

Transformational Priors Over Grammars

Jason Eisner
Johns Hopkins University
July 6, 2002 — EMNLP

The Big Concept

- Want to parse (or build a syntactic language model).
- Must estimate rule probabilities.
- Problem:** Too many possible rules!
 - Especially with lexicalization and flattening (which help).
 - So it's hard to estimate probabilities.

The Big Concept

- Problem:** Too many rules!
 - Especially with lexicalization and flattening (which help).
 - So it's hard to estimate probabilities.
- Solution:** Related rules tend to have related probs
 - POSSIBLE relationships are given a priori
 - LEARN which relationships are strong in this language (just like feature selection)
- Method has connections to:
 - Parameterized finite-state machines (Monday's talk)
 - Bayesian networks (inference, abduction, explaining away)
 - Linguistic theory (transformations, metarules, etc.)

Problem: Too Many Rules

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP
2 NP → NNP fund
2 NP → DT NPR NN fund
2 S → TO fund NP PP

NP → DT JJ NN fund
NP → DT NN fund PP
NP → DT ADJP NN fund
NP → DT JJ JJ NN fund
NP → DT NN fund SBAR
NP → DT NN fund PP
NP → DT NN fund VP
NP → DT NN fund PP
NP → DT JJ JJ NN fund
NP → DT NN fund SBAR
NP → DT NN fund PP
NP → DT JJ fund
NP → DT NN fund
NP → DT NN fund PP
NP → DT JJ JJ NN fund
NP → DT NN fund SBAR
NP → DT NN fund PP
NP → DT JJ fund
NP → DT NN fund

[Want To Multiply Rule Probabilities]

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP
2 NP → NNP fund
2 NP → DT NPR NN fund
2 S → TO fund NP PP

NP → DT JJ NN fund
NP → DT NN fund PP
NP → DT ADJP NN fund
NP → DT JJ JJ NN fund
NP → DT NN fund SBAR
NP → DT NN fund PP
NP → DT NN fund VP
NP → DT NN fund PP
NP → DT JJ fund
NP → DT NN fund
NP → DT NN fund PP
NP → DT JJ JJ NN fund
NP → DT NN fund SBAR
NP → DT NN fund PP
NP → DT JJ fund
NP → DT NN fund

$p(\text{tree}) = \dots p(\triangle | S) \times p(\triangle | TO) \times p(\triangle | NP) \times p(\triangle | SBAR) \times \dots$
(oversimplified)

Too Many Rules ... But Luckily ...

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP
2 NP → NNP fund
2 NP → DT NPR NN fund
2 S → TO fund NP PP

NP → DT JJ NN fund
NP → DT NN fund PP
NP → DT ADJP NN fund
NP → DT JJ JJ NN fund
NP → DT NN fund SBAR
NP → DT NN fund PP
NP → DT NN fund VP
NP → DT NN fund PP
NP → DT JJ fund
NP → DT NN fund
NP → DT NN fund PP
NP → DT JJ JJ NN fund
NP → DT NN fund SBAR
NP → DT NN fund PP
NP → DT JJ fund
NP → DT NN fund

All these rules for fund – & other, still unobserved rules – are connected by the deep structure of English.

Rules Are Related

- fund behaves like a typical singular noun ...
 - ... or transitive verb ...
- one fact!
though PCFG represents it as many apparently unrelated rules.

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP


1 NP → DT NN fund
1 NP → DT JJ NN fund
1 NP → DT NN fund SBAR
1 NP → fund
1 NP-PRD → DT NN fund VP
1 NP → DT NN fund PP
1 NP → DT ADJP NN fund ADJP
1 NP → DT ADJP fund PP
1 NP → DT JJ fund PP-TMP
1 NP-PRD → DT ADJP NN fund VP
1 NP → NNP fund , VP ,
1 NP → PRP fund
1 S-ADV → DT JJ fund
1 NP → DT NNP NN fund
1 SBAR → NP MD fund NP PP
1 NP → DT JJ fund SBAR
1 NP → DT JJ NN fund SBAR
1 NP → DT NN fund SBAR
1 NP → NPS JJ NN fund
1 NP → DT JJ fund

Rules Are Related

- fund behaves like a typical singular noun ...
 - ... or transitive verb ...
- one more fact!
even if several more rules.
Verb rules are RELATED.
Should be able to PREDICT the ones we haven't seen.

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP
2 NP → NNP fund
2 NP → DT NPR NN fund
2 S → TO fund NP PP

1 NP → DT NN fund
1 NP → DT JJ NN fund
1 NP → DT NN fund SBAR
1 NP → fund
1 NP-PRD → DT NN fund VP
1 NP → DT NN fund PP
1 NP → DT ADJP NN fund ADJP
1 NP → DT ADJP fund PP
1 NP → DT JJ fund PP-TMP
1 NP-PRD → DT ADJP NN fund VP
1 NP → NNP fund , VP ,
1 NP → PRP fund
1 S-ADV → DT JJ fund
1 NP → DT NNP NN fund
1 SBAR → NP MD fund NP PP
1 NP → DT JJ fund SBAR
1 NP → DT JJ NN fund SBAR
1 NP → DT NN fund SBAR
1 NP → NPS JJ NN fund
1 NP → DT JJ fund



Rules Are Related

- fund behaves like a typical singular noun ...
 - ... or transitive verb ...
 - ... but as noun, has an idiosyncratic fondness for purpose clauses ...
- the ACL fund to put proceedings online
the old ACL fund for students to attend ACL
one more fact!
predicts dozens of unseen rules

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP
2 NP → NNP fund
2 NP → DT NPR NN fund
2 S → TO fund NP PP

1 NP → DT NN fund
1 NP → DT JJ NN fund
1 NP → DT NN fund SBAR
1 NP → fund
1 NP-PRD → DT NN fund VP
1 NP → DT NN fund PP
1 NP → DT ADJP NN fund ADJP
1 NP → DT ADJP fund PP
1 NP → DT JJ fund PP-TMP
1 NP-PRD → DT ADJP NN fund VP
1 NP → NNP fund , VP ,
1 NP → PRP fund
1 S-ADV → DT JJ fund
1 NP → DT NNP NN fund
1 SBAR → NP MD fund NP PP
1 NP → DT JJ fund SBAR
1 NP → DT JJ NN fund SBAR
1 NP → DT NN fund SBAR
1 NP → NPS JJ NN fund
1 NP → DT JJ fund

Rules Are Related

- fund behaves like a typical singular noun ...
 - ... or transitive verb ...
 - ... but as noun, has an idiosyncratic fondness for purpose clauses ...
 - ... and maybe other idiosyncrasies to be discovered, like unaccusativity ...
- NSF issued the grant
The grant issued today
↓ ???
NSF funded the grant
The grant funded today
- unlikely sentence, but if we do see it, is unaccusativity plausible? (vs. other parse)

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP
2 NP → NNP fund
2 NP → DT NPR NN fund
2 S → TO fund NP PP

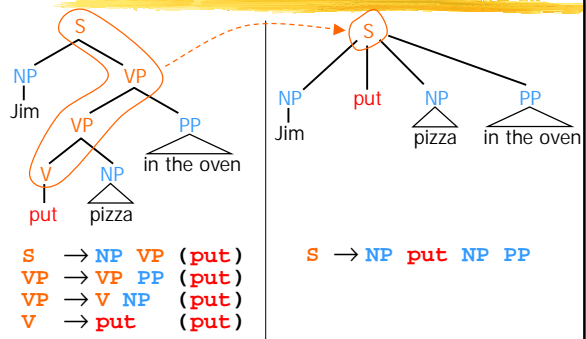
1 NP → DT NN fund
1 NP → DT JJ NN fund
1 NP → DT NN fund SBAR
1 NP → fund
1 NP-PRD → DT NN fund VP
1 NP → DT NN fund PP
1 NP → DT ADJP NN fund ADJP
1 NP → DT ADJP fund PP
1 NP → DT JJ fund PP-TMP
1 NP-PRD → DT ADJP NN fund VP
1 NP → NNP fund , VP ,
1 NP → PRP fund
1 S-ADV → DT JJ fund
1 NP → DT NNP NN fund
1 SBAR → NP MD fund NP PP
1 NP → DT JJ fund SBAR
1 NP → DT JJ NN fund SBAR
1 NP → DT NN fund SBAR
1 NP → NPS JJ NN fund
1 NP → DT JJ fund

All This Is Quantitative!

- fund behaves like a typical singular noun ...
 - ... or transitive verb ...
 - ... but as noun, has an idiosyncratic fondness for purpose clauses ...
 - ... and maybe other idiosyncrasies to be discovered, like unaccusativity ...
- how often?
and how does that tell us p(rule)?

26 NP → DT fund
24 NN → fund
8 NP → DT NN fund
7 NNP → fund
5 S → TO fund NP
2 NP → NNP fund
2 NP → DT NPR NN fund
2 S → TO fund NP PP
1 NP → DT NN fund
1 NP → DT JJ NN fund
1 NP → DT NN fund SBAR
1 NP → fund
1 NP-PRD → DT NN fund VP
1 NP → DT NN fund PP
1 NP → DT ADJP NN fund ADJP
1 NP → DT ADJP fund PP
1 NP → DT JJ fund PP-TMP
1 NP-PRD → DT ADJP NN fund VP
1 NP → NNP fund , VP ,
1 NP → PRP fund
1 S-ADV → DT JJ fund
1 NP → DT NNP NN fund
1 SBAR → NP MD fund NP PP
1 NP → DT JJ fund SBAR
1 NP → DT JJ NN fund SBAR
1 NP → DT NN fund SBAR
1 NP → NPS JJ NN fund
1 NP → DT JJ fund

Format of the Rules




Format of the Rules

Why use flat rules?

- Avoids silly independence assumptions: a win
 - Johnson 1998 →
 - New experiments

Our method likes them

- Traditional rules aren't systematically related
- But relationships exist among wide, flat rules that express different ways of filling same roles



$S \rightarrow NP \text{ put } NP \text{ PP}$


Format of the Rules

Why use flat rules?

- Avoids silly independence assumptions: a win
 - Johnson 1998 →
 - New experiments

Our method likes them

- Traditional rules aren't systematically related
- But relationships exist among wide, flat rules that express different ways of filling same roles



$S \rightarrow NP \text{ put } PP \text{ NP}$


Format of the Rules

Why use flat rules?

- Avoids silly independence assumptions: a win
 - Johnson 1998 →
 - New experiments

Our method likes them

- Traditional rules aren't systematically related
- But relationships exist among wide, flat rules that express different ways of filling same roles



$S \rightarrow NP, NP \text{ put } PP$


Format of the Rules

Why use flat rules?

- Avoids silly independence assumptions: a win
 - Johnson 1998 →
 - New experiments

Our method likes them

- Traditional rules aren't systematically related
- But relationships exist among wide, flat rules that express different ways of filling same roles



in short, flat rules are the locus of transformations

Format of the Rules

Why use flat rules?

- Avoids silly indep. assumptions: a win
 - Johnson 1998 →
 - New experiments

Our method likes them

- Traditional rules aren't systematically related
- But relationships exist among wide, flat rules that express different ways of filling same roles

flat rules are the locus of exceptions (e.g., put is exceptionally likely to take a PP, but not a second PP)

in short, flat rules are the locus of transformations

Intuition: Listing is costly and hard to learn. Most rules are derived.

Hey - Just Like Linguistics!

Lexicalized syntactic formalisms: CG, LFG, TAG, HPSG, LCFG ...

- Grammar = set of "lexical entries" very like flat rules
- Exceptional entries OK

flat rules are the locus of exceptions (e.g., put is exceptionally likely to take a PP, but not a second PP)

listed entries

derived entries

- Explain "coincidental" patterns of lexical entries: metarules/transformations/lexical redundancy rules

in short, flat rules are the locus of transformations

The Rule Smoothing Task

- **Input:** Rule counts (from parses or putative parses)
- **Output:** Probability distribution over rules
- **Evaluation:** Perplexity of held-out rule counts
 - That is, did we assign high probability to the rules needed to correctly parse test data?

The Rule Smoothing Task

- **Input:** Rule counts (from parses or putative parses)
- **Output:** Probability distribution over rules
- **Evaluation:** Perplexity of held-out rule counts

Rule probabilities: $p(s \rightarrow \text{NP put NP PP} \mid s, \text{put})$

Infinite set of possible rules; so we will estimate

$p(S \rightarrow \text{NP Adv PP put PP PP NP AdjP S} \mid s, \text{put})$
= a very tiny number > 0

Grid of Lexicalized Rules

$S \rightarrow \dots$	encourage	question	fund	merge	repay	remove
To — NP						
To — NP PP						
To AdvP — NP						
To AdvP — NP PP						
To — PP						
To — S						
NP — NP .						
NP — NP PP .						
NP Md — NP						
NP Md — NP PPTmp						
NP Md — PP PP						
NP — SBar .						
(etc.)						

$S \rightarrow \text{To fund NP PP}$ ("to fund projects with ease")
 $S \rightarrow \text{To merge NP PP}$ ("to merge projects with ease")

Training Counts

$S \rightarrow \dots$	encourage	question	fund	merge	repay	remove
To — NP	1	1	5	1	3	2
To — NP PP	1	1	2	2	1	1
To AdvP — NP						1
To AdvP — NP PP						1
NP — NP .		2				
NP — NP PP .	1					
NP Md — NP	1					
NP Md — NP PPTmp					1	
NP Md — PP PP						1
To — PP				1		
To — S	1					
NP — SBar .		2				
(other)						

Count of (word, frame)

Naive prob. estimates (MLE model)

$S \rightarrow \dots$	encourage	question	fund	merge	repay	remove
To — NP	200	167	714	250	600	333
To — NP PP	200	167	286	500	200	167
To AdvP — NP	0	0	0	0	0	167
To AdvP — NP PP	0	0	0	0	0	167
NP — NP .	0	333	0	0	0	0
NP — NP PP .	200	0	0	0	0	0
NP Md — NP	200	0	0	0	0	0
NP Md — NP PPTmp	0	0	0	0	200	0
NP Md — PP PP	0	0	0	0	0	167
To — PP	0	0	0	250	0	0
To — S	200	0	0	0	0	0
NP — SBar .	0	333	0	0	0	0
(other)	0	0	0	0	0	0

Estimate of $p(\text{frame} \mid \text{word}) * 1000$

TASK: counts \rightarrow probs ("smoothing")

$S \rightarrow \dots$	encourage	question	fund	merge	repay	remove
To — NP	142	117	397	210	329	222
To — NP PP	77	64	120	181	88	80
To AdvP — NP	0.55	0.47	1.1	0.82	0.91	79
To AdvP — NP PP	0.18	0.15	0.33	0.37	0.26	50
NP — NP .	22	161	7.8	7.5	7.9	7.5
NP — NP PP .	79	8.5	2.6	2.7	2.6	2.6
NP Md — NP	90	2.1	2.4	2.0	24	2.6
NP Md — NP PPTmp	1.8	0.16	0.17	0.16	69	0.19
NP Md — PP PP	0.1	0.027	0.027	0.038	0.078	59
To — PP	9.2	6.5	12	126	10	9.1
To — S	98	1.6	4.3	3.9	3.6	2.7
NP — SBar .	3.4	190	3.2	3.2	3.2	3.2
(other)	478	449	449	461	461	482

Estimate of $p(\text{frame} \mid \text{word}) * 1000$

Smooth Matrix via LSA / SVD, or SBS?

S → ...	encourage	question	fund	merge	repay	remove
To — NP	1	1	5	1	3	2
To — NP PP	1	1	2	2	1	1
To AdvP — NP						1
To AdvP — NP PP						1
NP — NP .		2				
NP — NP PP .	1					
NP Md — NP	1				1	
NP Md — NP PPTmp					1	
NP Md — PP PP						1
To — PP				1		
To — S	1					
NP — SBar .		2				
(other)						

Count of (word, frame)

Smoothing via a Bayesian Prior

- Choose grammar to maximize $p(\text{observed rule counts} \mid \text{grammar}) * p(\text{grammar})$
- grammar** = probability distribution over rules
- Our job:** Define $p(\text{grammar})$
- Question:** What makes a grammar likely, a priori?
- This paper's answer:** Systematicity. Rules are mainly derivable from other rules. Relatively few stipulations ("deep facts").

Only a Few Deep Facts

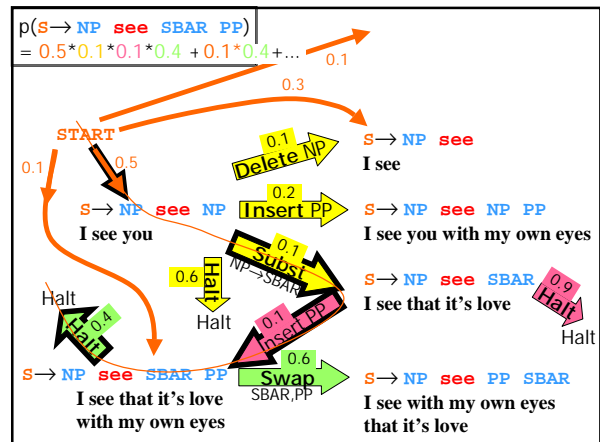
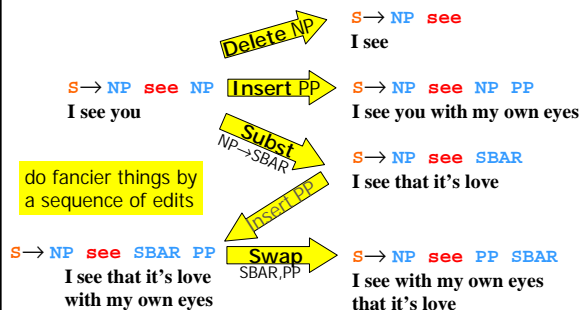
- fund** behaves like a transitive verb 10% of time ...
- and noun 90% of time ...
- ... takes purpose clauses 5 times as often as typical noun.

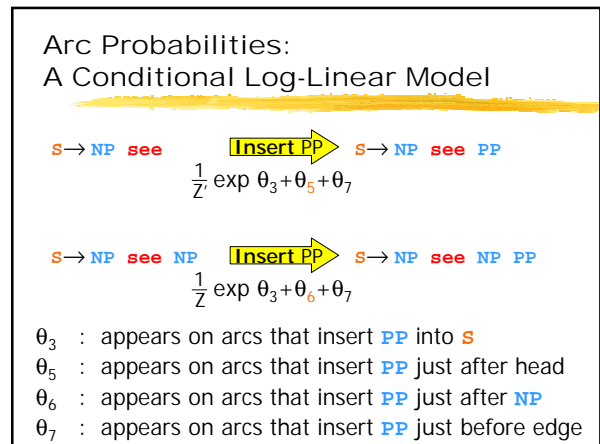
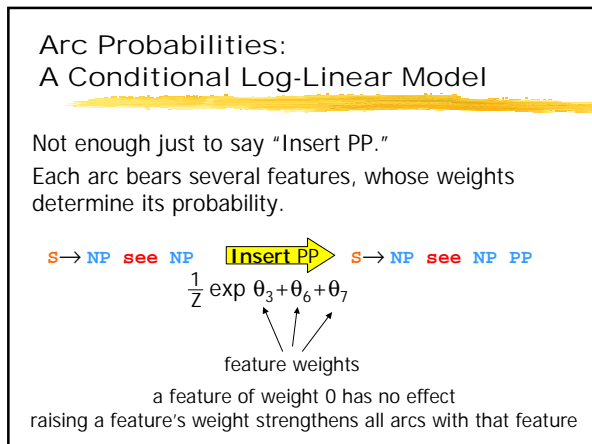
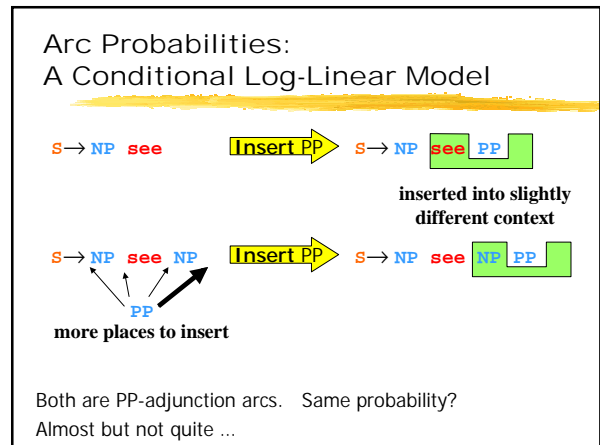
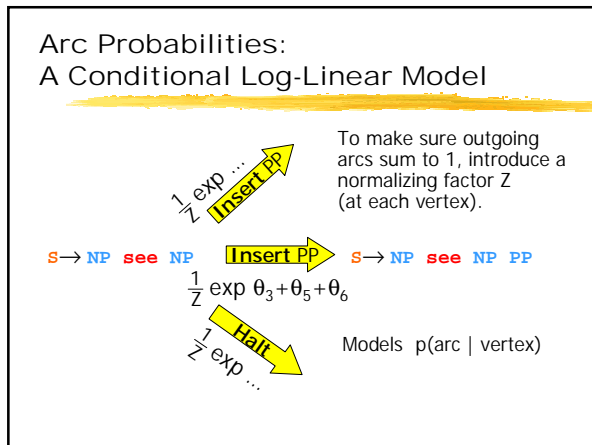
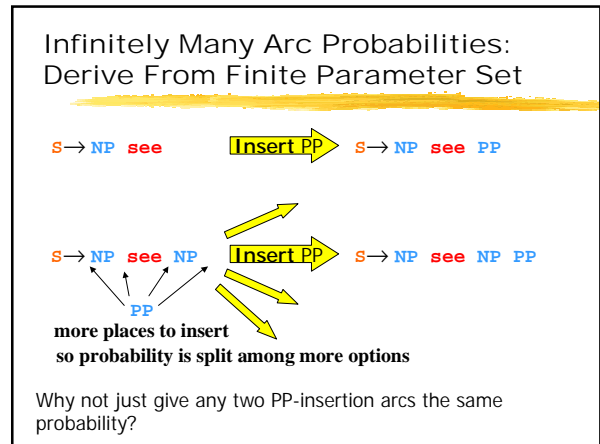
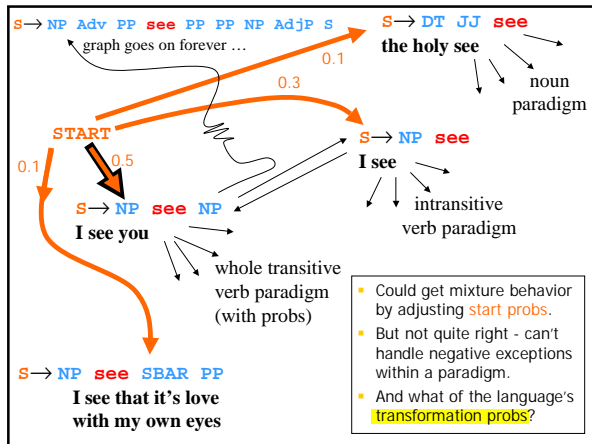
26	NP → DT fund
24	NN → fund
8	NP → DT NN fund
7	NNP → fund
5	S → TO fund NP
2	NP → NNP fund
2	NP → DT NPR NN fund
2	S → TO fund NP PP
1	NP → DT JJ NN fund
1	NP → DT NN JJ fund
1	NP → DT ADJP NP fund
1	NP → DT JJ JJ NN fund
1	NP → DT NN fund SBAR
1	NP → fund
1	NP-PHD → DT NN fund VP
1	NP → DT NN fund PP
1	NP → DT ADJP NN fund ADJP
1	NP → DT ADJP fund PP
1	NP → DT JJ fund VP-OMP
1	NP-PHD → DT ADJP NN fund VP
1	NP → NP fund VP
1	NP → PRPS fund
1	S-ADV → DT JJ fund
1	NP → DT NNP NNP fund
1	SBAR → NP NP fund NP PP
1	NP → DT JJ JJ fund SBAR
1	NP → DT JJ NP fund SBAR
1	NP → DT NP fund
1	NP → NP JJ NP fund
1	NP → DT JJ fund

Smoothing via a Bayesian Prior


- Previous work (several papers in past decade):
 - Rules should be few, short, and approx. equiprobable
 - These priors try to keep rules **out** of grammar
 - Bad idea for lexicalized grammars ...
- This work:
 - Prior tries to get related rules **into** grammar
 - transitive → passive at $\approx 1/20$ the probability
 - NSF spraggles the project → The project is spraggled by NSF
 - Would be weird for the passive to be missing, and prior knows it!
 - In fact, weird if $p(\text{passive})$ is too far from $1/20 * p(\text{active})$
 - Few facts, not few rules!

Simple Edit Transformations






Arc Probabilities: A Conditional Log-Linear Model



- θ_3 : appears on arcs that insert **PP** into **S**
- θ_5 : appears on arcs that insert **PP** just after head
- θ_6 : appears on arcs that insert **PP** just after **NP**
- θ_7 : appears on arcs that insert **PP** just before edge

Arc Probabilities: A Conditional Log-Linear Model



- These arcs share most features.
- So their probabilities tend to rise and fall together.
- To fit data, could manipulate them independently (via θ_5, θ_6).


Prior Distribution

- PCFG grammar is **determined** by $\theta_0, \theta_1, \theta_2, \dots$

Universal Grammar




Instantiated Grammar



Prior Distribution


- Grammar is **determined** by $\theta_0, \theta_1, \theta_2, \dots$
- Our prior: $\theta_i \sim N(0, \sigma^2)$, IID
- Thus: $-\log p(\text{grammar}) = c + (\theta_0^2 + \theta_1^2 + \theta_2^2 + \dots) / \sigma^2$
- So good grammars have few large weights.
- Prior prefers one generalization to many exceptions.

Arc Probabilities: A Conditional Log-Linear Model



To raise both rules' probs, cheaper to use θ_3 than both θ_5 & θ_6 .
This generalizes – also raises other cases of PP-insertion!

Arc Probabilities: A Conditional Log-Linear Model



To raise both probs, cheaper to use θ_3 than both θ_{82} & θ_{84} .
This generalizes – also raises other cases of PP-insertion!

Reparameterization

- Grammar is determined by $\theta_0, \theta_1, \theta_2, \dots$
- A priori, the θ_i are normally distributed
- We've reparameterized!
- The parameters are feature weights θ_i , not rule probabilities
- Important tendencies captured in big weights
 - Similarly: Fourier transform – find the formants
 - Similarly: SVD – find the principal components
 - It's on this deep level that we want to compare events, impose priors, etc.

```

{ (S
  (NP-SBJ (JJ big) (NN indexes) (NRP Bankers) (NRP Trust) (NNP Co.) )
  (ADVP (RB also) )
  (VP (VBE uses)
    (NP (NNS futures) )
    (PP-LOC (IN in)
      (NP
        (NP (DT a) (NN strategy) )
        (SBJ
          (S
            (PP (IN on)
              (NP (NN average) ))
            (NP-SBJ [-NONSE- +T=1] )
            (VP (VBE has)
              (PP (VBN added)
                (NP (CD one) (NN percentage) (NN point) )
                (PP-CLR (TO to)
                  (NP
                    (NP (PPOS its) (JJ enhanced) (NN fund) (POS 's) )
                    (NNS returns) ))))))))
      ( . . ) ) )

```

```

{ (S
  (NP (JJ big) (NN indexes) (NPA (NRP Bankers) (NRP Trust) Co.) )
  (ADVP (RB also)
  uses
  (NP futures)
  (PP-LOC in
    (NP
      (NP (DT a)
        strategy
        (SBJ that
          (S
            (PP on (NP average) )
            (VBE has)
            added
            (NP (CD one) (NN percentage) point)
            (PP to
              (NP
                (NPS (NP (PPOS its) (JJ enhanced) fund) 's)
                returns) ))))))
    ( . . ) )

```

keyword	lhs	rhs
big	JJ	→
indexes	NP	→ JJ → NPS
Bankers	NP	→
Trust	NP	→
Co.	NPS	→ NRP NP →
also	ADVP	→
uses	S	→ NP ADVP → NP PP-LOC →
futures	NP	→
in	PP-LOC	→ → NP
a	DT	→
strategy	NP	→ DT → SBJ
that	SBJ	→ → S
on	PP	→ → NP
average	NP	→
has	VBE	→
added	S	→ PP VBE → NP PP
one	CD	→
percentage	NN	→
point	NP	→ CD NP →
to	PP	→ → NP
its	PPOS	→
enhanced	NP	→ JJ →
fund	NP	→ PPOS JJ →
's	NPS	→ NP →
returns	NP	→ NPS →
.	.	→


Other models of this string:
max-likelihood
n-gram
Collins arg/adj
hybrids

Simple Bigram Model (Eisner 1996)

- A parser assumes tree is probable if its component rules are:



- Try assuming rule is probable if its component bigrams are:



$$p(A \mid \text{start}) \times p(B \mid A) \\ \times p(C \mid B) \times p(\text{---} \mid C) \\ \times p(D \mid \text{---}) \times p(\text{stop} \mid D)$$

- Markov process, 1 symbol of memory; conditioned on L, w, side of —
- One-count backoff to handle sparse data (Chen & Goodman 1996)

$$p(L \rightarrow A B C \text{---} D \mid w) = p(L \mid w) \cdot p(A B C \text{---} D \mid L, w)$$

headword	basic	flat	non-flat
big	JJ →		
indicate	VP → JJ (aw ---) NNS		
bankers	NP →		
trust	NP →		
co.	NP → NN NP (aw ---)		
edit	VP → (aw ---)		
name	N → VP (aw ---) NP PP (LOC)		
features	NP → (aw ---)		
is	VP-LOC → (aw ---) NP		
a	DT →		
strategy	NP → (aw ---) (aw ---) (aw ---)		
that	SBAR → (aw ---) (aw ---) (aw ---)		
on	PP → (aw ---) NP		
average	NP → (aw ---)		
has	VBZ →		
added	V → VP (aw ---) (aw ---) (aw ---) (aw ---) (aw ---) (aw ---)		
one	DT →		
percentage	NN →		
print	VP → NP NP (aw ---)		
to	PP → (aw ---) NP		
its	POS →		
released	VP → JJ		
fund	NP → NN (aw ---) (aw ---)		
's	POS → NP (aw ---)		
returns	VP → NP (aw ---)		
.	.		

Figure 8.7: A version of Fig. 8.3 if basic-internal brackets are retained at step 6 of data compression (98.2).

Use "non-flat" frames?
Extra training info.
For test, sum over all bracketings.

Perplexity: Predicting test frames

	basic	
	flat	non-flat ^b
Treebank	∞	∞
1-gram	1774.9	86435.1
2-gram	135.2	199.3
3-gram	136.5	177.4
Collins ^c	363.0	494.5
transformation	108.6	

20% further reduction

Can get big perplexity reduction just by flattening.



from previous lit.


Perplexity: Predicting test frames

	basic		Treebank/Markov	
	flat	non-flat ^b	Katz flat	one-count ^d flat
Treebank	∞	∞		
1-gram	1774.9	86435.1	340.9	160.0
2-gram	135.2	199.3	127.2	116.2
3-gram	136.5	177.4	132.7	123.5
Collins ^c	363.0	494.5	197.9	
transformation	108.6			
averaged ^e	102.3			

best model with transformations

from previous lit





Forced matching task

- Test model's ability to extrapolate novel frames for a word
- Randomly select **two** (word, frame) pairs from test data
 - ... ensuring that neither frame was ever seen in training
- Ask model to choose a matching:

word 1 — frame A

word 2 — frame B

word 1 \times frame A

word 2 \times frame B

i.e., does frame A look more like word 1's known frames or word 2's?
- 20% fewer errors than bigram model

Graceful degradation

	basic		non-associative backoff		
	flat	non-flat	flat	non-const	non-flat
1 gram	1891.2	96338.8	405.1	186.3	223.1
2 gram	182.3	236.8	185.2	138.6	286.8
3 gram	181.9	211.8	186.8	145.7	289.1
subtree	414.5	589.4	242.0		
lexical	138.6				
combined	118.6				

Twice as much data
But no transformations

	basic		non-associative backoff		
	flat	non-flat	flat	non-const	non-flat
1 gram	1774.8	96430.1	340.9	180.0	185.2
2 gram	138.3	189.3	127.2	138.2	154.7
3 gram	136.5	177.4	125.7	125.5	154.8
subtree	385.0	484.3	187.0		

Summary: Reparameterize PCFG in terms of deep transformation weights, to be learned under a simple prior.

- Problem:** Too many rules!
 - Especially with lexicalization and flattening (which help).
 - So it's hard to estimate probabilities.
- Solution:** Related rules tend to have related probs
 - POSSIBLE* relationships are given a priori
 - LEARN* which relationships are strong in this language (just like feature selection)
- Method has connections to:
 - Parameterized finite-state machines (Monday's talk)
 - Bayesian networks (inference, abduction, explaining away)
 - Linguistic theory (transformations, metarules, etc.)

FIN