Parsing:
Top-Down vs. Bottom-Up
Parsing Algorithms
Treebanks
Statistical Parsing
Partial Parsing – Chunking
Dependency Parsing

Parsing defined:

- The process of finding a derivation (i. e. sequence of productions) leading from the START symbol to a TERMINAL symbol
  - Shows how a particular sentence *could be* generated by the rules of the grammar
- If sentence is structurally ambiguous, more than one possible derivation is produced
- Can solve both the recognition and analysis problems
  - Is this sentence derived from this grammar?
  - Give the derivation(s) that can derive this sentence.
- Parsing algorithms give a strategy for finding a derivation by making choices among the derivation rules and deciding when the derivation is complete or not.

Top-down Parser

- Hypothesis-driven
  - At each stage, parser hypothesizes a structure, and tests whether data (next word in sentence) fits the hypothesis
- Looks at goal first (S) and then sees which rules can be applied
  - Typically progresses from top-to-bottom, left-to-right
  - Non-deterministic (can be rewritten in more than one way)
- When rules derive lexical elements (words), check with the input to see if the right sentence is being derived
- An algorithm may include a backtracking mechanism
  - When it is determined that the wrong rule has been used, it backs up and tries another rule

Example Grammar

- The flight grammar from the text:

$S \rightarrow NP\ VP$

$S \rightarrow Aux\ NP\ VP$

$S \rightarrow VP$

$NP \rightarrow Pronoun$

$NP \rightarrow Proper\text{-}Noun$

$NP \rightarrow Det\ Nominal$

$Nominal \rightarrow Noun$

$Nominal \rightarrow Nominal\ Noun$

$Nominal \rightarrow Nominal\ PP$

$VP \rightarrow Verb$

$VP \rightarrow Verb\ NP$

$VP \rightarrow Verb\ NP\ PP$

$VP \rightarrow Verb\ PP$

$VP \rightarrow VP\ PP$

$PP \rightarrow Preposition\ NP$

$Det \rightarrow that \mid this \mid a$

$Noun \rightarrow book \mid flight \mid meal \mid money$

$Verb \rightarrow book \mid include \mid prefer$

$Pronoun \rightarrow I \mid she \mid me$

$Proper\text{-}Noun \rightarrow Houston \mid TWA$

$Aux \rightarrow does$

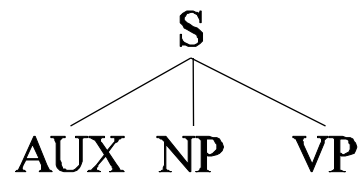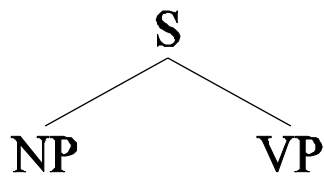$Preposition \rightarrow from \mid to \mid on \mid near \mid through$

Example Derivation
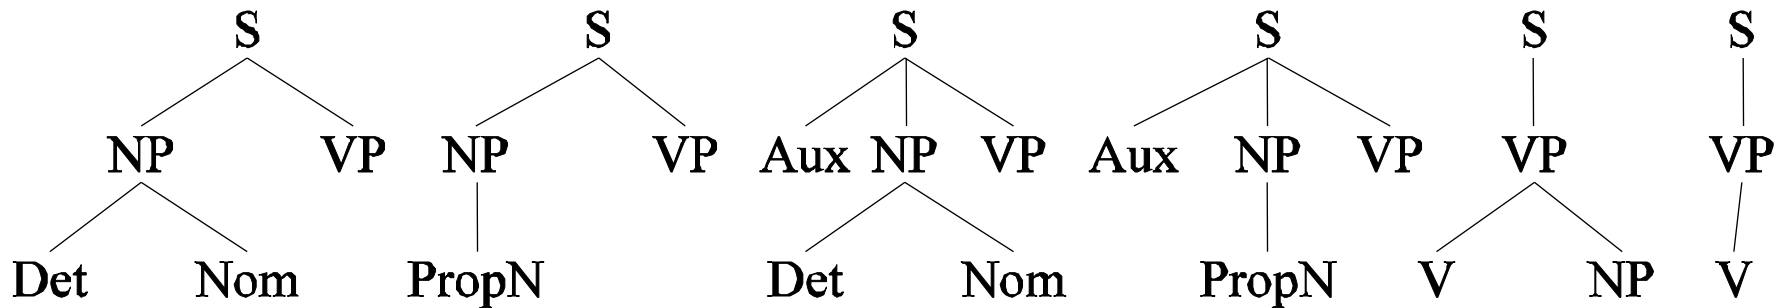
- Derivation for "Book that flight" (from the text)
  - The Start symbol

    S

  - Can derive 3 rules as follows:

```
        S                    S               S
       / \                 / | \             |
      NP  VP            AUX  NP  VP           VP
```

  - Each non-terminal can derive additional rules

```
      S              S            S              S           S      S
     / \            / \          / | \          / | \        |      |
   NP   VP        NP   VP     Aux NP  VP     Aux  NP  VP      VP     VP
   / \            |           / \           |            / \      |
 Det   Nom      PropN       Det   Nom     PropN        V    NP    V
```

  - Only the last two trees can derive the word "book" as first in the input

Bottom-up Parser

- Data-driven

- Looks at words in input string first, checks / assigns their category(ies), and tries to combine them into acceptable structures in the grammar

- Involves scanning the derivation so far for sub-strings which match the right-hand-side of grammar / production rules and using the rule that would show their derivation from the non-terminal symbol of that rule

Bottom-up Derivation

– Starts with input text

Book   that   flight

– derive the text from rules, in this case, two possible lexical rules
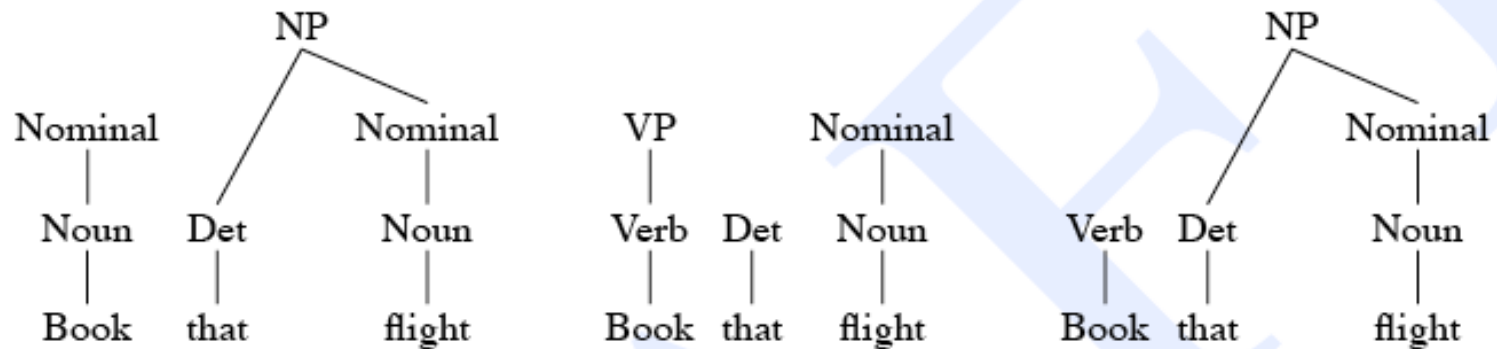
Noun   Det   Noun   Verb   Det   Noun
 | | | | | |
Book   that   flight   Book   that   flight

– Each of those can be derived from nonterminals

Nominal   Nominal   Nominal
 | | |
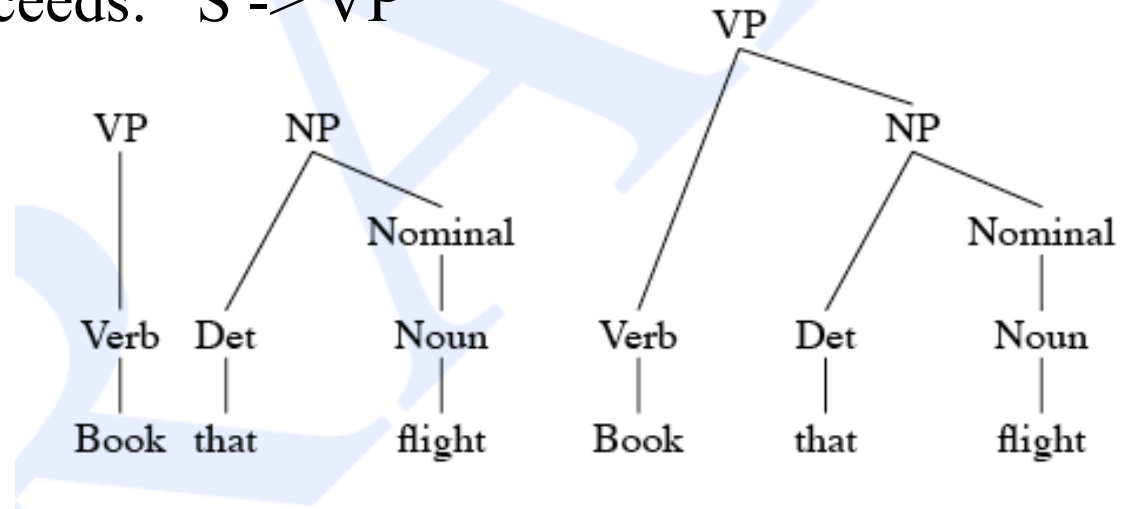Noun   Det   Noun   Verb   Det   Noun
 | | | | | |
Book   that   flight   Book   that   flight

Bottom-Up Derivation

- Only the rightmost tree can continue the derivation here:

- And only one succeeds: S -> VP

Bottom-up Parsing

- Algorithm called shift/reduce parsing
    - Scans the input from left to right and keeps a "stack" of the partial parse tree so far
    - The shift operation looks at the next input and shifts it onto the stack
    - The reduce operation looks at N symbols on the stack and if they match the RHS of a grammar rule, reduces the stack by replacing those symbols with the nonterminal
- Also must either incorporate back-tracking or must keep multiple possible parses

Parsing issues

- Top-down
  - Only searches for trees that can be answers (i.e. S's)
  - But also suggests trees that are not consistent with any of the words
- Bottom-up
  - Only forms trees consistent with the words
  - But suggest trees that make no sense globally
- Note that in the previous example, there was local ambiguity between "book" being a verb or a noun that was resolved at the end of the parse
- But examples with structural ambiguity will not be resolved, resulting in more than one possible derivation

Working with Parsing

- NLTK parsing demos
  - Top-down parsing using a recursive descent algorithm
    - Top down parsing with back-tracking
    - Must not have left-recursion in the grammar rules

   nltk.app.rdparser()

  - Bottom-up parsing using a shift-reduce algorithm
    - Instead of back-tracking or multiple parses, this NLTK implementation requires outside intervention to apply the correct rule when there is a choice

      nltk.app.srparser()

- Described in NLTK book, Chapter 8, Analyzing Sentence Structure

Structural Ambiguity

- *One morning I shot an elephant in my pajamas. How he got into my pajamas I don't know.* Groucho Marx, *Animal Crackers,* 1930.
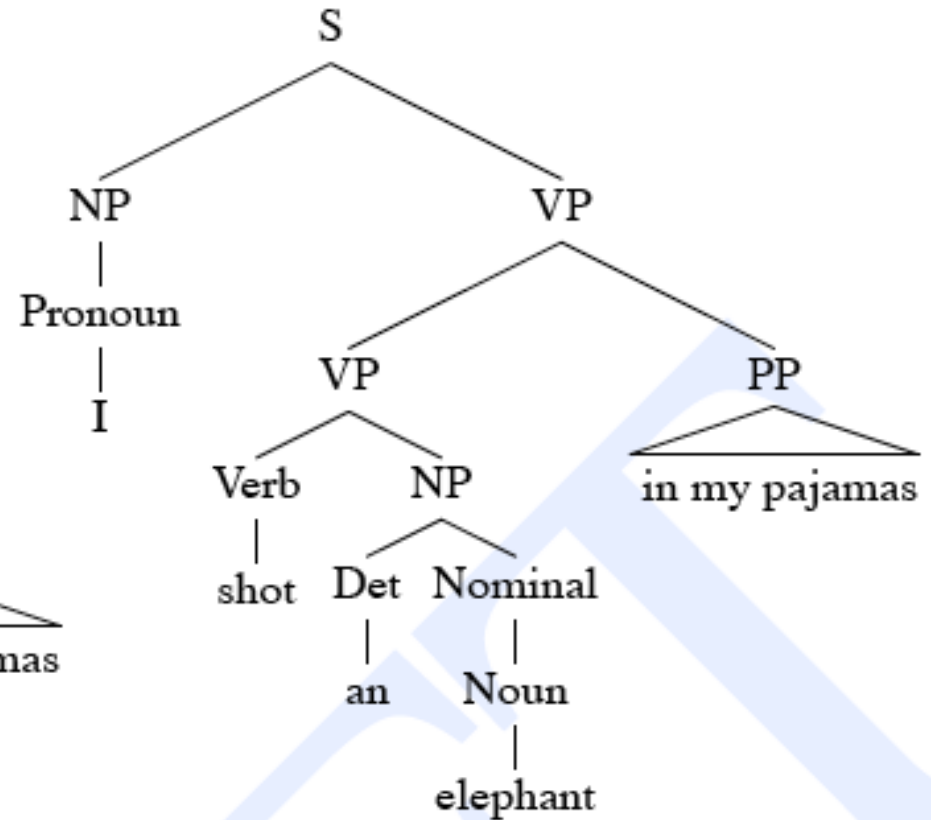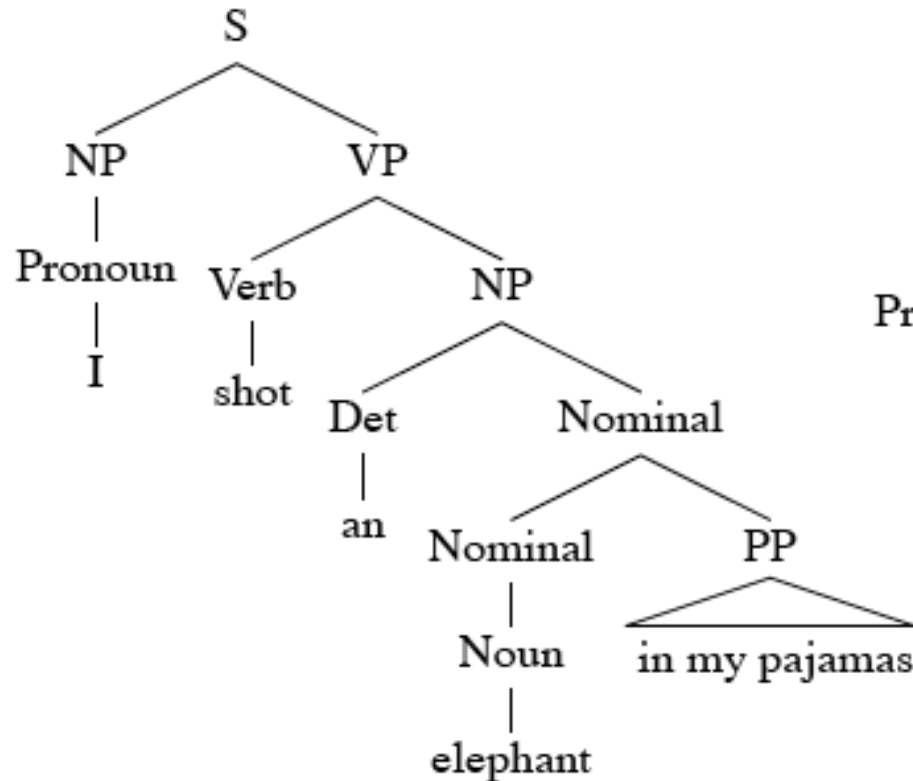
Classical Parsing

- In addition to the multiple parses due to structural ambiguity, typical grammars built before the 1990's would overgeneralize
  - Real broad-coverage language grammar could give rise to millions of parses on a single sentence
- Structuring grammar to restrict parses would leave up to 30% of the sentence without parses
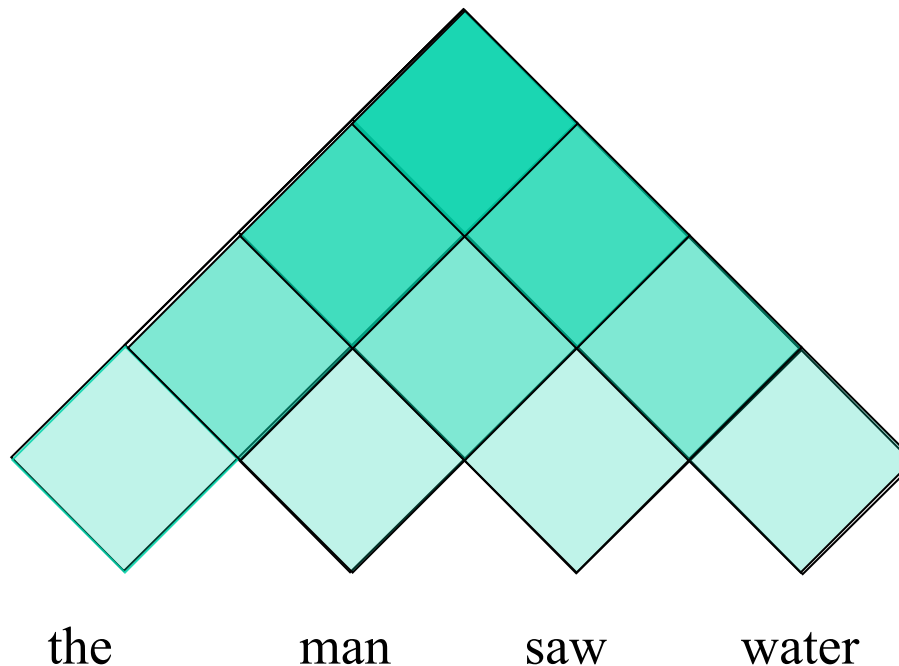
Parsing Algorithms

- The simple parsers that we have seen are exponential in time (recursive descent with back-tracking) and (shift reduce with back-tracking)

- Avoid back-tracking and re-doing subtrees
  - Recall that the backtracking recursive descent expanded some subtrees multiple times

- Use forms of dynamic programming to search for good parse trees
  - Attempt to perform exponential process in polynomial time

Chart Parsers

- CKY (Cocke-Kasami-Younger) algorithm is an example
  - Bottom-up parser
  - Requires grammar to be in Chomsky Normal Form, with only two symbols on the right-hand-side of each production
    - All CFG grammars have a Chomsky Normal Form
      - Grammar rule with 3 RHS symbols: NP -> Det NP PP
      - Transformed to equivalent grammar with only 2 RHS symbols:
        NP -> Det NPtemp
        Nptemp -> NP PP
  - Fills in a data structure called a chart or a parse triangle
- Binarization is key in reducing exponential process
  - Where binarization means only reducing rules with 2 RHS symbols
  - Other parsers, such as Earley's algorithm, use similar chart ideas to work on two subtrees at a time
  - Key idea in parser development from 1970 - 1990

CKY Parsing

- For input of length n, fills a parse table triangle of size (n+1, n+1) , where each element has the non-terminal production representing the span of text from position i to j.
  - Cells in first (bottom) layer describe trees of single words
  - Cells in second layer describes how rewrite rules can be used to combine trees in first layer for trees with two words
  - Etc.

the   man   saw   water

Grammar Rules

Rule	Prob
S → NP VP	0.9
S → VP	0.1
VP → V NP	0.5
VP → V	0.1
VP → V @VP_V	0.3
VP → V PP	0.1
@VP_V → NP PP	1.0
NP → NP NP	0.1
NP → NP PP	0.2
NP → N	0.7
PP → P NP	1.0
N → *people*	0.5
N → *fish*	0.2
N → *tanks*	0.2
N → *rods*	0.1
V → *people*	0.1
V → *fish*	0.6
V → *tanks*	0.3
P → *with*	1.0

CKY chart for sentence "fish people fish tanks":

	fish (0–1)	people (1–2)	fish (2–3)	tanks (3–4)
0–1	N → fish 0.2 V → fish 0.6 NP → N 0.14 VP → V 0.06 S → VP 0.006	NP → NP NP 0.0049 VP → V NP 0.105 S → VP 0.0105	NP → NP NP 0.0000686 VP → V NP 0.00147 S → NP VP 0.000882	NP → NP NP 0.0000009604 VP → V NP 0.00002058 S → NP VP 0.00018522
1–2		N → people 0.5 V → people 0.1 NP → N 0.35 VP → V 0.01 S → VP 0.001	NP → NP NP 0.0049 VP → V NP 0.007 S → NP VP 0.0189	NP → NP NP 0.0000686 VP → V NP 0.000098 S → NP VP 0.01323
2–3			N → fish 0.2 V → fish 0.6 NP → N 0.14 VP → V 0.06 S → VP 0.006	NP → NP NP 0.00196 VP → V NP 0.042 S → VP 0.0042
3–4				N → tanks 0.2 V → tanks 0.1 NP → N 0.14 VP → V 0.03 S → VP 0.003

Call buildTree(score, back) to get the best parse

Example showing filled-in CKY chart for a PCFG
for sentence "fish people fish tanks"

Need for Treebanks

- Before you can parse you need a grammar.
- So where do grammars come from?
  - Grammar Engineering
    - Hand-crafted decades-long efforts by humans to write grammars (typically in some particular grammar formalism of interest to the linguists developing the grammar).
  - TreeBanks
    - Semi-automatically generated sets of parse trees for the sentences in some corpus. Typically in a generic lowest common denominator formalism (of no particular interest to any modern linguist, but representing phrases of text in actual use).

Section on Treebanks and probabilistic parsing from Jim Martin's online slides.

The rise of annotated data

- Starting off, building a treebank seems a lot slower and less useful than building a grammar

- But a treebank gives us many things
    - Reusability of the labor
        - Many parsers, POS taggers, etc.
        - Valuable resource for linguistics
    - Broad coverage
    - Frequencies and distributional information
    - A way to evaluate systems on the same text

Penn Treebank

```
( (S
   (NP-SBJ (DT The) (NN move))
   (VP (VBD followed)
    (NP
      (NP (DT a) (NN round))
      (PP (IN of)
       (NP
         (NP (JJ similar) (NNS increases))
         (PP (IN by)
          (NP (JJ other) (NNS lenders)))
         (PP (IN against)
          (NP (NNP Arizona) (JJ real) (NN estate) (NNS loans))))))
    (, ,)
    (S-ADV
     (NP-SBJ (-NONE- *))
     (VP (VBG reflecting)
      (NP
        (NP (DT a) (VBG continuing) (NN decline))
        (PP-LOC (IN in)
         (NP (DT that) (NN market)))))))
   (. .)))
```

Getting grammar from a treebank

- Given an annotated sentence,

(11.10) [NP Shearson's] [JJ easy-to-film], [JJ black-and-white] "[SBAR Where We Stand]" [NNS commercials]

- We can make a grammar rule:

```
NP → NP JJ , JJ `` SBAR '' NNS
```

- And we'll make rules for sub-trees as well

Sample Rules for Noun Phrases

```
NP → DT JJ NNS
NP → DT JJ NN NN
NP → DT JJ JJ NN
NP → DT JJ CD NNS
NP → RB DT JJ NN NN
NP → RB DT JJ JJ NNS
NP → DT JJ JJ NNP NNS
NP → DT NNP NNP NNP NNP JJ NN
NP → DT JJ NNP CC JJ JJ NN NNS
NP → RB DT JJS NN NN SBAR
NP → DT VBG JJ NNP NNP CC NNP
NP → DT JJ NNS , NNS CC NN NNS NN
NP → DT JJ JJ VBG NN NNP NNP FW NNP
NP → NP JJ , JJ `` SBAR '' NNS
```

TreeBank Grammars

- Reading off the grammar…
- The grammar is the set of rules (local subtrees) that occur in the annotated corpus
- They tend to avoid recursion (and elegance and parsimony)
  - i.e. they tend to the flat and redundant
- Penn TreeBank (III) has about 17500 grammar rules under this definition.
- But the main use of the Treebank is to provide the probabilities to inform the statistical parsers, and the grammar does not actually have to be generated.
- The grammar hovers behind the Treebank; it is in the minds of the human annotators (and in the annotation manual!)

Probabilistic Context-Free Grammars

- By way of introduction to statistical parsers, we first introduce the idea of associating probabilities with grammar rewrite rules.

  - Attach probabilities to grammar rules
  - The expansions for a given non-terminal sum to 1

    VP -> Verb   .55

    VP -> Verb NP   .40

    VP -> Verb NP NP   .05

Getting the probabilities

- From a treebank of annotated data, get the probabilities that any non-terminal symbol is rewritten with a particular rule
  - So for example, to get the probability for a particular VP rule just count all the times the rule is used and divide by the number of VPs overall.
- The parsing task is to generate the parse tree with the highest probability (or the top n parse trees)
- The probability of a parse tree is the product of the probabilities of the rules used in the derivation

$$P(T,S) = \prod_{node \in T} P(rule(n))$$

Typical Approach

- Use CKY as the backbone of the algorithm
- Assign probabilities to constituents as they are completed and placed in the table
- Use the max probability for each constituent going up

Problems with PCFG Parsing

- But this typical approach always just picks the most likely rule in the derivation
  - For example, if it is more likely that a prepositional phrases attaches to the noun phrase that it follows instead of the verb, then the probabilistic parser will always attach prepositional phrases to the closest noun
- The probability model we're using is only based on the rules in the derivation…
  - Doesn't use the words in any real way
  - Doesn't take into account where in the derivation a rule is used
    - E.g. the parent of the non-terminal of the derivation
  - Doesn't really work
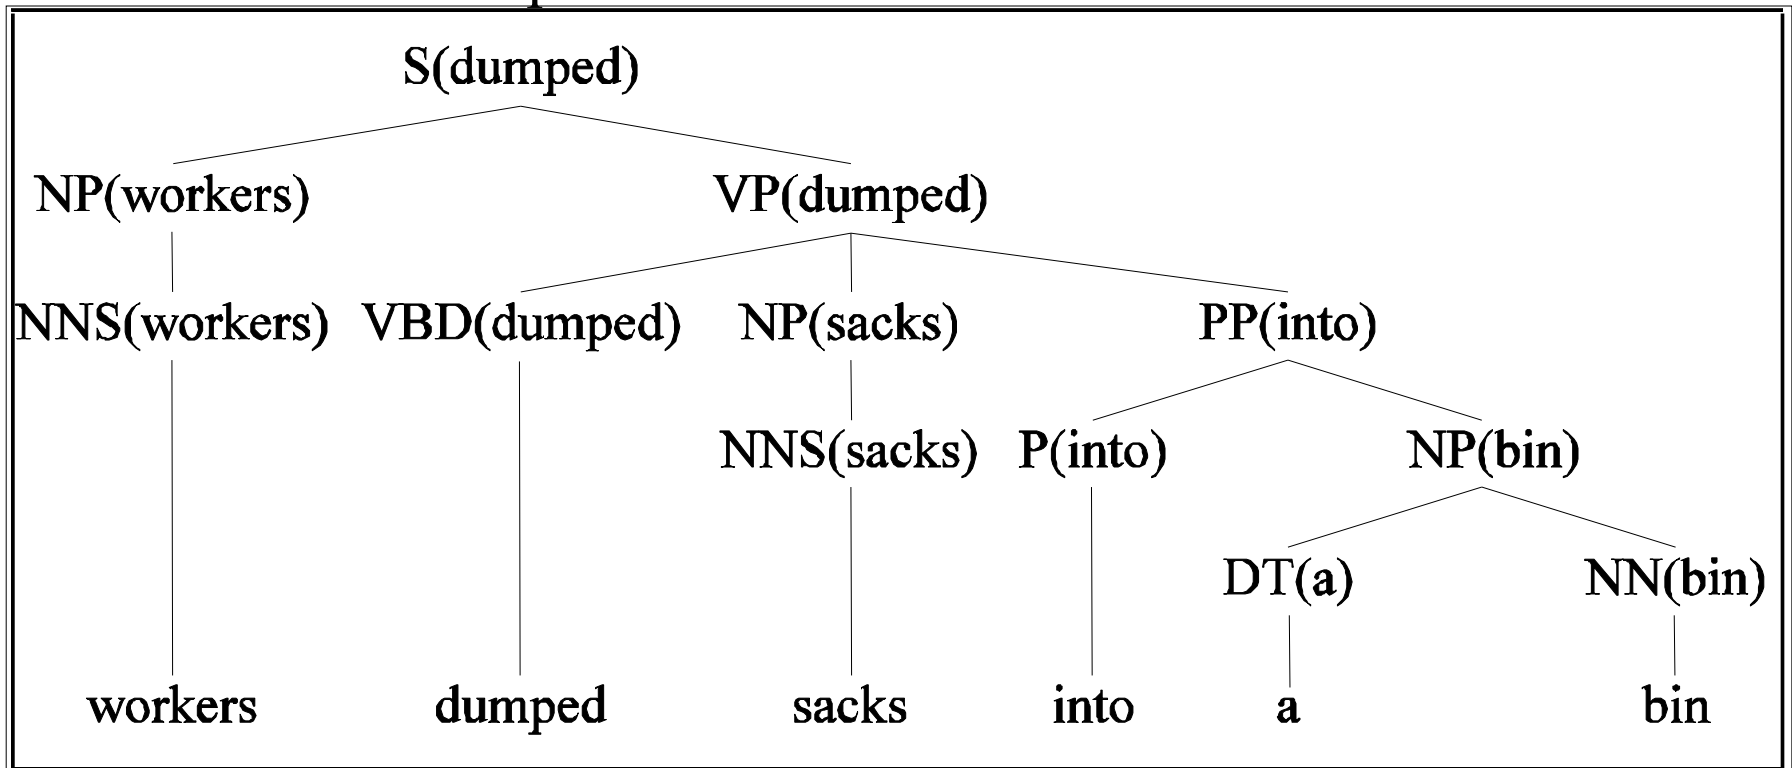   - Most probable parse isn't usually the right one (the one in the treebank test set).

Lexicalized Parsing

- Add lexical dependencies to the scheme…
  - Integrate the preferences of particular words into the probabilities in the derivation
  - i.e. Condition the rule probabilities on the actual words
- To do that we're going to make use of the notion of the head of a phrase
  - The head of an NP is its noun
  - The head of a VP is its verb
  - The head of a PP is its preposition

  (It's really more complicated than that but this will do.)
- Main parsing breakthrough idea of the 1990's
- Expand the set of phrase types with phrase type/word
 - In practice, we learn probabilities to automatically detect head words

Example (right)

- Should we attach the prepositional phrase with head "into" to the verb "dumped"?

S(dumped)

NP(workers)   VP(dumped)

NNS(workers) VBD(dumped) NP(sacks)   PP(into)

NNS(sacks) P(into)   NP(bin)

DT(a)   NN(bin)

workers   dumped   sacks   into   a   bin

- In this tree, each phrase type, such as NP or VP, is also shown with its attached head word.

Example (wrong)

- Or should we attach the prepositional phrase with head "into" to the noun "sacks"?
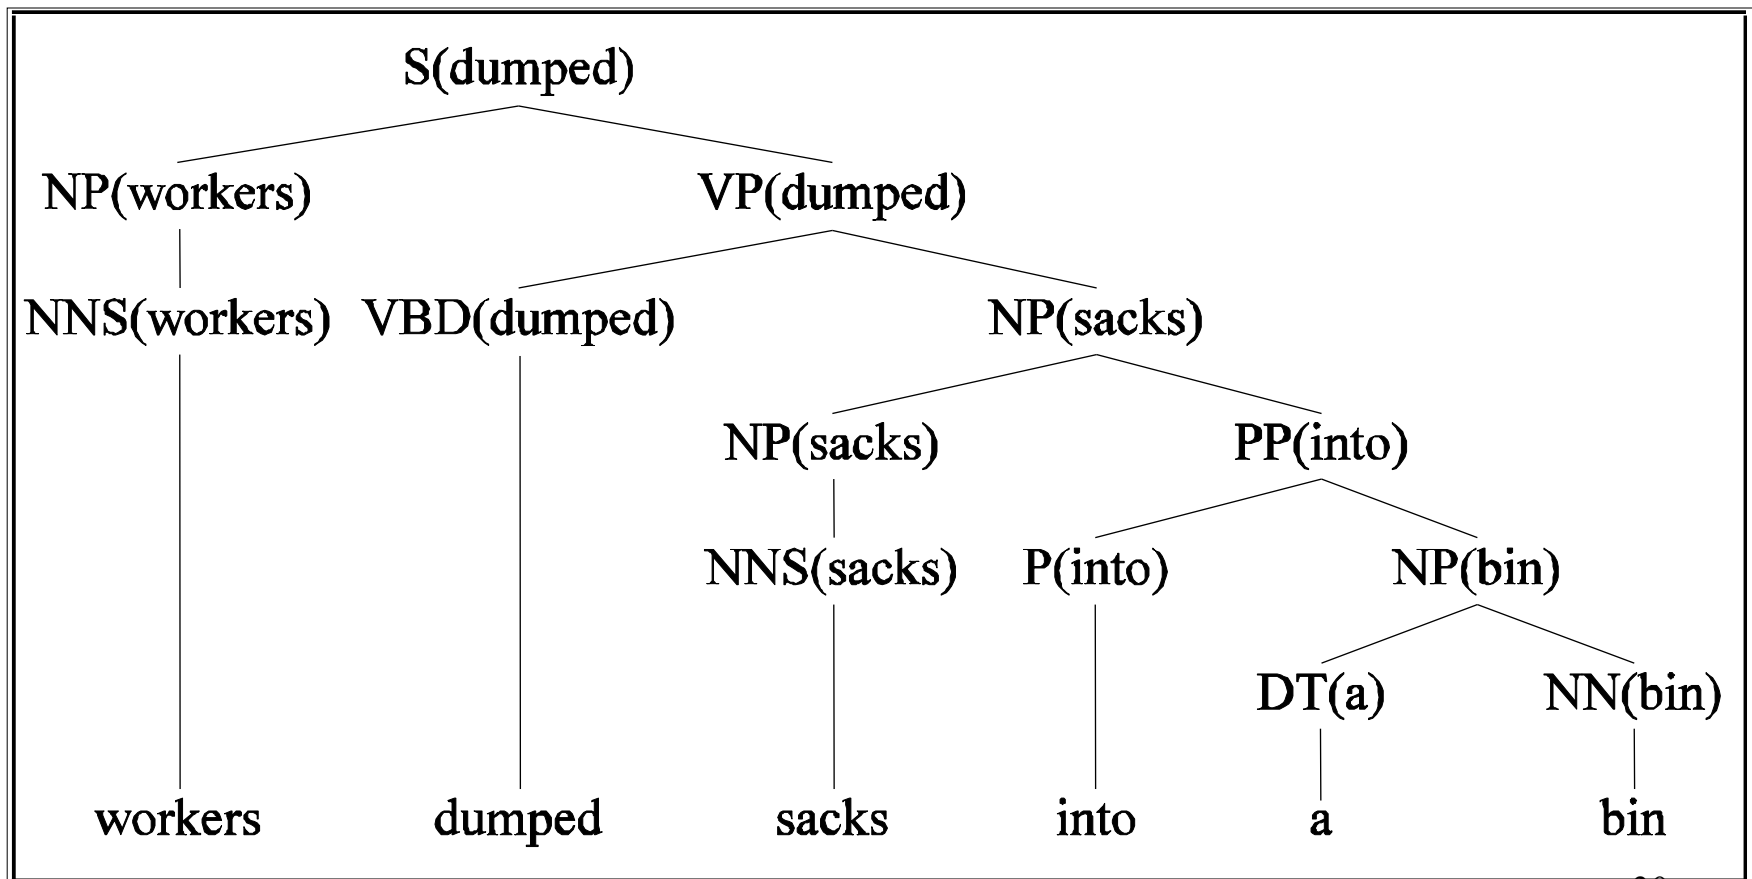
Preferences

- The issue here is the attachment of the PP. So the affinities we care about are the ones between dumped and into vs. sacks and into.
  - So count the places where dumped is the head of a constituent that has a PP daughter with into as its head and normalize
  - Vs. the situation where sacks is a constituent with into as the head of a PP daughter.
- In general, collect statistics on preferences (aka affinities)
  - Use verb subcategorization
    - Particular verbs have affinities for particular VPs
  - Objects affinities for their predicates (mostly their mothers and grandmothers)
    - Some objects fit better with some predicates than others

Preference example

- Consider the VPs
  - Ate spaghetti with gusto
  - Ate spaghetti with marinara
- The affinity of gusto for eat is much larger than its affinity for spaghetti
- On the other hand, the affinity of marinara for spaghetti is much higher than its affinity for ate

Preference Example (2)

- Note the relationship here is more distant and doesn't involve a headword since gusto and marinara aren't the heads of the PPs.

Vp (ate)

Vp(ate)   Pp(with)

v   np

Ate spaghetti with gusto

Vp(ate)

Np(spag)

np   Pp(with)

v

Ate spaghetti with marinara
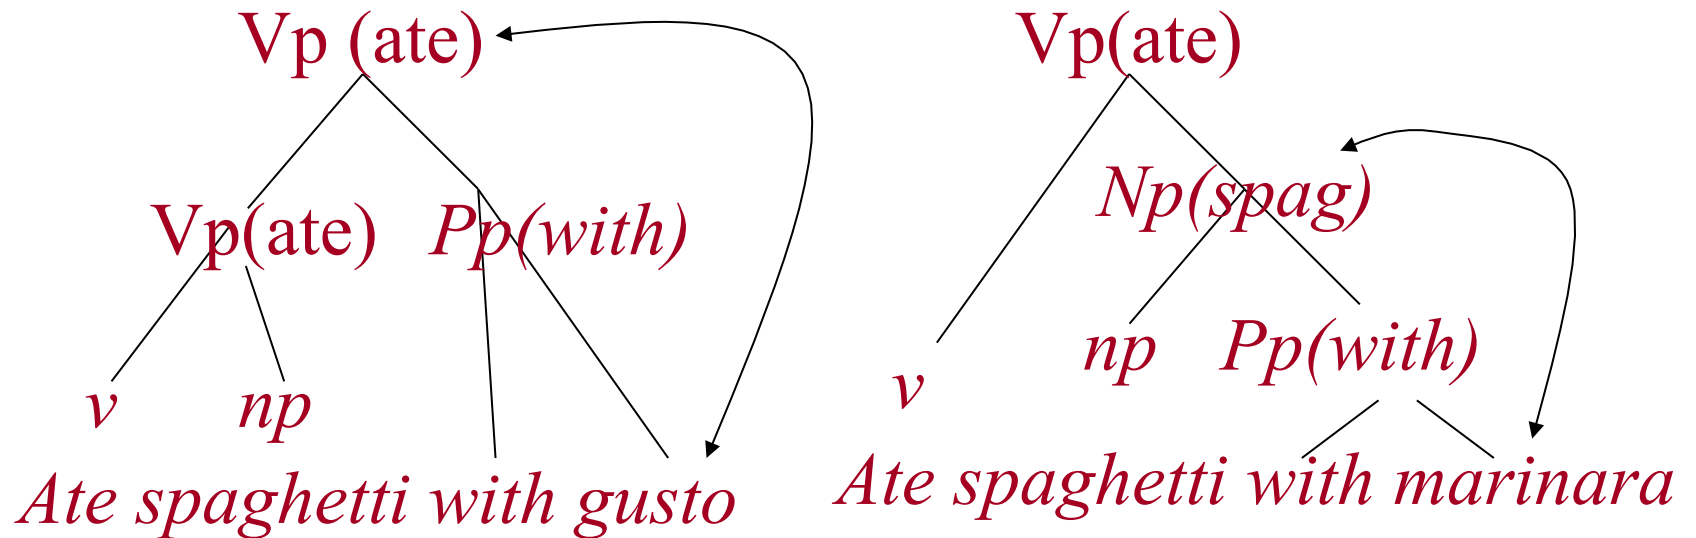
Note

- Jim Martin: "In case someone hasn't pointed this out yet, this lexicalization stuff is a thinly veiled attempt to incorporate semantics into the syntactic parsing process…
  - Duhh..,. Picking the right parse requires the use of semantics."

Last Points

- Statistical parsers are getting quite good, but it's still quite challenging to expect them to come up with the correct parse given only statistics from syntactic information.
- But if our statistical parser comes up with the top-N parses, then it is quite likely that the correct parse is among them.
- Lots of current work on
  - Re-ranking to make the top-N list even better.

- There are also grammar-driven parsers that are competitive with the statistical parsers, notably the CCG (Combinatory Categorial Grammar) parsers

Evaluation

- Given that it is difficult/ambiguous to produce the entire correct tree, look at how much of content of the trees are correctly produced
  - Evaluation measures based on the correct number of constituents (or sub-trees) in the system compared to the reference (gold standard)
- Precision
  - What fraction of the sub-trees in our parse matched corresponding sub-trees in the reference answer
    - How much of what we're producing is right?
- Recall
  - What fraction of the sub-trees in the reference answer did we actually get?
    - How much of what we should have gotten did we get?
- F-measure combines precision and recall to give an overall score.

Evaluation

- An additional evaluation measure that is often reported is that of Crossing Brackets errors, in which the subtrees are equal, but they are put together in a different order.

Parser hypothesis   Reference answer

((A B) C)   (A (B C))
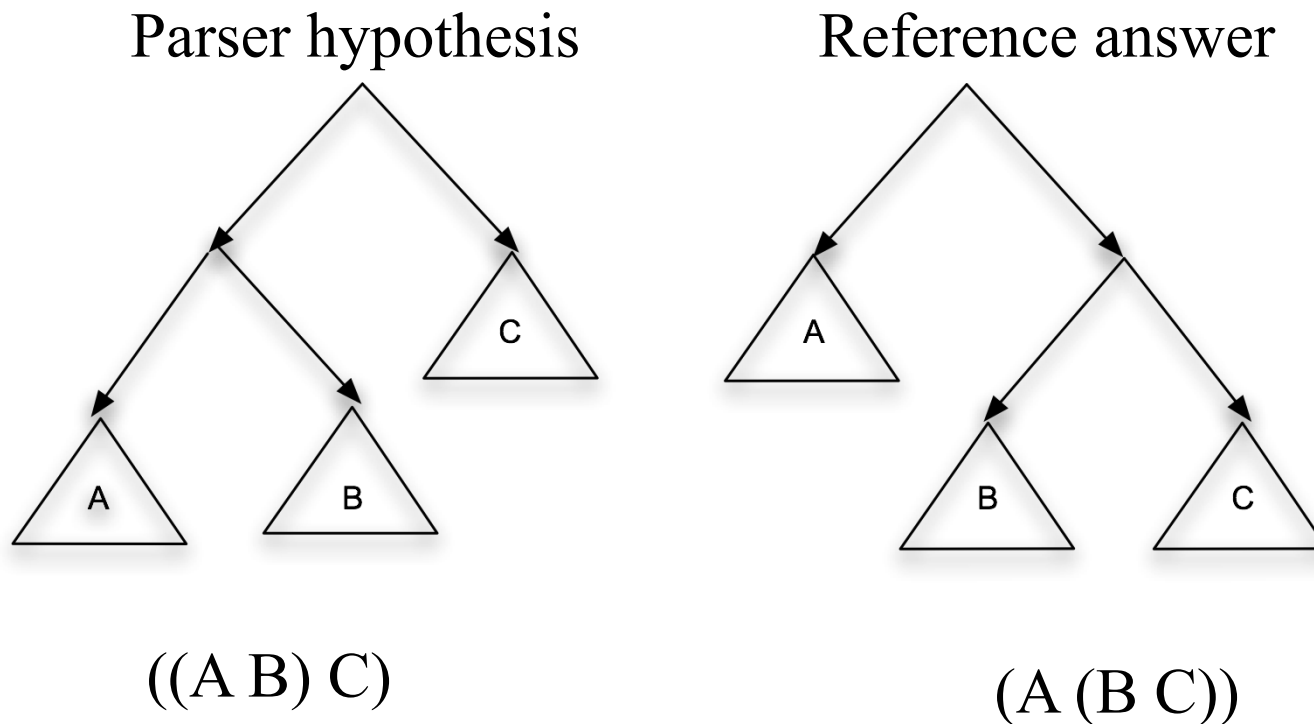
Available Parsers

- Among the family of lexicalized statistical parsers are the well-known Collins parser (Michael Collins 1996, 1999) and the Charniak parser (1997)
  - both are publicly available and widely used in NLP, for non-commercial purposes.
- The Charniak series of parsers is still under development, by Eugene Charniak and his group; it produces N-best parse trees.
  - Its evaluation is currently the best on the Penn Treebank at about 91% F measure.
- Another top performing parser, originally by Dan Klein and Christopher Manning, is available from the Stanford NLP group
  - combines "separate PCFG phrase structure and lexical dependency experts".
  - Demo at: http://nlp.stanford.edu:8080/parser/

Available Parsers

- The CCG parsers are available from their open source page
  - http://groups.inf.ed.ac.uk/ccg/software.html

- Parsers are also available through the OpenNLP project, with the OpenNLP API:
  - http://opennlp.sourceforge.net/

Partial Parsing

- For many applications you don't really need a full-blown syntactic parse. You just need a good idea of where the base syntactic units are.
  - Often referred to as chunks.
- For example, if you're interested in locating all the people, places and organizations in a text it might be useful to know where all the base NPs are.
- A full partial parse would have chunks for all the text, but with no hierarchical structure:

$[_{NP}$ The morning flight$]$ $[_{PP}$ from$]$ $[_{NP}$ Denver$]$ $[_{VP}$ has arrived.$]$

- A partial parse for just base NPs would be:

$[_{NP}$ The morning flight$]$ from $[_{NP}$ Denver$]$ has arrived.

Rule-Based Partial Parsing

- Restrict the form of rules to exclude recursion (make the rules flat).
- Group and order the rules so that the RHS of the rules can refer to non-terminals introduced in earlier rules but not later ones.
- Write regular expressions to recognize the right-hand-side of rules, starting from the later ones.
- For complete chunking, typical ordering:
  - Base syntactic phrases
  - Larger verb and noun groups
  - Sentential level rules

Partial Parsing

$NP \rightarrow (Det) \; Noun^* \; Noun$

$NP \rightarrow Proper\text{-}Noun$

$VP \rightarrow Verb$

$VP \rightarrow Aux \; Verb$

- No direct or indirect recursion allowed in these rules.

- That is you can't directly or indirectly reference the LHS of the rule on the RHS.

Evaluation

- For evaluation, we need a metric that works at the level of the chunks.

- Precision:
  - The fraction of chunks the system returned that were right
    - "Right" means the boundaries and the label are correct given some labeled test set.

- Recall:
  - The fraction of the chunks that system got from those that it should have gotten.

- F measure: Harmonic mean of those two numbers.

Dependency Parsing

- Dependency parsing has some resemblance to lexicalized parsing because of the importance of the lexical entities (words) to capturing the syntactic structure

- But dependency parsing produces a simpler representation of the structure.
    - Can be easier to use in some semantic applications

Transition-based parsers

- A typical parser of this type is that of Nivre 2004, which is a bottom-up "span" parser
  - A shift/reduce parser that adds dependency relations
- Operation of parser:
  - State: stack of partially processed items and a queue of remaining tokens
  - Transitions: add dependency arcs; stack or queue operations
    - Operations are
      - Build left arc
      - Build right arc
      - Shift
      - Reduce

Training the Parser

- How does the parser know which operation, and arc label, to apply?

- It learns these from the annotated corpus.

- For English, dependency grammar relations are derived from Penn Treebank

- Then a collection of examples is extracted from the text to train the parser
  - Each example consists of a set of features representing the state of the parser, including the next word, previous word, POS tags, etc.
  - A machine learning algorithm is applied to learn a classifier, which can assign a parsing operation to every parsing state

Non-Projective Parsing

- The bottom-up span parser is for projective parsing; but there is also a technique to add non-projective relations after the parse

  – Nivre 2008

- Alternate technique is due to Ryan McDonald, 2005, which converts the dependency parsing problem to that of finding a maximal spanning tree in the dependency graph.

  – Again, the dependencies in the graph are learned from the annotated corpus of Penn Treebank

- Non-projective dependency is still an active area of research