ADVANCED PARSING

David Kauchak
CS159 – Fall 2020

some slides adapted from Dan Klein

1

Admin

Assignment 3?

Assignment 4 (A and B)

Lab next Tuesday

2

Parsing evaluation

You've constructed a parser

You want to know how good it is

Ideas?

3

Parsing evaluation

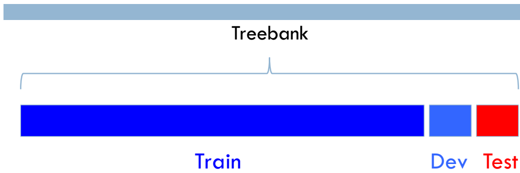

Learn a model using the training set

Parse the test set without looking at the "correct" trees

Compare our generated parse tree to the "correct" tree

4

Comparing trees

Computed Tree P

Correct Tree T

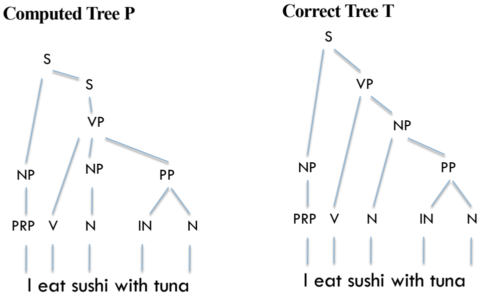

Ideas?

5

Comparing trees

Idea 1: see if the trees match exactly

- Problems?
  - Will have a low number of matches (people often disagree)
  - Doesn't take into account getting it *almost* right

Idea 2: compare the constituents

6

Comparing trees

Computed Tree P

Correct Tree T

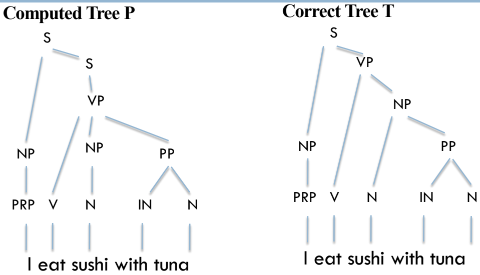

How many constituents match?

How can we turn this into a score?

7

Evaluation measures

Precision

$$\frac{\text{\# of correct constituents}}{\text{\# of constituents in the computed tree}}$$

Recall

$$\frac{\text{\# of correct constituents}}{\text{\# of constituents in the correct tree}}$$

F1

$$\frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}}$$

What does this favor?

8

Comparing trees

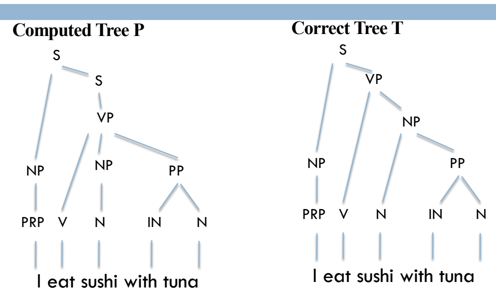

Constituents: 11   # Correct Constituents: 9   # Constituents: 10

Precision: 9/11   Recall: 9/10   F1: 0.857

9

Parsing evaluation

Corpus: Penn Treebank, WSJ

Training:	sections	02-21
Development:	section	22 (first 20 files)
Test:	section	23

Parsing has been fairly standardized to allow for easy comparison between systems

10

Treebank PCFGs

Use PCFGs for broad coverage parsing

Can take a grammar right off the trees (doesn't work well):

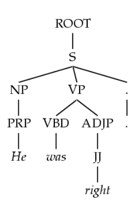

ROOT → S
S → NP VP .
NP → PRP
VP → VBD ADJP
…..

Model	F1
Baseline	72.0

11

Generic PCFG Limitations

PCFGs do not use any information about where the current constituent is in the tree

PCFGs do not rely on specific words or concepts, only general structural disambiguation is possible (e.g. prefer to attach PPs to Nominals)

MLE estimates are not always the best

12

Conditional Independence?

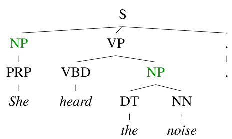

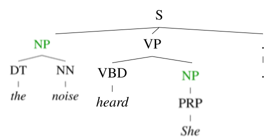

Will a PCFG differentiate between these?

What's the problem?

13

Conditional Independence?

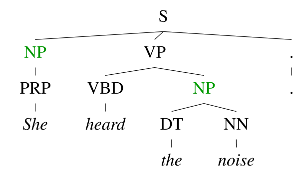

It treats all NPs as equivalent… but they're not!

- A grammar with symbols like "NP" won't be context-free
- Statistically, conditional independence too strong

14

Strong independence assumption

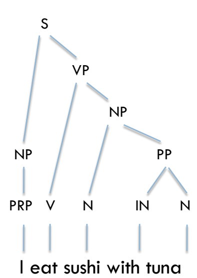

S -> NP VP
NP -> PRP
PRP -> I
VP -> V NP
V -> eat
NP -> N PP
N -> sushi
PP -> IN N
IN -> with
N -> tuna

We're making a strong independence assumption here!

15

Non-Independence

Independence assumptions are often too strong

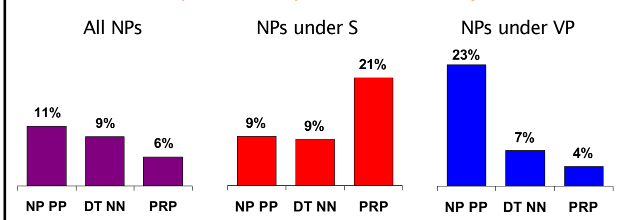

Example: the expansion of an NP is highly dependent on the parent of the NP (i.e., subjects vs. objects).

Also: the subject and object expansions are correlated

16

Grammar Refinement

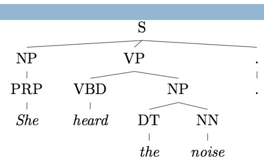

Idea: expand/refine our grammar

Challenges:

- Must refine in ways that facilitate disambiguation
- Must trade-offs between too little and too much refinement.
  - Too much refinement -> sparsity problems
  - To little -> can't discriminate (PCFG)

17

Grammar Refinement

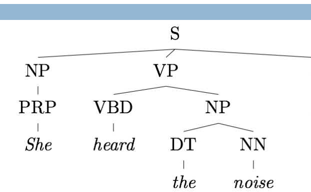

Ideas?

18

Grammar Refinement

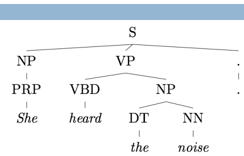

Structure Annotation [Johnson '98, Klein&Manning '03]

- Differentiate constituents based on their local context

Lexicalization [Collins '99, Charniak '00]

- Differentiate constituents based on the spanned words

Constituent splitting [Matsuzaki et al. 05, Petrov et al. '06]

- Cluster/group words into sub-constituents

19

Markovization

Except for the root node, every node in a parse tree has:

- A vertical history/context
- A horizontal history/context

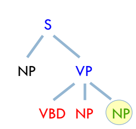

Traditional PCFGs use the full horizontal context and a vertical context of 1

20

Vertical Markovization

Vertical Markov order: rewrites depend on past *k* ancestor nodes.

Order 1 is most common: aka parent annotation

Order 1

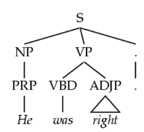

Order 2

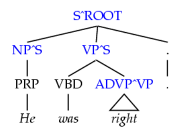

21

Allows us to make finer grained distinctions

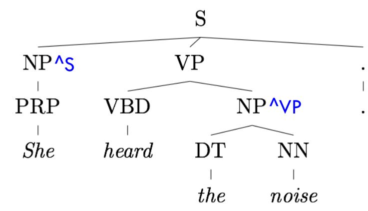

22

Vertical Markovization

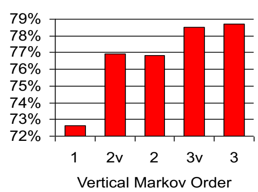

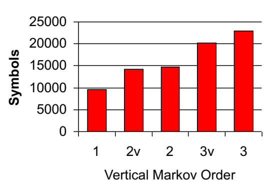

F1 performance

# of non-terminals

23

Horizontal Markovization

Horizontal Markov order: rewrites depend on past *k* sibling nodes

Order 1 is most common: condition on a single sibling

Order 1

Order ∞

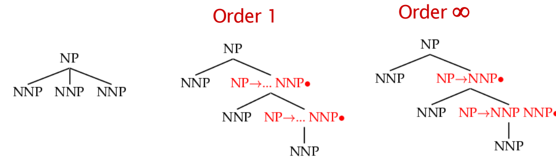

24

Horizontal Markovization

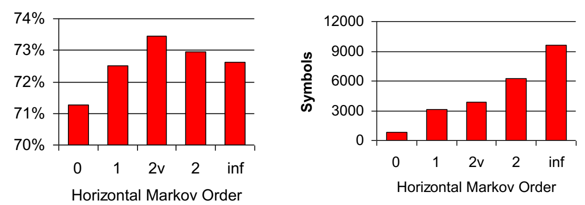

F1 performance

# of non-terminals

25

Problems with PCFGs

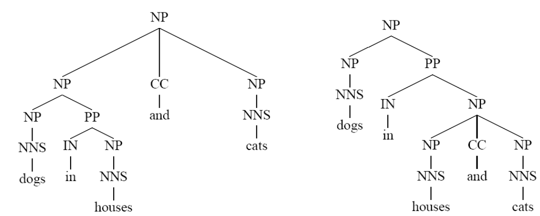

What's different between basic PCFG scores here?

26

Example of Importance of Lexicalization

A general preference for attaching PPs to NPs rather than VPs can be learned by an ordinary PCFG

But the desired preference can depend on specific words

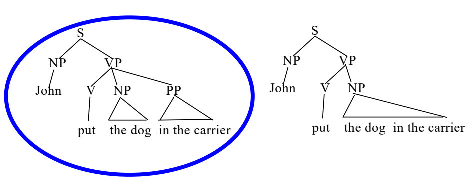

Which is correct?

27

Example of Importance of Lexicalization

A general preference for attaching PPs to NPs rather than VPs can be learned by an ordinary PCFG

But the desired preference can depend on specific words

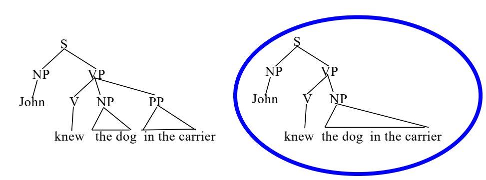

Which is correct?

28

Lexicalized Trees

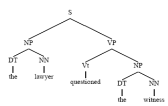

How could we lexicalize the grammar/tree?

29

Lexicalized Trees

Add "headwords" to each phrasal node

- Syntactic vs. semantic heads
- Headship not in (most) treebanks
- Usually *use head rules*, e.g.:
  - NP:
    - Take leftmost NP
    - Take rightmost N*
    - Take rightmost JJ
    - Take right child
  - VP:
    - Take leftmost VB*
    - Take leftmost VP
    - Take left child

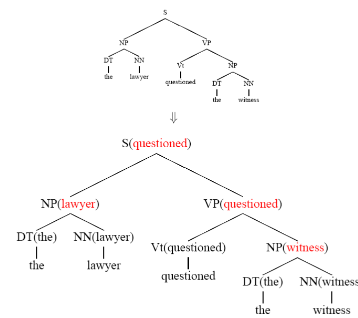

30

Lexicalized PCFGs?

Problem: we now have to estimate probabilities like

VP(put) → VBD(put) NP(dog) PP(in)

How would we estimate the probability of this rule?

$$\frac{\text{Count(VP(put)} \rightarrow \text{VBD(put) NP(dog) PP(in))}}{\text{Count(VP (put))}}$$

Never going to get these automatically off of a treebank

Ideas?

31

One approach

Combine this with some of the markovization techniques we saw

Collins' (1999) parser

Models productions based on context to the left and the right of the head child.

$$\text{LHS} \rightarrow L_n L_{n-1} \ldots L_1 H\ R_1 \ldots R_{m-1} R_m$$

32

One approach

$LHS \rightarrow L_n L_{n-1} \ldots L_1 H \; R_1 \ldots R_{m-1} R_m$

First generate the head (H) given the parent

Then repeatedly generate left symbols (L_i) until the beginning is reached

Then right (R_i) symbols until the end is reached

33

Sample Production Generation

$VP_{put} \rightarrow VBD_{put}\ NP_{dog}\ PP_{in}$

$VP_{put} \rightarrow$

34

Sample Production Generation

$VP_{put} \rightarrow VBD_{put}\ NP_{dog}\ PP_{in}$

$VP_{put} \rightarrow \underset{H}{VBD_{put}}$

$P_H(VBD \mid VP_{put})$

35

Sample Production Generation

$VP_{put} \rightarrow VBD_{put}\ NP_{dog}\ PP_{in}$

$VP_{put} \rightarrow \underset{L_1}{STOP}\ \ \underset{H}{VBD_{put}}$

$P_L(STOP \mid VP_{put})$

36

Sample Production Generation

$VP_{put} \rightarrow VBD_{put}\ NP_{dog}\ PP_{in}$

$VP_{put} \rightarrow$	STOP	VBD_{put}	NP_{dog}
	L_1	H	R_1

$P_R(NP_{dog} \mid VP_{put})$

37

Sample Production Generation

$VP_{put} \rightarrow VBD_{put}\ NP_{dog}\ PP_{in}$

$VP_{put} \rightarrow$	STOP	VBD_{put}	NP_{dog}	PP_{in}
	L_1	H	R_1	R_2

$P_R(PP_{in} \mid VP_{put})$

38

Sample Production Generation

$VP_{put} \rightarrow VBD_{put}\ NP_{dog}\ PP_{in}$

$VP_{put} \rightarrow$	STOP	VBD_{put}	NP_{dog}	PP_{in}	STOP
	L_1	H	R_1	R_2	R_3

$P_R(STOP \mid PP_{in})$

39

Sample Production Generation

$VP_{put} \rightarrow VBD_{put}\ NP_{dog}\ PP_{in}$

Note: Penn treebank tends to have fairly flat parse trees that produce long productions.

$VP_{put} \rightarrow$	STOP	VBD_{put}	NP_{dog}	PP_{in}	STOP
	L_1	H	R_1	R_2	R_3

$P_L(STOP \mid VP_{put}) * P_H(VBD \mid VP_{put})*$
$P_R(NP_{dog} \mid VP_{put})*$
$P_R(PP_{in} \mid VP_{put}) * P_R(STOP \mid PP_{in})$

40

Estimating Production Generation Parameters

Estimate P_H, P_L, and P_R parameters from treebank data

$$P_R(PP_{in} \mid VP_{put}) = \frac{\text{Count}(PP_{in} \text{ right of head in a } VP_{put} \text{ production})}{\text{Count(symbol right of head in a } VP_{put})}$$

$$P_R(NP_{dog} \mid VP_{put}) = \frac{\text{Count}(NP_{dog} \text{ right of head in a } VP_{put} \text{ production})}{\text{Count(symbol right of head in a } VP_{put})}$$

Smooth estimates by combining with simpler models conditioned on just POS tag or no lexical info

$$smP_R(PP_{in} \mid VP_{put-}) = \lambda_1 P_R(PP_{in} \mid VP_{put}) + (1-\lambda_1)\,(\lambda_2 P_R(PP_{in} \mid VP_{VBD}) + (1-\lambda_2)\, P_R(PP_{in} \mid VP))$$

41

Problems with lexicalization

We've solved the estimation problem

There's also the issue of performance

Lexicalization causes the size of the number of grammar rules to explode!

Our parsing algorithms take too long too finish

Ideas?

42

Pruning during search

We can no longer keep all possible parses around

We can no longer guarantee that we actually return the most likely parse

Beam search [Collins 99]

- In each cell only keep the **K** most likely hypotheses
- Disregard constituents over certain spans (e.g. punctuation)
- F1 of 88.6!

43

Pruning with a PCFG

The Charniak parser prunes using a two-pass approach [Charniak 97+]

- First, parse with the base (non-lexicalized) grammar
- For each X:[i,j] calculate P(X|i,j,s)
  - This isn't trivial, and there are clever speed ups
- Second, do the full CKY
  - Skip any X :[i,j] which had low (say, < 0.0001) posterior
- Avoids almost all work in the second phase!

F1 of 89.7!

44

Final Results

Parser	F1 ≤ 40 words	F1 all words
Klein & Manning '03	86.3	85.7
Matsuzaki et al. '05	86.7	86.1
Collins '99	88.6	88.2
Charniak & Johnson '05	90.1	89.6
Petrov et. al. 06	**90.2**	**89.7**

45

Human Parsing

How do humans do it?

How might you try and figure it out computationally/experimentally?

46

Human Parsing

Read these sentences

Which one was fastest/slowest?

John put the dog in the pen with a lock.

John carried the dog in the pen with a bone in the car.

John liked the dog in the pen with a bone.

47

Human Parsing

Computational parsers can be used to predict human reading time as measured by tracking the time taken to read each word in a sentence.

Psycholinguistic studies show that words that are more probable given the preceding lexical and syntactic context are read faster.

- John put the dog in the pen with a lock.
- John carried the dog in the pen with a bone in the car.
- John liked the dog in the pen with a bone.

Modeling these effects requires an ***incremental*** statistical parser that incorporates one word at a time into a continuously growing parse tree.

48

Garden Path Sentences

People are confused by sentences that seem to have a particular syntactic structure but then suddenly violate this structure, so the listener is "lead down the garden path".

- The horse raced past the barn fell.
  - vs. The horse raced past the barn broke his leg.
- The complex houses married students.
- The old man the sea.
- While Anna dressed the baby spit up on the bed.

Incremental computational parsers can try to predict and explain the problems encountered parsing such sentences.

49

More garden sentences

http://www.fun-with-words.com/ambiguous_garden_path.html

The prime number few.
Fat people eat accumulates.
The cotton clothing is usually made of grows in Mississippi.
Until the police arrest the drug dealers control the street.
The man who hunts ducks out on weekends.
When Fred eats food gets thrown.
Mary gave the child the dog bit a bandaid.
The girl told the story cried.
I convinced her children are noisy.
Helen is expecting tomorrow to be a bad day.
The horse raced past the barn fell.
I know the words to that song about the queen don't rhyme.
She told me a little white lie will come back to haunt me.
The dog that I had really loved bones.
That Jill is never here hurts.
The man who whistles tunes pianos.
The old man the boat.
Have the students who failed the exam take the supplementary.
The raft floated down the river sank.
We painted the wall with cracks.
The tycoon sold the offshore oil tracts for a lot of money wanted to kill JR.

50