The Proposition Bank

An Annotated Corpus of Semantic Roles

TzuYi Kuo

EMLCT

Saarland University

June 14, 2010

Outline

- Introduction
- Motivation
- PropBank
  - Semantic role
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Outline

- Introduction
- Motivation
- PropBank
  - Semantic Role
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Introduction

- Represent the full meaning of sentences

- Alternation
  - Syntactic realization of semantic arguments

(1) John broke the window.

(2) The window broke.

Introduction

- Represent the full meaning of sentences

- Alternation
  - Syntactic realization of semantic arguments

(1) John broke the window.

(2) The window broke.

same underlying semantic role

Introduction

Proposition Bank

Predicate-argument information

Penn Treebank

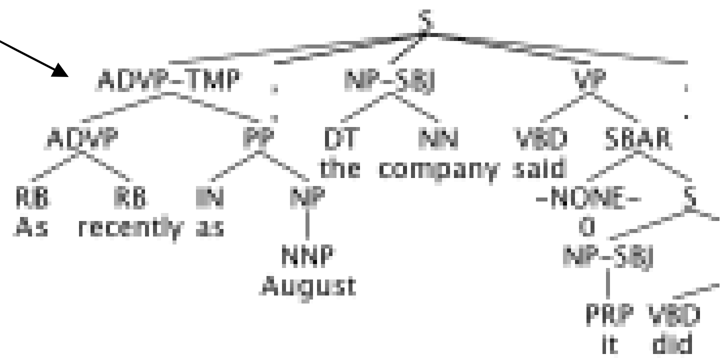

Introduction

- Focus on
  - Argument structure of verbs
  - Provide a complete corpus annotated with semantic roles

- Goal
  - Provide a broad-coverage hand-annotated corpus for supervised automatic role labelers
  - Show how and why these syntactic alternations take place

Outline

- Introduction
- Motivation
- PropBank
  - Semantic Role
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Motivation

- Inspired by Levin (1993)
  - Research into the linking between semantic roles and syntactic realization
  - Syntactic frames are a direct reflection of the underlying semantics
  - Define verb classes
    - Based on the ability of particular verbs
    - In syntactic frames

Motivation

- VerbNet (Kipper et al. ,2000)
  - Extend Levin’s classes
    - Adding an abstract representation of the syntactic frames for each class
    - Correspond between syntactic positions and the semantic roles they express

    Ex. Break

    Agent REL Patient

    Patient REL into pieces

Outline

- Introduction
- Motivation
- PropBank
  - Semantic Role
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

PropBank

- From sentences to propositions

John met Mary.

John and Mary met.

John met with Mary.

John and Mary had a meeting.

.
.
.

Proposition: meet(John, Mary)

Outline

- Introduction
- Motivation
- PropBank
  - Semantic Role
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Semantic Role

- Difficult to define a universal set of semantic roles covering all types of predicates

- Verb-by-verb basis
  - Arg0 $\rightarrow$ Agent
  - Arg1 $\rightarrow$ Prototypical Patient

Semantic Role

- Verb-specific numbered role

Frameset accept.01 "take willingly"

Arg0: Acceptor

Arg1: Thing accepted

Arg2: Accepted-from

Arg3: Attribute

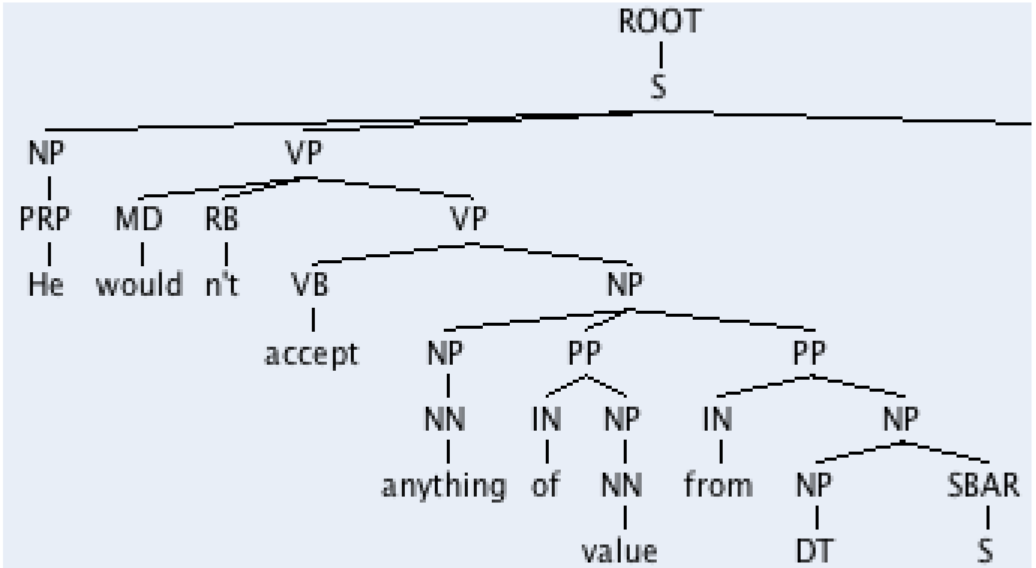

Semantic Role

- Verb-specific numbered role

Frameset accept.01 "take willingly"

Arg0: Acceptor

Arg1: Thing accepted

Arg2: Accepted-from

Arg3: Attribute

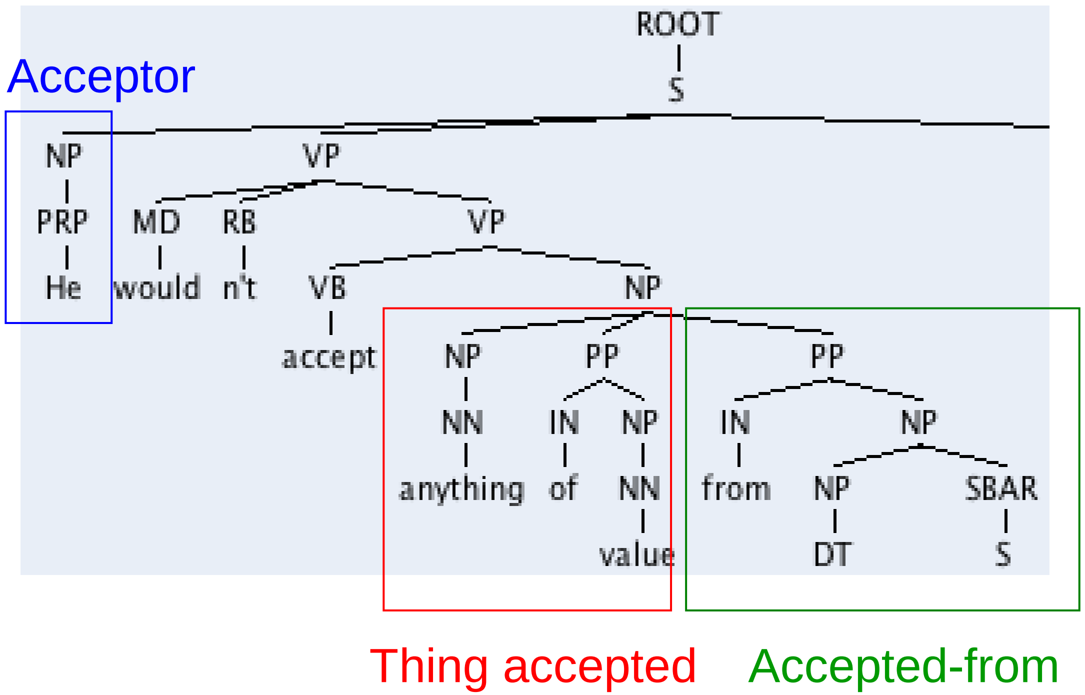

Semantic Role

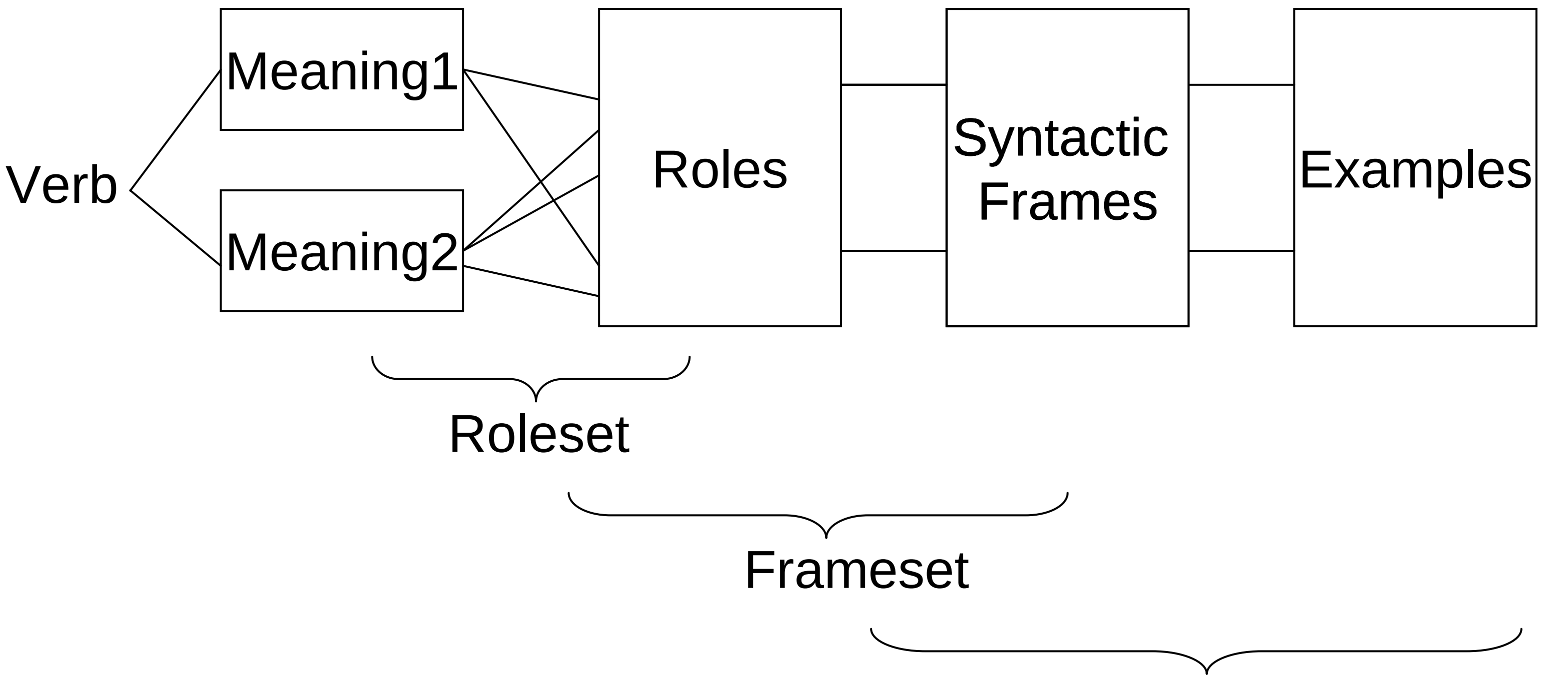

Frames File

Attempt to cover the range of syntactic alternations afforded by the usage

Outline

- Introduction
- PropBank
  - Semantic Role
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Framing

- Distinguishing Framesets
  - Different numbers of arguments

(15) Frameset **decline.01** "go down incrementally"

Arg1: entity going down

Arg2: amount gone down by, EXT

Arg3: start point

Arg4: end point

Ex: . . . [$_{\text{Arg1}}$ its net income] *declining* [$_{\text{Arg2-EXT}}$ 42%] [$_{\text{Arg4}}$ to \$121 million] [$_{\text{ArgM-TMP}}$ in the first 9 months of 1989]. (wsj_0067)

(16) Frameset **decline.02** "demure, reject"

Arg0: agent

Arg1: rejected thing

Ex: [$_{\text{Arg0}}$ A spokesman$_i$] *declined* [$_{\text{Arg1}}$ *trace*$_i$ to elaborate] (wsj_0038)

Framing

- Distinguishing Framesets
  - Verb-particle

Frameset cut.01 "slice"

Arg0: cutter

Arg1: thing cut

Arg2: medium, source

Arg3: instrument

Ex: [$_{Arg0}$ Longer production runs] [$_{ArgM-MOD}$ would] *cut* [$_{Arg1}$ inefficiencies from adjusting machinery between production cycles]. (wsj_0317)

Frameset cut.05 "cut back = reduce"

Arg0: cutter

Arg1: thing reduced

Arg2: amount reduced by

Arg3: start point

Arg4: end point

Ex: "Whoa," thought John, "[$_{Arg0}$ I$_i$]'ve got [$_{Arg0}$ *trace*$_i$] to start [$_{Arg0}$ *trace*$_i$] *cutting back* [$_{Arg1}$ my intake of chocolate].

Framing

- Distinguishing Framesets
  - Different syntactic type

Frameset **see.01** "view"

Arg0: viewer

Arg1: thing viewed

Ex1: [$_{\text{Arg0}}$ John] *saw* [$_{\text{Arg1}}$ the President] NP

Ex2: [$_{\text{Arg0}}$ John] *saw* [$_{\text{Arg1}}$ the President collapse] Clause object

Framing

- Secondary Predications

predicative reading	*ditransitive reading*
Mary called John a doctor.	Mary called John a doctor.[5]
(LABEL)	(SUMMON)
Arg0: Mary	Arg0: Mary
Rel: called	Rel: called
Arg1: John (item being labeled)	Arg2: John (benefactive)
Arg2-PRD: a doctor (attribute)	Arg1: a doctor (thing summoned)

Framing

- Traces
  - Empty category which known as trace
  - Coindex with other constituents in tree

[$_{Arg0}$ John$_i$] tried [$_{Arg0}$ *trace*$_i$] *to kick* [$_{Arg1}$ the football], but Mary pulled it away at the last moment.

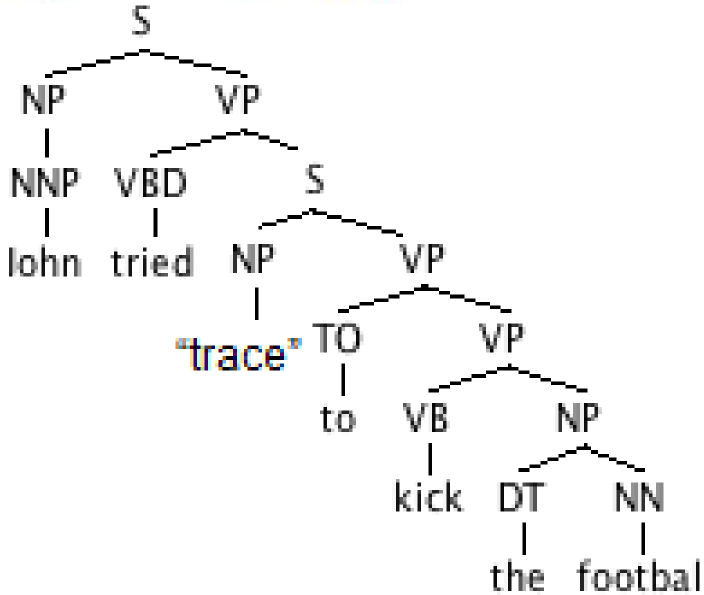

Framing

- Frames file
  - the collection of framesets for each lexeme

... [$_{Arg0}$ the company] to ... offer [$_{Arg1}$ a 15% to 20% stake] [$_{Arg2}$ to the public] (wsj_0345)[1]

... [$_{Arg0}$ Sotheby's] ... offered [$_{Arg2}$ the Dorrance heirs] [$_{Arg1}$ a money-back guarantee] (wsj_1928)

... [$_{Arg1}$ an amendment] offered [$_{Arg0}$ by Rep. Peter DeFazio] ... (wsj_0107)

... [$_{Arg2}$ Subcontractors] will be offered [$_{Arg1}$ a settlement] ... (wsj_0187)

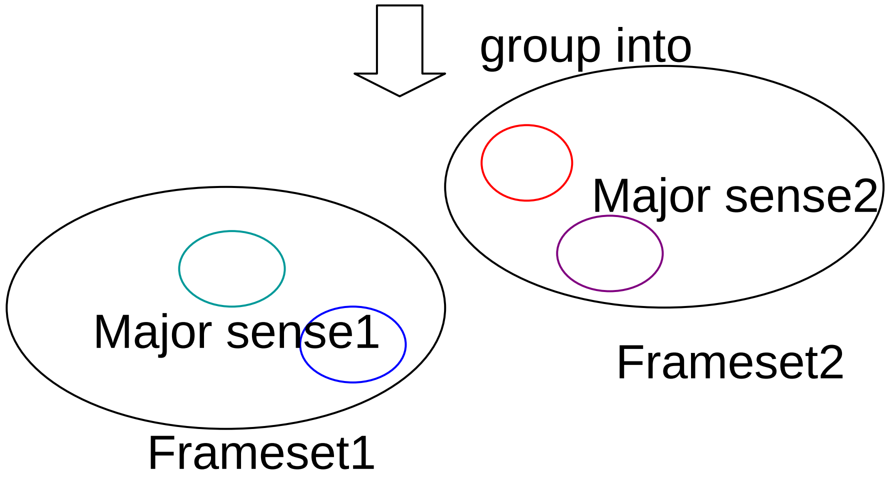

Framing

- In Wall Street Journal
  - Over 3,300 verbs framed
  - 4,500 framesets described
  - Average polysemy of 1.36
  - Each instance of a polysemous verb is marked as to which frameset it belongs to
  - Interannotator (ITA) agreement of 94%

Outline

- Introduction
- Motivation
- PropBank
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Development Process

- Annotation
  - Rule-based argument tagger (Palmer, Rosenzweig, and Cotton 2001)
    - Class-based mappings between grammatical and semantic roles
    - 83% accuracy
    - The output is then corrected by hand
      - Examining the descriptions of the arguments and the example tagged sentences

Development Process

- Annotation
  - Kappa statistic (Siegel and Castellan, 1988)
    - Measure agreement between annotators

$$\kappa = \frac{P(A) - P(E)}{1 - P(E)}$$

  - P(A) : the probability of inter-annotator agreement
  - P(E) : the agreement expected by chance

Development Process

- Annotation
  - Kappa statistic

Table 2
Interannotator agreement.

		$P(A)$	$P(E)$	κ
Including ArgM	Role identification	.99	.89	.93
	Role classification	.95	.27	.93
	Combined decision	.99	.88	.91
Excluding ArgM	Role identification	.99	.91	.94
	Role classification	.98	.41	.96
	Combined decision	.99	.91	.93

Outline

- Introduction
- PropBank
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Automatic Semantic Role Labeling

- Examine the importance of syntactic information for semantic-role labeling

- Comparing the performance of
  - System based on gold-standard parses
  - Automatically generated parser output

Automatic Semantic Role Labeling

- Gildea and Jurafsky (2002)
  - Statistical system trained on FrameNet project
    - Pass sentences through an automatic parser (Collins, 1999)
    - Extract syntactic features from the parses
    - Estimate probabilities for semantic roles from the syntactic and lexical features
    - Errors introduced by the parser no doubt negatively affected the results obtained

Outline

- Introduction
- PropBank
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Automatic Semantic Role Labeling

- Features
  - **Phrase type** : the syntactic type of the phrase expressing the semantic roles
  - **Parse tree path** : the path from the predicate through the parse tree to the constituent in question. In order to capture the syntactic relation of a constituent to the predicate

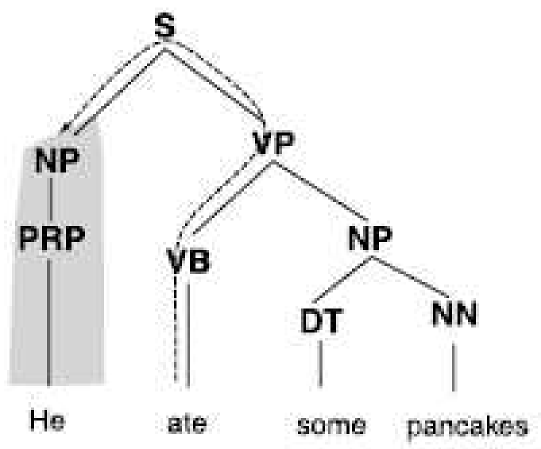

Automatic Semantic Role Labeling

- Features
  - **Position** : indicates whether the constituent to be labeled occurs before or after the predicate
  - **Voice** : distinguishes between active and passive, direct objects of active verbs correspond to subjects of passive verbs
  - **Headword** : a lexical feature and provides information about the semantic type of the role filler

Outline

- Introduction
- PropBank
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Automatic Semantic Role Labeling

- Predict argument roles

$$\mathrm{argmax}_{r_{1\ldots n}} P(r_{1\ldots n} | F_{1\ldots n}, p)$$

r_i : role of constituents i in the sentence

$F_i = \{pt_i, path_i, pos_i, v_i, h_i\}$: set of features at each constituent in the parse tree

Automatic Semantic Role Labeling

- Predict argument roles

$$P(r_{1\ldots n}|F_{1\ldots n},p) \approx P(\{r_{1\ldots n}\}|p)\prod_i \frac{P(r_i|F_i,p)}{P(r_i|p)}$$

$P(r_i \mid F_i,p)$: a constituent's role given our five features for the constituent and the predicate *p*

$P(\{r_{1\ldots n}\}|p)$: a set of roles appearing in a sentence given a predicate

Outline

- Introduction
- PropBank
  - Framing
  - Annotation
- Automatic Semantic-Role labeling
  - Features
  - Algorithm
  - Evaluation
- Conclusion

Automatic Semantic Role Labeling

- Data
  - PropBank (preliminary release version)
    - 72,109 predicate-argument structures
    - 190,815 individual arguments
    - examples from 2,462 lexical predicates (types)
  - Testing data : Penn Treebank Section 23

Automatic Semantic Role Labeling

- Results
  - Given the constituents which are arguments to the predicate and merely has to predict the correct role
  - Find the arguments in the sentence and label them correctly

	Accuracy		
	FrameNet	PropBank	PropBank > 10 examples
Automatic parses	82.0	79.9	80.9
Gold-standard parses		82.0	82.8

Accuracy of semantic-role prediction (in percentages) for known boundaries

Automatic Semantic Role Labeling

- Results
  - Adding Traces
    - Provide hints as to the semantics of individual clauses

	FrameNet		PropBank		PropBank > 10 examples	
	Precision	Recall	Precision	Recall	Precision	Recall
Automatic parses	64.6	61.2	68.6	57.8	69.9	61.1
Gold-standard parses			74.3	66.4	76.0	69.9
Gold-standard with traces			80.6	71.6	82.0	74.7

Accuracy of semantic-role prediction (in percentages) for unknown boundaries (the system must identify the correct constituents as arguments and give them the correct roles)

Automatic Semantic Role Labeling

- Results

Table 13
Accuracy of semantic-role prediction for unknown boundaries (the system must identify the correct constituents as arguments and give them the correct roles).

Role	Number	Precision	Labeled recall	Unlabeled recall
Arg0	1,197	94.2%	88.9%	92.2%
Arg1	1,436	95.4	82.5	88.9
Arg2	229	79.0	64.2	77.7
Arg3	61	71.4	49.2	54.1
Arg4	31	91.7	71.0	83.9
ArgM	127	59.6	26.8	52.0
ArgM-ADV	85	59.1	30.6	55.3
ArgM-DIR	49	76.7	46.9	61.2
ArgM-DIS	65	40.0	18.5	55.4
ArgM-EXT	18	81.2	72.2	77.8
ArgM-LOC	95	60.7	38.9	62.1
ArgM-MNR	80	62.7	40.0	63.8
ArgM-MOD	95	77.6	40.0	43.2
ArgM-NEG	40	63.6	17.5	40.0
ArgM-PRD	3	0.0	0.0	33.3
ArgM-PRP	54	70.0	25.9	37.0
ArgM-TMP	325	72.4	45.2	64.6

Labeled recall : how often the semantic-role label is correctly identified

Unlabeled recall : how often a constituent with the given role is correctly identified as being a semantic role, even if it is labeled with the wrong role

Automatic Semantic Role Labeling

- The relation of Syntactic Parsing and Semantic-Role labeling
  - Chunks
    - Do not build a full parse tree
    - Large advantage in speed
    - Contain basic-level constituent boundaries and labels
    - No dependencies between constituents

[$_{NP}$ Big investment banks] [$_{VP}$ refused to step] [$_{ADVP}$ up] [$_{PP}$ to]
[$_{NP}$ the plate] [$_{VP}$ to support] [$_{NP}$ the beleaguered floor traders] [$_{PP}$ by]
[$_{VP}$ buying] [$_{NP}$ bigblocks] [$_{PP}$ of] [$_{NP}$ stock], [$_{NP}$ traders] [$_{VP}$ say]. (wsj_2300)

Automatic Semantic Role Labeling

- The relation of Syntactic Parsing and Semantic-Role labeling

Table 14
Summary of results for unknown-boundary condition.

	Precision	Recall
Gold parse	74.3%	66.4%
Auto parse	68.6	57.8
Chunk	27.6	22.0
Chunk, relaxed scoring	49.5	35.1

Conclusion

- Consistent annotation has been achieved
- One step closer to a detailed semantic representation

- WSJ too domain specific, too financial, need broader coverage genres for more general annotation

Future work

- Add more informative thematic labels based on VerbNet
- Map annotation with FrameNet to merge two annotated data sets
- Explore
  - machine-learning approaches
  - Integration of semantic-role labeling and sense tagging with the parsing process

References

- Levin, B. (1993).

  English Verb Classes and Alternations: A preliminary Investigation.

  University of Chicago Press, Chicago.

- Kipper, K., Hoa T. D., and Martha, P. (2000).

  Class-based construction of a verb lexicon.

  Proceedings of the Seventh National Conference on Artificial Intelligence