules that would not add new facts to the database. This avoids trivial infinite loops.
- For each triggered rule, apply the substitution to the consequent (lhs) of the rule and add the resulting literal to database.
- Repeat until no rule is triggered.

tlp • Spring 02 • 47

Bottom Up Parsing

1. S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

John gave the book to Mary
0   1   2   3   4   5   6

tlp • Spring 02 • 48

Slide 9.2.48

Now, let's look at bottom-up parsing. We start with the facts indicating the positions of the words in the input, shown here graphically.

Slide 9.2.49

Note that all the rules indicating the lexical categories of the individual words, such as Name, Verb, etc, all trigger and can all be run to add the new facts shown here. Note that book is ambiguous, both a noun and a verb, and both facts are added.

Bottom Up Parsing

1. S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

Name Verb Art N V Prep Name
John gave the book to Mary
0   1   2   3   4   5   6

tlp • Spring 02 • 49

Bottom Up Parsing

1. S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

NP   VP   NP   NP
Name Verb Art N V Prep Name
John gave the book to Mary
0   1   2   3   4   5   6

tlp • Spring 02 • 50

Slide 9.2.50

Now these three rules (NP -> Name, VP-> Verb and NP -> Art N) all trigger and can be run.

Slide 9.2.51

Then, another three rules (S -> NP VP, VP -> Verb NP and PP -> Prep NP) trigger and can be run. Note that we now have an S fact, but it does not span the whole input.

Bottom Up Parsing

1. S → NP VP
- S → S Conjunction S
- NP → Pronoun
- NP → Name
- NP → Article Noun
- NP → Number
- NP → NP PP
- NP → NP RelClause
- VP → Verb
- VP → Verb NP
- VP → Verb Adj
- VP → VP PP
- PP → Prep NP
- RelClause → that VP

S   VP   PP
NP   VP   NP   NP
Name Verb Art N V Prep Name
John gave the book to Mary
0   1   2   3   4   5   6

tlp • Spring 02 • 51

Slide 9.2.52

Now, we trigger and run the S rule again as well as the VP->VP PP rule.

Slide 9.2.53

Finally, we run the S rule covering the whole input and we can stop.

Slide 9.2.54

Note that (not surprisingly) we generated some facts that did not make it into our final structure.

Slide 9.2.55

Bottom-up parsing, like top-down parsing, generates wasted work in that it generates structures that cannot be extended to the final sentence structure. Note, however, that bottom-up parsing has no difficulty with left-branching rules, as top-down parsing did. Of course, rules with an empty right hand side can always be used, but this is not a fundamental problem if we require that triggering requires that a rule adds a new fact. In fact, by adding all the intermediate facts to the data base, we avoid some of the potential wasted work of a pure search-based bottom-up parser.

Bottom Up Parsing

- Generates sub-trees that cannot be extended to S, for example, the interpretation of book as a verb in our example.
- No problem with left recursion, but potential problems with empty right hand side (empty antecedent).
- Saving all the facts makes this more efficient than a pure search-based bottom-up parser – does not have to redo sub-trees on failure.

Ambiguity

- A major problem in context-free parsing is ambiguity.
- Lexical class ambiguity: book is Noun and Verb
- Attachment ambiguity:
  - S → NP VP
  - NP → NP PP
  - VP → VP PP
  - VP → Verb NP
  - Mary (((saw (John)) (on the hill)) (with a telescope))
  - Mary ((saw (John)) (on the (hill with a telescope)))
  - Mary ((saw (John on the hill))) (with a telescope))
  - Mary (saw ((John on the hill)) (with a telescope)))
  - Mary (saw (John on the (hill with a telescope))))
- Generate all parses and discard semantically inconsistent ones
- Preferences during parsing

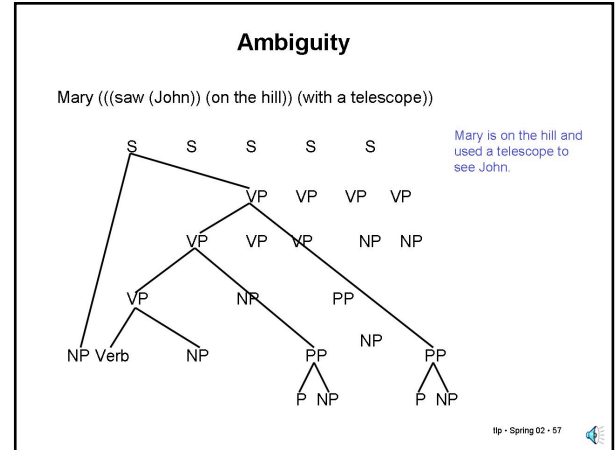tlp · Spring 02 · 56

Slide 9.2.56

One of the key facts of natural language grammars is the presence of ambiguity of many types. We have already seen one simple example of **lexical ambiguity**, the fact that the word book is both a noun and a verb. There are many classic examples of this phenomenon, such as "Time flies like an arrow", where all of "time", "flies" and "like" are ambiguous lexical items. If you can't see the ambiguity, think about "time flies" as analogous to "fruit flies".

Perhaps a more troublesome form of ambiguity is known as **attachment ambiguity**. Consider the simple grammar shown here that allows prepositional phrases to attach both to VPs and NPs. So, the sentence "Mary saw John on the hill with a telescope" has five different structurally different parses, each with a somewhat different meaning (we'll look at them more carefully in a minute).

Basically we have two choices. One is to generate all the legal parses and let subsequent phases of the analysis sort them out or somehow to select one - possibly based on learned preferences based on examples. We will assume that we simply generate all legal parses.

Slide 9.2.57

Here are the various interpretations of our ambiguous sentence. In this one, both prepositional phrases are modifying the verb phrase. Thus, Mary is on the hill she used a telescope to see John.

Ambiguity

Mary (((saw (John)) (on the hill)) (with a telescope))

Mary is on the hill and used a telescope to see John.

tlp · Spring 02 · 57

Ambiguity

Mary ((saw (John)) (on the (hill (with a telescope))))

Mary is on the hill that has a telescope and she saw John.

tlp · Spring 02 · 58
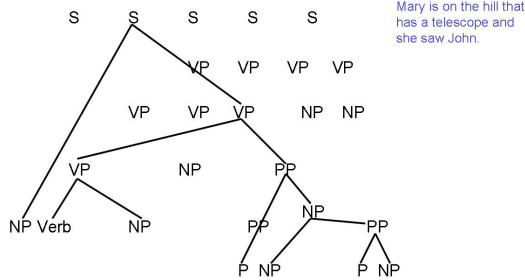
Slide 9.2.58

In this one, the telescope phrase has attached to the hill NP and so we are talking about a hill with a telescope. This whole phrase is modifying the verb phrase. Thus Mary is on the hill that has a telescope when she saw John.

Slide 9.2.59

In this one, the hill phrase is attached to John; this is clearer if you replace John with "the fool", so now Mary saw "the fool on the hill". She used a telescope for this, since that phrase is attached to the VP.

Ambiguity

Mary ((saw (John (on the hill))) (with a telescope))

Mary used a telescope to see John; John is on the hill.

tlp · Spring 02 · 59
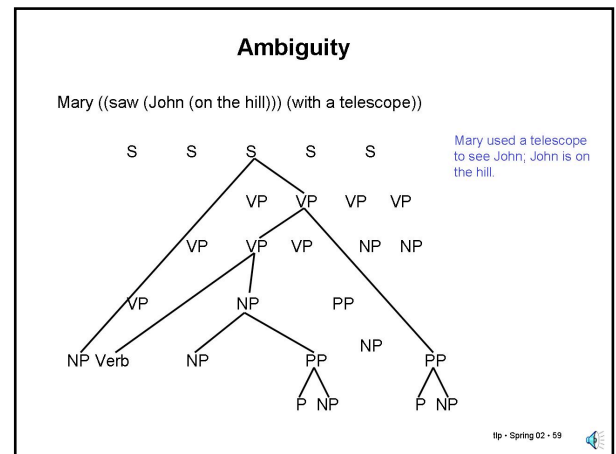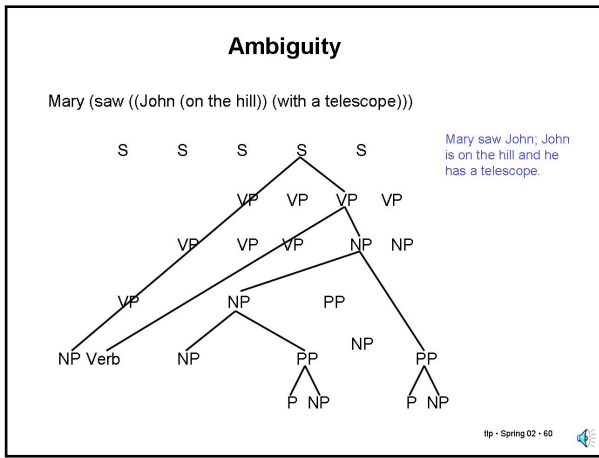
Slide 9.2.60

In this one, its the fool who is on the hill and who has the telescope that Mary saw.

Ambiguity

Mary (saw ((John (on the hill)) (with a telescope)))

Mary saw John; John is on the hill and he has a telescope.

tlp • Spring 02 • 60

Slide 9.2.61

Now its the fool who is on that hill with the telescope on it that Mary saw.

Note that the number of parses grows exponentially with the number of ambiguous prepositional phrases. This is a difficulty that only detailed knowledge of meaning and common usage can resolve.

Ambiguity

Mary (saw (John (on the (hill (with a telescope)))))

Mary saw John; John is on the hill that has a telescope.

tlp • Spring 02 • 61
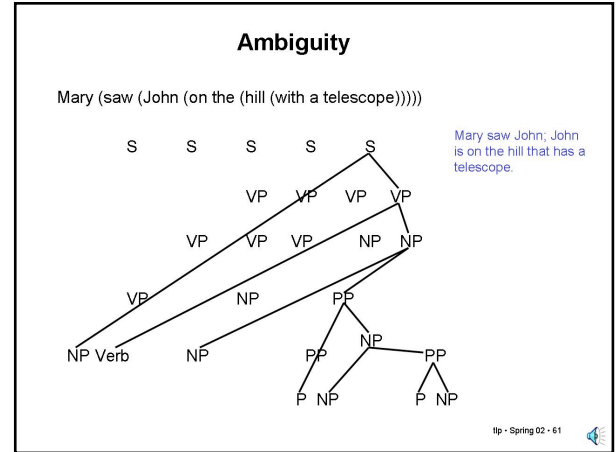
6.034 Notes: Section 9.3

Slide 9.3.1

Now, we move to consider the semantics phase of processing natural language.

6.034 Artificial Intelligence

- Natural Language Understanding
  - Getting at the meaning of text and speech
  - Not just pattern matching
- Overview
- Syntax
- Semantics

tlp • Spring 02 • 1

NLU Architecture

Input/Output data	Processing stage	Other data used
Frequency spectrogram	Speech Recognition	← Sound frequencies
Word sequence		
"He gave Mary..."	Syntactic Analysis	← Grammar
Sentence sstructure		
He gave Mary ...	Semantics Analysis	← Word meanings
Partial meaning 9x give(x,book,mary) Æ ...	Pragmatics	← Context of utterance
Sentence meaning give(john1,book2,Mary1) Æ ...		

tlp • Spring 02 • 2

Slide 9.3.2

Recall that our goal is to take in the parse trees produced by syntactic analysis and produce a meaning representation.
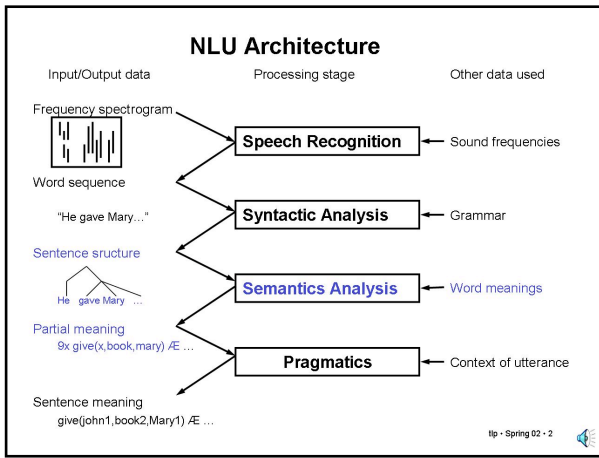
Slide 9.3.3

We want semantics to produce a representation that is somewhat independent of syntax. So, for example, we would like the equivalent active and passive voice versions of a sentence to produce equivalent semantics representations.

We will assume that the meaning representation is some variant of first order predicate logic. We will specify what type of variant later.

We have limited the scope of the role of semantics by ruling out context. So, for example, given the sentence "He gave her the book", we will be happy with indicating that some male gave the book to some female, without identifying who these people might be.

Semantics

- Represent meaning independent of surface syntax
  - He gave the book to her
  - The book was given to her by him
- Usually some variant of predicate calculus is used to represent meaning
- Does not represent context

> Give:
> agent: ?x (male)
> object: book
> recipient: ?y (female)

tlp • Spring 02 • 3

Semantics

- Represent meaning independent of surface syntax
  - He gave the book to her
  - The book was given to her by him
- Usually some variant of predicate calculus is used to represent meaning
- Does not represent context

> Give:
> agent: ?x (male)
> object: book
> recipient: ?y (female)

Pragmatics → John Smith, Mary Jones, Remembrance of things past

tlp • Spring 02 • 4

Slide 9.3.4

Part of the role of pragmatics, the next phase of processing, is to try to make those connections.

Slide 9.3.5

So, let's consider a very simple sentence "John hit Harry". We have here the simple parse tree. What should we expect the semantic representation to be?
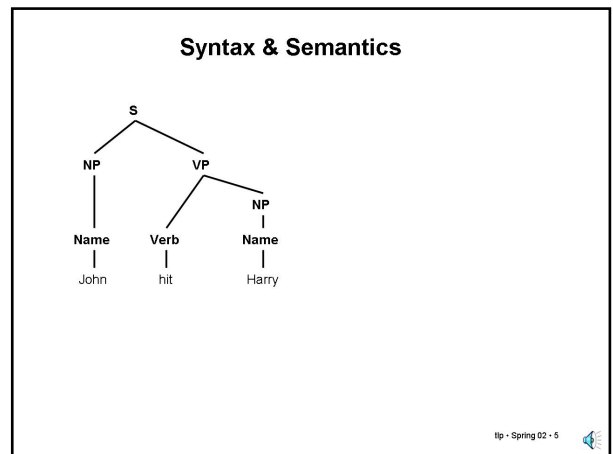
Syntax & Semantics

```
        S
       / \
      NP   VP
      |   /  \
    Name Verb  NP
      |   |     |
    John  hit  Name
                |
              Harry
```

tlp • Spring 02 • 5

Syntax & Semantics

S
NP   VP
Verb
NP

Name
John

Verb
hit

Name
Harry

(hit John Harry)
(hit ?x ?y)
John   Harry

Note, the semantics tree is not parallel in structure to the syntax tree.

Note: we will be using Lisp-like notation for logic throughout this section.

tlp • Spring 02 • 6

Slide 9.3.6

In this simple case, we might want something like this, where hit is a predicate and John and Harry are constant terms in the logical language. The key thing to notice is that even for this simple sentence the semantic structure produced is not perfectly parallel to the syntactic structure.

In this interpretation, the meaning of the verb is the center of the semantics. The meaning representation of the subject NP is embedded in the meaning representation of the verb phrase. This suggests that producing the semantics will not be a trivial variant of the parse tree. So, let's see how we can achieve this.

Slide 9.3.7

Our guiding principle will be that the semantics of a constituent can be constructed by composing the semantics of its constituents. However, the composition will be a bit subtle and we will be using feature values to carry it out.

Let's look at the sentence rule. We will be exploiting the "two way" matching properties of unification strongly here. This rule says that the meaning of the sentence is picked up from the meaning of the VP, since the second argument of the VP is the same as the semantics of the sentence as a whole. We already saw this in our simple example, so it comes as no surprise. Note also that the semantics of the subject NP is passed as the first argument of the VP (by using the same variable name).

Compositional Semantics

- The semantics of a constituent can be constructed by composing the semantics of its constituents.
  - `(S ?pred) :- (NP ?subj) (VP ?subj ?pred)`
    - The semantics of the subject noun phrase is ?subj, which is combined with the semantics of the verb phrase to produce the sentence semantics, ?pred.
  - `(VP ?subj ?pred) :- (Verb ?subj ?obj ?pred) (NP ?obj)`
  - `(NP ?sem) :- (Name ?sem)`

tlp • Spring 02 • 7

Compositional Semantics

- The semantics of a constituent can be constructed by composing the semantics of its constituents.
  - `(S ?pred) :- (NP ?subj) (VP ?subj ?pred)`
    - The semantics of the subject noun phrase is ?subj, which is combined with the semantics of the verb phrase to produce the sentence semantics, ?pred.
  - `(VP ?subj ?pred) :- (Verb ?subj ?obj ?pred) (NP ?obj)`
    - This rule is for a transitive verb that expects a single direct object noun phrase, whose semantics are ?obj.
    - The semantics of the VP will be constructed from the semantics of the verb, which will combine the semantics of the subject ?subj and the direct object ?obj to produce the VP semantics, ?pred.
  - `(NP ?sem) :- (Name ?sem)`
    - This rule is for proper names and the semantics of the NP is just that of the name.

tlp • Spring 02 • 8

Slide 9.3.8

The VP has two arguments, the semantics of the subject NP (which will be an input) and the resulting semantics of the VP. In the VP rule, we see that the result semantics is coming from the Verb, which is combining the semantics of the subject and the object NPs to produce the result for the VP (and ultimately the sentence).

Slide 9.3.9

Let's look at the rule for a particular Verb. Note that the first two arguments are simply variables which are then included in the expression for the verb semantics, the predicate hit with two arguments (the subject and the object).

Compositional Semantics

- The semantics of a constituent can be constructed by composing the semantics of its constituents.
  - `(S ?pred) :- (NP ?subj) (VP ?subj ?pred)`
  - `(VP ?subj ?pred) :- (Verb ?subj ?obj ?pred) (NP ?obj)`
  - `(NP ?sem) :- (Name ?sem)`
- The semantics of individual words are given in the lexicon.
  - `(Verb ?x ?y (hit ?x ?y)) :- hit`
    - The verb semantics for hit. Note that the subject will match ?x and the direct object will match ?y and the final semantics will be (hit ?x ?y)
  - `(Name John) :- John`
  - `(Name Harry) :- Harry`
    - Trivial semantics

tlp • Spring 02 • 9

Compositional Semantics

- The semantics of a constituent can be constructed by composing the semantics of its constituents.
  - `(S ?pred) :- (NP ?subj) (VP ?subj ?pred)`
  - `(VP ?subj ?pred) :- (Verb ?subj ?obj ?pred) (NP ?obj)`
  - `(NP ?sem) :- (Name ?sem)`
- The semantics of individual words are given in the lexicon.
  - `(Verb ?x ?y (hit ?x ?y)) :- hit`
  - `(Name John) :- John`
  - `(Name Harry) :- Harry`
- The sentence: "John hit Harry"
  - `(backchain '(S ?sem 0 3))`
  - `?sem = (hit John Harry)`
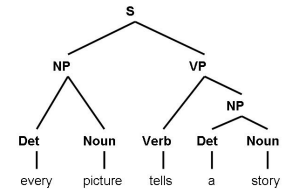
tlp · Spring 02 · 10

Slide 9.3.10

We can pull this altogether by simply calling backchain with the goal pattern for a successful parse. We will want to retrieve the value of the binding for ?sem, which is the semantics for the sentence.

Slide 9.3.11

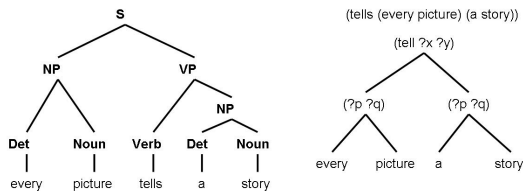Let's look at a somewhat more complex example - "Every picture tells a story". Here is the syntactic analysis.

Syntax & Semantics

tlp · Spring 02 · 11

Syntax & Semantics

Note, the semantics tree is not parallel in structure to the syntax tree.

tlp · Spring 02 · 12

Slide 9.3.12

This is one possible semantic analysis. Note that it follows the pattern of our earlier example. The top-level predicate is derived from the verb and it includes as arguments the semantics of the subject and direct object.

Slide 9.3.13

The only innovation in this grammar, besides the new words is a simple semantics for a noun phrase formed from a Determiner and a Noun - just placing them in a list. We can interpret the result as a quantifier operating on a predicate. But, what does this mean? It's certainly not legal logic notation.

Another Example

- The grammar
  - `(S ?pred) :- (NP ?subj) (VP ?subj ?pred)`
  - `(VP ?subj ?pred) :- (Verb ?subj ?obj ?pred) (NP ?obj)`
  - `(NP ?sem) :- (Name ?sem)`
  - `(NP (?detsem ?nsem)) :- (Det ?detsem) (Noun ?nsem)`
  - `(Verb ?x ?y (tells ?x ?y)) :- tells`
  - `(Noun picture) :- picture`
  - `(Noun story) :- story`
  - `(Det every) :- every`
  - `(Det a) :- a`
- The sentence: "Every picture tells a story"
  - `(backchain '(S ?sem 0 5))`
  - `?sem = (tell (every picture) (a story))`

tlp · Spring 02 · 13

Quantifiers

- (tell (every picture) (a story)) is ambiguous:
  - $\forall$ x Picture(x) $\rightarrow$ $\exists$ y Story(y) $\wedge$ Tell(x,y)
  - $\exists$ y Story(y) $\wedge$ $\forall$ x Picture(x) $\rightarrow$ Tell(x,y)
- The first of these is the usual interpretation, but consider:
  - Every US citizen has a president

tlp · Spring 02 · 14

Slide 9.3.14

Furthermore, even if we are generous and consider this a legal quantified expression, then it's ambiguous - in the usual sense that "Every man loves a woman" is ambitious. That is, is there one story per picture or do all the pictures tell the same story.
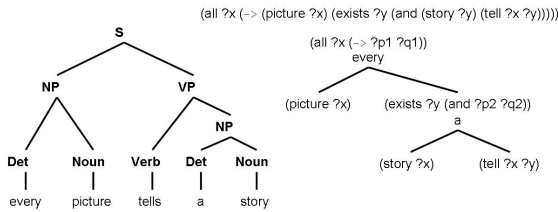
Slide 9.3.15

Let's pick one of the interpretations and see how we could generate it. At the heart of this attempt is a definition of the meaning of the determiners "every" and "a", which now become patterns for universally and existentially quantified statements. Note also that the nouns become patterns for predicate expressions.

Quantifiers

- (tell (every picture) (a story)) is ambiguous:
  - $\forall$ x Picture(x) $\rightarrow$ $\exists$ y Story(y) $\wedge$ Tell(x,y)
  - $\exists$ y Story(y) $\wedge$ $\forall$ x Picture(x) $\rightarrow$ Tell(x,y)
- The first of these is the usual interpretation, but consider:
  - Every US citizen has a president
- Let's consider how we could generate:
  - $\forall$ x Picture(x) $\rightarrow$ $\exists$ y Story(y) $\wedge$ Tell(x,y)
  - (all ?x (-> (picture ?x) (exists ?y (and (story ?y) (tell ?x ?y)))))
    - every = (all ?x (-> ?p1 ?q1))
    - picture = (picture ?x)
    - tells = (tell ?x ?y)
    - a = (exists ?y (and ?p2 ?q2))
    - story = (story ?x)

tlp · Spring 02 · 15

Syntax & Semantics

(all ?x (-> (picture ?x) (exists ?y (and (story ?y) (tell ?x ?y)))))

```
                S
             /     \
           NP       VP
          /  \      |  \
                   Verb  NP
                       /    \
        Det  Noun  Verb Det  Noun
         |    |     |    |    |
       every picture tells a story
```

(all ?x (-> ?p1 ?q1))
every

(picture ?x)   (exists ?y (and ?p2 ?q2))
                              a
                          /   \
               (story ?x)   (tell ?x ?y)

Note, the semantics tree is not parallel in structure to the syntax tree.

tlp · Spring 02 · 16

Slide 9.3.16

Our target semantic representation is shown here. Note that by requiring the semantics to be a legal logical sentence, we've had to switch the key role from the verb to the determiner. That is, the top node in the sentence semantics comes from the determiner, not the verb. The semantics of the verb is fairly deeply nested in the final semantics - but it still needs to combine the semantics of the subject and direct object NPs. Note, however, that it is incorporating them by using the quantified variable introduced by the determiners of the subject and object NPs.

Slide 9.3.17

Let's start with the definitions of the words. Here's the definition for the verb "tells". We have seen this before. It combines the semantics of the subject NP (bound to ?x) and the semantics of the object NP (bound to ?y) with the predicate representing the verb to produce the VP semantics.

Quantifiers

- **(Verb ?x ?y (tell ?x ?y)) :- tells**
  - ?x denotes the subject and ?y the direct object, the resulting semantics is (tell ?x ?y).

tlp · Spring 02 · 17

Quantifiers

- `(Verb ?x ?y (tell ?x ?y)) :- tells`
  - ?x denotes the subject and ?y the direct object, the resulting semantics is (tell ?x ?y).
- `(Noun ?x (picture ?x)) :- picture`
- `(Noun ?x (story ?x)) :- story`
  - ?x will typically be a variable, which we restrict to denote a picture or a story or (and (young ?x) (male ?x)) for boys, etc.

tlp · Spring 02 · 18

Slide 9.3.18

The nouns will be denoted by one of the quantified variables introduced by the quantifiers. The noun places a restriction on the entities that the variable can refer to. In this definition, the quantified variable will be bound to ?x and incorporated into the predicate representing the noun.

Slide 9.3.19

Finally, the determiners are represented by quantified formulas that combine the semantics derived from the noun with the semantics of the VP (for a subject NP) or of the Verb (for an object NP).

Quantifiers

- `(Verb ?x ?y (tell ?x ?y)) :- tells`
  - ?x denotes the subject and ?y the direct object, the resulting semantics is (tell ?x ?y).
- `(Noun ?x (picture ?x)) :- picture`
- `(Noun ?x (story ?x)) :- story`
  - ?x will typically be a variable, which we restrict to denote a picture or a story or (and (young ?x) (male ?x)) for boys, etc.
- `(Det ?x ?p ?q (all ?x (-> ?p ?q))) :- every`
- `(Det ?x ?p ?q (exists ?x (and ?p ?q))) :- a`
  - The ?x is the formal variable, ?p denotes the semantics of the noun and ?q the semantics of the predicate. For a subject NP, the predicate comes from the VP of the sentence. For an object NP, the predicate comes from the verb.

tlp · Spring 02 · 19

Quantifiers

- `(S ?sent) :- (NP ?x ?vp ?sent) (VP ?x ?vp)`
  - The semantics of the sentence will be derived from the NP, since the determiner provides the quantifier, which is the top node in the semantics.
  - ?x will be the "formal variable" for the quantifier, e.g. (all ?x ...)
- `(VP ?xs ?vp) :- (Verb ?xs ?xo ?verb) (NP ?xo ?verb ?vp)`
- `(NP ?x ?p ?np) :- (Det ?x ?noun ?p ?np) (Noun ?x ?noun)`
- `(Verb ?x ?y (tell ?x ?y)) :- tells`
- `(Noun ?x (picture ?x)) :- picture`
- `(Noun ?x (story ?x)) :- story`
- `(Det ?x ?p ?q (all ?x (-> ?p ?q))) :- every`
- `(Det ?x ?p ?q (exists ?x (and ?p ?q))) :- a`

tlp · Spring 02 · 20

Slide 9.3.20

The new sentence (S) rule reflects the difference in where the top-level semantics is being assembled. Before, we passed the semantics of the subject NP into the VP, now we go the other way. The semantics of the VP is an argument to the subject NP.

Note that the variable ?x here will not be bound to anything, it is the variable that will be used as the quantified variable by the determiner's semantics.

Slide 9.3.21

The VP rule is analogous. The semantics of the Verb will combine a reference to the subject and object semantics (through their corresponding quantified variables) and the resulting semantics of the Verb will be combined into the semantics of the NP (which will ultimately be derived from the semantics of the determiner).

Quantifiers

- `(S ?sent) :- (NP ?x ?vp ?sent) (VP ?x ?vp)`
  - The semantics of the sentence will be derived from the NP, since the determiner provides the quantifier, which is the top node in the semantics.
  - ?x will be the "formal variable" for the quantifier, e.g. (all ?x ...)
- `(VP ?xs ?vp) :- (Verb ?xs ?xo ?verb) (NP ?xo ?verb ?vp)`
  - Similarly, the semantics of the VP (?vp) will be derived from that of the direct object NP, e.g. (exists ?y (and (story ?y) (tell ?x ?y)))
  - Note that ?xs will be formal variable from subject NP and ?xo will be the variable from the object NP, they will be combined to form the Verb semantics (tell ?xs ?xo).
- `(NP ?x ?p ?np) :- (Det ?x ?noun ?p ?np) (Noun ?x ?noun)`
- `(Verb ?x ?y (tell ?x ?y)) :- tells`
- `(Noun ?x (picture ?x)) :- picture`
- `(Noun ?x (story ?x)) :- story`
- `(Det ?x ?p ?q (all ?x (-> ?p ?q))) :- every`
- `(Det ?x ?p ?q (exists ?x (and ?p ?q))) :- a`

tlp · Spring 02 · 21

Quantifiers

- `(S ?sent) :- (NP ?x ?vp ?sent) (VP ?x ?vp)`
  - The semantics of the sentence will be derived from the NP, since the determiner provides the quantifier, which is the top node in the semantics.
  - ?x will be the "formal variable" for the quantifier, e.g. (all ?x ...)
- `(VP ?xs ?vp) :- (Verb ?xs ?xo ?verb) (NP ?xo ?verb ?vp)`
  - Similarly, the semantics of the VP (?vp) will be derived from that of the direct object NP, e.g. (exists ?y (and (story ?y) (tell ?x ?y)))
  - Note that ?xs will be formal variable from subject NP and ?xo will be the variable from the object NP, they will be combined to form the Verb semantics (tell ?xs ?xo).
- `(NP ?x ?p ?np) :- (Det ?x ?noun ?p ?np) (Noun ?x ?noun)`
  - The semantics of the NP is produced by the determiner, which incorporates the semantics of the noun and that of the predicate.
- `(Verb ?x ?y (tell ?x ?y)) :- tells`
- `(Noun ?x (picture ?x)) :- picture`
- `(Noun ?x (story ?x)) :- story`
- `(Det ?x ?p ?q (all ?x (-> ?p ?q))) :- every`
- `(Det ?x ?p ?q (exists ?x (and ?p ?q))) :- a`

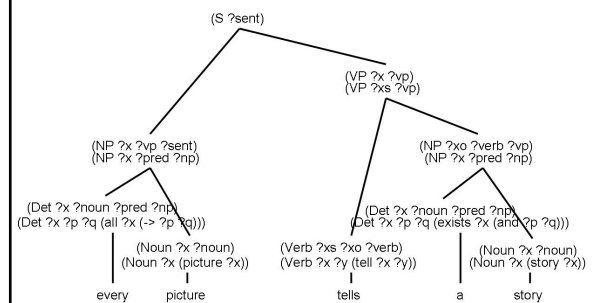tlp · Spring 02 · 22

Slide 9.3.22

The NP rule in fact takes ?p, which will be the semantics of the Verb phrase and combine them with the semantics of the noun in the semantics of the Determiner.

Slide 9.3.23

Here we see how the parse works out. You have to follow the bindings carefully to see how it all works out.

What is remarkable about this is that we were able to map from a set of words to a first-order logic representation (which does not appear to be very similar) with a relatively compact grammar and with quite generic mechanisms.

Parsing with Quantifiers

tlp · Spring 02 · 23

Quasi-Logical Form

- Semantics tries to capture sentence meaning independent of context. Producing the correct representation in First Order Logic usually requires context to resolve the ambiguity in language:
  - Syntactic ambiguity: "Mary saw John on the hill with a telescope"
  - Lexical ambiguity: "We went to the bank" {of the river? Fleet Bank?}
  - Quantifier scope ambiguity: "Every man loves a woman"
  - Referential ambiguity: "He gave her the book", "Stop that!"

tlp · Spring 02 · 24

Slide 9.3.24

The quantified expression we produced in the previous example is unambiguous, as required to be able to write an expression in first order logic. However, natural language is far from unambiguous. We have seen examples of syntactic ambiguity, lexical and attachment ambiguity in particular, plus there are many examples of semantic ambiguity, for example, ambiguity in quantifier scope and ambiguity on who or what pronouns refer to are examples.

Slide 9.3.25

One common approach to semantics is to have it produce a representation that is not quite the usual logical notation, sometimes called **quasi-logical form**, that preserves some of the ambiguity in the input, leaving it to the pragmatics phase to resolve the ambiguities employing contextual information.

Quasi-Logical Form

- Semantics tries to capture sentence meaning independent of context. Producing the correct representation in First Order Logic usually requires context to resolve the ambiguity in language:
  - Syntactic ambiguity: "Mary saw John on the hill with a telescope"
  - Lexical ambiguity: "We went to the bank" {of the river? Fleet Bank?}
  - Quantifier scope ambiguity: "Every man loves a woman"
  - Referential ambiguity: "He gave her the book", "Stop that!"
- Instead of producing FOL, produce quasi-logical form that preserves some of the ambiguity. Leave it for next phase to resolve the ambiguity.
  - (tell (every ?x (picture ?x)) (exists ?x (story ?x)))

tlp · Spring 02 · 25

Quasi-Logical Form

- Allow the use of quantified terms such as
  - (every ?x (picture ?x))
  - (exists ?x (story ?x))

tlp · Spring 02 · 26

Slide 9.3.26

One common aspect of quasi-logical notation is the use of **quantified terms**. These terms indicate the nature of the intended quantification but do not specify the scope of the quantifier in the sentence and thus preserves the ambiguity in the natural language. Note that we are treating these quantified expressions as terms, and using them as arguments to functions and predicates - which is not legal FOL.

Slide 9.3.27

In quasi-logical notation, one also typically extends the range of available quantifiers to correspond more closely to the range of determiners available in natural language. One important case, is the determiner "the", which indicates a unique descriptor.

Quasi-Logical Form

- Allow the use of quantified terms such as
  - (every ?x (picture ?x))
  - (exists ?x (story ?x))
- Allow a more general class of quantifiers
  - (the ?x (and (big ?x) (picture ?x) (author ?x "Sargent")))
  - (most ?x (child ?x))
  - (name ?x John)
  - (pronoun ?x he)

tlp · Spring 02 · 27

Quasi-Logical Form

- Allow the use of quantified terms such as
  - (every ?x (picture ?x))
  - (exists ?x (story ?x))
- Allow a more general class of quantifiers
  - (the ?x (and (big ?x) (picture ?x) (author ?x "Sargent")))
  - (most ?x (child ?x))
  - (name ?x John)
  - (pronoun ?x he)
- These will have to be converted to FOL and given an appropriate axiomatization.

tlp · Spring 02 · 28

Slide 9.3.28

These quantified terms and generalized quantifiers will require conversion to standard FOL together with a careful axiomatization of their intended meaning before the resulting semantics can be used for inference.

Slide 9.3.29

Let's illustrate how the type of language processing we have been discussing here could be used to build an extremely simple database system. We'll assume that we want to deal with a simple genealogy domain. We will have facts in our database describing the family relationships between some set of people. We will not restrict ourselves to just the minimal set of facts, such as parent, male and female, we will also keep derived relationships such as grandparent and cousin.

A very simple language system
The Database

- Genealogy database
  - (parent x y), (male x), (female x)
  - (grandparent x y), (aunt/uncle x y), (sibling x y), (cousin x y)

tlp · Spring 02 · 29

A very simple language system
The Database

- Genealogy database
  - `(parent x y)`, `(male x)`, `(female x)`
  - `(grandparent x y)`, `(aunt/uncle x y)`, `(sibling x y)`, `(cousin x y)`
- Assume relations explicit in database.
- Use forward-chaining of rules to expand relations when new facts added.

tlp • Spring 02 • 30

Slide 9.3.30

In fact, we will assume that all the relationships between people we know about are explicitly in the database. We can accomplish them by running a set of Prolog-like rules in forward chaining fashion whenever a new fact is added. We do this, rather than do deduction at retrieval time because of issues of equality, which we will discuss momentarily.

Slide 9.3.31

We will also allow assertions of the form (is x y) which indicate that two symbols denote the same person. We will assume that the forward chaining rules will propagate this equality to all the relevant facts. That is, we substitute equals for equals in each predicate, explicitly. This is not efficient, but it is simple.

A very simple language system
The Database

- Genealogy database
  - `(parent x y)`, `(male x)`, `(female x)`
  - `(grandparent x y)`, `(aunt/uncle x y)`, `(sibling x y)`, `(cousin x y)`
- Assume relations explicit in database.
- Use forward-chaining of rules to expand relations when new facts added.
- `(is x y)` indicates two symbols denote same person

tlp • Spring 02 • 31

A very simple language system
The Database

- Genealogy database
  - `(parent x y)`, `(male x)`, `(female x)`
  - `(grandparent x y)`, `(aunt/uncle x y)`, `(sibling x y)`, `(cousin x y)`
- Assume relations explicit in database.
- Use forward-chaining of rules to expand relations when new facts added.
- `(is x y)` indicates two symbols denote same person
- Retrieval query examples:
  - `(and (female ?x) (parent ?x John))`
  - `(and (male ?x) (cousin Mary ?x))`
  - `(grandparent Harry ?x)`

tlp • Spring 02 • 32

Slide 9.3.32

We can now do very simple retrieval from this database of facts using our backchaining algorithm. We initialize the goal stack in backchaining with the query. If the query is a conjunction, we initialize the stack with all the conjuncts.

Slide 9.3.33

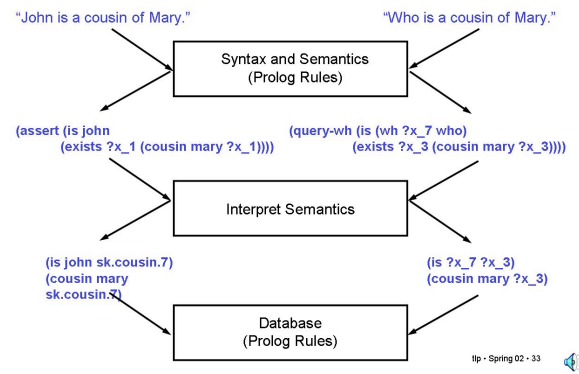Here we see a brief overview of the processing that we will do to interact with the genealogy database.

We will be able to accept declarative sentences, such as "John is a cousin of Mary". These sentences will be processed by a grammar to obtain a semantic representation. This representation will then be interpreted as a set of facts to be added to the database.

We can also ask questions, such as "Who is a cousin of Mary". Our grammar will produce a semantic representation. The semantics of this type of sentence is converted into a database query and passed to the database.

Let's look in more detail at the steps of this process.

A very simple language system
Processing Overview

"John is a cousin of Mary."   "Who is a cousin of Mary."

Syntax and Semantics
(Prolog Rules)

(assert (is john   (query-wh (is (wh ?x_7 who)
(exists ?x_1 (cousin mary ?x_1))))   (exists ?x_3 (cousin mary ?x_3))))

Interpret Semantics

(is john sk.cousin.7)   (is ?x_7 ?x_3)
(cousin mary   (cousin mary ?x_3)
sk.cousin.7)

Database
(Prolog Rules)

tlp • Spring 02 • 33

A very simple language system
The Grammar

- "John is a cousin of Mary."
- `(S (assert ?sem) …) :-`
  `(NP ?subj …)` → signal an assertion
  `(VP ?subj ?sem …)`

- "Is John a cousin of Mary?"
- `(S (query-is (is ?subj ?sem)) …) :-`
  `(is)`
  `(NP ?subj …)`
  `(NP ?sem …)` → signal a query

- "Who is a cousin of Mary?"
- `(S (query-wh ?sem) …) :-`
  `(NP ?subj … )`
  `(VP ?subj …)`

tlp · Spring 02 · 34

Slide 9.3.34

We will need a grammar built along the lines we have been discussing. One of the things the grammar does is classify the sentences into declarative sentences, such as "John is a cousin of Mary", which will cause us to assert a fact in our database, and questions, such as, "Is John a cousin of Mary" or "Who is a cousin of Mary", which will cause us to query the database.

Slide 9.3.35

Here we see one possible semantics for the declarative sentence "John is a cousin of Mary". The operation `assert` indicates the action to be taken. The body is in quasi-logical form; the quantified term `(exists ?x_1 (cousin mary ?x_1))` is basically telling us there exists a person that is in the cousin relationship to Mary. The outermost `is` assertion is saying that John denotes that person. This is basically interpreting this quasi-logical form as:

`] x . (is John x) ^ (cousin Mary x)`

A very simple language system
The Semantics

- "John is a cousin of Mary."
- `(assert (is john`
  `(exists ?x_1 (cousin mary ?x_1))))`

tlp · Spring 02 · 35

A very simple language system
The Semantics

- "John is a cousin of Mary."
- `(assert (is john`
  `(exists ?x_1 (cousin mary ?x_1))))`

- "Is John a cousin of Mary?"
- `(query-is (is john`
  `(exists ?x_5 (cousin mary ?x_5))))`

tlp · Spring 02 · 36

Slide 9.3.36

The semantics of the question "Is John a cousin of Mary?" is essentially identical to that of the declarative form, but it is prefixed by a query operation rather than an assertion. So, we would want to use this to query the database rather than for asserting new facts.

Slide 9.3.37

We can also have a question such as "Who is a cousin of Mary", which is similar except that John is replaced by a term indicating that we are interested in determining the value of this term.

A very simple language system
The Semantics

- "John is a cousin of Mary."
- `(assert (is john`
  `(exists ?x_1 (cousin mary ?x_1))))`

- "Is John a cousin of Mary?"
- `(query-is (is john`
  `(exists ?x_5 (cousin mary ?x_5))))`

- "Who is a cousin of Mary?"
- `(query-wh (is (wh ?x_7 who)`
  `(exists ?x_3 (cousin mary ?x_3))))`

tlp · Spring 02 · 37

A very simple language system
Using the Semantics

• "John is a cousin of Mary."
• `(assert (is john`
        `(exists ?x_1 (cousin mary ?x_1))))`
• Assign skolem constant for ?x_1, e.g. sk.cousin.7
• Convert body of exists into one or more facts
• Replace (exists ?x …) with skolem constant

tlp • Spring 02 • 38

Slide 9.3.38

Given the semantics, we have to actually decide how to add new facts and do the retrieval. Here we show an extremely simple approach that operates for these very simple types of queries (note that we are only using existentially quantified terms).

We are basically going to turn the assertion into a list of ground facts to add to the database. We will do this by skolemizing. Since we have only existentially quantified variables, this will eliminate all variables.

We replace the quantified terms with the corresponding skolem constant and we convert the body of the quantified term into a set of facts that describe the constant.

Slide 9.3.39

In this example, we get two new facts. One is from the outer `is` assertion which tells us that John denotes the same person as the skolem constant. The second fact comes from the body of the quantified term which tells us some properties of the person denote by the skolem constant.

A very simple language system
Using the Semantics

• "John is a cousin of Mary."
• `(assert (is john`
        `(exists ?x_1 (cousin mary ?x_1))))`
• Assign skolem constant for ?x_1, e.g. sk.cousin.7
• Convert body of exists into one or more facts
• Replace (exists ?x …) with skolem constant
• Add to the database:
  • `(is john sk.cousin.7)`
  • `(cousin mary sk.cousin.7)`

tlp • Spring 02 • 39

A very simple language system
Using the Semantics

• "Who is a cousin of Mary?"
• `(query-wh (is (wh ?x_7 who)`
            `(exists ?x_3 (cousin mary ?x_3))))`
• Convert body of exists into one or more additional conditions for query
• Replace (exists ?x …) with ?x
• Replace (wh ?y …) with ?y
• Retrieve from database:
  • `(and (is ?x_7 ?x_3) (cousin mary ?x_3))`
  • `?x_7/John`
  • `?x_3/sk.cousin.7`

tlp • Spring 02 • 40

Slide 9.3.40

We process the question in a similar way except that instead of using skolem constants we keep the variables, since we want those to match the constants in the database. When we perform the query, ?x_7 is bound to John as expected. In general, there may be multiple matches for the query, some may be skolem constants and some may be people names. We would want to return the specific names whenever possible.

Slide 9.3.41

Here are some examples that show that this approach can be used to do a little inference above and beyond what is explicitly stated. Note that the assertions do not mention cousin, uncle, sister or sibling relations, those are inferred. So, we are going beyond what an Internet search engine can do, that is, pattern match on the presence of particular words.

This example has been extremely simple but hopefully it illustrates the flavor of how such a system may be built using the tools we have been developing and what could be done with such a system.

Some Examples

• Assertions
    • John is the father of Tom.
    • Mary is the female parent of Tom.
    • Bill is the brother of John.
    • Jim is the male child of Bill.
    • Jane is the daughter of John.
    • Mary is the mother of Jane.
• Questions
    • Is Jim the cousin of Tom?) Yes
    • Who is the uncle of Tom?) Bill
    • Is Bill the uncle of Tom?) Yes
    • Is Jane the sister of Tom?) Yes
    • Who is a child of Mary?) Tom
    • Who is a sibling of Tom?) Jane

tlp • Spring 02 • 41

Discourse Context

- Anaphora = "use of a word referring to or replacing earlier words"
  - Jack lost his book. He looked for it for hours. Eventually he found it in his backpack.

tlp • Spring 02 • 42

Slide 9.3.42

At this point, we'll touch briefly on a set of phenomena that are beyond the scope of pure semantics because they start dealing with the issue of context.

One general class of language phenomena is called **anaphora**. this includes pronoun use, where a word is used to refer to other words appearing either elsewhere in the sentence or in another sentence.

Slide 9.3.43

Another phenomenon is called **ellipsis**, when words or phrases are missing and need to be filled in from context. In this example, the phrase "complete the job" is missing from the enf of the second conjoined sentence.

Discourse Context

- Anaphora = "use of a word referring to or replacing earlier words"
  - Jack lost his book. He looked for it for hours. Eventually he found it in his backpack.
- Ellipsis = "omission from a sentence of words needed to complete construction of meaning"
  - You did not complete the job as well as he did.

tlp • Spring 02 • 43

Discourse Context

- Anaphora = "use of a word referring to or replacing earlier words"
  - Jack lost his book. He looked for it for hours. Eventually he found it in his backpack.
- Ellipsis = "omission from a sentence of words needed to complete construction of meaning"
  - You did not complete the job as well as he did.
- Definite descriptions = "used to refer to uniquely identifiable entity (or entities)"
  - the tall man, the red book, the president

tlp • Spring 02 • 44

Slide 9.3.44

Another important mechanism in language is the use of **definite descriptions**, signaled by the determiner "the". The intent is that the listener be able to identify an entity previously mentioned or expected to be known.

All of these are linguistic mechanisms for incorporating context and require that a language understanding system that is engaged in an interaction with a human keep a context and be able to identify entities and actions in context based on the clues in the sentence. This is an area of active research and some systems with competence in this area have been built.

Slide 9.3.45

Even beyond conversational context, understanding human language requires access to the whole range of human knowledge. Even when speaking with a child, one assumes a great deal of "common sense" knowledge that computers are, as yet, sorely lacking in. The problem of language understanding at this point merges with the general problem of knowledge representation and use.

World Knowledge

- John needed money. He went to the bank.
- "bank of the river Charles?" "Fleet Bank?"
- Need to know that Fleet Bank has money but the bank of the Charles does not.

- John went to the store. He bought some bread.
- Did John go to the hardware store?
- Etc.

tlp • Spring 02 • 45

Applications

- Human computer interaction:
  - Restricted domains – flight reservations, classifying e-mails into a few classes, redirecting caller to one of a few destinations.
  - Limited syntax
  - Limited vocabulary
  - Limited context
  - Limited actions
  - It is very hard for humans to understand what the limits of the system are. Can be frustrating.

tlp • Spring 02 • 46

Slide 9.3.46

Real applications of natural language technology for human computer interfaces require a very limited scope so that the computer can get by with limited language skills and can have enough knowledge about the domain to be useful. However, it is difficult to keep people completely within the language and knowledge boundaries of the system. This is why the use of natural language interfaces is still limited.

Slide 9.3.47

There is, however, a rapidly increasing use of limited language processing in tasks that don't involve direct interaction with a human but do require some level of understanding of language. These tasks are characterized by situations where there is value in even limited capabilities, e.g. doing the first draft of a translation or a building a quick summary of a much longer news article.

I expect to see an explosion of applications of natural language technologies in the near future.

Applications

- Human computer interaction:
  - Restricted domains – flight reservations, classifying e-mails into a few classes, redirecting caller to one of a few destinations.
  - Limited syntax
  - Limited vocabulary
  - Limited context
  - Limited actions
  - It is very hard for humans to understand what the limits of the system are. Can be frustrating.
- Summarization, Search, Translation
  - Broader domain
  - Performance does not have to be perfect to be useful

tlp • Spring 02 • 47

Sources

- James Allen, *Natural Language Understanding*, Benjamin/Cummings
- Peter Norvig, *Paradigms of Artificial Intelligence Programming*, Morgan Kauffman
- Slides by Alison Cawsey (www.cee.hw.ac.uk/~alison/nl.html)

tlp • Spring 02 • 48

Slide 9.3.48

Here are some sources that were used in the preparation of these slides and which can serve as additional reading material.