

# Relation Extraction

Feiyu Xu

[feiyu@dfki.de](mailto:feiyu@dfki.de)

Language Technology-Lab  
DFKI, Saarbrücken

# Outline

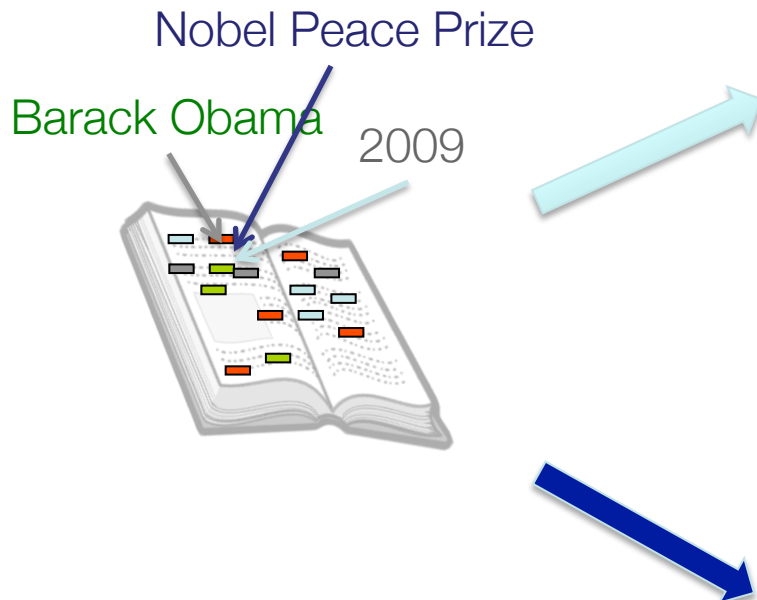
- Introduction to relation extraction
- Brief history of information extraction
- Machine learning approaches to relation extraction
  - State of the Art
  - DARE system (<http://dare.dfki.de>)
  - Web DARE: Distant Supervision
- Task-driven anaphora resolution for relation extraction
- References

# What is Relation Extraction



## ☆ Definition

- Given an unstructured text, relation extraction tool should be able to automatically recognize and extract relations among the relevant entities or concepts that are salient to the user's needs



- Walter Kohn, Nobel, Chemistry, 1998
- J.M. Coetzee, Nobel, Literature, 2003

### Patterns:

- *<prize>* be awarded to *<person>*
- *<person>* win *<prize>* in *<year>*
- .....

## General application task 1:



### ☆ Information access for information finder

mapping unstructured textual queries of users to more structured formal query for search and answer engines

The screenshot shows a Google search interface. At the top left is the Google logo. To its right is a search bar containing the text "Where is New York". Below the search bar, the word "Search" is displayed in red, followed by the text "About 4,460,000,000 results (0.28 seconds)". On the left side, there is a vertical menu with the following items: "Everything" (selected), "Images", "Maps", "Videos", "News", "Shopping", and "More". The main content area displays a map of New York City and its surrounding areas, including labels for Paterson, Yonkers, White Plains, Newark, Rochelle, Elizabeth, and Long Beach. To the right of the map, the text "New York, NY" is displayed in blue, with "maps.google.com" in green below it. At the bottom of the map, there is a copyright notice: "©2011 Google". Below the map, there is a list of links: "Hotels - Restaurants - Empire State Building - Top of the Rock Observation Deck - Metropolitan Museum of Art - Statue of Liberty - Rockefeller center - Nyc & Company". Below this list, there is a link to "New York City - Wikipedia, the free encyclopedia" with a magnifying glass icon, followed by the URL "en.wikipedia.org/wiki/New\_York\_City - Cached". At the very bottom, the text "New York is the most populous city in the United States and the center of the New York" is visible.

## General application task 2:



### ☆ Information acquisition for information provider

extract structured information from big amount free texts to construct knowledge database



**Relation  
Extraction**





# Relation in Information Extraction

## Information Extraction is ...

a technology that is futuristic from the user's point of view in the current information-driven world.

Rather than indicating which documents need to be read by a user, it extracts pieces of information that are salient to the user's needs ...

provided by NIST:

[[http://www-nlpir.nist.gov/related\\_projects/muc/](http://www-nlpir.nist.gov/related_projects/muc/)]



# Information Extraction: A Pragmatic Approach

- Identify the types of entities that are relevant to a particular task
- Identify the range of facts that one is interested in for those entities
- Ignore everything else

# Types of Information Extraction in LT

- Topic Extraction
- Term Extraction
- Named Entity Extraction
- Binary Relation Extraction
- N-ary Relation Extraction
- Event Extraction
- Answer Extraction
- Opinion Extraction
- Sentiment Extraction

**Relation Extraction** is the cover term for those IE tasks in which instances of semantic relations are detected in natural language texts.



# Types of Information Extraction in LT

- Topic Extraction
- Term Extraction
- Named Entity Extraction
- **Binary Relation Extraction**
- **N-ary Relation Extraction**
- **Event Extraction**
- **Answer Extraction**
- **Opinion Extraction**
- **Sentiment Extraction**

**Types of Relation Extraction**



# Extracting Job Openings from the Web:

## Semi-Structured Data

**foodscience.com-Job2**

JobTitle: Ice Cream Guru  
Employer: foodscience.com  
JobCategory: Travel/Hospitality  
JobFunction: Food Services  
JobLocation: Upper Midwest  
Contact Phone: 800-488-2611  
DateExtracted: January 8, 2001  
Source: www.foodscience.com/jobs\_midwest.html  
OtherCompanyJobs: foodscience.com-Job1

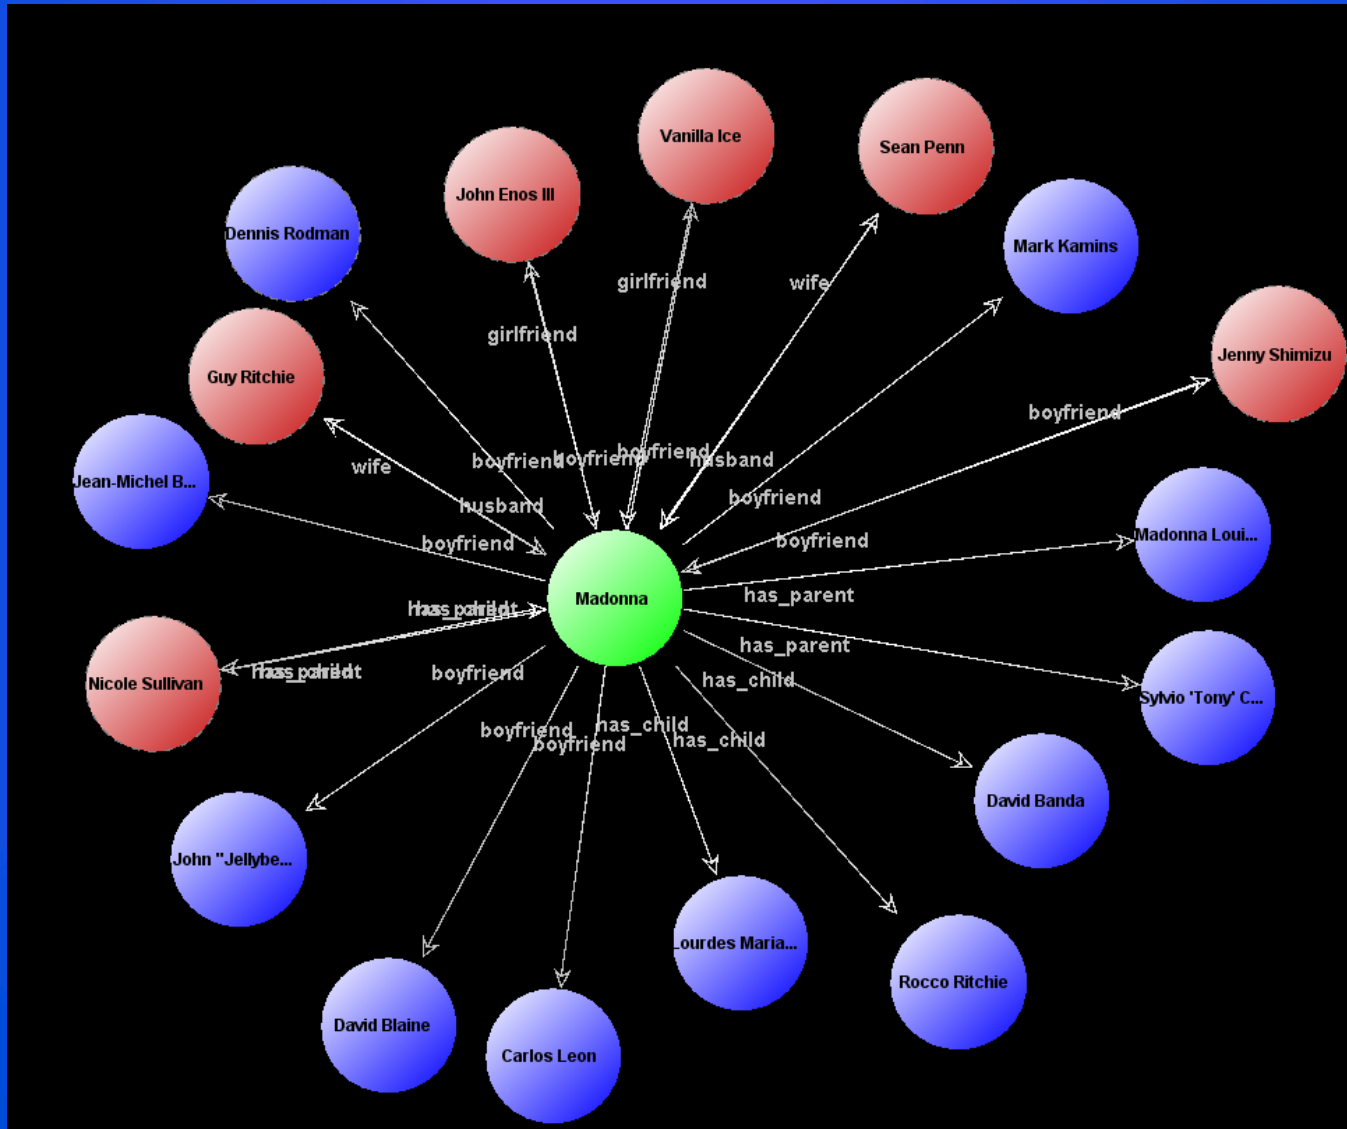
**Ice Cream Guru**

If you dream of cold creamy chocolate or coochy, coochy cookie, there's a great opportunity for you to maintain and expand this major corporation's high-end ice cream brand. Will be based in the Upper Midwest for about a year. After that, California here I come! Requires a BS in Food Science or dairy, plus ice cream formulation experience. Will consider entry level with an MS and an internship.  
Contact Susan: e-mail  
1-800-488-2611



# Example of Binary Social Relations

## Social Network of "Madonna" (Depth = 1)



# A Brief History of IE

# Message Understanding Conferences [MUC-7 98]

- U.S. Government sponsored conferences with the intention to coordinate multiple research groups seeking to improve IE and IR technologies (since 1987)
- defined several generic types of information extraction tasks (MUC Competition)
- MUC 1-2 focused on automated analysis of military messages containing textual information
- MUC 3-7 focused on information extraction from newswire articles
  - terrorist events
  - international joint-ventures
  - management succession event

# Evaluation of IE systems in MUC

- Participants receive description of the scenario along with the annotated *training corpus* in order to adapt their systems to the new scenario (1 to 6 months)
- Participants receive new set of documents (*test corpus*) and use their systems to extract information from these documents and return the results to the conference organizer
- The results are compared to the manually filled set of templates (*answer key*)

# Evaluation of IE systems in MUC

- precision and recall measures were adopted from the information retrieval research community

$$recall = \frac{N_{correct}}{N_{key}} \quad precision = \frac{N_{correct}}{N_{correct} + N_{incorrect}}$$

$$F = \frac{(\beta^2 + 1) \times precision \times recall}{\beta^2 \times precision + recall}$$

- Sometimes an  $F$ -measure is used as a combined recall-precision score

# Development Steps within IE Communities

- from attempts to use the methods of full text understanding to shallow text processing;
- from pure knowledge-based hand-coded systems to (semi-) automatic systems using machine learning methods;
- from complex domain-dependent event extraction to standardized domain-independent elementary entity identification, simple semantic relation and event extraction.

# Components of an IE Semantic Model (1)

- Entities - Individuals in the world *that are mentioned in a text*
  - Simple entities: singular objects
  - Collective entities: sets of objects of the same type *where the set is explicitly mentioned in the text*
- Relations – Properties that hold of tuples of entities.
- Complex Relations – Relations that hold among entities and relations
- Attributes – one place relations are attributes or individual properties

## Components of an IE Semantic Model (2)

- Temporal points and intervals
- Relations may be timeless or bound to time intervals
- Events – A particular kind of simple or complex relation among entities involving a change in relation state at the end of a time interval.

# Relations in Time

- time-dependent attribute:  $\text{age}(x)$
- timeless two-place relation:  $\text{father}(x, y)$
- time-dependent two-place relation:  $\text{boss}(x, y)$

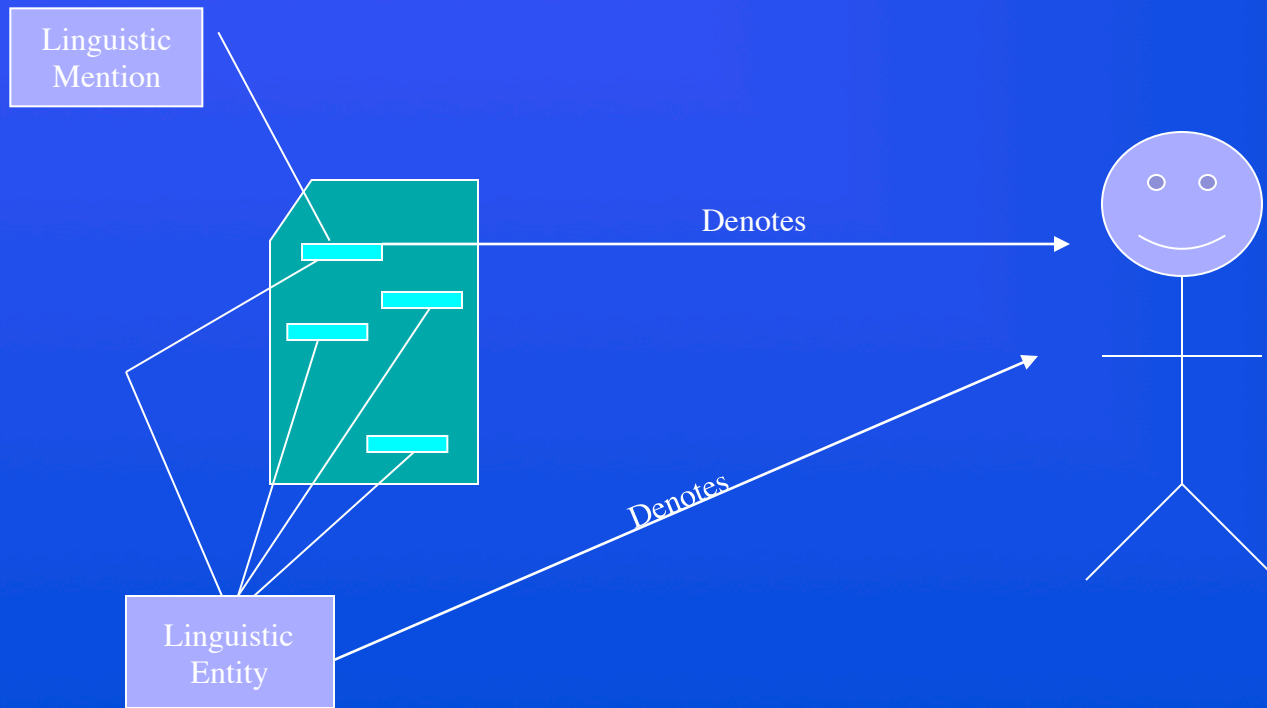
# IE as a Semantic Analysis: Relating Language to the Model

- Linguistic Mention
  - A particular linguistic phrase
  - Denotes a particular entity, relation, or event
    - A noun phrase, name, or possessive pronoun
    - A verb, nominalization, compound nominal, or other linguistic construct relating other linguistic mentions
- Linguistic Entity
  - Equivalence class of mentions with same meaning
    - Coreferring noun phrases
    - Relations and events derived from different mentions, but conveying the same meaning

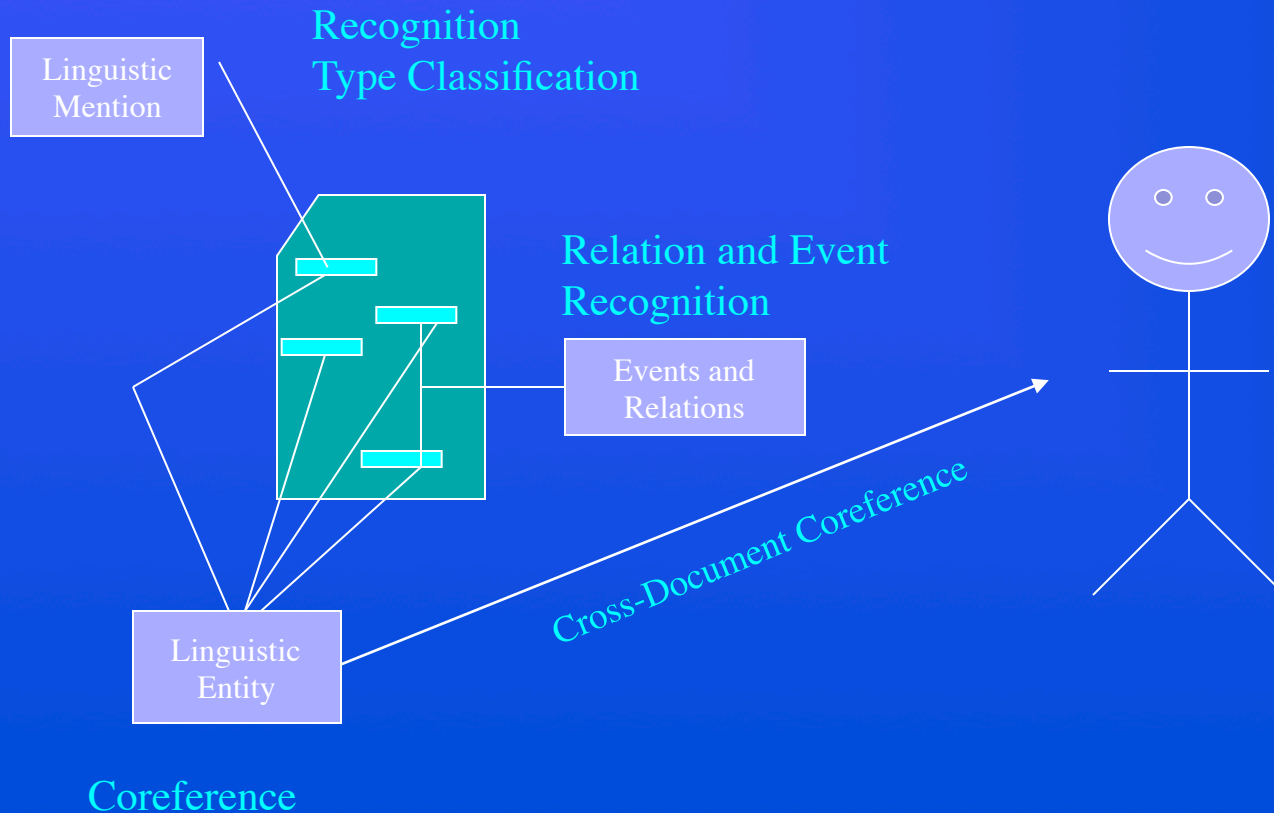
# The Basic Semantic Tasks of an IE System

- Recognition of linguistic mentions
- Classification of linguistic mentions into semantic types
- Identification of coreference equivalence classes of linguistic entities
- Identifying the actual individuals that are mentioned in an article
  - Associating linguistic entities with predefined individuals (e.g. a database, or knowledge base)
  - Forming equivalence classes of linguistic entities from different documents.

# Language and World Model



# NLP Tasks in an Extraction System



# Types of Linguistic Mentions

- Name mentions
  - The mention uses a proper name to refer to the entity
- Nominal mentions
  - The mention is a noun phrase whose head is a common noun
- Pronominal mentions
  - The mention is a headless noun phrase, or a noun phrase whose head is a pronoun, or a possessive pronoun

# Example of Linguistic Mentions

1. Three of the Nobel Prizes for Chemistry during the first decade **were awarded** for pioneering work in organic chemistry.
2. In **1902 Emil Fischer** (1852-1919), then in Berlin, **was given** the prize for **his** work on sugar and purine syntheses.
3. Another major influence from organic chemistry was the development of the chemical industry, and a chief contributor here was Fischer's teacher, **Adolf von Baeyer** (1835-1917) in Munich, **who was awarded** the prize in **1905**.

# Relation Extraction Example

A relation extraction task in the domain *management succession* (MUC-6)

< person\_in, person\_out, position, organisation >

- *person\_in*: the person who obtained the position
- *person\_out*: the person who left the position
- *position*: the job position that the two persons were involved in
- *organisation*: the organisation where the position was located

According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

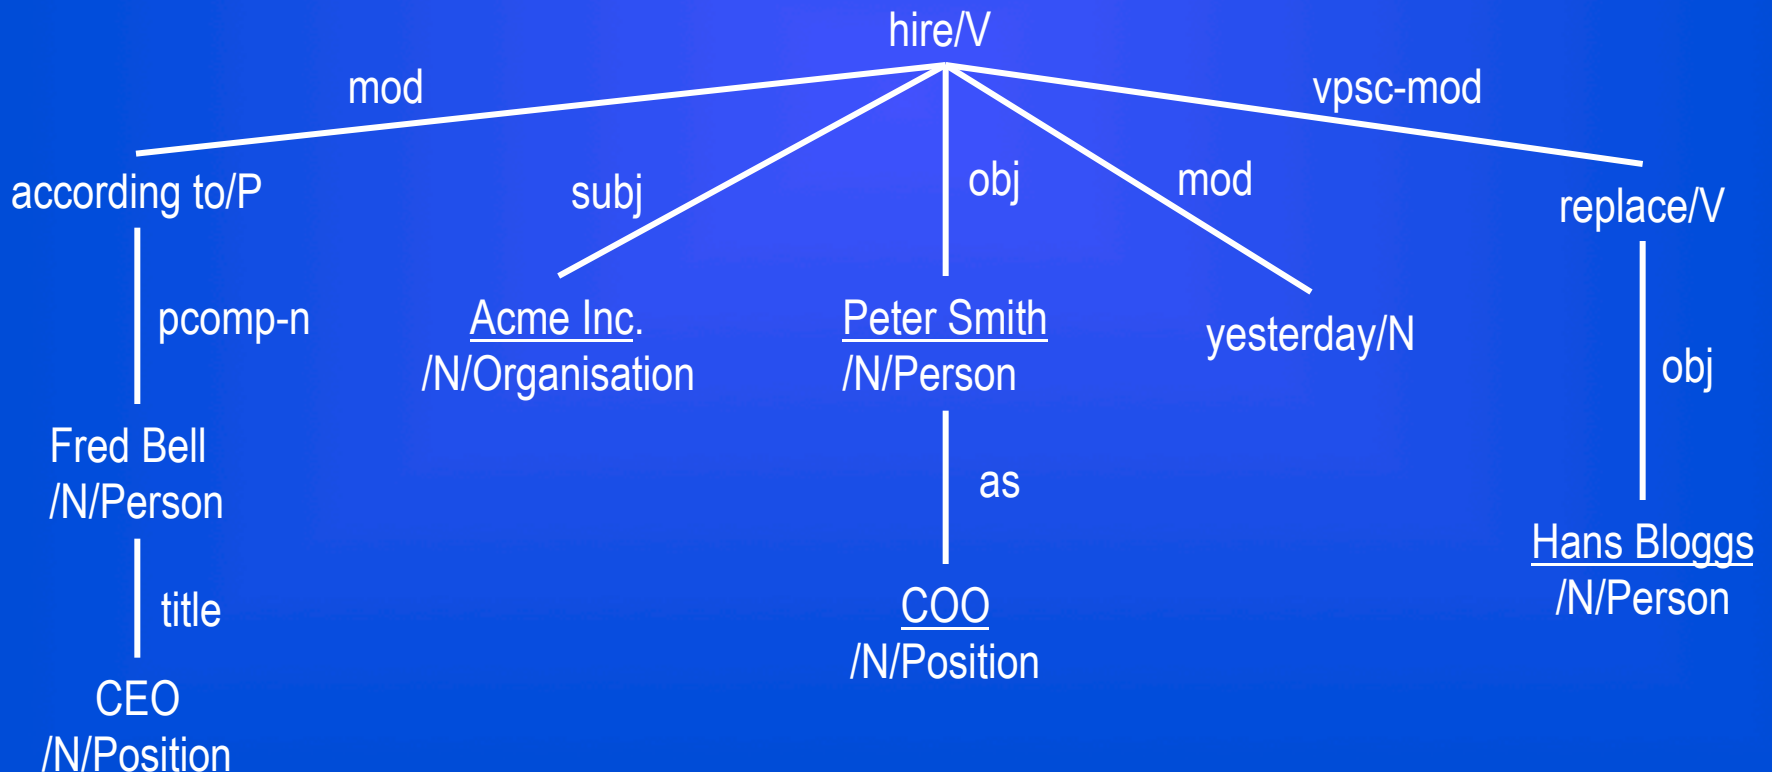
<person\_in, person\_out, position, organisation>

According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<person\_in, person\_out, position, organisation>

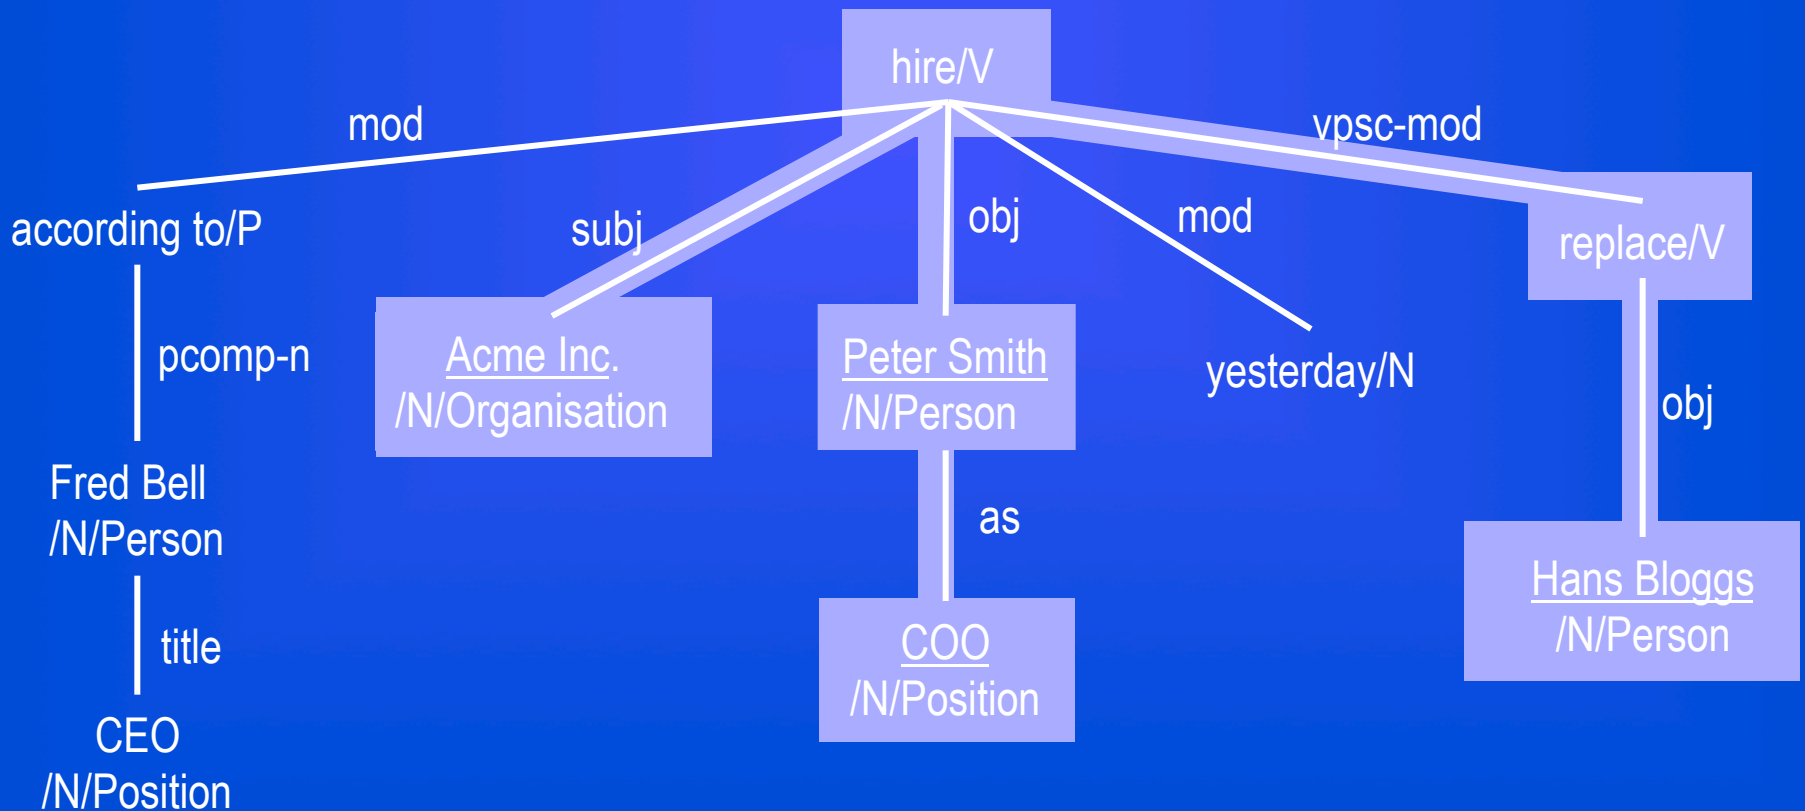
According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<person\_in, person\_out, position, organisation>




According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<person\_in, person\_out, position, organisation>



<person\_in, person\_out, position, organisation>



# Machine Learning for Relation Extraction

# Motivations of ML

- Porting to new domains or applications is expensive
- Current technology requires IE experts
  - Expertise difficult to find on the market
  - SME cannot afford IE experts
- Machine learning approaches
  - Domain portability is relatively straightforward
  - System expertise is not required for customization
  - “Data driven” rule acquisition ensures full coverage of examples

# Problems

- Training data may not exist, and may be very expensive to acquire
- Large volume of training data may be required
- Changes to specifications may require reannotation of large quantities of training data
- Understanding and control of a domain adaptive system is not always easy for non-experts

# Parameters of IE Real-World Tasks

- **Document structure**
  - Free text
  - Semi-structured
  - Structured
- **Linguistic annotation**
  - Shallow NLP
  - Deep NLP
- **Complexity and specificity of relation**
  - Unary
  - N-ary
- **Depth of extraction**
  - Recognition
  - Classification
  - Semantic role labelling

- **Degree of automation**
  - Semi-automatic
  - Supervised
  - Semi-Supervised
  - Minimally-Supervised
  - Distant Supervision
  - Unsupervised
- **Human interaction/contribution**
- **Data properties**
  - Domain relevance
  - Redundancy
  - Connectivity
- **Evaluation/validation**
  - With/without gold standard
  - Performance: recall & precision
  - Interaction among parameters


# State of the Art

# Binary Relation Only Approaches

- Extraction of binary relations only such as
  - author-book
  - company-location
- Do not employ the existing syntactic and semantic structures among  $n > 2$  arguments and rely on a later component to merge binary relations into complex relations
- Approaches
  - Ravichandran and Hovy, 2002
  - Pantel et al., 2004
  - Pasca et al., 2006a; Pasca et al., 2006b

According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<person\_in, person\_out, position, organisation>



# Surface-oriented Rule Representation

- Shallow linguistic analyses
- These formalisms are robust and efficient but only handle binary relations.
- Work best for relations whose arguments usually co-occur in close proximity within a sentence and whose mentions exhibit limited linguistic variation
- Approaches
  - Pasca et al., 2006a; Pasca et al., 2006b;
  - Kozareva et al., 2008; Hovy et al., 2009; Kozareva and Hovy, 2010

# Minimally Supervised (Bootstrapping)

- Based on iterative learning with limited initial knowledge
  - Start a small number of initial examples of relation instances (or patterns)
  - Label the free texts during iterations (e.g., Agichtein and Gravano, 2000; Yangarber et al., 2000; Ravichandran and Hovy, 2002; Stevenson and Greenwood, 2005).
- Often suffer from semantic drift or the propagation of errors occurring during iterations
- Performance depends on data properties

# Distant Supervision

- A massive seed-based, one step version of bootstrapping
- Rely on a large amount of trustworthy facts
- Their performance does not hinge on corpus data properties such as redundancy
- Approaches
  - (Mintz et al., 2009)
  - Others: (Wu and Weld, 2007); (Wu et al., 2008); (Weld et al., 2008); (Hoffmann et al., 2010); (Xu et al., 2011); (Nguyen and Moschitti, 2011)

# Open IE

- They do not target given relations.
- They are very useful for applications continuously faced with new relation or event types, e.g., online social media monitoring
- However, the results of these systems cannot be directly taken for filling knowledge databases, because the semantics of the new relations including the roles of the entities remains unknown.
- Example: TextRunner (Banko et al., 2007; Yates et al., 2007)

# Reality in IE Projects



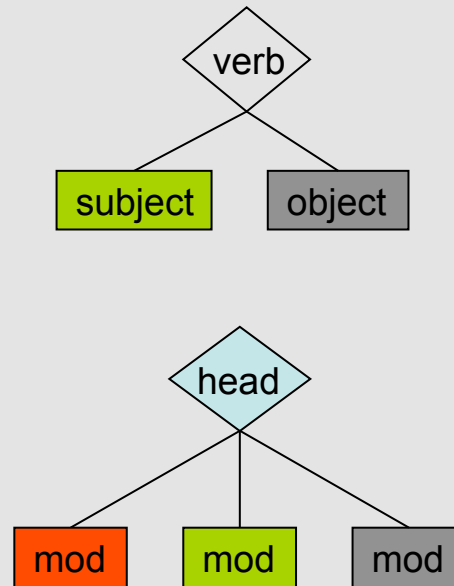


- Our IE users are often not domain experts
- IE experts have to develop methods and strategies for
  - Prospecting a domain
  - Proposing relevant relations
  - Finding relevant and suitable data

# DARE:

*Minimally Supervised Machine Learning  
of  
Relation Extraction Grammars  
for  
Relations of Various Complexity*

# Research Goal

Development of a general framework for automatically learning mappings between linguistic analyses and target semantic relations, with minimal human intervention.



# Challenges

- Easy adaptation to new relation types with varied complexity
- Automatic learning without annotated corpus
- Exhaustive discovery of relevant linguistic patterns
- Integration of semantic role information into linguistic patterns

# Example

A relation extraction task in the domain *management succession* (MUC-6)

< person\_in, person\_out, position, organisation >

- *person\_in*: the person who obtained the position
- *person\_out*: the person who left the position
- *position*: the job position that the two persons were involved in
- *organisation*: the organisation where the position was located

According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

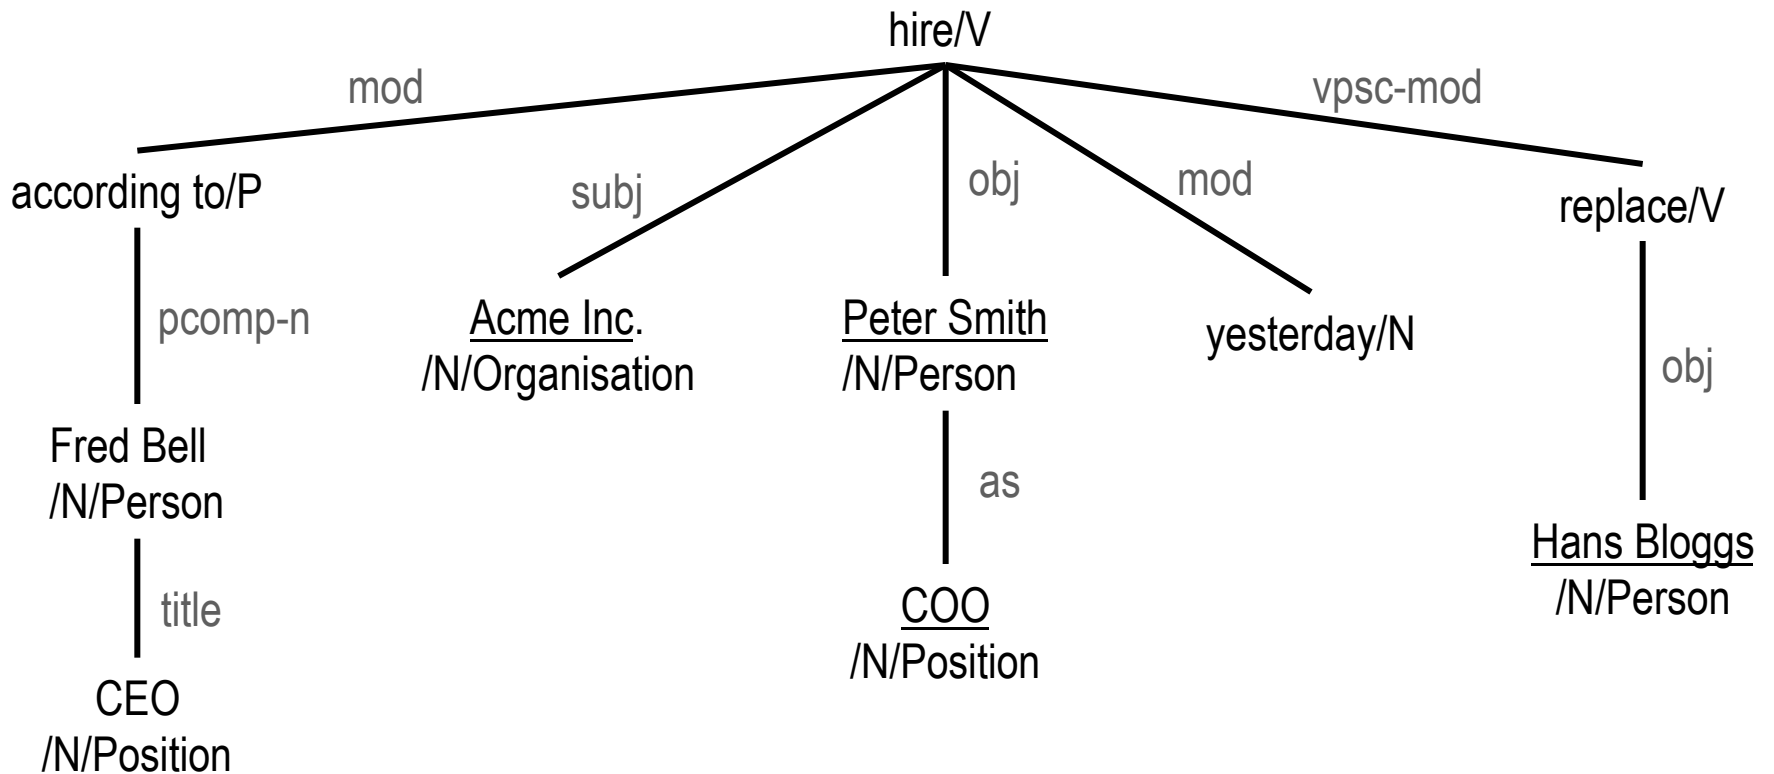
<person\_in, person\_out, position, organisation>

According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<person\_in, person\_out, position, organisation>

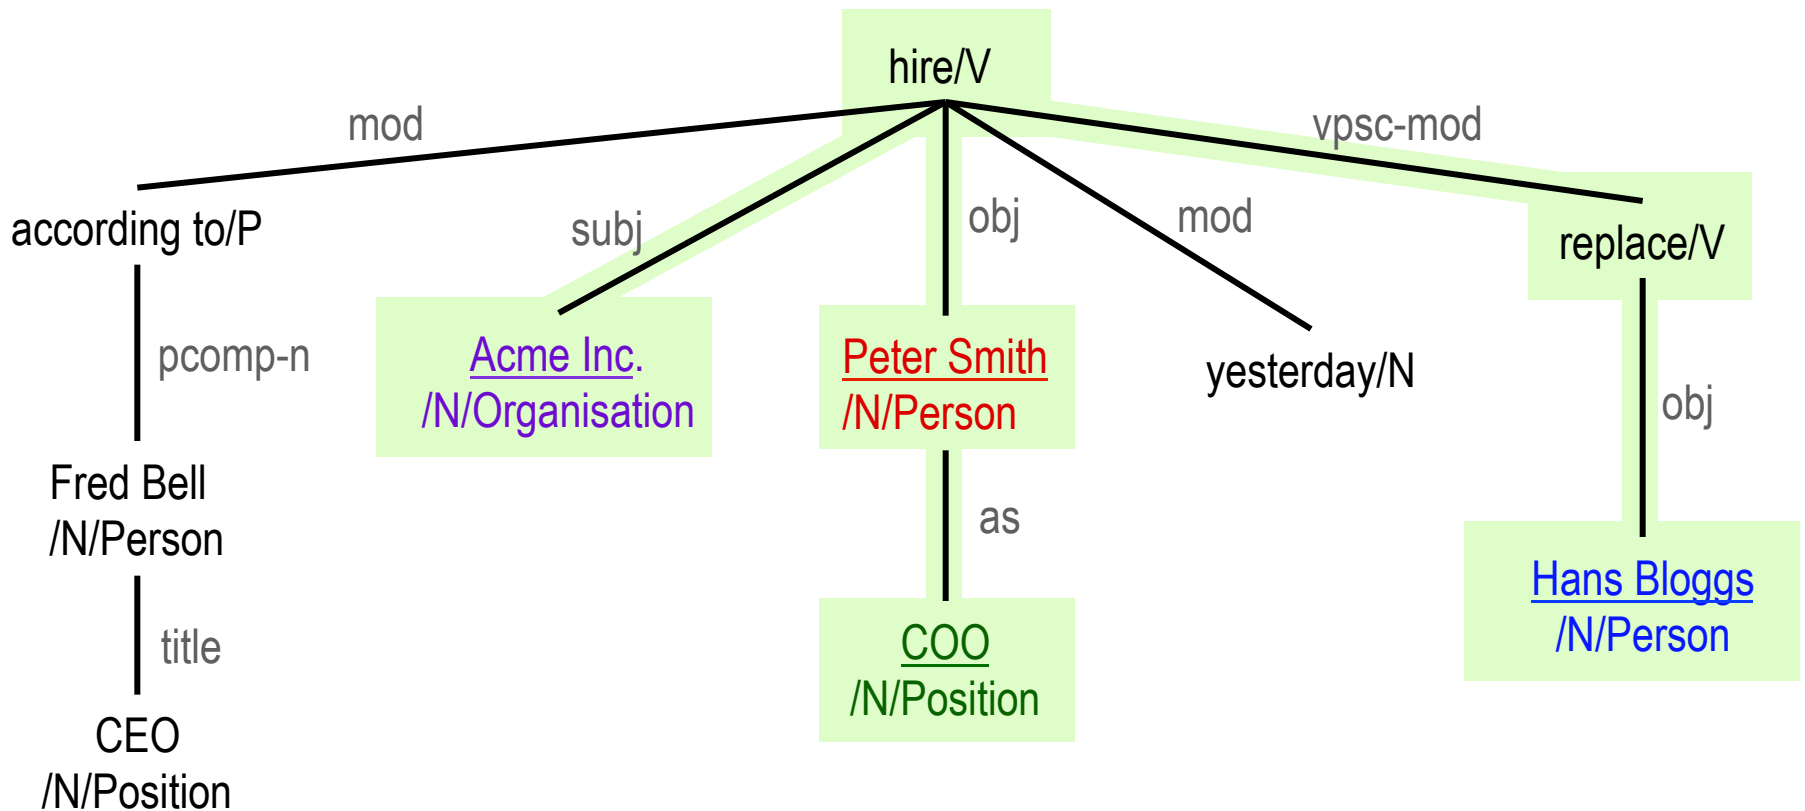
According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<person\_in, person\_out, position, organisation>

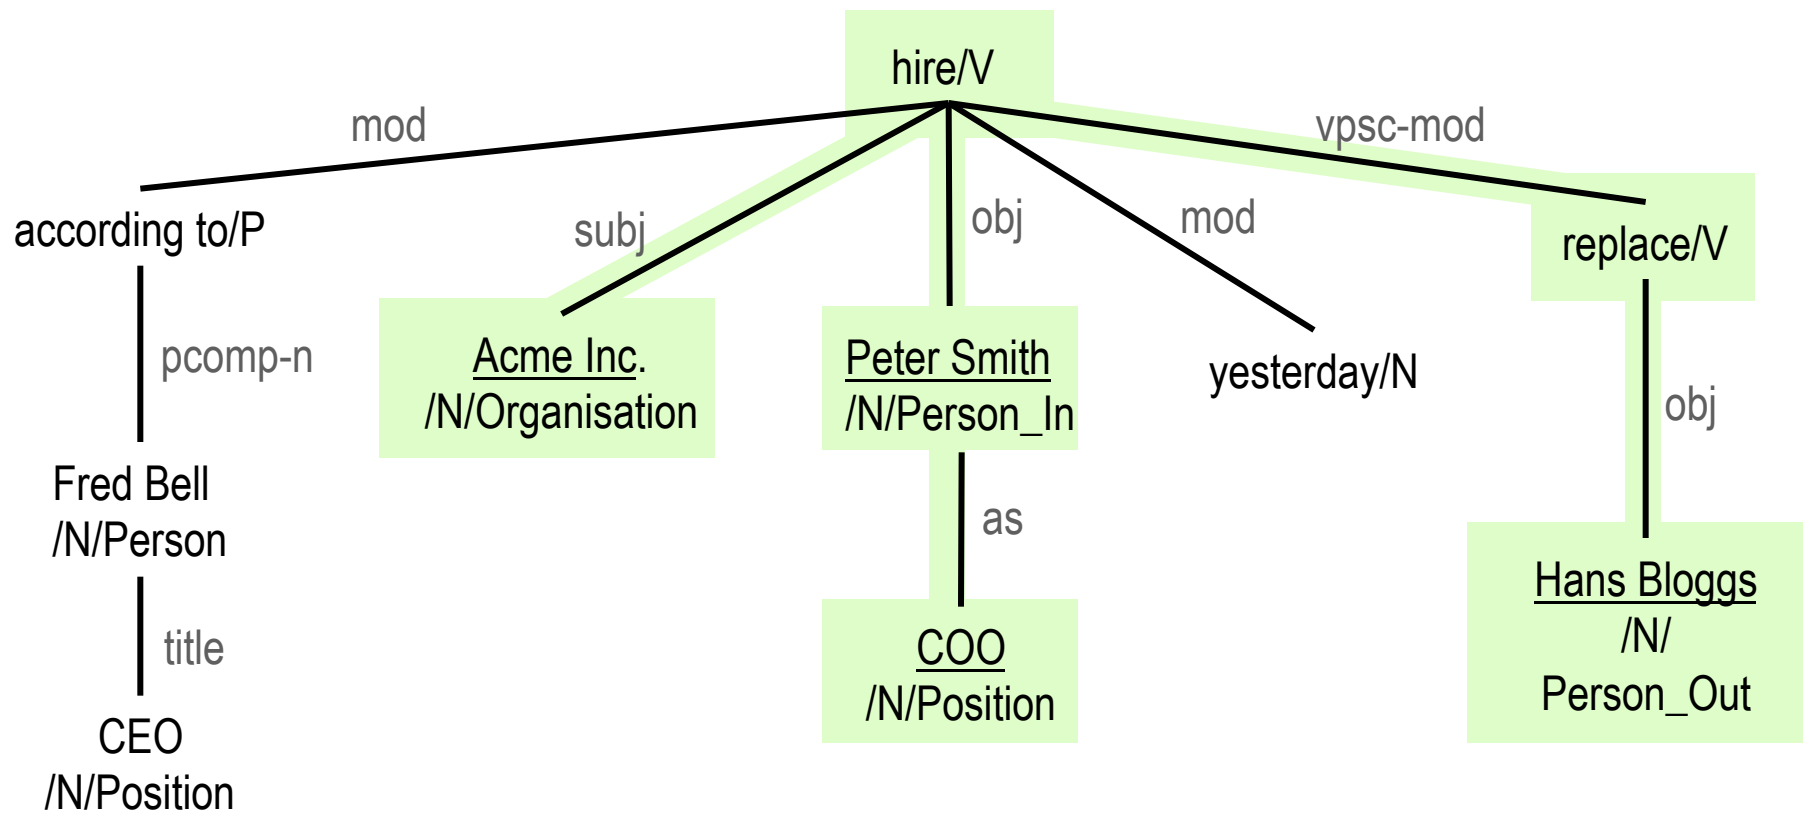


According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<person\_in, person\_out, position, organisation>



# Ideal Target Pattern



# DARE: *Domain Adaptive Relation Extraction*

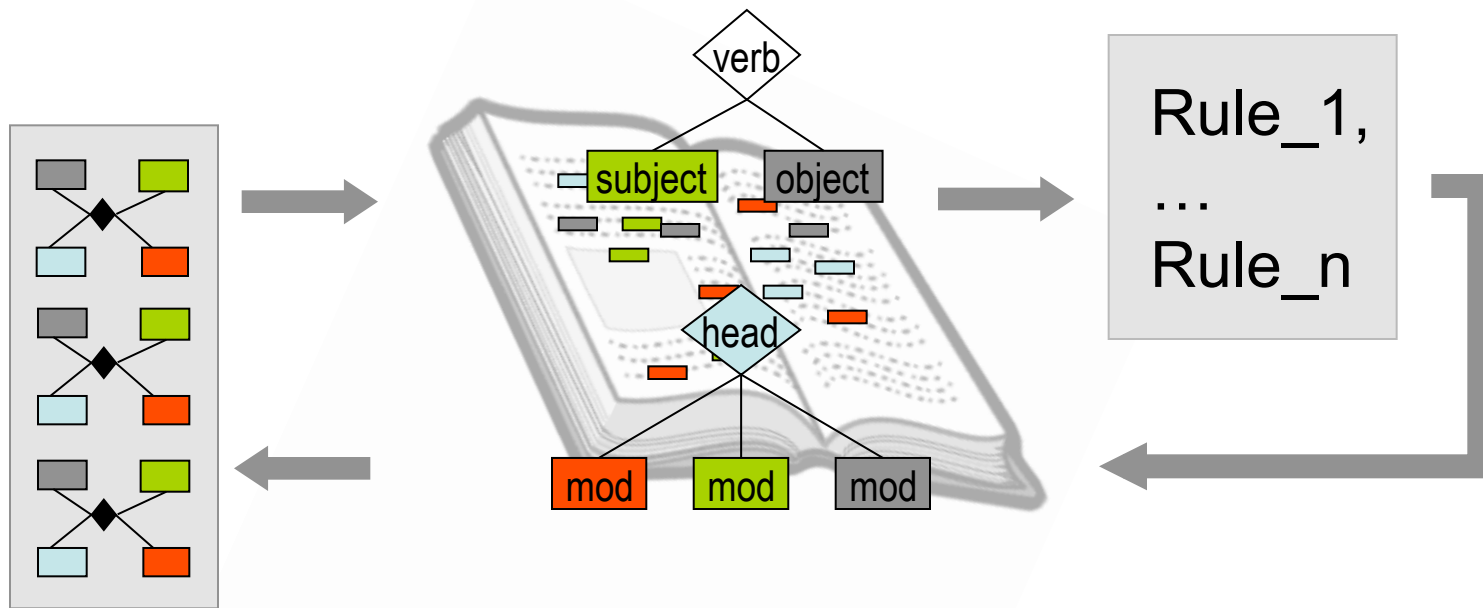
- Samples of target relation instances serve as semantic seed
- Systematic treatment of n-ary relations and their projections
- Exploitation of relation projections for pattern discovery
- Bottom-up compositional pattern discovery
- A recursive linguistic rule representation
- Rules contain semantic roles w.r.t. to target relation
- Bottom-up compression method to generalize rules

# Bootstrapping Relation Extraction with Semantic Seed

Adapted from

DIPRE (Brin, 1998) and Snowball (Agichtein & Gravano, 2000)

but extended and enriched with linguistic analysis



# Bootstrapping Relation Extraction with Semantic Seed

- DIPRE and Snowball


- binary relations only, no projections, no linguistic analysis

- DARE


- n-ary relations and their projections, deep linguistic analysis

(in the experiments I use MINIPAR by Dekan Lin 1999)


# Start of Bootstrapping (simplified)



# DARE Architecture

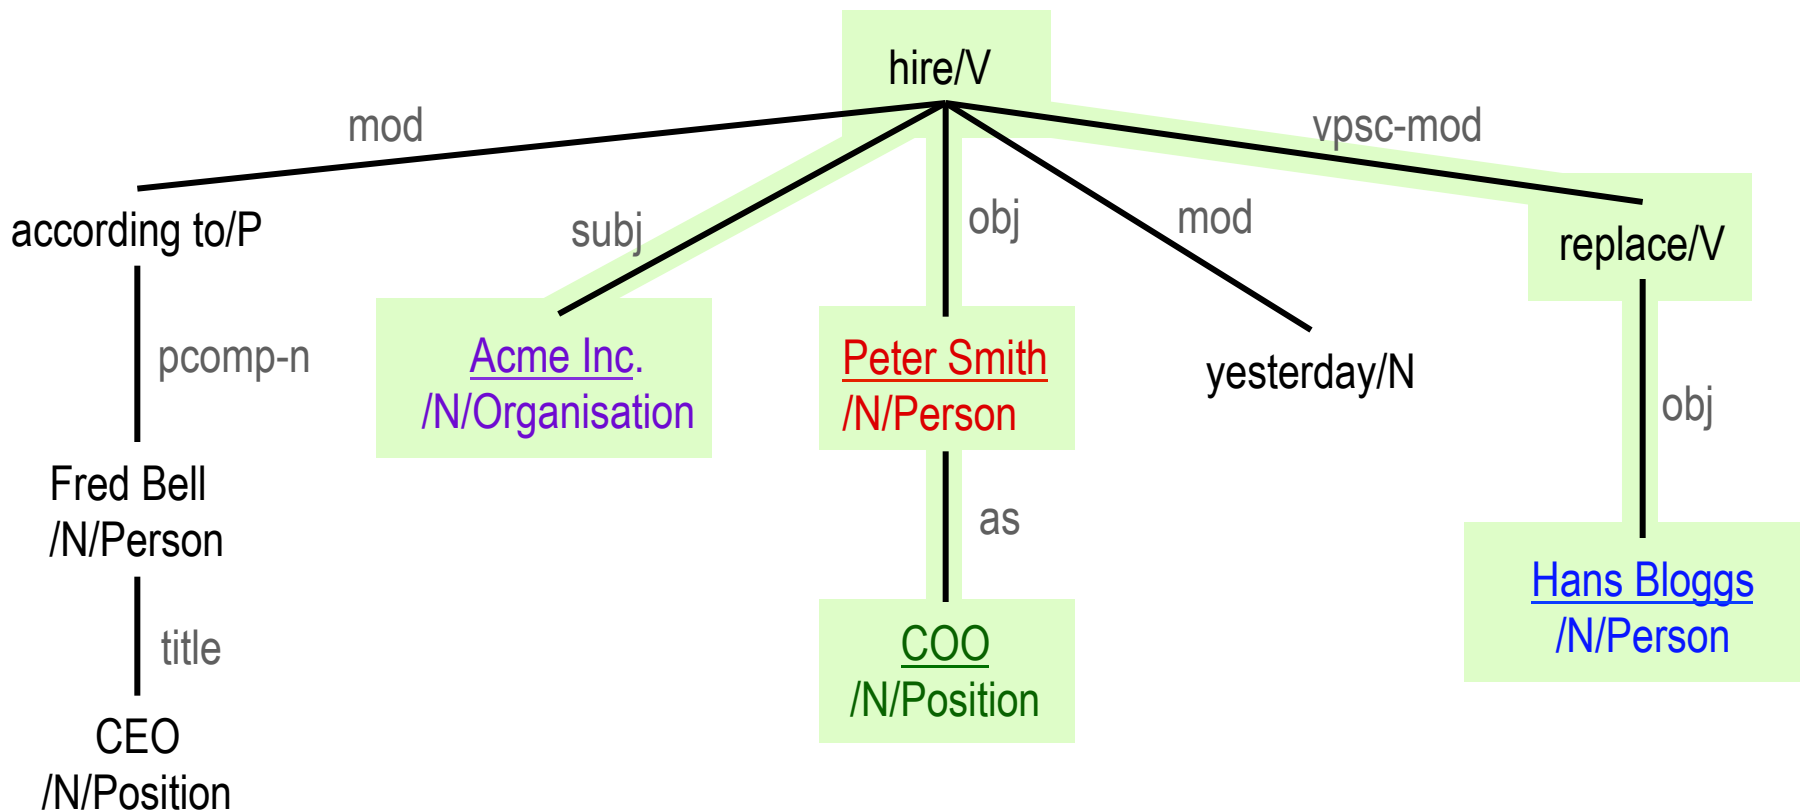


# DARE Architecture



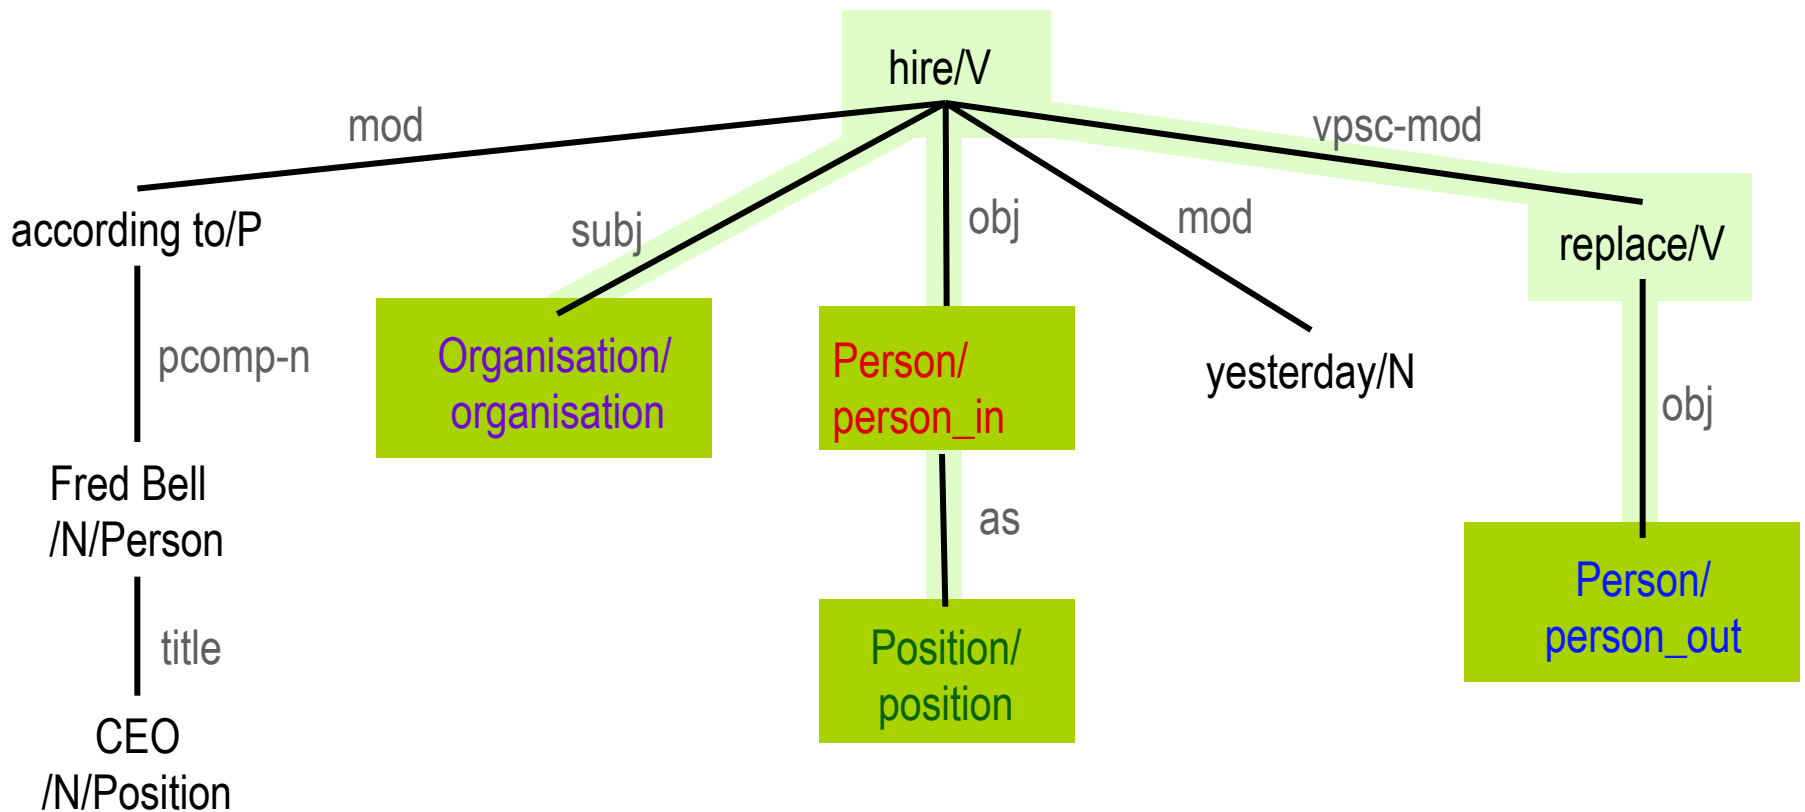
According to CEO Fred Bell, [Acme Inc.](#) hired [Peter Smith](#) as [COO](#) yesterday, replacing [Hans Bloggs](#).

<[Peter Smith](#)/person\_in, [Hans Bloggs](#)/person\_out, [COO](#) /position, [Acme Inc.](#) /organisation>



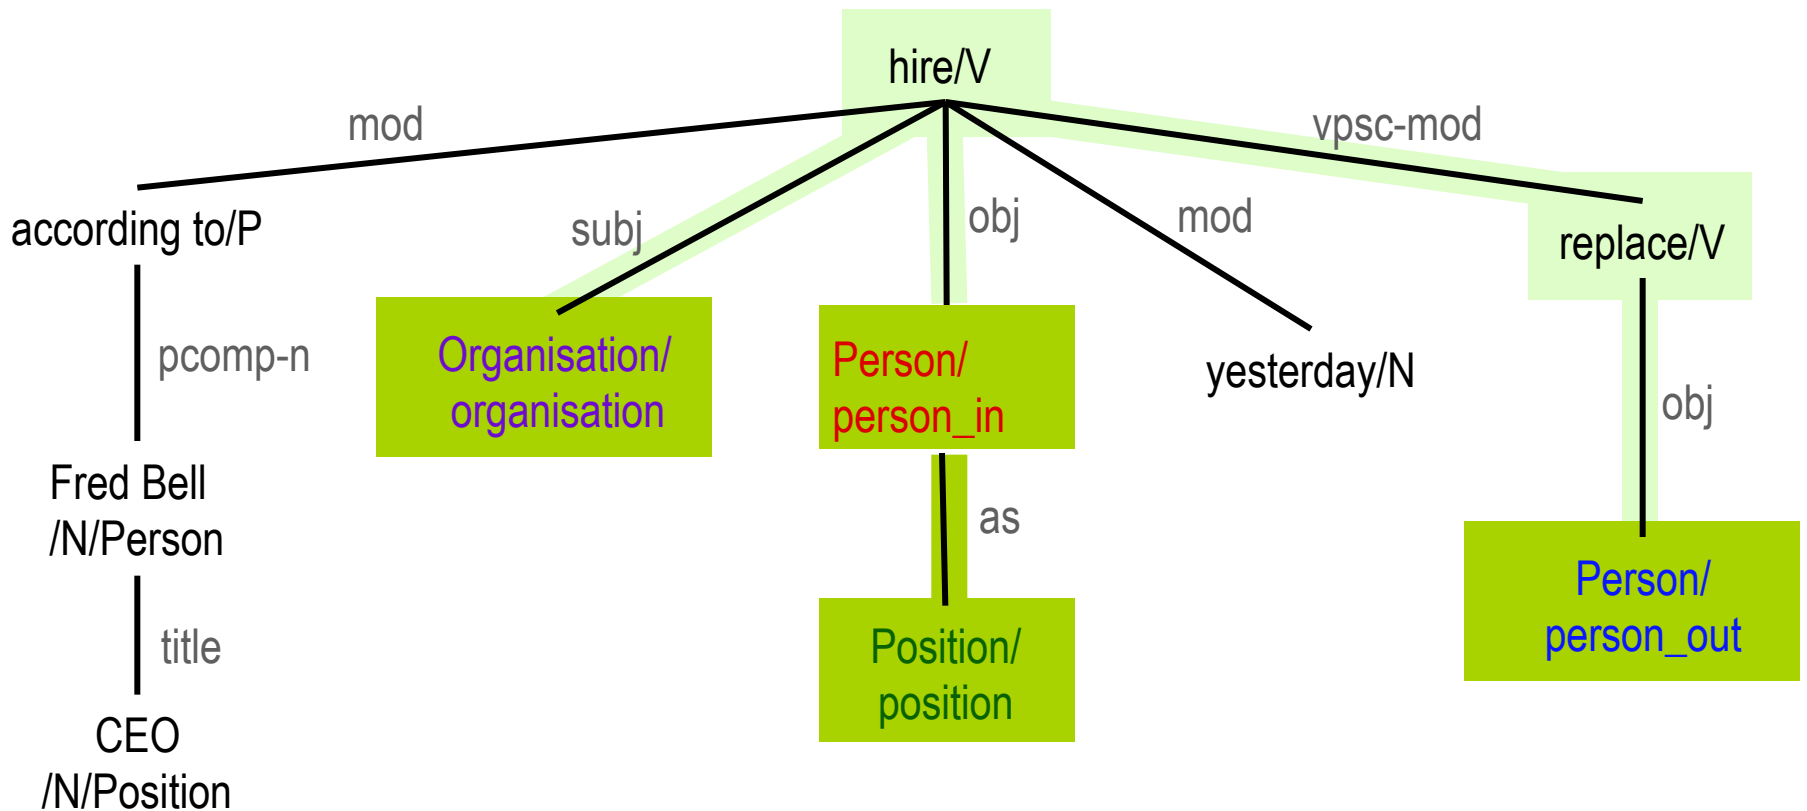
According to CEO Fred Bell, [Acme Inc.](#) hired [Peter Smith](#) as [COO](#) yesterday, replacing [Hans Bloggs](#).

<[Peter Smith](#)/person\_in, [Hans Bloggs](#)/person\_out, [COO](#) /position, [Acme Inc.](#) /organisation>



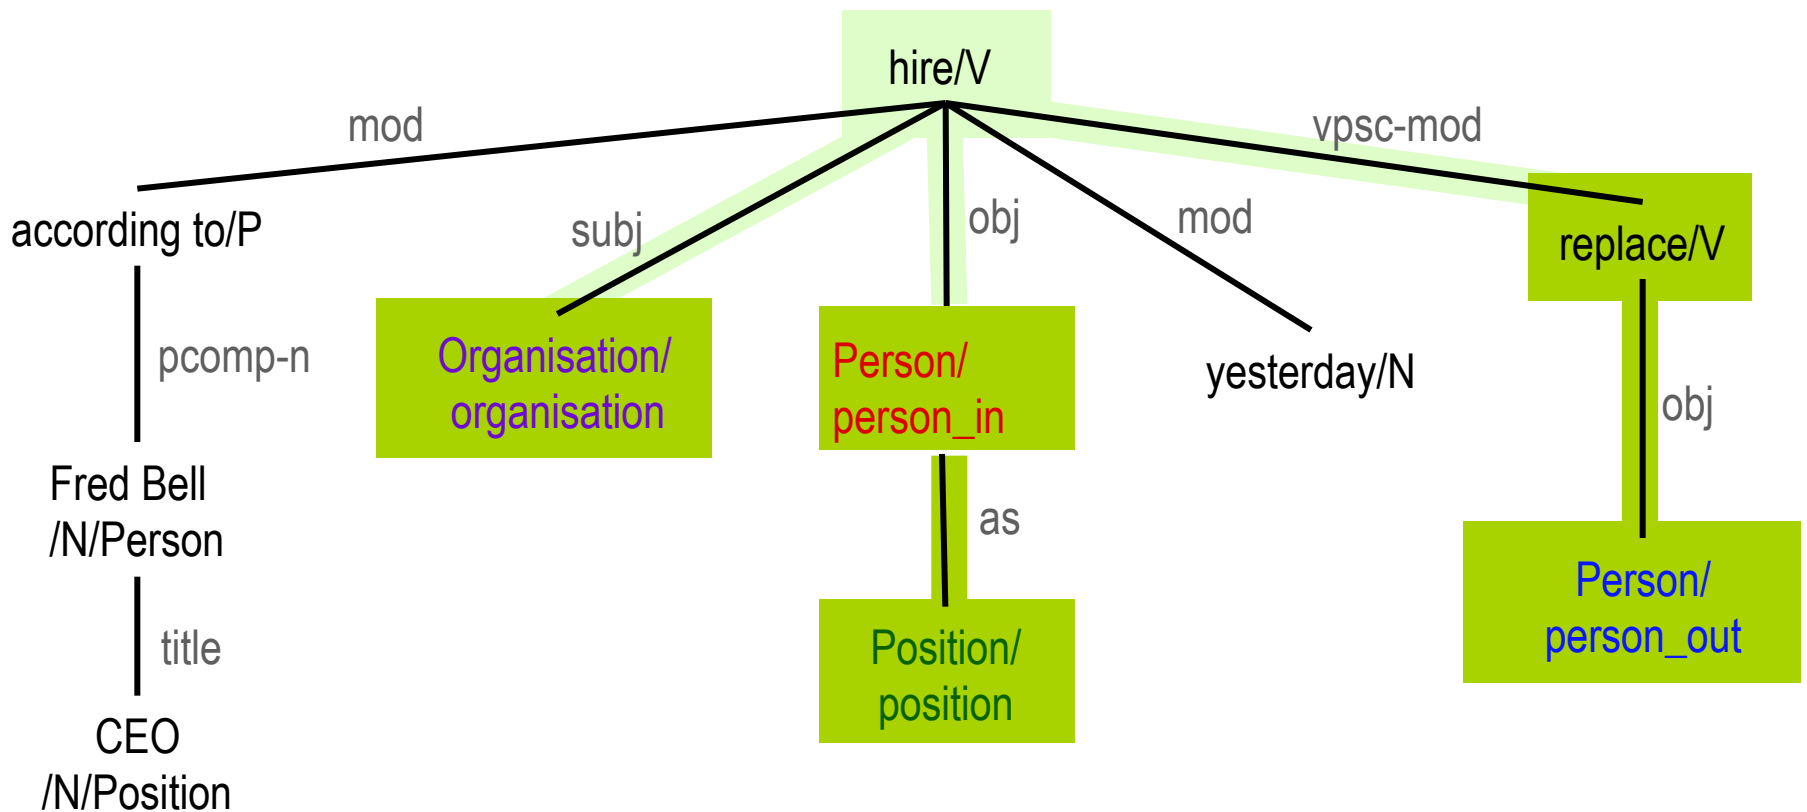
According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<Peter Smith/person\_in, Hans Bloggs/person\_out, COO /position, Acme Inc. /organisation>



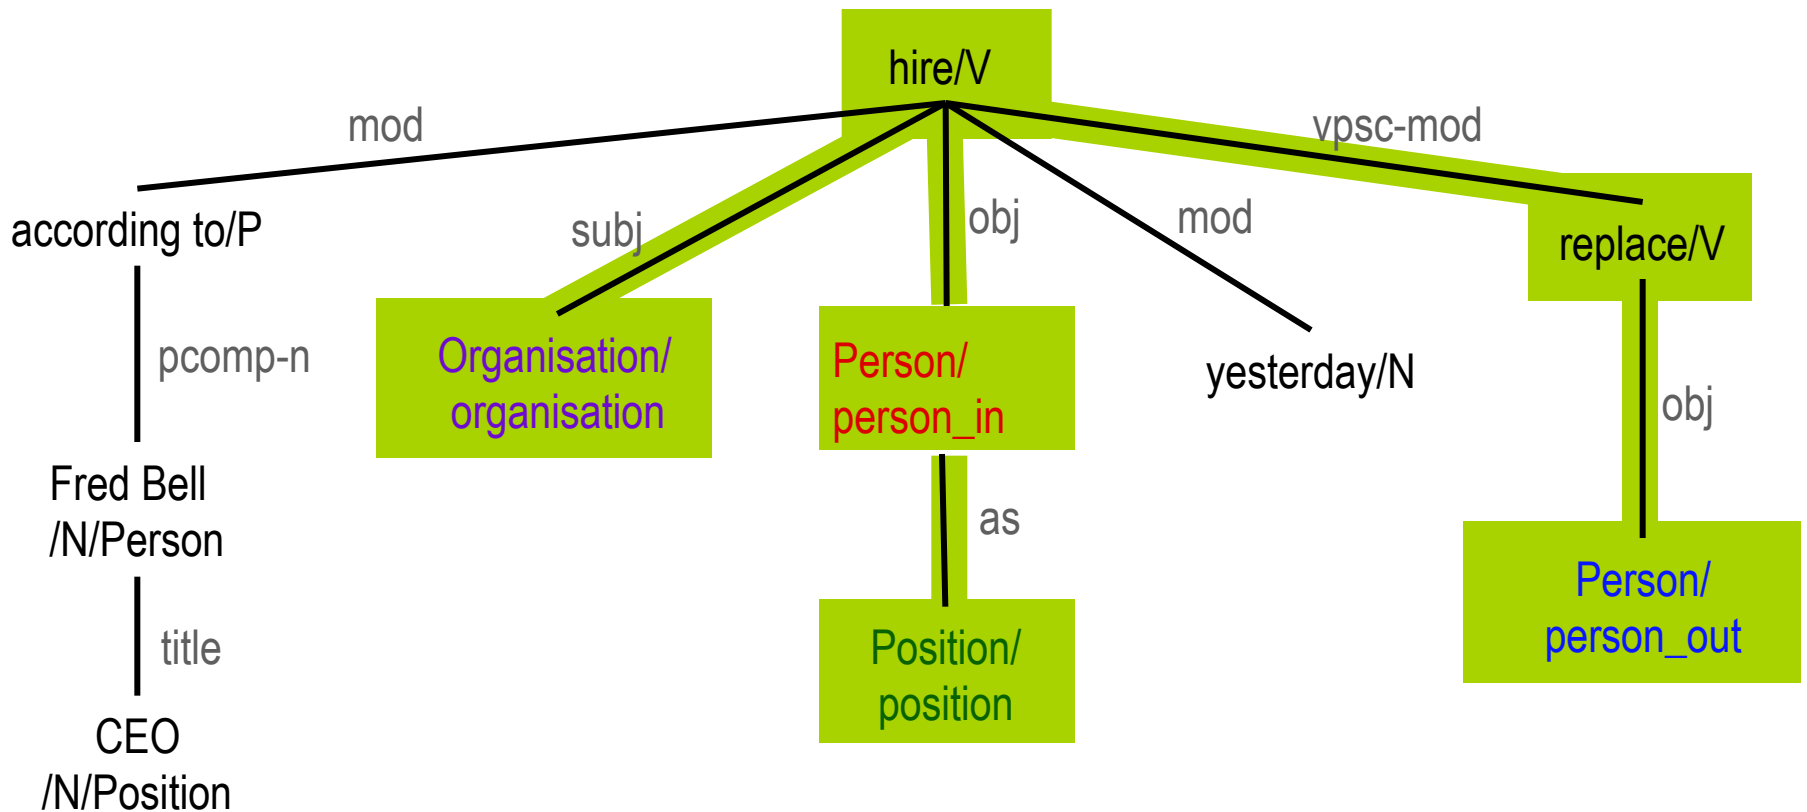
According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.


<Peter Smith/person\_in, Hans Bloggs/person\_out, COO /position, Acme Inc. /organisation>




According to CEO Fred Bell, Acme Inc. hired Peter Smith as COO yesterday, replacing Hans Bloggs.

<Peter Smith/person\_in, Hans Bloggs/person\_out, COO /position, Acme Inc. /organisation>




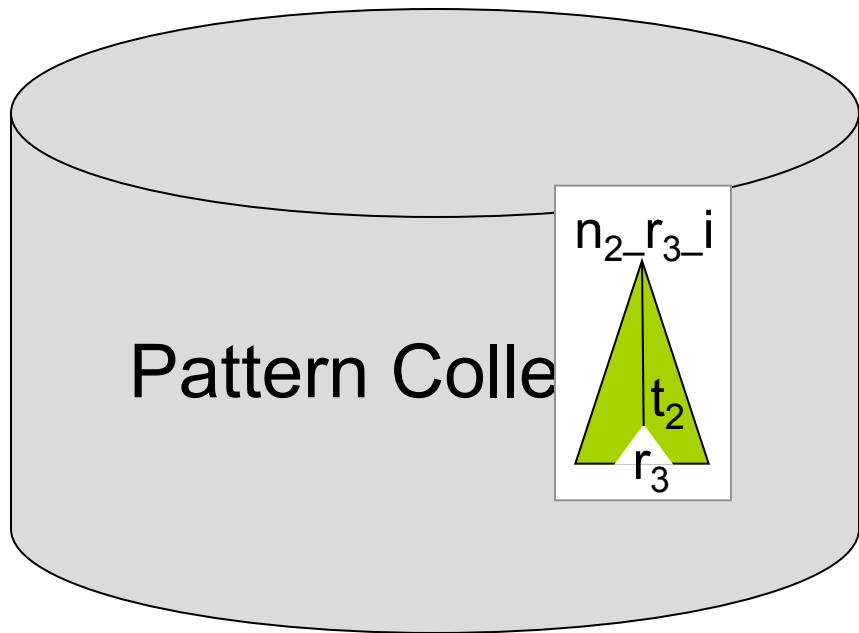



0. replace all nodes that are instantiated with the seed arguments by new nodes. Label these new nodes with the seed argument roles and their entity classes;



0. replace all nodes that are instantiated with the seed arguments by new nodes. Label these new nodes with the seed argument roles and their entity classes;






0. replace all nodes that are instantiated with the seed arguments by new nodes. Label these new nodes with the seed argument roles and their entity classes;

for  $i=1$  to  $n$

1. identify the set of the lowest non-terminal nodes  $N_1$  in  $t$  that dominate  $i$  arguments (possibly among other nodes).

2. substitute  $N_1$  by nodes labelled with the seed argument roles and their entity classes

3. prune the subtrees dominated by  $N_1$  from  $t$  and add these subtrees into the pattern collection. These subtrees are assigned the argument role information and a unique id.




0. replace all nodes that are instantiated with the seed arguments by new nodes. Label these new nodes with the seed argument roles and their entity classes;


for  $i=1$  to  $n$

1. identify the set of the lowest non-terminal nodes  $N_1$  in  $t$  that dominate  $i$  arguments (possibly among other nodes).

2. substitute  $N_1$  by nodes labelled with the seed argument roles and their entity classes

3. prune the subtrees dominated by  $N_1$  from  $t$  and add these subtrees into the pattern collection. These subtrees are assigned the argument role information and a unique id.






0. replace all nodes that are instantiated with the seed arguments by new nodes. Label these new nodes with the seed argument roles and their entity classes;


for  $i=1$  to  $n$

1. identify the set of the lowest non-terminal nodes  $N_1$  in  $t$  that dominate  $i$  arguments (possibly among other nodes).

2. substitute  $N_1$  by nodes labelled with the seed argument roles and their entity classes

3. prune the subtrees dominated by  $N_1$  from  $t$  and add these subtrees into the pattern collection. These subtrees are assigned the argument role information and a unique id.





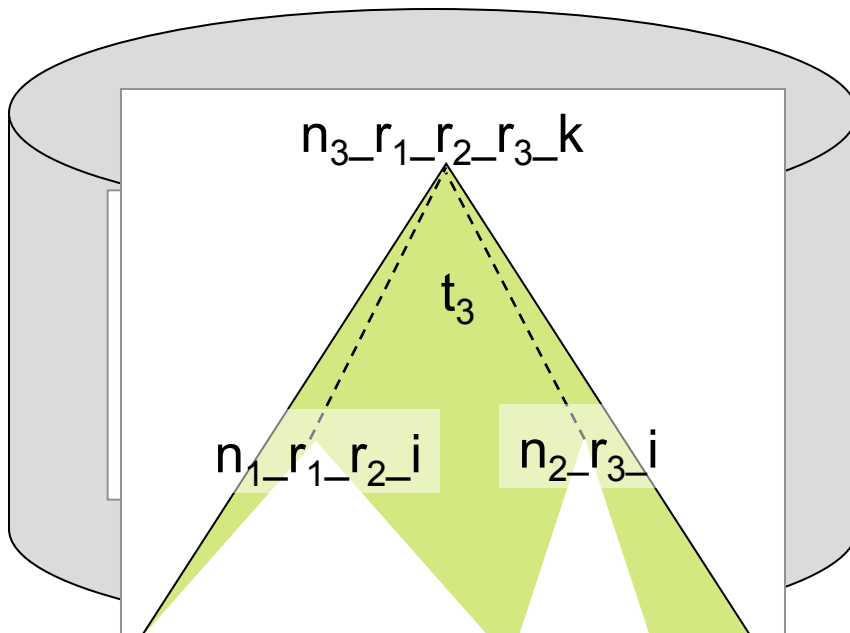
0. replace all nodes that are instantiated with the seed arguments by new nodes. Label these new nodes with the seed argument roles and their entity classes;


for  $i=1$  to  $n$

1. identify the set of the lowest non-terminal nodes  $N_1$  in  $t$  that dominate  $i$  arguments (possibly among other nodes).

2. substitute  $N_1$  by nodes labelled with the seed argument roles and their entity classes

3. prune the subtrees dominated by  $N_1$  from  $t$  and add these subtrees into the pattern collection. These subtrees are assigned the argument role information and a unique id.






0. replace all nodes that are instantiated with the seed arguments by new nodes. Label these new nodes with the seed argument roles and their entity classes;

for  $i=1$  to  $n$

1. identify the set of the lowest non-terminal nodes  $N_1$  in  $t$  that dominate  $i$  arguments (possibly among other nodes).

2. substitute  $N_1$  by nodes labelled with the seed argument roles and their entity classes

3. prune the subtrees dominated by  $N_1$  from  $t$  and add these subtrees into the pattern collection. These subtrees are assigned the argument role information and a unique id.



# Example in Nobel Prize Award Domain

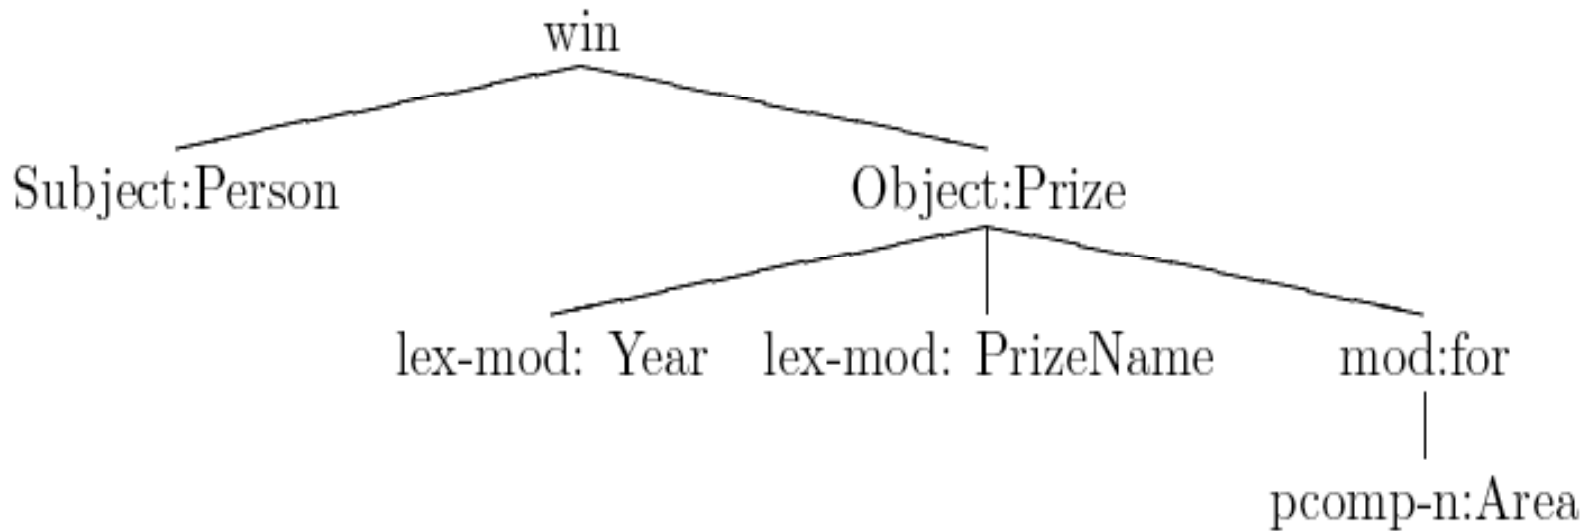
- Seed example

<Mohamed ElBaradei, Nobel, Peace, 2005>

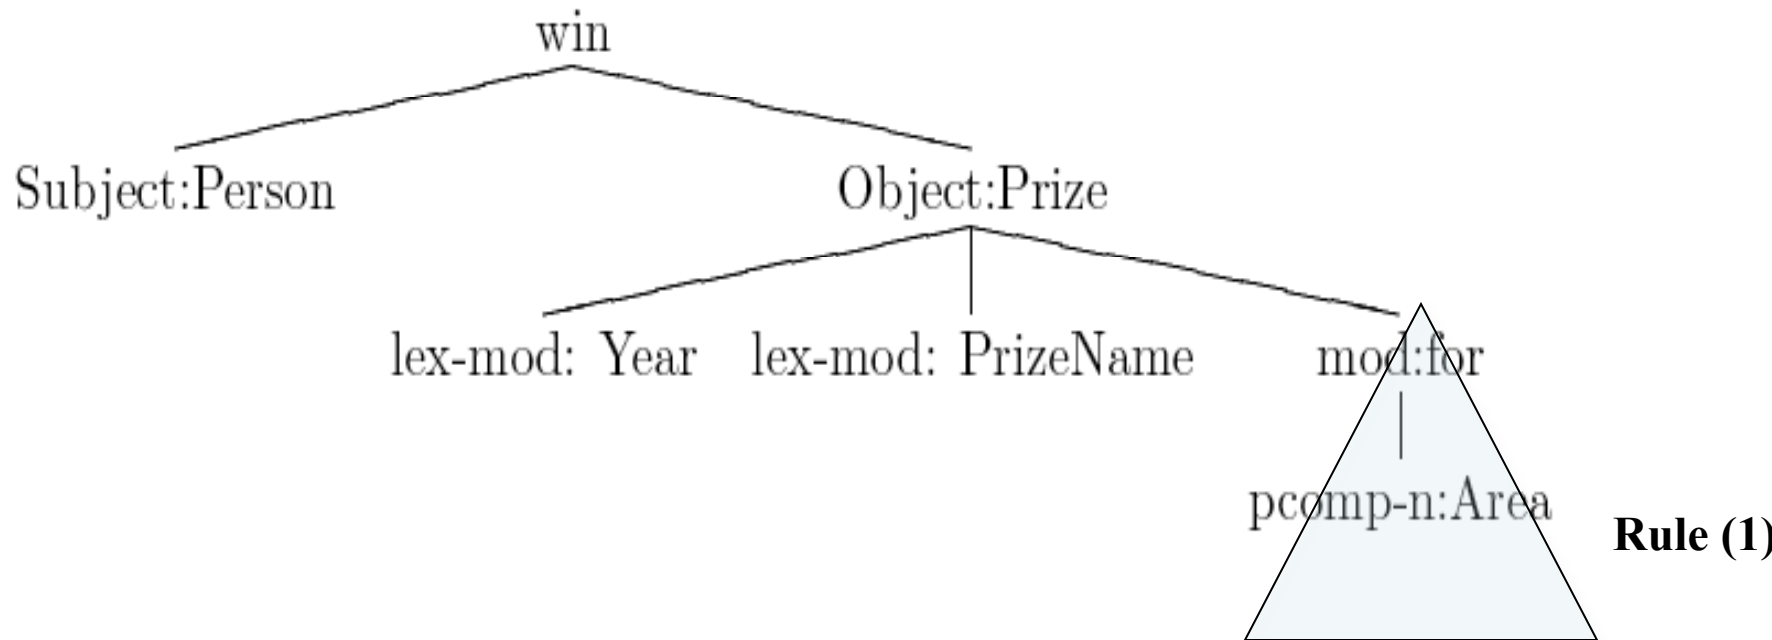
- Sentence matched with the seed

Mohamed ElBaradei, won the 2005 Nobel Prize for Peace on Friday for his efforts to limit the spread of atomic weapons.

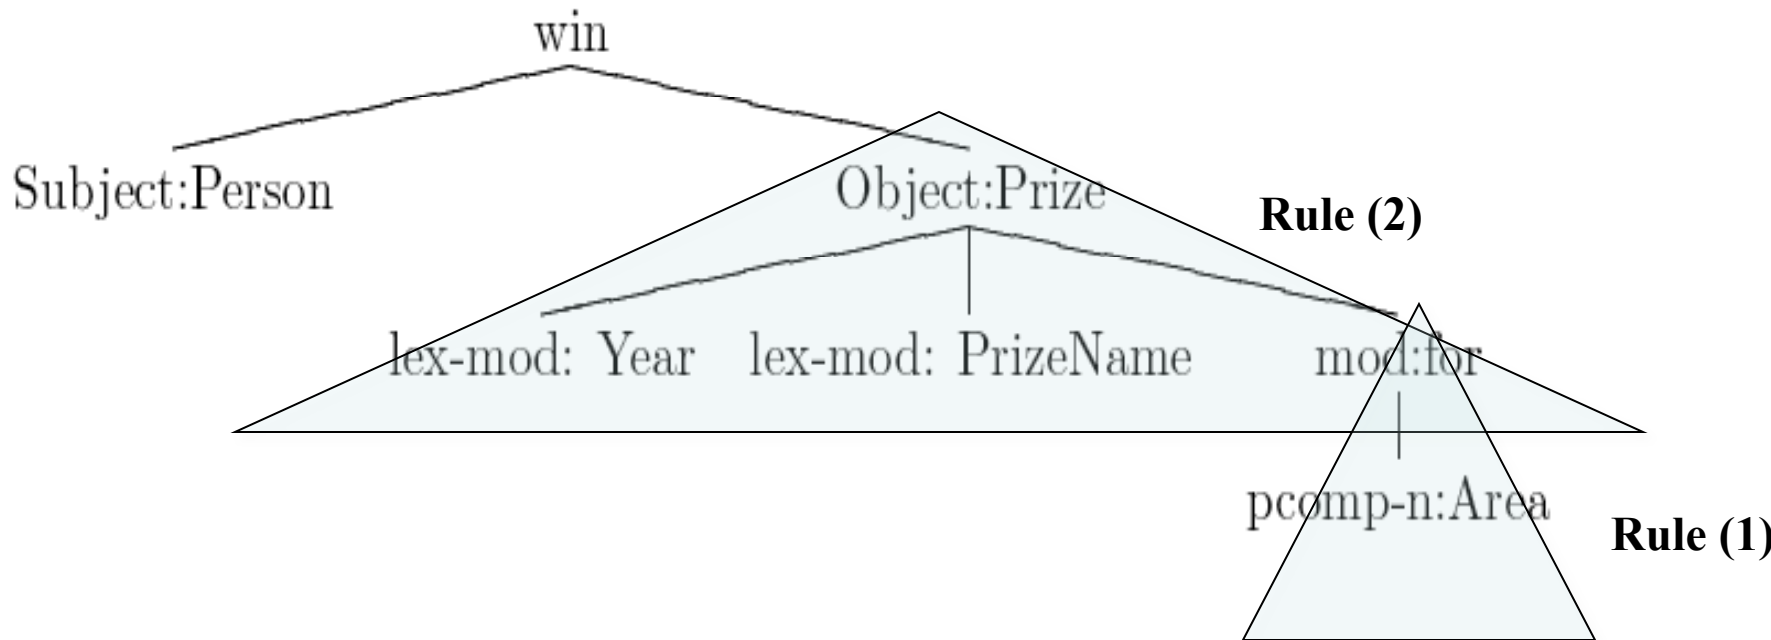
# Dependency Parse Result



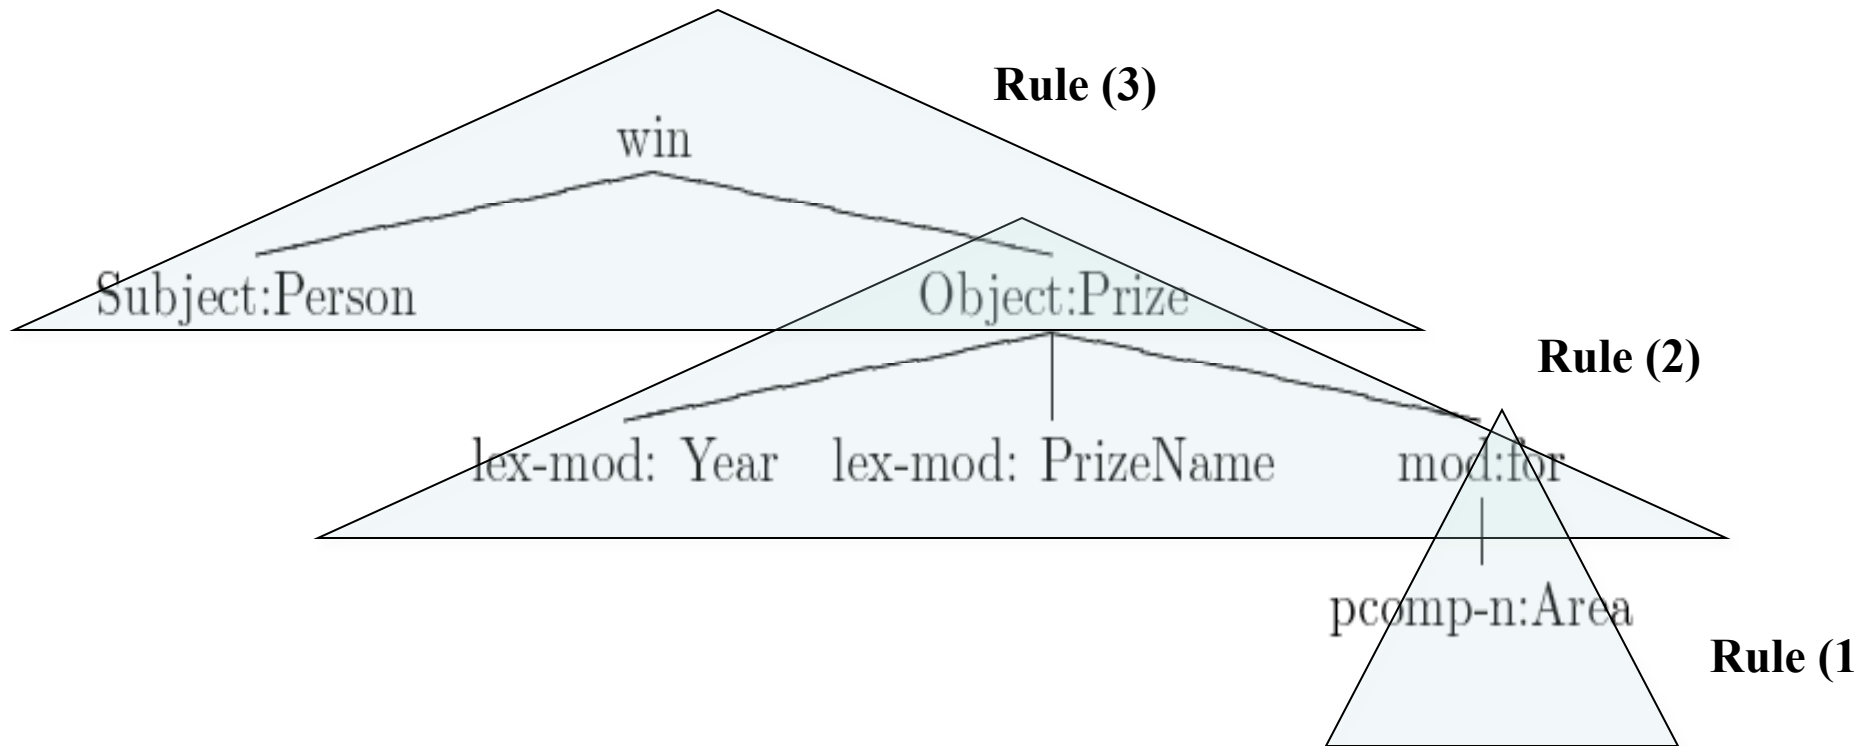
# Bottom Up Rule Learning



# Bottom Up Rule Learning



# Bottom Up Rule Learning



# Rule (1)

## **2005 Nobel Prize for Peace**

Rule name:: area\_1

Rule body::  $\left[ \begin{array}{l} \text{head} \left[ \begin{array}{l} \text{pos} \quad \text{preposition} \\ \text{lex-form} \quad \text{"for"} \end{array} \right] \\ \text{daughters} \left\langle \left[ \text{pcomp-n} \left[ \text{head} \quad \boxed{1} \text{Area} \right] \right] \right\rangle \end{array} \right]$

Output::  $\langle \boxed{1} \text{Area} \rangle$



# Rule (3)

Rule name:: recipient\_prize\_area\_year\_1

Rule body::

```
[ head [ pos      verb
        mode      active
        lex-form   "win"
      ],
  daughters < [ subject [ head [1] Person ],
              [ object [ head [ pos      noun
                              lex-form   "prize"
                            ],
                        rule  year_prize_area_1:: <[4]Year,[2]Prize,[3]Area> ] ] ] > ]
```

Output:: <[1]Recipient,[2]Prize,[3]Area,[4]Year>

# DARE Rule Components

Rule name:: recipient\_prize\_area\_year\_1

Rule body::

```
[ [ head [ pos      verb
         mode     active
         lex-form  "win" ]
  daughters < [ subject [ head [1] Person
                        rule  recipient_1:: <[1]Person > ] ],
              [ object [ head [ lex-form "prize" ]
                        rule  prize_area_year_1:: <[2]Prize, [3]Area, [4]Year > ] ] ] ] ]
```

Output:: <[1]Recipient, [2]Prize, [3]Area, [4]Year >

# Two Domains

- Award Events (start with subdomain Nobel Prizes)

reasons: good news coverage  
complete list of all award events  
good starting point for other award domains

- Management Succession Events

reason: comparison with previous work

# Experiments

- Two domains
  - Nobel Prize Awards: <recipient, prize, area, year>
  - Management Succession: <person\_in, person\_out, position, organisation>
- Test data sets

<b>Data Set Name</b>	<b>Doc Number</b>	<b>Data Amount</b>
<b>Nobel Prize</b>	<b>3328</b>	<b>18.4 MB</b>
<b>MUC-6</b>	<b>199</b>	<b>1MB</b>

# Relation Extraction without Coreference Resolution

domain	data size	initial seed no.	precision	recall
Nobel Prize	18.4 MB	1	86.5%	50.7%
MUC-6	1 MB	55	62%	48%

# Management Succession Domain

Initial Seed #	Precision	Recall
<b>1</b>	<b>12.6%</b>	<b>7.0%</b>
<b>1</b>	<b>15.1%</b>	<b>21.8%</b>
<b>20</b>	<b>48.4%</b>	<b>34.2%</b>
<b>55</b>	<b>62.0%</b>	<b>48.0%</b>

# The Dream

- Wouldn't it be wonderful if we could always automatically learn most or all relevant patterns of some relation from one single semantic instance!
- Or at least find all event instances.
- This sounds too good to be true!


# Research Questions

As scientists we want to know



- Why does it work for some tasks?
- Why doesn't it work for all tasks?
- How can we estimate the suitability of domains?
- How can we deal with less suitable domains?

Careful analysis confirmed the following assumption:  
redundancy, both on patterns and event mentions, helps.

Frequently reported events make rare patterns reachable




Popular patterns help to reach rarely mentioned events




# Instance to Pattern

## Nobel Prize vs. Management Succession



# Rule to Instances

(Nobel Prize vs. Management Succession)



# Insights

- Results from graph theory help to understand the requirements on data.

**Example:** small world property

- For data sets with continents and islands, we can sometimes exploit additional data or auxiliary domains to bridge the islands by learning rare patterns.

**Example:** use of Nobel prize domain for learning patterns for events concerning less popular prizes (many other prizes could be detected)

# Conclusion

- DARE is the first approach to combine the idea of bootstrapping IE systems with a linguistic grammar
- This can be illustrated by a simple formula:
  - reusable generic linguistic knowledge
  - + raw data
  - + a few examples (seed)
  - = domain specific relation extraction grammar
- In addition to the obvious practical advantages, the approach offers theoretical benefits: It supports a view of IE as a systematic gradual approximation of language understanding.

# Overcoming Obstacles

## ◆ Obstacles to Recall

- ◆ **missing bridges between islands/continents** ....use of auxiliary data
- ◆ **overly specific rules**..... better rule generalization
- ◆ **spread over several sentences**
  - ◆ **missing coreferences**..... coreference resolution

## ◆ Obstacles to Precision

- ◆ **intrusion of other relations**..... learning of negative rules
- ◆ **modality contexts**..... learning of negative rules
- ◆ **Integration of more linguistic context and structures** ..... deep NLP

# **Web-Scale Relation Extraction with Distant Supervision (Web-DARE)**



# Major Issues of Web-DARE Approach

- ☆ Discovery of large rule sets for n-ary relations for
  - covering most linguistic variations for a relation
  - thus solving the notorious long tail problem of real-world applications
  
- ☆ Parallel rule learning
  - A novel variant of distant supervision: many relations are learned in parallel, which enables a new method of rule filtering
  
- ☆ Freebase as seed and Web as learning corpus

## Why Distant Supervision (Mintz et al., 2009)

- ☆ There are large free semantic knowledge DBs available for automatically labeling an unannotated corpus with relation mentions
  - e.g. Freebase, Yago
  
- ☆ Generating training data for the RE systems
  - without the need of manual example manufacturing or
  - error-prone iterative learning algorithms, such as bootstrapping

# Essential Type

- ☆ All relations, whose  $k$  required arguments are the same NE types, are of the same *essential type*.

# Target Relations in People Domain

Relation	ARGUMENT NAMES & Entity Types				
	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5
<i>marriage</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	CEREMONY location	FROM date	TO date
<i>sibling relationship</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	-	-	-
<i>romantic relationship</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	FROM date	TO date	-
<i>person parent</i>	PERSON person <sup>Ⓡ</sup>	PARENT A person <sup>Ⓡ</sup>	PARENT B person	-	-
<i>place lived</i>	PERSON person <sup>Ⓡ</sup>	LOCATION location <sup>Ⓡ</sup>	FROM date	TO date	-
<i>person birth</i>	PERSON person <sup>Ⓡ</sup>	PLACE location	DATE date	-	-
<i>country of nationality</i>	PERSON person <sup>Ⓡ</sup>	COUNTRY location <sup>Ⓡ</sup>	-	-	-
<i>person death</i>	PERSON person <sup>Ⓡ</sup>	PLACE location	DATE date	CAUSE death cause	-
<i>sexual orientation phase</i>	PERSON person <sup>Ⓡ</sup>	ORIENTATION sxl orntn <sup>Ⓡ</sup>	FROM date	TO date	-
<i>education</i>	PERSON person <sup>Ⓡ</sup>	INSTITUTION organization <sup>Ⓡ</sup>	DEGREE degree	FIELD field of study	FROM date
<i>person alternate name</i>	PERSON person <sup>Ⓡ</sup>	ALTERNATE NAME alternate name <sup>Ⓡ</sup>	-	-	-

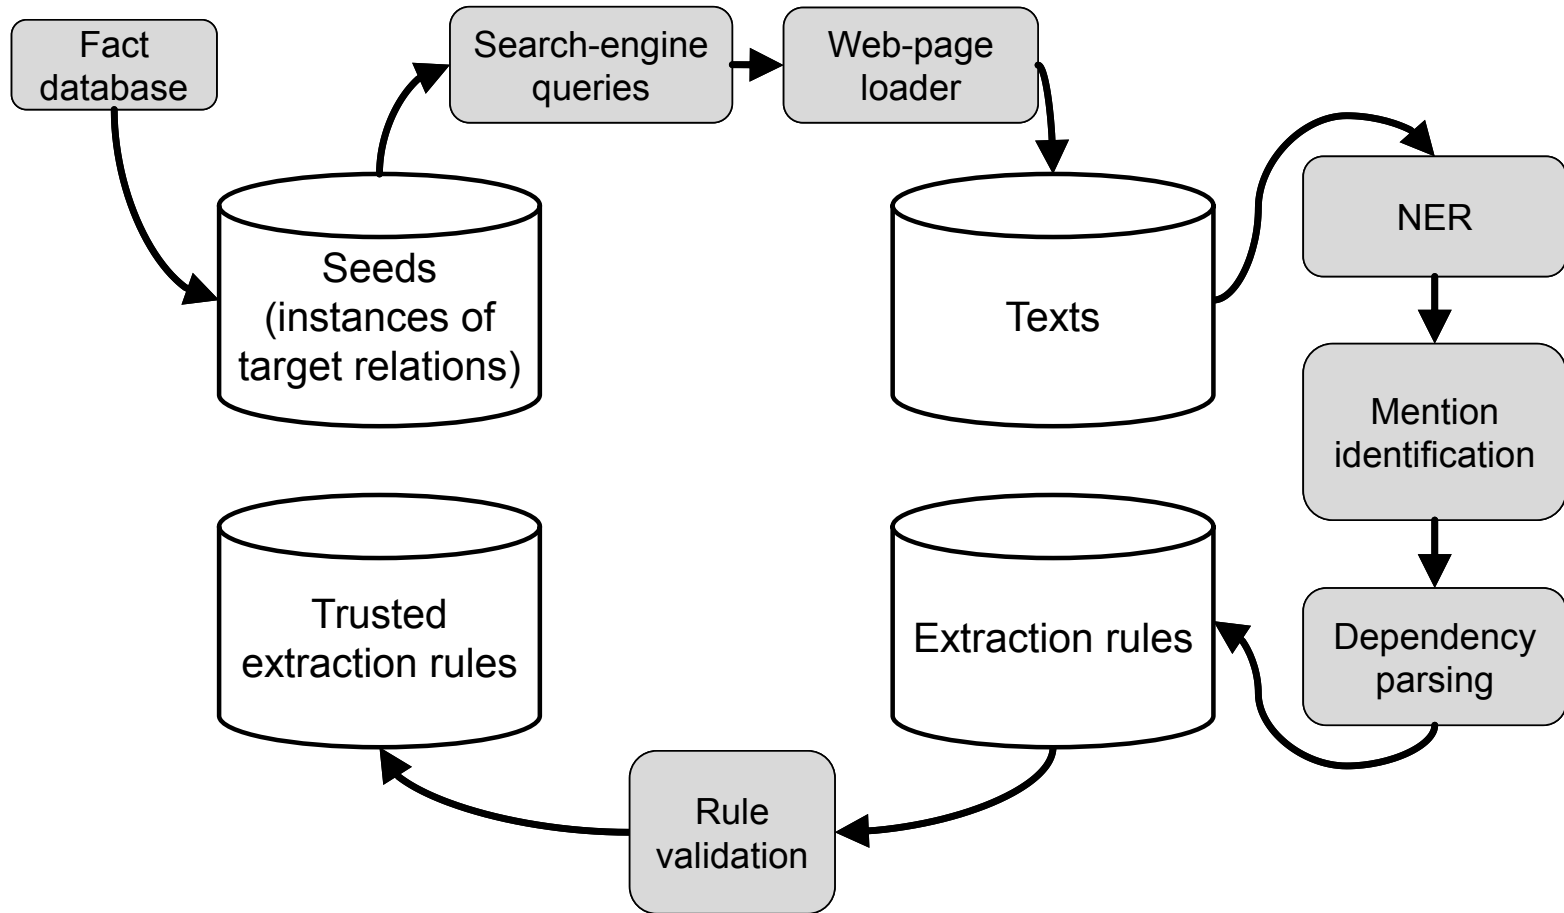


# Target Relations in People Domain

Relation	ARGUMENT NAMES & Entity Types				
	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5
<i>marriage</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	CEREMONY location	FROM date	TO date
<i>sibling relationship</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	-	-	-
<i>romantic relationship</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	FROM date	TO date	-
<i>person parent</i>	PERSON person <sup>Ⓡ</sup>	PARENT A person <sup>Ⓡ</sup>	PARENT B person	-	-
<i>place lived</i>	PERSON person <sup>Ⓡ</sup>	LOCATION location <sup>Ⓡ</sup>	FROM date	TO date	-
<i>person birth</i>	PERSON person <sup>Ⓡ</sup>	PLACE location	DATE date	-	-
<i>country of nationality</i>	PERSON person <sup>Ⓡ</sup>	COUNTRY location <sup>Ⓡ</sup>	-	-	-
<i>person death</i>	PERSON person <sup>Ⓡ</sup>	PLACE location	DATE date	CAUSE death cause	-
<i>sexual orientation phase</i>	PERSON person <sup>Ⓡ</sup>	ORIENTATION sxl orntn <sup>Ⓡ</sup>	FROM date	TO date	-
<i>education</i>	PERSON person <sup>Ⓡ</sup>	INSTITUTION organization <sup>Ⓡ</sup>	DEGREE degree	FIELD field of study	FROM date
<i>person</i>	PERSON person <sup>Ⓡ</sup>	ALTERNATE NAME name <sup>Ⓡ</sup>	-	-	-



# System Architecture (Web DARE)



# Example of Rule Extraction

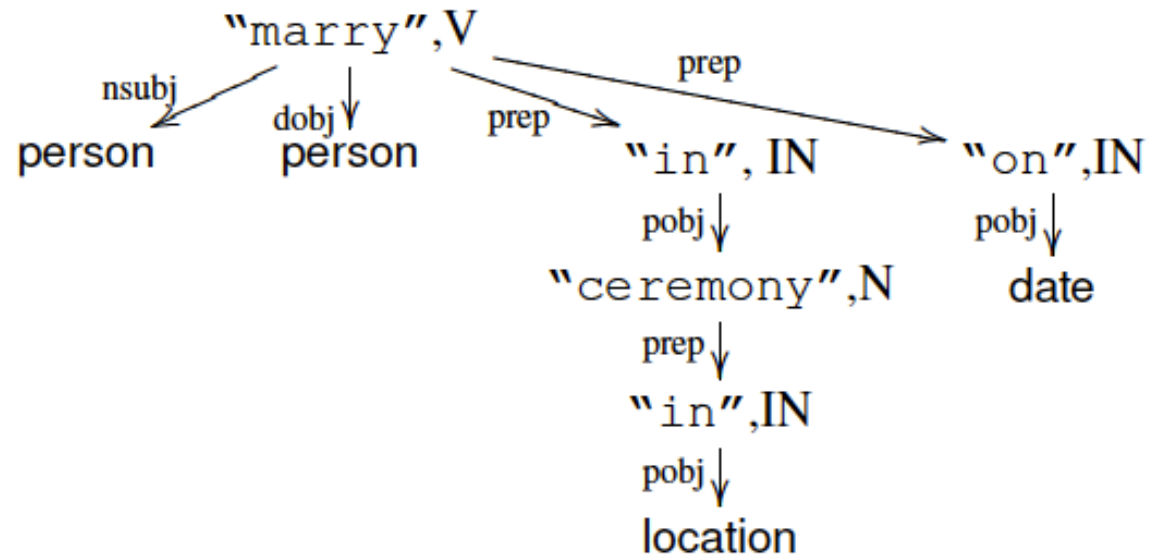
## ☆ Seed

<Brad Pitt, Jennifer Aniston, Malibu, 2000, 2005>

## ☆ Text Mention

Brad Pitt married Jennifer Aniston in a private wedding ceremony in Malibu on July 29, 2000.

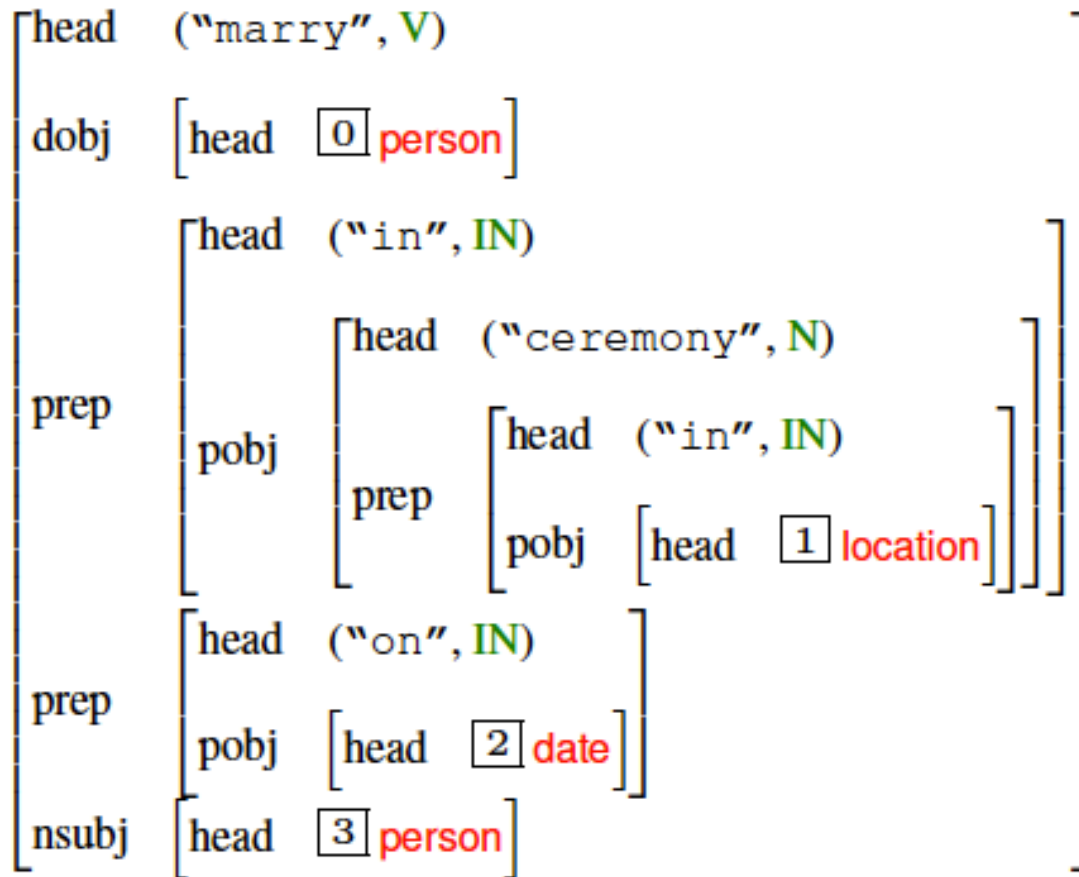
## ☆ Dependency parsing



# Example of Learned Rule

**Rule name** :: PersonA\_PersonB\_Ceremony\_From

**Rule body** ::



**Output** ::  $\langle [0] \text{ PERSON A}, [3] \text{ PERSON B}, [1] \text{ CEREMONY}, [2] \text{ FROM}, \text{---} \rangle$

# Experiment and Evaluation

## ☆ Training

- learning rules for 39 relations
- three domains: business, award and people
- 2.8 million relation instances retrieved from Freebase as seed
- 20 million web documents
- 3 million sentence mentions

## ☆ Testing

- ACE 2005 corpus
- Celebrity corpus

# Statistics of Document Acquisition

<b>Relation</b>	<b># Seeds</b>	<b>#Seeds used</b>	<b>#Seeds w/ search hits</b>	<b>#Search hits</b>
<i>award nomination</i>	86,087	12,969	12,730	1,000,141
<i>award honor</i>	48,917	11,013	10,643	1,000,021
<i>hall of fame induction</i>	2,208	2,208	2,200	443,416
<i>organization relationship</i>	219,583	70,946	51,538	1,000,009
<i>acquisition</i>	1,768	1,768	1,663	308,650
<i>organization merger</i>	154	154	145	25,976
<i>company name change</i>	1,051	1,051	939	124,612
<i>spin off</i>	222	222	216	32,613
<i>marriage</i>	16,616	6,294	5,993	1,000,174
<i>sibling relationship</i>	8,246	8,246	7,630	914,582
<i>romantic relationship</i>	544	544	544	280,508
<i>person parent</i>	23,879	3,447	3,379	1,000,023
<b>average of 39 relations</b>	<i>72.576</i>	<i>6.095</i>	<i>5.374</i>	<i>635.927</i>

# Statistics of Document Acquisition

<b>Relation</b>	<b># Seeds</b>	<b>#Seeds used</b>	<b>#Seeds w/ search hits</b>	<b>#Search hits</b>
<i>award nomination</i>	86,087	12,969	12,730	1,000,141
<i>award honor</i>	48,917	11,013	10,643	1,000,021
<i>hall of fame induction</i>	2,208	2,208	2,200	443,416
<i>organization relationship</i>	219,583	70,946	51,538	1,000,009
<i>acquisition</i>	1,768	1,768	1,663	308,650
<i>organization merger</i>	154	154	145	25,976
<i>company name change</i>	1,051	1,051	939	124,612
<i>spin off</i>	222	222	216	32,613
<i>marriage</i>	16,616	6,294	5,993	1,000,174
<i>sibling relationship</i>	8,246	8,246	7,630	914,582
<i>romantic relationship</i>	544	544	544	280,508
<i>person parent</i>	23,879	3,447	3,379	1,000,023
<b>average of 39 relations</b>	<b>72.576</b>	<b>6.095</b>	<b>5.374</b>	<b>635.927</b>

# Statistics of Rule Learning

<b>Relation</b>	<b># Sentences used</b>	<b># Sentences w/ a learned rule</b>	<b># Rules</b>	<b># Rules w/o duplicates</b>
<i>award nomination</i>	13,966	13,149	23,987	7,800
<i>award honor</i>	50,550	49,001	106,550	40,578
<i>hall of fame induction</i>	31,244	28,278	44,920	17,450
<i>organization relationship</i>	46,331	42,824	60,379	28,903
<i>acquisition</i>	63,967	60,903	96,747	50,544
<i>organization merger</i>	2,996	1,521	3,243	1,758
<i>company name change</i>	9,433	9,132	15,619	6,910
<i>spin off</i>	5,247	5,094	8,319	4,798
<i>marriage</i>	342,895	335,313	557,478	176,949
<i>sibling relationship</i>	167,611	160,893	255,788	69,596
<i>romantic relationship</i>	155,335	152,878	229,393	74,895
<i>person parent</i>	192,610	186,834	390,878	119,238
<b>average of 39 relations</b>	<b>66,545</b>	<b>66,509</b>	<b>109,435</b>	<b>41,620</b>

# Statistics of Rule Learning

<b>Relation</b>	<b># Sentences used</b>	<b># Sentences w/ a learned rule</b>	<b># Rules</b>	<b># Rules w/o duplicates</b>
<i>award nomination</i>	13,966	13,149	23,987	7,800
<i>award honor</i>	50,550	49,001	106,550	40,578
<i>hall of fame induction</i>	31,244	28,278	44,920	17,450
<i>organization relationship</i>	46,331	42,824	60,379	28,903
<i>acquisition</i>	63,967	60,903	96,747	50,544
<i>organization merger</i>	2,996	1,521	3,243	1,758
<i>company name change</i>	9,433	9,132	15,619	6,910
<i>spin off</i>	5,247	5,094	8,319	4,798
<i>marriage</i>	342,895	335,313	557,478	176,949
<i>sibling relationship</i>	167,611	160,893	255,788	69,596
<i>romantic relationship</i>	155,335	152,878	229,393	74,895
<i>person parent</i>	192,610	186,834	390,878	119,238
<b>average of 39 relations</b>	66,545	66,509	109,435	41,620

# Processing Time

Average time for all relations: 19 hours

<b>Relation</b>	<b>Search engine querying</b>	<b>Doc. download &amp; text extraction</b>	<b>Mention finding (incl. NER)</b>	<b>Rule learning (incl. parsing)</b>
<i>award nomination</i>	12:01	8:27	4:00	0:35
<i>award honor</i>	12:29	8:14	4:57	0:50
<i>hall of fame induction</i>	4:22	3:22	2:43	0:15
<i>organization relationship</i>	12:53	18:43	23:34	0:27
<i>acquisition</i>	2:44	4:31	1:25	0:20
<i>organization merger</i>	0:16	0:35	0:07	0:01
<i>company name change</i>	1:03	1:16	2:14	0:04
<i>spin off</i>	0:20	1:01	0:08	0:03
<i>marriage</i>	9:47	20:34	4:50	1:05
<i>sibling relationship</i>	5:17	11:25	4:13	1:05
<i>romantic relationship</i>	1:13	3:32	1:20	0:23
<i>person parent</i>	4:59	12:07	4:39	0:45
<b>average of 39 relations</b>	<b>6:17</b>	<b>8:40</b>	<b>3:11</b>	<b>0:40</b>

## Marriage Relation and Learned Rules

Relation	ARGUMENT NAMES & Entity Types				
	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5
<i>marriage</i>	PERSON A person <sup>®</sup>	PERSON B person <sup>®</sup>	CEREMONY location	FROM date	TO date

Arity	# Rules	Min. Freq.	Avg. Freq.	Med. Freq.	Max. Freq.
<b>2</b>	145,598	1	3.21	1	64,015
<b>3</b>	26,294	1	2.90	1	2,655
<b>4</b>	4,350	1	3.07	1	603
<b>5</b>	40	1	1.40	1	10

# Test Corpora

	<b>ACE 2005</b>	<b>Celebrity</b>
# Documents	599	150
# Gold-standard mentions	112	498
# Sentences used	3,612	7,463
# Gold-standard mentions used	14	259

## Baseline Performance

	<b>ACE 2005</b>	<b>Celebrity</b>
<b>Precision</b>	0.68 %	3.05 %
<b>Recall</b>	71.43 %	49.42 %
<b>F-measure</b>	1.36 %	5.74 %
<b># Correct Mentions</b>	11	125
<b># Wrong Mentions</b>	1,595	3,978
<b># Good Rules</b>	25	110
<b># Dangerous Rules</b>	1	105
<b># Bad Rules</b>	1,164	1,971
<b># Useless Rules</b>	175,759	174,763

## Bootstrapping vs. Web DARE (celebrity corpus)

- DARE learns rules from a comparable corpus with bootstrapping: same size and same source as the celebrity test corpus
- Comparison of the performance of
  - bootstrapping: the learned rules with bootstrapping
  - Web rules: the learned rules from Web with distant supervision


	<b>Bootstrapping</b>	<b>Web Rules</b>
<b>Precision</b>	2.88 %	3.05 %
<b>Recall</b>	30.5 %	49.42 %
<b>F-measure</b>	5.26 %	5.74 %
<b># Learned Rules</b>	3013	176,949
<b># Good+Danger. Rules</b>	39	215

# Target Relations in People Domain

Relation	ARGUMENT NAMES & Entity Types				
	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5
<i>marriage</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	CEREMONY location	FROM date	TO date
<i>sibling relationship</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	-	-	-
<i>romantic relationship</i>	PERSON A person <sup>Ⓡ</sup>	PERSON B person <sup>Ⓡ</sup>	FROM date	TO date	-
<i>person parent</i>	PERSON person <sup>Ⓡ</sup>	PARENT A person <sup>Ⓡ</sup>	PARENT B person	-	-
<i>place lived</i>	PERSON person <sup>Ⓡ</sup>	LOCATION location <sup>Ⓡ</sup>	FROM date	TO date	-
<i>person birth</i>	PERSON person <sup>Ⓡ</sup>	PLACE location	DATE date	-	-
<i>country of nationality</i>	PERSON person <sup>Ⓡ</sup>	COUNTRY location <sup>Ⓡ</sup>	-	-	-
<i>person death</i>	PERSON person <sup>Ⓡ</sup>	PLACE location	DATE date	CAUSE death cause	-
<i>sexual orientation phase</i>	PERSON person <sup>Ⓡ</sup>	ORIENTATION sxl orntn <sup>Ⓡ</sup>	FROM date	TO date	-
<i>education</i>	PERSON person <sup>Ⓡ</sup>	INSTITUTION organization <sup>Ⓡ</sup>	DEGREE degree	FIELD field of study	FROM date
<i>person</i>	PERSON person <sup>Ⓡ</sup>	ALTERNATE NAME name <sup>Ⓡ</sup>	-	-	-



# Euler Diagram for Four People Relations



# Rule Filtering

☆ Merged Filter:

$$valid_m^{\mathcal{R}}(r) = valid_{freq}^{\mathcal{R}}(r) \wedge valid_{inter}^{\mathcal{R}}(r)$$

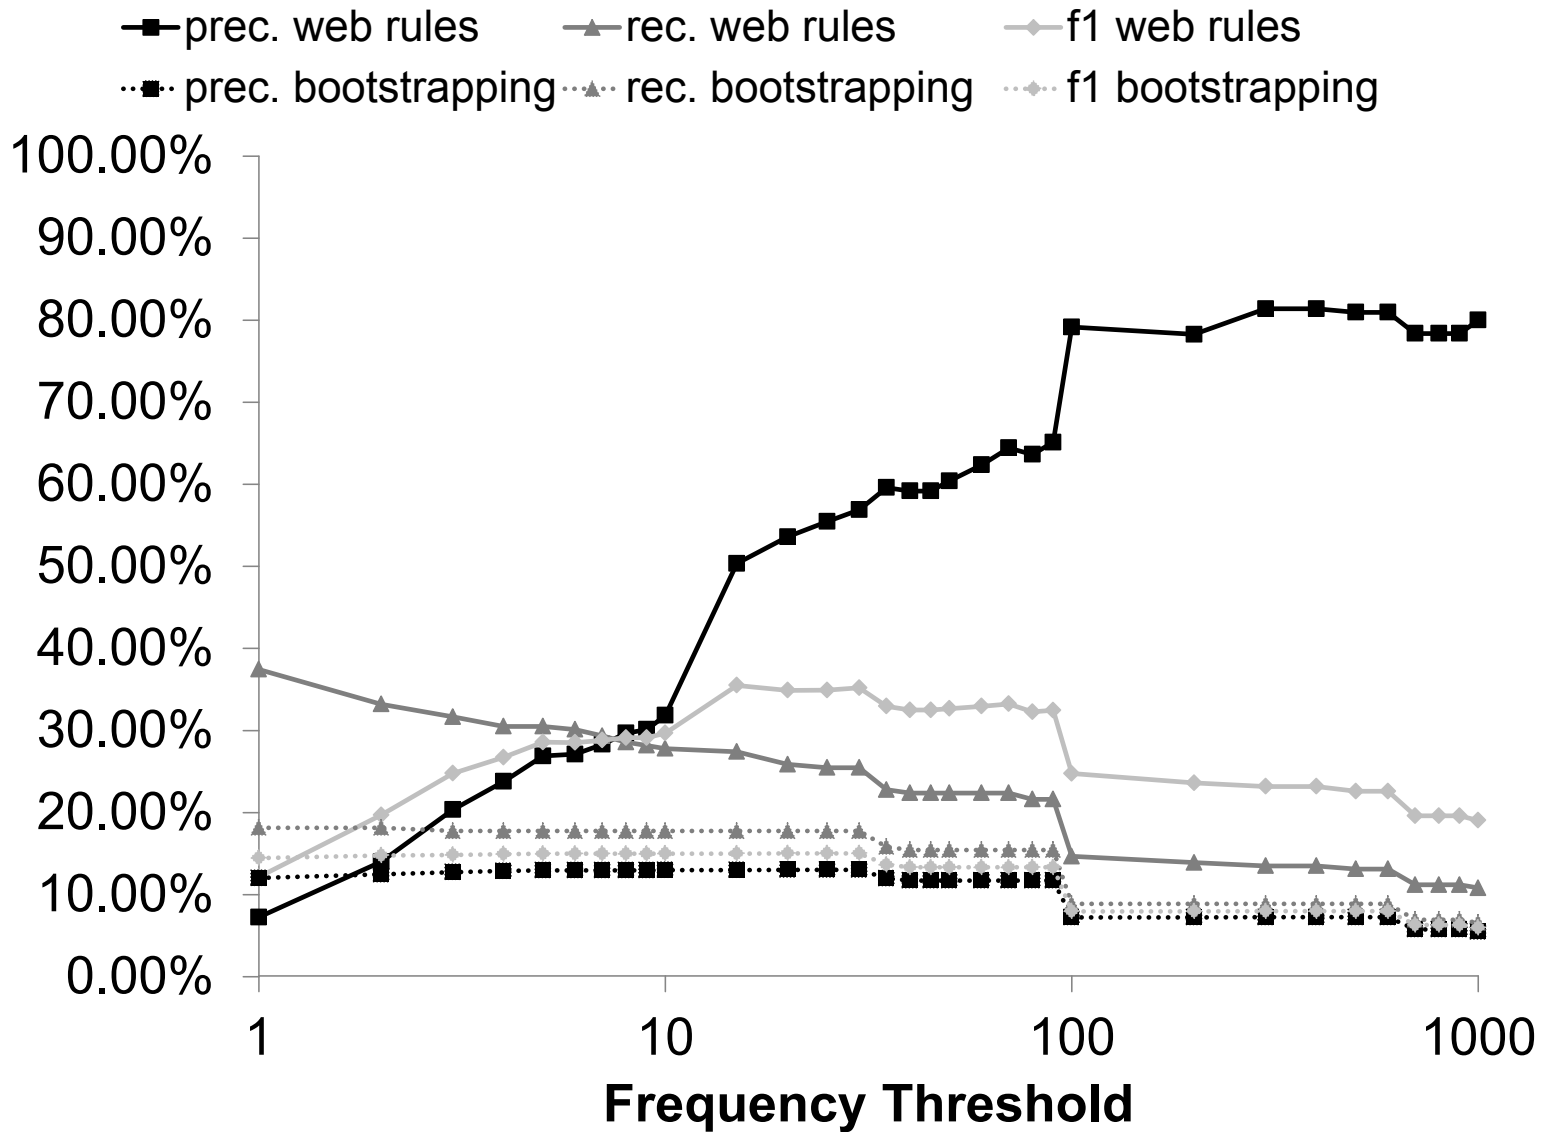
- 1) absolute frequency filtering
- 2) inter-relation filter

Let  $\mathbb{R}$  be a set of relations of the same essential type. The rule  $r$  is *valid* for the relation  $\mathcal{R} \in \mathbb{R}$  if the relative frequency of  $r$  in  $\mathcal{R}$  is higher than its relative frequencies for all other relations in  $\mathbb{R}$ :

$$valid_{inter}^{\mathcal{R}}(r) = \begin{cases} true & \text{if } \forall \mathcal{R}' \in \mathbb{R} \setminus \{\mathcal{R}\} : rf_{r,\mathcal{R}} > rf_{r,\mathcal{R}'} \\ false & \text{otherwise} \end{cases} \quad (4)$$



# Performance after Rule Filtering



## Error Analysis of False Negatives (Missing Mentions)

<b>AE</b>	<b>PE</b>	<b>NE</b>	<b>PE <math>\cup</math> NE</b>	<b>SU</b>	<b>MR</b>
X	–	–	–	–	–
–	X	X	X	X	–
–	X	X	X	–	X
6	78	54	111	28	7–34
4.58 %	59.54 %	41.22 %	84.73 %	21.37 %	5.34 %–25.95 %

- AE: annotation error
- PE: dependency parsing
- NE: named entity recognition
- SU: semantic understand
- MR: no rule learned

## Example of Semantic Understanding

*“...that sounded good to a tired mom like me,” says Ruiz, 34, who has two children, James, 8, and Justine, 6, with husband Donald, 42, ...*

the syntactic dependencies do not link "husband" directly to "Ruiz".

# Conclusion

- Our system exploits the Web for training corpus and Freebase as seed for the extraction of n-ary relations
- The large amount of learned rules can lead to improvement of recall
- There is still a big potential for precision improvement by exploiting relations of the same essential type
- The error analysis showed that the linguistic preprocessing has strong influence on our recall performance

## Future Work

- Integration of semantic constraints in the rule discovery
- Integration of logical relations among the relations for rule filtering
  - Super/subset relations
  - Mutual exclusive relations
- Improvement of NER
- Large-scale domain adaptive parsing
  - Parse reranking for many relations in parallel

# Task-Driven Anaphora Resolution

# Example

1. Three of the Nobel Prizes for Chemistry during the first decade **were awarded** for pioneering work in organic chemistry.
2. In **1902** **Emil Fischer** (1852-1919), then in Berlin, **was given** the prize for his work on sugar and purine syntheses.
3. Another major influence from organic chemistry was the development of the chemical industry, and a chief contributor here was Fischer's teacher, **Adolf von Baeyer** (1835-1917) in Munich, who **was awarded** the prize in **1905**.

# Anaphora in Texts

He/The scientist won the 2005 Nobel Prize for Peace on Friday for his efforts to limit the spread of atomic weapons.




<?PERSON, Nobel, Peace, 2005>

# Coreference Relations and Indicators

- Complex linguistic phenomena, influenced by lexical, syntactic, semantic and discourse constraints
  
- The indicators shared by many approaches are
  - Distance: coreference expressions are often close to each other in the surface structure;
  - Syntactic: pronominal resolution constraints within sentence
  - Semantic: same or compatible semantic category, agreement in number, gender and person;
  - Discourse: parallelism, repetition, apposition, name alias.

# Recency Indicator in Nobel Prize Domain

- News reports from New York Times, online BBC and CCN (18.4 MB, 3328 documents)



- Two Americans have won the 2002 Nobel Prize in Economic Sciences.
- The two scientists, Daniel Kahneman and Vernon L. Smith, received the honour on Wednesday for their work using psychological research and laboratory experiments in economic analysis.

- Egypt honours its Nobel Prize chemist.
- President Hosni Mubarak of Egypt has awarded the country's most prestigious prize - the Nile Necklace - to the Egyptian-born chemist Ahmed Zewail.

# Repetition and Elaboration

- Cohesion indicator *repetition* is often used as indicator for semantic similarity and semantic consistency, e.g.,
  - „two Americans“ and „two scientists“
  - „chemist“ and „chemist“
  
- Elaboration phenomena are normal in newspaper texts
  - S1 is an Elaboration of S0 if a proposition P follows from the assertions of both S0 and S1, but S1 contains a property of one of the elements of P that is not in S0 (Hobbs, 1979)

# Relation Argument as a Complex Semantic Object

- A complex noun phrase contains often more than one property about an argument: e.g.

*Egyptian-born chemist Ahmed Zewail*

- Relevant properties of a winner in Nobel Prize domain
  - Nationality/origin/inhabitant: e.g., two Americans, the Egyptian-born, a Dutch
  - Profession/occupation: e.g., novelist, chemist, scientist, researcher
  - Title/position: e.g., professor, president
  - Domain description: e.g., recipient, winner, Nobel Laureate
  - General description: e.g., the man, a woman, the team

*„two Americans“*

*sentence\_id : i*

*number : [ type : plural  
amount : 2 ]*

*definite : indef*

*grammarrole : subject*

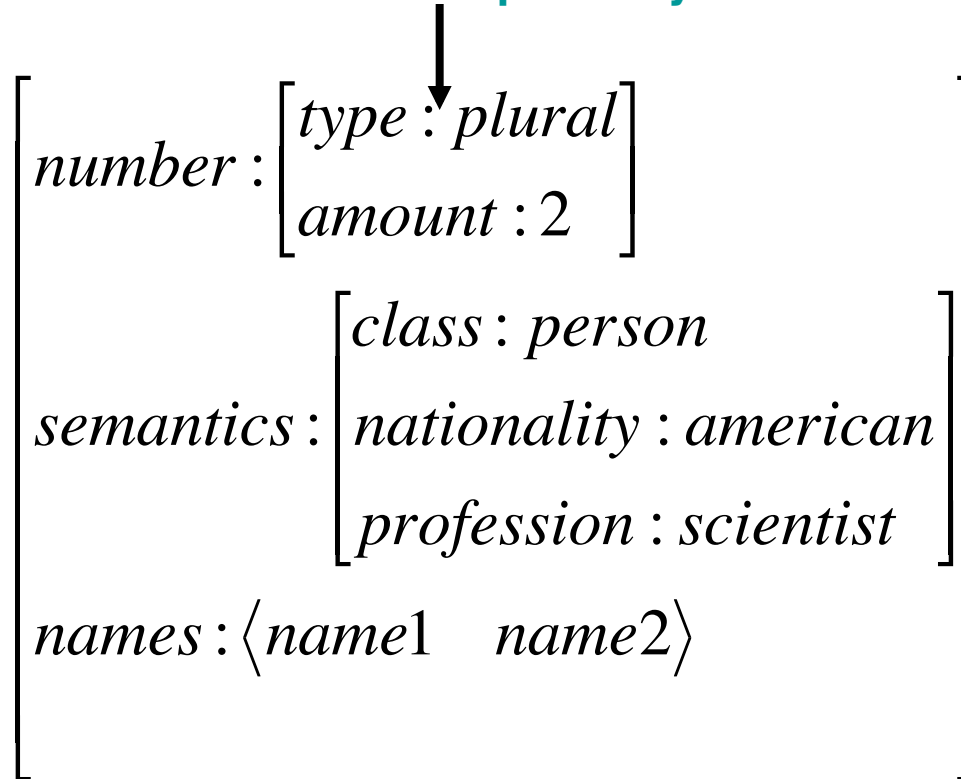
*semantics : [ nationality : american ]*

„the two scientists, ....“

$$\left[ \begin{array}{l} \textit{sentence\_id} : i + 1 \\ \textit{number} : \left[ \begin{array}{l} \textit{type} : \textit{plural} \\ \textit{amount} : 2 \end{array} \right] \\ \textit{definite} : \textit{def} \\ \textit{grammarrole} : \textit{subject} \\ \textit{semantics} : \left[ \begin{array}{l} \textit{class} : \textit{person} \\ \textit{profession} : \textit{scientist} \end{array} \right] \\ \textit{names} : \langle \textit{name1} \quad \textit{name2} \rangle \end{array} \right]$$

Unification of  
„two Americans“ and „the two scientists,...“

- **Recency and valide chain**
- **Parallel subject**
- **Repetition: number**
- **Semantic compatibility**



## References

1. N. Kushmerick. Wrapper induction: Efficiency and Expressiveness, Artificial Intelligence, 2000.
2. I. Muslea. Extraction Patterns for Information Extraction. AAI-99 Workshop on Machine Learning for Information Extraction.
3. Riloff, E. and R. Jones. Learning Dictionaries for Information Extraction by Multi-Level Bootstrapping. In Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99) , 1999, pp. 474-479.
4. R. Yangarber, R. Grishman, P. Tapanainen and S. Huttunen. Automatic Acquisition of Domain Knowledge for Information Extraction. In Proceedings of the 18th International Conference on Computational Linguistics: COLING-2000, Saarbrücken.
5. F. Xu, H. Uszkoreit and Hong Li. Automatic Event and Relation Detection with Seeds of Varying Complexity. In Proceedings of AAAI 2006 Workshop Event Extraction and Synthesis, Boston, July, 2006.
6. F. Xu, D Kurz, J Piskorski, S Schmeier. A Domain Adaptive Approach to Automatic Acquisition of Domain Relevant Terms and their Relations with Bootstrapping. In Proceedings of LREC 2002.
7. W. Drozdowski, H.U. Krieger, J. Piskorski, U. Schäfer and F. Xu. Shallow Processing with Unification and Typed Feature Structures -- Foundations and Applications. In KI (Artificial Intelligence) journal 2004.
8. Feiyu Xu, Hans Uszkoreit, Hong Li. A Seed-driven Bottom-up Machine Learning Framework for Extracting Relations of Various Complexity. In Proceedings of ACL 2007, Prague
9. <http://www.dfki.de/~neumann/ie-essli04.html>

# DARE and Extensions

(<http://dare.dfki.de>)

- Xu, Feiyu, Hans Uszkoreit, and Hong Li. 2007. A seed-driven bottom-up machine learning framework for extracting relations of various complexity. *ACL 2007*
- Xu, Feiyu. 2007. *Bootstrapping Relation Extraction from Semantic Seeds*. Phd-thesis.
- Feiyu Xu, Hans Uszkoreit, Hong Li. *Task driven coreference resolution for relation extraction*. *ECAI 2008*.
- Xu, Feiyu, Hans Uszkoreit, Hong Li, and Niko Felger. *Adaptation of relation extraction rules to new domains*. *LREC 2008*.
- Hans Uszkoreit, Feiyu Xu, Hong Li. *Analysis and Improvement of Minimally Supervised Machine Learning for Relation Extraction*. *NLDB 2009*. *Keynote*.
- Xu, Feiyu, Hans Uszkoreit Sebastian Krause and Hong Li. *Boosting relation extraction with limited closed-world knowledge*. *COLING 2010*.

# DARE and Extensions

(<http://dare.dfki.de>)

- Feiyu Xu, Hong Li, Yi Zhang, Hans Uszkoreit, Sebastian Krause. [Minimally Supervised Domain-Adaptive Parse Reranking for Relation Extraction](#). In Proceedings of International Conference on Parsing Technologies (IWPT 2011), Dublin, Ireland.
- Sebastian Krause, Hong Li, Hans Uszkoreit, Feiyu Xu. [Large-Scale Learning of Relation-Extraction Rules with Distant Supervision from the Web](#). In Proceedings of the 11th International Semantic Web Conference, Boston, USA, Springer, 11/2012.