

# General techniques for creating treebanks

Fei Xia

University of Washington

[fxia@u.washington.edu](mailto:fxia@u.washington.edu)

Jan 6, 2008

# Motivation

- Treebanks are valuable resources for NLP:
  - Word segmentation
  - POS tagging
  - Chunking
  - Parsing
  - Grammar extraction
  - ...
- Problem: Creating treebanks is still an art, not a science.

# My experience with treebanks

- As a member of the Chinese Penn Treebank (CTB) project: 1998-2000
  - Designed annotation guidelines for segmentation, POS tagging, and bracketing (with Nianwen Xue).
  - Project manager in the first year
  - Organized several workshops on Chinese NLP
- As a user of treebanks
  - Work on grammar extraction
  - Work on DS=>PS conversion, Chinese POS tagging, etc.

# Current work

- RiPLEs project:
  - Plan to build mini-treebanks for 5-10 languages
  - Each treebank has 300-1000 sentences
- The Hindi/Urdu treebank project:
  - Joint work with IIIT, Univ of Colorado, Columbia Univ, and UMass

# Outline

- Treebank overview
- Main issues
- Case study: The Chinese Penn Treebank
- Creating a Hindi/Urdu treebank

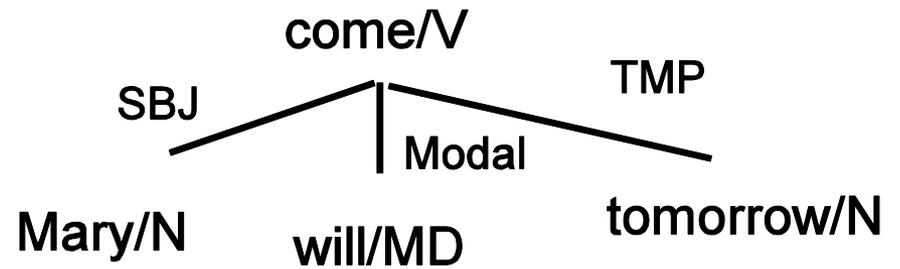
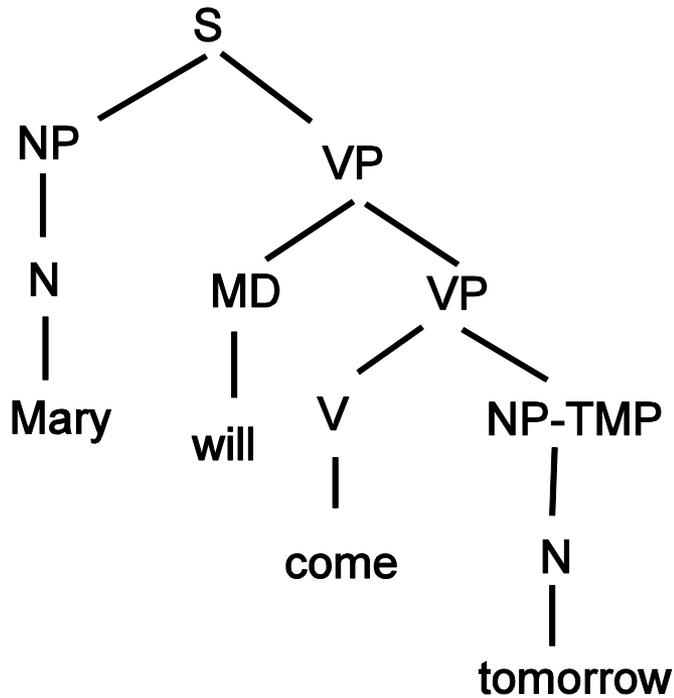
# Treebank overview

# Types of treebanks

- (Syntactic) treebanks:
    - Phrase-structure treebanks
    - Dependency treebanks
  - PropBank: predicate-argument structure for verbs
  - NomBank: predicate-argument structure for nouns
  - Discourse Treebank: discourse structure
- ➔ In this talk, we will focus on syntactic treebanks

# Syntactic treebank

Mary will come tomorrow.



# A Propbank example

Roles for “give”:

Arg0: giver

Arg1: thing given

Arg2: entity given to

*The executives gave the chefs a standing ovation.*

Arg0: *The executives*

REL: *gave*

Arg2: *the chefs*

Arg1: *a standing ovation*

# A NomBank example

Roles for “gift” (the nominalization of “give”)

Arg0: giver

Arg1: thing given

Arg2: entity given to

Nancy's gift from her cousin was a complete surprise.

Arg0: *her cousin*

REL: *gave*

Arg2: *Nancy*

Arg1: *gift*

# A Discourse Treebank example

Argument Structure of Explicit/Implicit Connectives (spans):

She hasn't played any music since the earthquake hit.

**She hasn't played any music** since **the earthquake hit.** (temporal)

We asked police to investigate why they are allowed to distribute the flag in this way. It should be considered against the Law.

**We asked police to investigate why they are allowed to distribute the flag in this way.** Implicit=because **It should be considered against the Law.** (Causal)

# Existing treebanks

- Brown corpus, English Penn Treebank
- Chinese, Arabic, Korean Penn Treebanks
- Prague Dependency Treebank (Czech)
- Tiger Treebank (German)
- Treebanks for Bulgarian, French, Italian, Japanese, Polish, Portuguese, Spanish, Swedish, Turkish, ...

# Beyond syntactic treebanks

- PropBank: Penn English, Chinese, ...
- NomBank: Penn English, Chinese, ...
- Discourse Treebank: Penn English
- Parallel Treebanks: Penn Chinese-English, Arabic-English treebanks

# Outline

- Treebank overview
- Main issues
- Case study: The Chinese Penn Treebank
- Future work: Creating a Hindi/Urgu treebank

# Main issues

- Creating guidelines
- Involving the community
- Forming a team
- Selecting data
  
- Role of processing NLP tools
- Quality control
- Distributing the data
- Future expansion of the treebanks

# Highlights

- Detailed, “searchable” guidelines are important
  - Ex: the CTB’s guidelines have 266 pages
- Guidelines take a lot time to create, and changes to the guidelines after annotation starts are inevitable.
  - An important issue: How to update the annotation when the guidelines changes?
- It is a good idea to involve the annotators while creating the guidelines
- Define high-level guiding principles, which lower-level decisions should follow naturally
  - ➔ reduce the number of decisions that annotators have to memorize

# A high-quality treebank should be

- Informative: it provides the info needed by its users
  - Morphological analysis: lemma, derivation, inflection
  - Tagging: POS tags
  - Parsing: phrase structure, dependency relation, etc.
  - ...
- Consistent: The consistency is important for
  - training
  - evaluation
  - conversion
- Reasonable annotation speed
- Some tradeoff is needed:
  - Ex: walked/VBD vs. walk/V+ed/pastTense

# An example: the choice of the tagset

- Large tagset vs. small tagset
- Types of tags:
  - POS tags: e.g., N, V, Adj
  - Syntactic tags: e.g., NP, VP, AdjP
  - Function tags: e.g., -TMP, -SBJ
    - Temporal NPs vs. object NPs
    - Adjunct/argument distinction
  - Empty categories: e.g., \*T\*, \*pro\*
    - Useful if you want to know subcategorization frames, long-distance dependency, etc.

# When there is no consensus

- Very often, there is no consensus on various issues
- Try to be theory-neutral: linguistic theories keep changing.
- Study existing analyses and choose the best ones
- Make the annotation rich enough so that it is easy to convert the current annotation to something else

# Two common questions

- Grammars vs. annotation guidelines
- Phrase structure vs. dependency structure

# Writing grammar vs. creating annotation guidelines

- Similarity:
  - Both require a thorough study of the linguistic literature and a careful selection of analyses for common constructions
- Differences:
  - Annotation guidelines can leave certain issues undecided.
    - Ex: argument / adjunct distinction
  - Annotation guidelines need to have a wide coverage, including the handling of issues that are not linguistically important
    - Ex: attachment of punctuation marks
- Currently, they do not interact much. We should increase the interaction between the two.
- Existing work:
  - Treebanking with existing grammars
  - Extracting grammars from treebanks

# Treebanking with a pre-existing grammar

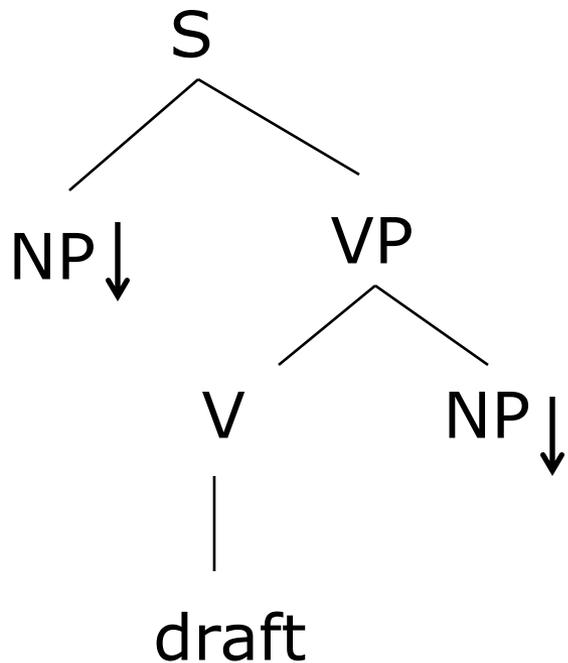
- Ex: Redwoods HPSG treebank
- Procedure:
  - Use the grammar to parse the sentences
  - Correct the parsing output
- Advantage:
  - The analyses used by the treebank are as well-founded as the grammar.
  - As the grammar changes, the treebank could potentially be automatically updated.
- Disadvantage:
  - It requires a large-scale grammar.
  - The treebank could be heavily biased by the grammar

# Extracting grammars from treebanks

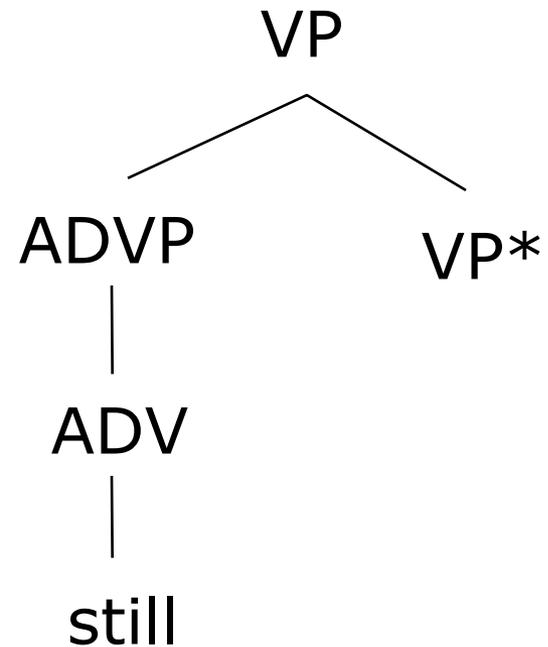
- A lot of work on grammar extraction
  - Different grammar formalisms: e.g., CFG, LTAG, CCG, LFG
- Compared to hand-crafted grammars
  - Extracted grammars have better coverage and include statistical information, both are useful for parsing.
  - Extracted grammars are more noisy and lack rich features.

# Extracting LTAGs from Treebanks

Initial tree:

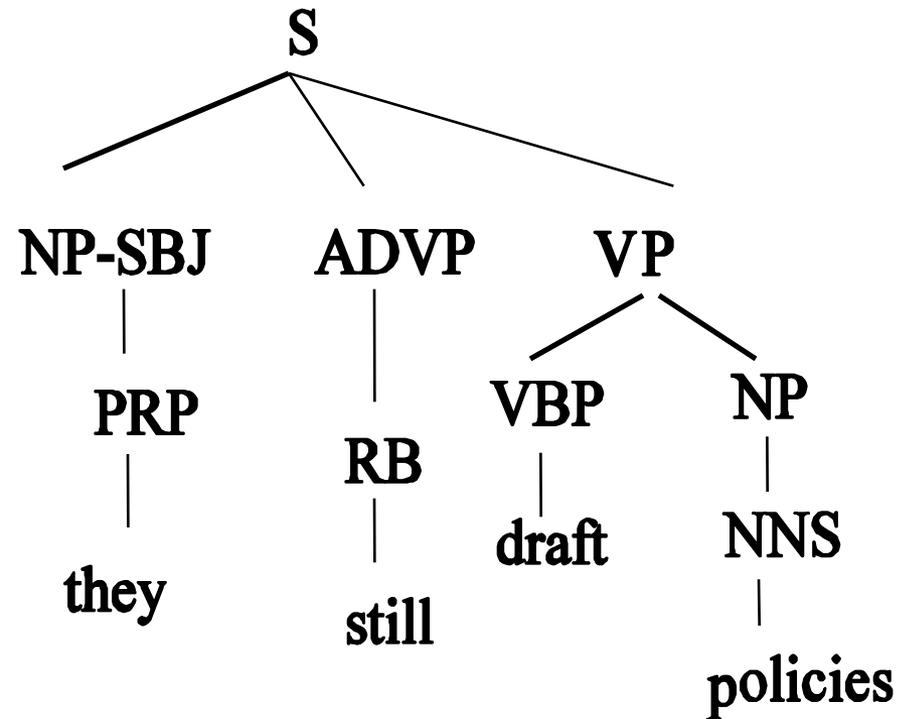


Auxiliary tree:

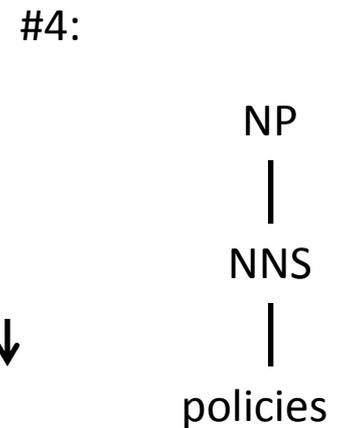
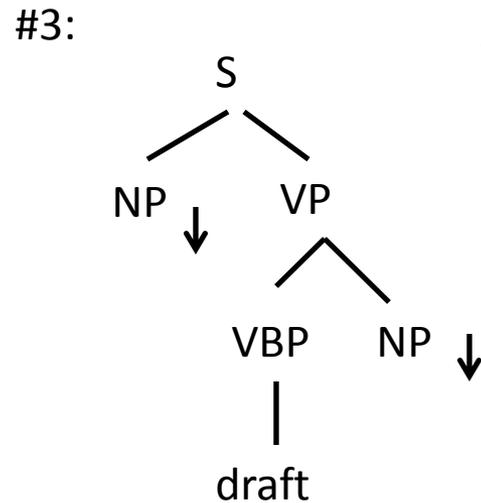
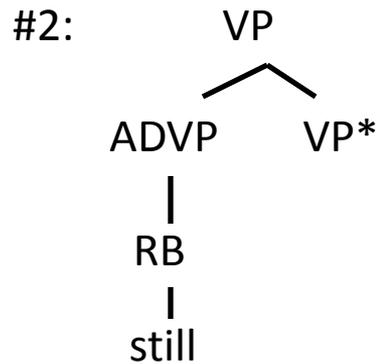
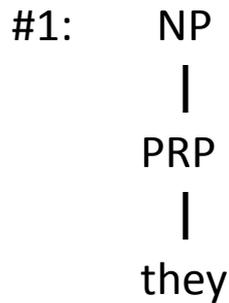


➔ Arguments and adjuncts are in different types of elementary trees

# The treebank tree



# Extracted grammar



We ran the system (LexTract) to convert treebanks into the data that can be used to train and test LTAG parsers.

# Two common questions

- Grammars vs. annotation guidelines
  - Grammars and treebank guidelines are closely related.
  - There should be more interaction between the two.
- Phrase structure vs. dependency structure

# Information in PS and DS

	PS (e.g., PTB)	DS (some target DS)
POS tag	yes	yes
Function tag (e.g., -SBJ)	yes	yes
Syntactic tag	yes	no
Empty category and co-indexation	Often yes	Often no
Allowing crossing	Often no	Often yes

# PS or DS for treebanking?

- PS treebank is good for phrase structure parsing
- Dependency treebank is good for dependency parsing.
- Ideally, we want to have both. But annotating both would be too expensive.
  
- Conversion algorithms between the two have been proposed, but they are far from perfect.
  
- Remedy: Make annotations (just) rich enough to support both.
  - Ex: mark the head in PS

# PS → DS

- For each internal node in the PS
  - (1) Find the head child
  - (2) Make the non-head child depend on head-child
- For (1), very often people use a head percolation table and functional tags.

# DS → PS

- (Collins, Hajič, Ramshaw and Tillmann, 1999)
- (Xia and Palmer, 2001)
- Both are based on heuristics.
- Need to handle non-projectivity and ambiguity.

# Main issues

- Creating guidelines
  - Involving the community
  - Forming the team
  - Selecting data
- 
- Role of processing NLP tools
  - Quality control
  - Distributing the data
  - Future expansion of the treebanks

# Community involvement

- Before the project starts, find out
  - what the community needs
  - whether there are existing resources (guidelines, tools, etc.)
- During the project, ask for feedback on
  - new guidelines
  - annotation examples
  - tools trained on preliminary release
- Don't be discouraged by negative feedback

# Forming the team

- Computational linguists:
  - Create annotation guidelines
  - Make/use NLP tools for preprocessing, final cleaning, etc.
- Linguistics experts
  - Help to create annotation guidelines
- Annotators
  - Training on linguistics and NLP is a big plus
- Advisory board: experts in the field

# Annotators

- Linguists do make good annotators!
- Training annotators well takes a very long time
- Keeping trained annotators is not easy
  - Full time is good (combo annotation and scripting, error searching, workflow, etc.)
- Good results are possible:
  - Ex: IAA for CTB is 94%

# Selecting data

- Permission for distribution
- The data should be a good sample of the language.
- Data from multiple genres?
  - Ex: 500K words from one genre, 250K from one genre and 250K from another, or other combinations?
- Active learning?
  - To select the hardest sentences for annotation.

# Roles of tools

- Annotation tools
- Preprocessing tools
- Other tools:
  - Corpus search tools: e.g., tgrep
  - Conversion tools:
  - Error detection tools:

# Preprocessing tools (e.g., taggers, parsers)

- Use pre-existing tools or train new ones:
  - train a tool with existing data
  - preprocess new data with the tool
  - manually check and correct errors
  - Add the new data to the training data
  - Repeat the procedure
- It can speed up annotation and improve consistency
- However, the tools introduce a big bias to the treebanks, as annotators often fail to correct the mistakes introduced by the tools.
- Quality control is essential.

# Quality control

- Human errors are inevitable
- Good guidelines, well-trained annotators, easy-to-use annotation tools, search tools, ...
- Inter-annotator agreement should be monitored throughout the project.
- Detecting annotation errors using NLP tools
- Feedback from the user
  - From parsing work
  - From PropBank work
  - From grammar extraction work

# Inter-annotator agreement

- Procedure:
  - Randomly select some data for double annotation
  - Compare double annotation results and create gold standard
  - Calculate annotation accuracy (e.g., f-measure) and inter-annotator agreement
- Possible reasons of the disagreement:
  - Human errors
  - Problems in annotation guidelines
    - ➔ modify the guidelines if needed

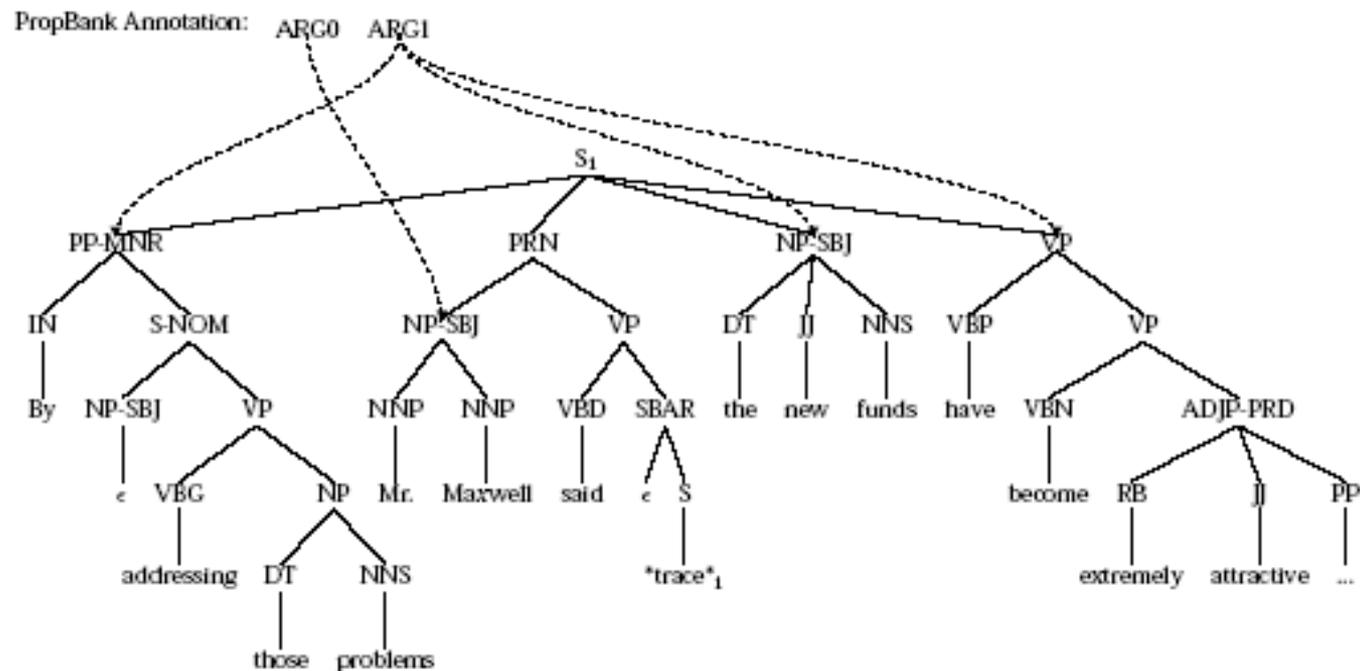
# Distributing the data

- Find a good collaborator: e.g., LDC
- Multiple releases
  - Preliminary releases for feedback
  - Later release with more data and/or fewer errors
- Presentations at major conferences

# Expanding the treebank

- More data
- More genres
- Other layers of information
  - Ex: PropBank, NomBank, Discourse Treebank on top of treebanks
  - The choice made by the treebank could affect new layers

# Treebank-PropBank Reconciliation



Problem: One PropBank argument can involve many parse nodes

Solution: Single argument – single parse node analysis

# Chinese Treebank (CTB)

# CTB: overview

- Website: <http://verbs.colorado.edu/chinese>
- Started in 1998 at Penn
- Supported by DOD, NSF, DARPA
- Now a nearly 1M word Chinese corpus
  - Segmented, POS-tagged, syntactically bracketed
  - Phrase structure annotation
  - Inter-annotator agreement: 94%
  - On-going expansion, 1.1M words planned
- Additional layers of annotation
  - Propbank/Nombank, Discourse annotation

# CTB: Milestones

Version	Year	Quantity (words)	Source	Propbank/ Nombank	Discourse annotation
CTB1.0	2001	100K	Xinhua	yes	Pilot
CTB3.0	2003	250K	+HK News	yes	no
CTB4.0	2004	400K	+Sinorama	yes	no
CTB5.0	2005	500K	+Sinorama	yes	no
CTB6.0	2007	780K	+BN	yes	no
CTB7.0*	2008*	950K	+BC	no	no
CTB8.0*	2009*	1.1M	+??	no	no

# An example

Raw data:

他还提出一系列具体措施和政策要点。

A tree in CTB-1:

```
(IP (NP-SBJ (PN 他))
  (VP (ADV (AD 还))
    (VP (VV 提出)
      (NP-OBJ (QP (CD 一)
        (CLP (M 系列)))
        (NP (NP (ADJP (JJ 具体))
          (NP (NN 措施)))
          (CC 和)
          (NP (NN 政策)
            (NN 要点))))))
    (PU 。 ))
```

# CTB-1

- The tasks:
  - Laying the good foundation for the whole project: creating guidelines, forming the team, getting feedback from the community, etc.
  - Annotating 100K-word Xinhua News
- Main stages:
  - Stage 0 (6/98 - 8/98): Feasibility study
  - Stage 1 (9/98 – 3/99): Word segmentation and POS tagging.
  - Stage 2 (4/99 – 9/00): Bracketing
  - Stage 3 (6/00 – 12/00): Preliminary release of CTB-1

# The team for CTB1

- PIs: Martha Palmer, Mitch Marcus, Tony Kroch
- Linguistic consultants: Tony Kroch, Shizhe Huang
- Project managers and guideline designers: Fei Xia, Nianwen Xue
- Annotators: Nianwen Xue, Fu-dong Chiou
- Programming support: Zhibiao Wu

# Community involvement

- Two workshops:
  - 06/1998: 3-day workshop at UPenn
  - 10/2000: 1-day workshop at Hong Kong (during ACL-2000)
- Three meetings:
  - 08/1998: At ACL98 in Montreal, Canada
  - 11/1998: At ICCIP98 in Beijing, China
  - 06/1999: At ACL99 in Maryland, US
- Two preliminary releases: in 6/00 and 12/00 by LDC

# Challenges in designing guidelines for Chinese

- No natural delimiters between words in written text
- Very little, if any, inflectional morphology
  - Ex: No (explicit) tense, gender, person, number, agreement morphology
- Many open questions about syntactic constructions
- Little consensus on standards and analyses within the Chinese linguistics/NLP community

# Word segmentation

日文章鱼怎么说？

日文 章鱼 怎么说？

Japanese octopus how say

“How to say octopus in Japanese?”

日 文章 鱼 怎么说？

Japan article fish how say

“? How to say fish in Japanese articles?”

# POS: verb or noun

美国 将 与 中国 讨论 贸易 赤字

U.S. will with China discuss/discussion trade deficit

“The U.S. will discuss trade deficit with China.”

美国 将 与 中国 就 贸易 赤字 进行 讨论

U.S. will with China regarding trade deficit engage discuss/discussion

“The U.S. will engage in a discussion on the trade deficit with china.”

# Verb or preposition?

Google 用 30 亿 现金 收购 Double Click  
Google use/with 30 100-million cash buy Double Click

Google used 3 billion cash to buy Double Click  
Google bought Double Click with 3 billion cash

# Main issue in POS tagging

Should POS tags be determined by distribution or by meaning?

Our approach:

- Use distribution (not meaning) for POS tagging
- Provide detailed tests for confusing tag pairs: e.g., (noun, verb)

# Bracketing example: Sentential complement or object control?

他            希望    她            做            作业  
he/him      hope      she/her      do            homework

*“He hopes that she will do her homework.”*

他            逼            她            做            作业  
he/him      force      she/her      do            homework

*“He forced her to do her homework.”*

NP

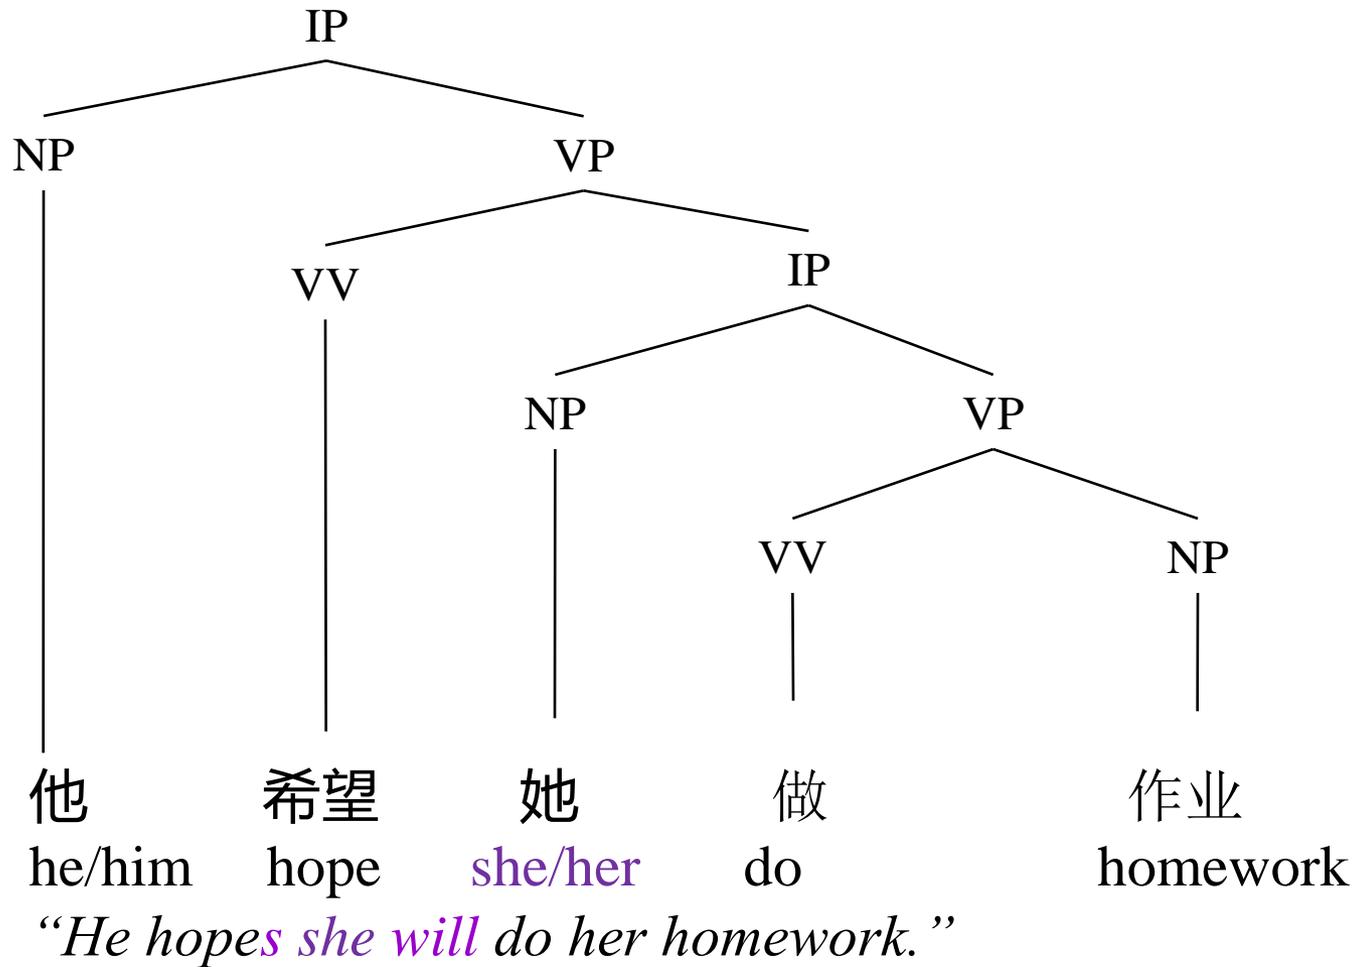
V

NP

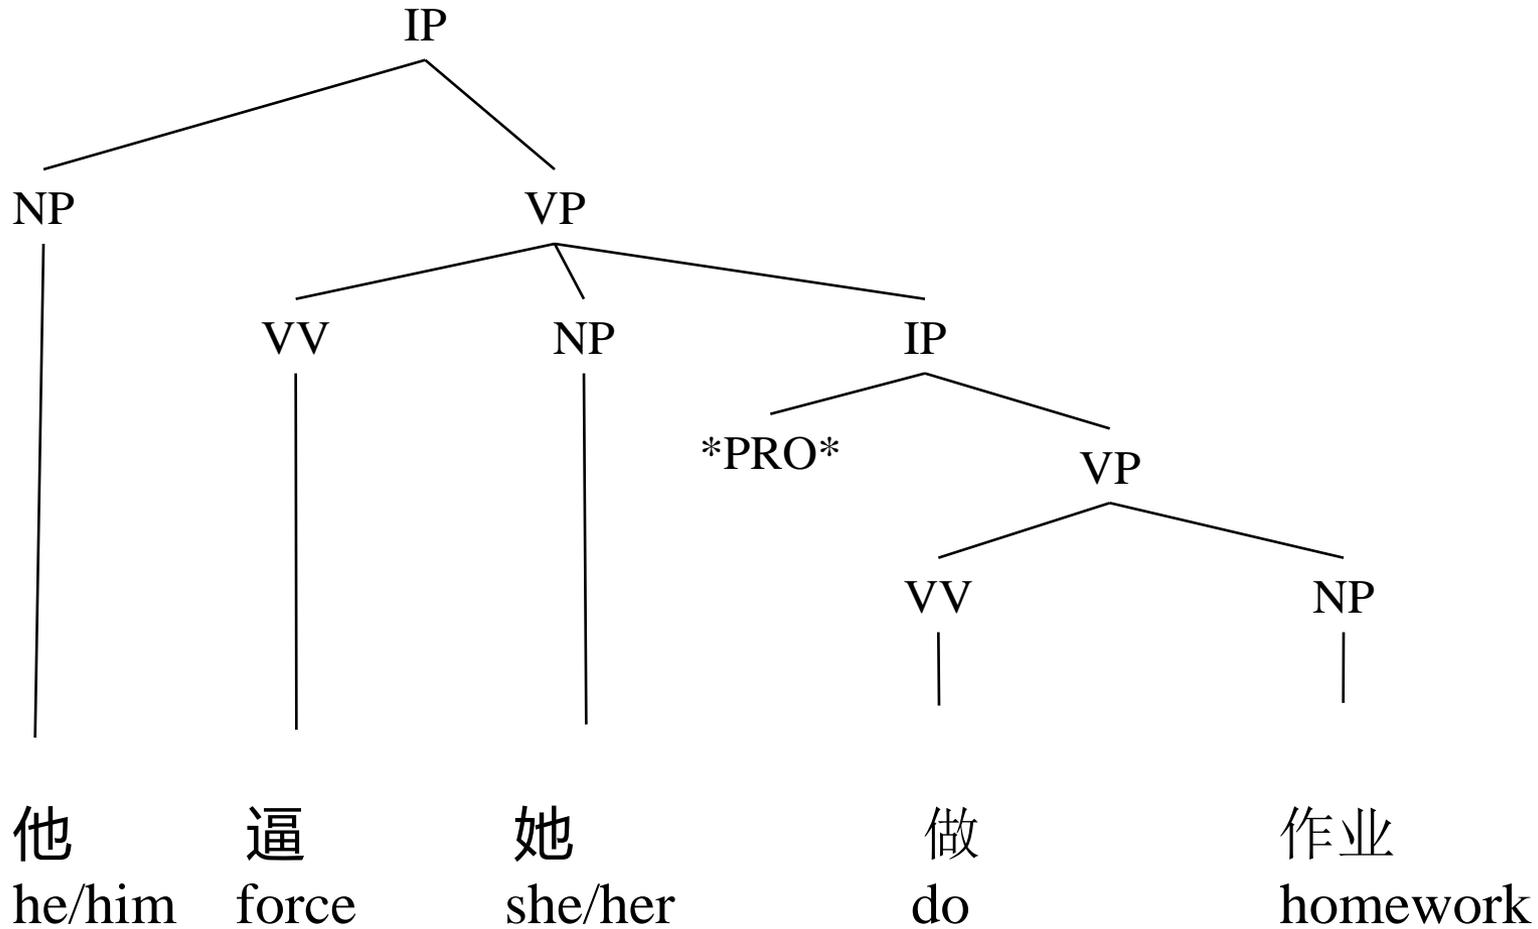
V

NP

# Sentential complement



# Object control



*"He forced her to do her homework."*

# Tests for sentential complement vs object control

For verb v1 in “NP1 v1 NP2 v2 NP3”:

- Can it take an existential construction as its complement?
- Can it take an idiom as its complement?
- Can it take a BEI construction as its complement?
- Can it take a topic construction as its complement?
- Can the complement clause have an aspectual marker?

Yes      Sentential complement

No      Object control

# Good annotation guidelines

- Correctness / plausibility
- Convertibility
- Consistency
- Searchability
- Wide coverage
- Annotation speed

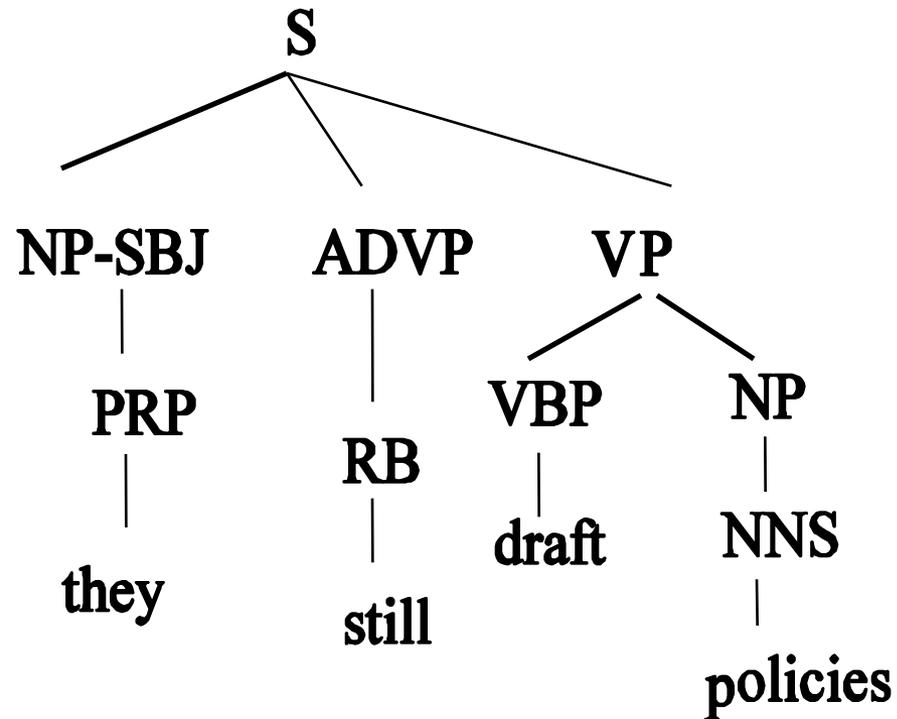
# Revision of guidelines

- First draft before annotation starts
- Second draft after the 1<sup>st</sup> pass of annotation
- Final version after the 2<sup>nd</sup> pass of annotation
- Three sets of guidelines
  - Segmentation: 31 pages
  - POS tagging: 44 pages
  - Bracketing: 191 pages

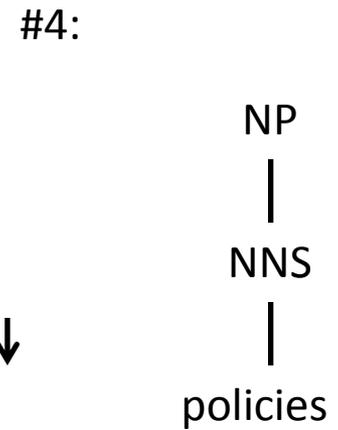
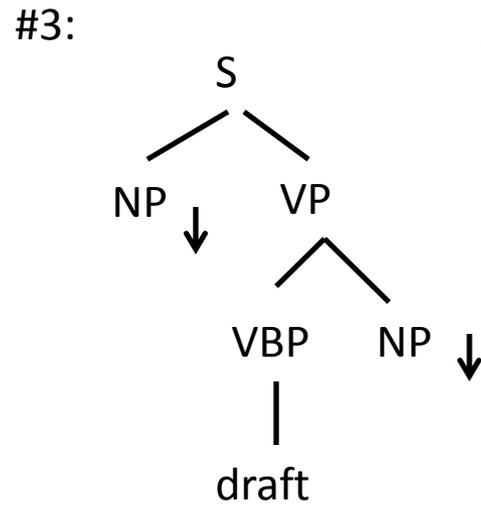
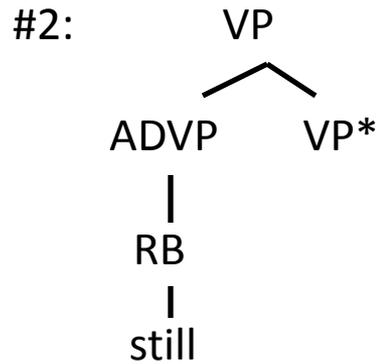
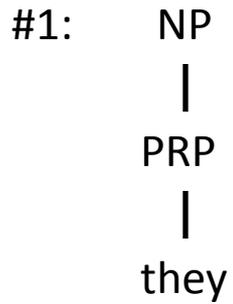
# Quality control

- Inter-annotator agreement:
  - Double annotation:
  - Inter-annotation agreement: 94%
  - Compared against the gold standard: 95-99%

# The treebank tree



# Extracted grammar



# Detecting annotation errors using NLP tools

- A tool, LexTract, that extracts tree-adjoining grammars from treebanks
- Experiments:
  - run LexTract on the treebank and get a grammar  $G$
  - mark each “rule” in  $G$  as correct or incorrect
  - correct trees in the treebank that generate the wrong “rules” in  $G$
- Results:
  - Detect about 550 errors in CTB-1
  - A good grammar with statistical info

# Preprocessing

	preprocessing	prec/recall	speed
set 1	–	–	240 words/hr
set 2	with parser	76.7%/75.4%	412 words/hr
set 3	with revised parser	82.8%/81.4%	478 words/hr

- The data: 20K-word Xinhua News, segmented and POS tagged.
- A stochastic TAG parser: trained on tested on CTB-1

# Uses

- Segmentation
  - International Chinese word segmentation bake-offs: (2003, 2005, 2006, 2008)
- POS tagging
  - Tseng et al 2005, Hillard et al 2006, Xia and Cheung 2006, ...
- BaseNP chunking
  - Liang et al 2006, Xu et al 2006, Chen et al 2006...
- Empty category recovery
  - Zhao and Ng 2007

# More on uses

- Constituent structure parsing
  - Chiang and Bikel 2002, Levy and Manning 2003, Luo 2003, Hearne and Way 2004, Bikel 2004, Xiong et al 2005, Bod 2006, ...
- Dependency structure parsing
  - Ma et al 2004, Jin et al 2005, Cheng et al 2006, Xu and Zhang 2006, Duan et al 2007, Wang 2007, Wang, Lin and Schuurmans 2007, Nivre 2007, ...

# More on uses

- Grammar extraction
  - Xia et al 2000; Burke et al 2004; Guo et al 2007
- Classifier Assignment
  - Guo and Zhong 2005
- Machine Translation
  - Wang, Collins and Koehn 2007,

# The formation of SIGHAN

- A special interest group of ACL, formed in 2000
- A direct result of the two Chinese NLP workshops and three meetings in 1998-2000.
- 6 SIGHAN workshops, 4 bakeoffs so far
- A community consisting of researchers from all over the world

# Chinese PropBank (CPB)

Version	CPB1.0	CPB2.0	*CPB3.0
Date	2005	2007	2008
Words	250K	500K	780K
Predicates	4,865	11,765	13,534

# Future expansion

- Discourse relations
  - Pilot study (Xue 2005)
  - Need to start with sense tagging of discourse connectives
- Temporal relations
  - Annotating tense in a tense-less language

# Outline

- Treebank overview
- Main issues
- Case study: The Chinese Penn Treebank
- Creating a Hindi/Urdu treebank

# Future work

- To build a Hindi/Urdu treebank
  - Collaboration between IIIT and four US sites (UW, Colorado, Columbia, and UMass)
  - 3-year project
  - (?) 350K Hindi, 150K Urdu: 25K are parallel treebanks
  - Part of the data will be PropBank-ed
  - Starting with DS annotation, converting DS to PS
- UW will focus on the DS => PS conversion

# Previous work on DS => PS conversion

- (Collins, Hajič, Ramshaw and Tillmann, 1999)
- (Xia and Palmer, 2001)
- Both are based on heuristics.
- Need to handle non-projectivity and ambiguity.

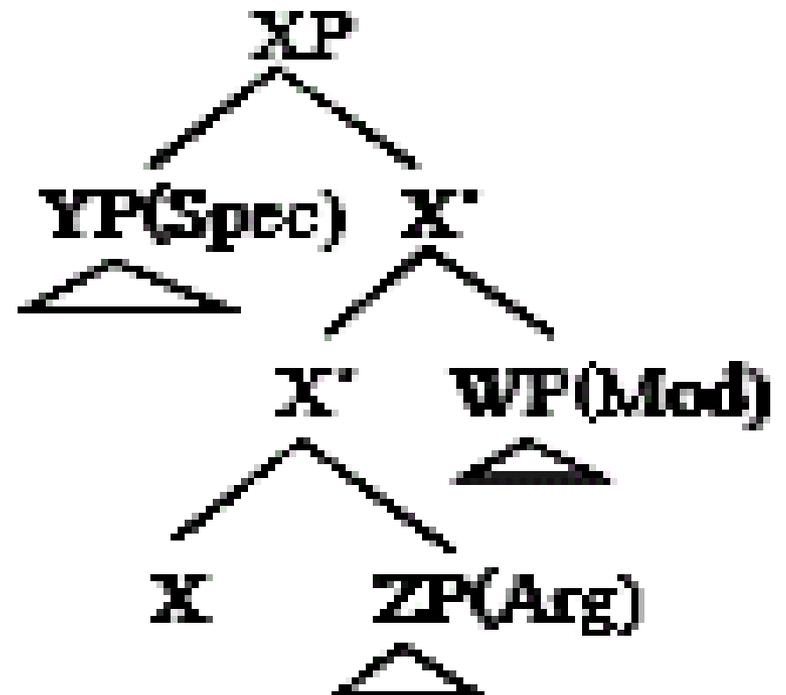
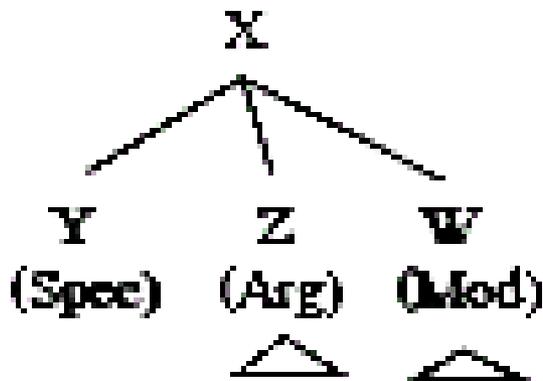
# Algorithm 1: Applying X-bar theory

X-bar theory:

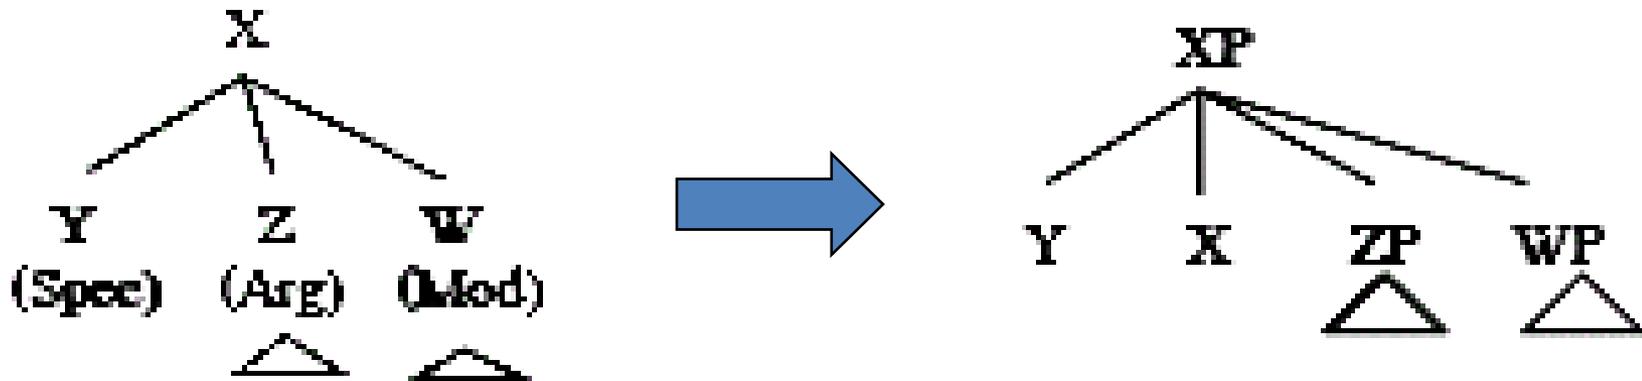
(1)  $XP \rightarrow YP X'$

(2)  $X' \rightarrow X' WP$

(3)  $X' \rightarrow X ZP$



# Algorithm 2: (Collins et al., 1999)



# Algorithm 3: (Xia and Palmer, 2001)

- The user provides an argument table and a projection table.
- Ex:
  - The argument table: the argument of a verb can be NP/S
  - The projection table: a noun can project to NP, and a verb can project to VP.

# Results

	Unlabeled recall	Unlabeled precision	Ratio of test/gold
Algorithm 1	81.3	32.8	2.48
Algorithm 2	54.2	91.5	0.59
Algorithm 3	86.2	88.7	0.98

Test data: Section 0 of the English Penn Treebank

# Remaining issues

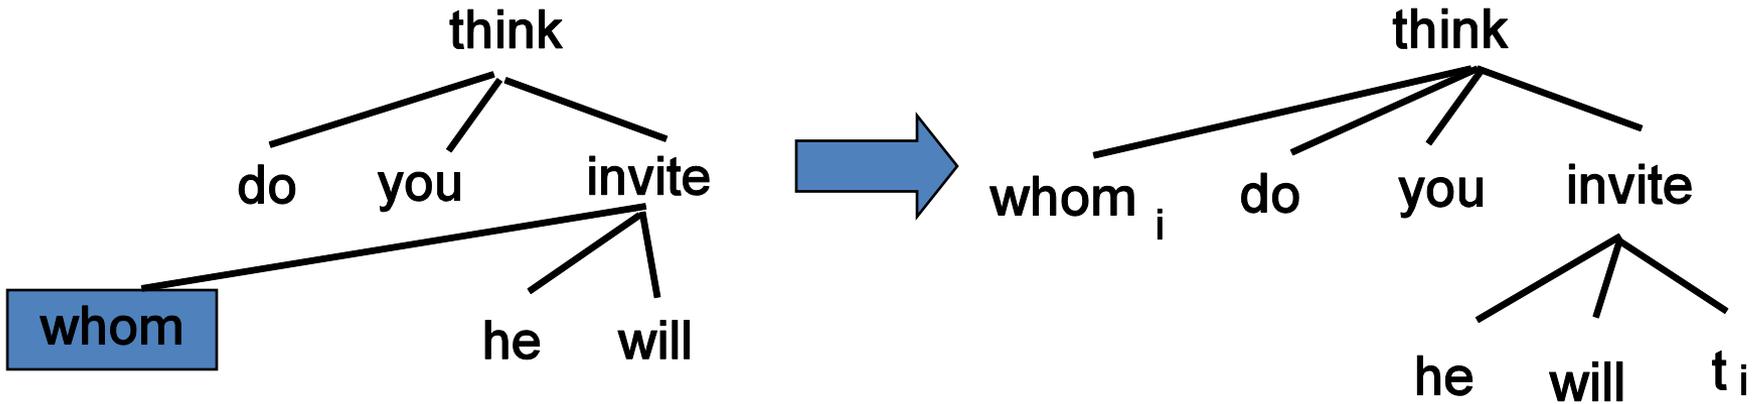
- Need to provide the argument table and the projection table.
- The experiment is artificial.
  - PS → DS → PS
- The f-measure is still low: about 87%
  - Information missing in DS
  - Inconsistency in the PS
  - Can we do much better if we can decide what the DS and the PS should look like?

# The new approach

- $DS \rightarrow DS^+$  : e.g., removing non-projectivity by introducing trace and co-indexation.
- $DS^+ \rightarrow PS^+$ : We prefer to keep this step simple and general.
- $PS^+ \rightarrow PS$ : e.g., choose one or more phrase structures stored in  $PS^+$ .

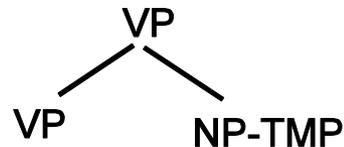
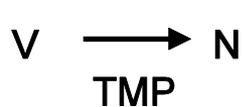
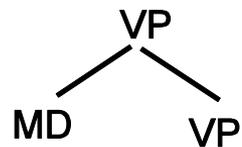
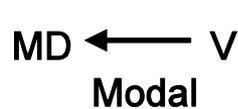
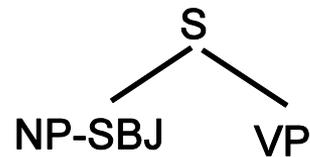
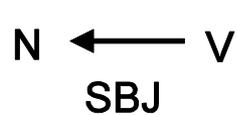
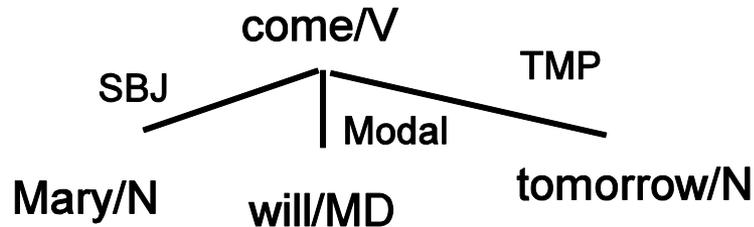
DS  $\rightarrow$  DS<sup>+</sup>

Whom do you think he will invite?

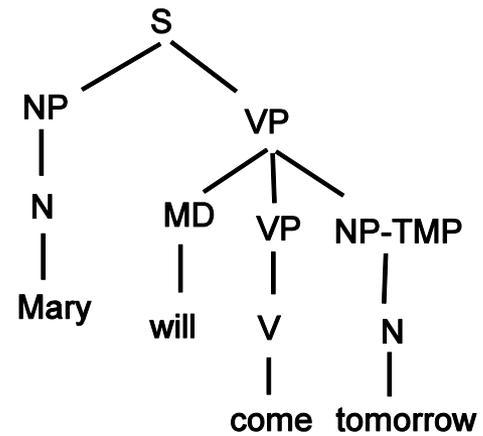
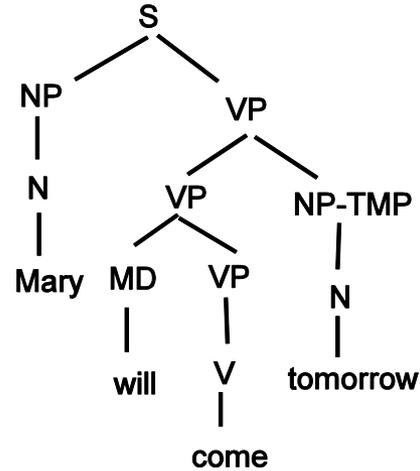
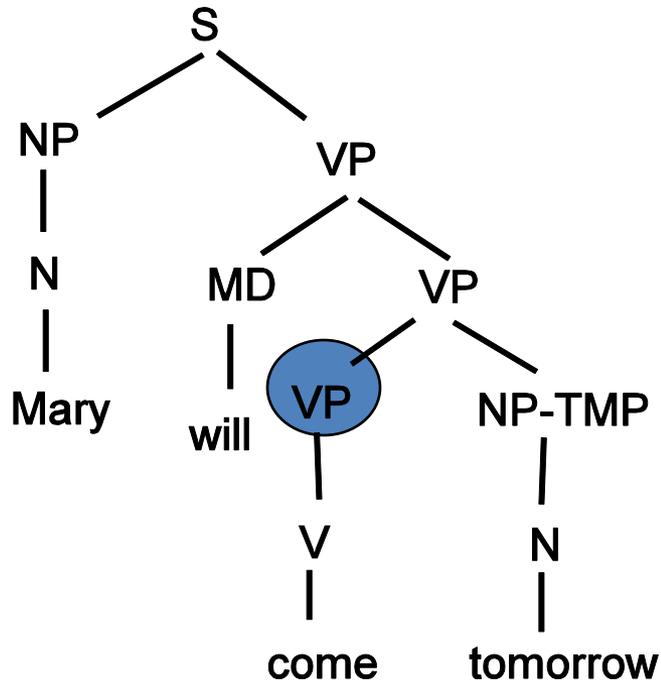


# DS<sup>+</sup> → PS<sup>+</sup>

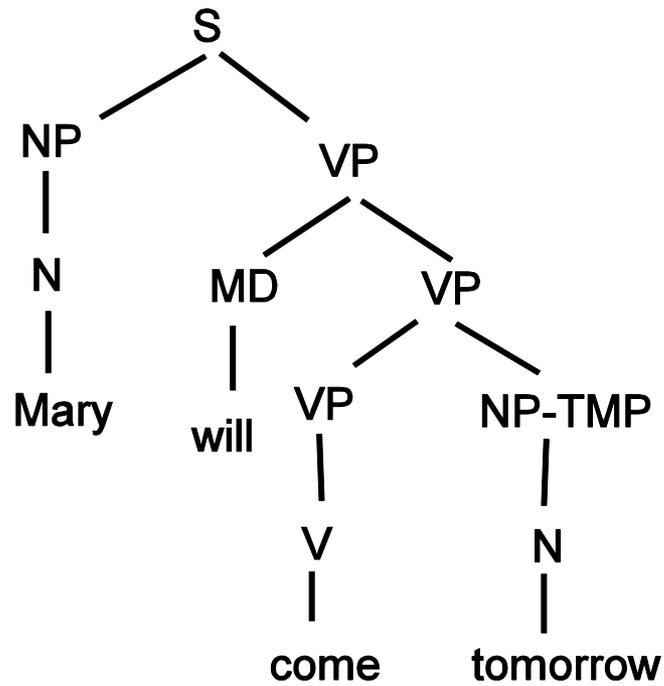
- Mary will come tomorrow



# Resulting phrase structures

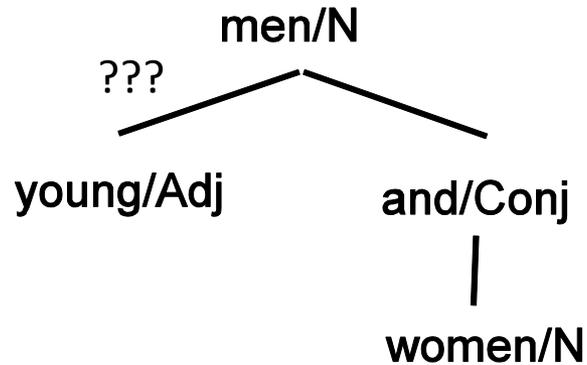


PS<sup>+</sup> → PS



# Ambiguity in DS

Ex1: young men and women



➔ Need to eliminate ambiguity in DS

# Preliminary results

- Learned the rules from Sect 19, tested on Section 22 of PTB.



DS pattern type	# of rules	Labeled precision	Labeled recall
(N, V)	1841	81.87	86.37
(N, V, SBJ)	2507	83.63	88.93
(N, V, SBJ, is_leaf)	2826	86.93	90.53
(N, V, SBJ, is_leaf, yes)	3513	88.11	90.67

# Observations

- The new approach outperforms Algorithm 3 by 3%.
- To reduce the number of rules, we can reduce the POS tagset.
  - Merging tags for nouns does not hurt the performance. The same is true for adjectives and adverbs
  - Merging tags for verbs hurts the f-measure by more than 2% (Ex: According/VBG to/IN ....)

# Error analysis

- Current result: 89.4
- No missing rules: 92.25
- Each link in the input DS is given the correct rule: 99.1

=> Is it possible to have one rule per DS pattern?

# The next step

- While designing DS and PS guidelines
  - the analyses in the two should be consistent
  - the conversion rules will be created at the same time (manually or extracted from annotated data)
- We will test
  - whether we can maintain one rule per DS pattern
  - whether DS pattern needs to look at more than one DS link
- Applying this algorithm to other languages (e.g., Czech)

# Conclusion

# Annotation procedure

- Selecting data
- Creating guidelines
- Training annotators
  
- Tokenization / Word segmentation
- POS tagging
- Bracketing
  
- Quality control
- Preliminary and final release
  
- ➔ Train preprocessing tools to speed up annotation.
- ➔ Revision is needed at various stages

# Lessons learned from treebanking

- Good annotation guidelines:
  - A treebank should be informative, and the annotation should be consistent.
  - There should be more interaction between grammar development and treebank development.
- Good, trained people:
  - Linguists for guideline design
  - Computational linguists for preprocessing and system support
  - Well-trained annotators
  - The large community for feedback

# Lessons learned (cont)

- Quality control
  - Routine double annotation
  - Tools for detecting annotation errors
  - Feedback from parsing, PropBank, etc.
- Use of NLP tools
  - Preprocessing speeds up annotation, but could potentially biases the treebank.
  - Other tools: search, conversion, etc.
- There should be more coordination between different layers of annotation (e.g., treebank and PropBank)

# The next step

- To build Hindi/Urdu treebank
  - We believe that DS  $\Rightarrow$  PS is possible if the DS is carefully designed.
- More information:  
<http://faculty.washington.edu/fxia/treebank.htm>

# Acknowledgments

- Martha Palmer, Univ. of Colorado
- Nianwen Xue, Univ. of Colorado
- Owen Rambow, Columbia University
- Dipti Misra Sharma and Rajeev Sangal, IIIT
  
- The whole CTB team and all the CTB users
- DOD, NSF, and DARPA
- Participants to our previous workshops