Predicting
a Structured Semantic
Representation for Text
(And What We Can Do With One)

Vivek Srikumar

Stanford University Natural Language Processing

THE UNIVERSITY OF UTAH

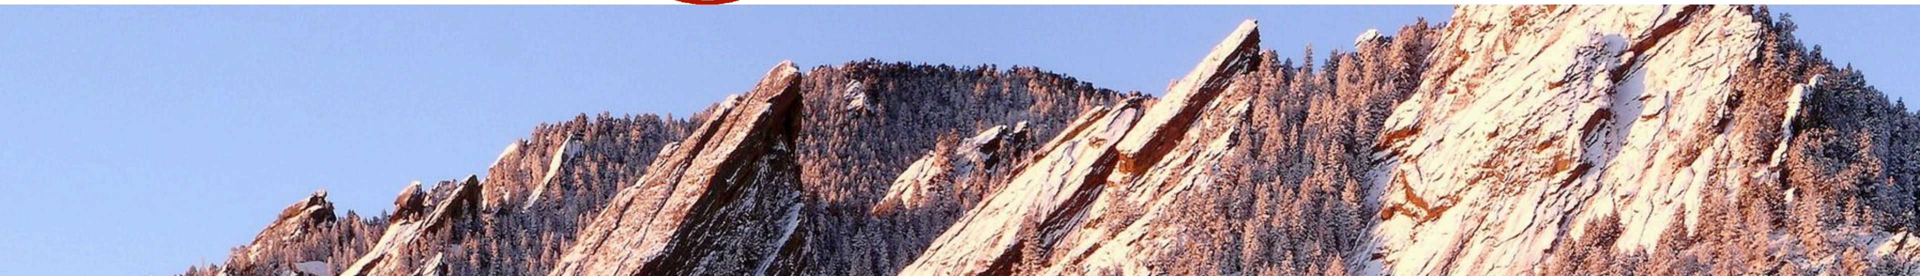

Outline

1. Semantic Role Labeling
   – Looking beyond verbs and nominalizations
   – Preposition relations

2. Using a semantic representation
   – Looking beyond sentences
   – Reading comprehension

Understanding text

On Feb. 20, 1962, a Marine Corps fighter pilot from small-town America stepped forward in response to the country's need. The astronaut was John Glenn, whom the author Tom Wolfe has called "the last true national hero America has ever had".
On Monday and Tuesday, Glenn will be honored with a dinner and a spaceflight forum at Ohio State University, home of the John Glenn School of Public Affairs.

- Tom Wolfe is an author.
- The spaceflight forum is located in the Ohio State University.
- John Glenn will be honored on Monday.
- John Glenn School of Public Affairs is located at the Ohio State University.

Understanding text

On Feb. 20, 1962, a Marine Corps fighter pilot from small-town America stepped forward in response to the country's need. The astronaut was John Glenn, whom the author Tom Wolfe has called "the last true national hero America has ever had".
On Monday and Tuesday, Glenn will be honored with a dinner and a spaceflight forum at Ohio State University, home of the John Glenn School of Public Affairs.

Who called John Glenn "the last true national hero America has ever had"?

~~A: The astronaut~~
B: The author Tom Wolfe

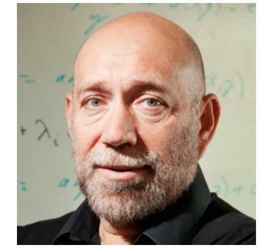

Dan Roth

Representing Predicates and Arguments

Extending Semantic Role Labeling: Prepositions and more

[Srikumar &Roth ’11, ‘13]

Semantic Role Labeling (SRL)

Given a sentence, identifies who does what to whom, where and when.

The astronaut was John Glenn, whom the author Tom Wolfe has **called**
"the last true national hero America has ever had."

Predicate

Arguments

Relation: `Call`

Caller[**A0**]: the author Tom Wolfe

Item being labeled [**A1**]: John Glenn

Attribute [**A2**]: the last true national hero America has ever had

"**the author Tom Wolfe**" fulfills the role of *Caller* for the predicate `Call`

Who called John Glenn "the last true national hero America has ever had"?

Verbs do not give the full picture

On Feb. 20, 1962, a Marine Corps fighter pilot from small-town America stepped forward in response to the country's need. The astronaut was John Glenn, whom the author Tom Wolfe has called “the last true national hero America has ever had”.
On Monday and Tuesday, Glenn will be honored with a dinner and a spaceflight forum at Ohio State University, home of the John Glenn School of Public Affairs.

- ✗ Tom Wolfe is an author.
- ✗ The spaceflight forum is located in the Ohio State University.
- ✓ John Glenn will be honored on Monday.
- ✗ John Glenn School of Public Affairs is located at the Ohio State University.

Prepositions trigger relations

John enjoyed the visit to $\boxed{\text{the zoo in NYC}}$.

- Enjoy
  - **`Agent/Enjoyer`**: John
  - **`Cause/Thing-enjoyed`**: the visit to the zoo in NYC
- Visit
  - **`Agent`**: John
  - **`Destination`**: the zoo in NYC

Q: Where is the zoo located?

A: NYC.

Ontology of Preposition Relations

Examples of preposition relations

Possessor

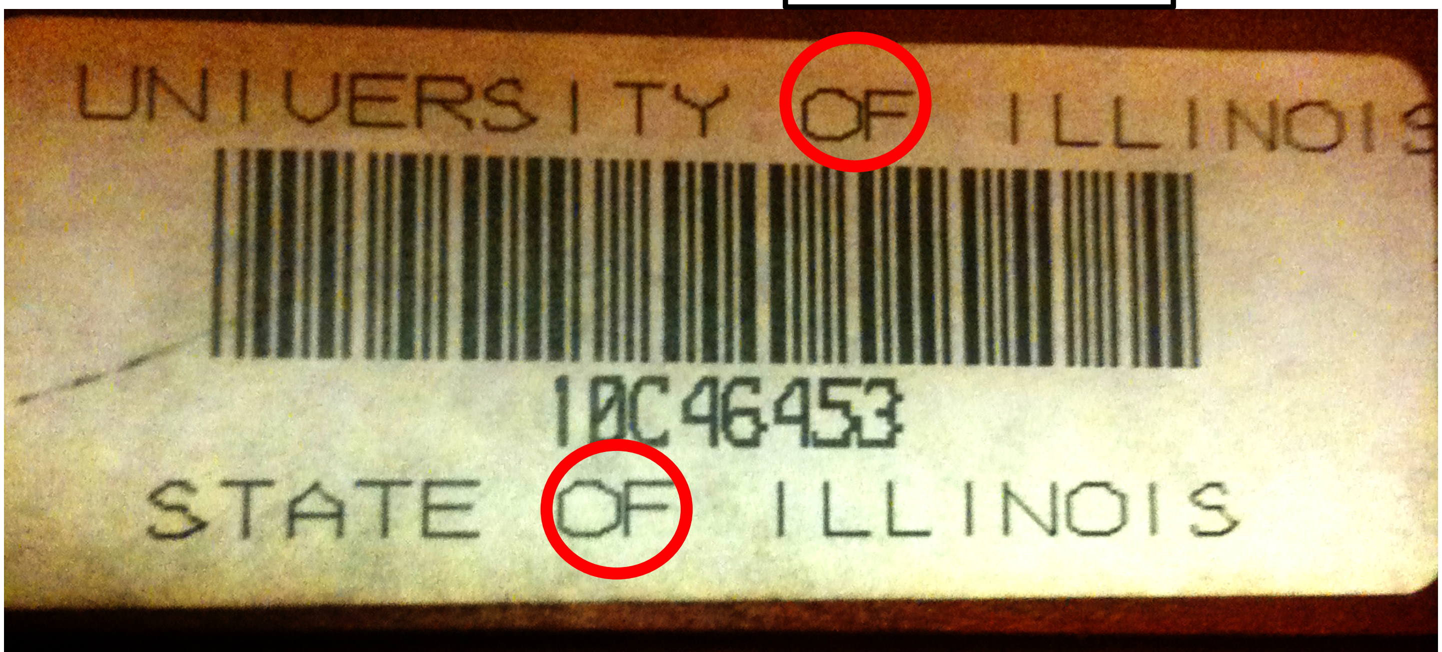

Species

Preposition Sense Disambiguation

Eg. *State **of** Illinois* vs. *University **of** Illinois*

- The Preposition Project [Litkowski and Hargraves, 2005]
  - Word sense for 34 prepositions
  - Based on preposition definitions in Oxford Dictionary of English

Mapping from senses to relations

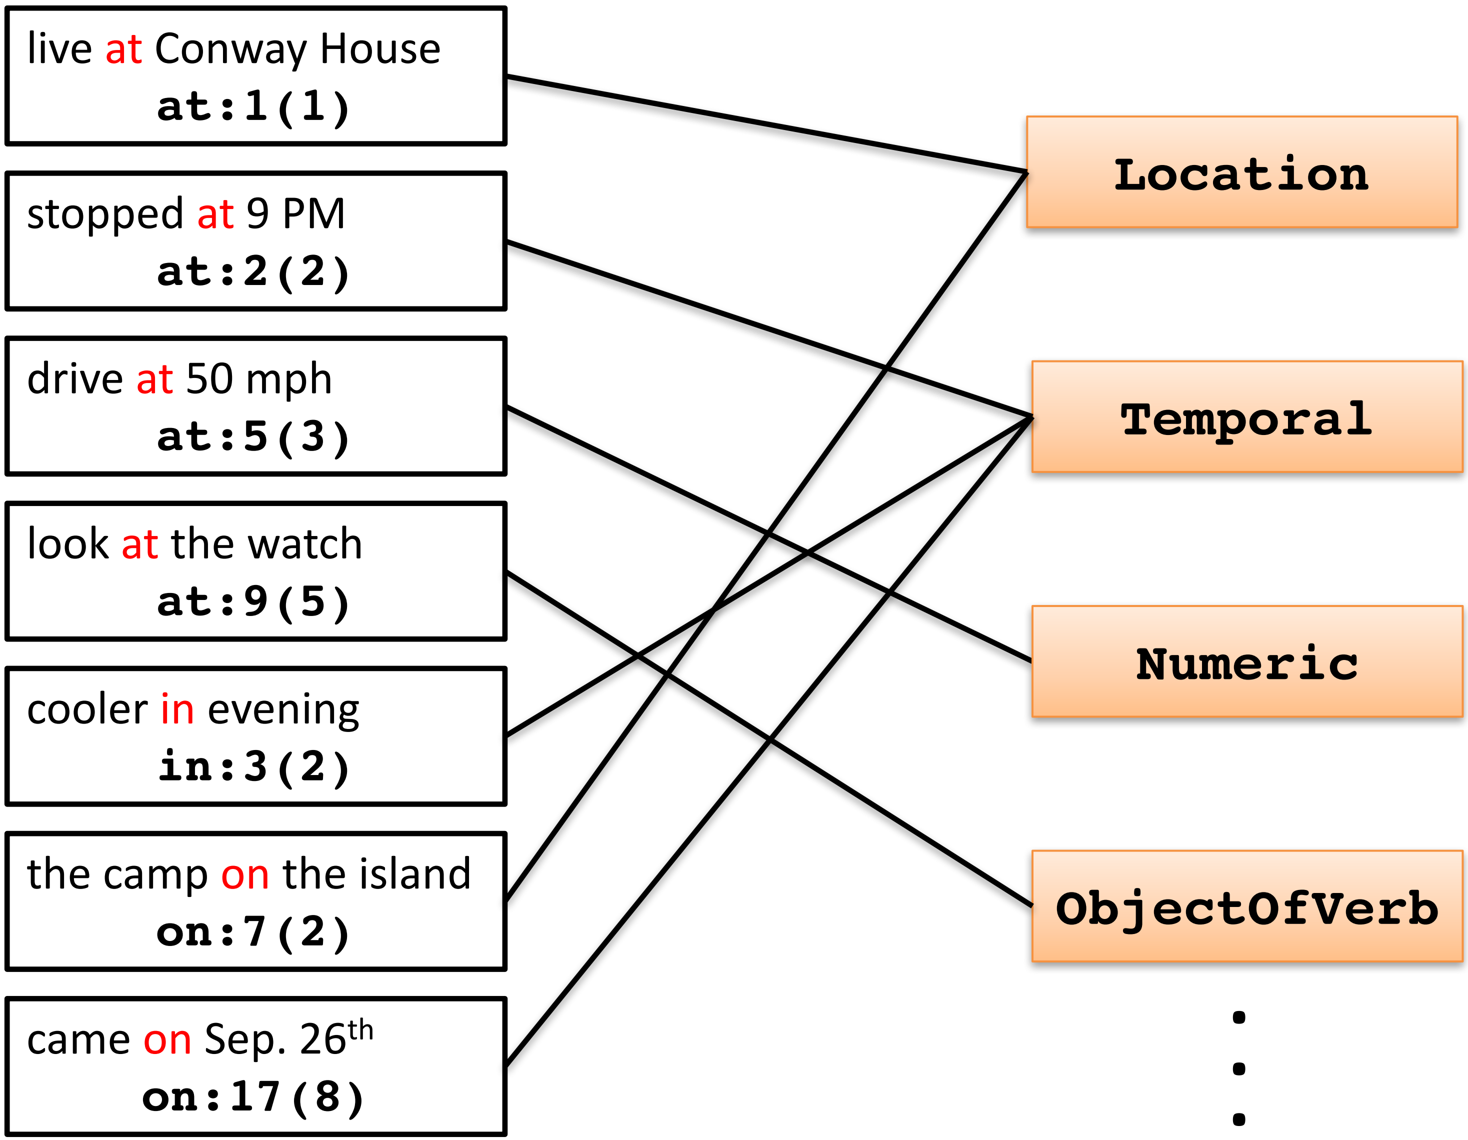

An inventory of preposition relations

- Labels that act as the predicate
  - Semantically related senses of prepositions merged
  - ~250 senses → 32 relation labels
- Word sense disambiguation data, re-labeled
  - SemEval 2007 shared task gives relation labeled data
    - ~16K training and ~8K test instances
    - 34 prepositions

List of preposition relations

Relation	Train	Test
Activity	63	39
Agent	367	159
Attribute	510	266
Beneficiary	205	105
Cause	591	289
Co-Participants	112	58
Destination	1054	526
Direction	909	441
EndState	188	104
Experiencer	116	54
Instrument	565	290
Location	3096	1531
Manner	457	245
MediumOfCommunication	57	30
Numeric	113	48
ObjectOfVerb	1801	882

Relation	Train	Test
Opponent/Contrast	233	131
Other	72	42
PartWhole	958	471
Participant/Accompanier	292	142
PhysicalSupport	399	202
Possessor	508	269
ProfessionalAspect	45	22
Purpose	261	113
Recipient	378	190
Separation	345	172
Source	740	357
Species	394	198
StartState	69	37
Temporal	331	157
Topic	886	462
Via	61	26

“zoo in NYC” $\longrightarrow$ `Location`(zoo, NYC)

Two Models for Predicting Preposition Relations

Structure of prepositions

Poor care led to her death **from** flu.

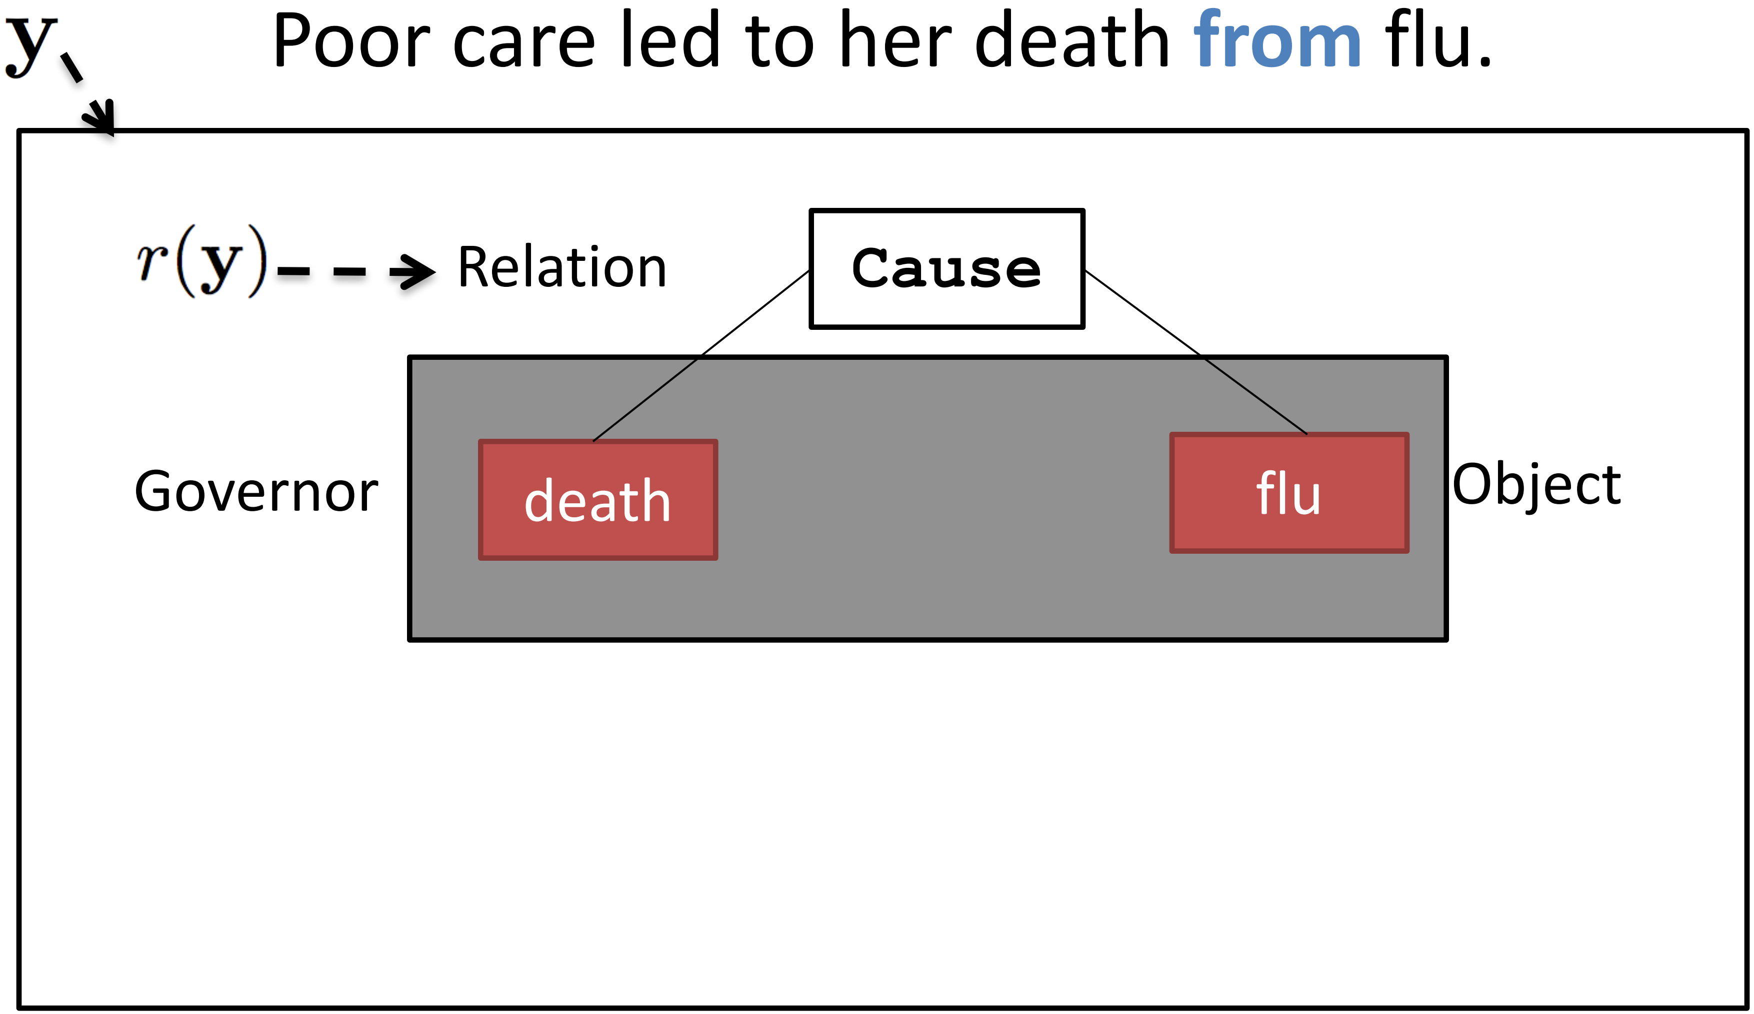

Parser not sufficient for arguments

Oracle: Selects the true argument if it is among the {parser's, heuristic} candidates

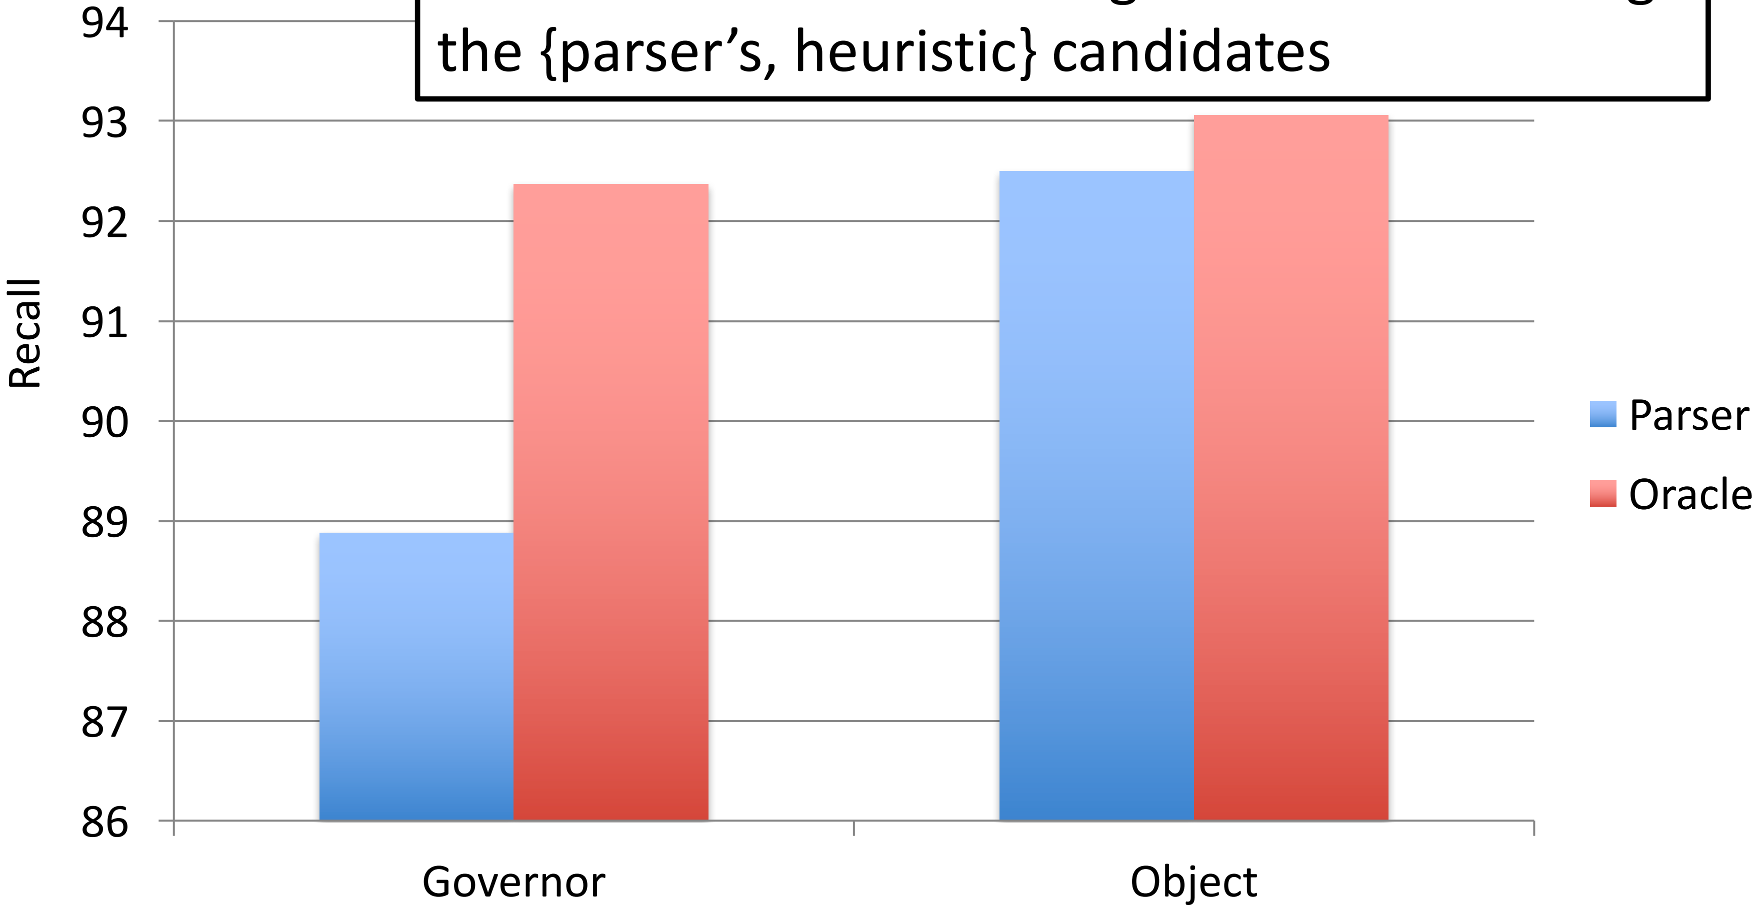

Easy-first dependency parser [Goldberg & Elhadad, 2010]

Relation depends on argument types

Poor care led to her death **from** flu.

`Cause`(death, flu)

Poor care led to her death **from** pneumonia.

How do we generalize the classifier to unseen arguments in the same "type"?

Why are types important?

- Goes beyond words
  - Abstract *flu and pneumonia* into the same group

- Some semantic relations hold only for certain types of entities

- Two notions of type
    - WordNet hypernyms
    - Distributional word clusters
  - Allow for multiple meanings and concept hierarchies

WordNet IS-A hierarchy

```
pneumonia
  => respiratory disease
    => disease
      => illness
        => ill health
          => pathological state
            => physical condition
              => condition
                => state
                  => attribute
                    => abstraction
                      => entity
```

Picking the right level in this hierarchy can generalize pneumonia and flu

More general, but less discrimniative

Picking incorrectly will over-generalize

Type #2: Word clusters

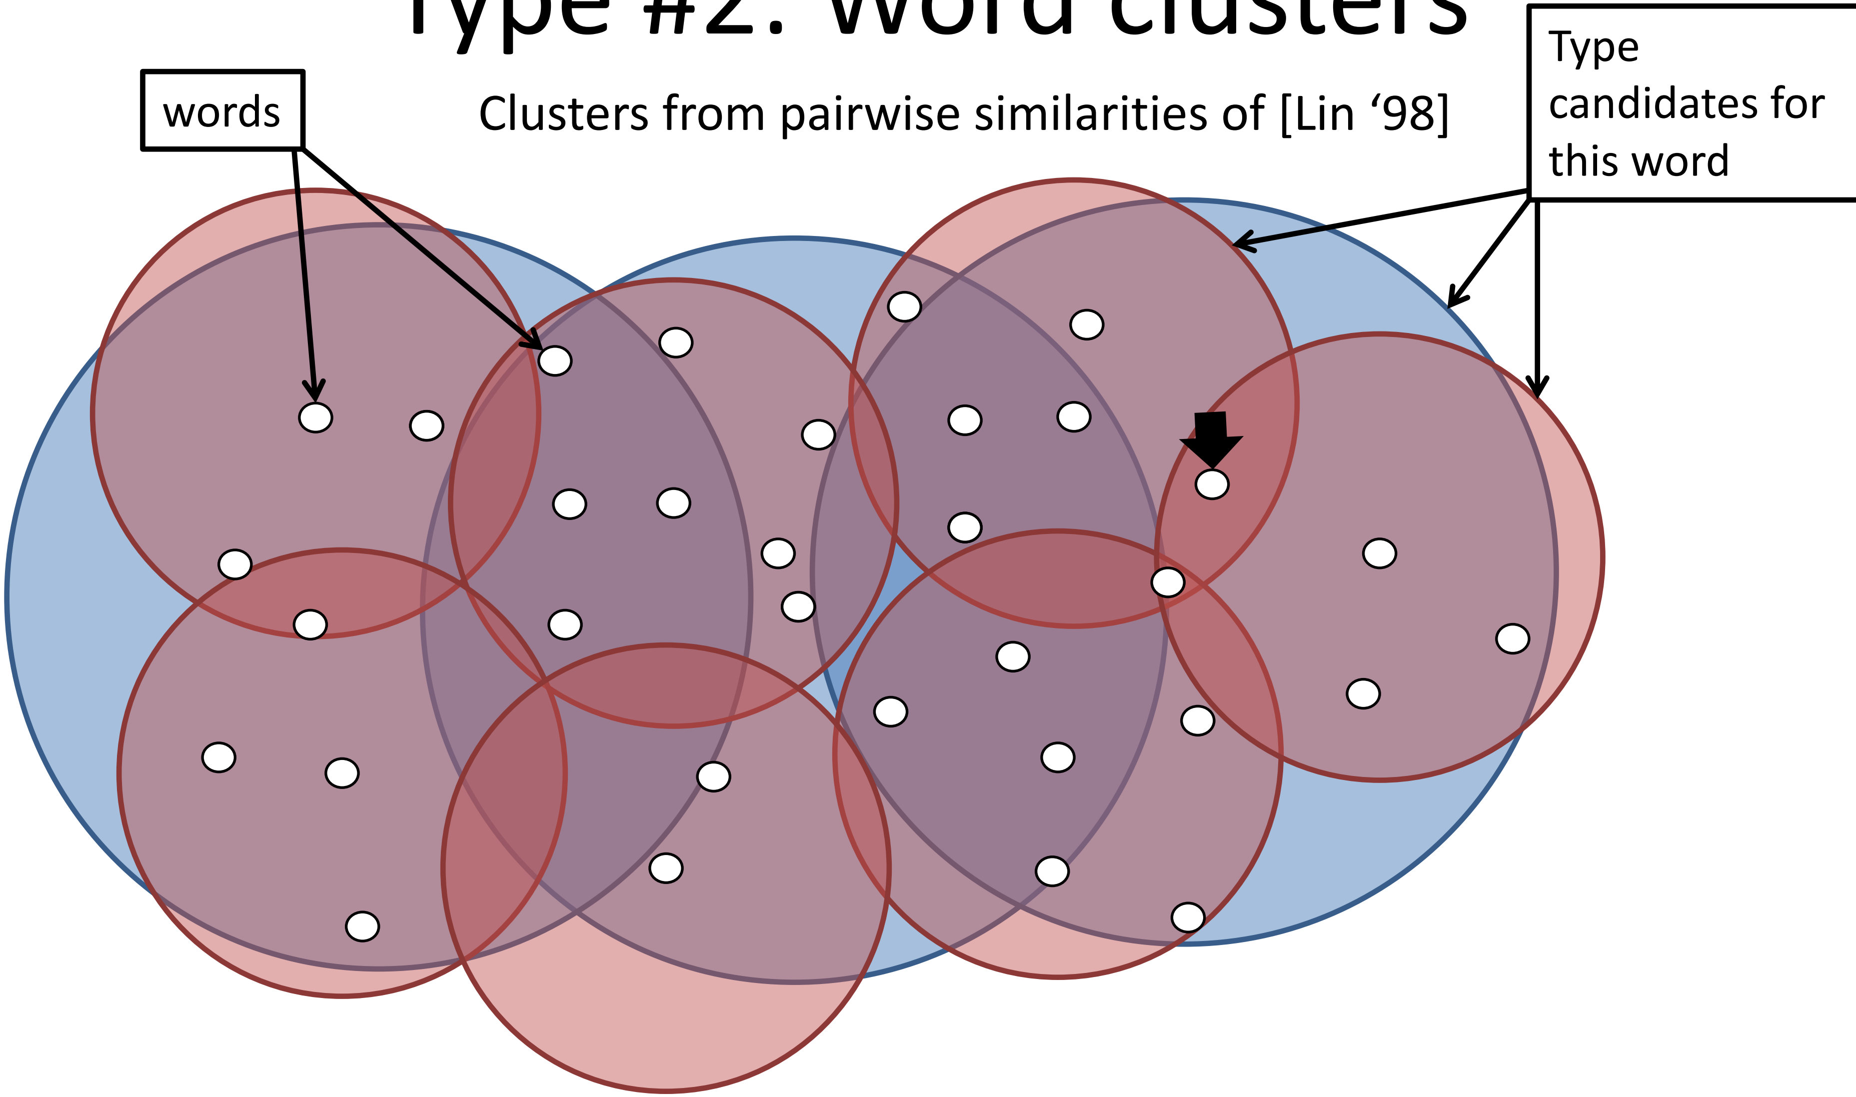

Cover all points with fixed radius clusters. Repeat with a different radii.

Example clusters

Jimmy Carter	metalwork	fox news channel	expert
Ronald Reagan	porcelain	NBC News	Wall Street analyst
richard nixon	handicraft	MSNBC	analyst
George Bush	jade	Fox News	economist
Lyndon Johnson	bronzeware	CNBC	telecommunications analyst
Richard M.	carving	CNNfn	strategist
Nixon	pottery	C-Span	media analyst
Gerald Ford	ceramic		
	earthenware		
	jewelry		
	stoneware		
	lacquerware		

Structure of prepositions

Poor care led to her death **from** flu.

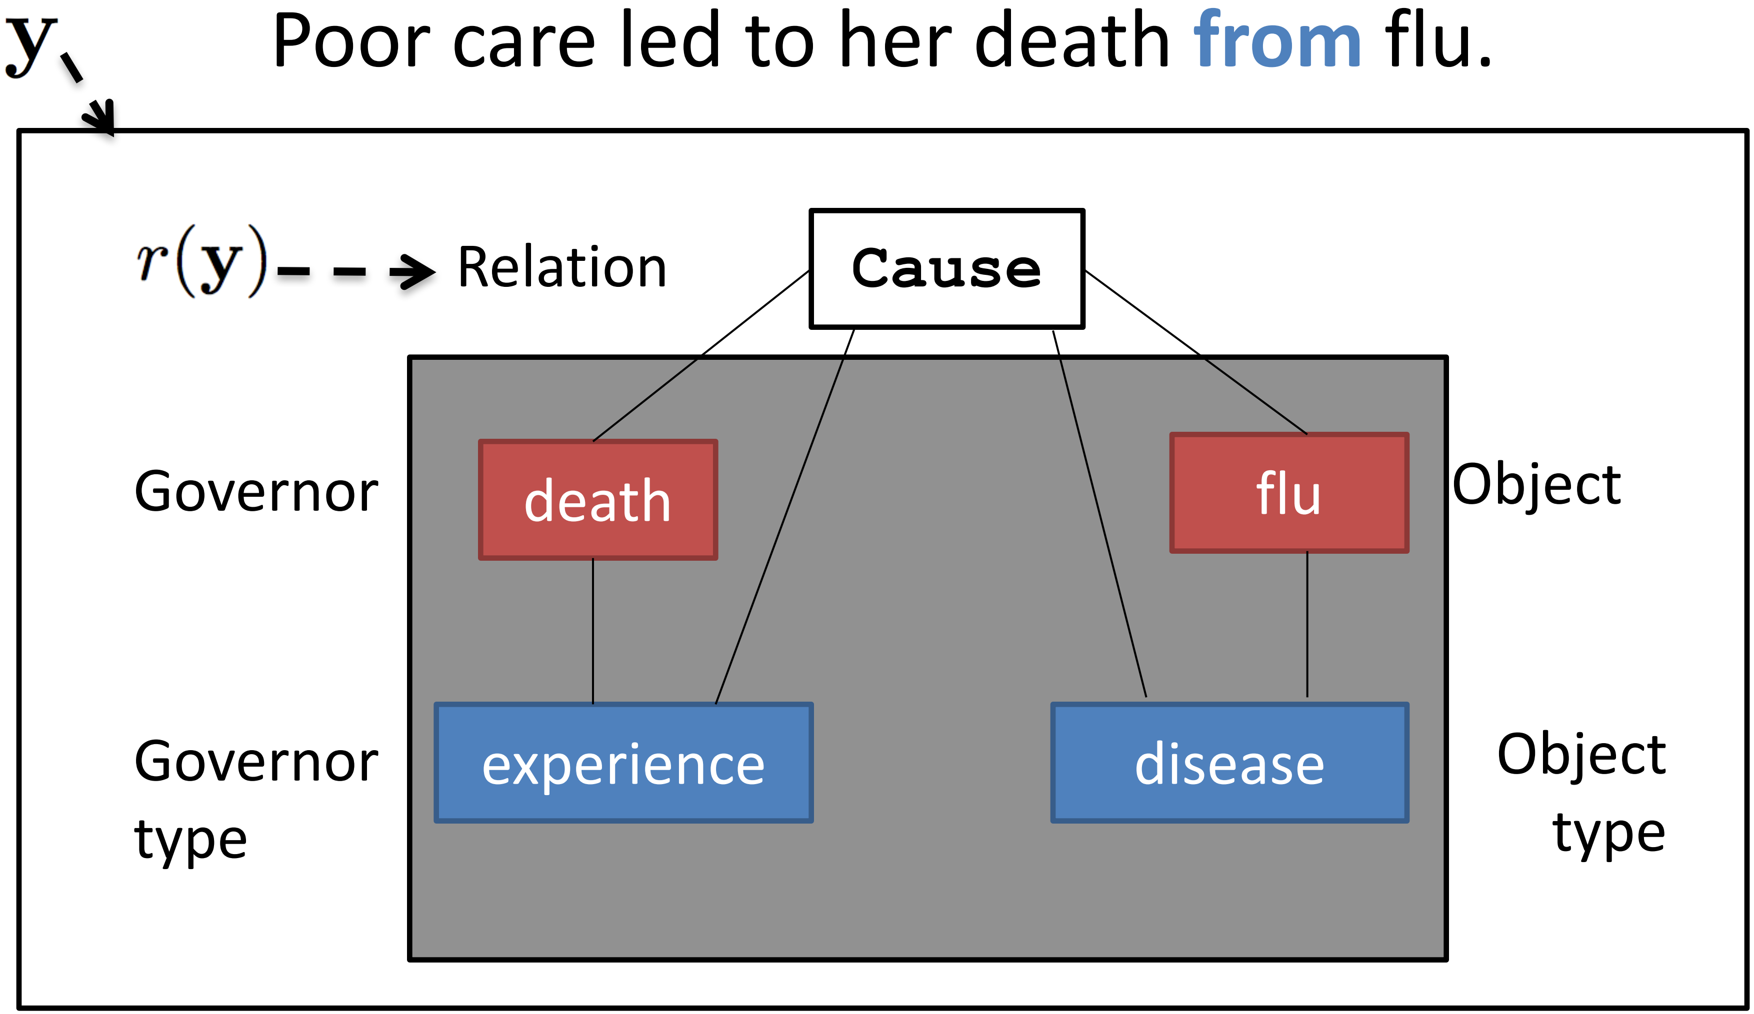

Governor and object base features

1. Word, POS, lemma and capitalization indicator
2. Conflated part-of-speech (Noun/Verb/Adjective/Adverb/Other)
3. Indicator for existence in WordNet,
4. WordNet synsets for the first and all senses,
5. WordNet lemma, lexicographer file names and part, member and substance holonyms,
6. Roget thesaurus divisions for the word,
7. The first and last two and three letters, and
8. Indicators for known affixes.

Two models

- Model 1
  - Predict only relation label: Multi-class
  - Use features from all possible governor and object candidates
    - Also types
- Model 2 uses features from the structure
  - Predict full structure: relation and arguments
    - Also types

Model 1: Predict relation label

Poor care led to her death **from** flu.

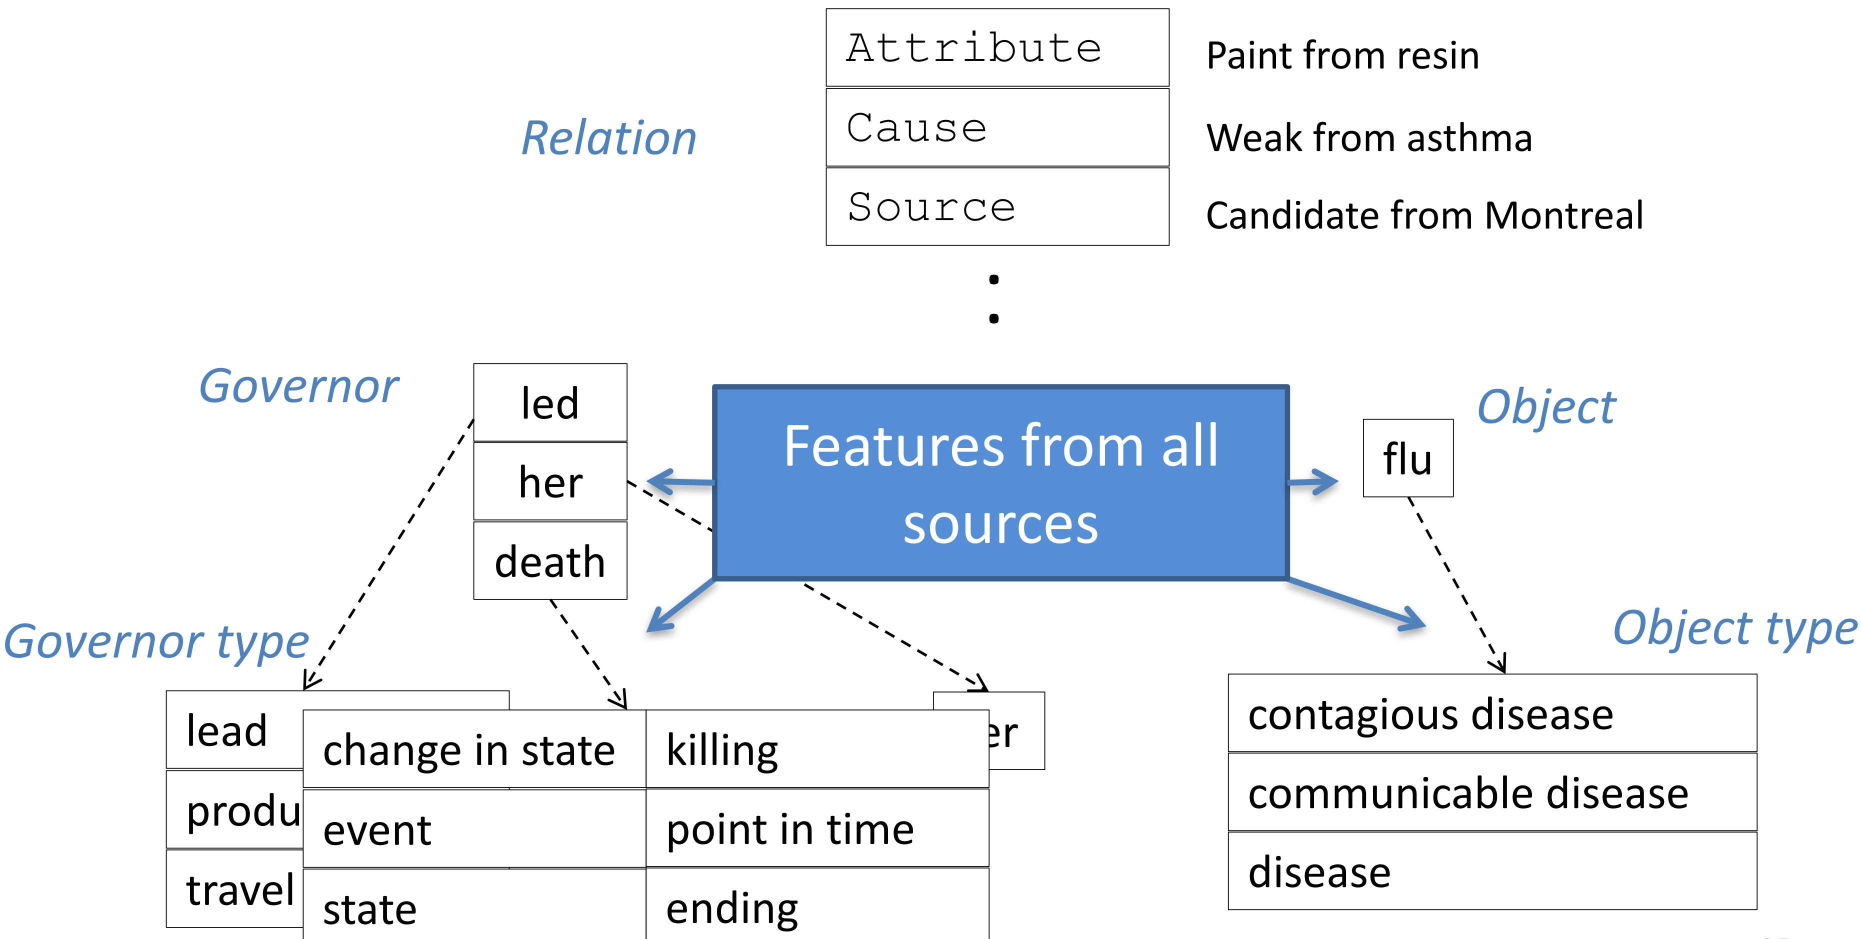

Model 2: Predict full structure

Poor care led to her death **from** flu.

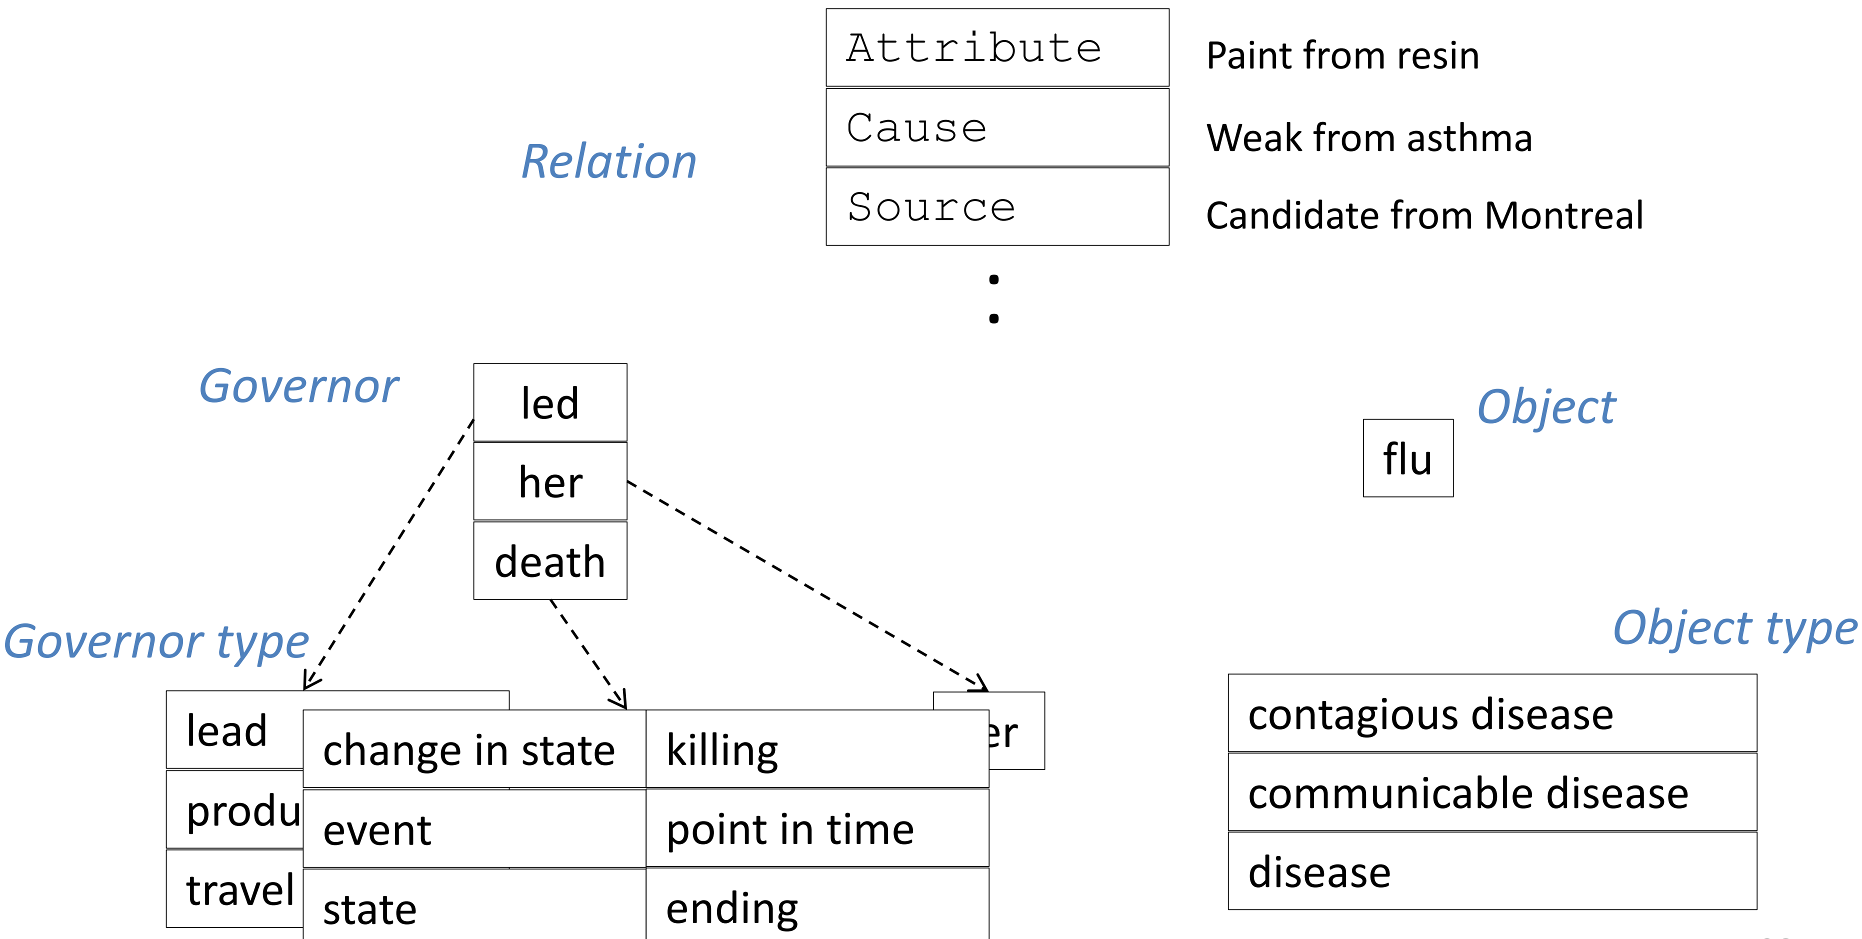

Structure of prepositions

$\mathbf{y}$

Poor care led to her death **from** flu.

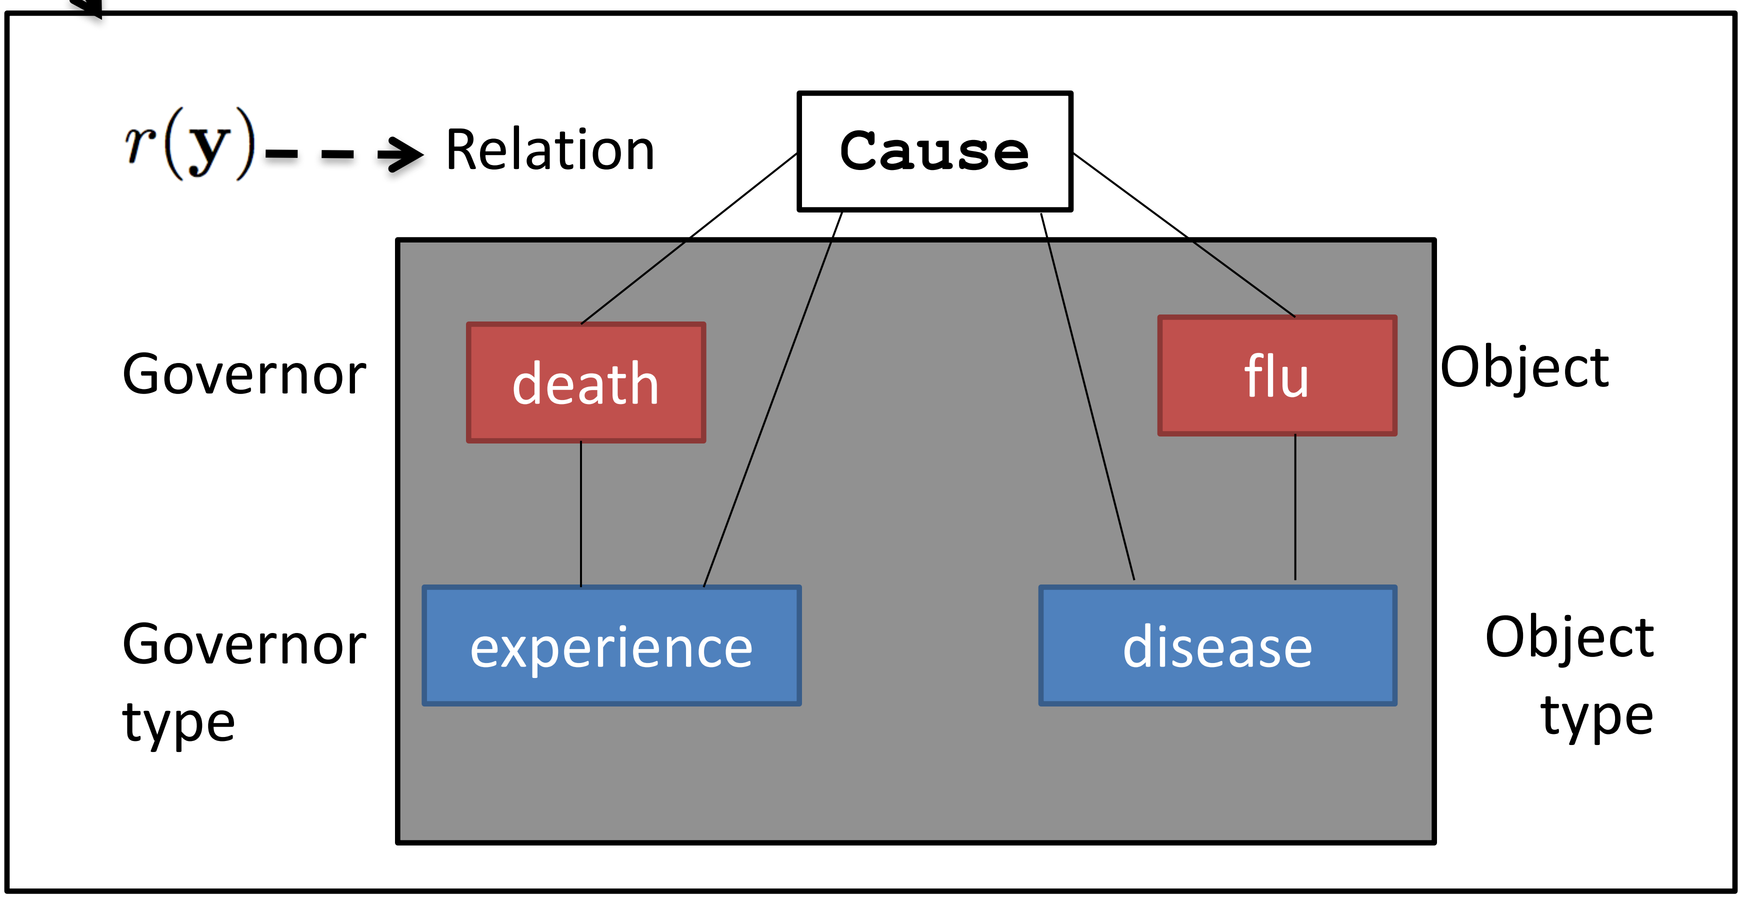

Learning Model 2: Latent inference

- ***Standard inference***: Find an assignment to the full structure

$$\max_{\mathbf{y}} \mathbf{w}^T \Phi(\mathbf{x}, \mathbf{y})$$

- ***Latent inference***: Given an example with annotated $r(\mathbf{y}^*)$

$$\begin{aligned} \max_{\mathbf{y}} \quad & \mathbf{w}^T \Phi(\mathbf{x}, \mathbf{y}) \\ \text{s.t.} \quad & r(\mathbf{y}^*) = r(\mathbf{y}) \end{aligned}$$

- *"Complete the structure given current model"*

Standard Supervised Learning

Structured SVM: [Tsochantaridis, et al 2004]

- Given a collection of labeled examples $\{(\mathbf{x}_i, \mathbf{y}_i)\}$ find a scoring function that minimizes loss

$$\min_{\mathbf{w},\xi} \frac{1}{2}\mathbf{w}^T\mathbf{w} + C\sum_i \xi_i$$

Score of annotated structure $\geq$ **Score of any other structure** $+$ **Penalty for predicting other structure**

- Several algorithms for learning.
  - We use the SVM solver from [Chang et al, 2010]

Learning Model 2

- Initialize weight vector using Model 1
- Repeat
  - Use latent inference with current weight to "complete" all missing pieces
  - Train with Structured SVM
    - During training, the learning algorithm is penalized more if it makes a mistake on $r(\mathbf{y})$

Generalization of Latent Structure SVM [Yu & Joachims '09]

Penalizing mistakes during learning

$$\min_{\mathbf{w},\xi} \frac{1}{2}\mathbf{w}^T\mathbf{w} + C\sum_i \xi_i$$

Score of annotated structure $\geq$ Score of any other structure $+ \Delta(\mathbf{y}, \mathbf{y}_i) - \xi_i$

- Penalty for making a mistake must not be the same for the labeled and inferred parts
  - Reduced penalty for mistakes made for anything except $r(\mathbf{y})$
  - Modified loss-augmented inference

Coherence with word sense

- Each sense corresponds to a relation

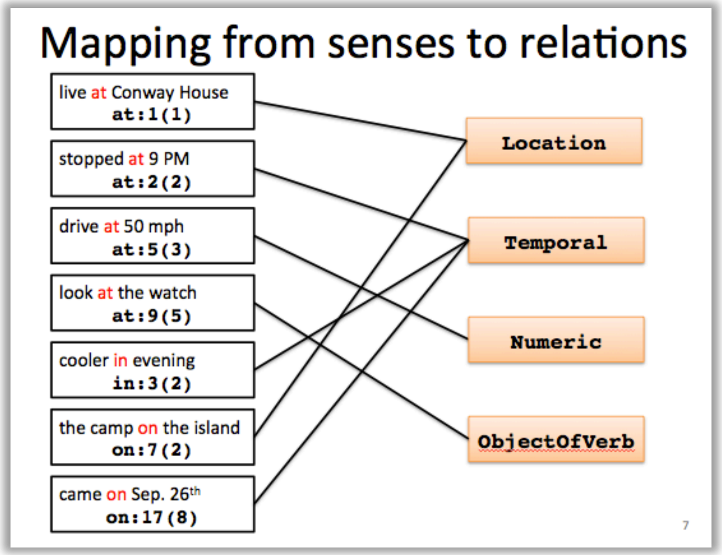

- Can explicitly enforce coherence
  – Model 1 or 2 for predicting relations
  – State-of-the-art preposition WSD [Hovy et al, 2010]
  – Coherence constraints from relation definition

Preposition Sense and Relations

Poor care led to her death from flu.

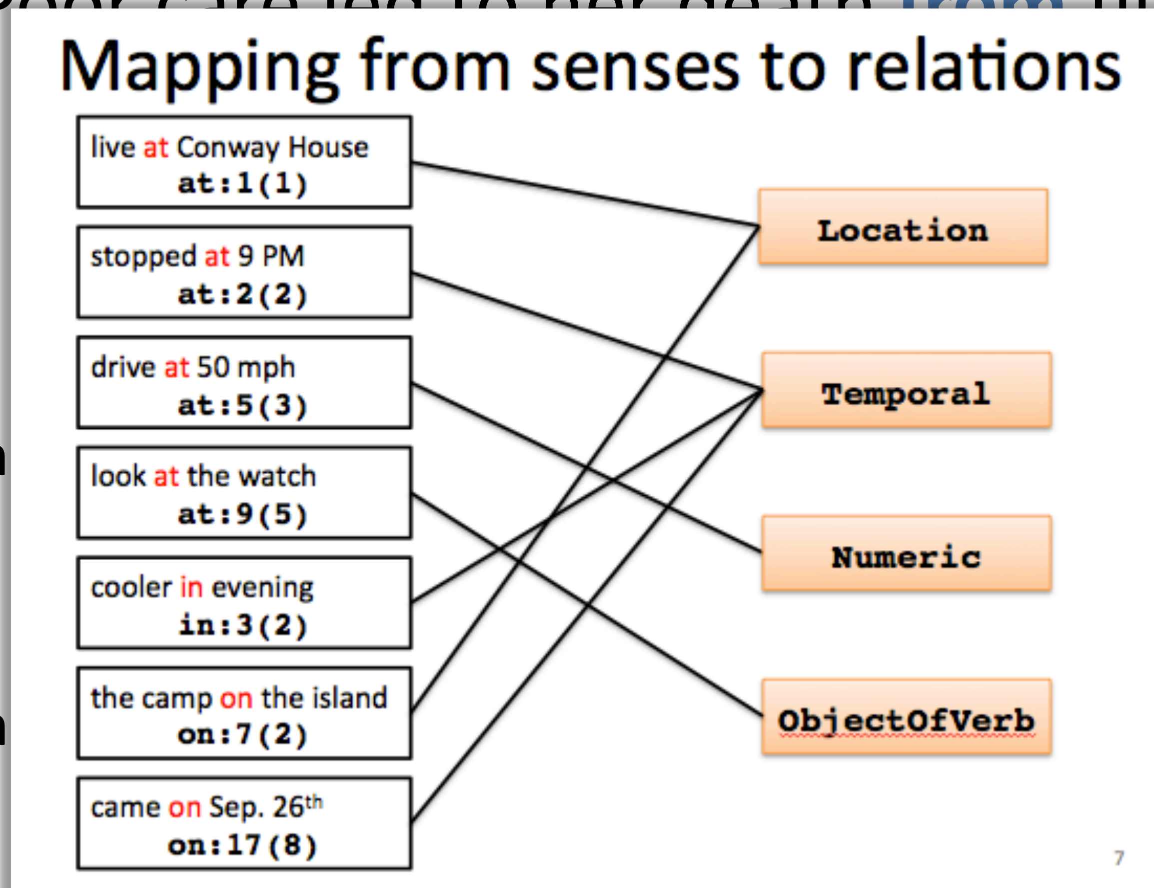

[Hovy et al, 2010]

n:12(9)

Govern

bject

Govern
type

Object
type

Experiments

Accuracy of relation labeling

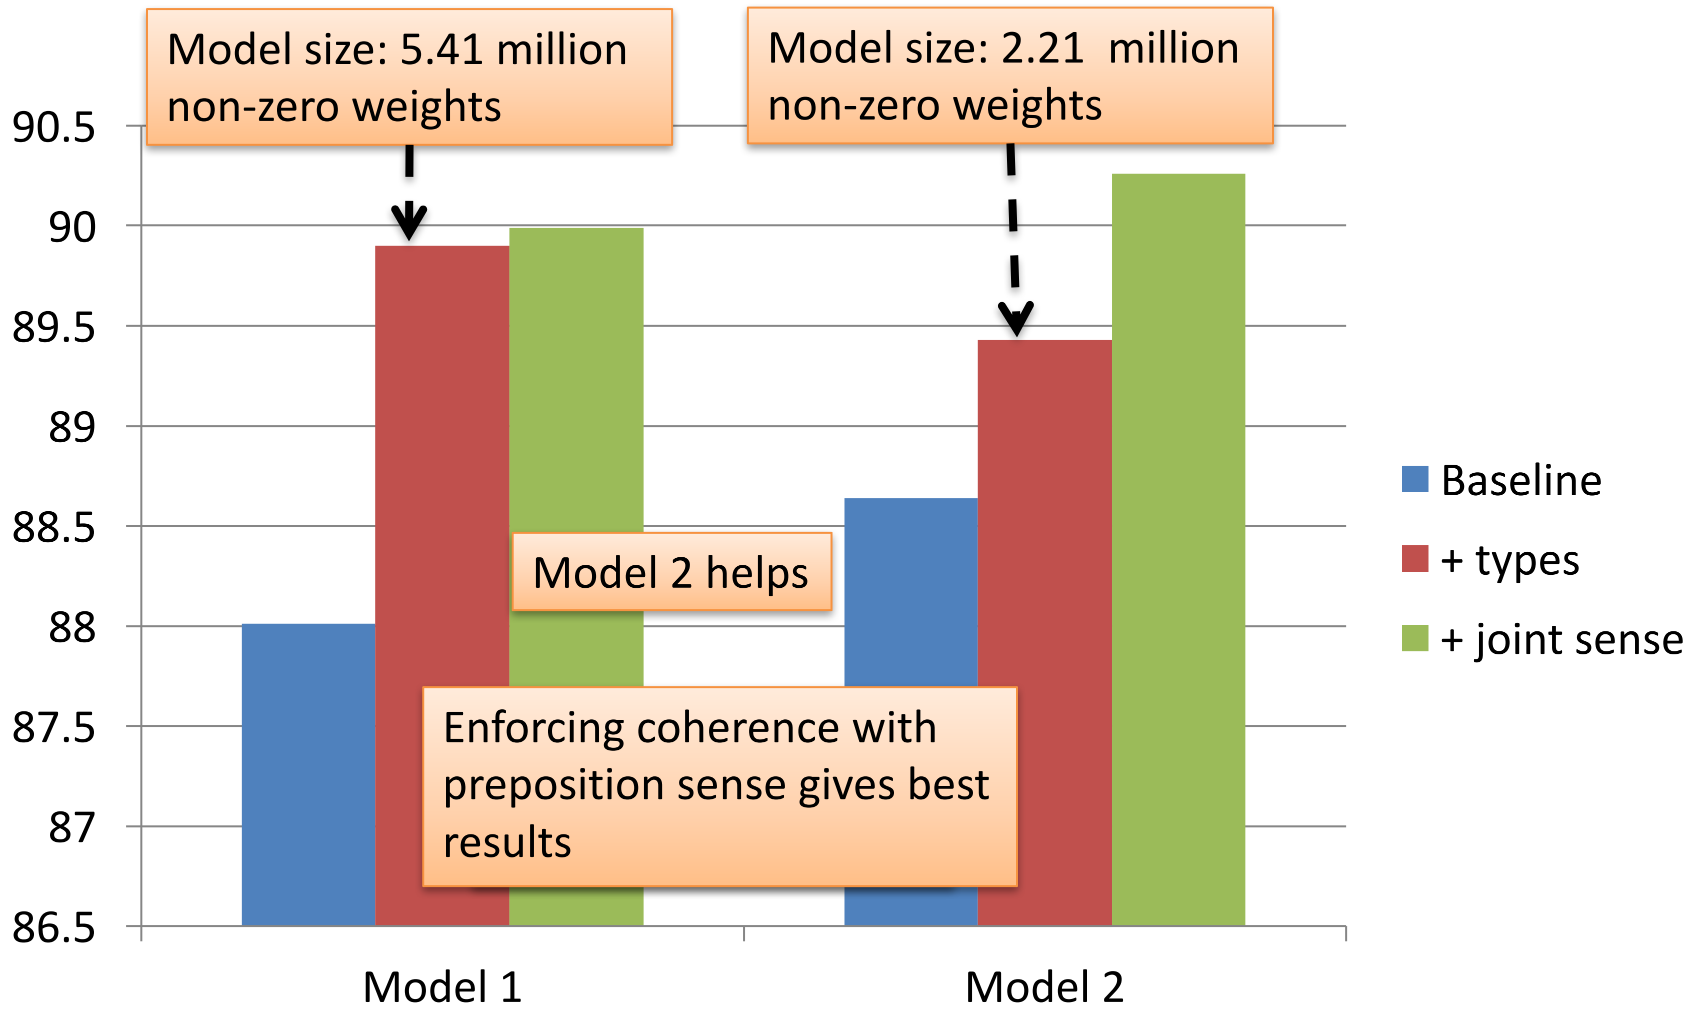

Analysis

- Sharing features across prepositions helps
  - Possible because the label space is the same
  - Not the case for WSD
- Parser or heuristics alone is not good enough
  - Combining them gives the best performance

Results summary

- Selecting a single assignment to arguments and types helps
- Enforcing coherence with preposition sense also helps
- Not fully believing the parser is useful
- Types are important

What do we have?

Input	Relation	Governor type	Object type
Died of pneumonia	`Cause`	Experience	Disease
Suffering from flu	`Cause`	Experience	Disease
Recovered from flu	`StartState`	Change	Disease

Governor, object and their types as a certificate for the choice of relation label

Summary

- Prepositions express a diverse set of relations
  - An ontology of preposition relations
  - Can enrich existing PropBank/FrameNet representation
- Models for predicting preposition relations
  - Arguments and types help

Data, word clusters, software available

Outline

1. Semantic Role Labeling
   - Looking beyond verbs and nominalizations
   - Preposition relations

2. Using a semantic representation
   - Looking beyond sentences
   - Reading comprehension

Jonathan Berant

Heather Chen

Chris Manning

What can we do with a semantic representation?

Modeling Biological Processes for Reading Comprehension

In review

Understanding text

On Feb. 20, 1962, a Marine Corps fighter pilot from small-town America stepped forward in response to the country's need. The astronaut was John Glenn, whom the author Tom Wolfe has called "the last true national hero America has ever had".
On Monday and Tuesday, Glenn will be honored with a dinner and a spaceflight forum at Ohio State University, home of the John Glenn School of Public Affairs.

Who called John Glenn "the last true national hero America has ever had"?

~~A: The astronaut~~
B: The author Tom Wolfe

Reading comprehension can be hard!

Water is split, providing a source of electrons and protons (hydrogen ions, H^+) and giving off O2 as a by-product. Light absorbed by chlorophyll drives a transfer of the electrons and hydrogen ions from water to an acceptor called $NADP^+$.

What does the splitting of water lead to?

A: Light absorption

B: Transfer of ions

Reading comprehension can be hard!

Enable

Cause

Water is split, providing a source of electrons and protons (hydrogen ions, H^+) and giving off O2 as a by-product. ***Light absorbed*** by chlorophyll drives a ***transfer of the electrons and hydrogen ions*** from water to an acceptor called $NADP^+$.

What does the splitting of water lead to?

A: Light absorption

B: Transfer of ions

Section Outline

1. A new reading comprehension task that requires reasoning over process structure

2. A new dataset `ProcessBank` consisting of descriptions of biological processes with
    1. Rich process structure annotated
    2. Multiple-choice questions

3. A new reading comprehension method
    1. Predict structure and treat it as a knowledge base
    2. Parse question as query to this KB

A new data set

Motivation

- Macro-reading vs. Micro-reading
  - Redundancy vs. reasoning

  [Etzioni et al., 2006; Carlson et al., 2010; Fader et al., 2011] vs [Richardson et al., 2013; Kushman et al., 2014]

- Non-factoid questions

- Modeling process structures
  - Events, entities and event-causality
  - Generally applicable to many situations

Non-factoid questions: AP Exam

In the development of a seedling, which of the following will be the last to occur?

A. Initiation of the breakdown of the food reserve
B. Initiation of cell division in the root meristem
C. Emergence of the root
D. Emergence and greening of the first true foliage leaves
E. Imbibition of water by the seed

Actual order: E A C B D

Interlude: Winograd questions

Q: *The town councillors refused to give the angry demonstrators a permit because they feared violence. Who feared violence?*
A1: The town councillors
A2: The angry demonstrators

Q: *The town councillors refused to give the angry demonstrators a permit because they advocated violence. Who advocated violence?*
A1: The town councillors
A2: The angry demonstrators

- "*Google-proof*" questions
- "*No cheap tricks*" questions [Levesque '13]

Making a difficult reading comprehension data set

- 200 paragraphs from the textbook *Biology*

[Scaria, et al. 2013]

  – Each paragraph is a single biological process

- What we want:
  1. Questions should test an understanding of inter-relations between events and entities
  2. Both answers should have similar lexical overlap to trump shallow approaches

Reading comprehension annotation

- 200 paragraphs: 585 questions
  - Created by a biologist
- Randomly chosen half validated by second annotator
  - 98.1% agreement
- Often requires complex reasoning about chains of event-event and event-entity relations

Examples of annotated questions

Dependencies between events/entities

Q: *What does the splitting of water lead to?*

A: Light absorption

B: Transfer of ions

Temporal ordering of events

Q: *What is the correct order of events?*

A: PDGF binds to tyrosine kinases, then cells divide, then wound healing

B: Cells divide, then PDGF binds to tyrosine kinases, then wound healing

True-False questions

Q: *Cdk associates with MPF to become cyclin*

A: True

B: False

The Structure of a Biological Process

Water is split, providing a source of electrons and protons (hydrogen ions, H^+) and giving off O2 as a by-product. ***Light absorbed*** by chlorophyll drives a ***transfer of the electrons and hydrogen ions*** from water to an acceptor called $NADP^+$.

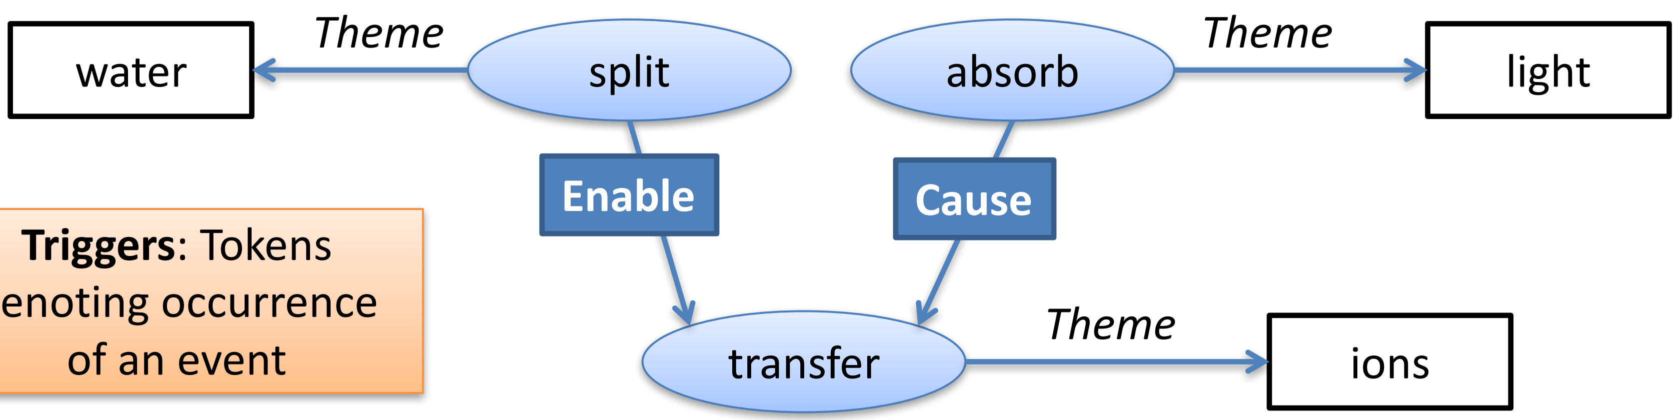

Triggers: Tokens denoting occurrence of an event

Arguments: *Agent, Theme, Source, Destination, Location, Result* and *Other*

E-E relations: *Cause, Enable, Prevent, Super, Same*
Also disjunctions

Process structure annotation

- Same 200 paragraphs

- Three annotators
  - Biologists, independent from QA annotator
    - Potentially conflicting

Triggers	Arguments	Relations
7.0	11.3	7.9

Water is split, providing a source of electrons and protons (hydrogen ions, H^+) and giving off O2 as a by-product. Light absorbed by chlorophyll drives a transfer of the electrons and hydrogen ions from water to an acceptor called $NADP^+$.

Process Structure

Answer: B

What does the splitting of water lead to?

A: Light absorption

B: Transfer of ions

Modeling Processes to Answer Questions

System overview

Process Structure Prediction

"... **Water is split**, providing a source of electrons and protons (hydrogen ions, H+) and giving off O2 as a by-product. **Light absorbed** by chlorophyll drives a **transfer of the electrons and hydrogen ions** from water to an acceptor called NADP+ ..."

Answering Question

Q What does the splitting of water lead to?
a Light absorption
b Transfer of ions

Question Parsing

Step 1: Structure prediction

- Event triggers
  - Classifier trained over all verbs, nouns, adjectives
- Arguments
  - Basically SRL
  - Argument identification + classification
- Event-event relations
  - Label all pairs of triggers (directed edges)
- Joint inference between arguments and relations
  - Framed as integer linear program

Constraints

1. Unique labels for all candidates
2. No overlapping arguments
3. Relation symmetry
   - R(a,b) $\rightarrow$ no R(b, a), except for SAME, where R(a,b) $\rightarrow$ R(b,a)
4. Maximum number of arguments per trigger
5. Maximum number of triggers per entity
6. Connectivity
7. Events that share arguments must be related
8. Unique SUPER parent

Structure Prediction Performance

- Learned with structured perceptron
- Performance

	Precision	Recall	F1
Triggers	78.8	73.9	74.8
Arguments	42.5	33.9	37.7
Relations	26.5	21.8	23.9

System overview

Process Structure Prediction

"... **Water is split**, providing a source of electrons and protons (hydrogen ions, H+) and giving off O2 as a by-product. **Light absorbed** by chlorophyll drives a **transfer of the electrons and hydrogen ions** from water to an acceptor called NADP+ ..."

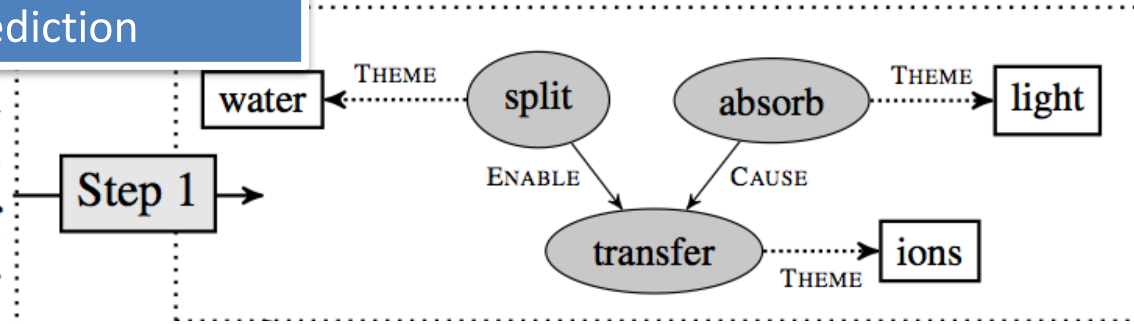

Q What does the splitting of water lead to?
a Light absorption
b Transfer of ions

Question Parsing

Step 2: Question parsing

- Task: Given a question and two answers, produce two queries
  - One for each answer

- Question + answer $\rightarrow$ path template in structure

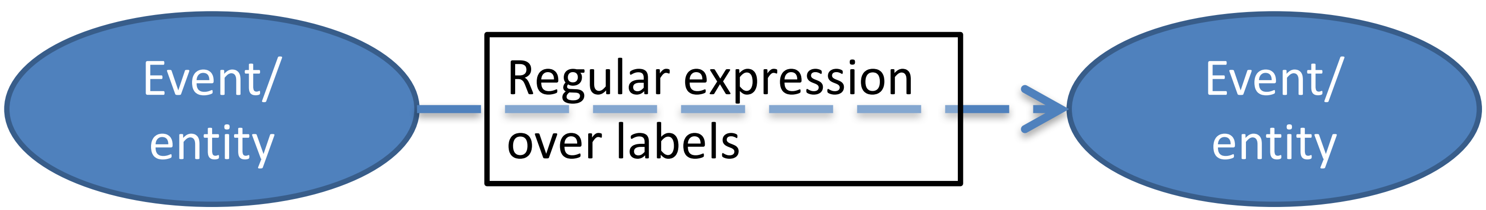

Deterministic query parsing

1. Align Q&A to structure
   – Recall that we limit lexical variablility
2. Identify source and target
   – If multiple candidates, try all combinations
3. Identify regular expressions
   – Chose from small set (~10) of regular expressions

Identifying regular expressions

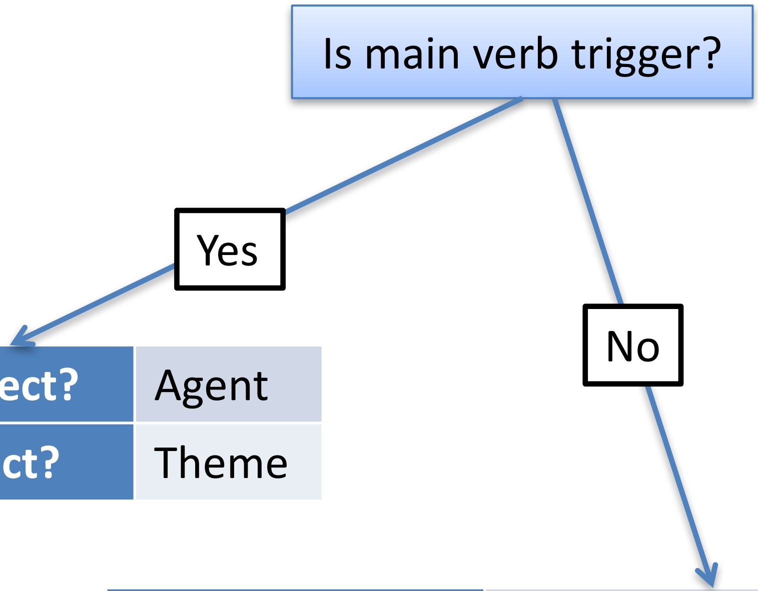

Wh-word subject?	Agent
Wh-word object?	Theme

default	(Enable|Super)+
DIRECT?	(Enable|Super)
PREVENT?	(Enable|Super)*Prevent(Enable|Super)*

System overview

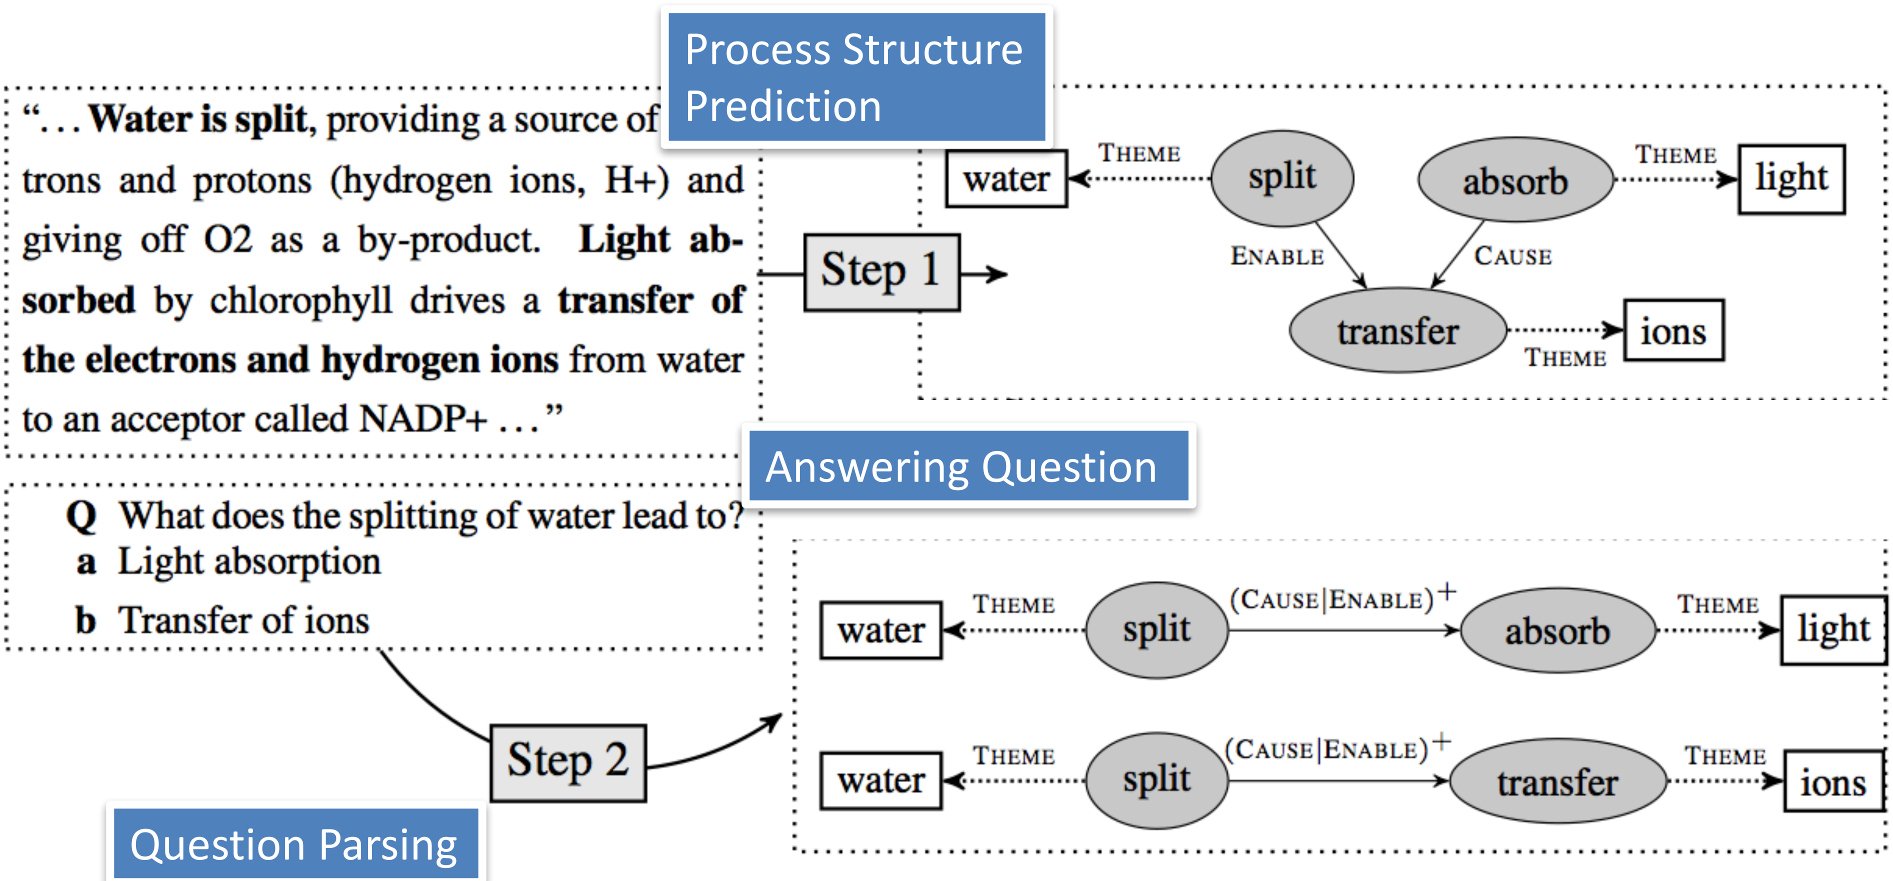

Step 3: Answering questions

- Given
  - Process structure
  - Two queries

- Answering algorithm
  1. Find matching path (valid proof)
  2. Else, find contradiction of causality (refutation)
  3. Else, find any path from source to target
  4. Else, back off to dependency based baseline

Experiments and Results

Experimental setup

- Train structure predictor on 150 paragraphs, test on 50 paragraphs

- Baselines
  - Bag of words
  - Text proximity of answer to question
  - Dependency based proximity

Question Answering Accuracy

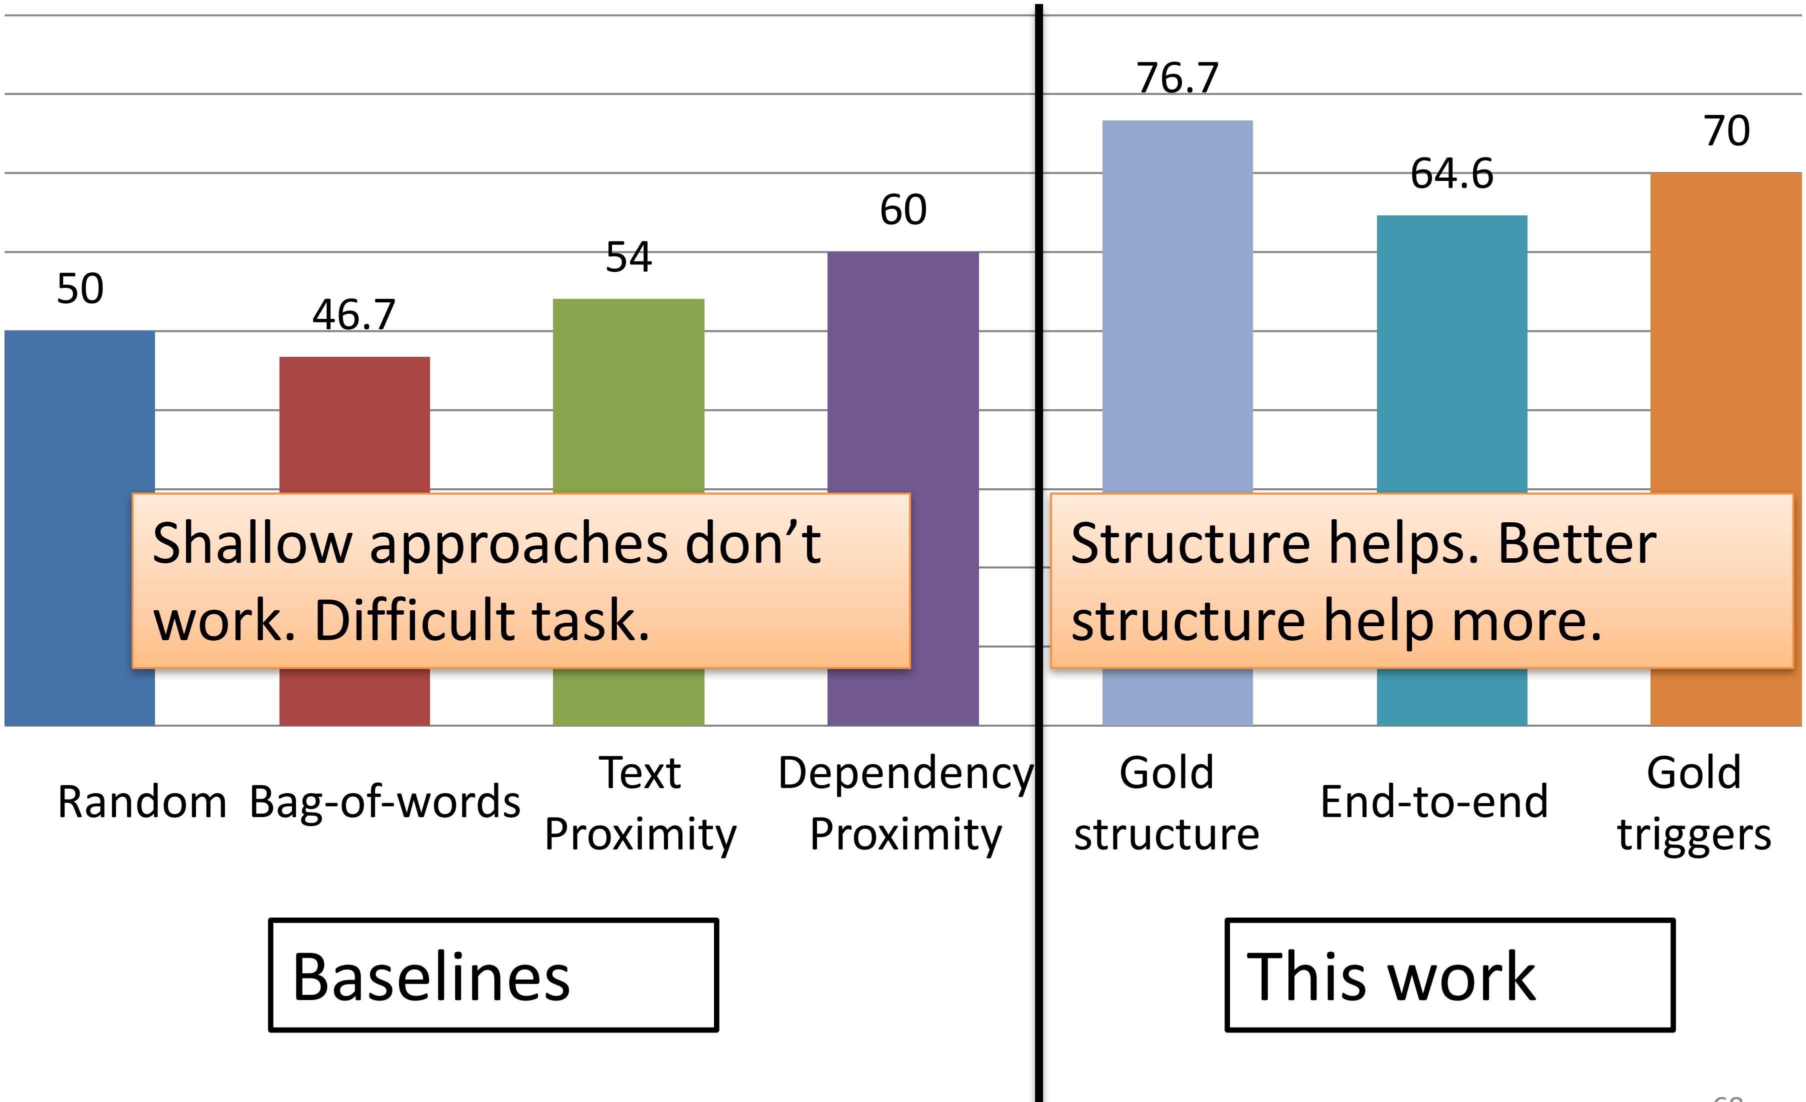

Analysis

- Errors with gold structure mainly because of
  - Alignment problems
  - Relevant part of structure not annotated
  - Not modeling entity coreference

- Errors with predicted structure
  - Mostly because of prediction errors
  - Sometimes, we get the answer even if the structure is incorrect in other places

Summary

1. A new reading comprehension task that requires reasoning over process structure

2. A new dataset `ProcessBank` consisting of descriptions of biological processes with
   - Rich process structure annotated
   - Multiple-choice questions

3. A new reading comprehension method
   - Predict structure and treat it as a knowledge base
   - Parse question as query to this KB

Final words

- Extending SRL
  - Data and system available
  - Guiding spirit is similar to AMR
  - Also, ongoing work here

- What can we do with something like SRL
  - Process structure (further extends beyond sentence boundaries)
  - Can tackle difficult reading comprehension setup
  - Data and system will be available

Questions, Comments?

Also, another talk tomorrow