

# Semantic Role Labeling. Generalizing Lexical Features using Selectional Preferences

**Lluís Màrquez**

TALP Research Center  
Technical University of Catalonia

ILCC/HCRC seminar series at the University of  
Edinburgh's School of Informatics  
Edinburgh, June 3, 2011

- 1 Semantic Role Labeling
- 2 The Statistical Approach to SRL
- 3 Semantic Features for SRL
- 4 Conclusions

# Talk Overview

- 1 Semantic Role Labeling
- 2 The Statistical Approach to SRL
- 3 Semantic Features for SRL
- 4 Conclusions

# The Problem

## Semantic Role Labeling

SRL  $\stackrel{\text{def}}{=}$  identify the *arguments* of a given proposition and assign them *semantic labels* describing the *roles* they play in the predicate (i.e., recognize predicate argument structures)

# The Problem

IE point of view

SRL <sup>def</sup> = detecting basic event structures such as *who* did *what* to *whom*, *when* and *where*

[The luxury auto maker]<sub>AGENT</sub> [last year]<sub>TEMP</sub> sold<sub>P</sub> [1,214 cars]<sub>OBJECT</sub>  
[in the U.S.]<sub>LOCATIVE</sub>

# The Problem

IE point of view

SRL <sup>def</sup> = detecting basic event structures such as *who* did *what* to *whom*, *when* and *where*

[The luxury auto maker]<sub>AGENT</sub> [last year]<sub>TEMP</sub> sold<sub>P</sub> [1,214 cars]<sub>OBJECT</sub>  
[in the U.S.]<sub>LOCATIVE</sub>

# The Problem

## Syntactic variations

TEMP     HITTER     THING HIT     INSTRUMENT  
Yesterday, Kristina hit Scott with a baseball

- Scott was hit by Kristina yesterday with a baseball
- Yesterday, Scott was hit with a baseball by Kristina
- Yesterday Scott was hit by Kristina with a baseball
- Kristina hit Scott with a baseball yesterday

⇒ All of them share the same semantic representation:

hit(Kristina, Scott, yesterday, with a baseball)

Example from (Yih & Toutanova, 2006)

# The Problem

## Syntactic variations

TEMP     HITTER     THING HIT     INSTRUMENT  
Yesterday, Kristina hit Scott with a baseball

- Scott was hit by Kristina yesterday with a baseball
- Yesterday, Scott was hit with a baseball by Kristina
- Yesterday Scott was hit by Kristina with a baseball
- Kristina hit Scott with a baseball yesterday

⇒ All of them share the same semantic representation:

**hit**(Kristina, Scott, yesterday, with a baseball)

Example from (Yih & Toutanova, 2006)

# The Problem

## Structural view

Mapping from input to output structures:

- **Input** is *text* (enriched with morpho-syntactic information)
- **Output** is a *sequence of labeled arguments*
- **Sequential** segmenting/labeling problem

“ Mr. Smith *sent* the report to me this morning . ”

[Mr. Smith]<sub>AGENT</sub> *sent* [the report]<sub>OBJ</sub> [to me]<sub>RECIP</sub> [this morning]<sub>TMP</sub> .

Mr.<sub>B-AGENT</sub> Smith<sub>I</sub> *sent* the<sub>B-OBJ</sub> report<sub>I</sub> to<sub>B-RECIP</sub> me<sub>I</sub> this<sub>B-TMP</sub>  
morning<sub>I</sub> .O

# The Problem

## Structural view

Mapping from input to output structures:

- **Input** is *text* (enriched with morpho-syntactic information)
- **Output** is a *sequence of labeled arguments*
- **Sequential** segmenting/labeling problem

“ Mr. Smith *sent* the report to me this morning . ”

[Mr. Smith]<sub>AGENT</sub> *sent* [the report]<sub>OBJ</sub> [to me]<sub>RECIP</sub> [this morning]<sub>TMP</sub> .

Mr.<sub>B-AGENT</sub> Smith<sub>I</sub> *sent* the<sub>B-OBJ</sub> report<sub>I</sub> to<sub>B-RECIP</sub> me<sub>I</sub> this<sub>B-TMP</sub>  
morning<sub>I</sub> .O

# The Problem

## Structural view

Mapping from input to output structures:

- **Input** is *text* (enriched with morpho-syntactic information)
- **Output** is a *sequence of labeled arguments*
- **Sequential** segmenting/labeling problem

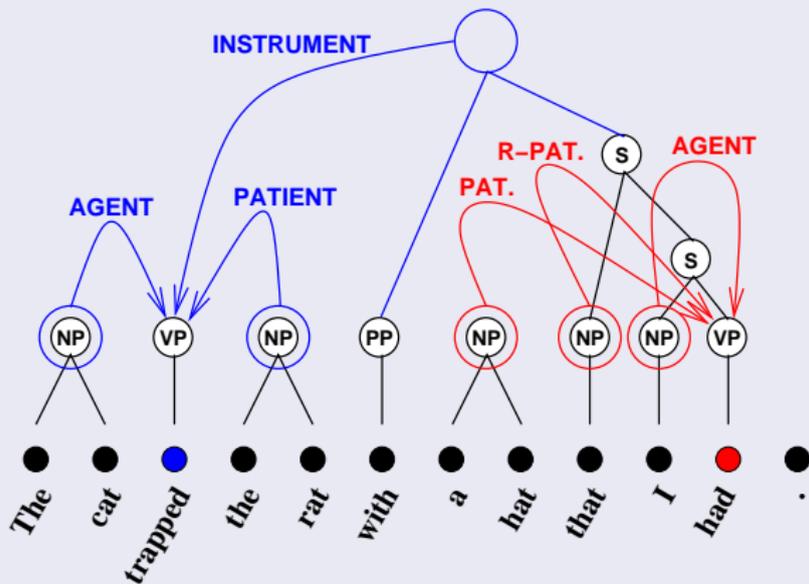
“ Mr. Smith *sent* the report to me this morning . ”

[Mr. Smith]<sub>AGENT</sub> *sent* [the report]<sub>OBJ</sub> [to me]<sub>RECIP</sub> [this morning]<sub>TMP</sub> .

Mr.<sub>B-AGENT</sub> Smith<sub>I</sub> *sent* the<sub>B-OBJ</sub> report<sub>I</sub> to<sub>B-RECIP</sub> me<sub>I</sub> this<sub>B-TMP</sub>  
morning<sub>I</sub> .O

# The Problem

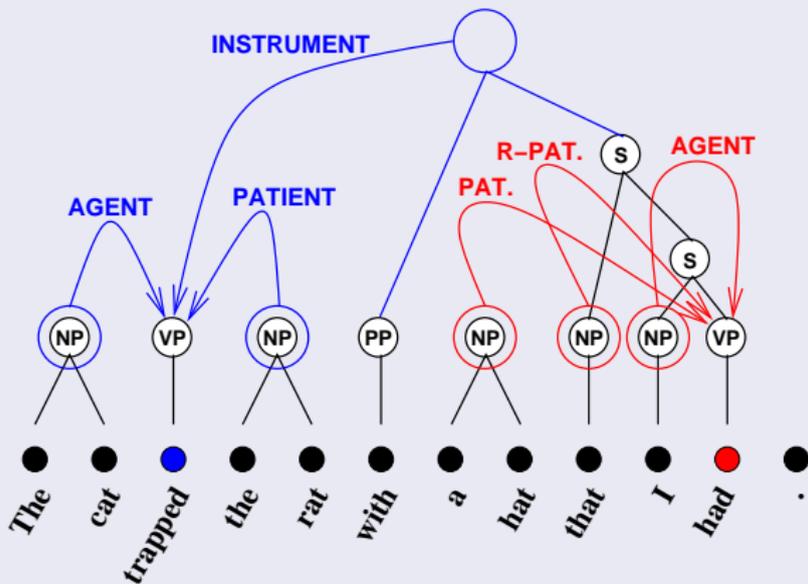
## Structural View



Output is a hierarchy of labeled arguments

# The Problem

## Structural View

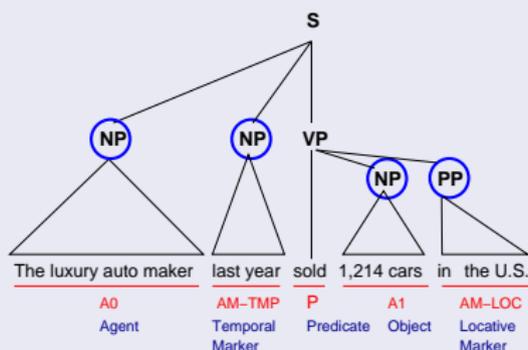


**Output** is a *hierarchy of labeled arguments*

# The Problem

## Linguistic nature of the problem

- Argument identification is strongly related to syntax

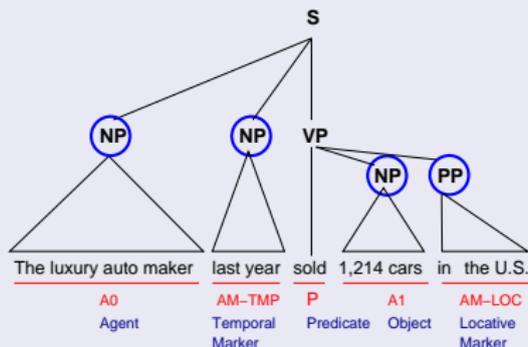


- Role labeling is a semantic task  
(e.g., selectional preferences could play an important role)

# The Problem

## Linguistic nature of the problem

- Argument identification is strongly related to syntax



- Role labeling is a semantic task  
(e.g., selectional preferences could play an important role)

# SRL Systems Available

- **ASSERT** (Automatic Statistical SEmantic Role Tagger)  
<http://cemantix.org/assert.html>
- **UIUC** system demo  
<http://l2r.cs.uiuc.edu/~cogcomp/srl-demo.php>
- **SwiRL**: state-of-the-art system from CoNLL-2005  
<http://www.surdeanu.name/mihai>
- **Shalmaneser**: FrameNet-based system from SALSA project  
<http://www.coli.uni-saarland.de/projects/salsa/shal/>
- **SEMAFOR**: Probabilistic Frame(Net)-Semantic Parser  
<http://www.ark.cs.cmu.edu/SEMAFOR/>
- **Brutus**: A CCG-based Semantic Role Labeler  
<http://www.ling.ohio-state.edu/~boxwell/software/brutus.html>

# Corpora Resources

- (English) PropBank  
<http://verbs.colorado.edu/~mpalmer/projects/ace.html>
- FrameNet  
<http://framenet.icsi.berkeley.edu>
- Korean PropBank  
<http://www ldc.upenn.edu/>
- Chinese PropBank  
<http://verbs.colorado.edu/chinese/cpb/>
- AnCora corpus: Spanish and Catalan  
<http://http://clic.ub.edu/ancora/>
- Prague Dependency Treebank: Czech  
<http://ufal.mff.cuni.cz/pdt2.0/>
- Penn Arabic TreeBank: Arabic  
<http://www.ircs.upenn.edu/arabic/>

# Corpora Resources

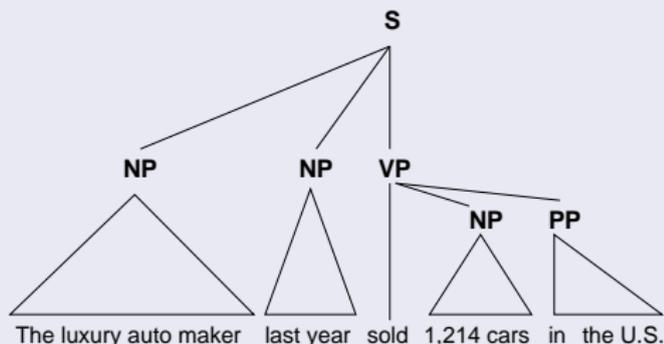
- (English) **PropBank**  
<http://verbs.colorado.edu/~mpalmer/projects/ace.html>
- FrameNet  
<http://framenet.icsi.berkeley.edu>
- Korean PropBank  
<http://www ldc.upenn.edu/>
- Chinese PropBank  
<http://verbs.colorado.edu/chinese/cpb/>
- AnCora corpus: Spanish and Catalan  
<http://http://clic.ub.edu/ancora/>
- Prague Dependency Treebank: Czech  
<http://ufal.mff.cuni.cz/pdt2.0/>
- Penn Arabic TreeBank: Arabic  
<http://www.ircs.upenn.edu/arabic/>

# Corpora Resources

## PropBank

(Palmer et al., 2005)

- **Syntax**-based approach: explaining the varied expression of verb arguments within syntactic positions
- Annotation of all verbal predicates in WSJ (Penn Treebank)
- <http://verbs.colorado.edu/~mpalmer/projects/ace.html>
- Add a semantic layer to the Syntactic Trees

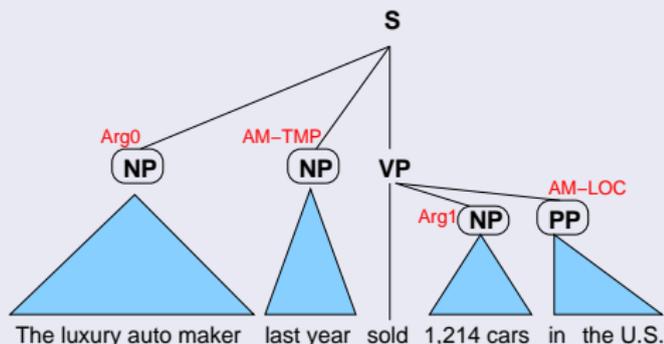


# Corpora Resources

## PropBank

(Palmer et al., 2005)

- **Syntax**-based approach: explaining the varied expression of verb arguments within syntactic positions
- Annotation of all verbal predicates in WSJ (Penn Treebank)
- <http://verbs.colorado.edu/~mpalmer/projects/ace.html>
- Add a semantic layer to the Syntactic Trees



# Corpora Resources

## PropBank

(Palmer et al., 2005)

- Theory neutral numbered core roles (Arg0, Arg1, etc.)
  - ⇒ Interpretation of roles: verb-specific **framesets**
  - ⇒ **Arg0** and **Arg1** usually correspond to prototypical **Agent** and **Patient/Theme** roles. Other arguments do not consistently generalize across verbs
  - ⇒ Different senses have different framesets
  - ⇒ Syntactic alternations that preserve meaning are kept together in a single frameset
- Closed set of 13 general labels for Adjuncts (e.g., Temporal, Manner, Location, etc.)

# Corpora Resources

## PropBank

(Palmer et al., 2005)

- Theory neutral numbered core roles (Arg0, Arg1, etc.)
  - ⇒ Interpretation of roles: verb-specific **framesets**
  - ⇒ **Arg0** and **Arg1** usually correspond to prototypical **Agent** and **Patient/Theme** roles. Other arguments do not consistently generalize across verbs
  - ⇒ Different senses have different framesets
  - ⇒ Syntactic alternations that preserve meaning are kept together in a single frameset
- Closed set of 13 general labels for Adjuncts (e.g., Temporal, Manner, Location, etc.)

# Corpora Resources

## PropBank

(Palmer et al., 2005)

- Theory neutral numbered core roles (Arg0, Arg1, etc.)
  - ⇒ Interpretation of roles: verb-specific **framesets**
  - ⇒ **Arg0** and **Arg1** usually correspond to prototypical **Agent** and **Patient/Theme** roles. Other arguments do not consistently generalize across verbs
  - ⇒ Different senses have different framesets
  - ⇒ Syntactic alternations that preserve meaning are kept together in a single frameset
- Closed set of 13 general labels for Adjuncts (e.g., Temporal, Manner, Location, etc.)

## Corpora Resources

## PropBank: Frame files

(Palmer et al., 2005)

- **sell.01**: commerce: seller  
 Arg0= "seller" (*agent*); Arg1= "thing sold" (*theme*); Arg2= "buyer" (*recipient*); Arg3= "price paid"; Arg4= "benefactive"  
 [Al Brownstein]<sub>Arg0</sub> **sold** [it]<sub>Arg1</sub> [for \$60 a bottle]<sub>Arg3</sub>
- **sell.02**: give up  
 Arg0= "entity selling out"  
 [John]<sub>Arg0</sub> **sold out**
- **sell.03**: sell until none is/are left  
 Arg0= "seller"; Arg1= "thing sold"; ...  
 [The new Harry Potter]<sub>Arg1</sub> **sold out** [within 20 minutes]<sub>ArgM-TMP</sub>

# Corpora Resources

## PropBank: Frame files

(Palmer et al., 2005)

- sell.01:** commerce: seller  
 Arg0= “seller” (*agent*); Arg1= “thing sold” (*theme*); Arg2= “buyer” (*recipient*); Arg3= “price paid”; Arg4= “benefactive”  
 [Al Brownstein]<sub>Arg0</sub> **sold** [it]<sub>Arg1</sub> [for \$60 a bottle]<sub>Arg3</sub>
- sell.02:** give up  
 Arg0= “entity selling out”  
 [John]<sub>Arg0</sub> **sold out**
- sell.03:** sell until none is/are left  
 Arg0= “seller”; Arg1= “thing sold”; ...  
 [The new Harry Potter]<sub>Arg1</sub> **sold out** [within 20 minutes]<sub>ArgM-TMP</sub>

# Applications

## Examples of applications of SRL (I)

- Information Extraction (Surdeanu et al., 2003)
- Question & Answering (Narayanan and Harabagiu, 2004; Frank et al., 2007; Shen and Lapata, 2007)
- Automatic Summarization (Melli et al., 2005)
- Coreference Resolution (Ponzetto and Strube, 2006)
- Text Categorization (Person et al., 2010)
- Opinion Expression Detection (Johansson and Moschitti, 2010)

# Applications

## Examples of applications of SRL (II)

- Machine Translation Evaluation (Giménez and Màrquez, 2007)
- Machine Translation (Boas, 2002; Wu and Fung, 2009a;2009b)
- Textual Entailment (Tatu & Moldovan, 2005; Burchardt et al., 2007)
- Modeling Early Language Acquisition (Connor et al., 2008;2009)
- Pictorial Communication Systems (Goldberg, et al., 2008)

# Empirical Evaluation of SRL Systems

## Evaluation Exercises

- Up to 10 evaluation exercises in the last 7 years
  - ⇒ CoNLL-2004/2005 shared tasks  
(Carreras & Màrquez, 2004; 2005)
  - ⇒ Senseval-3 (Litkowski, 2004)
  - ⇒ SemEval-2007 (Pradhan et al., 2007; Màrquez et al., 2007)  
(Baker et al., 2007; Litkowski & Hargraves, 2007)
  - ⇒ CoNLL-2008 shared task (Surdeanu et al., 2008)
  - ⇒ CoNLL-2009 shared task (Hajič et al., 2009)
  - ⇒ SemEval-2010 (Ruppenhofer et al., 2010)

# Talk Overview

- 1 Semantic Role Labeling
- 2 The Statistical Approach to SRL**
- 3 Semantic Features for SRL
- 4 Conclusions

# SRL Architecture: Step by Step

## Step 1: Select argument candidates

- Given a sentence and a designated predicate
- Parse the sentence
- Identify candidates in tree constituents (filtering/pruning)
  - ⇒ Simple heuristic rules can be used, which maintain a high recall (Xue & Palmer, 2004)
- **Key point:** 95% of semantic arguments coincide with unique syntactic constituents in the gold parse tree (PropBank)
  - ⇒ Matching is still ~90% when using automatic parsers

# SRL Architecture: Step by Step

## Step 1: Select argument candidates

- Given a sentence and a designated predicate
- Parse the sentence
- Identify candidates in tree constituents (filtering/pruning)
  - ⇒ Simple heuristic rules can be used, which maintain a high recall (Xue & Palmer, 2004)
- **Key point:** 95% of semantic arguments coincide with unique syntactic constituents in the gold parse tree (PropBank)
  - ⇒ Matching is still ~90% when using automatic parsers

# SRL Architecture: Step by Step

## Step 2: Local scoring of candidates

- Apply classifiers to **assign confidence scores** to argument candidates (all labels + 'non-argument')
- Candidates are **treated independently** of each other
- *Identification* and *Classification* may be performed separately
  - Computational reasons but also modularity in feature engineering
- Many ML paradigms have been used: not big differences
- Features are more important

# SRL Architecture: Step by Step

## Step 2: Local scoring of candidates

- Apply classifiers to **assign confidence scores** to argument candidates (all labels + 'non-argument')
- Candidates are **treated independently** of each other
- *Identification* and *Classification* may be performed separately
  - ⇒ Computational reasons but also modularity in feature engineering
- Many ML paradigms have been used: not big differences
- Features are more important

# SRL Architecture: Step by Step

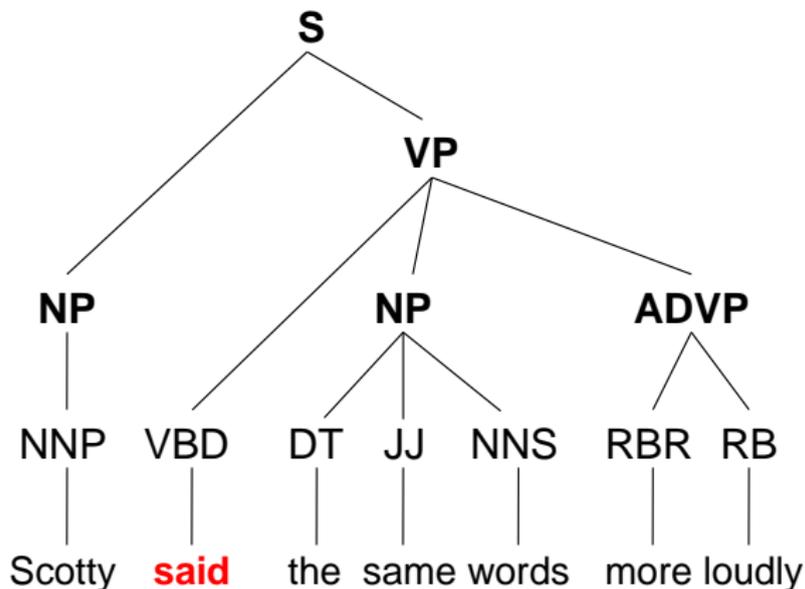
## Step 2: Local scoring of candidates

- Apply classifiers to **assign confidence scores** to argument candidates (all labels + 'non-argument')
- Candidates are **treated independently** of each other
- *Identification* and *Classification* may be performed separately
  - ⇒ Computational reasons but also modularity in feature engineering
- Many ML paradigms have been used: not big differences
- Features are more important

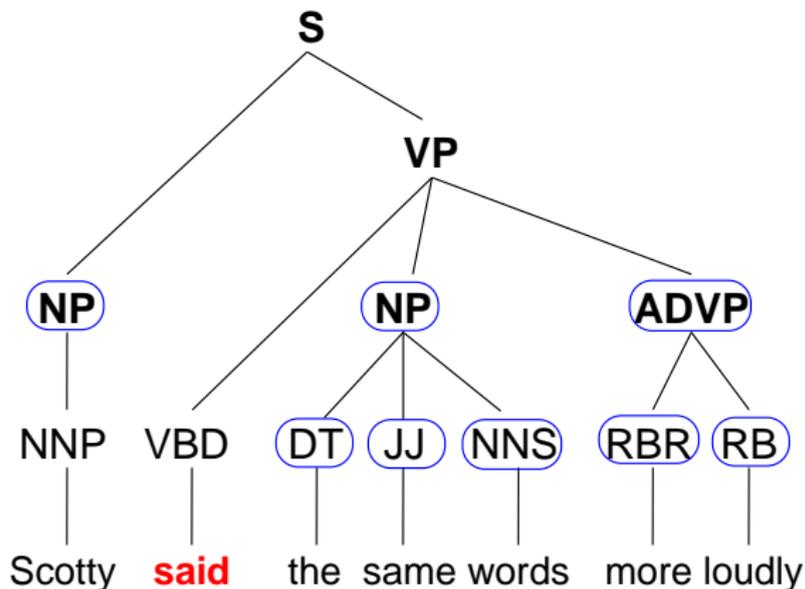
# SRL Architecture: Steps 1 + 2

Scotty **said** the same words more loudly

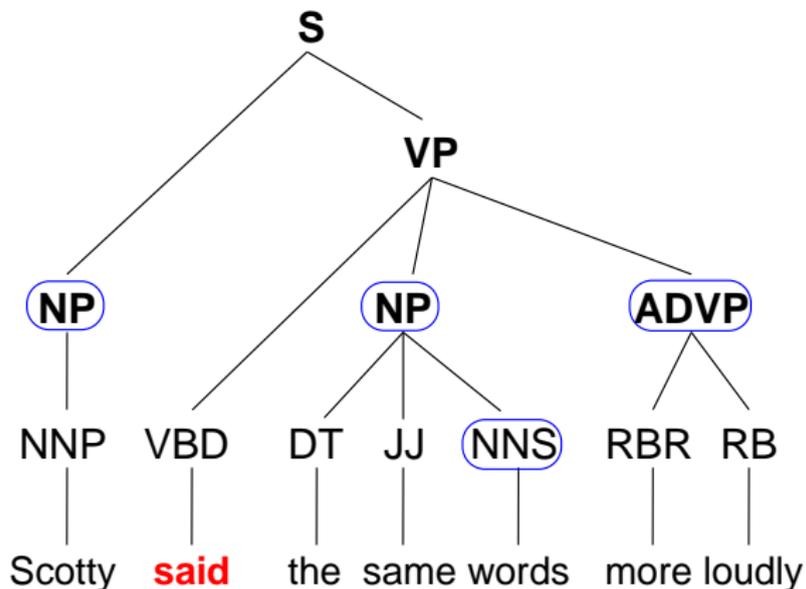
## SRL Architecture: Steps 1 + 2



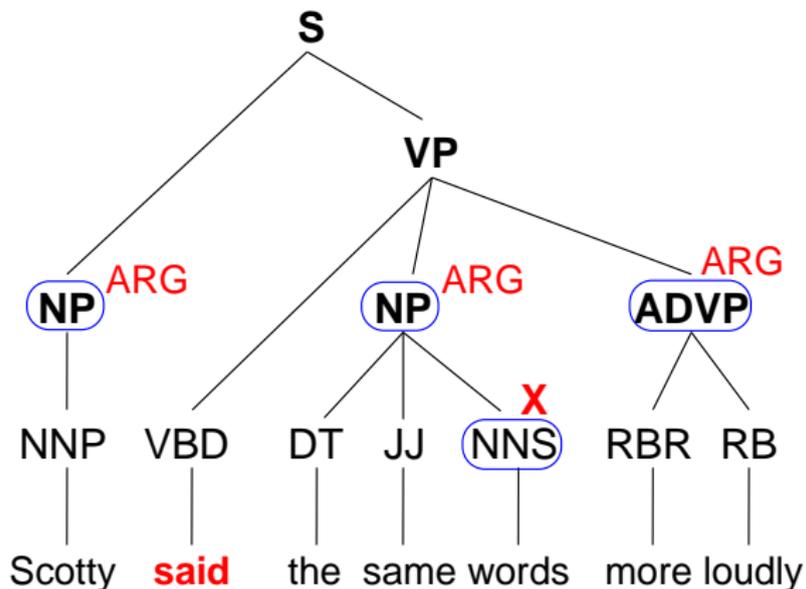
## SRL Architecture: Steps 1 + 2



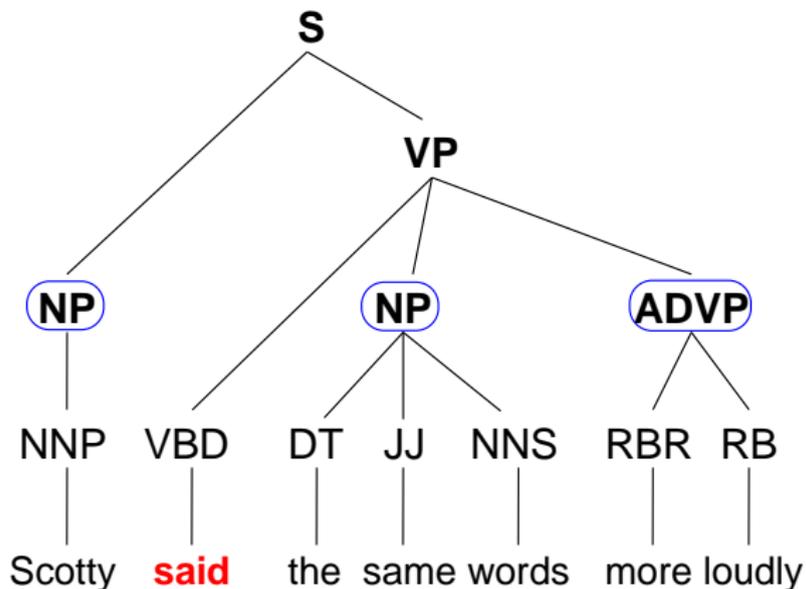
## SRL Architecture: Steps 1 + 2



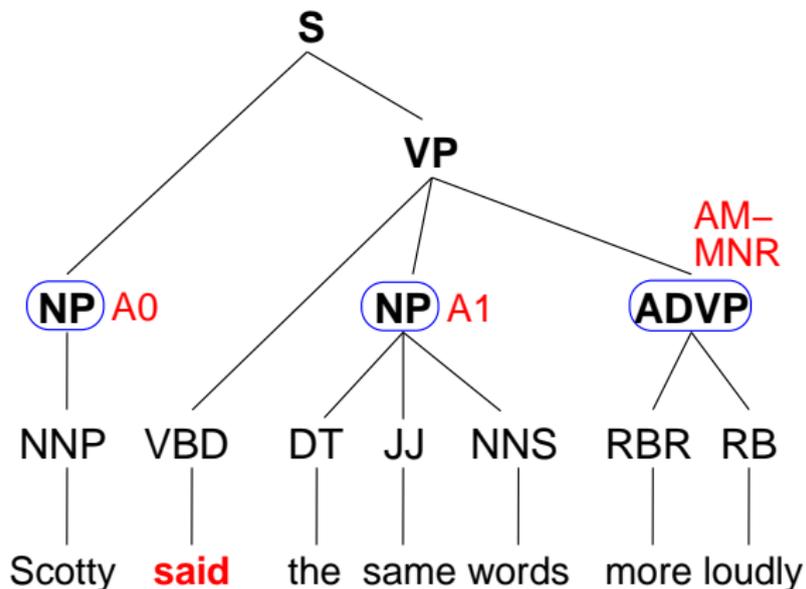
## SRL Architecture: Steps 1 + 2



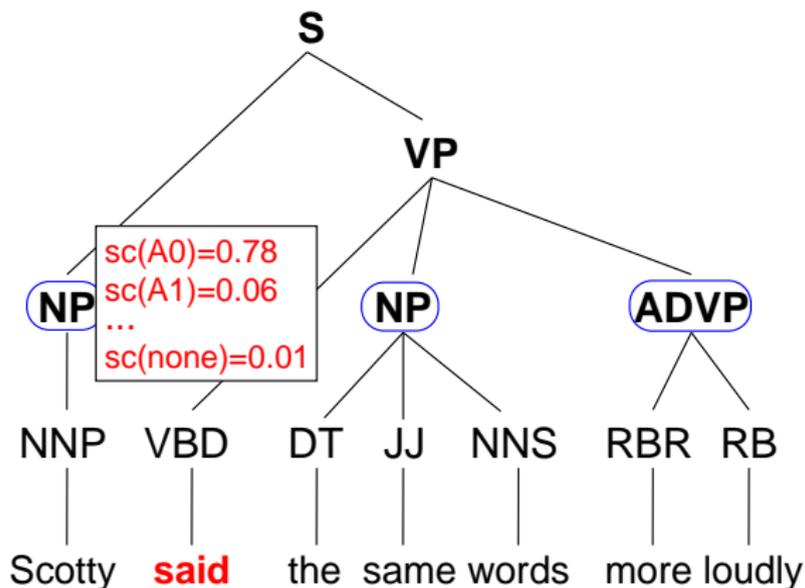
## SRL Architecture: Steps 1 + 2



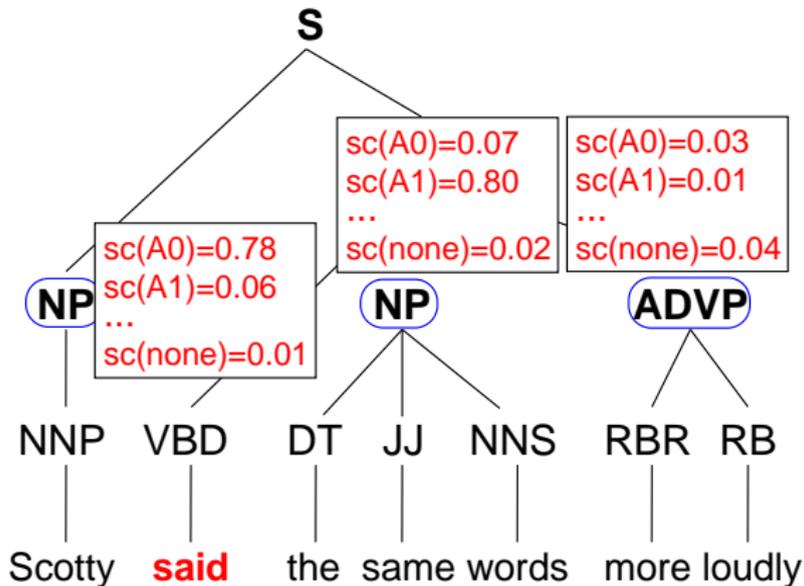
## SRL Architecture: Steps 1 + 2



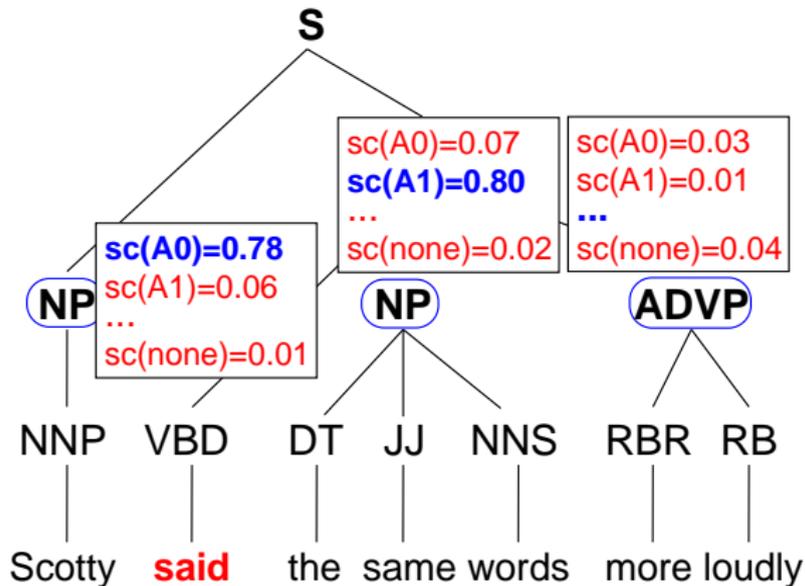
## SRL Architecture: Motivating next step (joint scoring)



## SRL Architecture: Motivating next step (joint scoring)

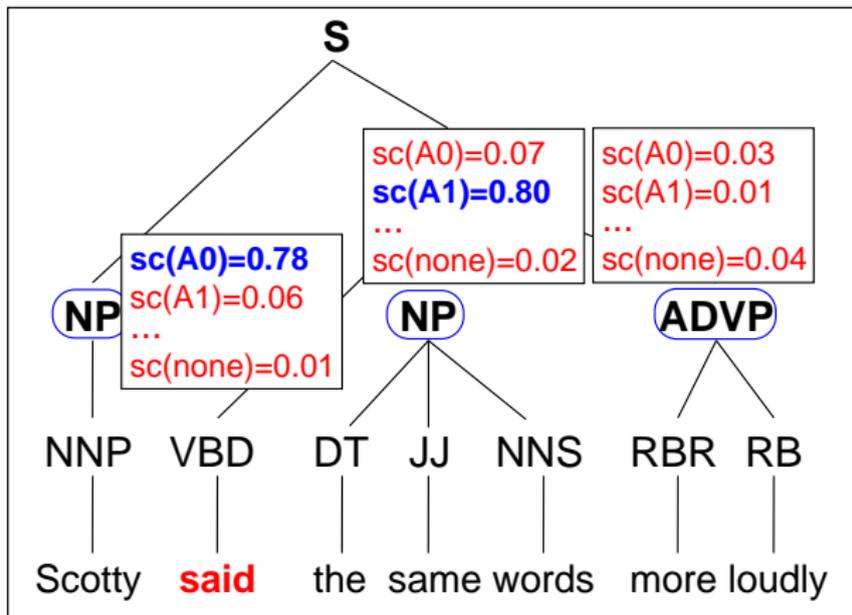


## SRL Architecture: Motivating next step (joint scoring)

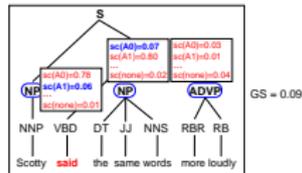
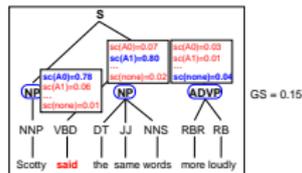
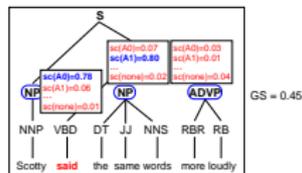


## SRL Architecture: Motivating next step (joint scoring)

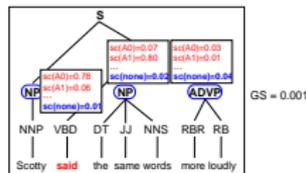
Global Score = 0.30



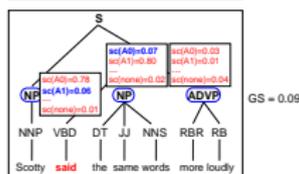
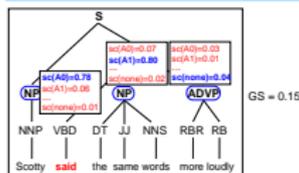
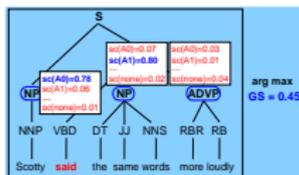
## SRL Architecture: Motivating next step (joint scoring)



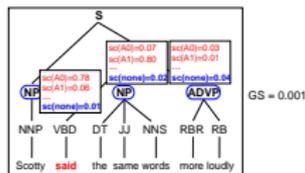
...



## SRL Architecture: Motivating next step (joint scoring)



...



# SRL Architecture: Step by Step

## Step 3: Joint scoring — Paradigmatic examples

- Combine local predictions through ILP to find the best solution according to structural and linguistic constraints  
(Koomen et al., 2005; Punyakanok et al., 2008)

-learning +features +search

- Re-ranking of several candidate solutions  
(Haghighi et al., 2005; Toutanova et al., 2008)

+learning +features -search

- Global search integrating joint scoring: Tree CRFs  
(Cohn & Blunsom, 2005)

+learning +/-features +/-search

# SRL Architecture: Step by Step

## Step 3: Joint scoring — Paradigmatic examples

- Combine local predictions through ILP to find the best solution according to structural and linguistic constraints  
(Koomen et al., 2005; Punyakanok et al., 2008)

`-learning +features +search`

- Re-ranking of several candidate solutions  
(Haghighi et al., 2005; Toutanova et al., 2008)

`+learning +features -search`

- Global search integrating joint scoring: Tree CRFs  
(Cohn & Blunsom, 2005)

`+learning +/-features +/-search`

# SRL Architecture: Step by Step

## Step 4: Post-processing

- Application of a set of heuristic rules to:
  - Correct **frequent errors**
  - Enforce consistency** in the solution

## Architecture

- Use a probabilistic local SRL model to produce multiple ( $n$ -best) candidate solutions for the predicate structure
- Use a feature-rich reranking model to select the best solution among them

**Main goal:** is to build a rich model for joint scoring, which takes into account the dependencies among the labels of argument phrases

## Local Steps

- i. Parse the sentence and apply pruning (Xue & Palmer, 2004) to filter argument candidates for a given predicate  $p$
- ii. Apply a simple local scoring model trained with log-linear classifiers (MaxEnt):  $P(\text{label}_i | \text{node}, p)$  probability distribution
- iii. Consider a simple global scoring scheme assuming independence of local assignments:  
$$P_{\text{LOCAL}}(L | \text{tree}, p) = \prod_{\text{node}_i \in \text{tree}} P(\text{label}_i | \text{node}_i, p)$$
- iv. Use dynamic programming to find the  $n$ -most probable non-overlapping complete labelings for predicate  $p$

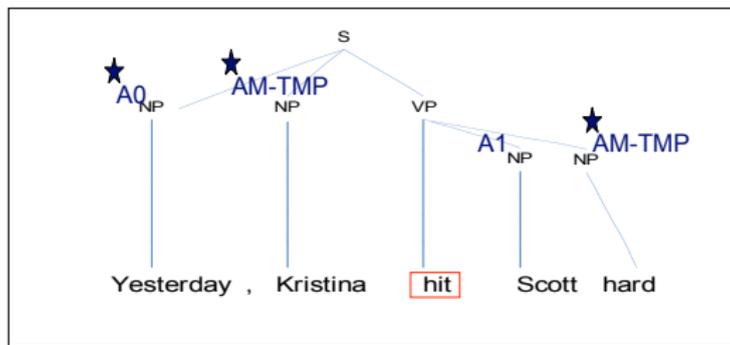
## Reranking Step

- i. Consider a reranking model trained to select the best among the  $n$ -most probable complete labelings; again a log-linear model:  $P_{JOINT}(L_i|tree, p)$
- ii. Consider the following combination of local and joint scoring models:  $\log(P_{SRL}(L|tree, p)) = \log(P_{JOINT}(L|tree, p)) + \lambda \log(P_{LOCAL}(L|tree, p))$
- iii. Select the complete labeling ( $L_i \in \{L_1, L_2, \dots, L_n\}$ ) that maximizes the previous formula (reranking)

Features: joint scoring

slide from (Yih &amp; Toutanova, 2006)

## Joint Model Features



**Repetition features:** count of arguments with a given label  $c(\text{AM-TMP})=2$

**Complete sequence syntactic-semantic features for the core arguments:**

[NP\_A0 hit NP\_A1], [NP\_A0 VBD NP\_A1] (backoff)

[NP\_A0 hit] (left backoff)

[NP\_ARG hit NP\_ARG] (no specific labels)

[1 hit 1] (counts of left and right core arguments)

## Enhancement by using multiple trees

- For top  $k$  trees from Charniak's parser,  $t_1, t_2, \dots, t_k$ , find corresponding best SRL assignments  $L_1, L_2, \dots, L_k$  and choose the tree and assignment that maximize the score (approx. joint probability of tree and assignment)

$$\text{score}(L_i, t_i) = \alpha \log(P(t_i)) + \log(P_{SRL}(L_i|t_i))$$

- **Final Results** (2nd best at CoNLL):  
WSJ-23: 78.45 ( $F_1$ ), 79.54 (Prec.), 77.39 (Rec.)  
Brown: 67.71 ( $F_1$ ), 70.24 (Prec.), 65.37 (Rec.)  
Bug-fixed post-evaluation: **80.32**  $F_1$  (WSJ) **68.81**  $F_1$  (Brown)
- Improvement due to the joint model:  $> 2 F_1$  points

## Enhancement by using multiple trees

- For top  $k$  trees from Charniak's parser,  $t_1, t_2, \dots, t_k$ , find corresponding best SRL assignments  $L_1, L_2, \dots, L_k$  and choose the tree and assignment that maximize the score (approx. joint probability of tree and assignment)  
 $score(L_i, t_i) = \alpha \log(P(t_i)) + \log(P_{SRL}(L_i|t_i))$
- **Final Results** (2nd best at CoNLL):  
WSJ-23: 78.45 ( $F_1$ ), 79.54 (Prec.), 77.39 (Rec.)  
Brown: 67.71 ( $F_1$ ), 70.24 (Prec.), 65.37 (Rec.)  
Bug-fixed post-evaluation: **80.32**  $F_1$  (WSJ) **68.81**  $F_1$  (Brown)
- Improvement due to the joint model:  $> 2 F_1$  points

## Enhancement by using multiple trees

- For top  $k$  trees from Charniak's parser,  $t_1, t_2, \dots, t_k$ , find corresponding best SRL assignments  $L_1, L_2, \dots, L_k$  and choose the tree and assignment that maximize the score (approx. joint probability of tree and assignment)  
 $score(L_i, t_i) = \alpha \log(P(t_i)) + \log(P_{SRL}(L_i|t_i))$
- **Final Results** (2nd best at CoNLL):  
WSJ-23: 78.45 ( $F_1$ ), 79.54 (Prec.), 77.39 (Rec.)  
Brown: 67.71 ( $F_1$ ), 70.24 (Prec.), 65.37 (Rec.)  
Bug-fixed post-evaluation: **80.32**  $F_1$  (WSJ) **68.81**  $F_1$  (Brown)
- Improvement due to the joint model:  $> 2 F_1$  points

# SRL Architecture

## Exceptions to the standard architecture

- SRL as sequential tagging  
(Hacioglu et al., 2004; Màrquez et al., 2005; Surdeanu et al., 2007)
- Joint treatment of all predicates in the sentence  
(Carreras et al., 2004; Surdeanu et al., 2008)

# SRL Architecture

## Exceptions to the standard architecture

- Parsing variations for SRL
  - ⇒ Syntactic parser trained to predict argument candidates (Yi & Palmer, 2005)
  - ⇒ Joint parsing and SRL: semantic parsing (Musillo & Merlo, 2006; Merlo & Musillo, 2008)
  - ⇒ SRL based on dependency parsing (Johansson & Nugues, 2007)
  - ⇒ Systems from the CoNLL-2008 and 2009 shared tasks (Surdeanu et al., 2008; Hajič et al., 2009)
  - ⇒ CCG parser (Gildea and Hockenmaier, 2005; Boxwell et al., 2009)
  - ⇒ HPSG parsers with handcrafted grammars (Zhang et al., 2008; 2009)
- SRL using Markov Logic Networks (Meza-Ruiz & Riedel, 2008; 2009)

# SRL Architecture

## Exceptions to the standard architecture

- Parsing variations for SRL
  - ⇒ Syntactic parser trained to predict argument candidates (Yi & Palmer, 2005)
  - ⇒ Joint parsing and SRL: semantic parsing (Musillo & Merlo, 2006; Merlo & Musillo, 2008)
  - ⇒ SRL based on dependency parsing (Johansson & Nugues, 2007)
  - ⇒ Systems from the CoNLL-2008 and 2009 shared tasks (Surdeanu et al., 2008; Hajič et al., 2009)
  - ⇒ CCG parser (Gildea and Hockenmaier, 2005; Boxwell et al., 2009)
  - ⇒ HPSG parsers with handcrafted grammars (Zhang et al., 2008; 2009)
- SRL using Markov Logic Networks (Meza-Ruiz & Riedel, 2008; 2009)

# Feature Engineering

Features: local scoring

(Gildea & Jurafsky, 2002)

- Highly influential for the SRL work. They characterize:
  - i. The candidate argument (constituent) and its context: **phrase type**, **head word**, **governing category** of the constituent
  - ii. The verb predicate and its context: **lemma**, **voice**, **subcategorization pattern** of the verb
  - iii. The relation between the constituent and the predicate: **position** of the constituent with respect to the verb, **category path** between them.

# Feature Engineering

## Features: local scoring — extensions

- “Brute force” features. Applied to the constituent and possibly to parent and siblings:
  - ⇒ **First and last words/POS** in the constituent, **bag-of-words**, ***n*-grams** of POS, and **sequence of top syntactic elements** in the constituent.
- Linguistically-inspired features
  - ⇒ **Content word**, **named entities** (Surdeanu et al., 2003), **syntactic frame** (Xue & Palmer, 2004), **path variations**, **semantic compatibility between constituent head and predicate** (Zapirain et al., 2007, 2009), etc.
- Significant (and cumulative) increase in performance

# Feature Engineering

## Features: local scoring — extensions

- “Brute force” features. Applied to the constituent and possibly to parent and siblings:
  - ⇒ **First and last words/POS** in the constituent, **bag-of-words**, ***n*-grams** of POS, and **sequence of top syntactic elements** in the constituent.
- Linguistically-inspired features
  - ⇒ **Content word**, **named entities** (Surdeanu et al., 2003), **syntactic frame** (Xue & Palmer, 2004), **path variations**, **semantic compatibility** between constituent head and predicate (Zapirain et al., 2007; 2009), etc.
- Significant (and cumulative) increase in performance

# Feature Engineering

## Features: local scoring — extensions

- “Brute force” features. Applied to the constituent and possibly to parent and siblings:
  - ⇒ **First and last words/POS** in the constituent, **bag-of-words**, ***n*-grams** of POS, and **sequence of top syntactic elements** in the constituent.
- Linguistically-inspired features
  - ⇒ **Content word**, **named entities** (Surdeanu et al., 2003), **syntactic frame** (Xue & Palmer, 2004), **path variations**, **semantic compatibility** between constituent head and predicate (Zapirain et al., 2007; 2009), etc.
- Significant (and cumulative) increase in performance

# Feature Engineering

## Features: joint scoring

- Richer features taking into account information from several arguments at a time
- Best example: when doing re-ranking one may codify patterns on the whole candidate argument structure (Hiaghighi et al., 2005; Toutanova et al., 2008)
- Good for capturing **global preferences**

# Feature Engineering

## Features: the Kernel approach

- **Knowledge poor** approach
- Let the kernel function to compute the similarity/differences between examples by considering all possible substructures as features
- Motivation: avoid intense knowledge engineering
- Potentially useful for rapid system development and working with under resourced languages
- Mostly variants of Collins' **all-subtrees** convolution kernel (Moschitti et al., 2008; Pighin & Moschitti, 2009; 2010)

# Feature Engineering

## Features: the Kernel approach

- **Knowledge poor** approach
- Let the kernel function to compute the similarity/differences between examples by considering all possible substructures as features
- Motivation: avoid intense knowledge engineering
- Potentially useful for rapid system development and working with under resourced languages
- Mostly variants of Collins' **all-subtrees** convolution kernel (Moschitti et al., 2008; Pighin & Moschitti, 2009; 2010)

# Feature Engineering

## Features: the Kernel approach

- **Knowledge poor** approach
- Let the kernel function to compute the similarity/differences between examples by considering all possible substructures as features
- Motivation: avoid intense knowledge engineering
- Potentially useful for rapid system development and working with under resourced languages
- Mostly variants of Collins' **all-subtrees** convolution kernel (Moschitti et al., 2008; Pighin & Moschitti, 2009; 2010)

# Talk Overview

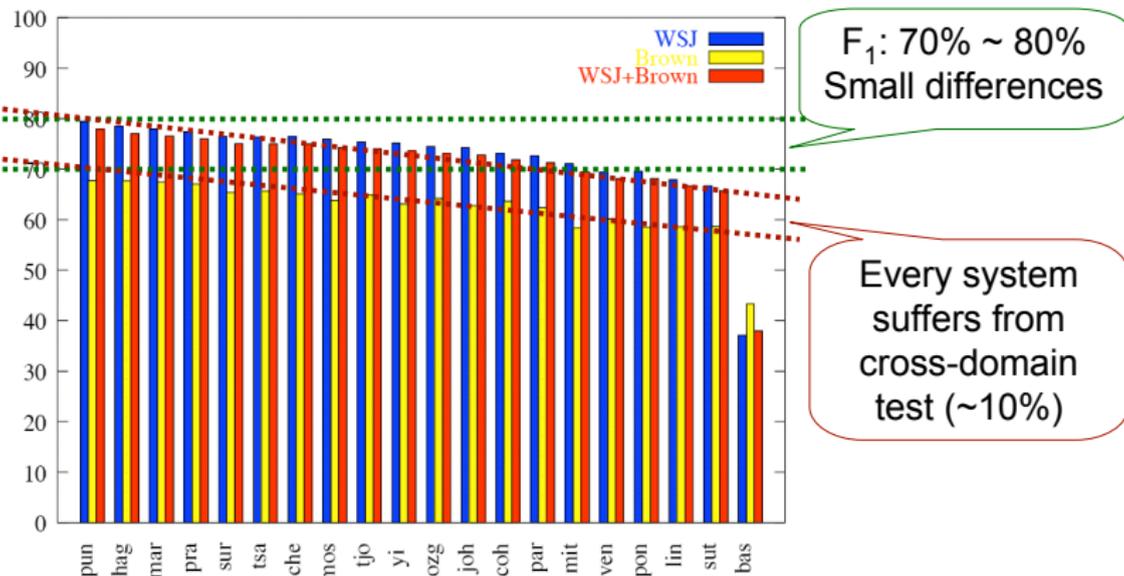
- 1 Semantic Role Labeling
- 2 The Statistical Approach to SRL
- 3 Semantic Features for SRL
  - Joint work with **E. Agirre**, **M. Surdeanu** and **B. Zapirain**
- 4 Conclusions

# Talk Overview

- 1 Semantic Role Labeling
- 2 The Statistical Approach to SRL
- 3 Semantic Features for SRL
  - Joint work with **E. Agirre**, **M. Surdeanu** and **B. Zapirain**
- 4 Conclusions

## Results from CoNLL-2005 shared task

## Results on WSJ and Brown Tests



# Results from CoNLL-2005 shared task

## Reasons for the low generalization ability

- The training corpus is not representative and large enough (and it will never be)
- Taggers and syntactic parsers also experience a significant drop in performance
- The main loss in performance takes place in role classification, not identification — **semantic explanation**  
(Pradhan et al., 2008)

# Semantic Features for SRL

## Motivation

- Most current systems capture semantics through lexicalized features on the predicate and the head word of the argument to be classified

- But lexical features are *sparse* and *generalize badly*

[JFK]<sub>Patient</sub> *was\_assassinated* [in Dallas]<sub>LOC</sub>

[JFK]<sub>Patient</sub> *was\_assassinated* [in November]<sub>TMP</sub>

- [in Texas]???, [in autumn]???

# Semantic Features for SRL

## Motivation

- Most current systems capture semantics through lexicalized features on the predicate and the head word of the argument to be classified
- But lexical features are **sparse** and **generalize badly**

[JFK]<sub>Patient</sub> was\_assassinated [in Dallas]<sub>LOC</sub>

[JFK]<sub>Patient</sub> was\_assassinated [in November]<sub>TMP</sub>

- [in Texas]???, [in autumn]???

# Semantic Features for SRL

## Motivation

- Most current systems capture semantics through lexicalized features on the predicate and the head word of the argument to be classified
- But lexical features are **sparse** and **generalize badly**

[JFK]<sub>Patient</sub> was\_assassinated [in Dallas]<sub>LOC</sub>

[JFK]<sub>Patient</sub> was\_assassinated [in November]<sub>TMP</sub>

- [in Texas]???, [in autumn]???

# Semantic Features for SRL

## Motivation

Selectional Preferences and distributional similarity techniques should help us to classify arguments with low-frequency or unknown head words

[Dallas  $\approx$  Texas]*Location*, [November  $\approx$  autumn]*Temporal*

# Previous Work

## Selectional Preferences

- Modeling semantic preferences that predicates impose on their arguments
- Long tradition of automatic acquisition of selectional preferences (SPs) from corpora. WordNet-based and distributional models of SPs

(Resnik, 1993; Pantel and Lin, 2000; Brockmann and Lapata, 2003)  
(Erk 2007; Erk et al., 2011; etc.)

- ⇒ e.g., estimate plausibility of triples:  
(verb, argument, head-word)
- ⇒ useful for syntactic-semantic disambiguation

# Previous Work

## SPs applied to Semantic Role Labeling

- (Gildea and Jurafsky, 2002) – FrameNet
  - ⇒ First researchers to apply selectional preferences to SRL
  - ⇒ Distributional clustering and WordNet-based techniques to generalize argument heads
  - ⇒ Slight improvement in role classification (NP arguments)
- Zapirain et al. (2009; 2010) – PropBank
  - ⇒ Show that selectional preferences can improve semantic role classification in a state-of-the-art SRL system

## Two types of selectional preferences (SP)

- i. **verb**-*role*: list of heads of NP arguments of the predicate **verb** that are labeled with the role *role*

write-Arg0: Angrist anyone baker ball bank Barlow Bates ...  
 write-Arg1: abstract act analysis article asset bill book ...  
 write-Arg2: bank commander hundred jaguar Kemp member ...  
 write-AM-LOC: paper space ...

...

- ii. **prep**-*role*: list of nominal heads of PP arguments with preposition **prep** that are labeled with the role *role*

from-Arg2: academy account acquisition activity ad ...  
 from-Arg3: activity advertising agenda airport ...  
 from-Arg4: europe Golenbock system Vizcaya west  
 from-AM-TMP: april august beginning bell day dec. half ...  
 from-AM-LOC: agency area asia body bureau orlando ...

...

SP models:  $SP_{sim}(p, r, w)$  plausibility score

- **Discriminative approach:** given a new argument of a predicate  $p$ , we compare its head ( $w$ ) to the selectional preference of each possible role label  $r$ , i.e., we want to find the role with the selectional preference that fits the head best
- We compute the compatibility scores using two different methods
  - ⇒ WordNet based —using (Resnik, 1993)
  - ⇒ Based on distributional similarity —a la Erk (2007)

SP models:  $SP_{sim}(p, r, w)$  plausibility score

- **Discriminative approach**: given a new argument of a predicate  $p$ , we compare its head ( $w$ ) to the selectional preference of each possible role label  $r$ , i.e., we want to find the role with the selectional preference that fits the head best
- We compute the compatibility scores using two different methods
  - ⇒ WordNet based —using (Resnik, 1993)
  - ⇒ Based on distributional similarity —a la Erk (2007)

## WordNet SP models

- Resnik formula (1993) is used to precalculate a weighted list of relevant synsets for the lists of words contained in the SPs

SP **write-Arg0**: Angrist anyone baker ball bank Barlow Bates ...

n#00002086 5.875 **life form** organism being living thing "any living entity"  
 n#00001740 5.737 **entity** something "anything having existence (living or nonliving)"  
 n#00009457 4.782 **object** physical object "a physical (tangible and visible) entity;"  
 n#00004123 4.351 **person** individual someone somebody mortal human soul "a human being;"  
 ...

SP **write-Arg1**: abstract act analysis article asset bill book ...

n#00019671 7.956 **communication** "something that is communicated between people or groups"  
 n#04949838 4.257 **message** content subject matter substance "what a communication that ..."  
 n#00018916 3.848 **relation** "an abstraction belonging to or characteristic of two entities"  
 n#00013018 3.574 **abstraction** "a concept formed by extracting common features from examples"  
 ...

## WordNet SP models

- At test time, for a new argument of the predicate **write** with head word **book**:
  - ⇒ consider  $S = \{\langle \text{book} \rangle\} \cup$  “all its hypernyms in WordNet” (for all senses of book)
  - ⇒  $SP_{Res}(\text{write}, \text{Arg1}, \text{book})$  returns the sum of the weights of the sysnsets in  $S$  matching the sysnsets in the list corresponding to the SP **write-Arg1**

Distributional SP models: based on Erk's (2007) setting

JFK was assassinated [in Texas]???

SP **in-TMP**: *November, century, month*

SP **in-LOC**: *Dallas, railway, city*

$$SP_{sim}(p, r, w) = \sum_{w_i \in \text{Seen}(p, r)} sim(w, w_i) \cdot weight(p, r, w_i)$$

$$SP(in, TMP, Texas) = sim(Texas, November) \cdot weight(in, TMP, November) + \\ sim(Texas, century) \cdot weight(in, TMP, century) + \\ sim(Texas, month) \cdot weight(in, TMP, month)$$

Distributional SP models: based on Erk's (2007) setting

JFK was assassinated [in **Texas**]???

SP **in**-*TMP*: *November, century, month*

SP **in**-*LOC*: *Dallas, railway, city*

$$SP_{sim}(p, r, w) = \sum_{w_i \in \text{Seen}(p, r)} sim(w, w_i) \cdot weight(p, r, w_i)$$

$$SP(\text{in}, \text{TMP}, \text{Texas}) = sim(\text{Texas}, \text{November}) \cdot weight(\text{in}, \text{TMP}, \text{November}) + \\ sim(\text{Texas}, \text{century}) \cdot weight(\text{in}, \text{TMP}, \text{century}) + \\ sim(\text{Texas}, \text{month}) \cdot weight(\text{in}, \text{TMP}, \text{month})$$

Distributional SP models: based on Erk's (2007) setting

JFK was assassinated [in **Texas**]???

SP **in-TMP**: *November, century, month*

SP **in-LOC**: *Dallas, railway, city*

$$SP_{sim}(p, r, w) = \sum_{w_i \in Seen(p, r)} sim(w, w_i) \cdot weight(p, r, w_i)$$

$$SP(in, LOC, \mathbf{Texas}) = sim(\mathbf{Texas}, Dallas) \cdot freq(in, LOC, Dallas) + \\ sim(\mathbf{Texas}, railway) \cdot freq(in, LOC, railway) + \\ sim(\mathbf{Texas}, city) \cdot freq(in, LOC, city)$$

$$SP(in, \mathbf{LOC}, \mathbf{Texas}) > SP(in, \mathbf{TMP}, \mathbf{Texas})$$

Distributional SP models: based on Erk's (2007) setting

JFK was assassinated [in Texas]???

SP **in-TMP**: *November, century, month*

SP **in-LOC**: *Dallas, railway, city*

$$SP_{sim}(p, r, w) = \sum_{w_i \in Seen(p, r)} sim(w, w_i) \cdot weight(p, r, w_i)$$

$$SP(in, LOC, Texas) = sim(Texas, Dallas) \cdot freq(in, LOC, Dallas) + \\ sim(Texas, railway) \cdot freq(in, LOC, railway) + \\ sim(Texas, city) \cdot freq(in, LOC, city)$$

$$SP(in, LOC, Texas) > SP(in, TMP, Texas)$$

Distributional SP models: various instantiations for *sim*

- Using Padó and Lapata's software (2007) for computing distributional similarity measures
  - ⇒ Run on the British National Corpus
  - ⇒ Optimal parameterization as described in the paper
  - ⇒ Jaccard, cosine and Lin's similarity measures:  $sim_{Jac}$ ,  $sim_{cos}$  and  $sim_{Lin}$
- Using the already available Lin's thesaurus (Lin, 1998)
  - ⇒ Direct and second order similarity:  $sim_{Lin}^{th}$ ,  $sim_{Jac}^{th2}$  and  $sim_{cos}^{th2}$
  - ⇒ Average of both directions similarity

**Setting:** Assign role labels to argument head words based solely on SP scores

⇒ For each head word ( $w$ ), select the role ( $r$ ) of the predicate or preposition ( $p$ ) which fits best the head word:

$$R_{sim}(p, w) = \arg \max_{r \in Roles(p)} SP_{sim}(p, r, w)$$

⇒ SPs based on  $(p, r, w)$  triples from CoNLL-2005 data

⇒ In-domain (WSJ) and out-of-domain (Brown) test sets  
CoNLL-2005

⇒ **Lexical baseline** model: for a test pair  $(p, w)$ , assign the role under which the head ( $w$ ) occurred most often in the training data given the predicate ( $p$ )

**Setting:** Assign role labels to argument head words based solely on SP scores

⇒ For each head word ( $w$ ), select the role ( $r$ ) of the predicate or preposition ( $p$ ) which fits best the head word:

$$R_{sim}(p, w) = \arg \max_{r \in Roles(p)} SP_{sim}(p, r, w)$$

⇒ SPs based on  $(p, r, w)$  triples from CoNLL-2005 data

⇒ In-domain (WSJ) and out-of-domain (Brown) test sets  
CoNLL-2005

⇒ **Lexical baseline** model: for a test pair  $(p, w)$ , assign the role under which the head ( $w$ ) occurred most often in the training data given the predicate ( $p$ )

**Setting:** Assign role labels to argument head words based solely on SP scores

⇒ For each head word ( $w$ ), select the role ( $r$ ) of the predicate or preposition ( $p$ ) which fits best the head word:

$$R_{sim}(p, w) = \arg \max_{r \in Roles(p)} SP_{sim}(p, r, w)$$

⇒ SPs based on  $(p, r, w)$  triples from CoNLL-2005 data

⇒ In-domain (WSJ) and out-of-domain (Brown) test sets  
CoNLL-2005

⇒ **Lexical baseline** model: for a test pair  $(p, w)$ , assign the role under which the head ( $w$ ) occurred most often in the training data given the predicate ( $p$ )

## Evaluation of SPs in isolation

(Zapirain et al., 2009; 2010)

	WSJ-test			Brown		
	prec.	rec.	F <sub>1</sub>	prec.	rec.	F <sub>1</sub>
lexical	<b>82.98</b>	<b>43.77</b>	57.31	<b>68.47</b>	<b>13.60</b>	22.69
$SP_{Res}$	63.47	53.24	57.91	55.12	44.15	49.03
$SP_{sim_{Jac}}$	61.83	61.40	61.61	55.42	53.45	54.42
$SP_{sim_{cos}}$	64.67	64.22	64.44	56.56	54.54	55.53
$SP_{sim_{Jac}^{th2}}$	70.82	<b>70.33</b>	<b>70.57</b>	62.37	<b>60.15</b>	<b>61.24</b>
$SP_{sim_{cos}^{th2}}$	70.28	69.80	70.04	62.36	60.14	61.23

- ⇒ Lexical features have a high precision but very low recall
- ⇒ SPs are able to effectively generalize lexical features
- ⇒ SPs based on distributional similarity are better
- ⇒ Second-order similarity variants (Lin) attain the best results

## Evaluation of SPs in isolation

(Zapirain et al., 2009; 2010)

	WSJ-test			Brown		
	prec.	rec.	F <sub>1</sub>	prec.	rec.	F <sub>1</sub>
lexical	<b>82.98</b>	43.77	57.31	<b>68.47</b>	<b>13.60</b>	22.69
$SP_{Res}$	63.47	53.24	57.91	55.12	44.15	49.03
$SP_{sim_{Jac}}$	61.83	61.40	61.61	55.42	53.45	54.42
$SP_{sim_{cos}}$	64.67	64.22	64.44	56.56	54.54	55.53
$SP_{sim_{Jac}^{th2}}$	70.82	<b>70.33</b>	<b>70.57</b>	62.37	<b>60.15</b>	<b>61.24</b>
$SP_{sim_{cos}^{th2}}$	70.28	69.80	70.04	62.36	60.14	61.23

- ⇒ Lexical features have a high precision but very low recall
- ⇒ SPs are able to effectively generalize lexical features
- ⇒ SPs based on distributional similarity are better
- ⇒ Second-order similarity variants (Lin) attain the best results

- *SwiRL* system for SRL (Surdeanu et al., 2007)
  - ⇒ System from CoNLL-2005 shared task (PropBank)
  - ⇒ Standard architecture (ML based on AdaBoost and SVMs)
  - ⇒ Best results from single (non-combined) systems at CoNLL-2005
- Simple approach: extending *SwiRL* features with SP predictions
  - ⇒ We train several extended *SwiRL-SP<sub>i</sub>* models, one per selectional preferences model *SP<sub>i</sub>*
  - ⇒ For each example  $(p, w)$  of *SwiRL-SP<sub>i</sub>*, we add a single new feature whose value is the predicted role label  $R_i(p, w)$

## Results

	WSJ-test			Brown		
	Core	Adj	All	Core	Adj	All
<i>SwiRL</i>	93.25	81.31	90.83	84.42	57.76	79.52
<i>SwiRL+SP<sub>Res</sub></i>	93.17	81.08	90.76	84.52	59.24	79.86
<i>SwiRL+SP<sub>sim<sub>Jac</sub></sub></i>	93.37	80.30	90.86	84.43	59.54	79.83
<i>SwiRL+SP<sub>sim<sub>cos</sub></sub></i>	93.33	80.92	90.87	85.14	60.16	80.50
<i>SwiRL+SP<sub>sim<sub>Jac</sub><sup>th2</sup></sub></i>	93.03	82.75	90.95	85.62	59.63	80.75
<i>SwiRL+SP<sub>sim<sub>cos</sub><sup>th2</sup></sub></i>	93.78	80.56	91.23	84.95	61.01	80.48

- ⇒ Slight improvements, especially noticeable on Brown corpus
- ⇒ Weak signal of a single feature?

- Simple combinations of the individual *SwiRL*+ $SP_i$  classifiers worked quite well (**majority voting**)
- We also trained a **meta-classifier** to combine the *SwiRL*+ $SP_i$  classifiers and the stand-alone  $SP_i$  models:
  - ⇒ Binary classification approach:  
“is a proposed role correct or not?”
  - ⇒ Features are based on the predictions of base  $SP_i$  and *SwiRL*+ $SP_i$  models
  - ⇒ Trained with a SVM with a quadratic polynomial kernel

## Results (II)

	WSJ-test			Brown		
	Core	Adj	All	Core	Adj	All
<i>SwiRL</i>	93.25	81.31	90.83	84.42	57.76	79.52
+ $SP_{sim_{COS}^{th2}}$	93.78	80.56	91.23	84.95	61.01	80.48
Meta	<b>94.37</b>	<b>83.40</b>	<b>92.12</b>	<b>86.20</b>	<b>63.40</b>	<b>81.91</b>

- Statistically significant improvements (99%) for both core and adjunct arguments, both in domain and out of domain

## Results (II)

	WSJ-test			Brown		
	Core	Adj	All	Core	Adj	All
<i>SwiRL</i>	93.25	81.31	90.83	84.42	57.76	79.52
+ $SP_{sim_{COS}^{th2}}$	93.78	80.56	91.23	84.95	61.01	80.48
Meta	<b>94.37</b>	<b>83.40</b>	<b>92.12</b>	<b>86.20</b>	<b>63.40</b>	<b>81.91</b>

- Statistically significant improvements (99%) for both core and adjunct arguments, both in domain and out of domain

## Output analysis

- Manual inspection of 50 cases in which the meta classifier corrects SwiRL:
  - ⇒ Usually cases with low frequency verbs or argument heads
  - ⇒ In  $\sim 58\%$  of the cases, syntax does not disambiguate, seems to suggest a wrong role label or it is confusing SwiRL because it is incorrect. However, most of the SP predictions are correct.
  - ⇒  $\sim 30\%$  of the cases: unclear source of the SwiRL error but still several SP models suggest the correct role
  - ⇒  $\sim 12\%$  of the cases: chance effect

## Output analysis

- Manual inspection of 50 cases in which the meta classifier corrects SwiRL:
  - ⇒ Usually cases with low frequency verbs or argument heads
  - ⇒ In  $\sim 58\%$  of the cases, syntax does not disambiguate, seems to suggest a wrong role label or it is confusing SwiRL because it is incorrect. However, most of the SP predictions are correct.
  - ⇒  $\sim 30\%$  of the cases: unclear source of the SwiRL error but still several SP models suggest the correct role
  - ⇒  $\sim 12\%$  of the cases: chance effect

## SPs in a SRL System

(Zapirain et al., 2009; 2010)

## Output analysis: example 1

		Several	JJ	(S1(S(NP*
		traders	NNS	*)
		could	MD	(VP*
		be	VB	(VP*
		seen	VBN	(VP*
		shaking	VBG	(S(VP*
		their	PRP\$	(NP*
		heads	NNS	*))
		when	WRB	(SBAR(WHADVP*)
A1	A0	the	DT	(S(NP*
A1	A0	news	NN	*)
	(P)	flashed	VBD	(VP*))))))
		.	.	*)

## SPs in a SRL System

(Zapirain et al., 2009; 2010)

## Output analysis: example 2

		Italian	NNP	(S1(S(NP*
		President	NNP	*
		Francesco	NNP	*
		Cossiga	NNP	*)
	(P)	<b>promised</b>	VBD	(VP*
A2	A1	a	DT	(NP(NP*
A2	A1	quick	JJ	*
A2	A1	<b>investigation</b>	NN	*)
A2	A1	into	IN	(PP*
A2	A1	whether	IN	(SBAR*
A2	A1	Olivetti	NNP	(S(NP*)
A2	A1	broke	VBD	(VP*
A2	A1	Cocom	NNP	(NP*
A2	A1	rules	NNS	*)))))))
		.	.	*)

## Output analysis: example 3

		Annual	JJ	(S(NP*
		payments	NNS	*)
		will	MD	(VP*
		more	RBR	(VP(ADVP*
		than	IN	*)
	(P)	double	VB	*
A3	TMP	from	IN	(PP*
A3	TMP	a	DT	(NP*
A3	TMP	year	NN	*
A3	TMP	ago	RB	*))
		to	TO	(PP*
		about	RB	(NP(QP*
		\$240	CD	*
		million	CD	*))
		...		

## SPs in a SRL System

(Zapirain et al., 2009; 2010)

## Output analysis: example 4

		Procter	NNP	(S1(S(NP*
		&	CC	*
		Gamble	NNP	*
		Co.	NNP	*)
		plans	VBZ	(VP*
		to	TO	(S(VP*
		begin	VB	(VP*
	(P)	testing	VBG	(S(VP*
		next	JJ	(NP*
		month	NN	*))
A1	A0	a	DT	(NP(NP*
A1	A0	superco.	JJ	*
A1	A0	detergent	NN	*)
A1	A0	that	WDT	(SBAR(WHNP*)
		...		
A1	A0	washload	NN	(NP*))))))))))
		.	.	*)

# Talk Overview

- 1 Semantic Role Labeling
- 2 The Statistical Approach to SRL
- 3 Semantic Features for SRL
- 4 Conclusions**

# Summary

- i. SRL is an important problem in NLP, strongly related to applications requiring some degree of semantic interpretation
- ii. It is an active topic of research, which has generated an important body of work in the last 8 years  
⇒ techniques, resources, applications

Some news are good but...

- ⇒ SRL still has to resolve important problems before we see a spread usage in real open-domain applications
- ⇒ A jump is needed from the laboratory conditions to the real world.

# Summary

- i. SRL is an important problem in NLP, strongly related to applications requiring some degree of semantic interpretation
- ii. It is an active topic of research, which has generated an important body of work in the last 8 years  
⇒ techniques, resources, applications

Some news are good but...

- ⇒ SRL still has to resolve important problems before we see a spread usage in real open-domain applications
- ⇒ A jump is needed from the laboratory conditions to the real world.

# Remarks

- i. Generalization to new predicates/out-of-domain corpora is a weak point of statistical SRL systems
  - ⇒ **System portability** must be improved (e.g., domain adaptation, appropriate role sets, lexical semantic generalization, etc.)
  
- ii. **System complexity** is increasing at a higher scale than **performance**
  - ⇒ SRL systems have to be more efficient for massive text processing

# Remarks

- i. Generalization to new predicates/out-of-domain corpora is a weak point of statistical SRL systems
  - ⇒ **System portability** must be improved (e.g., domain adaptation, appropriate role sets, lexical semantic generalization, etc.)
  
- ii. **System complexity** is increasing at a higher scale than **performance**
  - ⇒ SRL systems have to be more efficient for massive text processing

# Remarks

- iii. SRL **Systems for languages other than English** should be developed and made available to the NLP community
- iv. **Reducing the cost of producing semantically annotated corpora** for under resourced languages (e.g., making use of semi-supervised learning, projections from corpora in other languages, etc.)

## Some Research Lines

- i. Unsupervised and semi-supervised approaches to SRL
- ii. SRL adapted to other languages types, genres, etc. (e.g., tweets, MT output, etc.)
- iii. SRL applied to assess semantic similarity (e.g., MT evaluation)

## Some Research Lines

- iv. Investigate learning architectures that take advantage of the joint resolution of several syntactic–semantic levels (esp. parsing–SRL, but also WSD, NEs, coreference, etc.)
  
- v. SemEval-2010 task #10: “Linking Events and their Participants in Discourse” (cross-sentence links between argument structures: null instantiation linking)

# Semantic Role Labeling. Generalizing Lexical Features using Selectional Preferences

**Lluís Màrquez**

TALP Research Center  
Technical University of Catalonia

ILCC/HCRC seminar series at the University of  
Edinburgh's School of Informatics  
Edinburgh, June 3, 2011