Discourse Processing

Jessy Li
Dept. of Linguistics, UT Austin
jessy@austin.utexas.edu

With slides/material from Greg Durrett, Dan Jurafsky, Kenton Lee, Diane Litman, and Ani Nenkova

Computational Discourse

- Text is more than the sum of its individual sentences/utterances.
- Discourse processing: NLP beyond the sentence/ utterance boundary.
  - Monologue
  - Dialogue
    - Maybe multi-party
    - Maybe human-machine

What do you think of this text?

- “Consider, for example, the difference between passages (18.71) and (18.72). Almost certainly not. The reason is that these utterances, when juxtaposed, will not exhibit coherence. Do you have a discourse? Assume that you have collected an arbitrary set of well-formed and independently interpretable utterances, for instance, by randomly selecting one sentence from each of the previous chapters of this book.”

Or this?

- “Assume that you have collected an arbitrary set of well-formed and independently interpretable utterances, for instance, by randomly selecting one sentence from each of the previous chapters of this book. Do you have a discourse? Almost certainly not. The reason is that these utterances, when juxtaposed, will not exhibit coherence. Consider, for example, the difference between passages (18.71) and (18.72).”

What makes a text coherent?

- Discourse/Topic structure
  - In a coherent text the parts of the discourse exhibit a sensible ordering and/or hierarchical relationship
- Entity structure (“Focus”)
  - A coherent text is **about** some entity or entities, and the entity/entities is/are referred to in a structured way throughout the text.
- Rhetorical structure (“coherence/discourse relations”)
  - The elements in a coherent text are related via meaningful relations

What makes a text coherent?

- Discourse/Topic structure
  - In a coherent text the parts of the discourse exhibit a sensible ordering and/or hierarchical relationship
- Entity structure (“Focus”)
  - A coherent text is **about** some entity or entities, and the entity/ entities is/are referred to in a structured way throughout the text.
- Rhetorical structure (“coherence/discourse relations”)
  - The elements in a coherent text are related via meaningful relations

Outline

- Reference resolution
- Discourse relations

Outline

- **Reference resolution**
- Discourse relations

- Gracie: Oh yeah ... and then Mr. and Mrs. Jones were having matrimonial trouble, and my brother was hired to watch Mrs. Jones.
- George: Well, I imagine **she** was a very attractive woman.
- Gracie: **She** was, and my brother watched her day and night for six months.
- George: Well, what happened?
- Gracie: **She** finally got a divorce.
- George: Mrs. Jones?
- Gracie: No, my brother's wife.

Reference Resolution

Guiliani left **Bloomberg** to be mayor of a city with a big budget problem. **It**'s unclear how **he**'ll be able to handle **it** during **his** term.

- Process of associating Bloomberg/he/his with particular person and big budget problem/it with a concept
- Referring exprs.: Guilani, Bloomberg, he, it, his
- Presentational it, there: non-referential
- Referents: the person named Bloomberg, the concept of a big budget problem

Reference Resolution

Guiliani left **Bloomberg** to be mayor of a city with a big budget problem. **It**’s unclear how **he**’ll be able to handle **it** during **his** term.

- Co-referring referring expressions:
  - Bloomberg, he, his
- Antecedent: Bloomberg
- Anaphors: he, his

Discourse Model

- Needed to model reference because referring expressions (e.g. Guiliani, Bloomberg, he, it budget problem) encode information about beliefs about the referent
- When a referent is first mentioned in a discourse, a representation is evoked in the model
  - Information predicated of it is stored also in the model
  - On subsequent mention, it is accessed from the model

Types of Reference

- Entities, concepts, places, propositions, events, ...
- ***According to John, Bob bought Sue an Integra, and Sue bought Fred a Legend.***
  - But **that** turned out to be a lie. (a speech act)
  - But **that** was false. (proposition)
  - **That** struck me as a funny way to describe the situation. (manner of description)
  - **That** caused Sue to become rather poor. (event)
  - **That** caused them both to become rather poor. (combination of multiple events)

Reference Phenomena

- Indefinite NPs
  - **A homeless man** hit up Bloomberg for a dollar.
  - **Some homeless guy** hit up Bloomberg for a dollar.
- Definite NPs
  - **The poor fellow** only got a lecture.
- Demonstratives
  - **This homeless man** got a lecture but **that one** got carted off to jail.
- Names
  - **Tom** is afraid of **Jerry**.

Pronouns

- ***A large tiger escaped from the Central Park zoo chasing a tiny sparrow. It was recaptured by a brave policeman.***

- Referents of pronouns usually require some degree of ***salience*** in the discourse (as opposed to definite and indefinite NPs, e.g.)

- How do items become ***salient*** in discourse?

E: So you have the engine assembly finished. Now attach the rope. By the way, did you buy the gas can today?

A: Yes.

E: Did it cost much?

A: No.

E: OK, good. Have you got it attached yet?

What does "it" refer to? Why?

But things get complicated really fast…

Inferables

I almost bought an Acura Integra today,

but **a door** had **a dent** and **the engine** seemed noisy.

Mix the flour, butter, and water.

Knead **the dough** until smooth and shiny.

Discontinuous Sets

John has a St. Bernard and Mary has a Yorkie.

They arouse some comment when **they** walk **them** in the park.

Generics

I saw two Corgis and their seven puppies today.
They are the funniest dogs!

Coreference Resolution

Input document
A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized. Most of the deceased were killed in the crush as workers tried to flee the blaze in the four-story building.

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Coreference Resolution

Input document
A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized. Most of the deceased were killed in the crush as workers tried to flee the blaze in the four-story building.

Cluster #1	A fire in a Bangladeshi garment factory	the blaze in the four-story building

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Coreference Resolution

Input document
A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized. Most of the deceased were killed in the crush as workers tried to flee the blaze in the four-story building.

Cluster #1	A fire in a Bangladeshi garment factory	the blaze in the four-story building
Cluster #2	a Bangladeshi garment factory	the four-story building

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Coreference Resolution

Input document
A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized. Most of the deceased were killed in the crush as workers tried to flee the blaze in the four-story building.

Cluster #1	A fire in a Bangladeshi garment factory	the blaze in the four-story building
Cluster #2	a Bangladeshi garment factory	the four-story building
Cluster #3	at least 37 people	the deceased

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Two Subproblems

Input document
A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized. Most of the deceased were killed in the crush as workers tried to flee the blaze in the four-story building.

Mention detection

A fire in a Bangladeshi garment factory
at least 37 people
…
the four-story building

Mention clustering

Cluster #1	A fire in a Bangladeshi garment factory	the blaze in the four-story building
Cluster #2	a Bangladeshi garment factory	the four-story building
Cluster #3	at least 37 people	the deceased

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Rule-based pipeline

Input document

A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized.

Syntactic parser

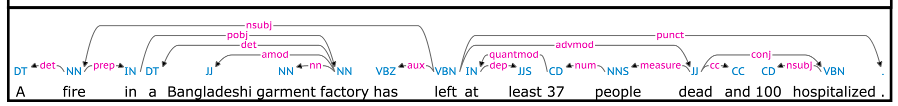

Hand-engineered rules

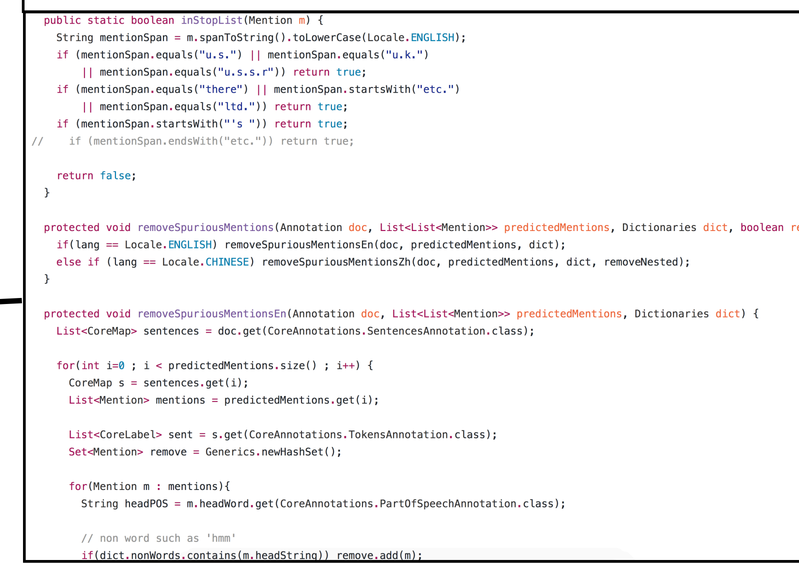

Candidate mentions
A fire in a Bangladeshi garment factory
garment
factory
at least 37 people dead and 100 hospitalized
…

Mention #1	Mention #2	Coreferent?
A fire in a Bangladeshi garment factory	garment	✓/✗
garment	factory	✓/✗
factory	at least 37 people dead and 100 hospitalized	✓/✗
…	…	✓/✗

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Easy Victories and Uphill Battles in Coreference Resolution

Greg Durrett and Dan Klein

UC Berkeley

Basic Architecture

Voters agree when they are given a chance to decide if they ...

PROPERS: NER chunks
PRONOUNS: PRP, PRP$
NOMINALS: all other maximal NPs

[Voters] agree when [they] are given a [chance] to decide if [they] ...

[Voters]$_1$ agree when **[they]**$_1$ are given a chance to decide if **[they]**$_1$ ...

[Voters]$_1$ agree when [they]$_1$ are given a [chance]$_2$ to decide if [they]$_1$ ...

Mention-Ranking Architecture

$$Pr(A_i = a|x) \propto \exp(w^\top f(i,a,x))$$

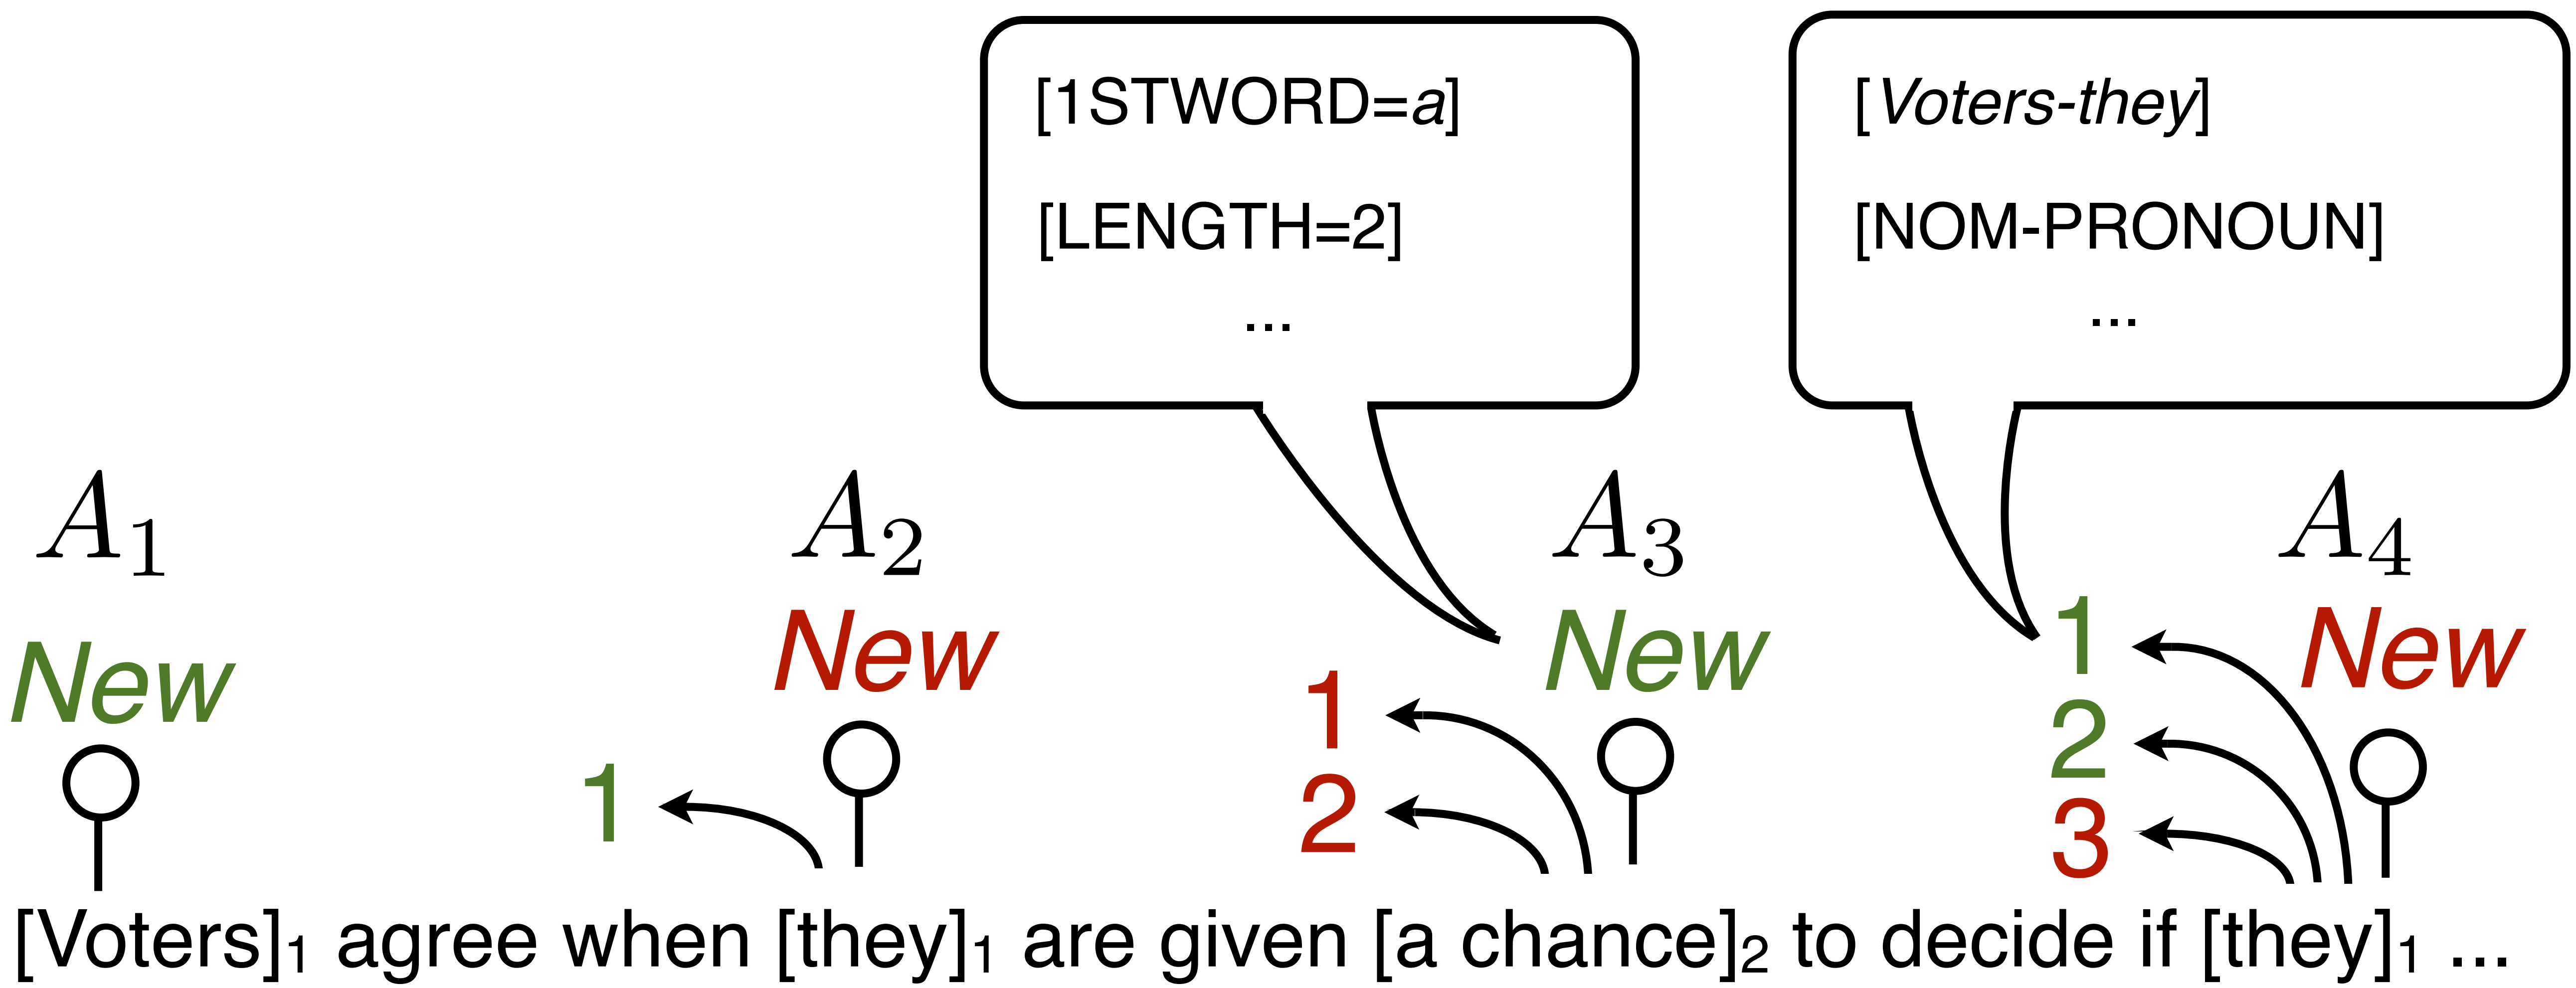

Denis and Baldridge (2008), Durrett et al. (2013)

Learning

$$A_i \in \{1, 2, \ldots, i-1, New\}$$

$$Pr(A_i = a|x) \propto \exp(w^\top f(i, a, x))$$

Maximize: $\sum_{i=1}^{n} [\log Pr(A_i \in \{\text{correct}\}|x)] + \|w\|_1$

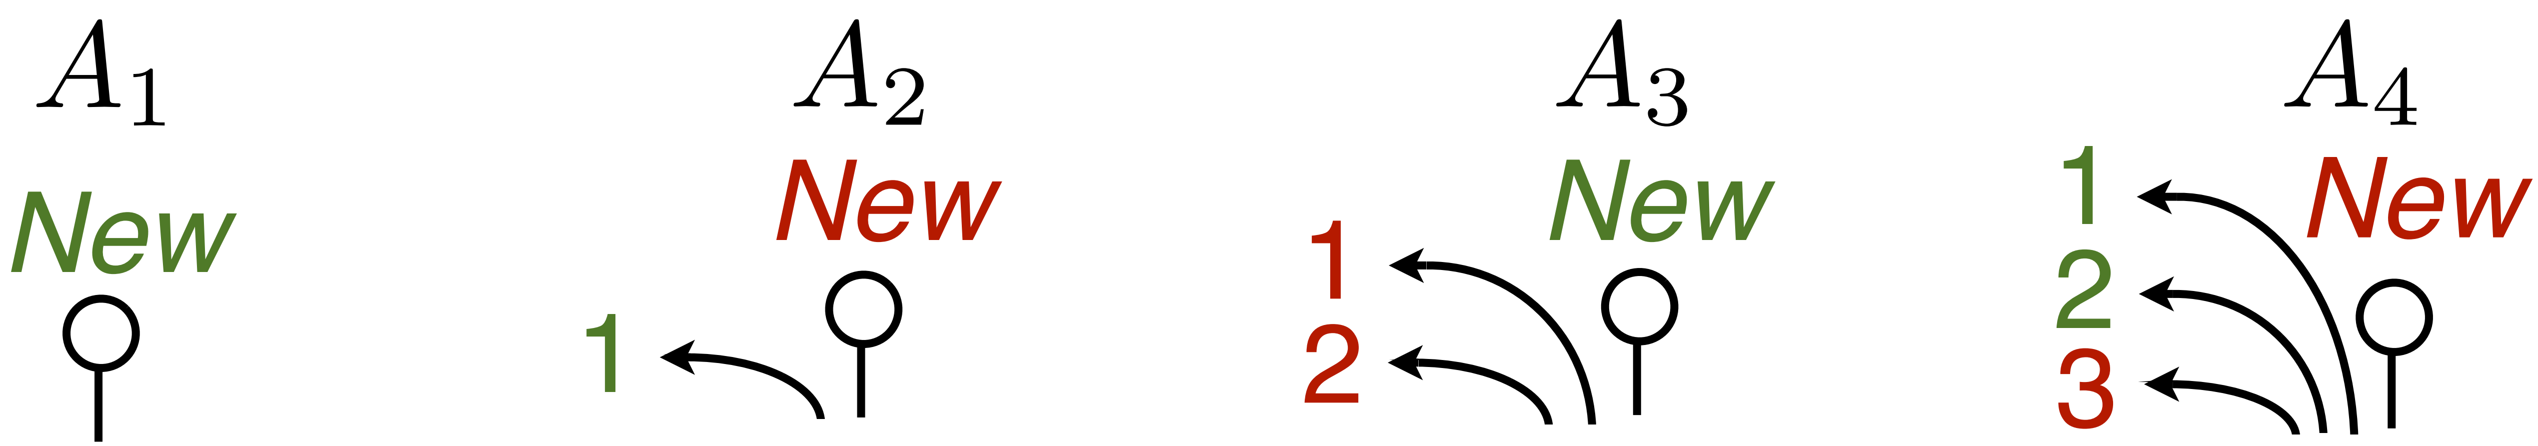

[Voters]$_1$ agree when [they]$_1$ are given a [chance]$_2$ to decide if [they]$_1$ ...

Gimpel and Smith (2010), Durrett et al. (2013)

Let's think about rules…

- Number agreement
  - ***John's parents like opera.***
  - *John hates **it.***
  - *John hates **them**.*
- Person and case agreement
  - Nominative: I, we, you, he, she, they
  - Accusative: me, us, you, him, her, them
  - Genitive: my, our, your, his, her, their
  - ***George and Edward brought bread and cheese.***
  - ***They** shared **them.***

Let’s think about rules…

- Gender agreement
  - **John has a Porsche. *He/it/she* is attractive.**
- Syntactic constraints: binding theory
  - **John bought *himself* a new Volvo.** (himself = John)
  - **John bought *him* a new Volvo.** (him = not John)
- Selectional restrictions
  - **John left his plane in the hangar.**
  - **He had flown *it* from Memphis this morning.**

Centering

SUBJECT–SUBJECT

OBJECT–SUBJECT

[*Barack Obama*]$_1$ *met with* [*David Cameron*]$_2$. [*He*]$_1$ *said ...*

Grosz al. (1995)

Slide from Durrett and Klein, Easy Victories and Uphill Battles in Coreference Resolution, EMNLP 2013

Data and Evaluation

- Ontonotes (5.0)
  - 2.9 million words, 3.5K documents
  - Layered annoations
  - News, broadcast news, broadcast conversations, blogs, Old and New Testaments
- CoNLL 2011 & 2012 shared tasks
- Precision, recall and F on pairs of mentions (MUC), links (B-cubed), and entity (CEAF)

Coreference Results

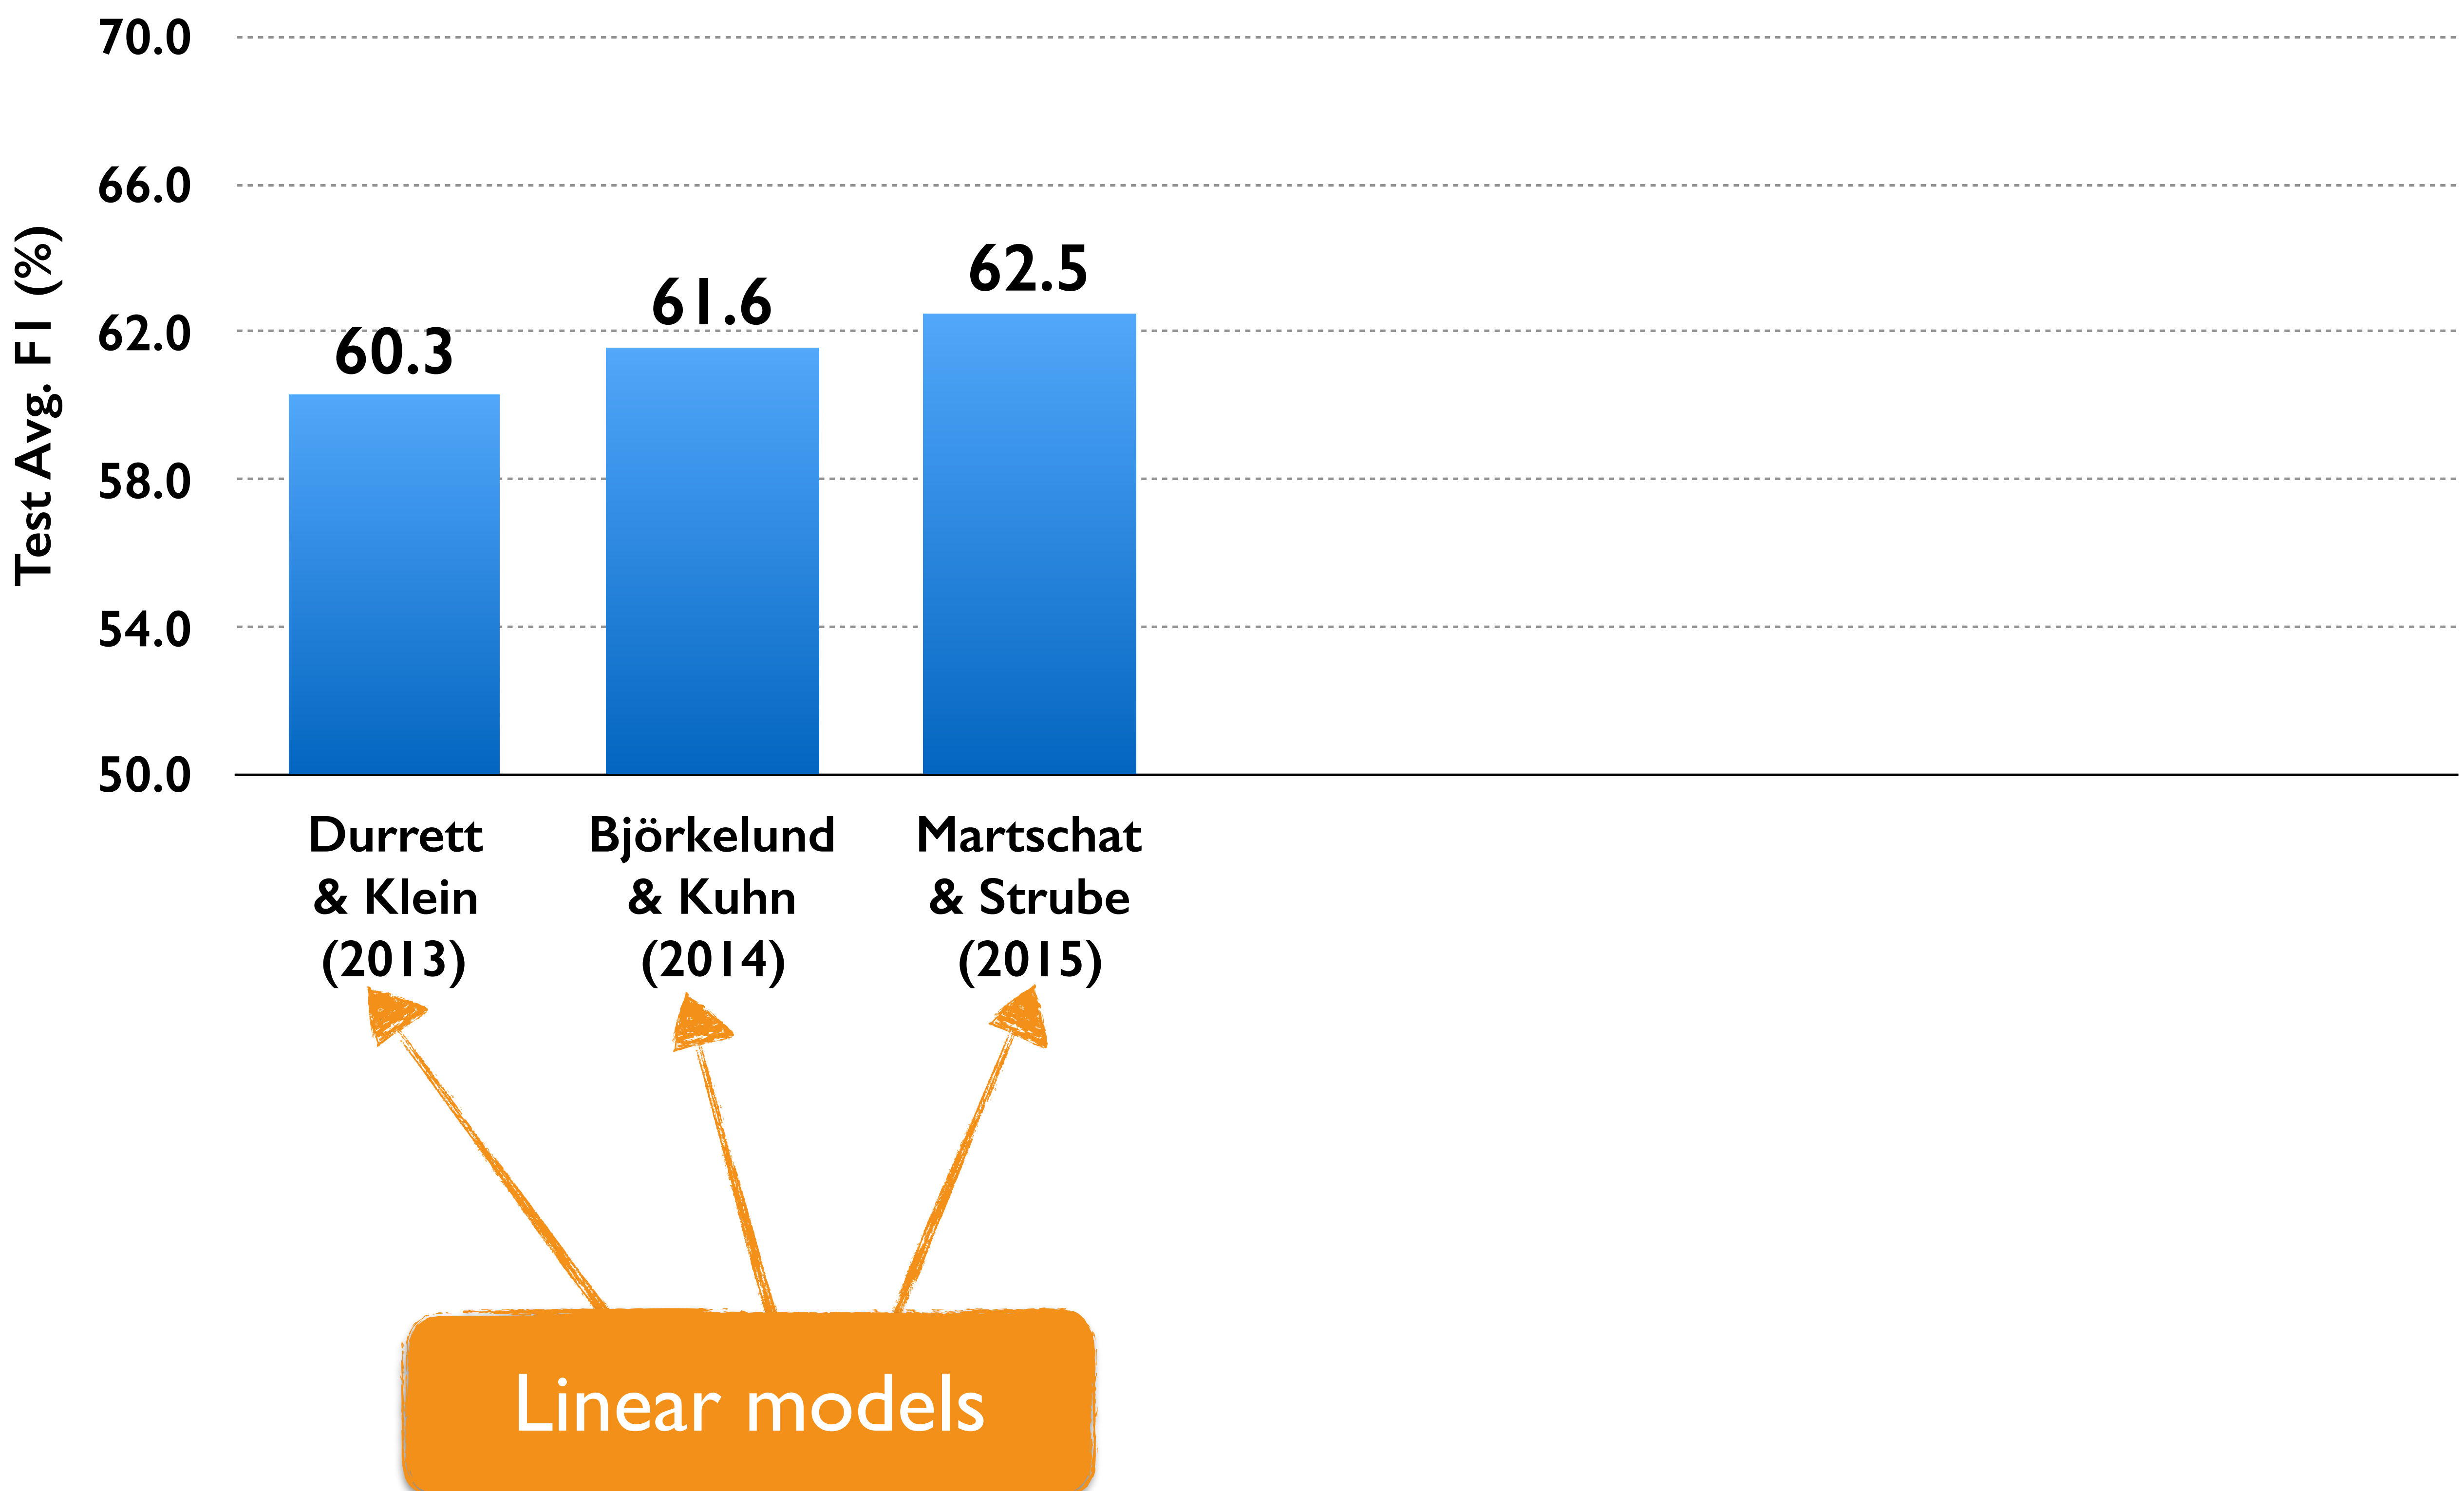

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Coreference Results

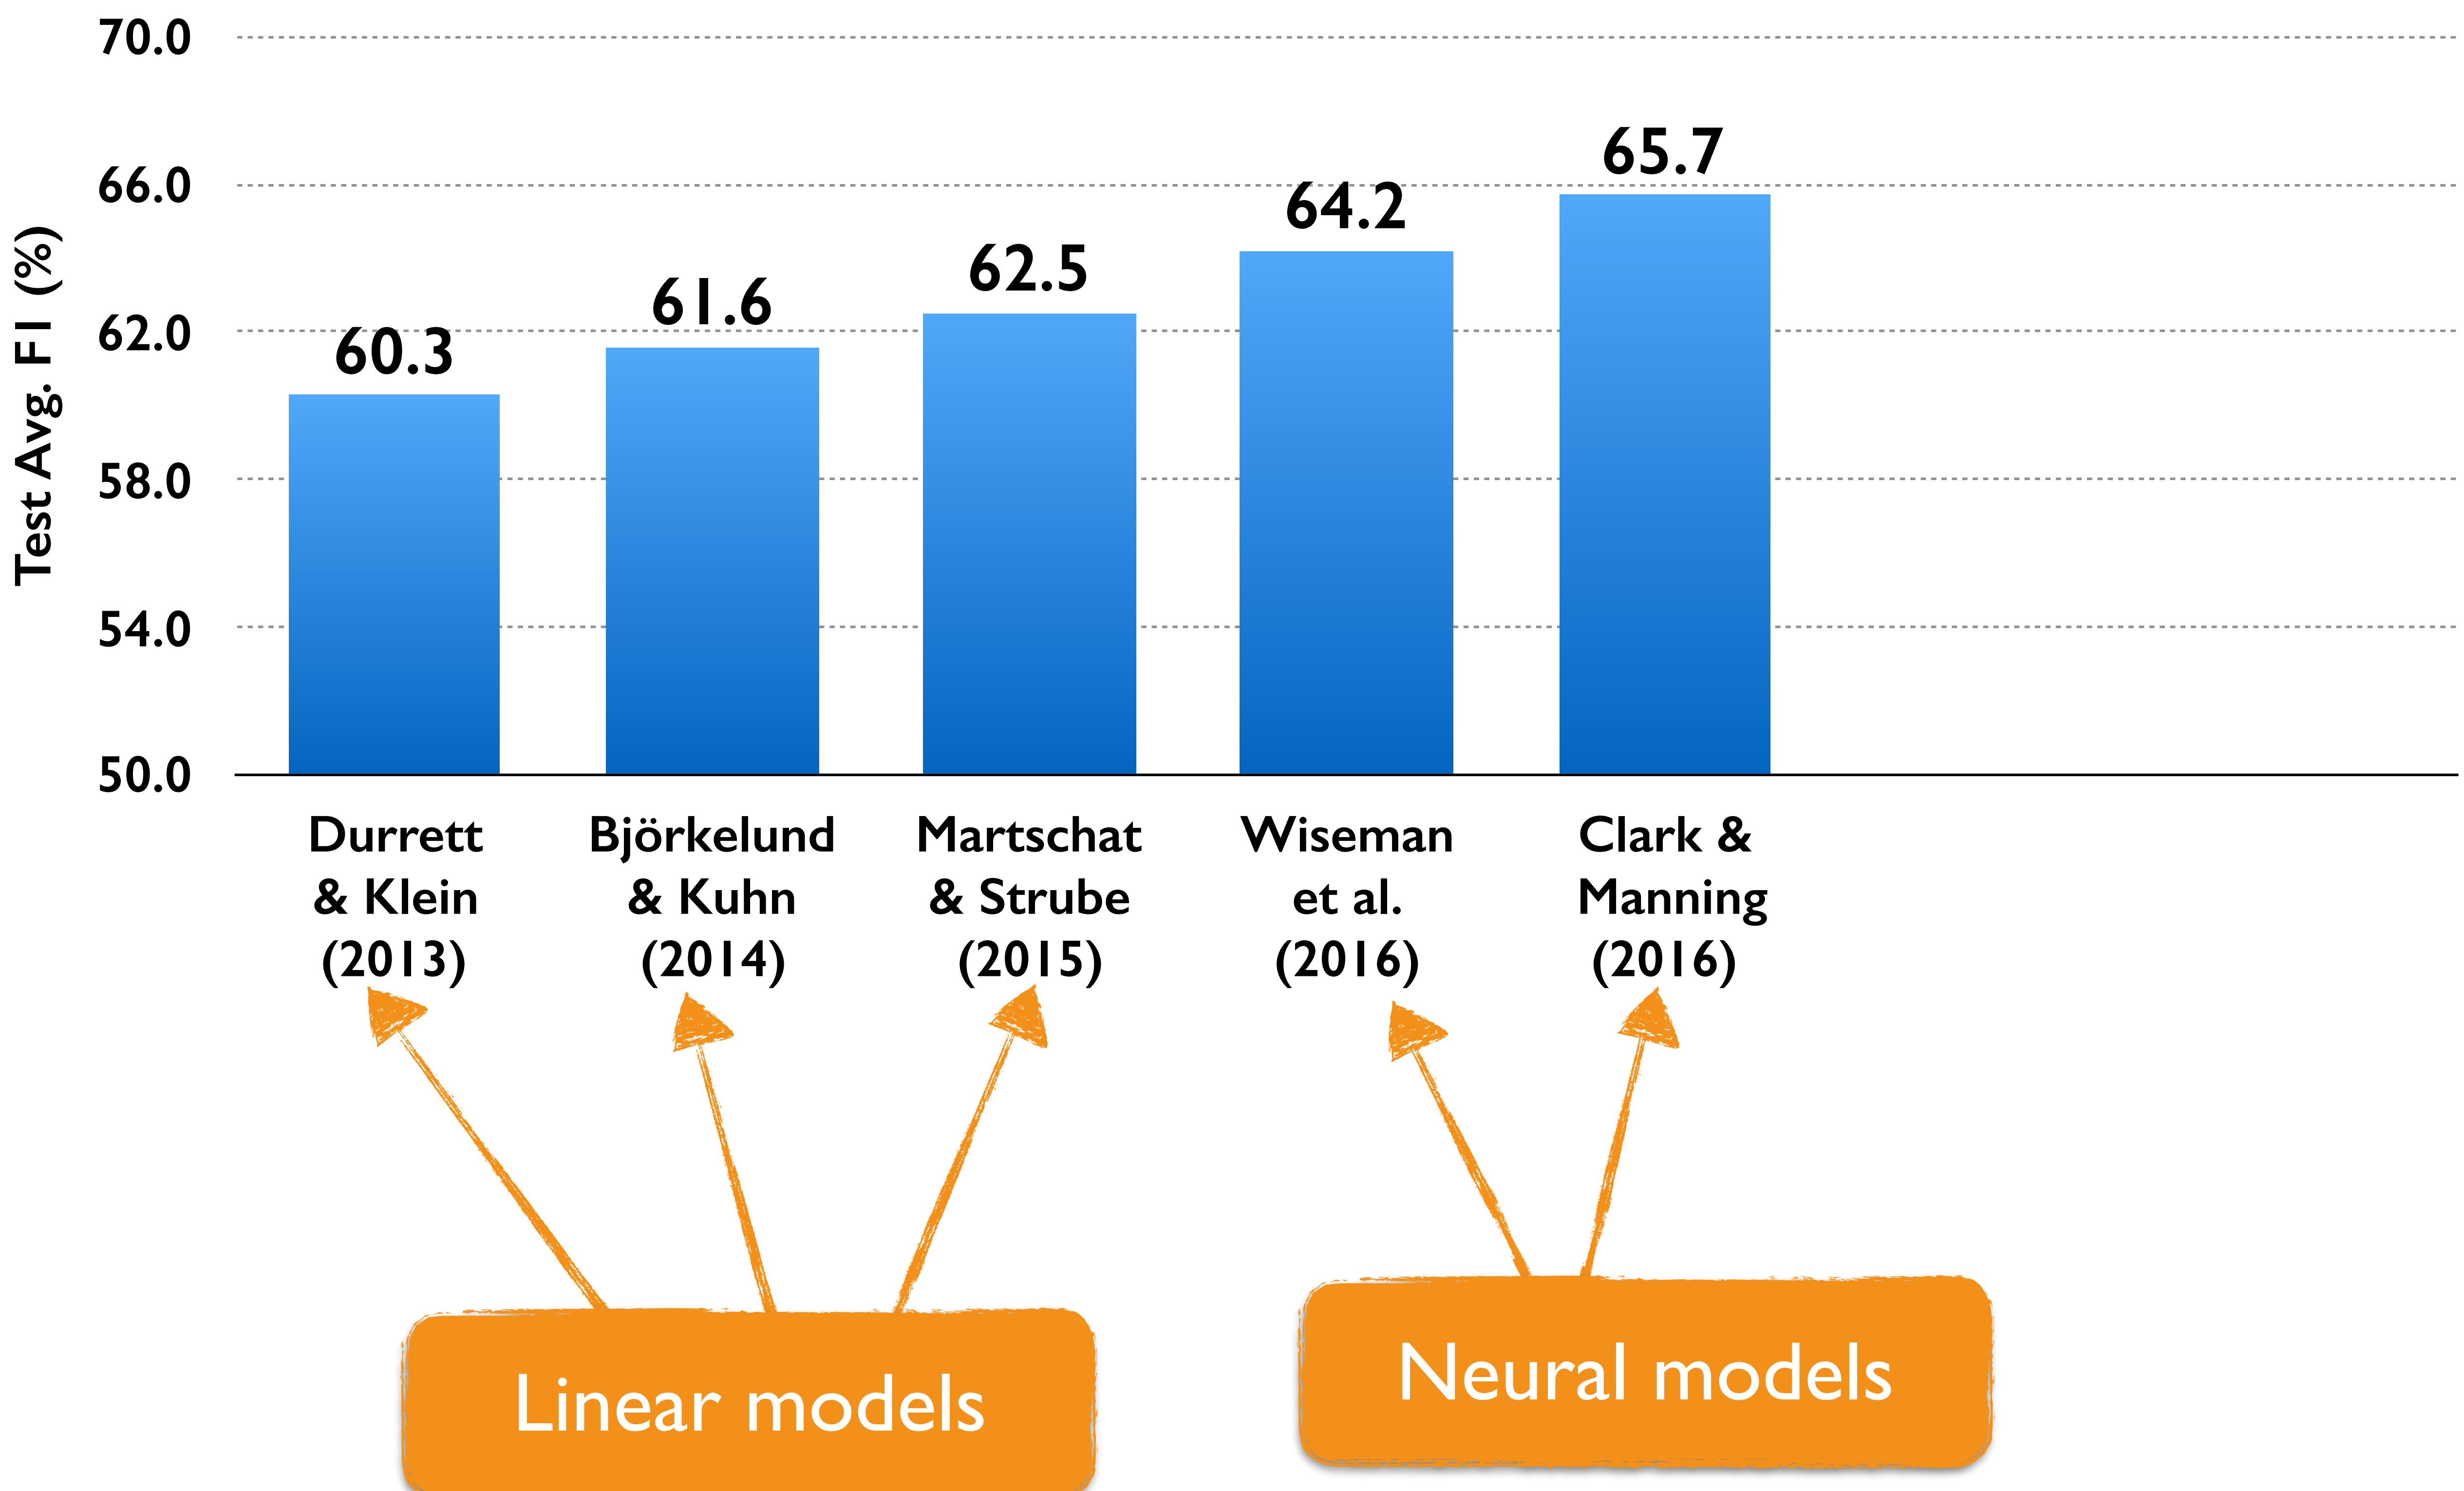

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

Coreference Results

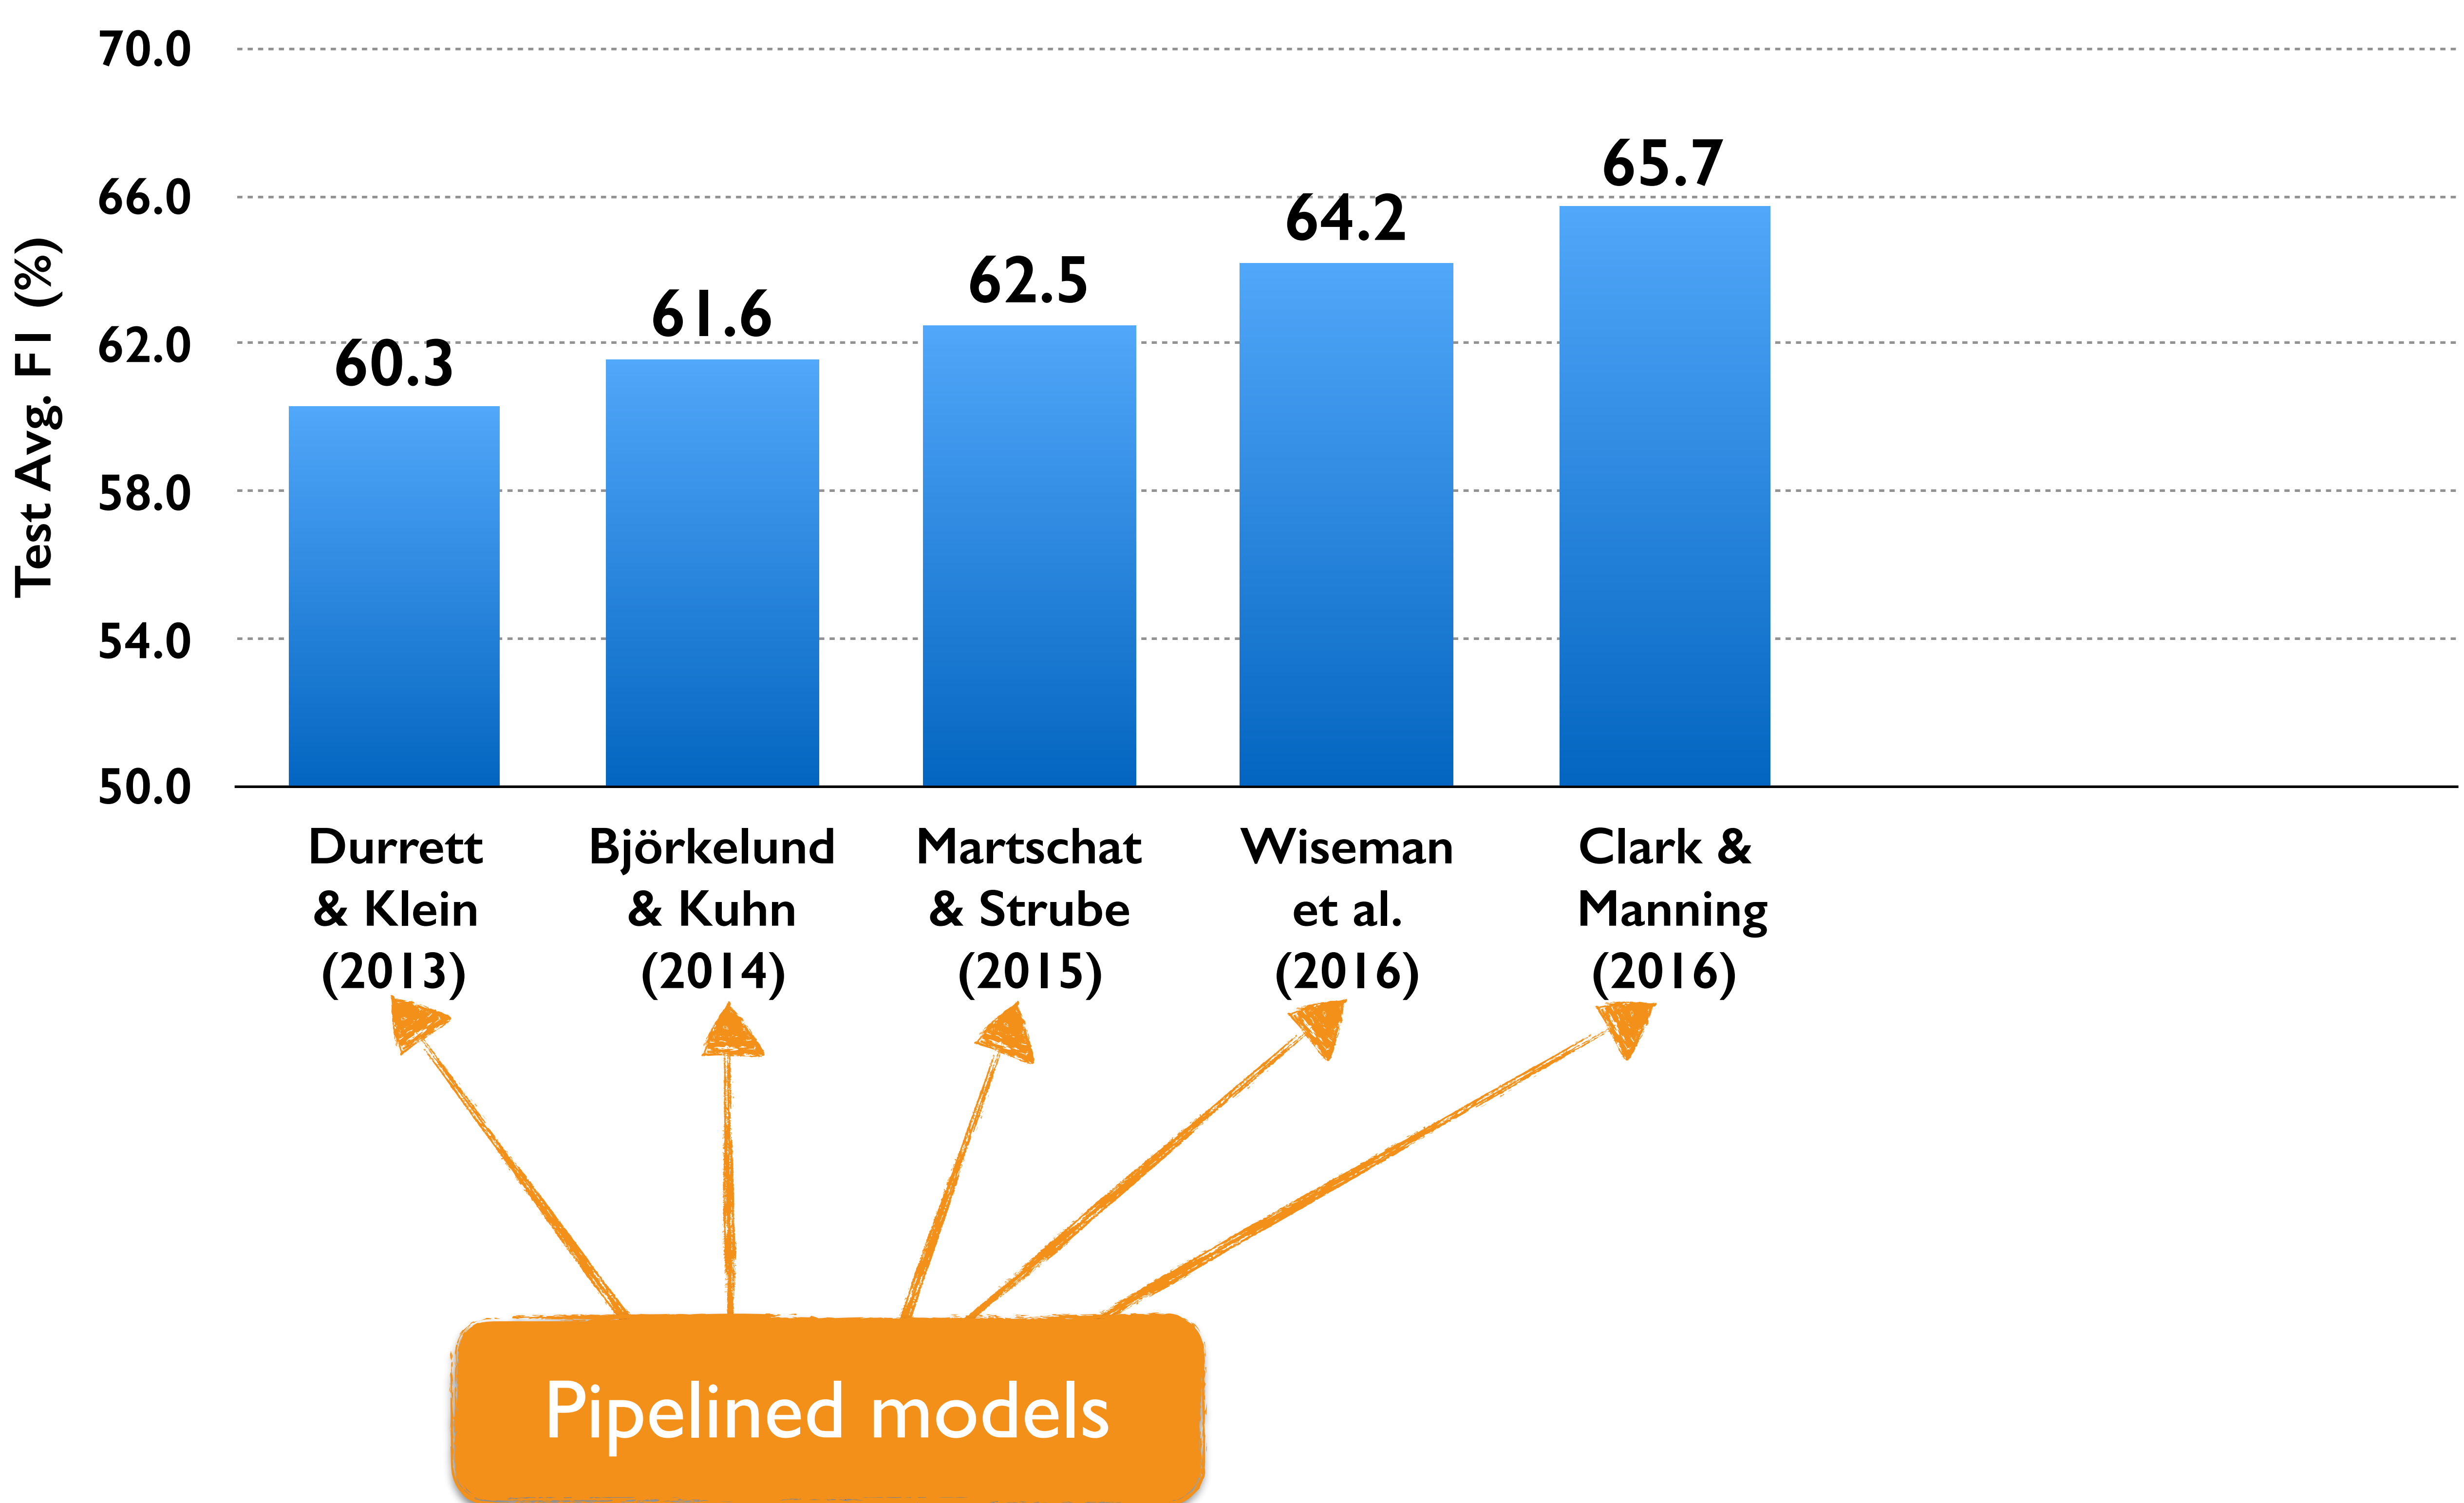

Slide from Lee et al., End-to-end Neural Coreference Resolution, EMNLP 2017

End-to-end Neural Coreference Resolution

Kenton Lee Luheng He Mike Lewis Luke Zettlemoyer

University of Washington

Facebook AI Research

Allen Institute for Artificial Intelligence

Previous Approach: Rule-based pipeline

Input document

A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized.

Syntactic parser

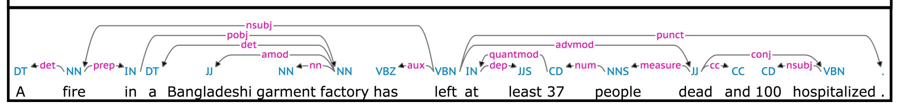

Hand-engineered rules

```
public static boolean inStopList(Mention m) {
  String mentionSpan = m.spanToString().toLowerCase(Locale.ENGLISH);
  if (mentionSpan.equals("u.s.") || mentionSpan.equals("u.k.")
      || mentionSpan.equals("u.s.s.r")) return true;
  if (mentionSpan.equals("there") || mentionSpan.startsWith("etc.")
      || mentionSpan.equals("ltd.")) return true;
  if (mentionSpan.startsWith("'s ")) return true;
//    if (mentionSpan.endsWith("etc.")) return true;

  return false;
}

protected void removeSpuriousMentions(Annotation doc, List<List<Mention>> predictedMentions, Dictionaries dict, boolean remov
  if(lang == Locale.ENGLISH) removeSpuriousMentionsEn(doc, predictedMentions, dict);
  else if (lang == Locale.CHINESE) removeSpuriousMentionsZh(doc, predictedMentions, dict, removeNested);
}

n>> predictedMentions, Dictionaries dict) {
ss);

ass);
```

Relies on parser for:

- mention detection
- syntactic features for clustering (e.g. head words)

A fire in a Bangladeshi garment factory
garment
factory
at least 37 people dead and 100 hospitalized
…

		Coreferent?
A fire in a Bangladeshi garment factory	garment	✓/✗
garment	factory	✓/✗
factory	at least 37 people dead and 100 hospitalized	✓/✗
…	…	✓/✗

Our Contribution: End-to-end Approach

- Joint mention detection and clustering
- No preprocessing (no parser, no POS-tagger etc.)

End-to-end Approach

- Consider all possible spans
- Learn to rank antecedent spans
- Factored model to prune search space

O(N^4) pairwise decisions

Input document (N words)
A fire in a Bangladeshi garment factory has left at least 37 people dead and 100 hospitalized. Most of the deceased were killed in the crush as workers tried to flee the blaze in the four-story building. Witnesses say the only exit door was on the ground floor, and that it was locked when the fire broke out.

Span #1	**Span #2**	**Coreferent?**
A	A fire	✓/✗
A fire	A fire in	✓/✗
A fire in	A fire in a	✓/✗
…	…	✓/✗

Neural Span Representations

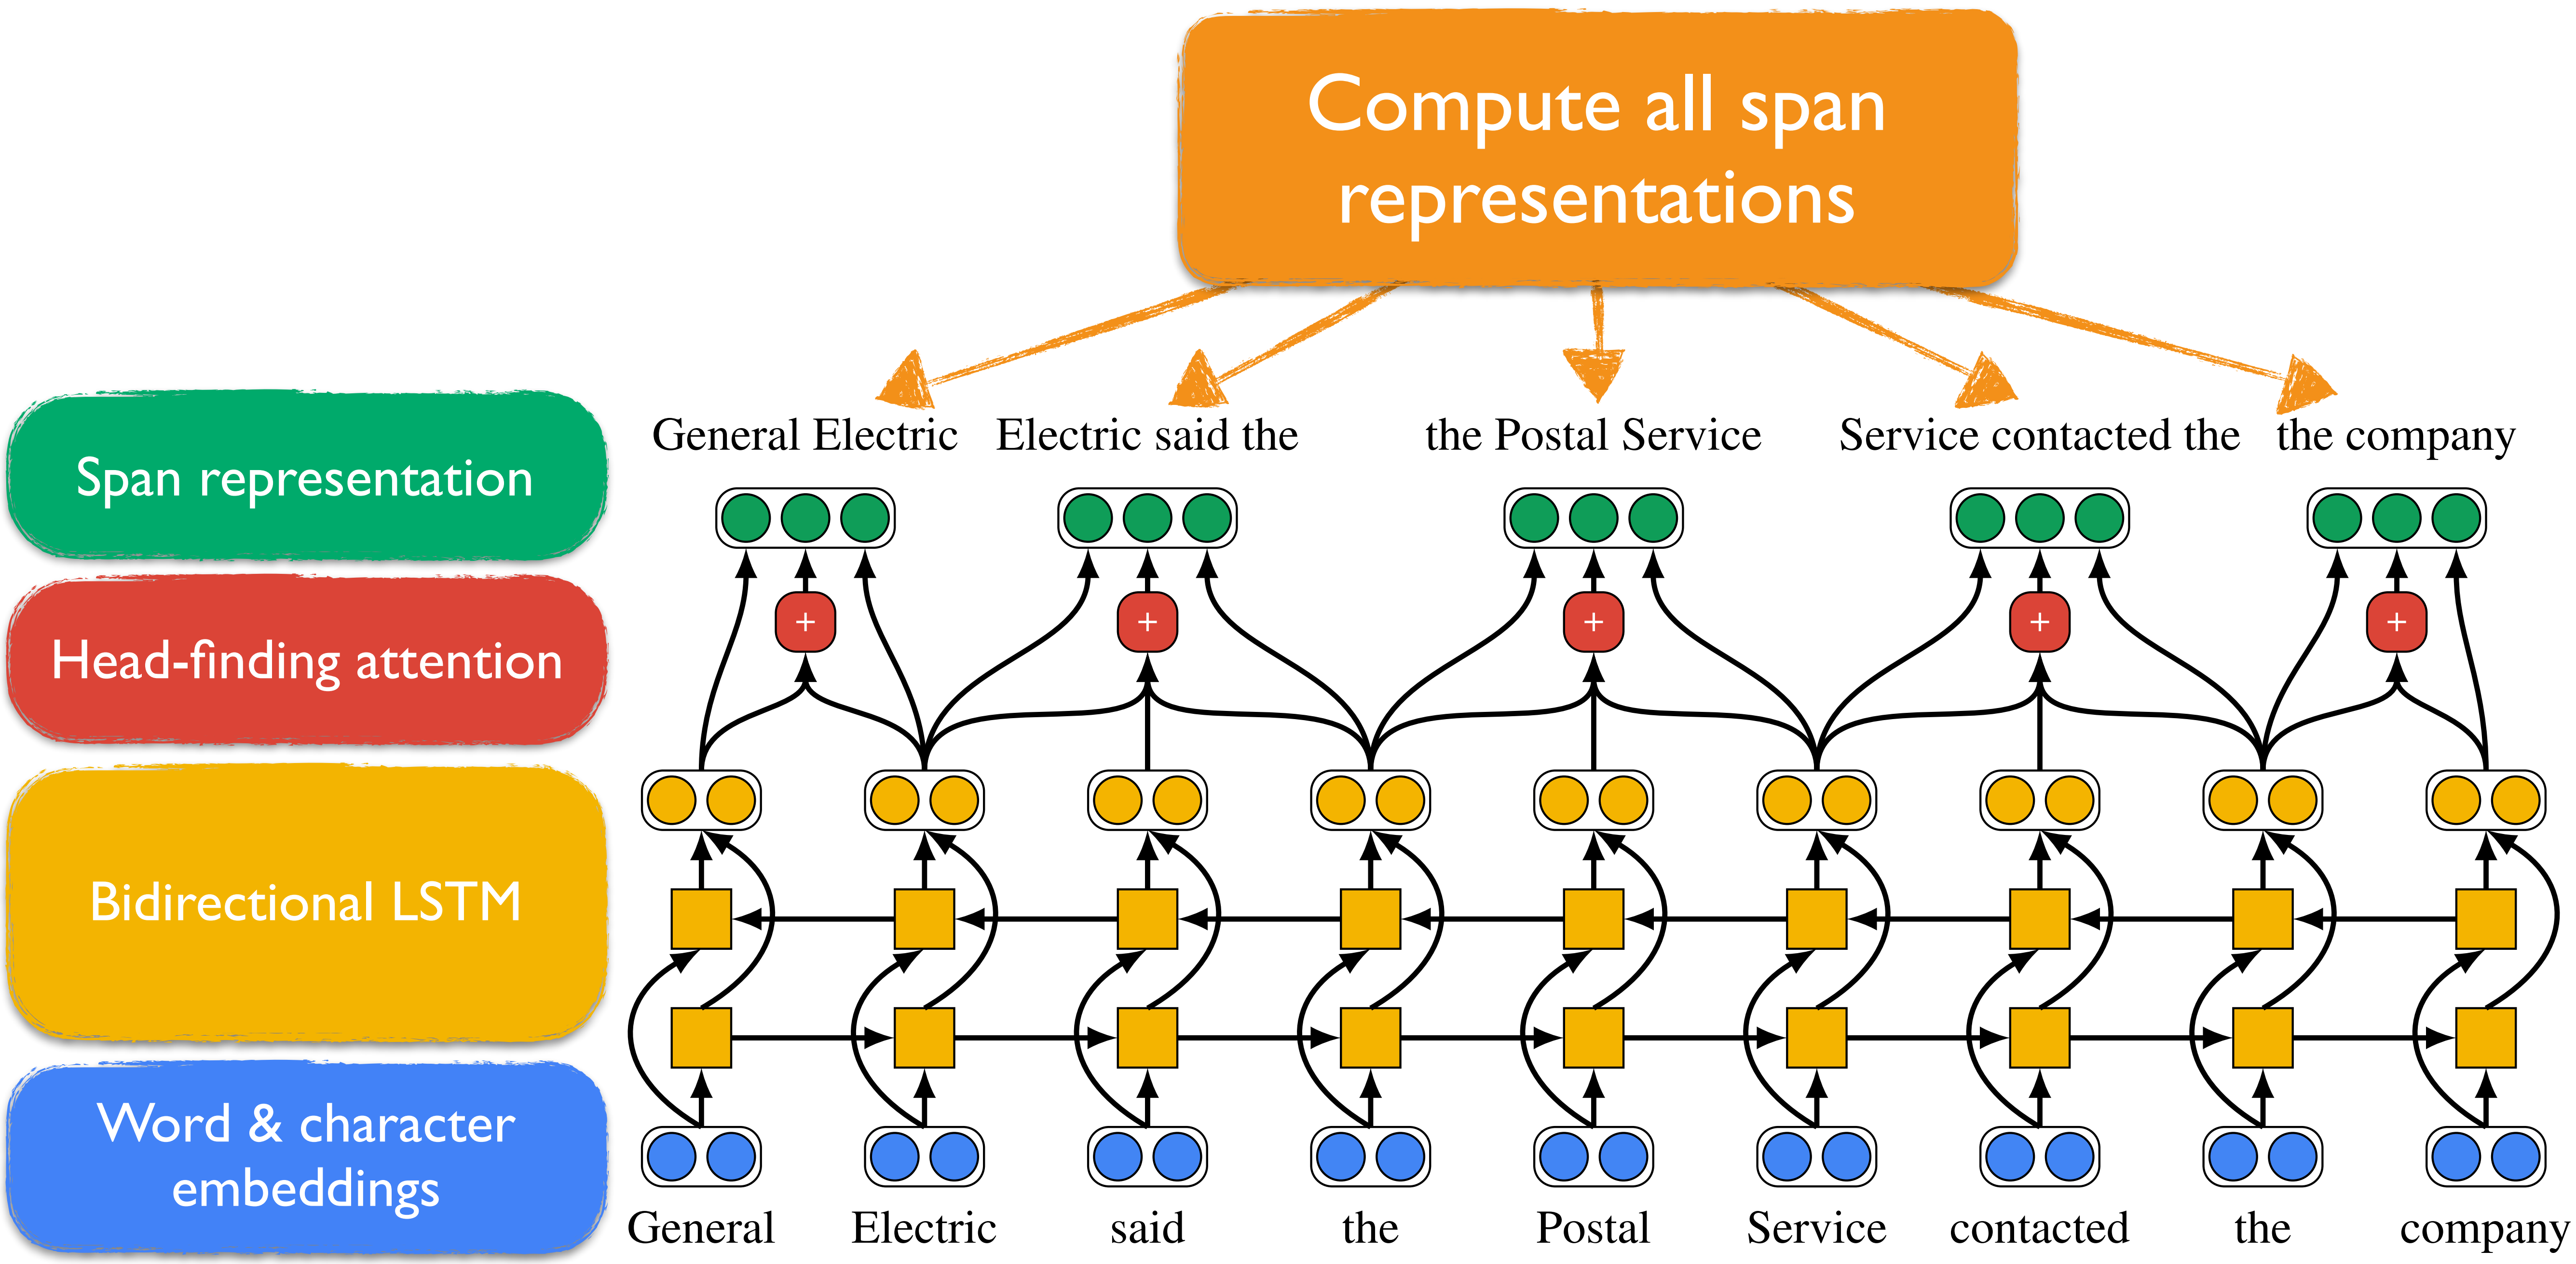

Coreference Architecture

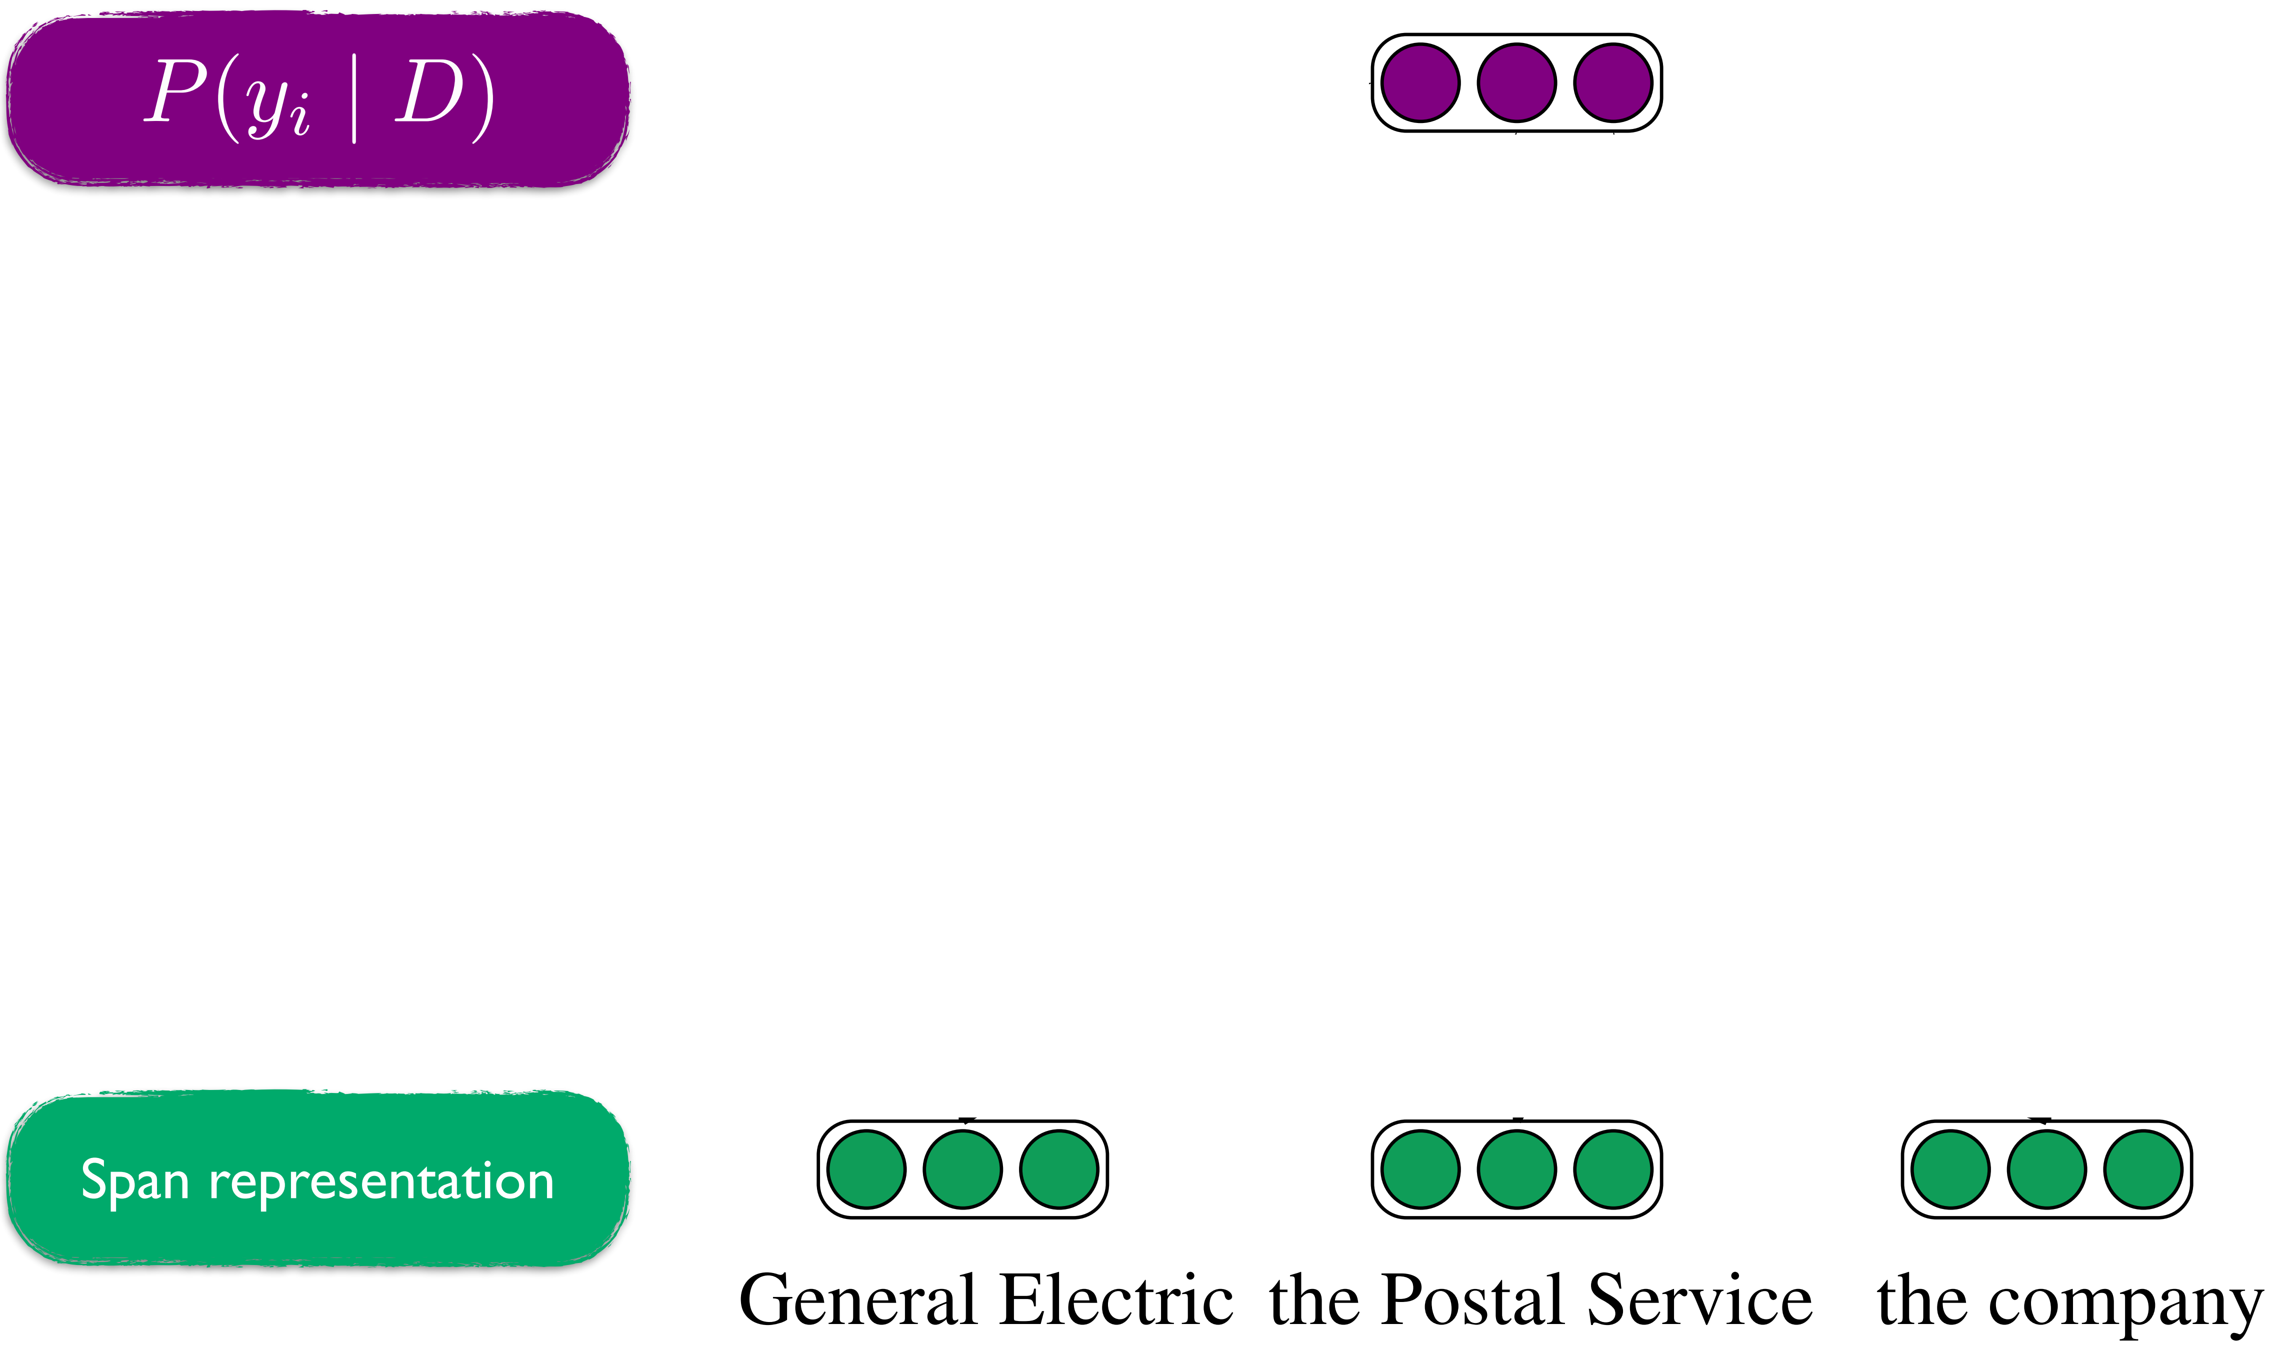

Coreference Architecture

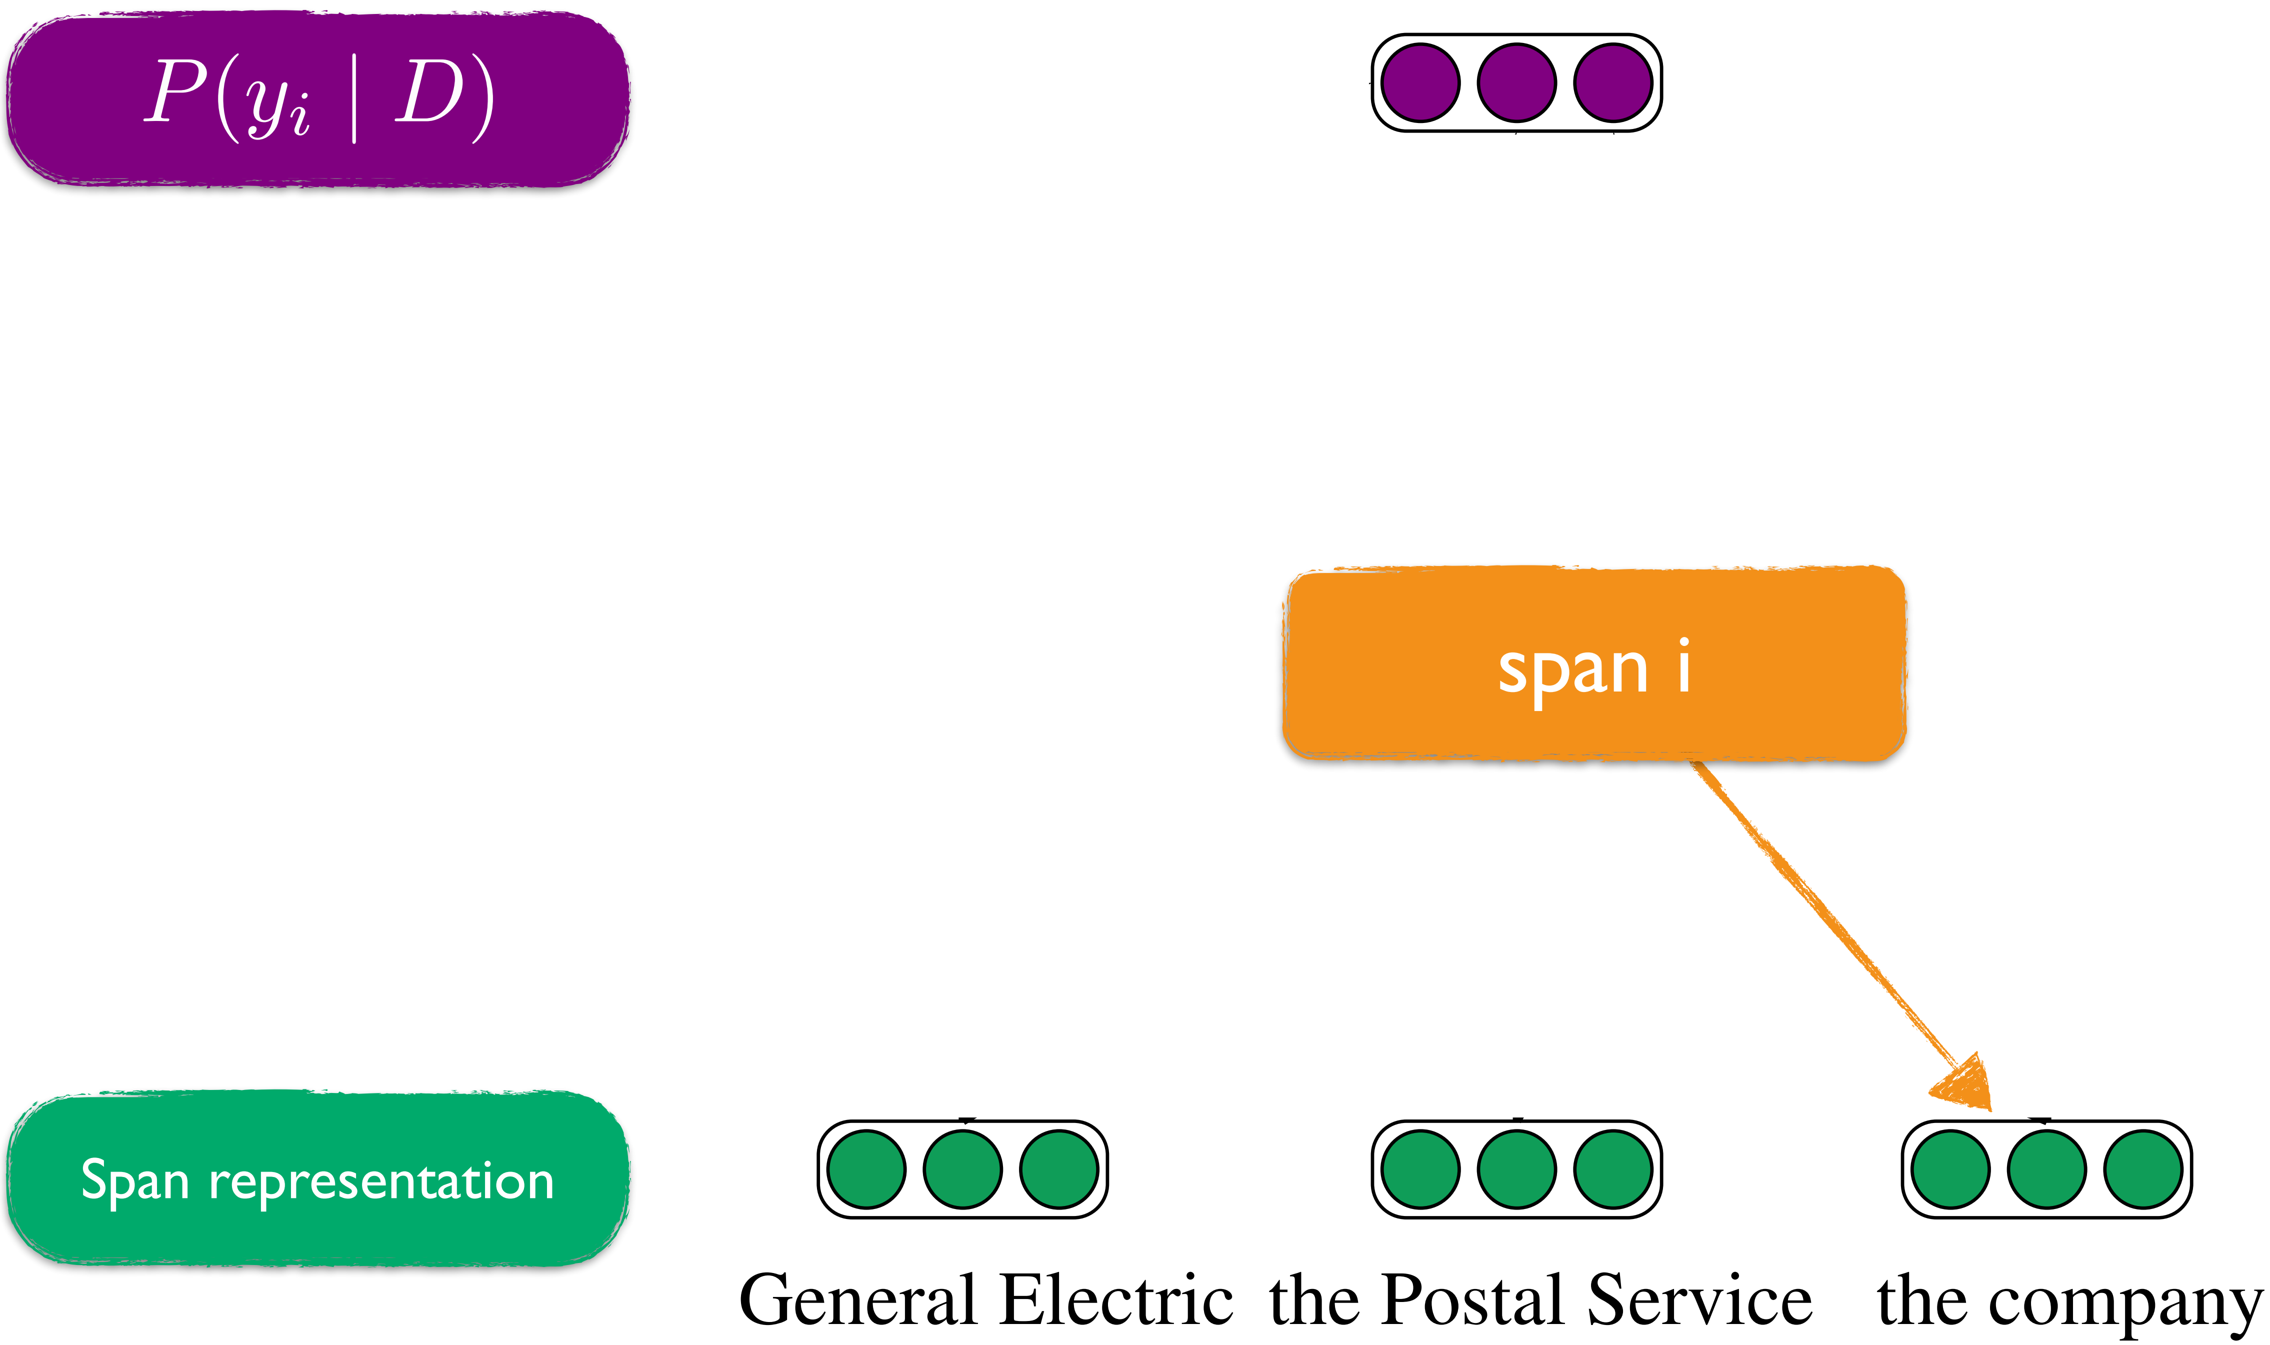

Coreference Architecture

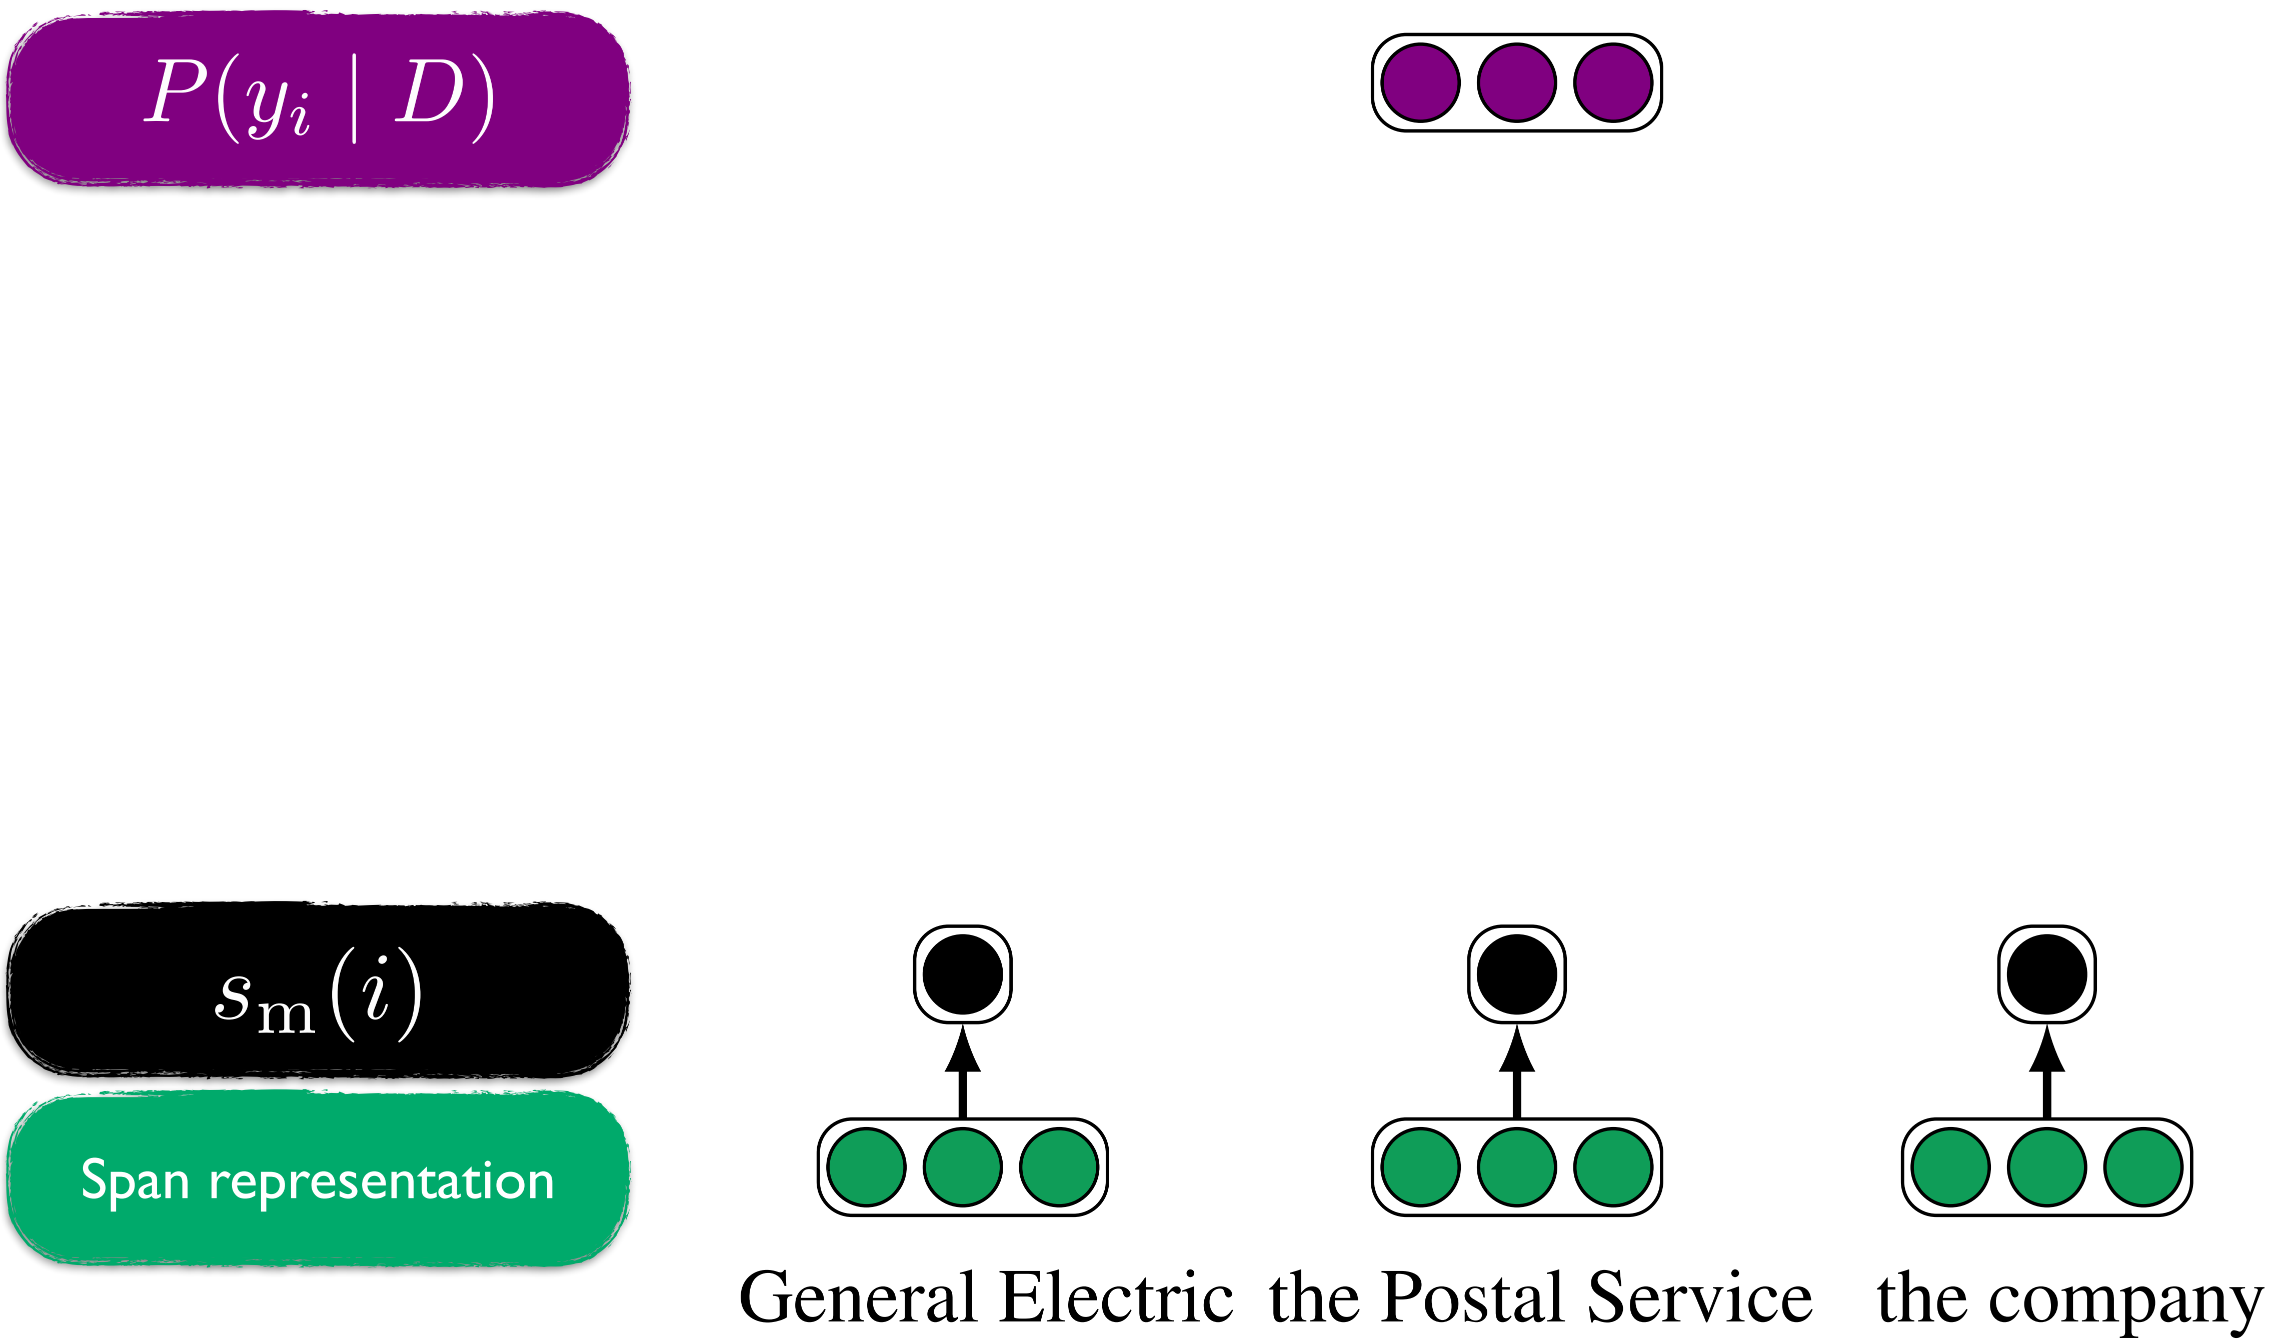

Coreference Architecture

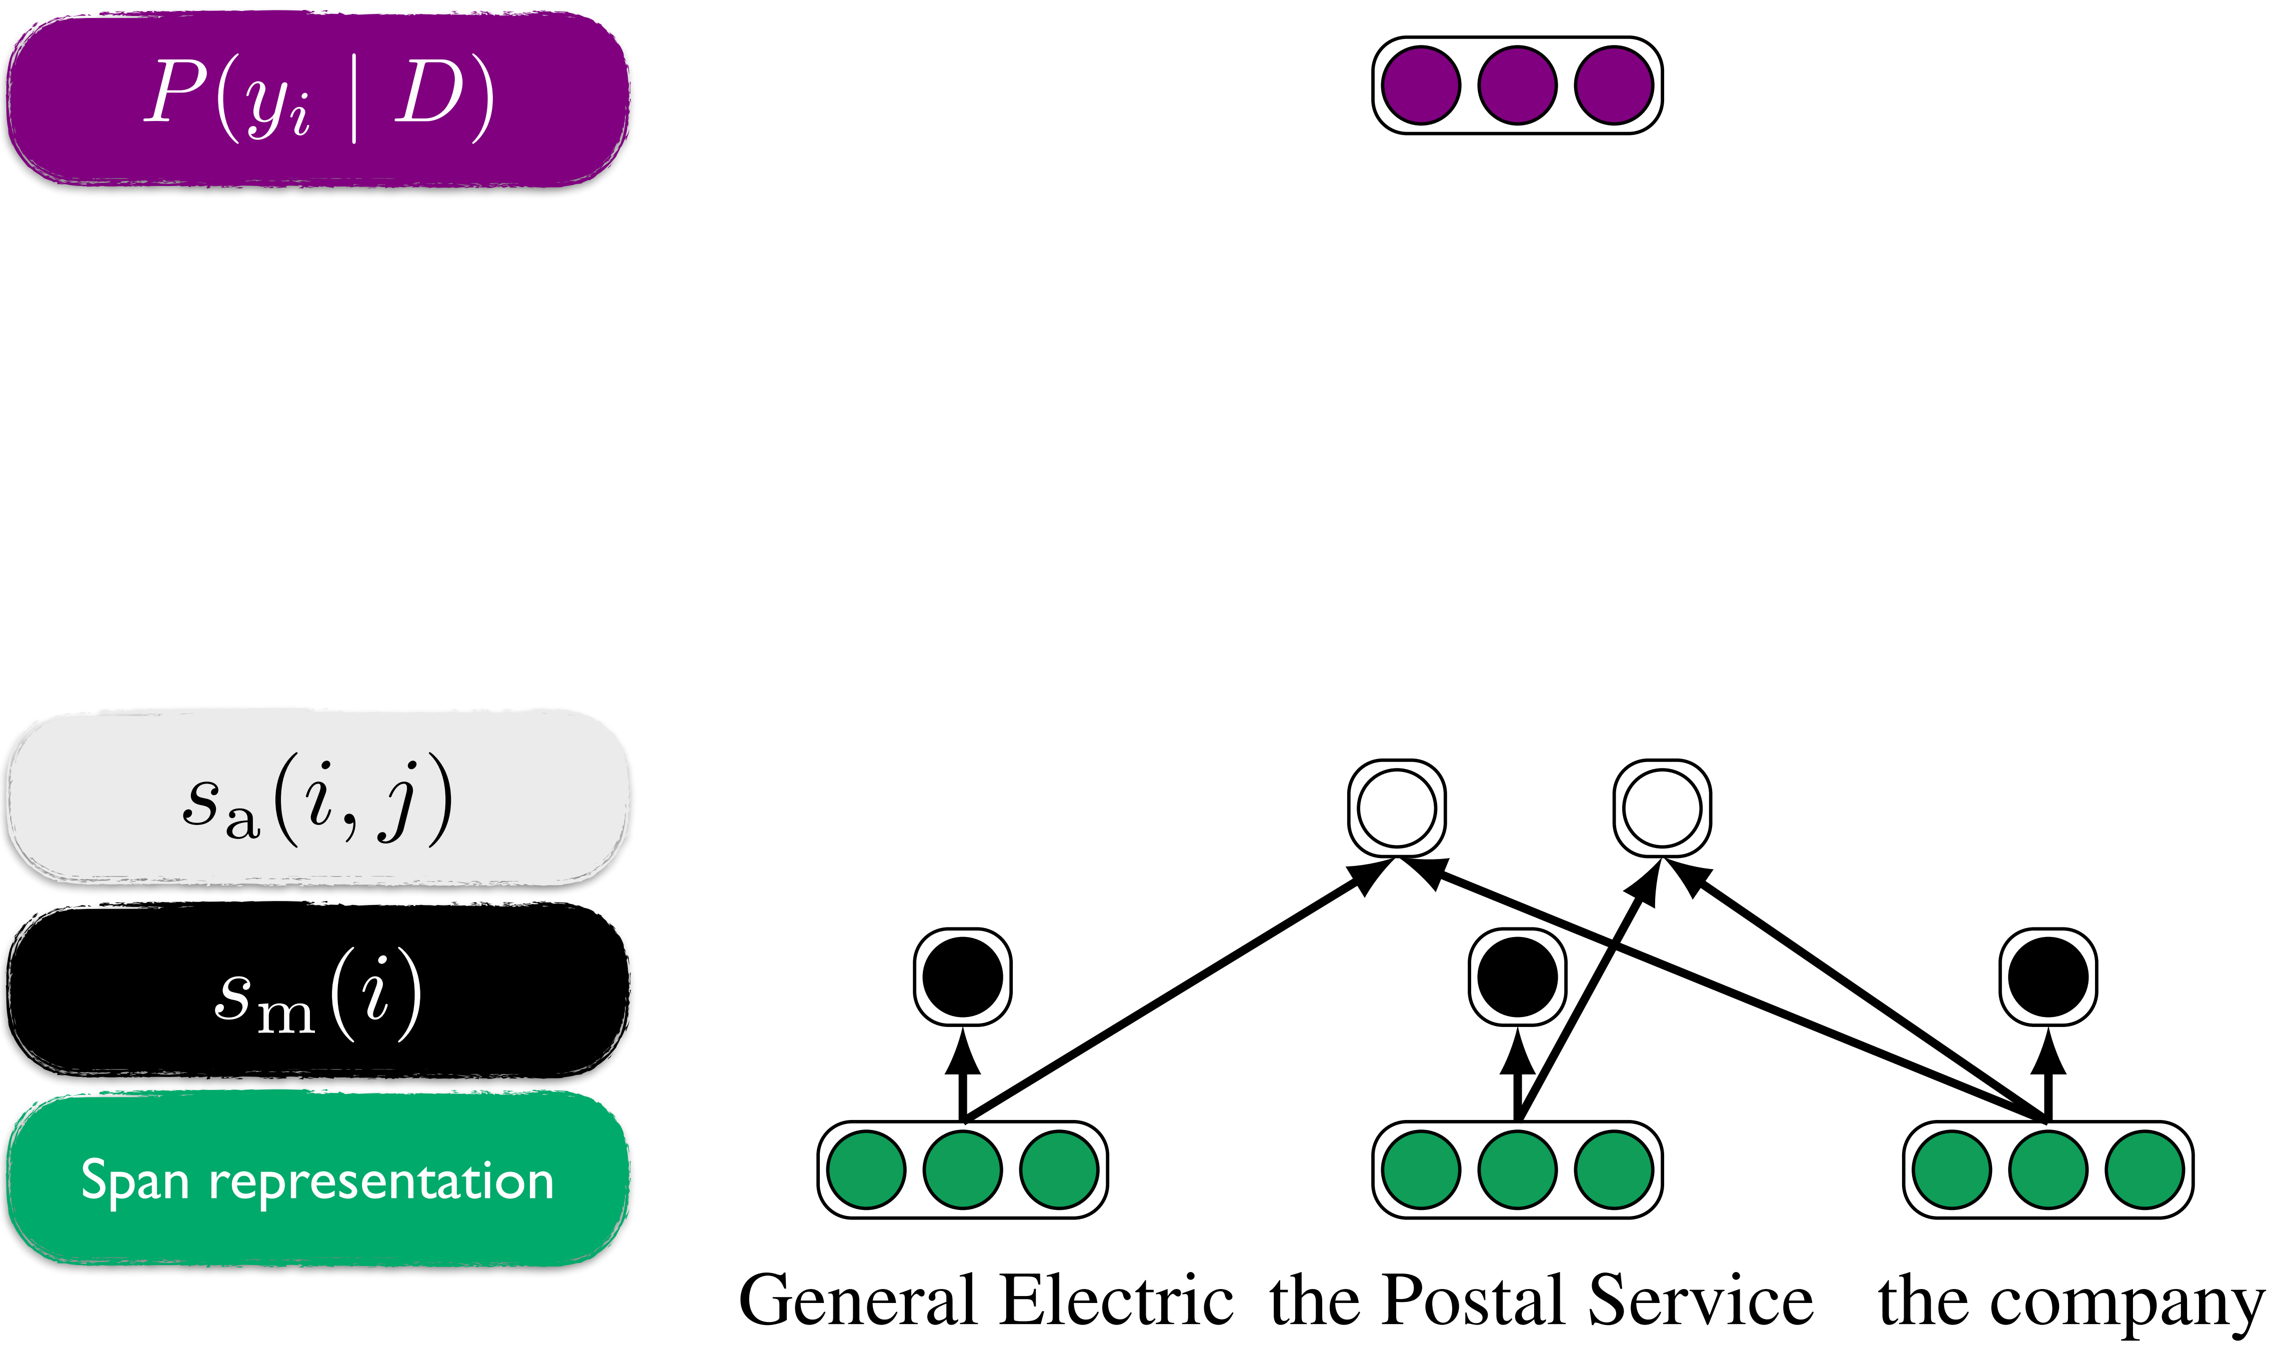

Coreference Architecture

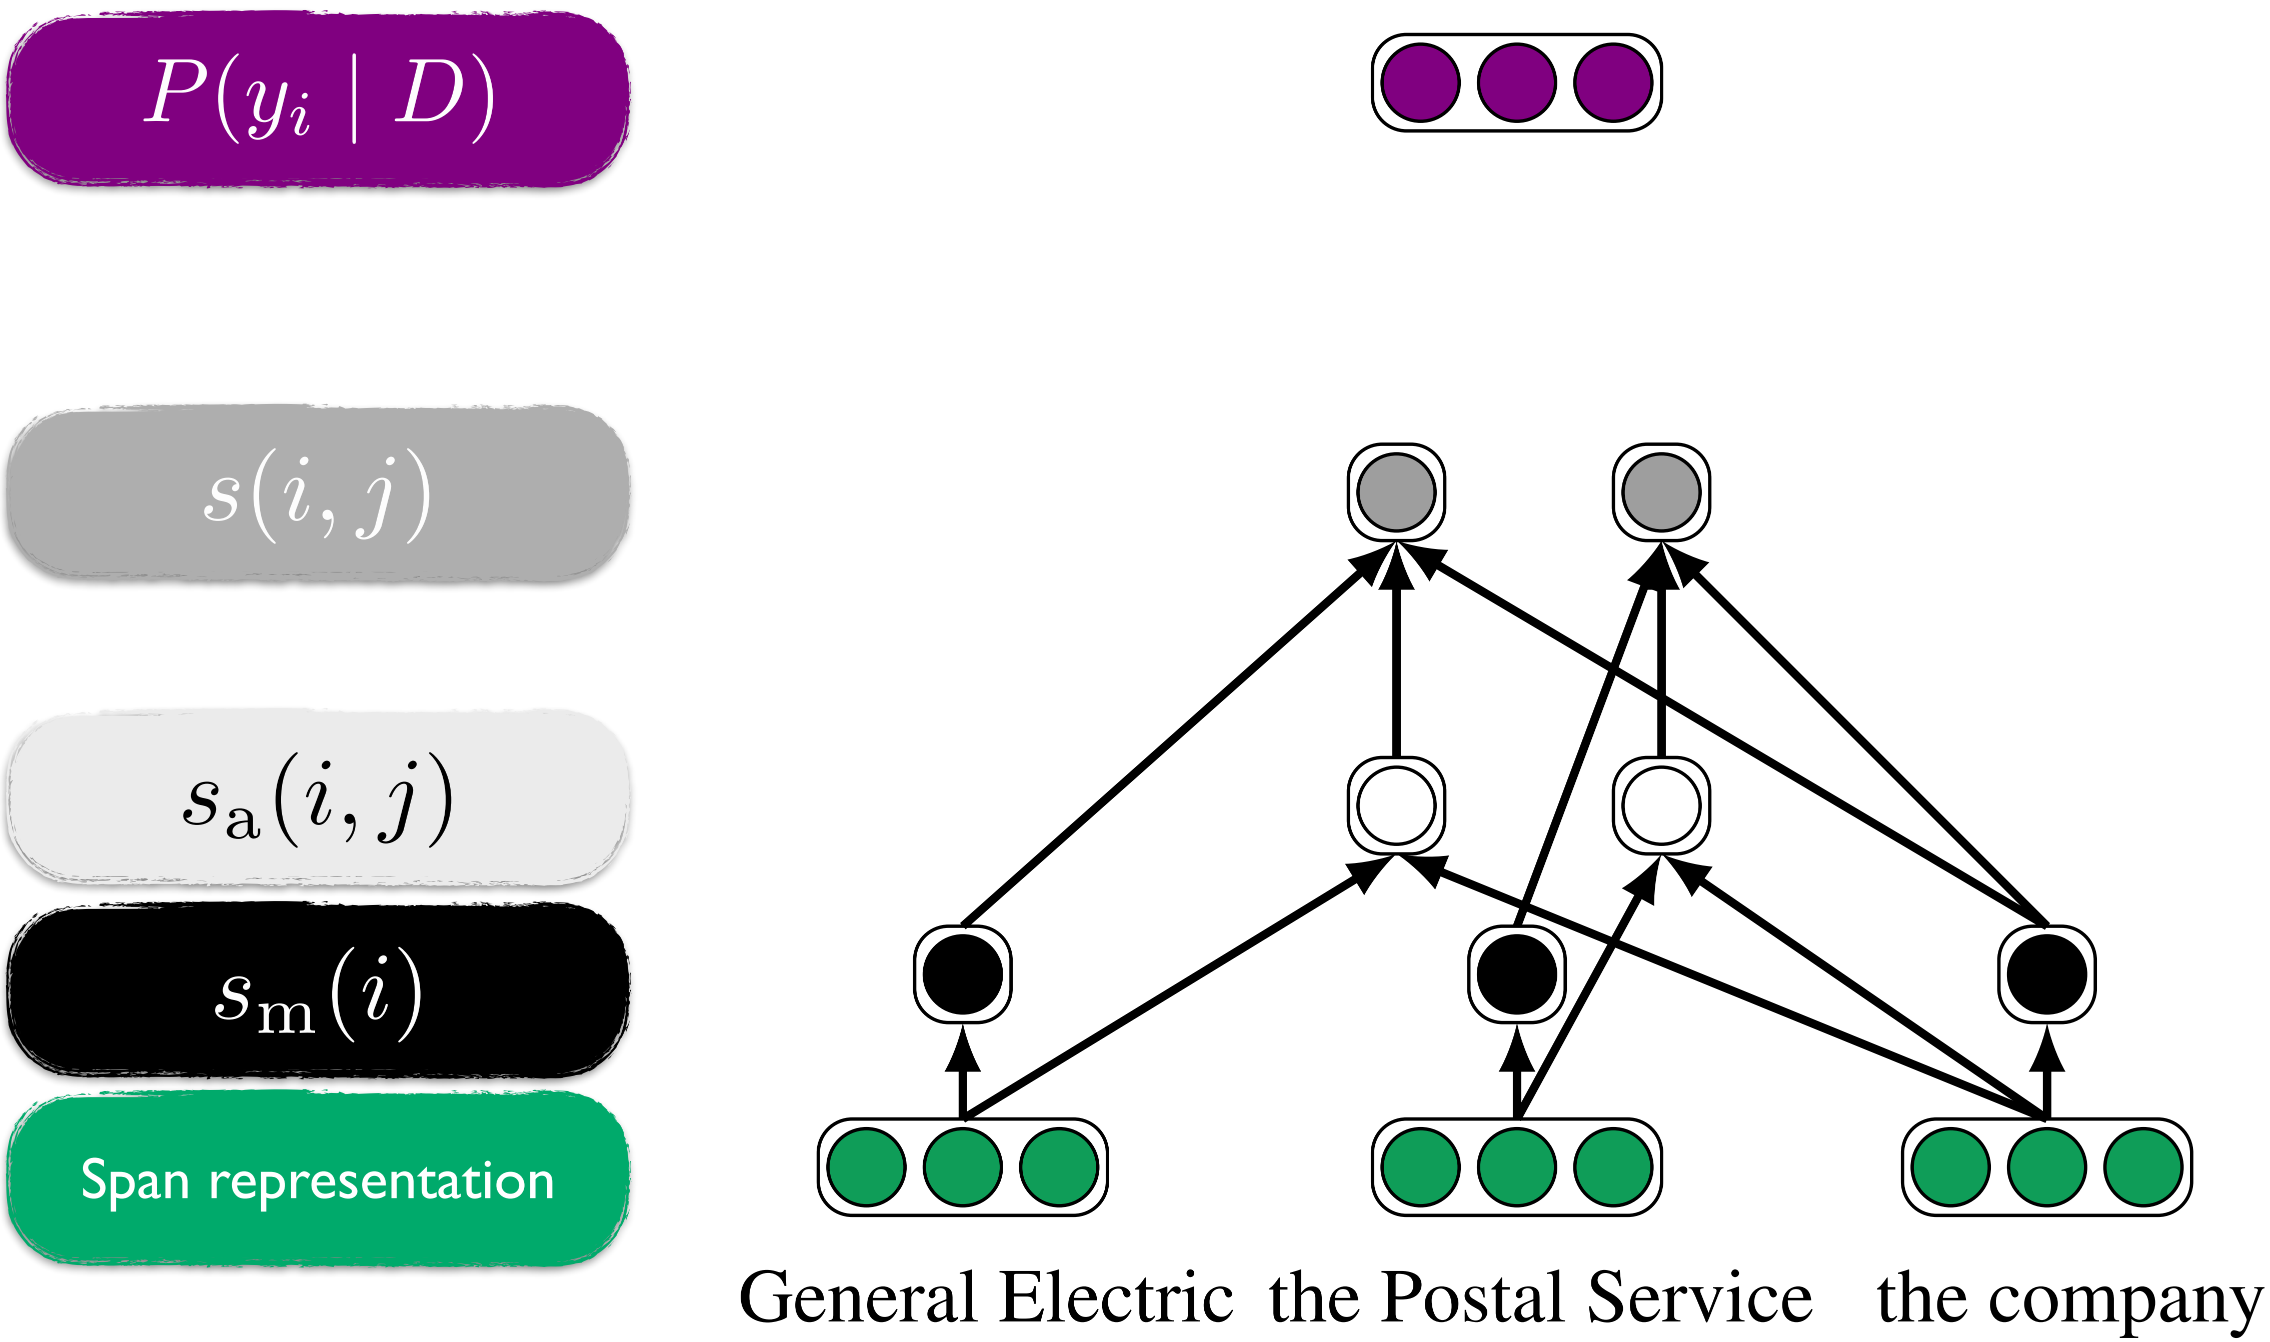

Coreference Architecture

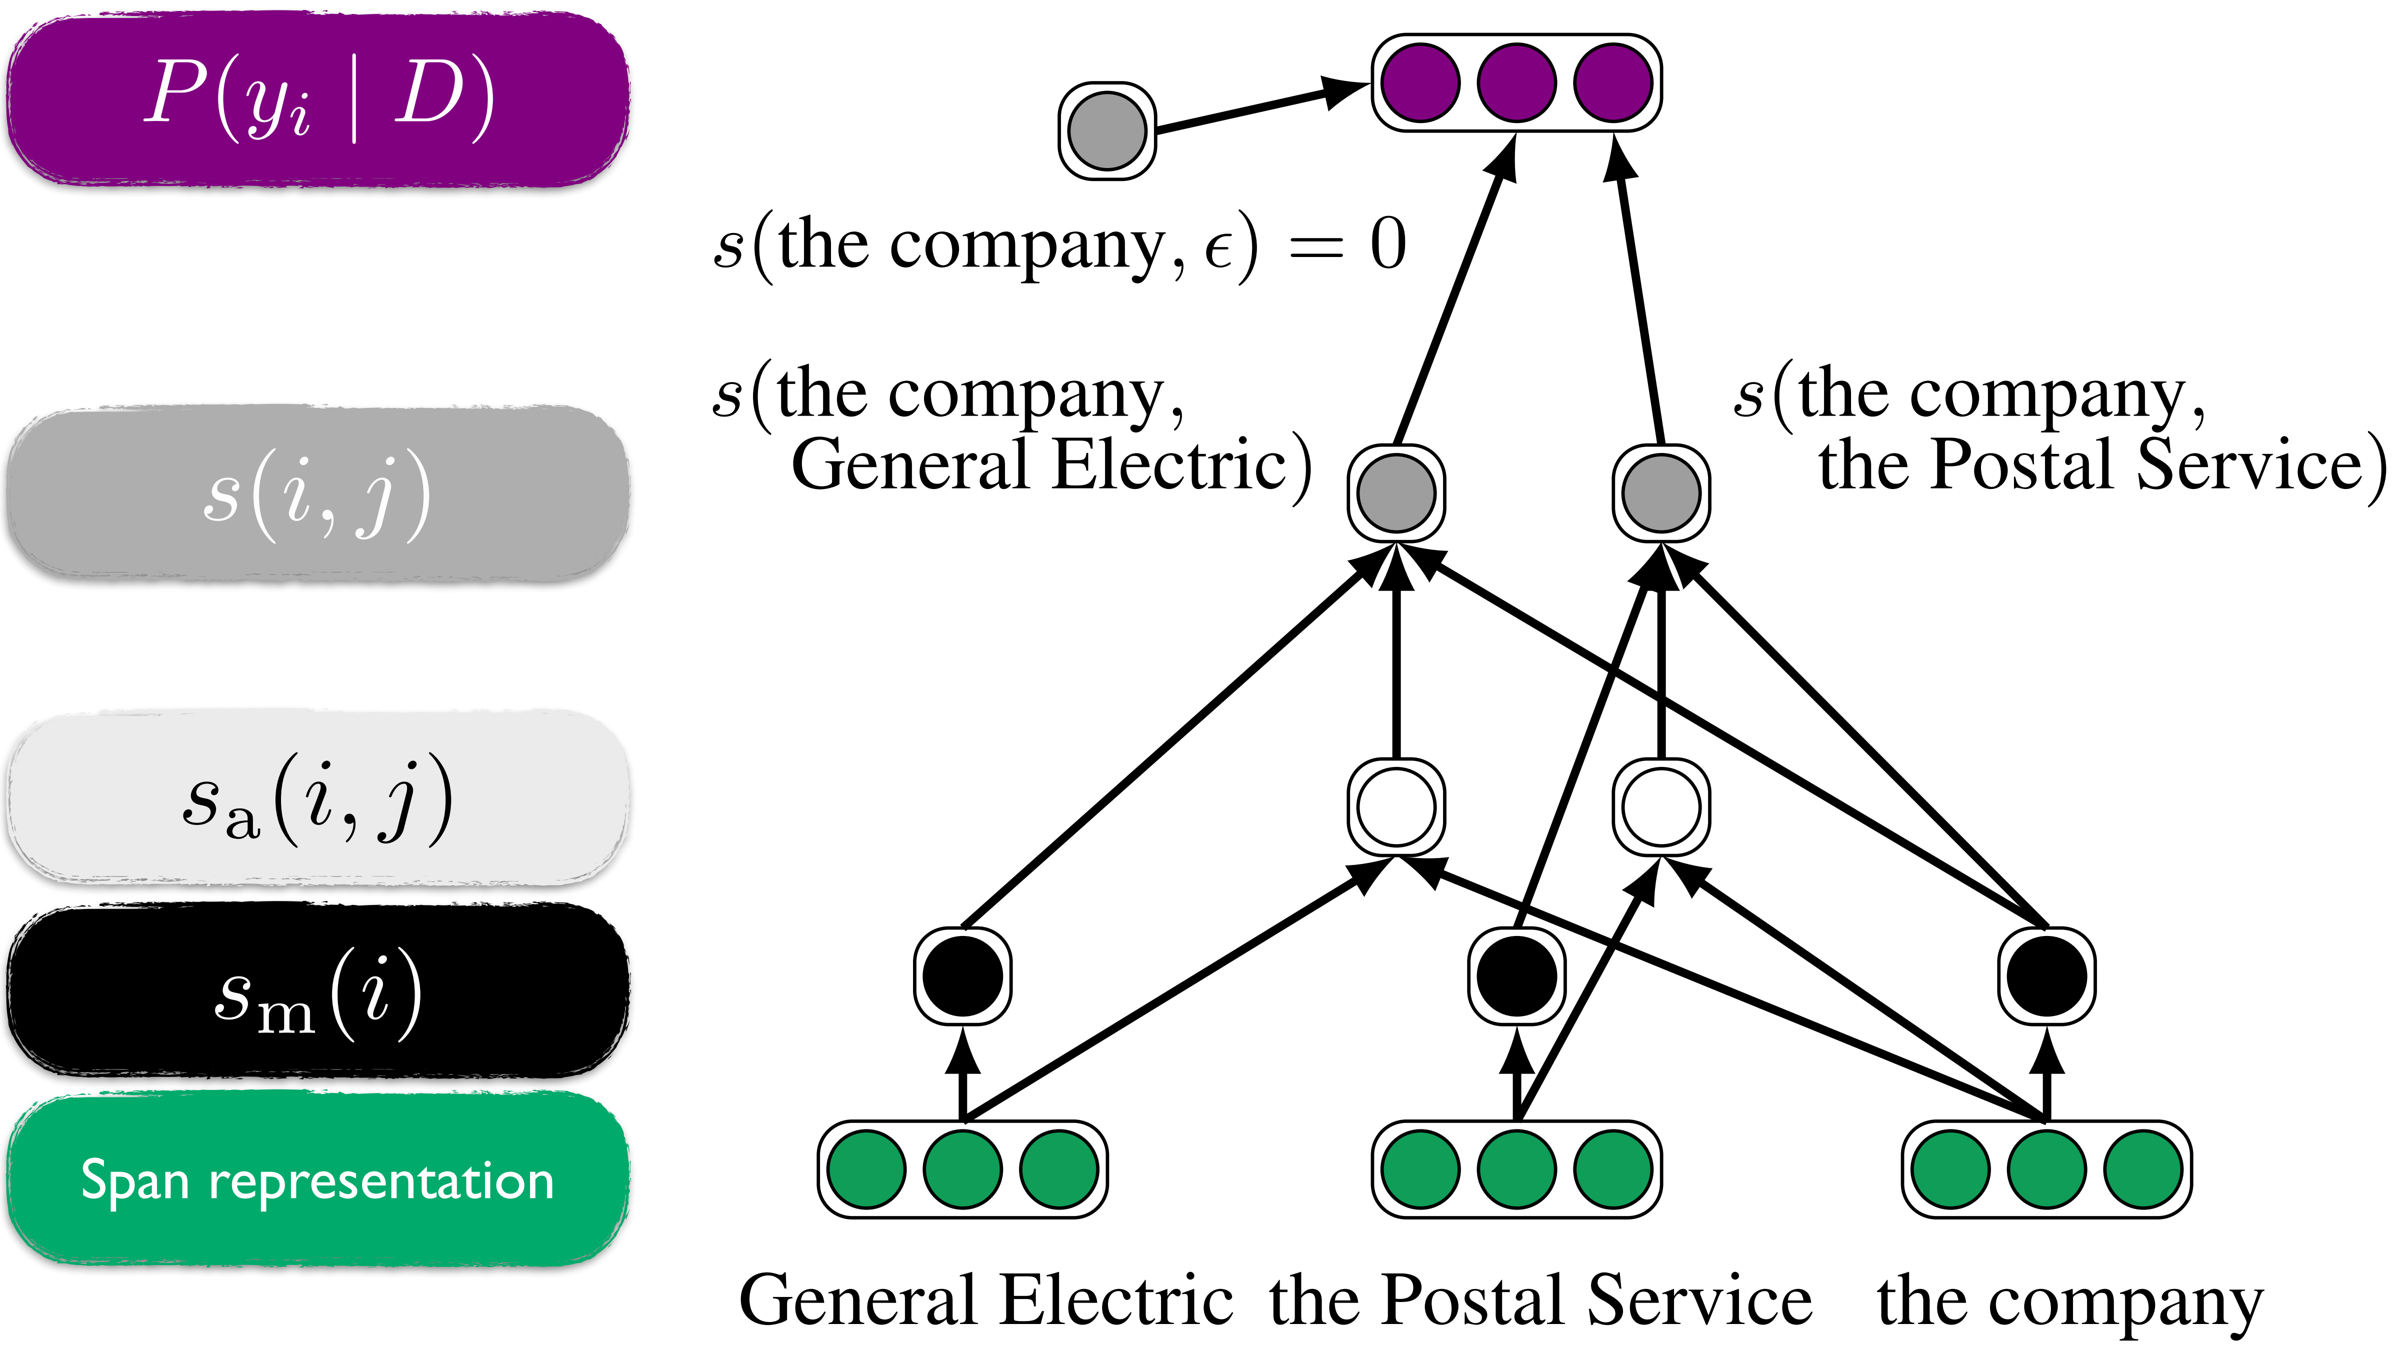

Coreference Results

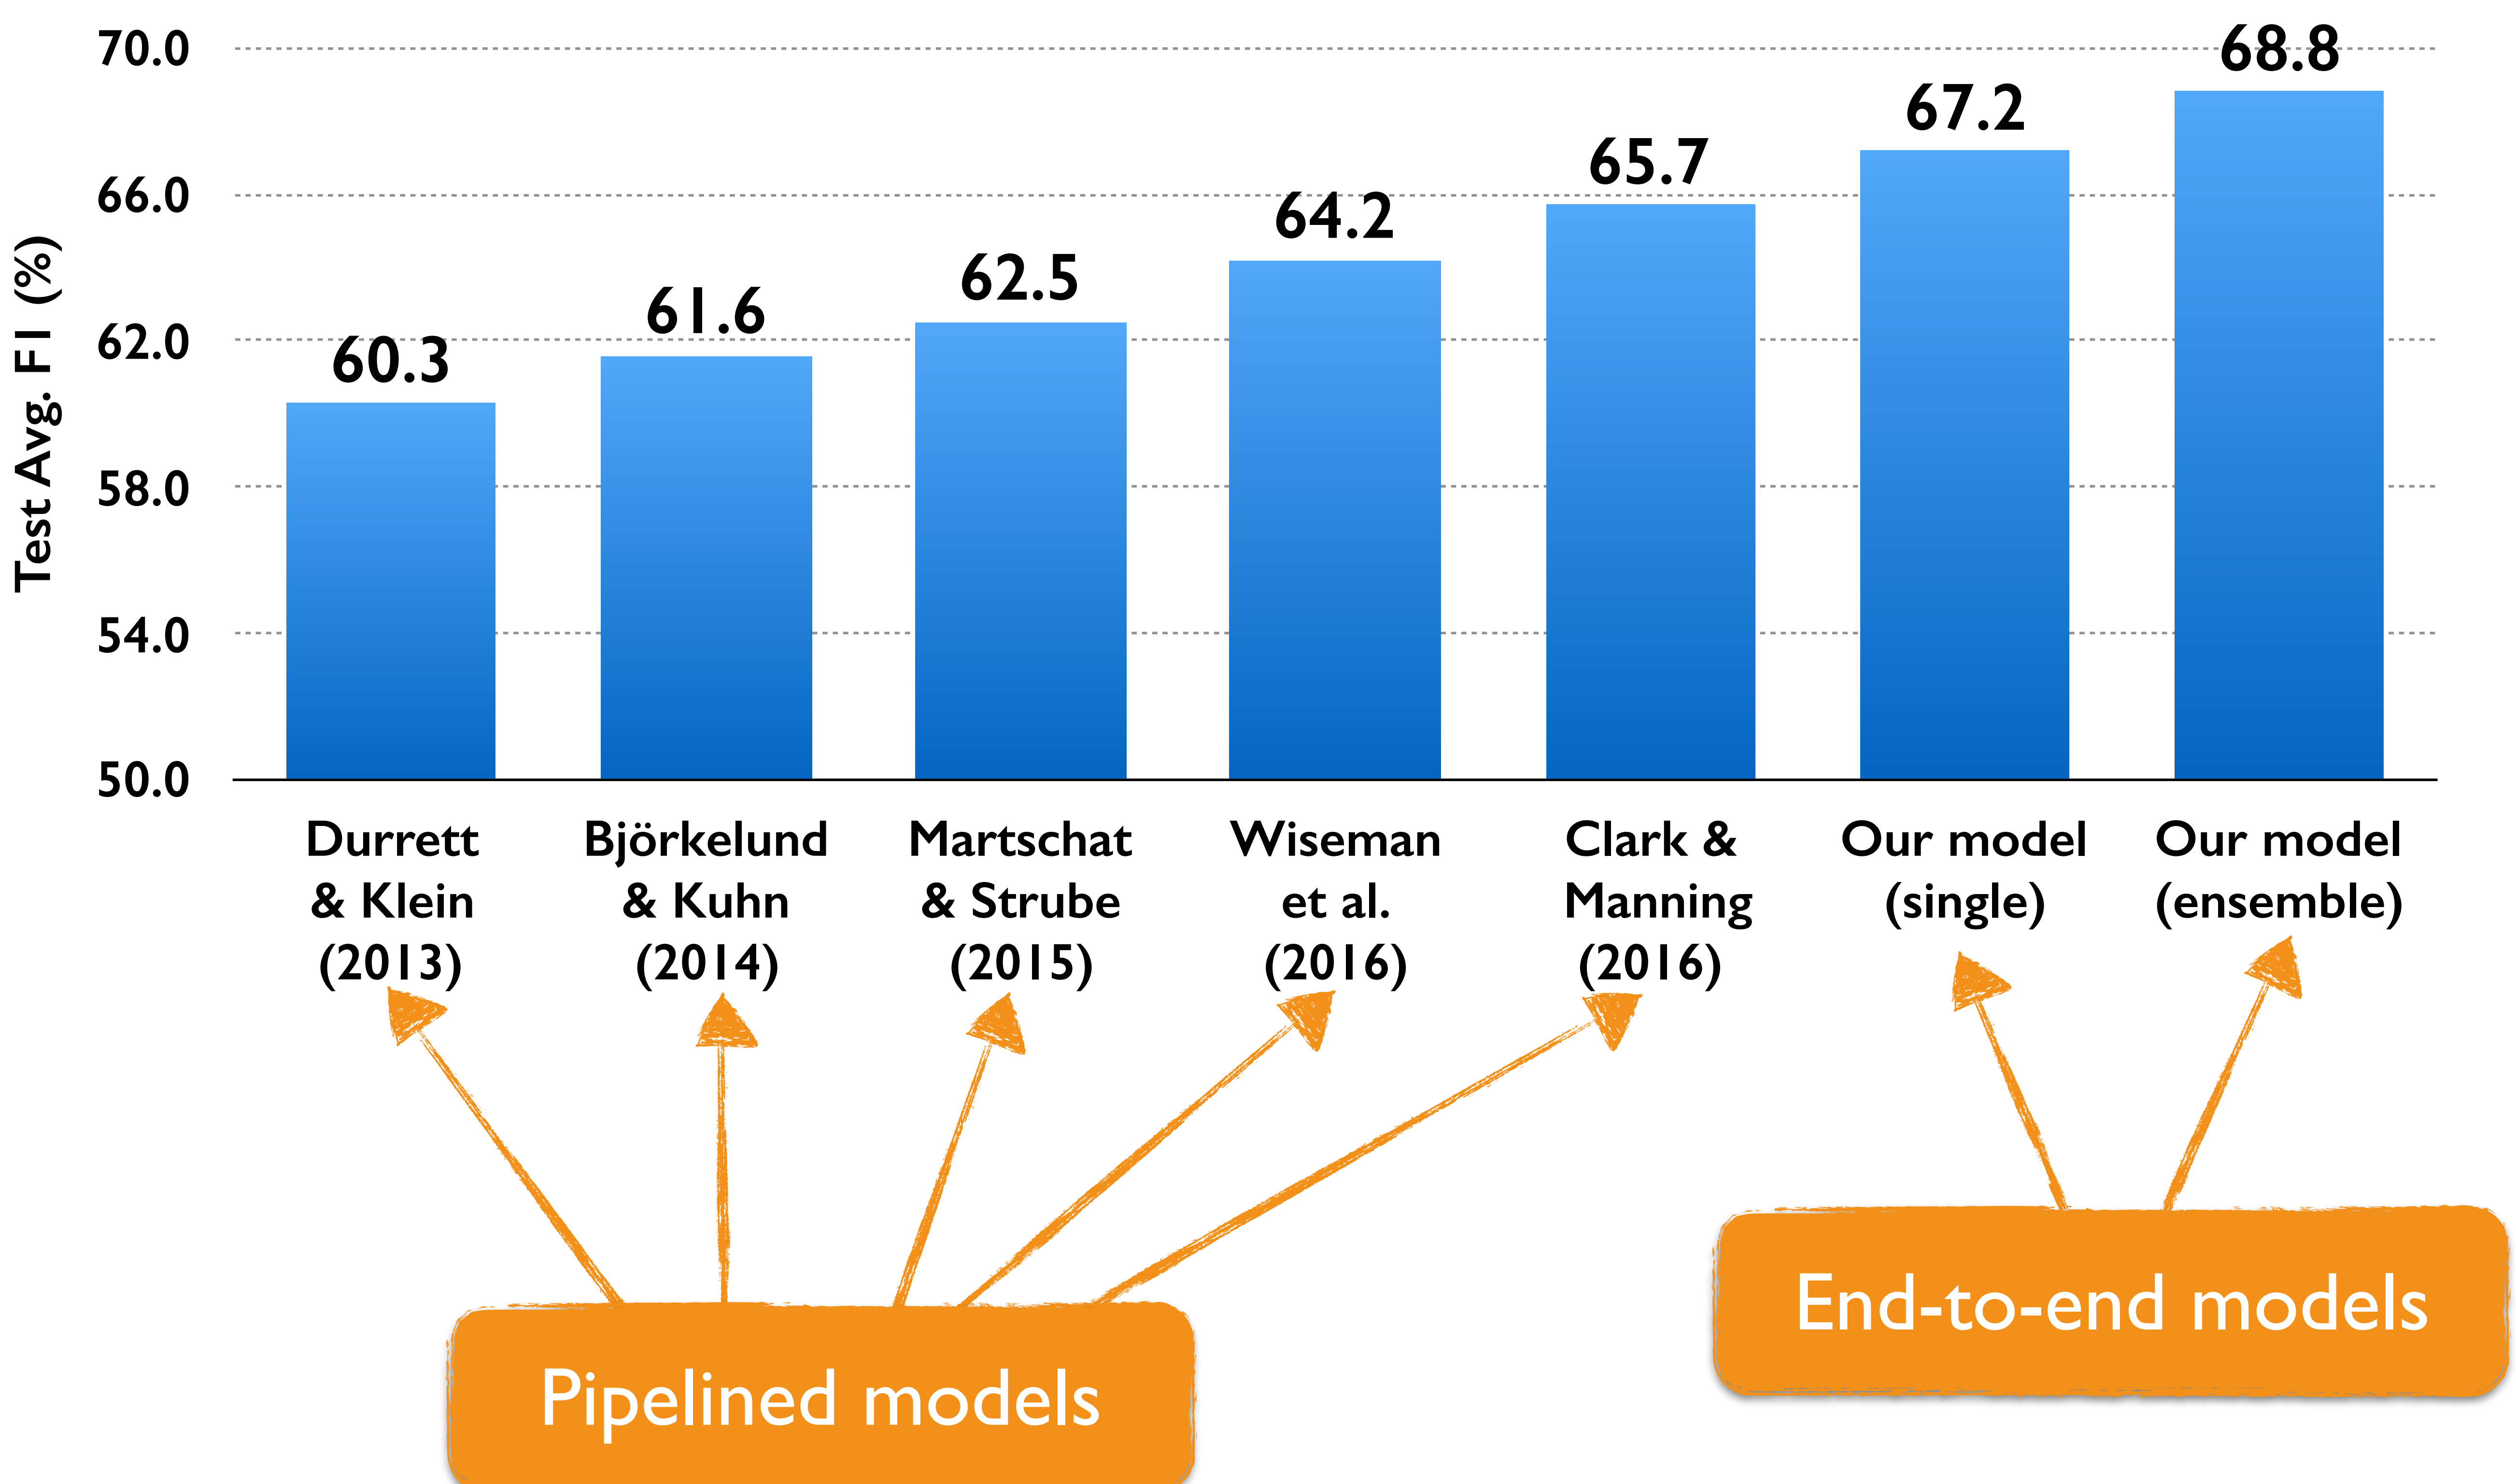

What makes a text coherent?

- Discourse/Topic structure
  - In a coherent text the parts of the discourse exhibit a sensible ordering and/or hierarchical relationship
- **Entity structure (“Focus”)**
  - **A coherent text is about some entity or entities, and the entity/entities is/are referred to in a structured way throughout the text.**
- Rhetorical structure (“coherence/discourse relations”)
  - The elements in a coherent text are related via meaningful relations

Which one is more coherent?

John went to his favorite music store to buy a piano.

He had frequented the store for many years.

He was excited that he could finally buy a piano.

He arrived just as the store was closing for the day.

John went to his favorite music store to buy a piano.

It was a store John had frequented for many years.

He was excited that he could finally buy a piano.

It was closing just as John arrived.

Entity-centric Local Coherence

- Centering Theory (Grosz et al 1995): The way entities are introduced and discussed influences coherence
  - Entities in an utterance are ranked according to salience.
    - Is an entity pronominalized or not?
    - Is an entity in a prominent syntactic position?
- Each utterance has one center (≈topic or focus).
  - Coherent discourses have utterances with common centers.
- Entity transitions capture degrees of coherence
  - (e.g., in Centering theory CONTINUE > SHIFT).
- Computational model: Entity Grid (Barzilay and Lapata 2005)

Coherence

- John hid Bill’s car keys. He was drunk.
- John hid Bill’s car keys. He likes spinach.

Outline

- Reference resolution
- **Discourse relations**

Discourse Relations

Cause

With the national announcement last week of plans to sell some breakfast items all day long, the company expects to buy even more eggs. For example, the Egg McMuffin, which uses one egg per sandwich, is among the company's most popular menu items.

Discourse Relations

Instantia
tion

With the national announcement last week of plans to sell some breakfast items all day long, the company expects to buy even more eggs.

For example, the Egg McMuffin, which uses one egg per sandwich, is among the company's most popular menu items.

Discourse Relations

With the national announcement last week of plans to sell some breakfast items all day long, the company expects to buy even more eggs. For example, the Egg McMuffin, which uses one egg per sandwich, is among the company's most popular menu items.

Two theories

Rhetorical Structure Theory (RST)
(Mann and Thompson 1988)

Corpus: RST Discourse Treebank
(Carlson et al., 2001)
385 documents from the Penn Treebank

Penn Discourse Treebank (PDTB)
(Prasad et al., 2008)

Corpus: Wall Street Journal portion of the Penn Treebank

Two theories

Rhetorical Structure Theory (RST)
(Mann and Thompson 1988)

Corpus: RST Discourse Treebank (Carlson et al., 2001)
385 documents from the Penn Treebank

Penn Discourse Treebank (PDTB)
(Prasad et al., 2008)

Corpus: Wall Street Journal portion of the Penn Treebank

Rhetorical Structure Theory (RST)

- Discourse relations "describe the relations between text parts in functional terms, identifying both the transition point of a relation and the extent of the terms related." *(Mann and Thompson 1988)*

- Document as a tree structure.

Rhetorical Structure Theory (RST)

Only the midday sun at tropical latitudes is warm enough to thaw ice on occasion, but any liquid water formed in this way would evaporate almost instantly because of the low atmospheric pressure.

Rhetorical Structure Theory (RST)

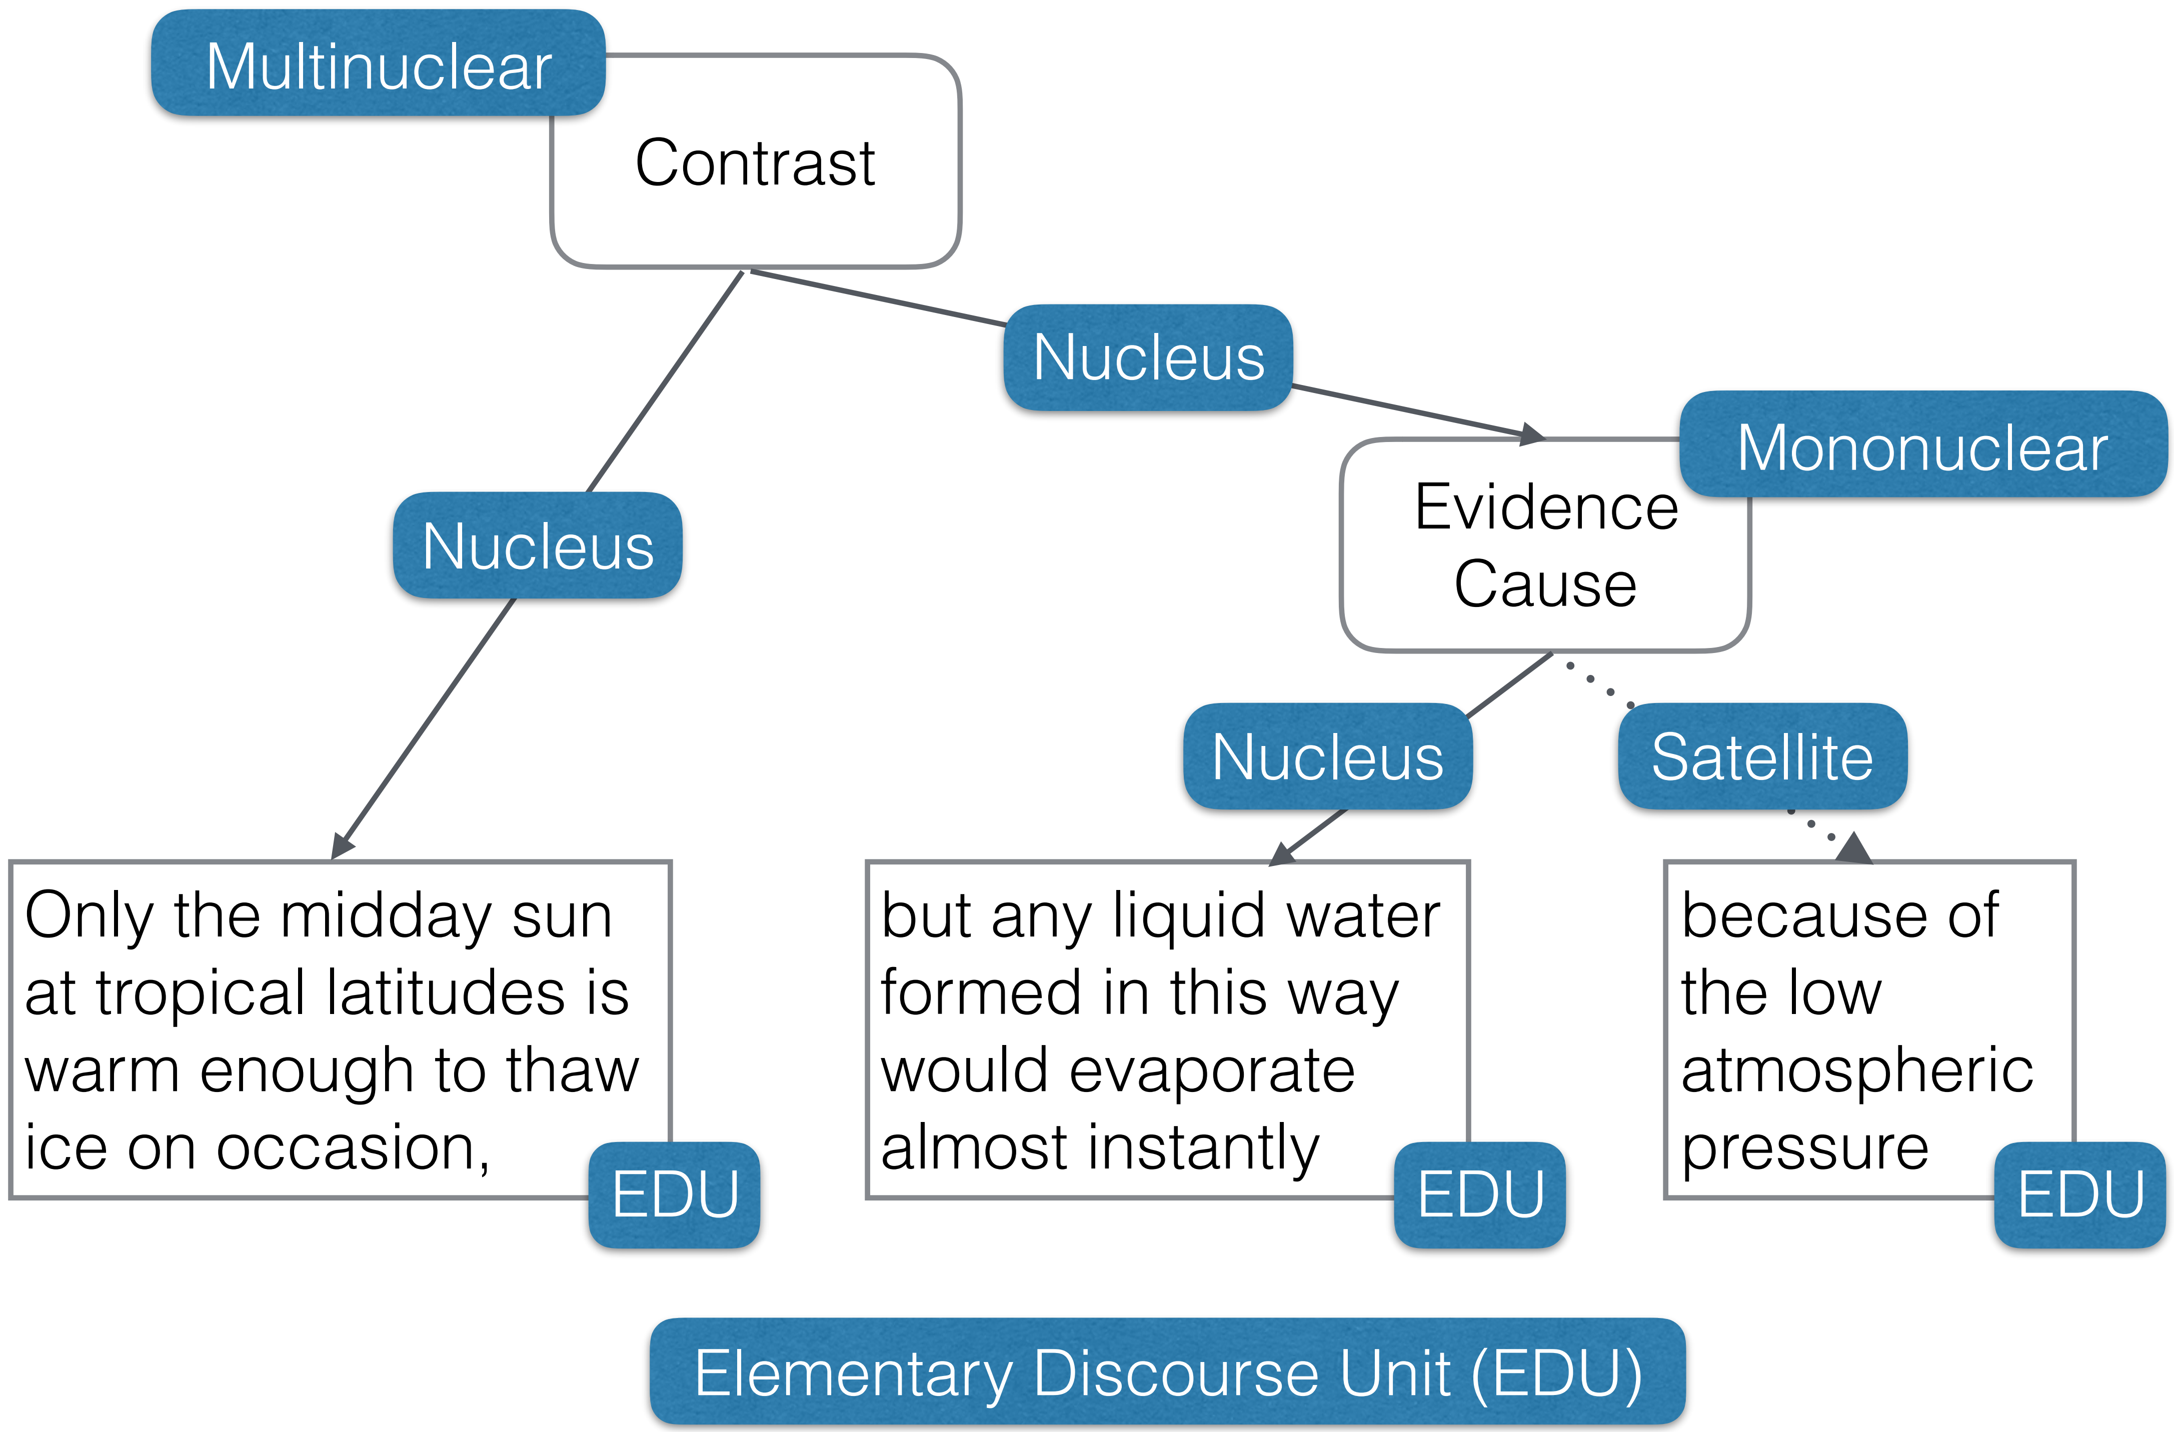

Rhetorical Structure Theory (RST)

- ***Attribution:*** attribution, attribution-negative
- ***Background:*** background, circumstance
- ***Cause:*** cause, result, consequence
- ***Comparison:*** comparison, preference, analogy, proportion
- ***Condition:*** condition, hypothetical, contingency, otherwise
- ***Contrast:*** contrast, concession, antithesis
- ***Elaboration:*** elaboration-additional, elaboration-general-specific, elaboration-part-whole, elaboration-process-step, elaboration-object-attribute, elaboration-set-member, example, definition
- ***Enablement:*** purpose, enablement
- ***Evaluation:*** evaluation, interpretation, conclusion, comment
- ***Explanation:*** evidence, explanation-argumentative, reason
- ***Joint:*** list, disjunction
- ***Manner-Means:*** manner, means
- ***Topic-Comment:*** problem-solution, question-answer, statement-response, topic-comment, comment-topic, rhetorical-question
- ***Summary:*** summary, restatement
- ***Temporal:*** temporal-before, temporal-after, temporal-same-time, sequence, inverted-sequence
- ***Topic Change:*** topic-shift, topic-drift

Units of discourse: RST

- The syntactic constructions that encode a minimum unit of meaning and/or discourse function interpretable relative to a set of contexts. (Polanyi et al., 2004)

- In practice: EDUs ~ primary clauses

[Such trappings suggest a glorious past] [but give no hint of a troubled present.]

[Xerox Corp.'s third-quarter net income grew 6.2% on 7.3% higher revenue,]
[earning mixed reviews from Wall Street analysts.]

✘ [<u>Deciding what constitutes "terrorism"</u> can be a legalistic exercise.]

✘ [He said] [the thrift will try to <u>get regulators to reverse the decision</u>.]

✘ [Once inside, she spends fours hours <u>measuring and diagramming</u> each room in the 80-year-old house,...]

Units of discourse: RST

- Same content can be packaged into varied # of EDUs
  - [Xerox Corp.'s third-quarter net income grew 6.2% on 7.3% higher revenue.] [This earned mixed reviews from Wall Street analysts.]
  - [Xerox Corp's third-quarter net income grew 6.2% on 7.3% higher revenue,] [which earned mixed reviews from Wall Street analysts.]
  - [Xerox Corp's third-quarter net income grew 6.2% on 7.3% higher revenue,] [earning mixed reviews from Wall Street analysts.]
  - [The 6.2% growth of Xerox Corp.'s third-quarter net income on 7.3% higher revenue earned mixed reviews from Wall Street analysts.]

Information salience

- Certain spans are more important and is manifested in discourse structure

Nuclearity example

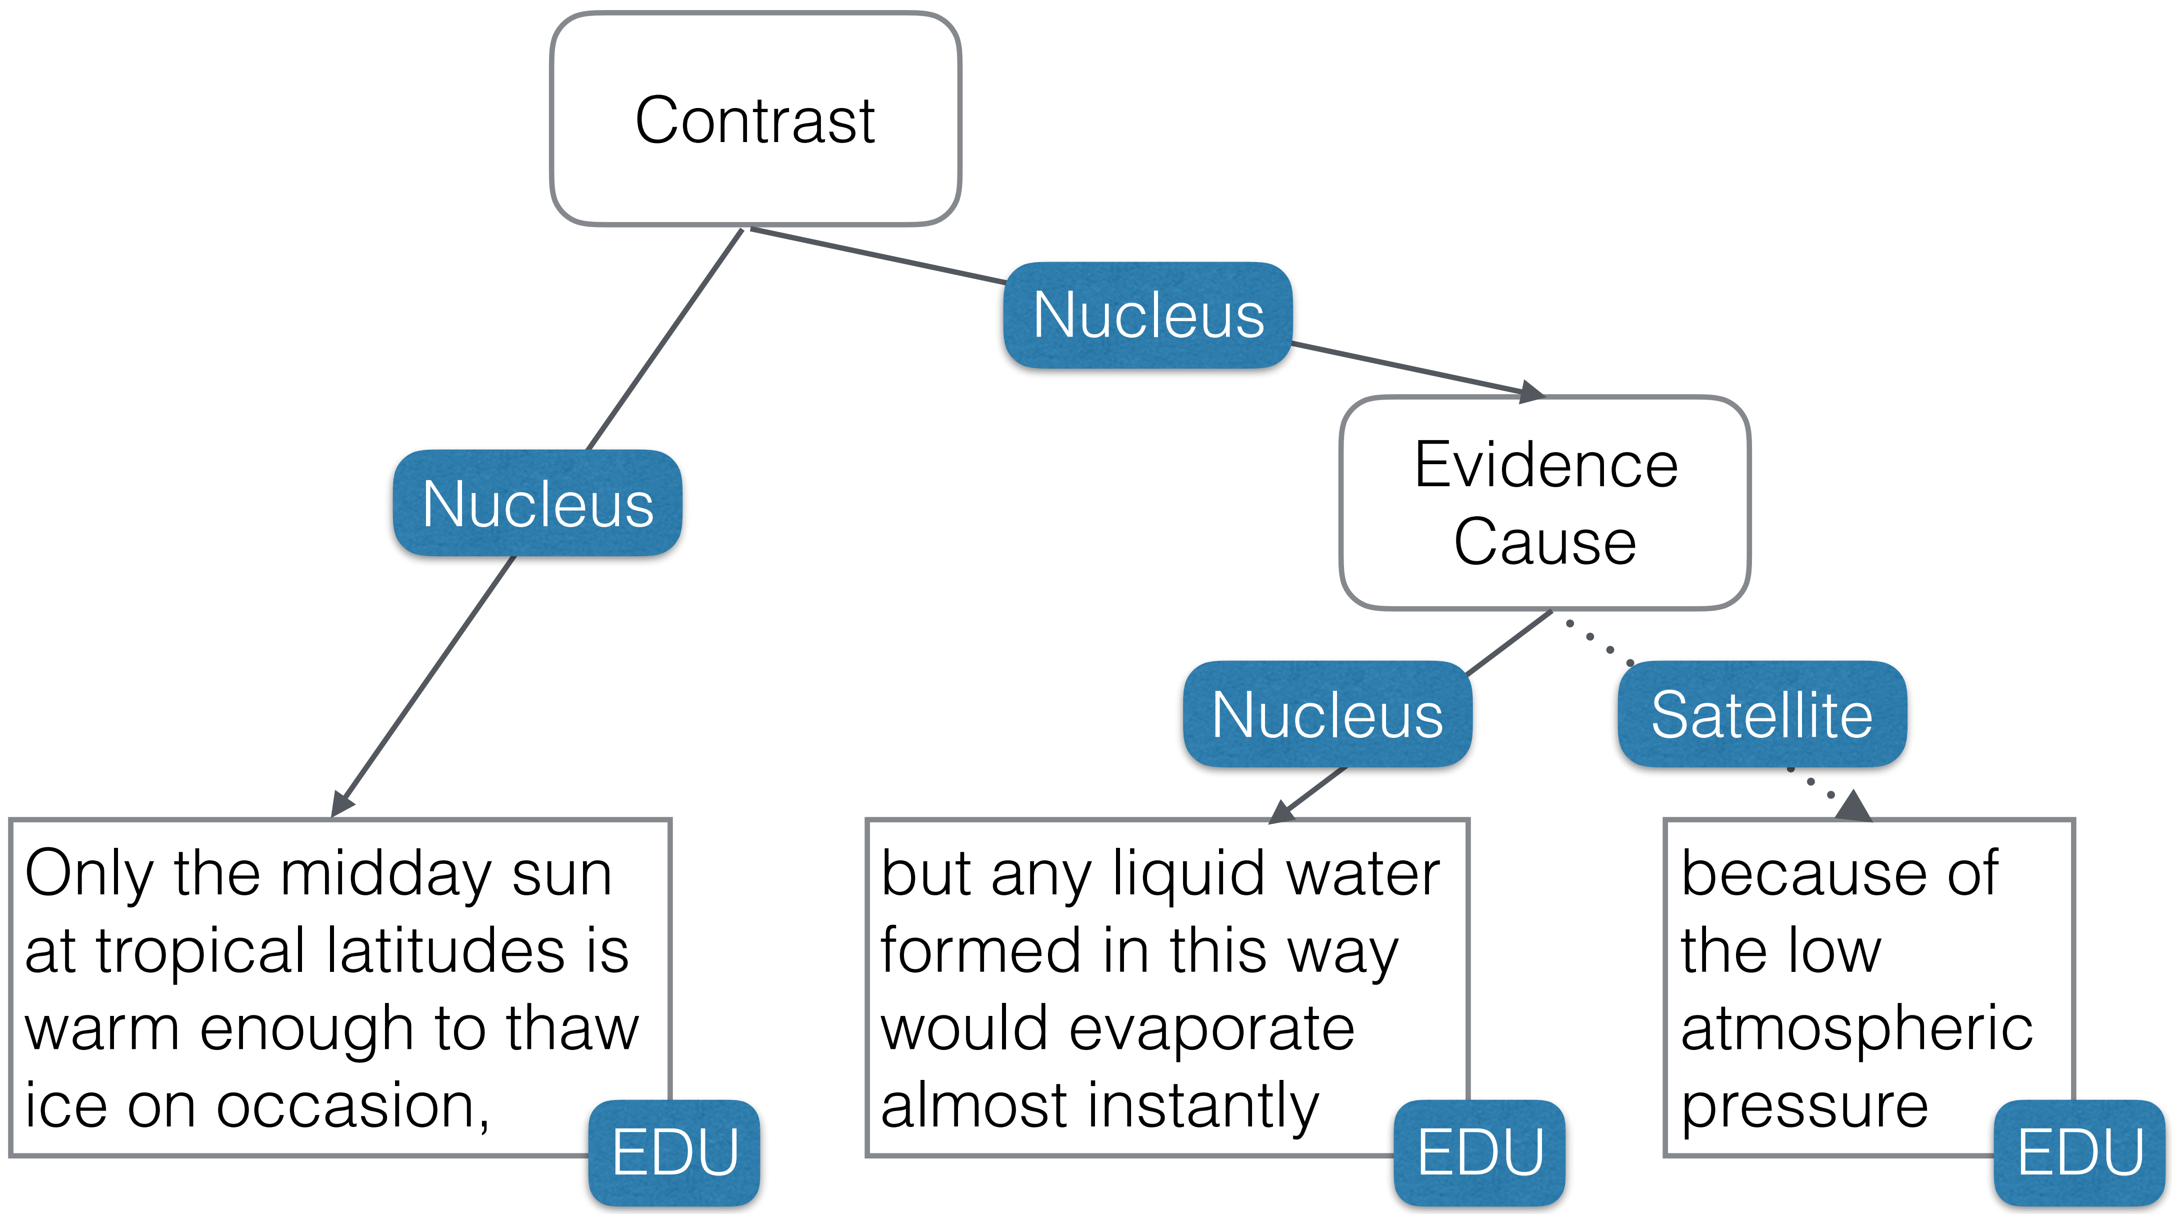

Nuclearity example

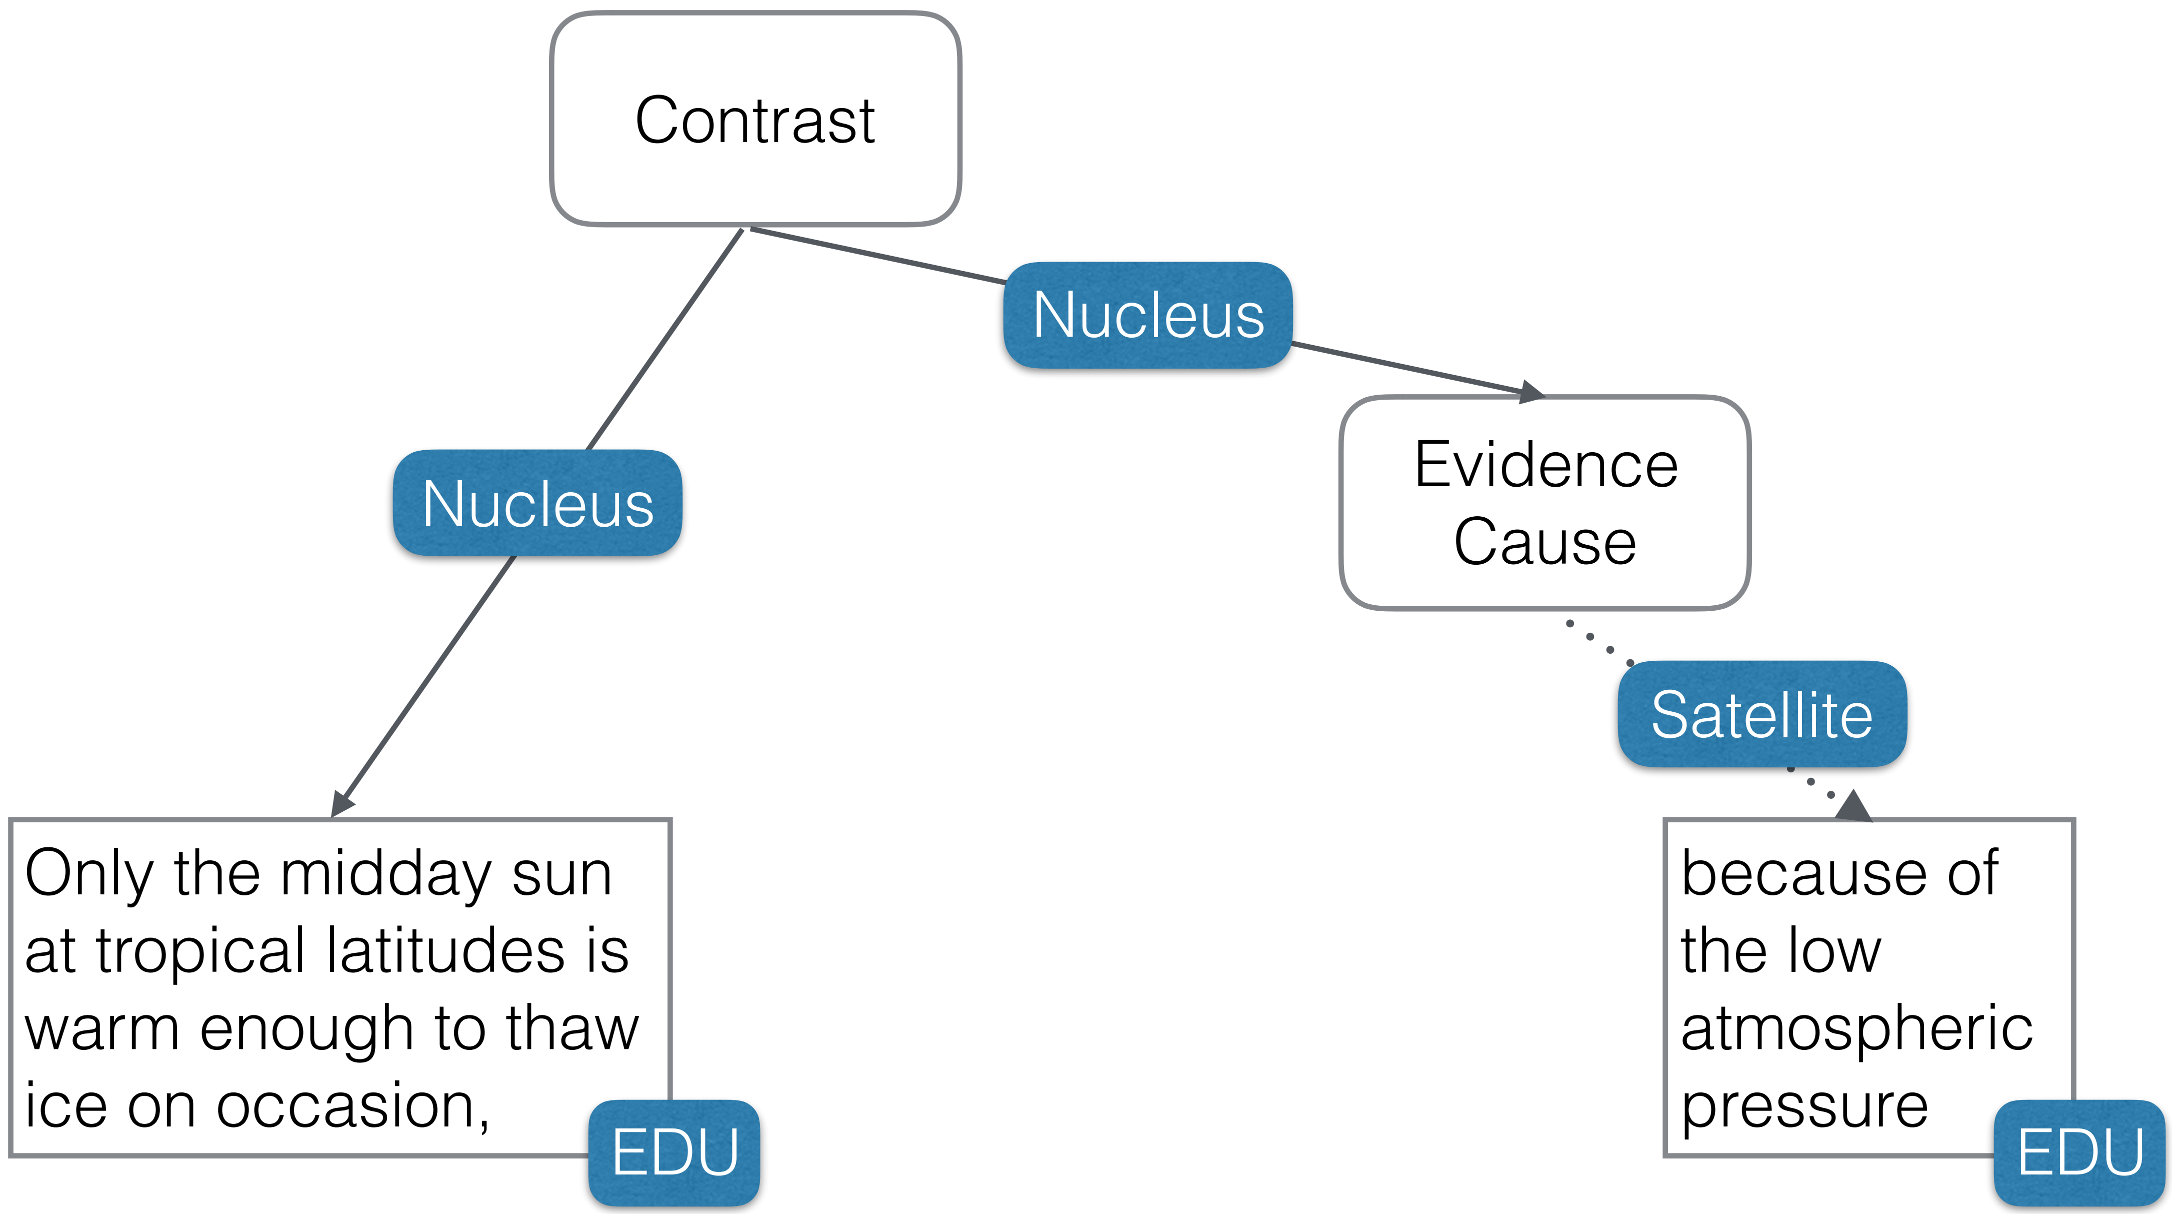

RST Parsing

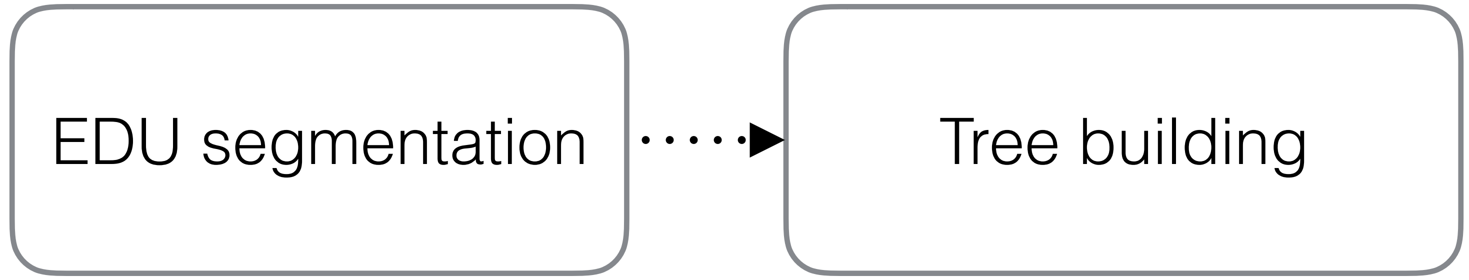

Constructing discourse trees: first attempt

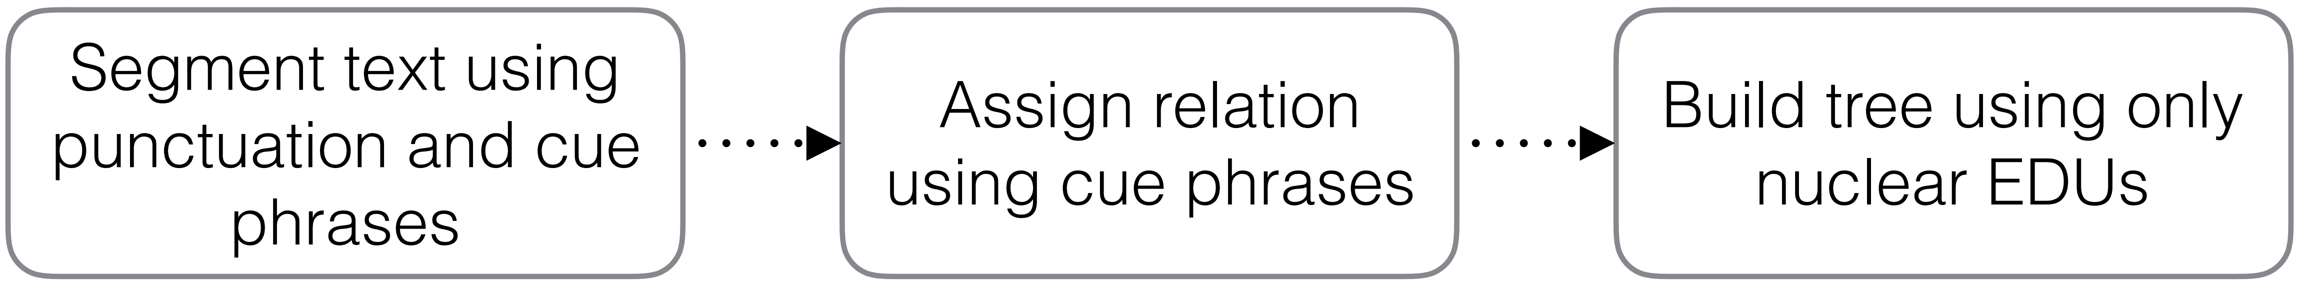

- Problems?
  - Implicit relations
  - Cue phrases have non-discourse use
  - Ambiguous connectives

EDU segmentation

[Some analysts are concerned , however ,] [that Banco Exterior may have

C C C C C C B C C C C C

waited too long] [to diversify from its traditional export-related activities .]

C C B C C C C C C C

- As a binary classification task
  - Make a decision for each token (Soricut and Marcu, 2003; Fisher and Roark, 2007; Subba and Di Eugenio, 2009; Joty et al., 2015)

- As a sequential labeling task
  - Find the most likely sequence
  - Conditional Random Fields (Hernault et al., 2010, Xuan Bach et al., 2012, Feng and Hirst 2014)

Neural approach

- Wang et al., EMNLP 2018

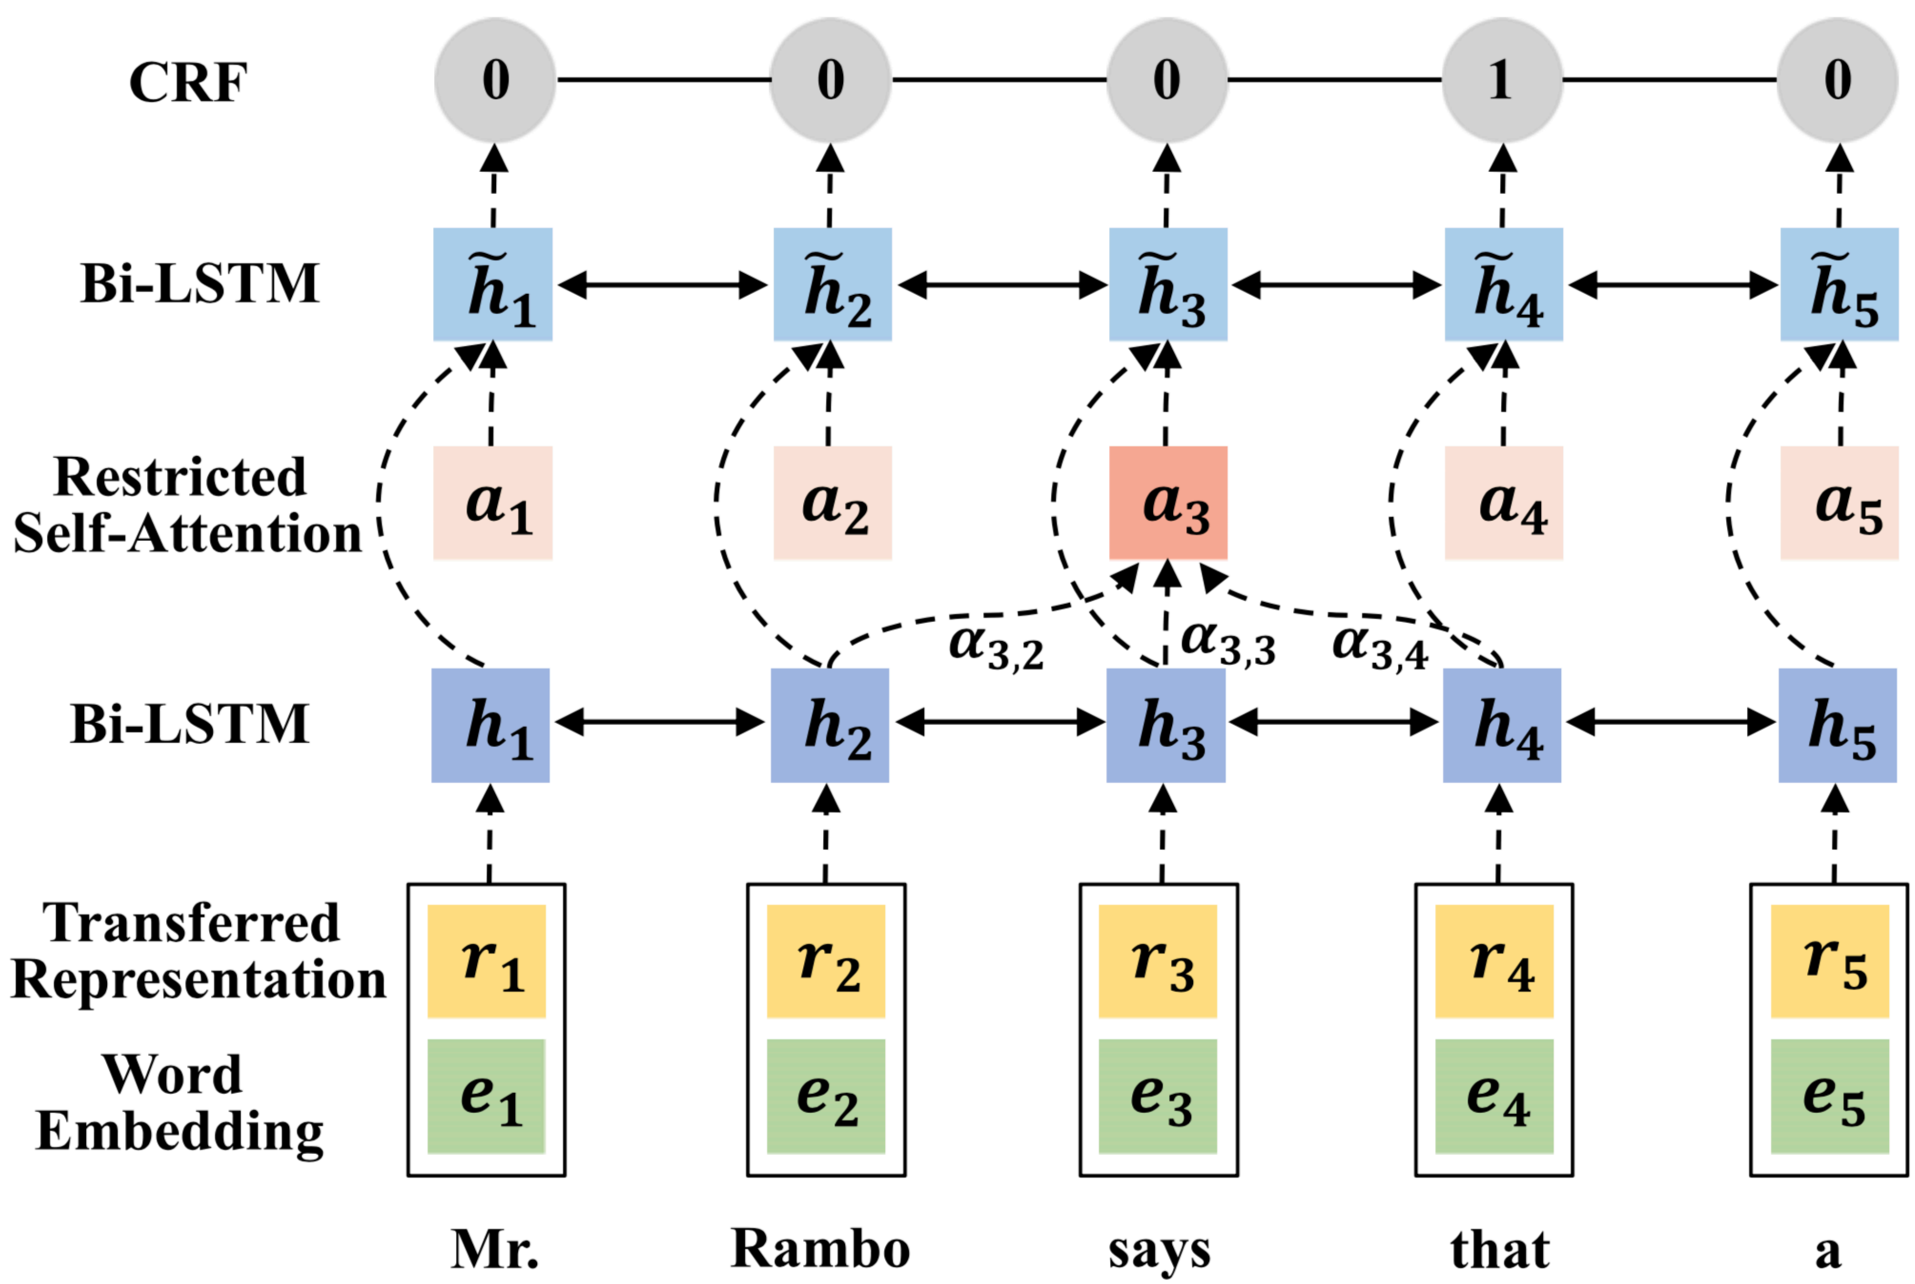

Performance

Model	Tree	P(%)	R(%)	F1(%)
SPADE	Gold	84.1	85.4	84.7
NNDS	Gold	85.5	86.6	86.0
CRFSeg	Gold	92.7	89.7	91.2
Reranking	Gold	**93.1**	94.2	93.7
CRFSeg	Stanford	91.0	87.2	89.0
CODRA	BLLIP	88.0	92.3	90.1
Reranking	Stanford	91.5	90.4	91.0
Two-Pass	BLLIP	92.8	92.3	92.6
Our Model	No	92.9	**95.7**	**94.3**
- Attention	No	92.4	94.8	93.6
- ELMo	No	87.9	84.5	86.2
- Both	No	87.0	82.8	84.8
Human	No	98.5	98.2	98.3

Constructing discourse trees

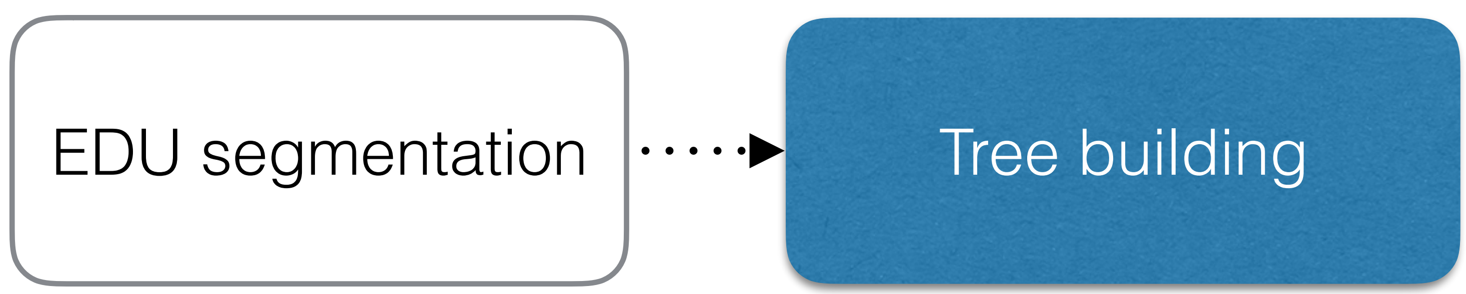

Modern RST Parsing

- Classification+CKY (Hernault et al., 2010; Feng & Hirst, 2012)
- Sequence labeling (Ghosh et al., 2012; Joty et al., 2013, Feng & Hirst, 2014)
- Shift-reduce (Sagae, 2009; Ji & Eisenstein, 2014; Heilman & Sagae, 2015)
- Representation learning (Ji & Eisenstein, 2014; Li et al., 2014)

CIDER

- Liu and Lapata, EMNLP 2017
- Two-stage approach:
  - parse each sentence in a document into a tree whose leaves correspond to EDUs,
  - then parse the document into a tree whose leaves correspond to already pre- processed sentences

Intra-sentential parser

- Linear-chain CRF
- Separate CRFs for structure and relations

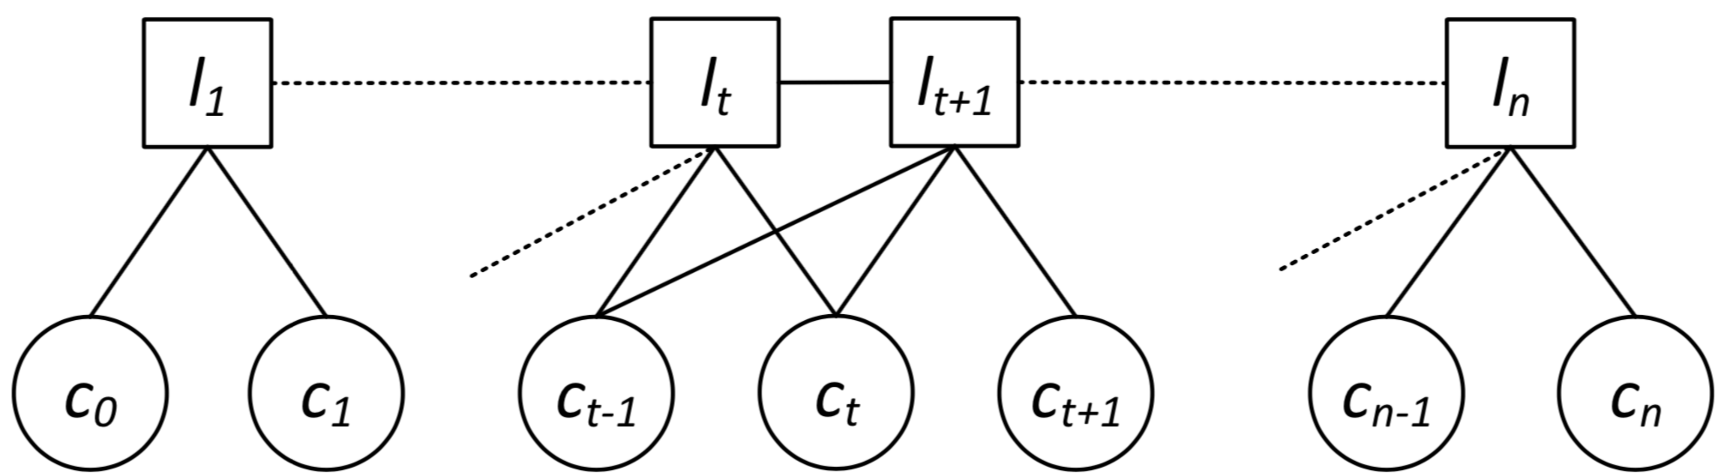

Structure parser.

C: constituent; L: latent structure nodes

Intra-sentential parser

- Linear-chain CRF
- Separate CRFs for structure and relations

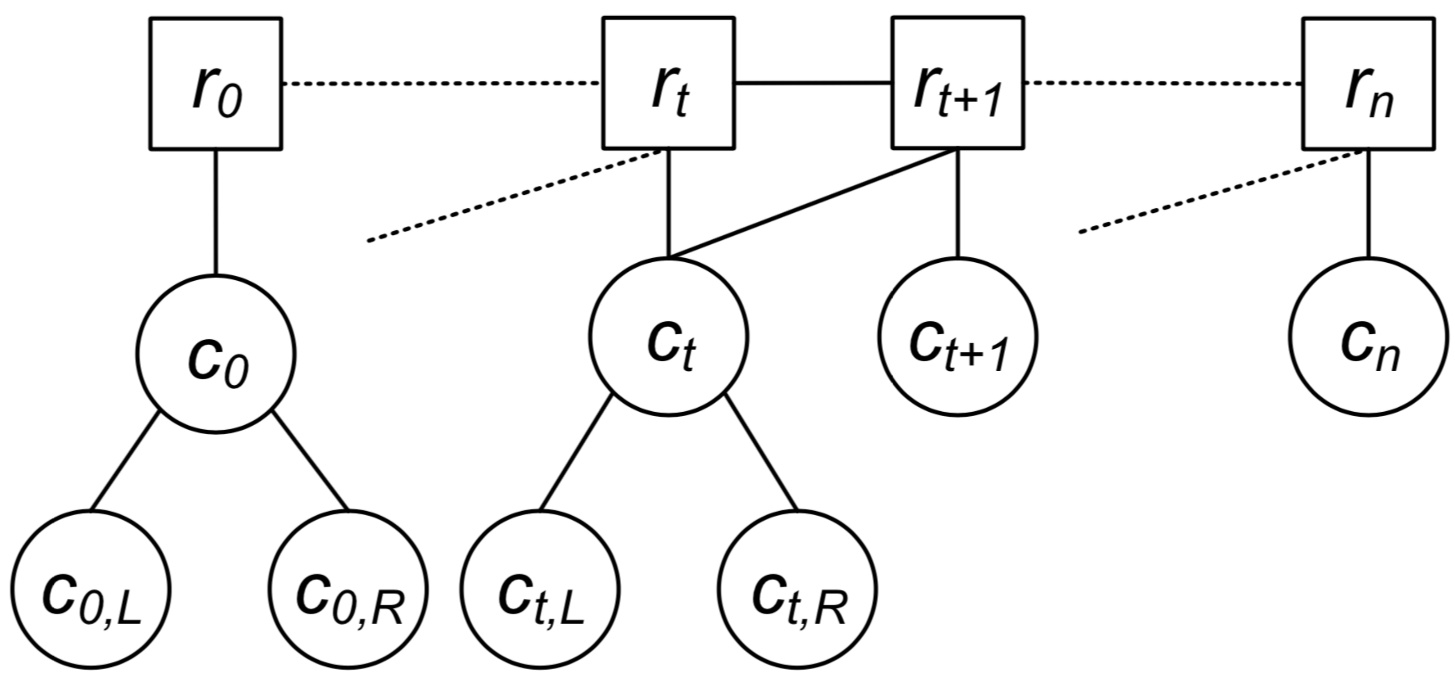

Relation parser.

C: constituent; R: relation nodes

Representation Learning: intra-sentential

- Learn EDU representations via LSTMs

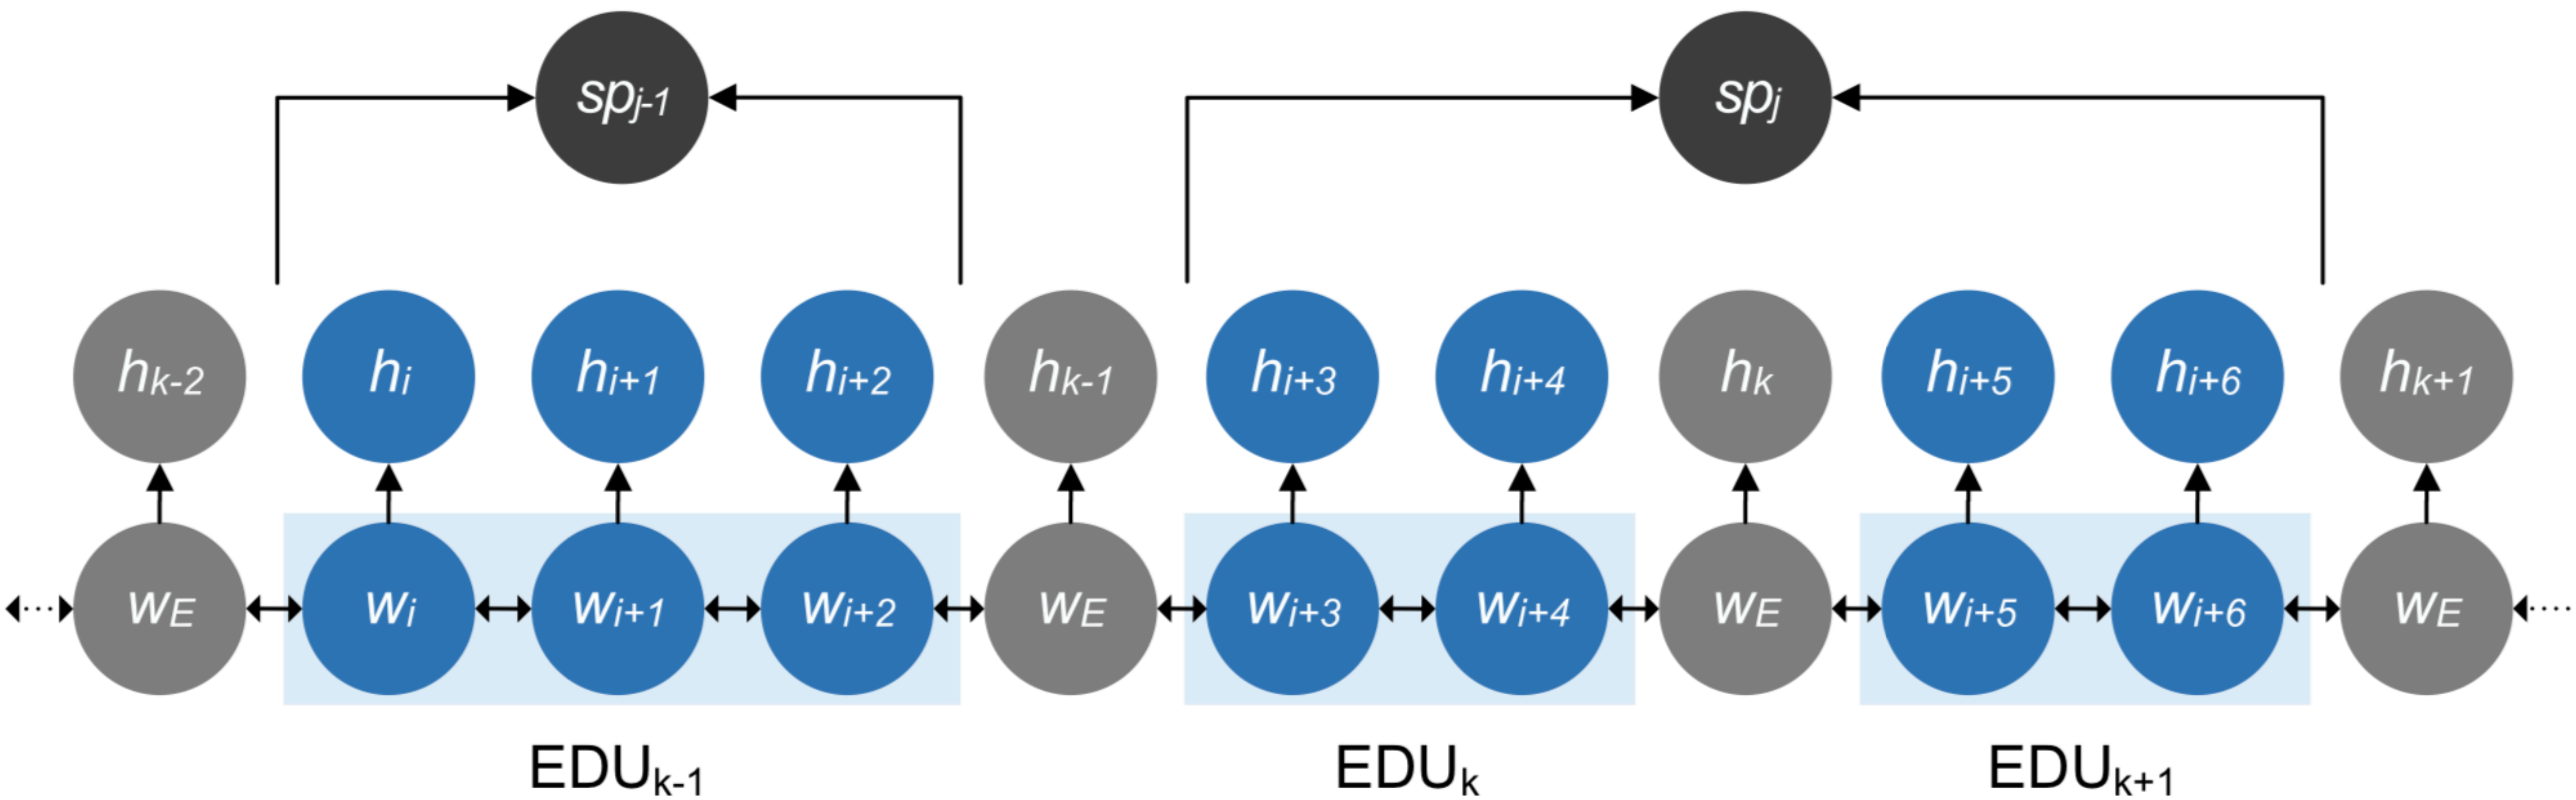

Document level parser

- Can we do something similar?

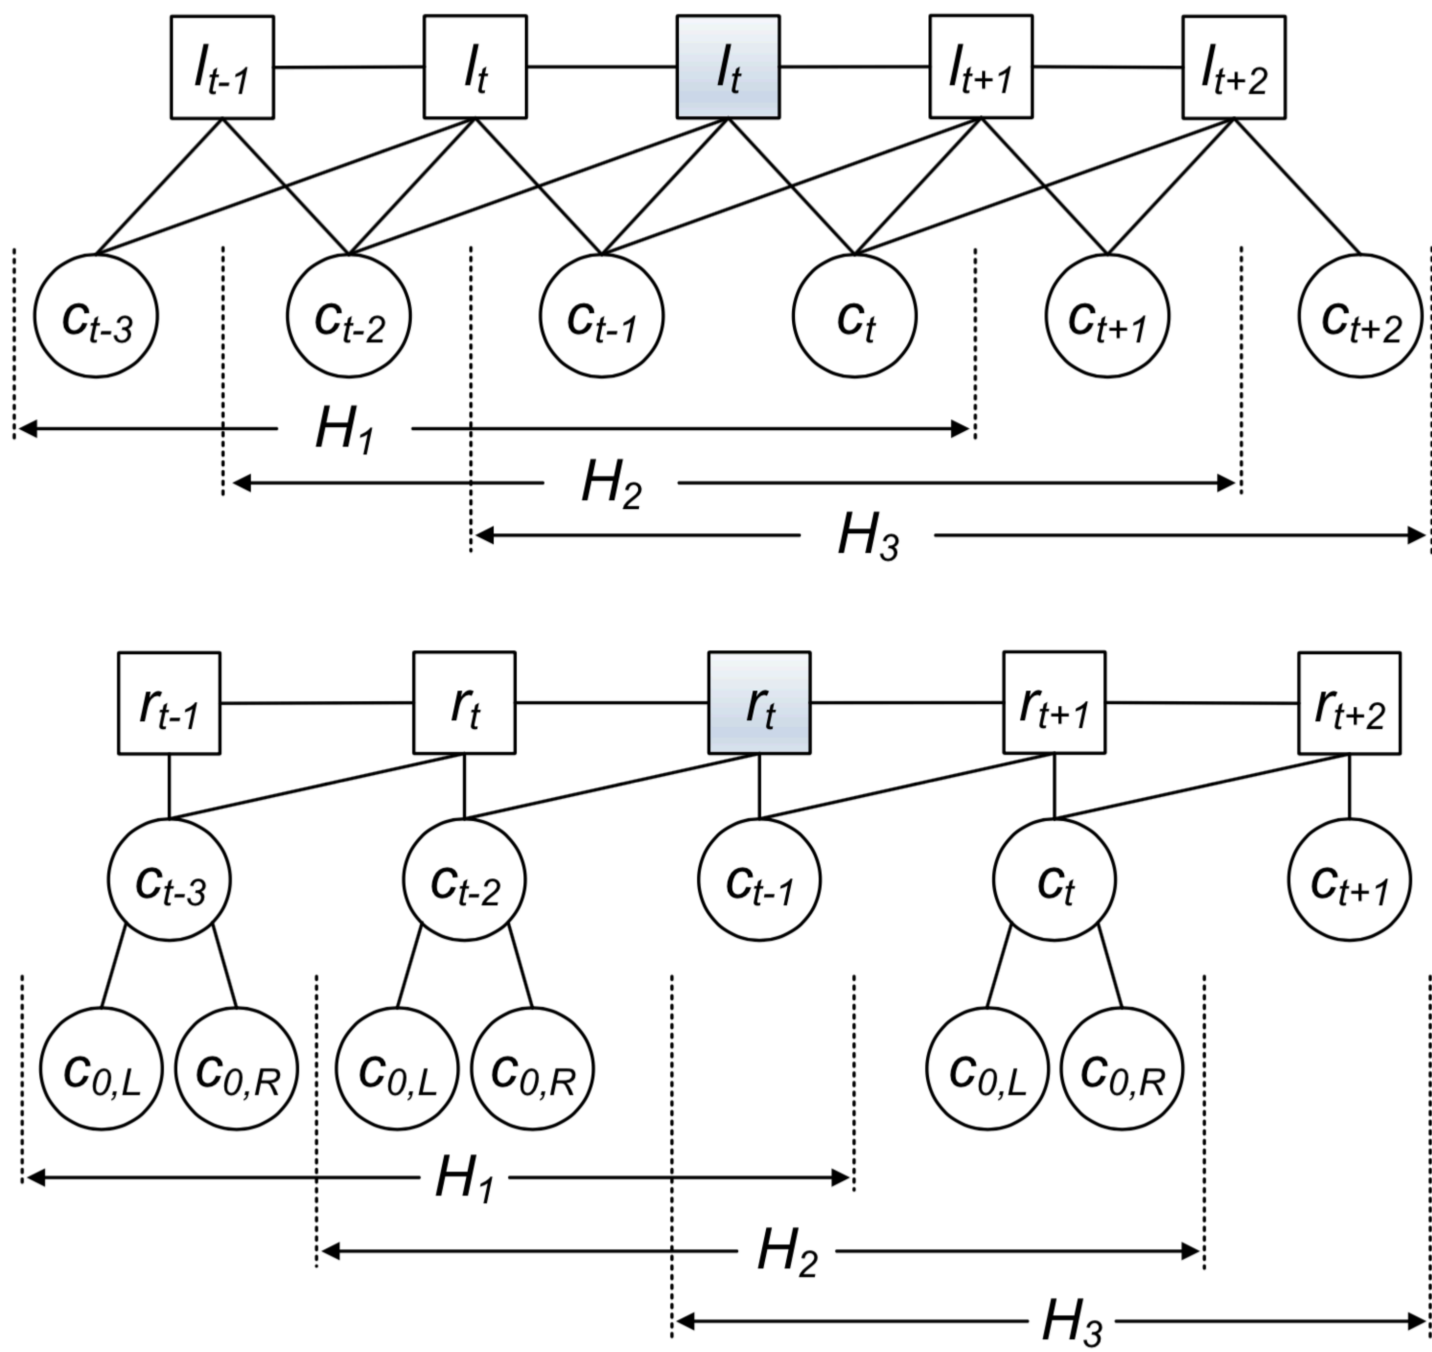

Representation Learning: doc-level

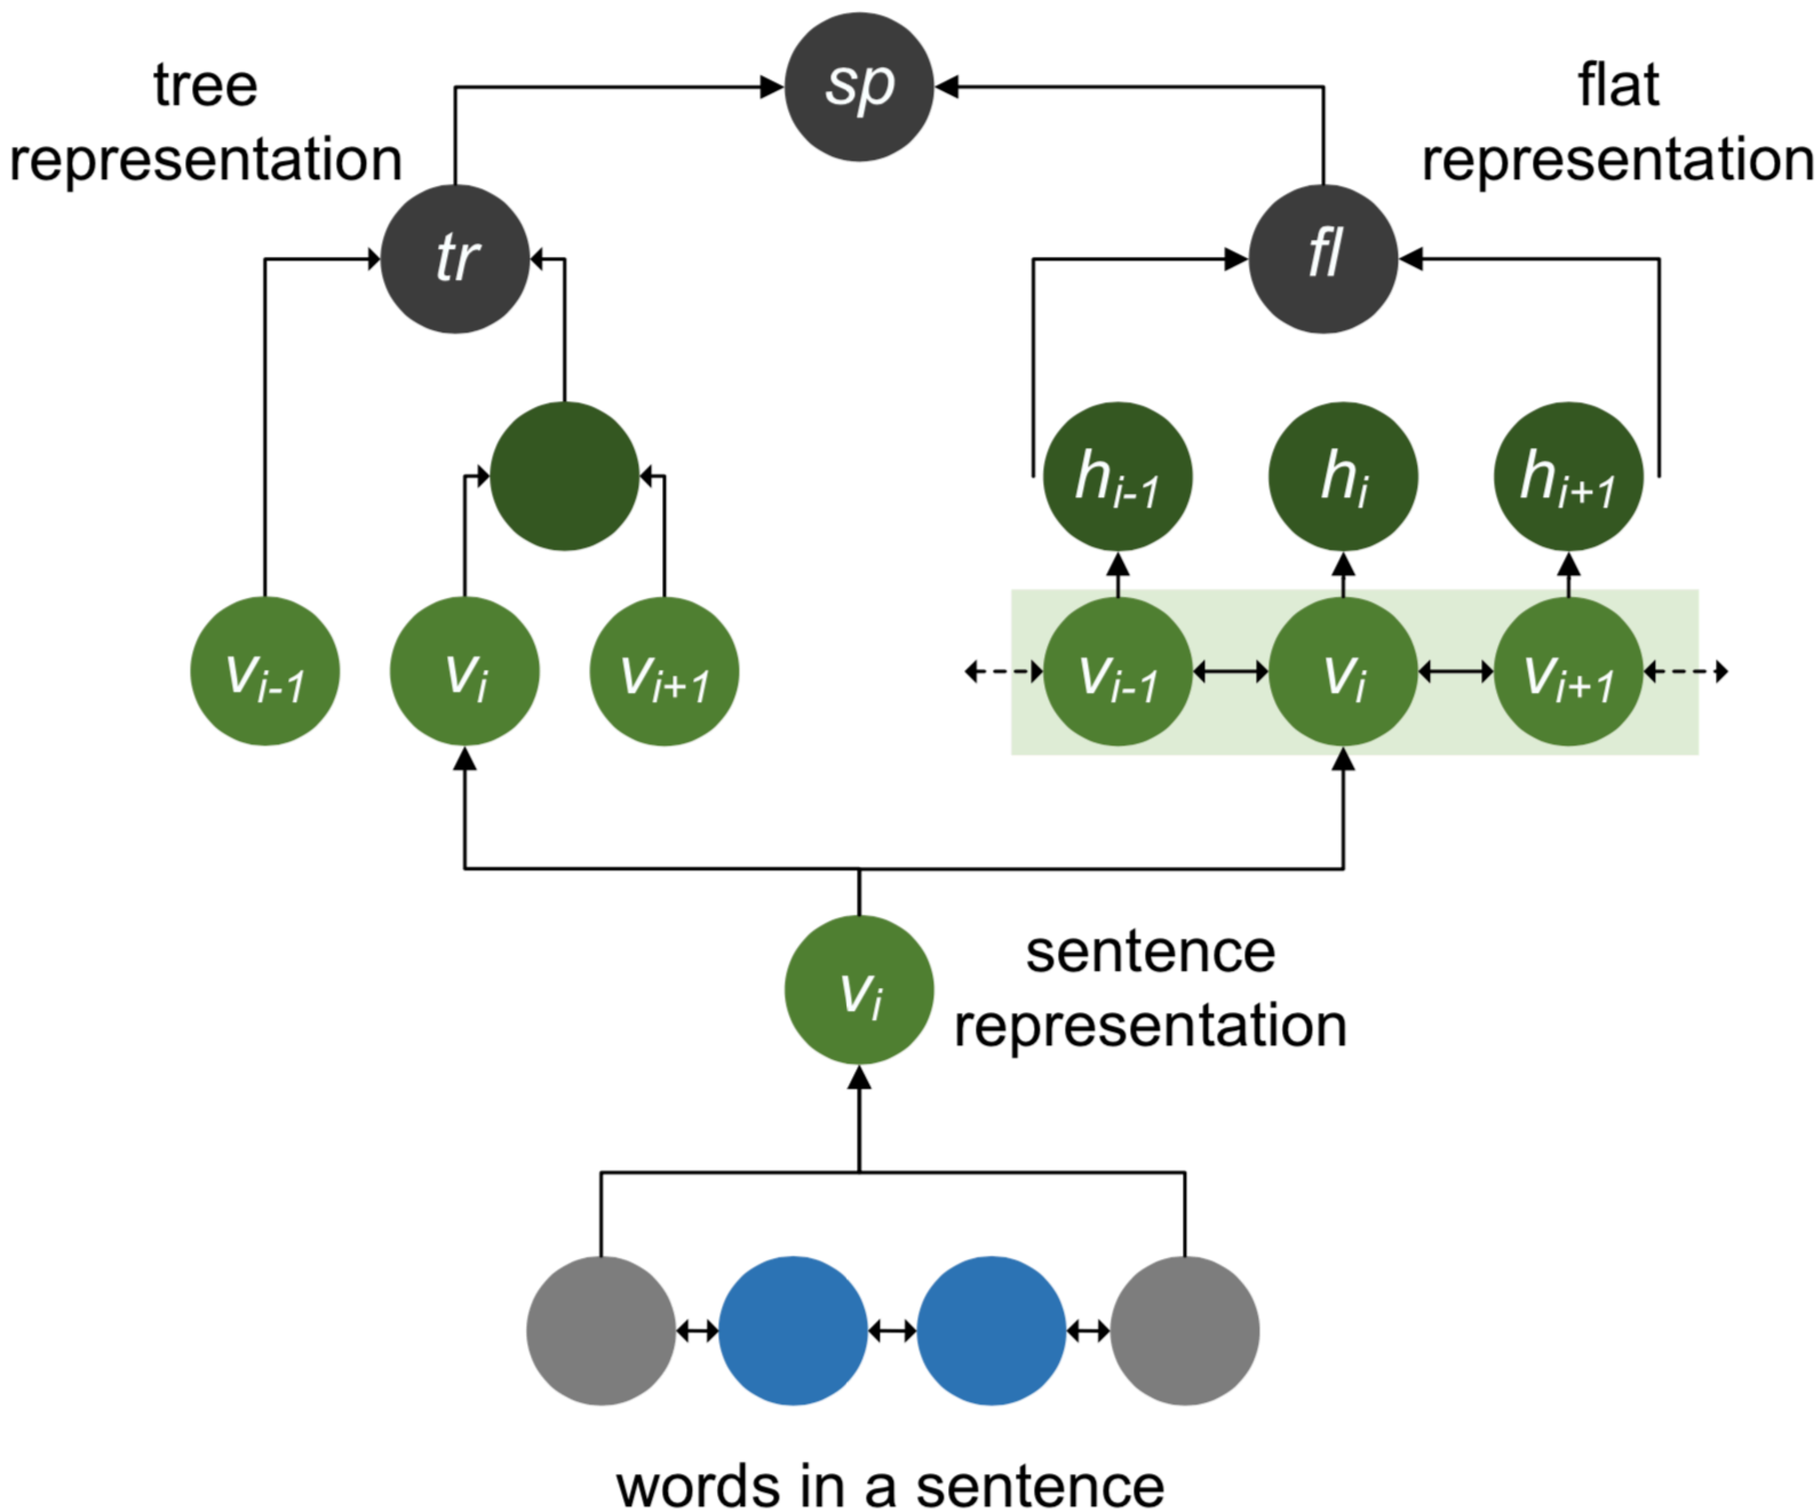

Performance

- How many things should we evaluate?
  - Tree skeleton (each subtree span)
  - Discourse relations
  - Separate nuclearity from discourse relations
- Should we use gold-standard human EDU segmentation? Why/why not?

Performance

Discourse Parsers	S	N	R	Speed
Ji and Eisenstein (2014)	82.1	71.1	61.6	0.21
Feng and Hirst (2014)	85.7	71.0	58.2	9.88
Heilman and Sagae (2015)	83.5	68.1	55.1	0.40
CIDER (−AF)	83.6	71.1	57.3	3.80
CIDER (+AF)	85.0	71.1	59.0	3.80
Li et al. (2014)	84.0	70.8	58.6	26.00
Li et al. (2016)	85.8	71.7	58.9	–
Human	88.7	77.7	65.8	–

Catch!

- What is it that this system will **fail to** handle?
- Leaky units!

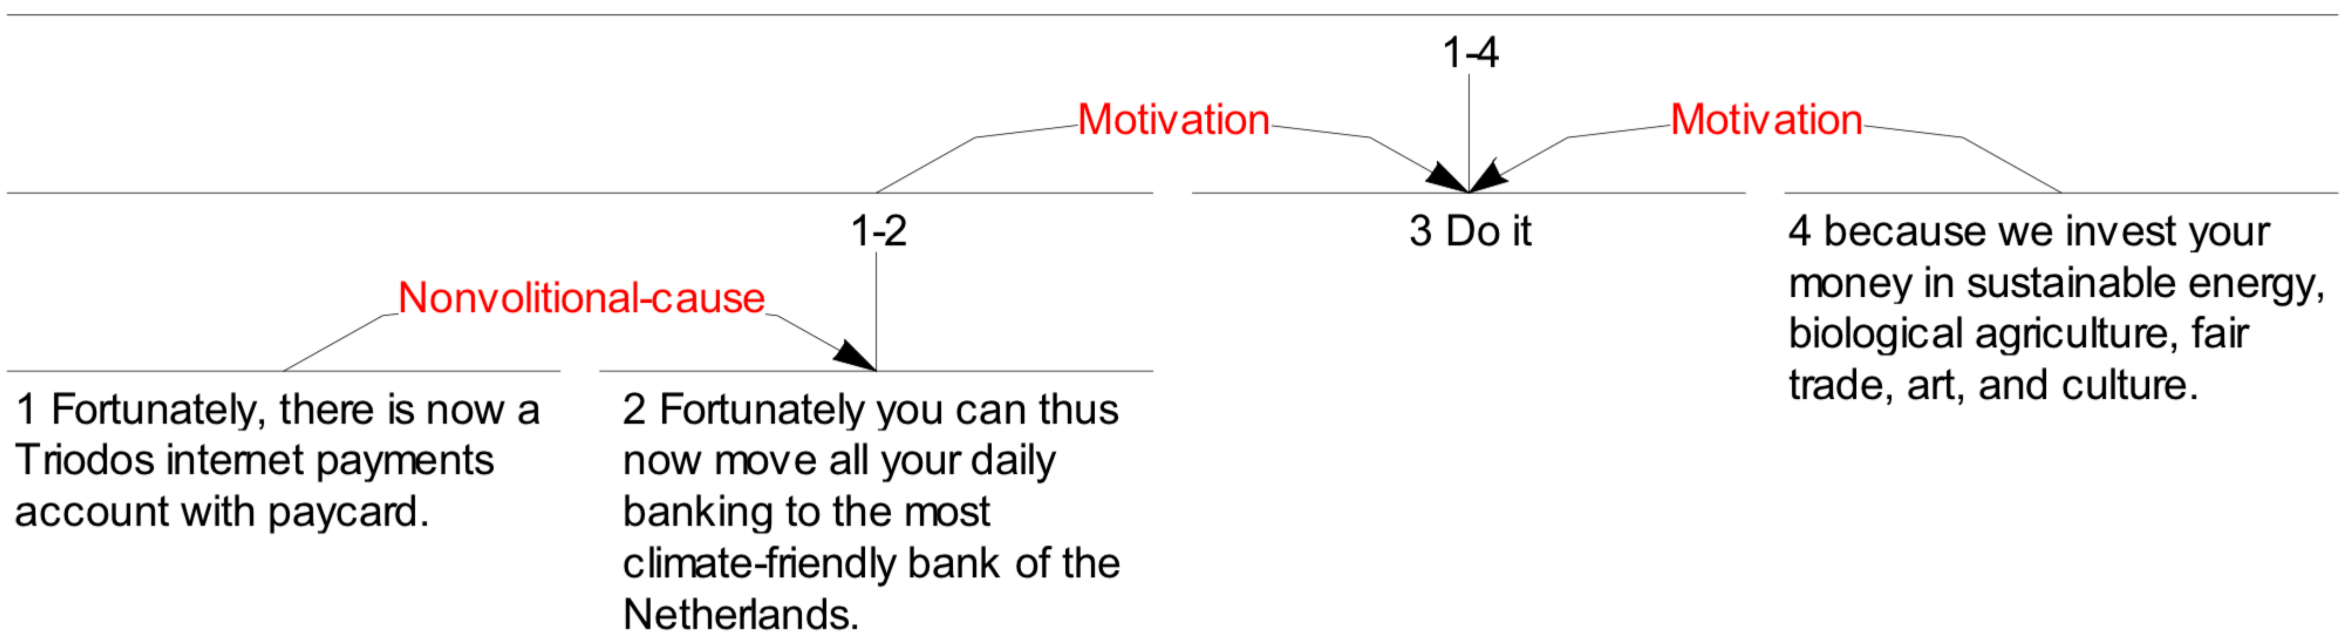

~5% in RST-DT, but much much more in other genres!

Two theories

Rhetorical Structure Theory (RST)
(Mann and Thompson 1988)

Corpus: RST Discourse Treebank
(Carlson et al., 2001)
385 documents from the Penn Treebank

Penn Discourse Treebank (PDTB)
(Prasad et al., 2008)

Corpus: Wall Street Journal portion of the Penn Treebank

Penn Discourse Treebank (PDTB)

- Discourse relations are “predicate-argument relations between two abstract objects such as events, states and propositions.” *Mitsakaki et al. (2008)*

- Theory-neutral: a flat, linear structure

- Lexically grounded

Penn Discourse Treebank (PDTB)

Arg1

Explicit

The federal government suspended sales of U.S. savings bonds **because** Congress hasn't lifted the ceiling on government debt.

Arg2

Penn Discourse Treebank (PDTB)

Explicit

Non-explicit

Implicit

I never gamble too far.
Implicit=In particular I quit after one try, whether I win or lose.

Penn Discourse Treebank (PDTB)

Explicit

Non-explicit

Implicit

AltLex

So Seita has introduced blonde cigarettes under the Gauloises label, and intends to relaunch the unsuccessful Gitanes Blondes in new packaging. AltLex **The aim is** to win market share from imported cigarettes, and to persuade smokers who are switching to blonde cigarettes to keep buying French.

Penn Discourse Treebank (PDTB)

Explicit

Non-explicit

Implicit

AltLex

EntRel

Proceeds from the offering are expected to be used for remodeling the company*'s Desert Inn resort in Las Vegas, refurbishing certain aircraft of the MGM Grand Air unit, and to acquire the property for the new resort.* EntRel **The company said it estimates the Desert Inn remodeling will cost about $32 million, and the refurbishment of the three DC-8-62 aircraft, made by McDonnell Douglas Corp., will cost around $24.5 million.**

Penn Discourse Treebank (PDTB)

Explicit

Non-explicit

- Implicit
- AltLex
- EntRel

NoRel

Jacobs is an international engineering and construction concern. NoRel **Total capital investment at the site could be as much as $400 million, according to Intel.**

Penn Discourse Treebank (PDTB)

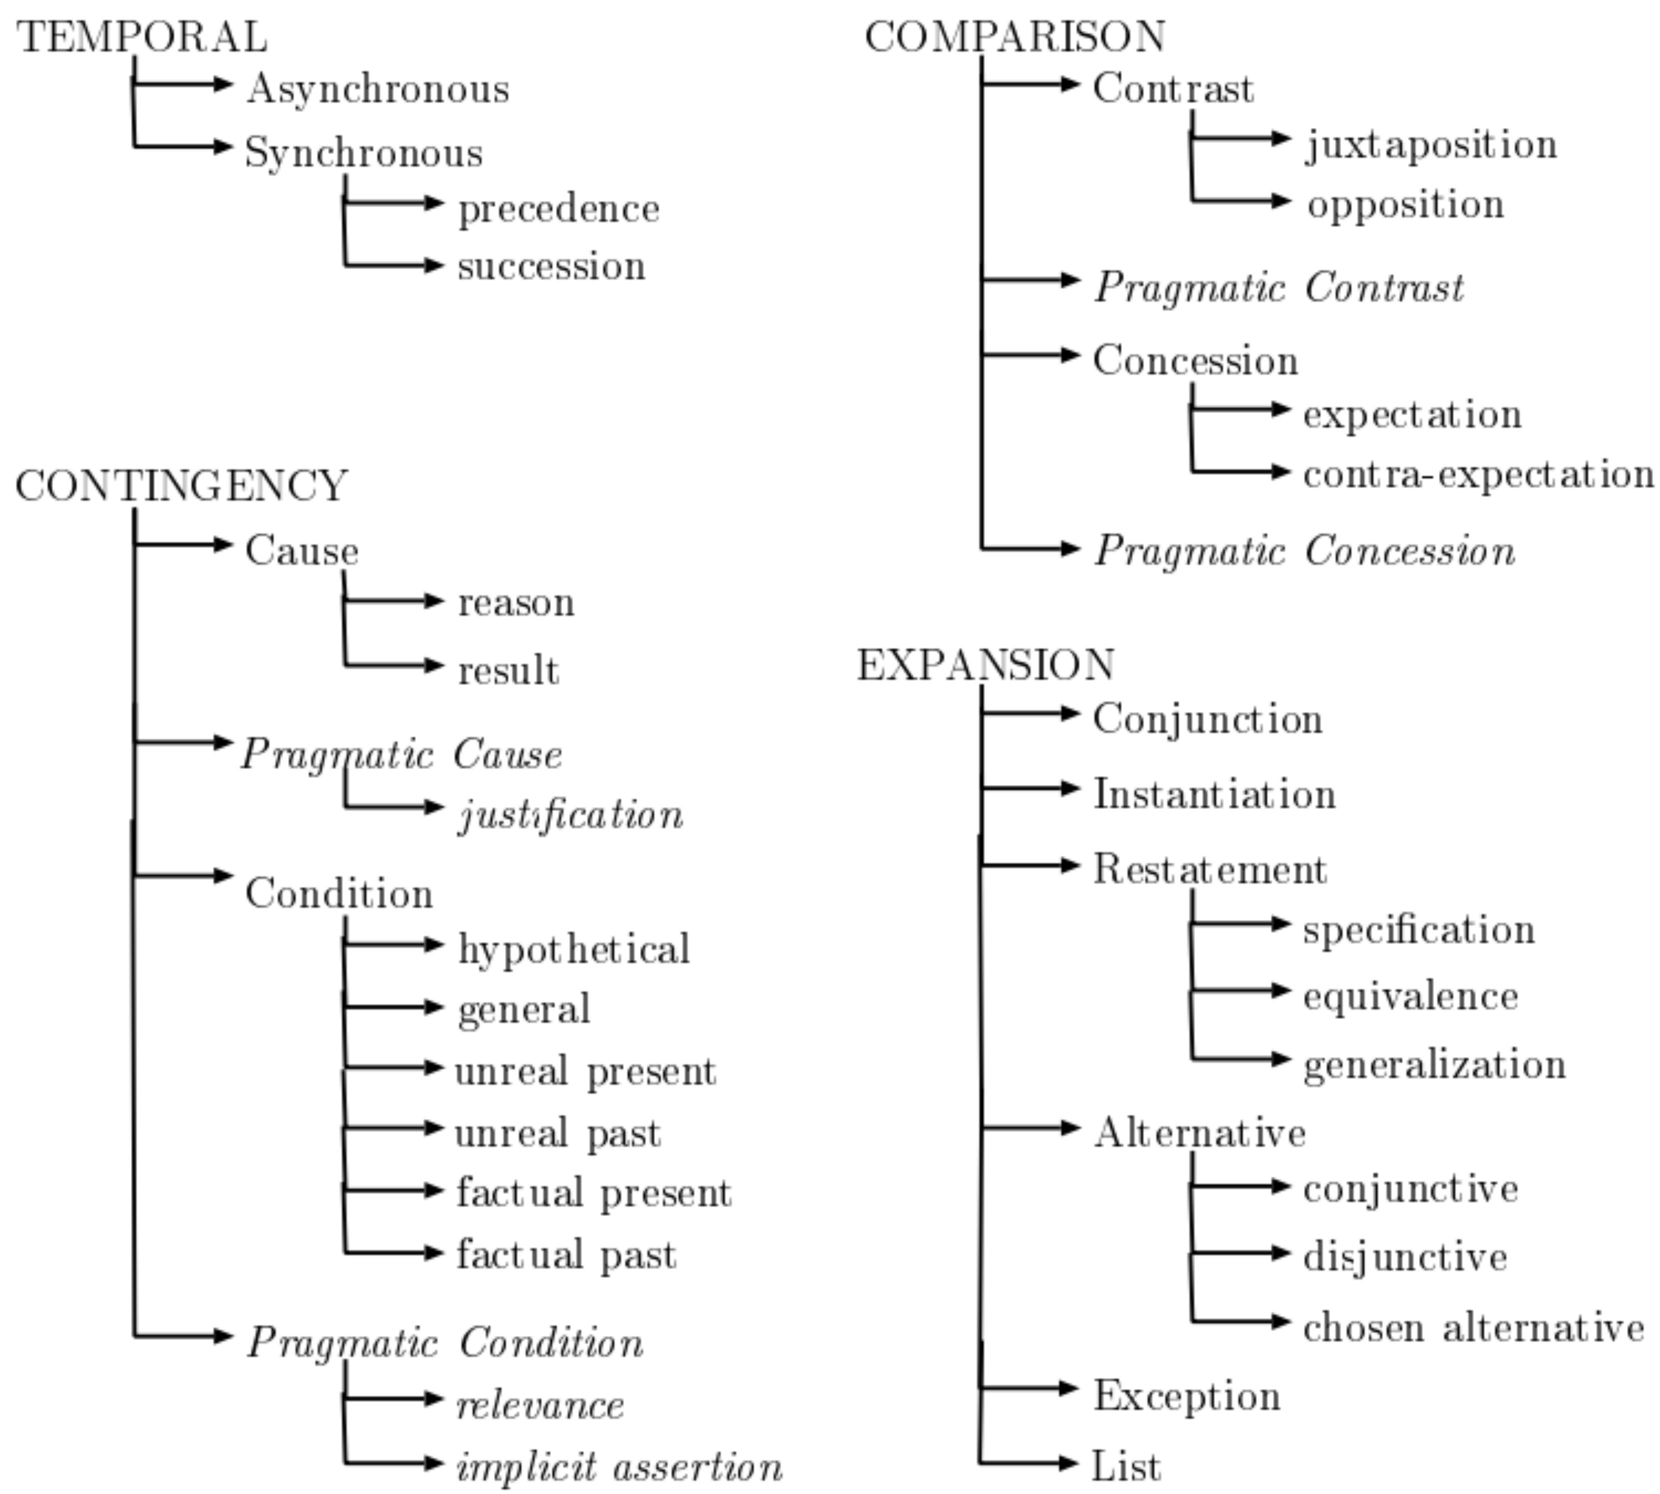

Shallow discourse parsing: PDTB

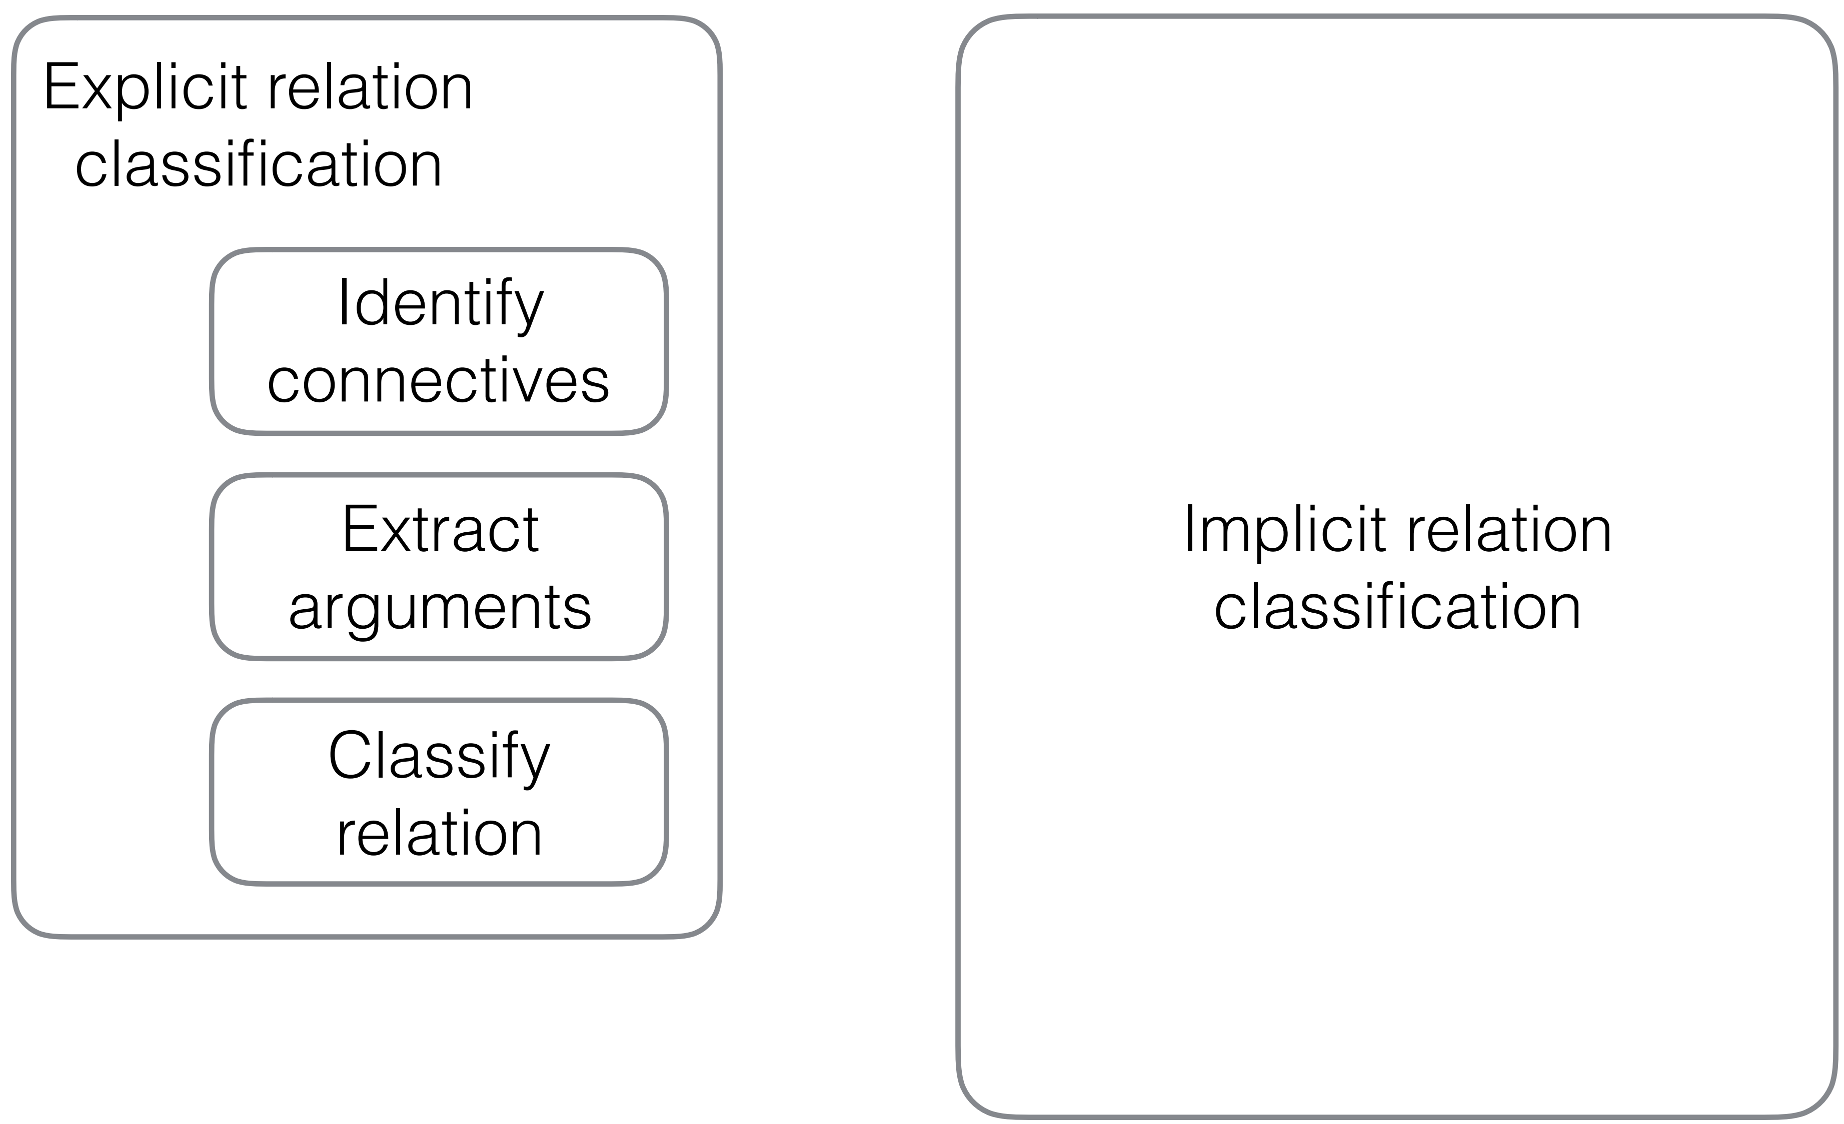

Explicit relation classification

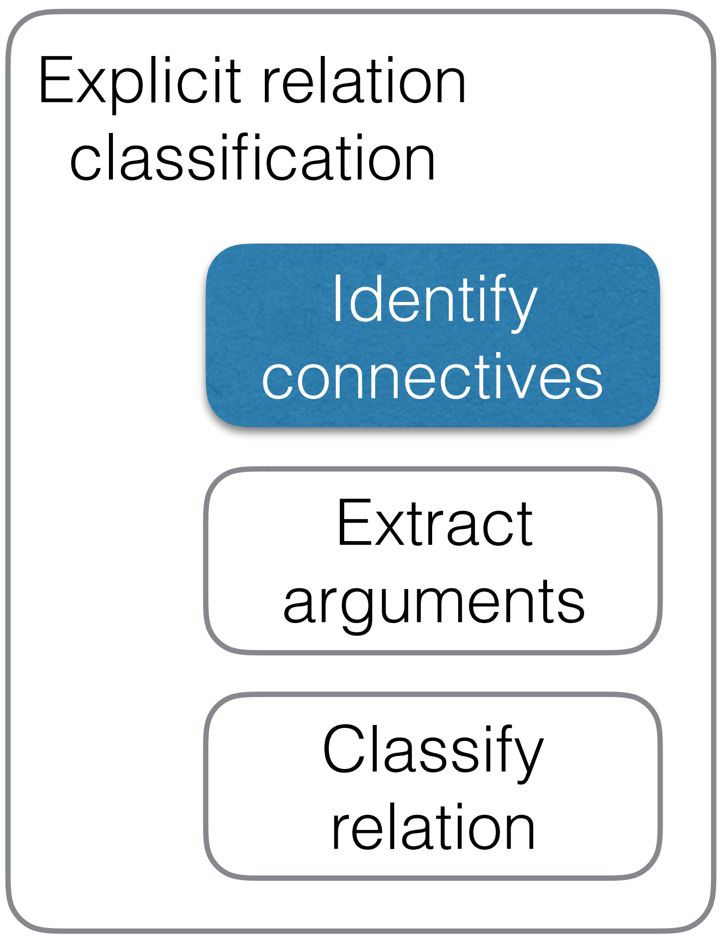

- Cue phrases have non-discourse usage:
    - John and Mary went to the theatre and saw a nice play.
- Connectives in PDTB with discourse usage: 30% of the time
- They vary in frequency of non-discourse usage
    - or: 2.8%
    - although: 91.4%
- Binary classification problem
    - Achieved accuracy of >95% (Pitler and Nenkova 2009, Lin et al. 2014)

Explicit relation classification

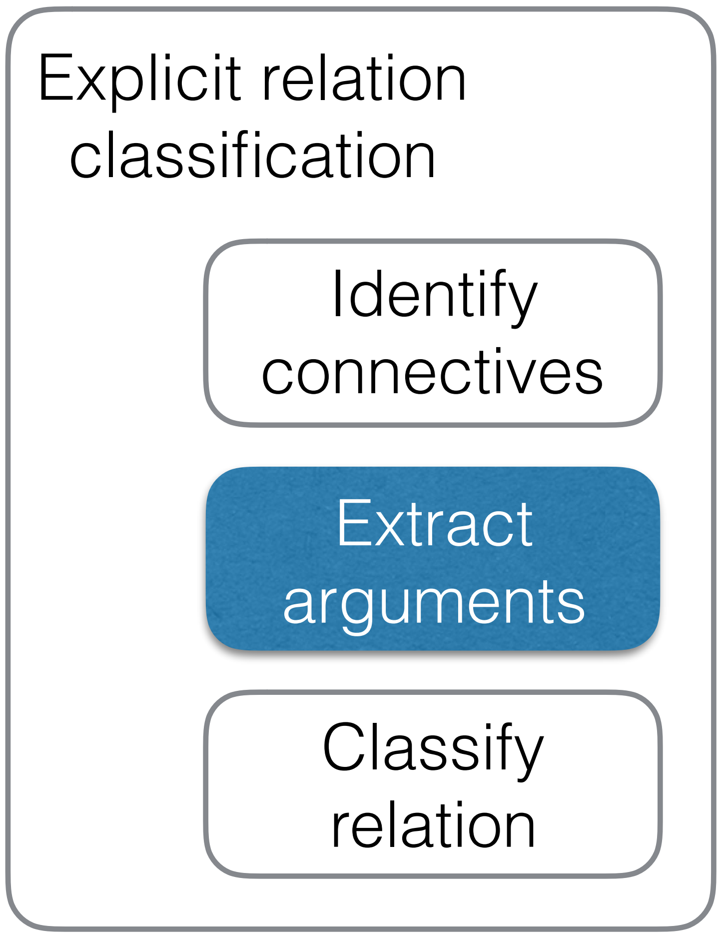

- Argument arrangement:
  - 60.9% in the same sentence
  - 30.1% arg1 in previous adjacent sentence of arg2
  - 9% arg1 in previous non-adjacent sentence of arg2
- Same-sentence arrangement
  - arg1 arg2
  - arg2 arg1
  - [arg1 [arg2]]
  - [arg2 [arg1]]

Explicit relation classification

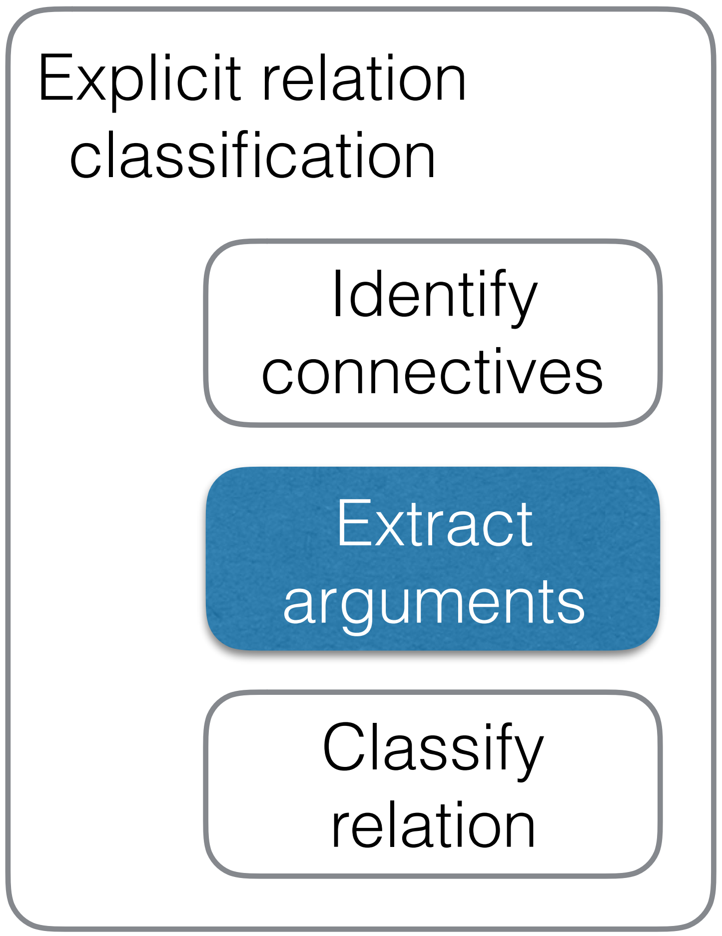

Lin et al. 2014:

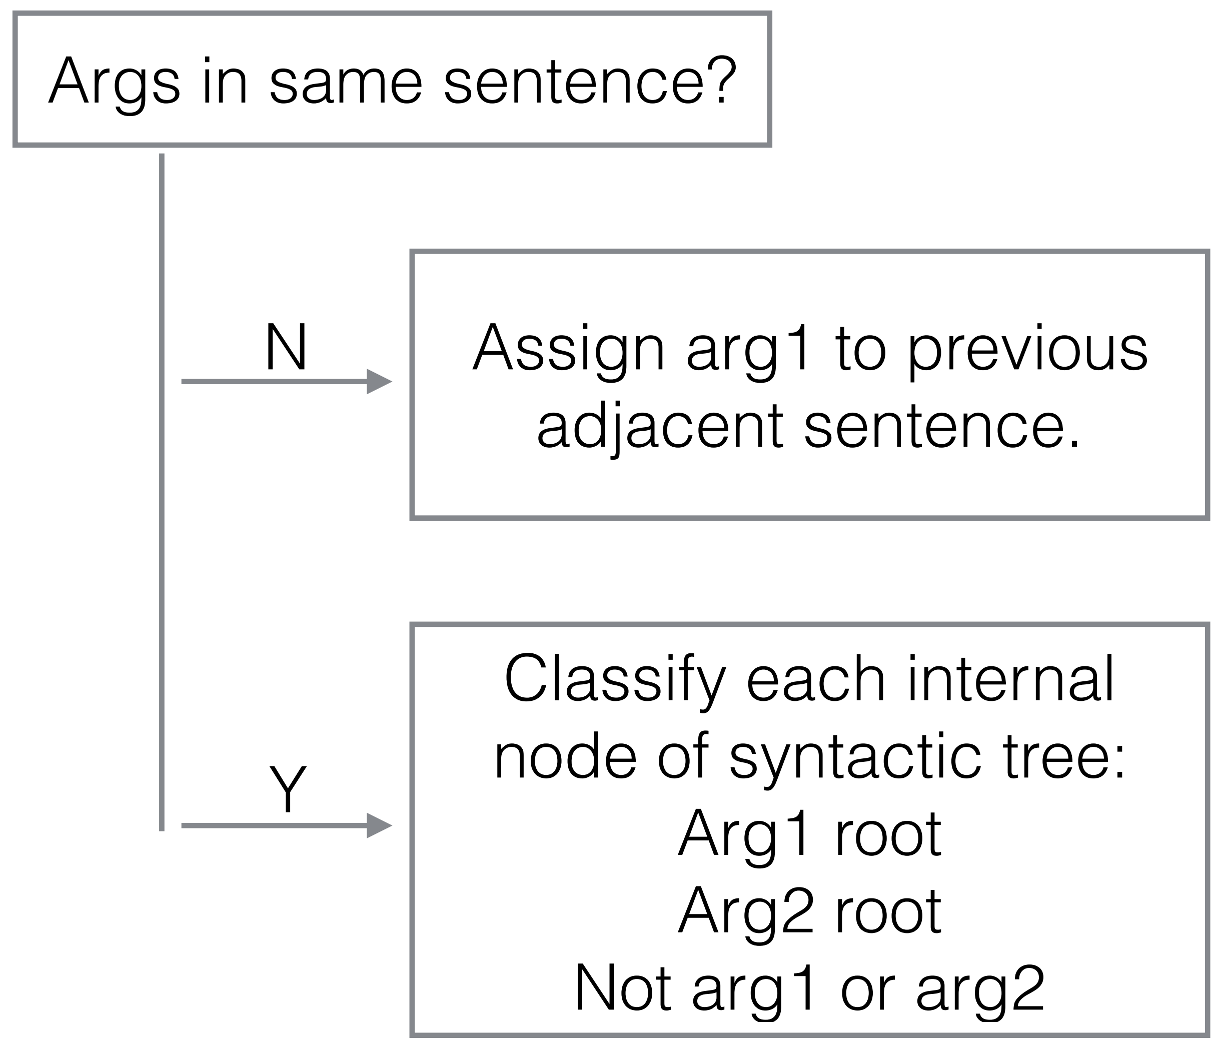

~92% argument position F-score with error propagation

Explicit relation classification

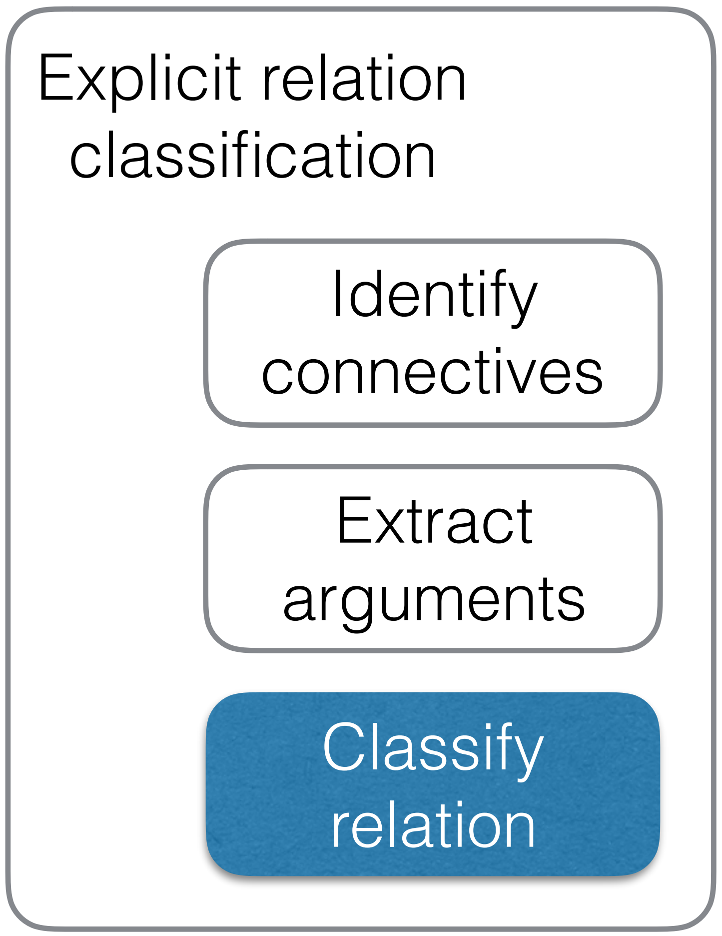

- Most connectives are not ambiguous (Pitler and Nenkova 2008)
- Some frequent connectives are highly ambiguous
  - while, Comparison 66.1%
  - since, Contingency 52.2%
  - as, Temporal 70.3%
  - meanwhile, Temporal 48.7%
- Effective classification (86% F-score) using connective and neighboring token POS (Lin et al. 2014)

Implicit relation classification

- Classification on adjacent sentences within the same paragraph.
  - x: feature vector from two arguments
  - y: PDTB relations (1st level or 2nd level) + EntRel + NoRel
- Neural networks with…
  - Sophisticated attention, multi-task learning, highway networks…
- Very hard task; best system F-scores:
  - Comparison: ~48
  - Contingency: ~59
  - Expansion: ~73
  - Temporal: ~39
  - 11 2nd level relation average F: 45-48

Applications

Automatic Summarization

- Content selection in compressive summarization (Hirao et al., 2013; Kikuchi et al., 2014)
- Connecting product aspects in review summarization (Gerani et al., 2014)

Automatic Summarization

Question Generation and Answering

- Discourse structure helpful for answering non-factoid ("how" and "why") questions (Jensen et al., 2014)
- Generating "why" questions using discourse annotation (Agarwal et al., 2011; Prasad and Joshi, 2008)

Automatic Summarization

Question Generation and Answering

Machine Translation

- Sentential discourse structure for MT evaluation (Guzman et al., 2014)
- Discourse factors highly impact the quality of MT outputs (Li et al., 2014)

Automatic Summarization

Question Generation and Answering

Machine Translation

Sentiment Analysis

- Deep analysis of rhetorical structure helps with content weighting in document polarity classification (Hogenboom et al., 2015)
- Certain relations contain more sentiment expressions than others (Trnavac and Taboada 2013)