

Statistical NLP

Spring 2010



Lecture 19: Semantic Roles

Dan Klein – UC Berkeley

Includes examples from Johnson, Jurafsky and Gildea, Luo, Palmer

Semantic Role Labeling (SRL)

- Characterize clauses as *relations with roles*:


[*Judge* She] **blames** [*Evaluate* the Government] [*Reason* for failing to do enough to help] .

Holman would characterise this as **blaming** [*Evaluate* the poor] .

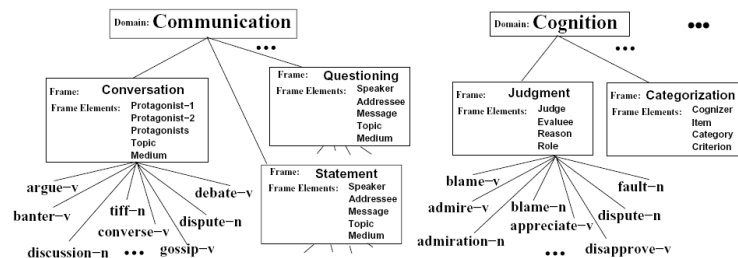
The letter quotes Black as saying that [*Judge* white and Navajo ranchers] misrepresent their livestock losses and **blame** [*Reason* everything] [*Evaluate* on coyotes] .

- Want to more than which NP is the subject (but not much more):
- Relations like *subject* are syntactic, relations like *agent* or *message* are semantic
- Typical pipeline:
 - Parse, then label roles
 - Almost all errors locked in by parser
 - Really, SRL is quite a lot easier than parsing

SRL Example



PropBank / FrameNet



- FrameNet: roles shared between verbs
- PropBank: each verb has its own roles
- PropBank more used, because it's layered over the treebank (and so has greater coverage, plus parses)
- Note: some linguistic theories postulate even fewer roles than FrameNet (e.g. 5-20 total: agent, patient, instrument, etc.)

PropBank Example

fall.01 sense: move downward
roles: Arg1: thing falling
 Arg2: extent, distance fallen
 Arg3: start point
 Arg4: end point

Sales fell to \$251.2 million from \$278.7 million.
arg1: Sales
rel: fell
arg4: to \$251.2 million
arg3: from \$278.7 million

PropBank Example

rotate.02 sense: shift from one thing to another
roles: Arg0: causer of shift
 Arg1: thing being changed
 Arg2: old thing
 Arg3: new thing

Many of Wednesday's winners were losers yesterday as investors quickly took profits and rotated their buying to other issues, traders said. (wsj_1723)
arg0: investors
rel: rotated
arg1: their buying
arg3: to other issues

PropBank Example

aim.01 sense: intend, plan
 roles: Arg0: aimer, planner
 Arg1: plan, intent


The Central Council of Church Bell Ringers aims *trace* to improve relations with vicars. (wsj_0089)
 arg0: The Central Council of Church Bell Ringers
 rel: aims
 arg1: *trace* to improve relations with vicars

aim.02 sense: point (weapon) at
 roles: Arg0: aimer
 Arg1: weapon, etc.
 Arg2: target


Banks have been aiming packages at the elderly.
 arg0: Banks
 rel: aiming
 arg1: packages
 arg2: at the elderly

Shared Arguments

(NP-SBJ (JJ massive) (JJ internal) (NN debt))
 (VP (VBZ has)
 (VP (VBN forced)
 (S
 (NP-SBJ-1 (DT the) (NN government))
 (VP
 (VP (TO to)
 (VP (VB borrow)
 (ADVP-MNR (RB massively))...



Path Features



<i>Path</i>	<i>Description</i>
VB↓VP↓PP	PP argument/adjunct
VB↑VP↑S↓NP	subject
VB↑VP↓NP	object
VB↑VP↑VP↑S↓NP	subject (embedded VP)
VB↑VP↓ADVP	adverbial adjunct
NN↑NP↑NP↓PP	prepositional complement of noun

Results

- **Features:**
 - Path from target to filler
 - Filler's syntactic type, headword, case
 - Target's identity
 - Sentence voice, etc.
 - Lots of other second-order features


- **Gold vs parsed source trees**

- SRL is fairly easy on gold trees
- Harder on automatic parses

CORE		ARGM	
F1	Acc.	F1	Acc.
92.2	80.7	89.9	71.8

CORE		ARGM	
F1	Acc.	F1	Acc.
84.1	66.5	81.4	55.6

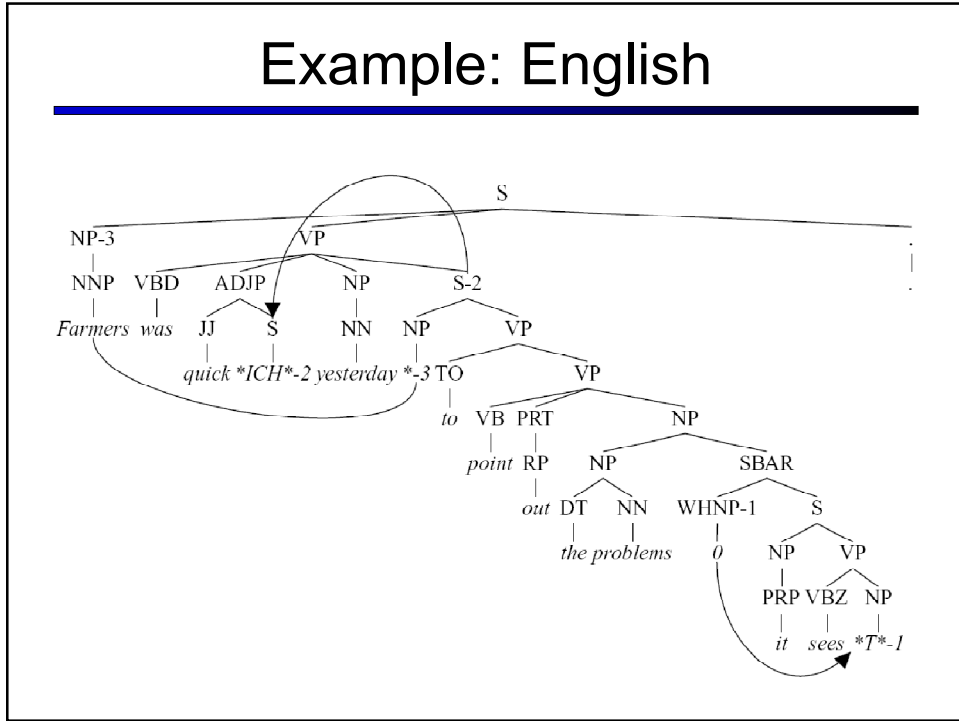
Interaction with Empty Elements



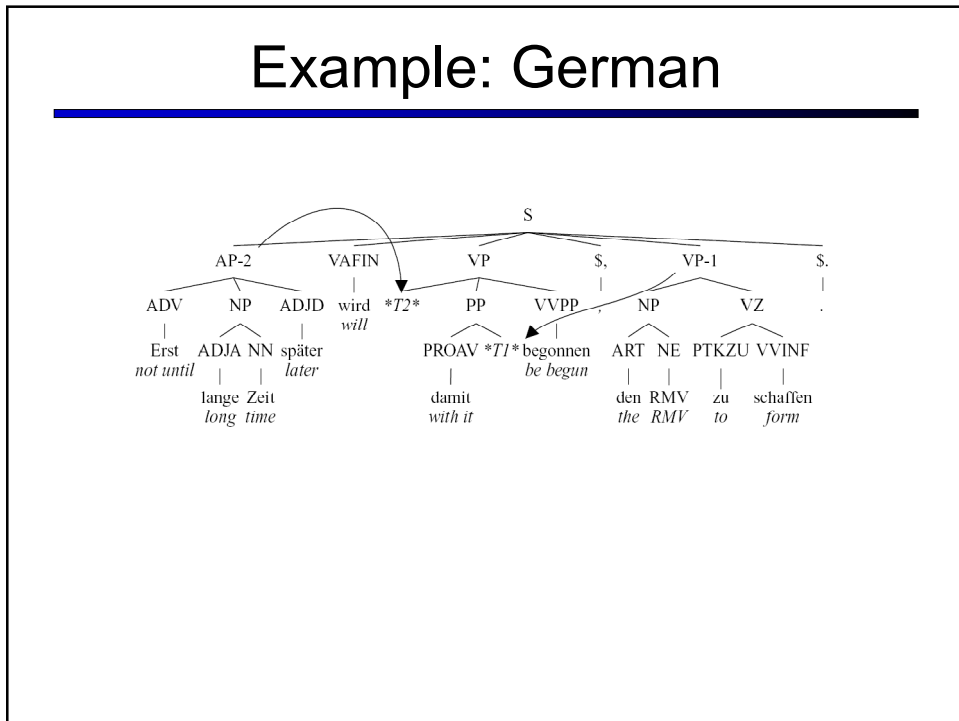
Empty Elements

- In the PTB, three kinds of empty elements:
 - Null items (usually complementizers)
 - Dislocation (WH-traces, topicalization, relative clause and heavy NP extraposition)
 - Control (raising, passives, control, shared argumentation)
- Need to reconstruct these (and resolve any indexation)

Example: English





Example: German




Types of Empties

Antecedent	POS	Label	Count	Description
NP	NP	*	18,334	NP trace (e.g., <i>Sam was seen</i> *)
	NP	*	9,812	NP PRO (e.g., <i>* to sleep is nice</i>)
WHNP	NP	*T*	8,620	WH trace (e.g., <i>the woman who you saw</i> *T*)
		U	7,478	Empty units (e.g., <i>\$ 25</i> *U*)
		0	5,635	Empty complementizers (e.g., <i>Sam said 0 Sasha snores</i>)
S	S	*T*	4,063	Moved clauses (e.g., <i>Sam had to go, Sasha explained</i> *T*)
WHADVP	ADVP	*T*	2,492	WH-trace (e.g., <i>Sam explained how to leave</i> *T*)
	SBAR		2,033	Empty clauses (e.g., <i>Sam had to go, Sasha explained</i> (SBAR))
	WHNP	0	1,759	Empty relative pronouns (e.g., <i>the woman 0 we saw</i>)
	WHADVP	0	575	Empty relative pronouns (e.g., <i>no reason 0 to leave</i>)



A Pattern-Matching Approach

- [Johnson 02]



Pattern-Matching Details

- Something like transformation-based learning
- Extract patterns
 - Details: transitive verb marking, auxiliaries
 - Details: legal subtrees
- Rank patterns
 - Pruning ranking: by correct / match rate
 - Application priority: by depth
- Pre-order traversal
- Greedy match

Top Patterns Extracted

Count	Match	Pattern
5816	6223	(S (NP (-NONE- *)) VP)
5605	7895	(SBAR (-NONE- 0) S)
5312	5338	(SBAR WHNP-1 (S (NP (-NONE- *T*-1)) VP))
4434	5217	(NP QP (-NONE- *U*))
1682	1682	(NP \$ CD (-NONE- *U*))
1327	1593	(VP VBN_t (NP (-NONE- *)) PP)
700	700	(ADJP QP (-NONE- *U*))
662	1219	(SBAR (WHNP-1 (-NONE- 0)) (S (NP (-NONE- *T*-1)) VP))
618	635	(S S-1 , NP (VP VBD (SBAR (-NONE- 0) (S (-NONE- *T*-1)))) .)
499	512	(SINV `` S-1 , '' (VP VBZ (S (-NONE- *T*-1))) NP .)
361	369	(SINV `` S-1 , '' (VP VBD (S (-NONE- *T*-1))) NP .)
352	320	(S NP-1 (VP VBZ (S (NP (-NONE- *-1)) VP)))
346	273	(S NP-1 (VP AUX (VP VBN_t (NP (-NONE- *-1)) PP)))
322	467	(VP VBD_t (NP (-NONE- *)) PP)
269	275	(S `` S-1 , '' NP (VP VBD (S (-NONE- *T*-1))) .)

Results

Empty node		Section 23			Parser output		
POS	Label	<i>P</i>	<i>R</i>	<i>f</i>	<i>P</i>	<i>R</i>	<i>f</i>
(Overall)		0.93	0.83	0.88	0.85	0.74	0.79
NP	*	0.95	0.87	0.91	0.86	0.79	0.82
NP	*T*	0.93	0.88	0.91	0.85	0.77	0.81
	0	0.94	0.99	0.96	0.86	0.89	0.88
	U	0.92	0.98	0.95	0.87	0.96	0.92
S	*T*	0.98	0.83	0.90	0.97	0.81	0.88
ADVP	*T*	0.91	0.52	0.66	0.84	0.42	0.56
SBAR		0.90	0.63	0.74	0.88	0.58	0.70
WHNP	0	0.75	0.79	0.77	0.48	0.46	0.47

A Machine-Learning Approach

- [Levy and Manning 04]
- Build two classifiers:
 - First one predicts where empties go
 - Second one predicts if/where they are bound
 - Use syntactic features similar to SRL (paths, categories, heads, etc)

	Performance on gold trees							Performance on parsed trees					
	ID			Rel		Combo		ID			Combo		
	P	R	F1	Acc	P	R	F1	P	R	F1	P	R	F1
WSJ(full)	92.0	82.9	87.2	95.0	89.6	80.1	84.6	34.5	47.6	40.0	17.8	24.3	20.5
WSJ(sm)	92.3	79.5	85.5	93.3	90.4	77.2	83.2	38.0	47.3	42.1	19.7	24.3	21.7
NEGRA	73.9	64.6	69.0	85.1	63.3	55.4	59.1	48.3	39.7	43.6	20.9	17.2	18.9

Semantic Interpretation

- **Back to meaning!**
 - A very basic approach to computational semantics
 - Truth-theoretic notion of semantics (Tarskian)
 - Assign a “meaning” to each word
 - Word meanings combine according to the parse structure
 - People can and do spend entire courses on this topic
 - We’ll spend about an hour!
- **What’s NLP and what isn’t?**
 - Designing meaning representations?
 - Computing those representations?
 - Reasoning with them?
- **Supplemental reading will be on the web page.**

Meaning


- **“Meaning”**
 - What is meaning?
 - “The computer in the corner.”
 - “Bob likes Alice.”
 - “I think I am a gummi bear.”
 - Knowing whether a statement is true?
 - Knowing the conditions under which it’s true?
 - Being able to react appropriately to it?
 - “Who does Bob like?”
 - “Close the door.”
- **A distinction:**
 - Linguistic (semantic) meaning
 - “The door is open.”
 - Speaker (pragmatic) meaning
- **Today: assembling the semantic meaning of sentence from its parts**

Entailment and Presupposition

- Some notions worth knowing:
 - Entailment:
 - A entails B if A being true necessarily implies B is true
 - ? “Twitchy is a big mouse” → “Twitchy is a mouse”
 - ? “Twitchy is a big mouse” → “Twitchy is big”
 - ? “Twitchy is a big mouse” → “Twitchy is furry”
 - Presupposition:
 - A presupposes B if A is only well-defined if B is true
 - “The computer in the corner is broken” presupposes that there is a (salient) computer in the corner


Truth-Conditional Semantics

- Linguistic expressions:
 - “Bob sings”
- Logical translations:
 - sings(bob)
 - Could be $p_{1218}(e_{397})$
- Denotation:
 - [[bob]] = some specific person (in some context)
 - [[sings(bob)]] = ???
- Types on translations:
 - bob : e (for entity)
 - sings(bob) : t (for truth-value)




Truth-Conditional Semantics

- Proper names:
 - Refer directly to some entity in the world
 - Bob : bob $[[\text{bob}]]^W \rightarrow ???$
- Sentences:
 - Are either true or false (given how the world actually is)
 - Bob sings : sings(bob)
- So what about verbs (and verb phrases)?
 - sings must combine with bob to produce sings(bob)
 - The λ -calculus is a notation for functions whose arguments are not yet filled.
 - sings : $\lambda x.\text{sings}(x)$
 - This is *predicate* – a function which takes an entity (type e) and produces a truth value (type t). We can write its type as $e \rightarrow t$.
 - Adjectives?



Compositional Semantics

- So now we have meanings for the words
- How do we know how to combine words?
- Associate a combination rule with each grammar rule:
 - $S : \beta(\alpha) \rightarrow NP : \alpha \quad VP : \beta$ (function application)
 - $VP : \lambda x . \alpha(x) \wedge \beta(x) \rightarrow VP : \alpha \quad \text{and} : \emptyset \quad VP : \beta$ (intersection)
- Example:




Denotation

- What do we do with logical translations?
 - Translation language (logical form) has fewer ambiguities
 - Can check truth value against a database
 - Denotation (“evaluation”) calculated using the database
 - More usefully: assert truth and modify a database
 - Questions: check whether a statement in a corpus entails the (question, answer) pair:
 - “Bob sings and dances” → “Who sings?” + “Bob”
 - Chain together facts and use them for comprehension

Other Cases

- Transitive verbs:
 - likes : $\lambda x.\lambda y.likes(y,x)$
 - Two-place predicates of type $e \rightarrow (e \rightarrow t)$.
 - likes Amy : $\lambda y.likes(y,Amy)$ is just like a one-place predicate.
- Quantifiers:
 - What does “Everyone” mean here?
 - Everyone : $\lambda f.\forall x.f(x)$
 - Mostly works, but some problems
 - Have to change our NP/VP rule.
 - Won't work for “Amy likes everyone.”
 - “Everyone likes someone.”
 - This gets tricky quickly!




Indefinites

- First try

- “Bob ate a waffle” : $\text{ate}(\text{bob}, \text{waffle})$
- “Amy ate a waffle” : $\text{ate}(\text{amy}, \text{waffle})$

- Can't be right!

- $\exists x : \text{waffle}(x) \wedge \text{ate}(\text{bob}, x)$
- What does the translation of “a” have to be?
- What about “the”?
- What about “every”?



Grounding

- Grounding

- So why does the translation $\text{likes} : \lambda x. \lambda y. \text{likes}(y, x)$ have anything to do with actual liking?
- It doesn't (unless the denotation model says so)
- Sometimes that's enough: wire up **bought** to the appropriate entry in a database

- Meaning postulates

- Insist, e.g. $\forall x, y. \text{likes}(y, x) \rightarrow \text{knows}(y, x)$
- This gets into lexical semantics issues


- Statistical version?

Tense and Events

- In general, you don't get far with verbs as predicates
- Better to have event variables e
 - "Alice danced" : $\text{danced}(\text{alice})$
 - $\exists e : \text{dance}(e) \wedge \text{agent}(e, \text{alice}) \wedge (\text{time}(e) < \text{now})$
- Event variables let you talk about non-trivial tense / aspect structures
 - "Alice had been dancing when Bob sneezed"
 - $\exists e, e' : \text{dance}(e) \wedge \text{agent}(e, \text{alice}) \wedge$
 $\text{sneeze}(e') \wedge \text{agent}(e', \text{bob}) \wedge$
 $(\text{start}(e) < \text{start}(e') \wedge \text{end}(e) = \text{end}(e')) \wedge$
 $(\text{time}(e') < \text{now})$

Adverbs

- What about adverbs?
 - "Bob sings terribly"
 - $\text{terribly}(\text{sings}(\text{bob}))?$
 - $(\text{terribly}(\text{sings}))(\text{bob})?$
 - $\exists e \text{ present}(e) \wedge$
 $\text{type}(e, \text{singing}) \wedge$
 $\text{agent}(e, \text{bob}) \wedge$
 $\text{manner}(e, \text{terrible}) ?$
 - It's really not this simple..



Propositional Attitudes

- “Bob thinks that I am a gummi bear”
 - `thinks(bob, gummi(me))` ?
 - `thinks(bob, “I am a gummi bear”)` ?
 - `thinks(bob, ^gummi(me))` ?
- Usual solution involves intensions ($\wedge X$) which are, roughly, the set of possible worlds (or conditions) in which X is true
- Hard to deal with computationally
 - Modeling other agents models, etc
 - Can come up in simple dialog scenarios, e.g., if you want to talk about what your bill claims you bought vs. what you actually bought

Trickier Stuff


- Non-Intersective Adjectives
 - `green ball` : $\lambda x.[\text{green}(x) \wedge \text{ball}(x)]$
 - `fake diamond` : $\lambda x.[\text{fake}(x) \wedge \text{diamond}(x)]$? $\longrightarrow \lambda x.[\text{fake}(\text{diamond}(x))]$
- Generalized Quantifiers
 - `the` : $\lambda f.[\text{unique-member}(f)]$
 - `all` : $\lambda f. \lambda g [\forall x.f(x) \rightarrow g(x)]$
 - `most`?
 - Could do with more general second order predicates, too (why worse?)
 - `the(cat, meows)`, `all(cat, meows)`
- Generics
 - “Cats like naps”
 - “The players scored a goal”
- Pronouns (and bound anaphora)
 - “If you have a dime, put it in the meter.”
- ... the list goes on and on!

Multiple Quantifiers

- **Quantifier scope**
 - Groucho Marx celebrates quantifier order ambiguity:
“In this country a woman gives birth every 15 min.
Our job is to find that woman and stop her.”
- **Deciding between readings**
 - “Bob bought a pumpkin every Halloween”
 - “Bob put a warning in every window”
 - Multiple ways to work this out
 - Make it syntactic (movement)
 - Make it lexical (type-shifting)

Implementation, TAG, Idioms

- Add a “sem” feature to each context-free rule
 - $S \rightarrow NP \text{ loves } NP$
 - $S[\text{sem}=\text{loves}(x,y)] \rightarrow NP[\text{sem}=x] \text{ loves } NP[\text{sem}=y]$
 - Meaning of S depends on meaning of NPs
- TAG version:



- Template filling: $S[\text{sem}=\text{showflights}(x,y)] \rightarrow$
I want a flight from $NP[\text{sem}=x]$ to $NP[\text{sem}=y]$

Modeling Uncertainty

- Gaping hole warning!
- Big difference between statistical disambiguation and statistical reasoning.

The scout saw the enemy soldiers with night goggles.

- With probabilistic parsers, can say things like “72% belief that the PP attaches to the NP.”
 - That means that *probably* the enemy has night vision goggles.
 - However, you can’t throw a logical assertion into a theorem prover with 72% confidence.
 - Not clear humans really extract and process logical statements symbolically anyway.
 - Use this to decide the expected utility of calling reinforcements?
- In short, we need probabilistic reasoning, not just probabilistic disambiguation followed by symbolic reasoning!

CCG Parsing

- **Combinatory
Categorial
Grammar**

- Fully (mono-) lexicalized grammar
- Categories encode argument sequences
- Very closely related to the lambda calculus
- Can have spurious ambiguities (why?)

$John \vdash NP : john'$

$shares \vdash NP : shares'$

$buys \vdash (S \backslash NP) / NP : \lambda x. \lambda y. buys' xy$

$sleeps \vdash S \backslash NP : \lambda x. sleeps' x$

$well \vdash (S \backslash NP) \backslash (S \backslash NP) : \lambda f. \lambda x. well' (fx)$

