

Easy Victories and Uphill Battles in Coreference Resolution



Greg Durrett and Dan Klein

UC Berkeley




Aspects of Coreference



Aspects of Coreference


Binding theory (Chomsky, 1981)





Aspects of Coreference

Binding theory (Chomsky, 1981)




Pronoun agreement (Hobbs, 1977)

MALE FEMALE FEMALE
John₁ talked to Jane₂. She₂ asked...



Aspects of Coreference

Binding theory (Chomsky, 1981)



Pronoun agreement (Hobbs, 1977)

MALE FEMALE FEMALE
John₁ talked to Jane₂. She₂ asked...


Centering (Grosz et al., 1995)

SUBJ. OBJ. SUBJ.
John₁ talked to Bill₂. He₁ asked...



Aspects of Coreference

Binding theory (Chomsky, 1981)



Pronoun agreement (Hobbs, 1977)

MALE FEMALE FEMALE
John₁ talked to Jane₂. She₂ asked...

Centering (Grosz et al., 1995)

SUBJ. OBJ. SUBJ.
John₁ talked to Bill₂. He₁ asked...


Semantic compatibility

The president₁ ... The leader₁



Aspects of Coreference

Binding theory (Chomsky, 1981)



Pronoun agreement (Hobbs, 1977)

MALE FEMALE FEMALE
John₁ talked to Jane₂. She₂ asked...

Centering (Grosz et al., 1995)

SUBJ. OBJ. SUBJ.
John₁ talked to Bill₂. He₁ asked...

Semantic compatibility

The president₁ ... The leader₁

Definiteness

The president

A president



Definiteness: Classical

Soon et al. (2001), Ng et al. (2002), etc.



Definiteness: Classical

The president

Soon et al. (2001), Ng et al. (2002), etc.



Definiteness: Classical

The president

```
if (startsWith("the"))  
    DEFINITE  
else if (startsWith("a"))  
    INDEFINITE  
else  
    NO_ART
```

Soon et al. (2001), Ng et al. (2002), etc.



Definiteness: Classical

The president

```
if (startsWith("the"))  
  DEFINITE  
else if (startsWith("a"))  
  INDEFINITE  
else  
  NO_ART
```

DEFINITE

INDEFINITE NO_ART

Soon et al. (2001), Ng et al. (2002), etc.



Definiteness: Data-Driven

The president



Definiteness: Data-Driven

The president



The



Definiteness: Data-Driven

The president



The these
that U.S. all his
this a no
an some more John those
their its
Barack Israeli new

...



Approaches

Classical approach



Approaches

Classical approach

- ▶ Learning with heuristic features { Soon et al. (2001), inter alia



Approaches

Classical approach

- ▶ Learning with heuristic features { Soon et al. (2001), inter alia
- ▶ Rule-based { Haghghi and Klein (2009),
Lee et al. (2011)



Approaches

Classical approach

- ▶ Learning with heuristic features { Soon et al. (2001), inter alia
- ▶ Rule-based { Haghghi and Klein (2009),
Lee et al. (2011)

Mixed approach

- ▶ Add data-driven features on a few axes { Bengtson and Roth (2008),
Rahman and Ng (2011),
Björkelund and Nugues (2011)



Approaches

Classical approach

- ▶ Learning with heuristic features { Soon et al. (2001), inter alia
- ▶ Rule-based { Haghghi and Klein (2009),
Lee et al. (2011)

Mixed approach

- ▶ Add data-driven features on a few axes { Bengtson and Roth (2008),
Rahman and Ng (2011),
Björkelund and Nugues (2011)

Data-driven approach (this work)

- ▶ >400,000 comprehensive, uniform features



Mention-Ranking Architecture

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



Mention-Ranking Architecture

[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...

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



Mention-Ranking Architecture

A_1

New



[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...

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



Mention-Ranking Architecture




[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...

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



Mention-Ranking Architecture




[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...

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




Mention-Ranking Architecture




[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...

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



Mention-Ranking Architecture

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



[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...



Mention-Ranking Architecture

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

[Voters-they]
[NOM-PRONOUN]
...

A_1
New
○

A_2
New
○
1 ←

A_3
New
○
1
2 ←

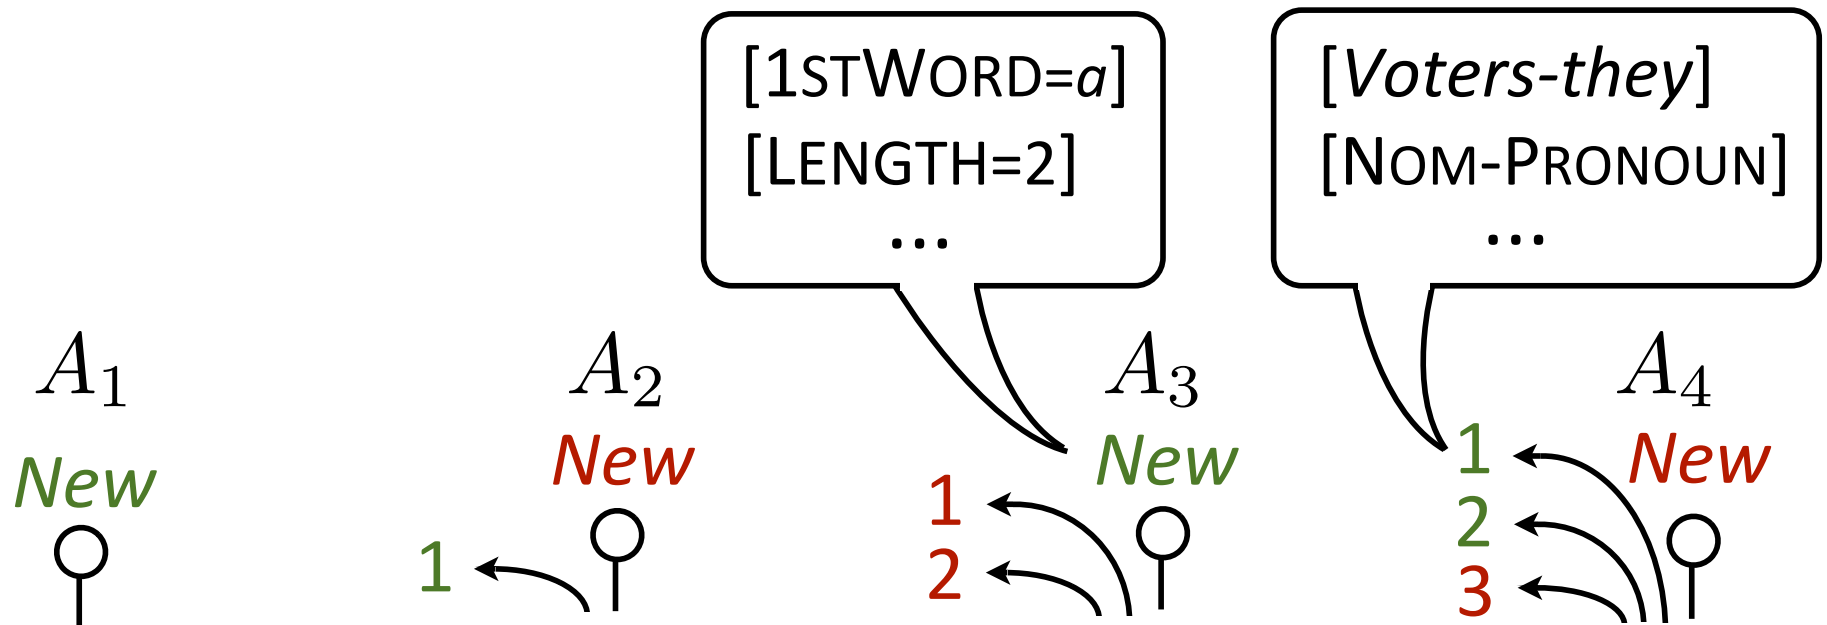
A_4
New
○
1
2
3 ←

[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...



Mention-Ranking Architecture

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



[Voters]₁ agree when [they]₁ are given [a chance]₂ to decide if [they]₁ ...



BASIC Features

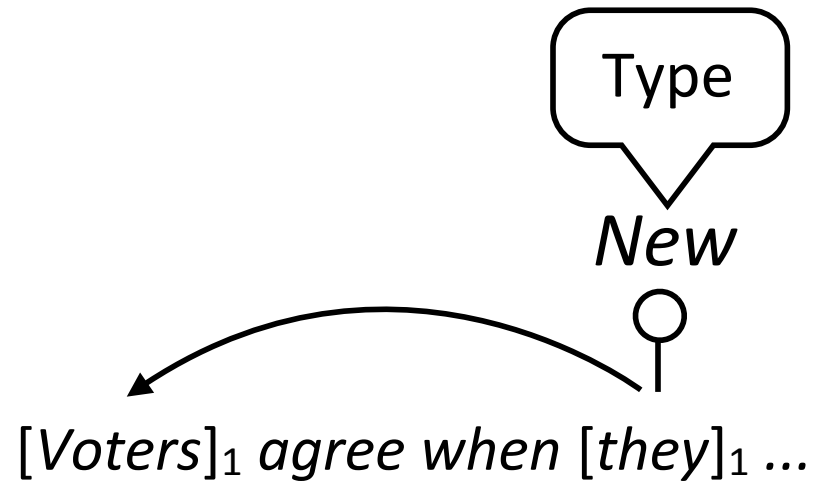


BASIC Features

New
○
|
↙
[Voters]₁ agree when [they]₁ ...



BASIC Features





BASIC Features

Types

Sentence distance (+ types)

Mention distance (+ types)

Head match

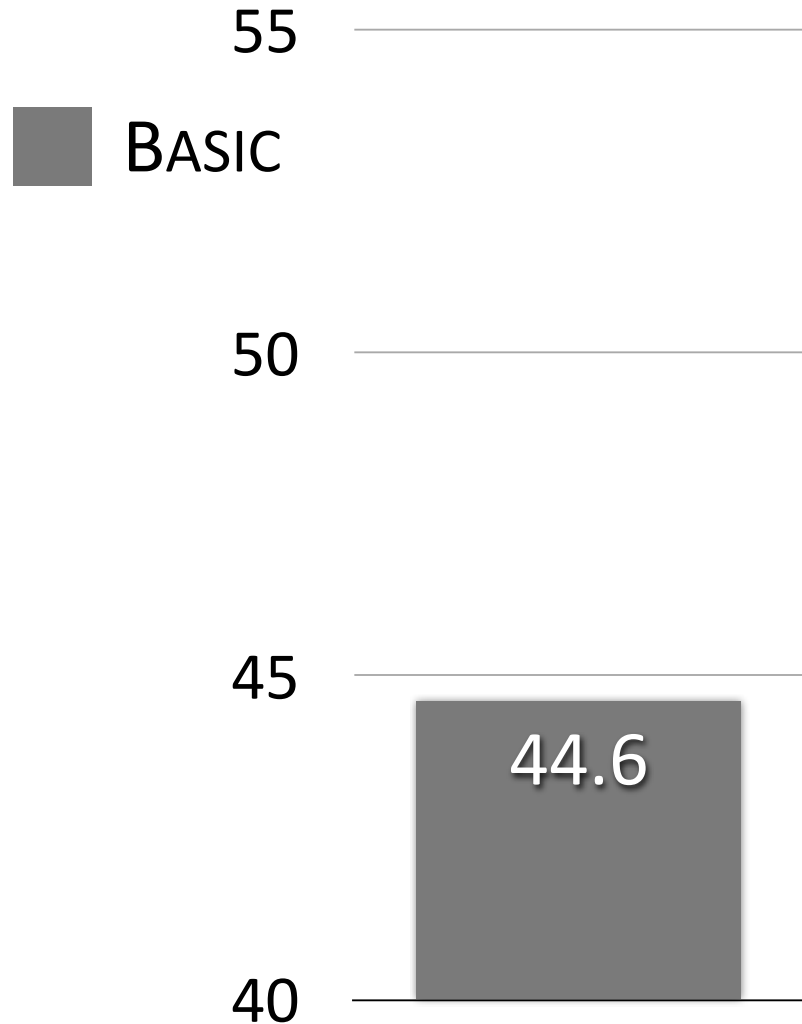
Type

New

[*Voters*]₁ agree when [*they*]₁ ...



BASIC Features



(Dev set CoNLL-F1, predicted mentions)



Definiteness

Soon et al. (2001), Ng et al. (2002), etc.



Definiteness

New




The president

Soon et al. (2001), Ng et al. (2002), etc.



Definiteness



Soon et al. (2001), Ng et al. (2002), etc.



Definiteness

[1STWORD=*The*]


New



The president



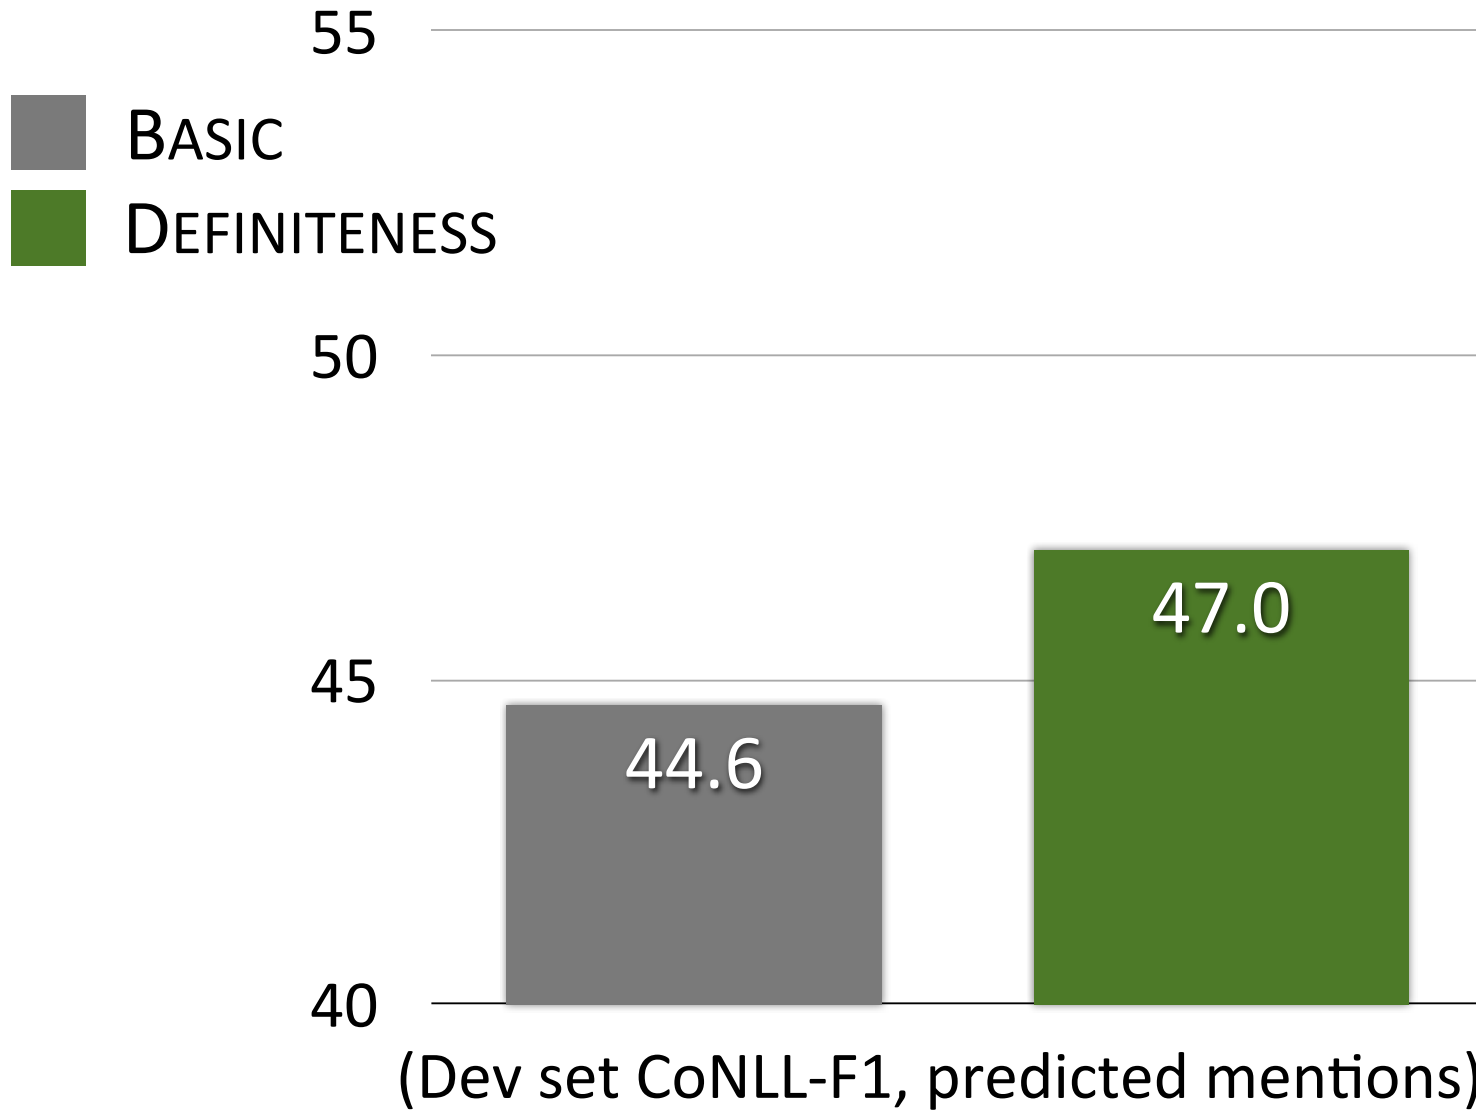
Definiteness



(Dev set CoNLL-F1, predicted mentions)




Definiteness





Definiteness





Definiteness

DEFINITENESS

47.0



Definiteness

DEFINITENESS


47.0

FIRST WORD: "the, a, an"

47.0




Definiteness



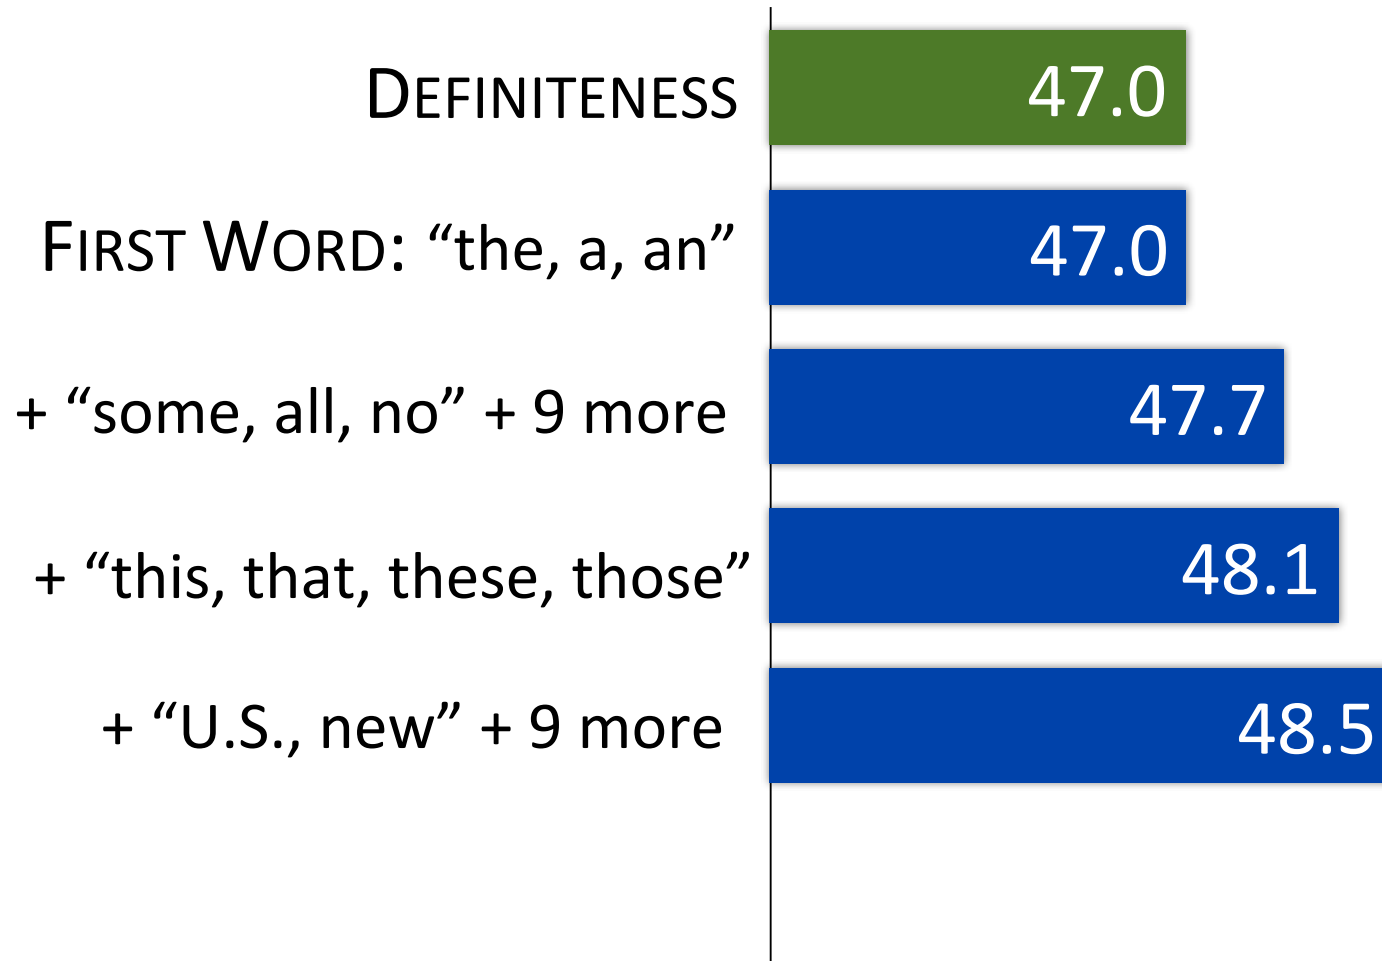


Definiteness






Definiteness





Definiteness





Centering

Grosz al. (1995)



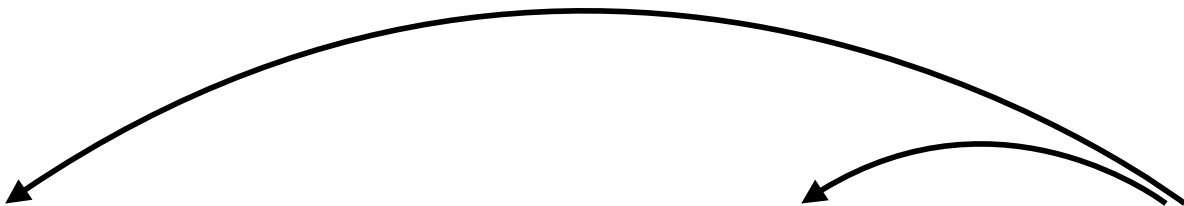
Centering

[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...



Centering

[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...





Centering

SUBJECT–SUBJECT


OBJECT–SUBJECT

[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...



Centering

SUBJECT–SUBJECT



[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...



Centering

[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...



Centering

[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...

SUBJECT

OBJECT

SUBJECT



Centering

[OBJECT-SUBJECT]

[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...

SUBJECT

OBJECT

SUBJECT



Centering

[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...





Centering

*[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...
with [X] .*



Centering

*[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...
with [X] . . [X] said*



Centering

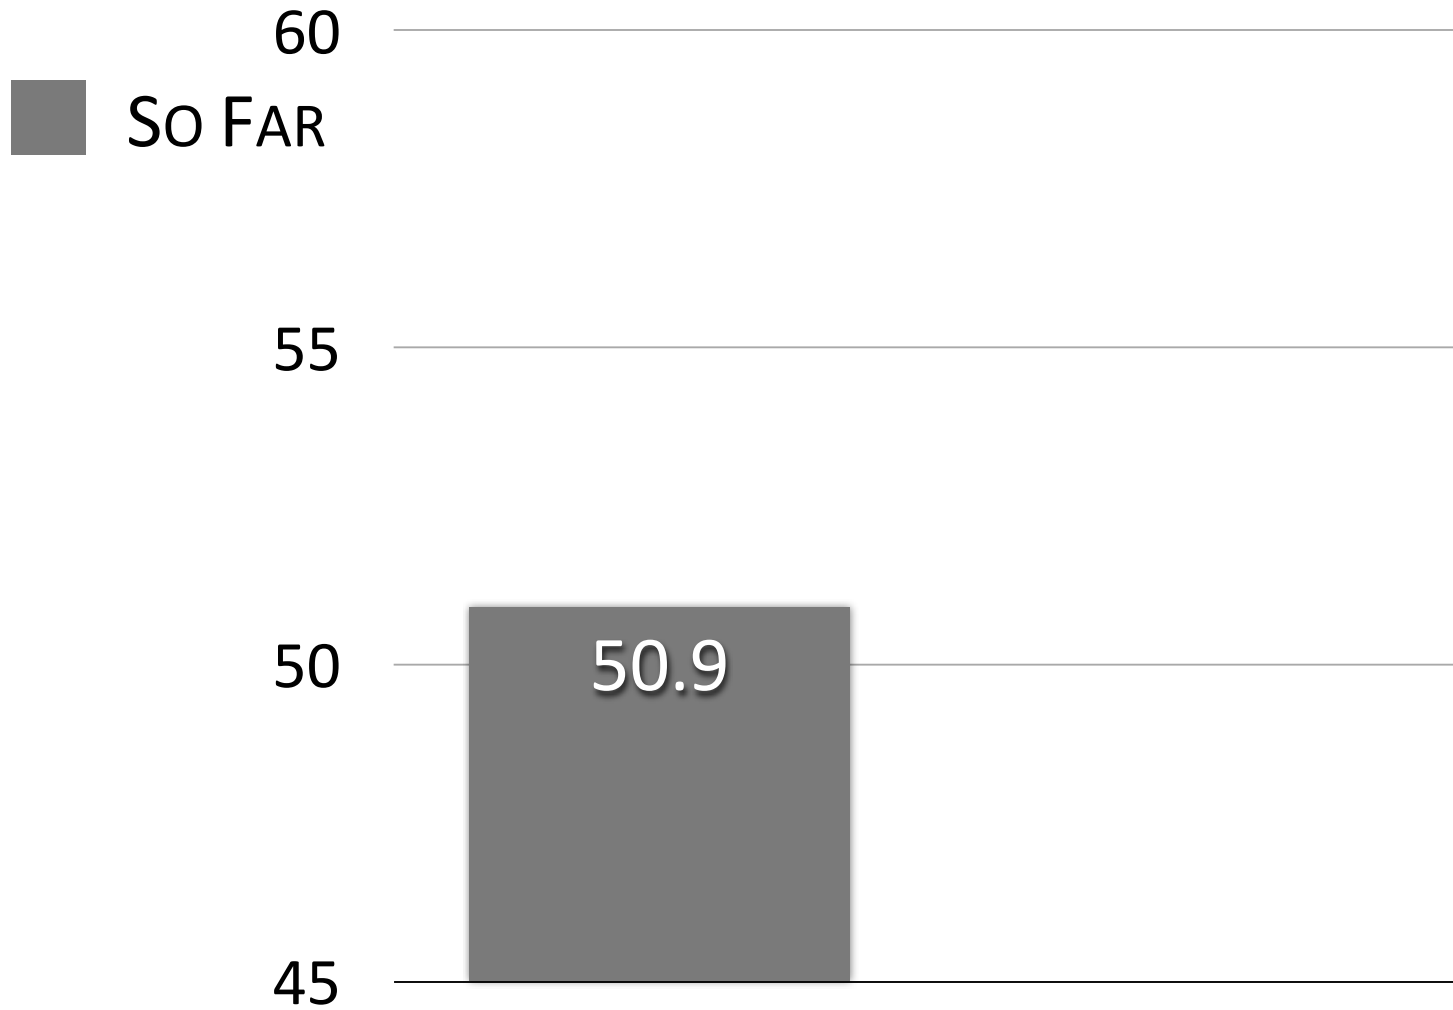
[with X - . Y]
[with X - Y said]
...



[Barack Obama]₁ met with [David Cameron]₂ . [He]₁ said ...
with [X] . . [X] said




Centering



(Dev set CoNLL-F1, predicted mentions)



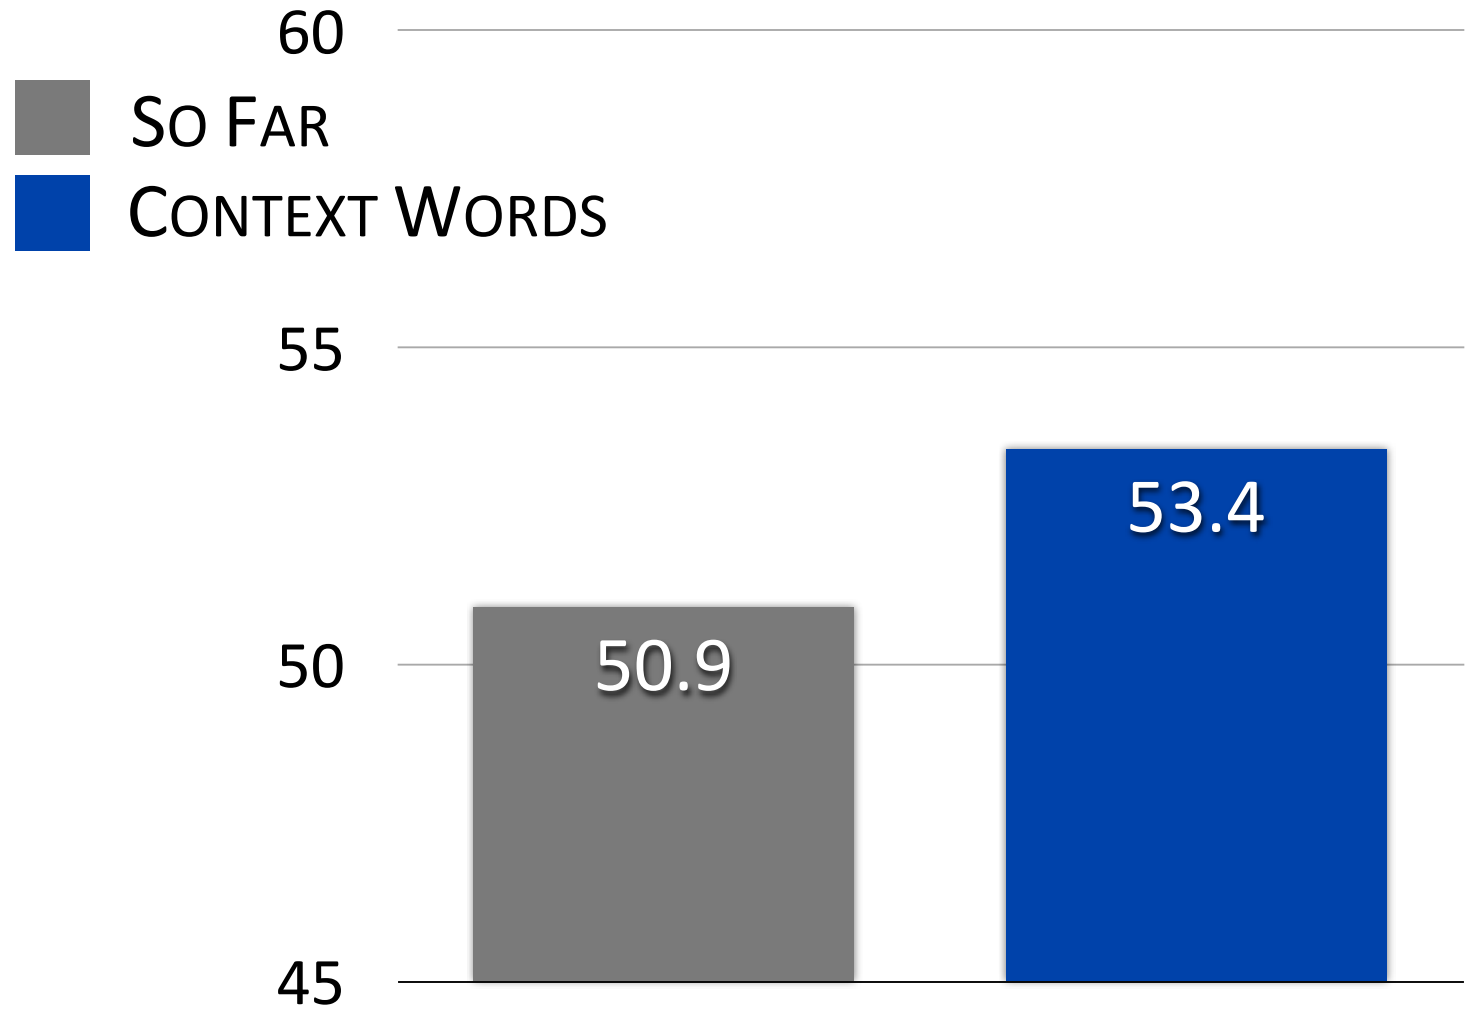
Centering




(Dev set CoNLL-F1, predicted mentions)



Centering



(Dev set CoNLL-F1, predicted mentions)



Pronoun Agreement

Bergsma and Lin (2006)

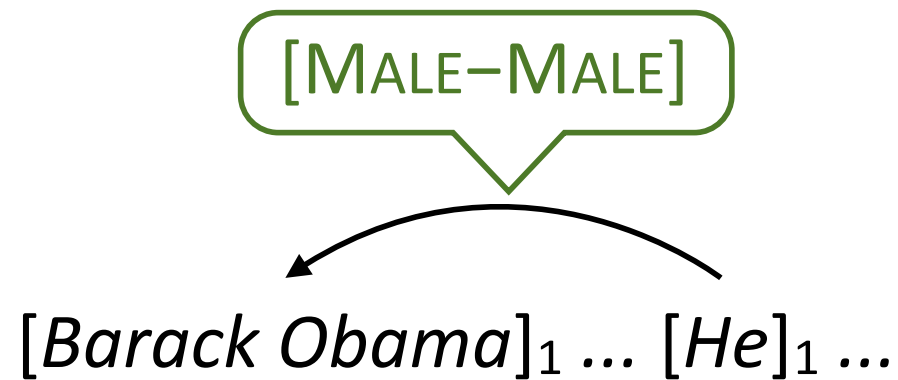


Pronoun Agreement

[Barack Obama]₁ ... [He]₁ ...

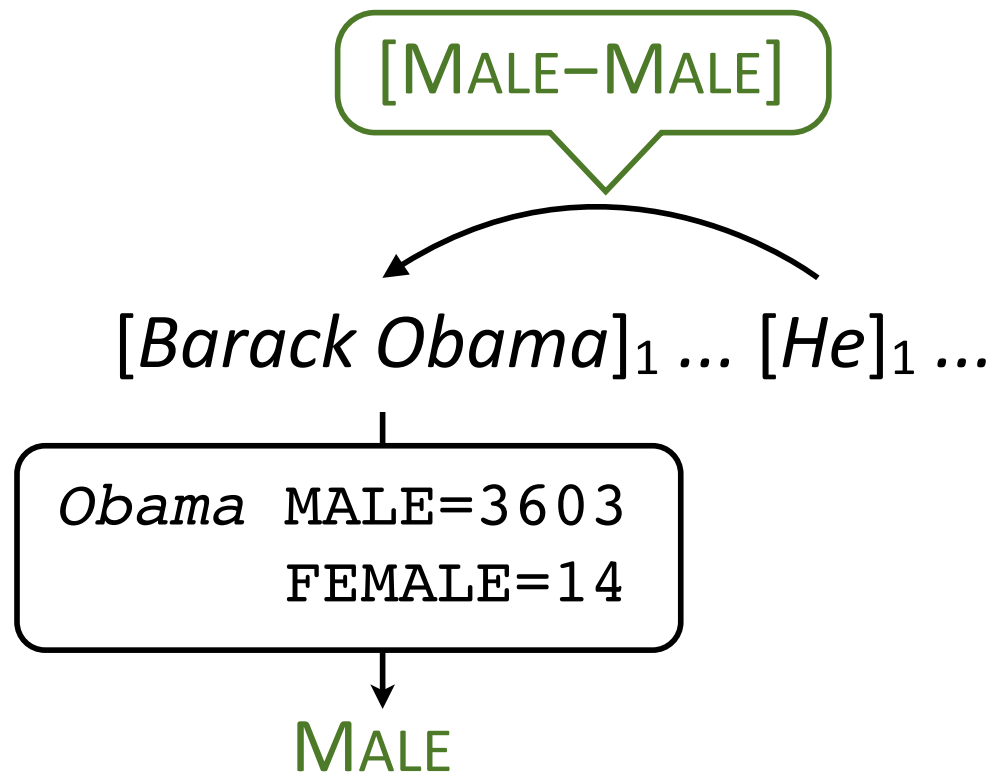


Pronoun Agreement



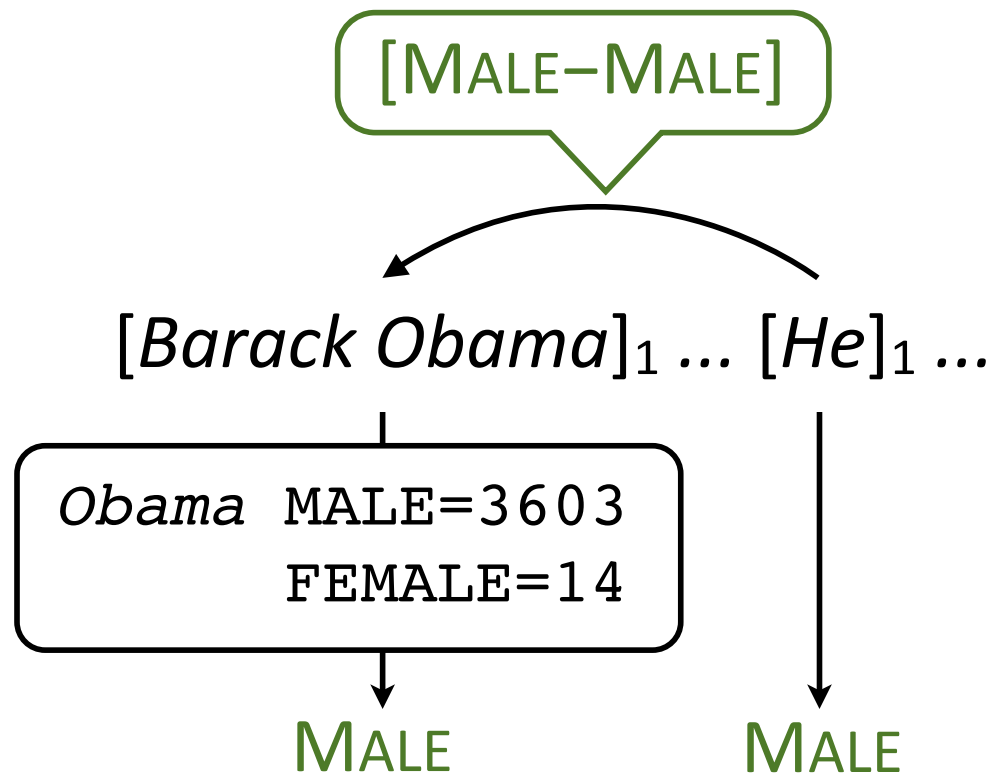


Pronoun Agreement





Pronoun Agreement





Pronoun Agreement

[Barack Obama]₁ ... [He]₁ ...

A curved arrow pointing from the right to the left, indicating the relationship between the two phrases. The arrow starts above the second phrase and points towards the first phrase.



Pronoun Agreement


[*Obama-He*]



[*Barack Obama*]₁ ... [*He*]₁ ...




Pronoun Agreement



(Dev set CoNLL-F1, predicted mentions)

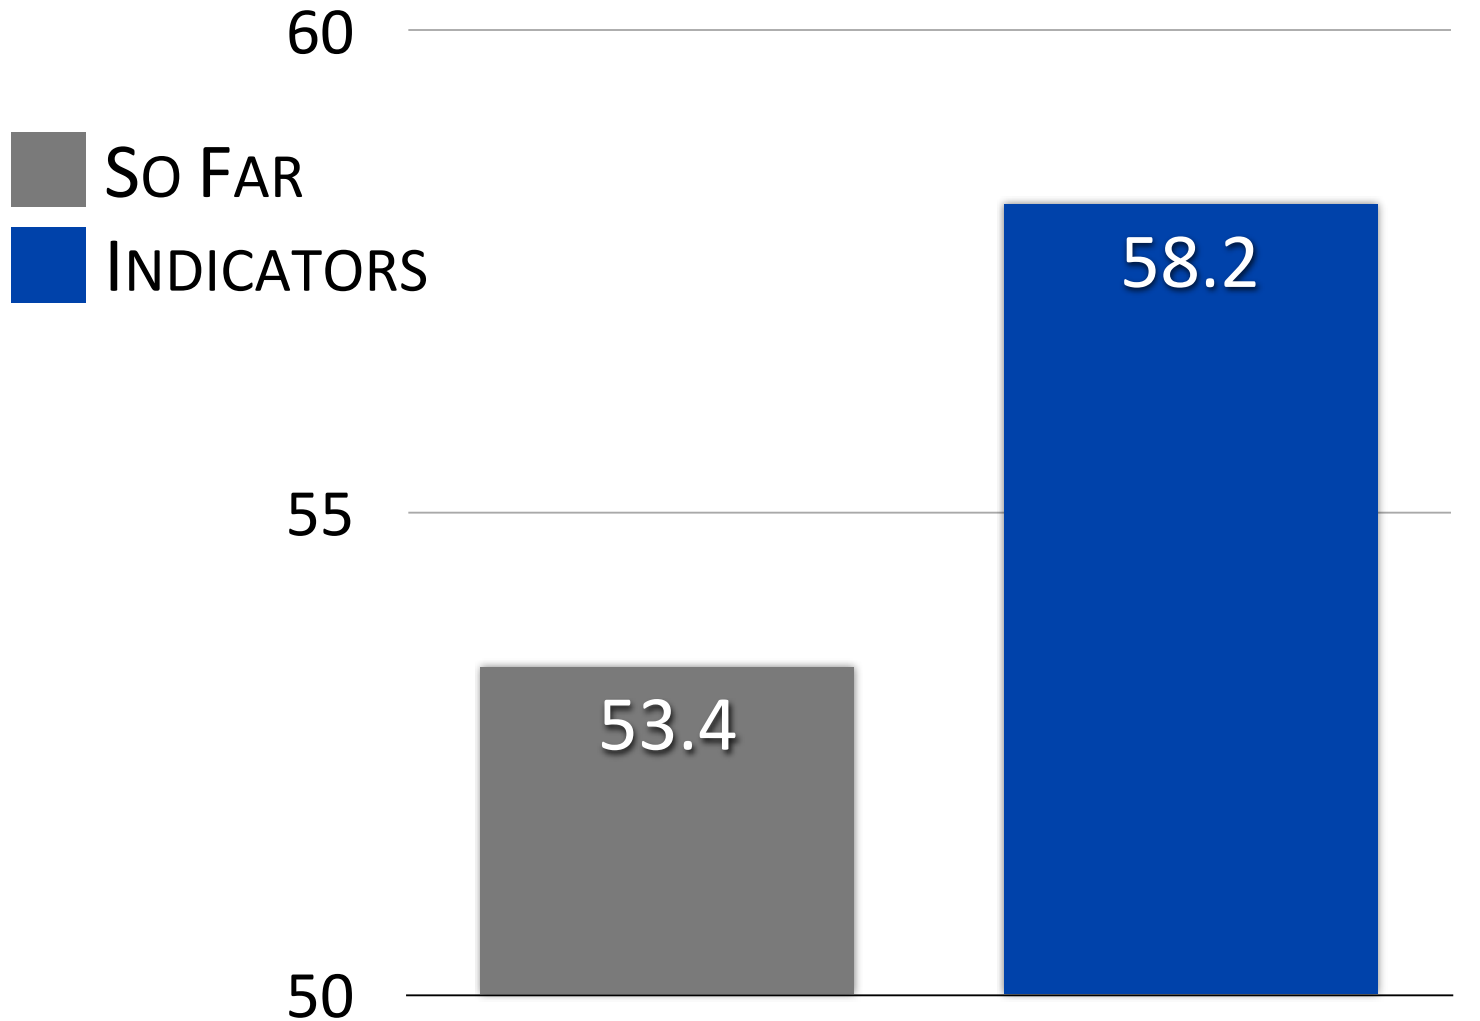


Pronoun Agreement





Pronoun Agreement



(Dev set CoNLL-F1, predicted mentions)



SURFACE Information

- ▶ Features conjoin surface-level mention attributes



SURFACE Information

- ▶ Features conjoin surface-level mention attributes

<s> [*President Barack Obama*] signed the bill ... Afterwards [*he*] said ...



SURFACE Information

- ▶ Features conjoin surface-level mention attributes

<s> [*President Barack Obama*] signed the bill ... Afterwards [*he*] said ...



SURFACE Information


- ▶ Features conjoin surface-level mention attributes

<s> [*President Barack Obama*] signed the bill ... Afterwards [*he*] said ...



SURFACE Information

- ▶ Features conjoin surface-level mention attributes



<s> [*President Barack Obama*] *signed the bill ... Afterwards* [*he*] *said ...*

LENGTH = 3
PROPER

LENGTH = 1
PRONOUN



SURFACE Information

- ▶ Features conjoin surface-level mention attributes

Mention distance
Sentence distance

<s> [*President Barack Obama*] *signed the bill ... Afterwards* [*he*] *said ...*

LENGTH = 3
PROPER

LENGTH = 1
PRONOUN



SURFACE Information

- ▶ Features conjoin surface-level mention attributes

Mention distance
Sentence distance
Head match
Exact match

<s> [*President Barack Obama*] signed the bill ... Afterwards [*he*] said ...


LENGTH = 3
PROPER

LENGTH = 1
PRONOUN



SURFACE Accuracy



■ SO FAR



(Dev set CoNLL-F1, predicted mentions)



SURFACE Accuracy



(Dev set CoNLL-F1, predicted mentions)



Error Analysis



Error Analysis

Anaphoric pronouns

Obama ← he



Error Analysis

Anaphoric pronouns

Obama ← he

72.0%



Error Analysis

Anaphoric pronouns

Obama ← he

72.0%

Referring: head match

the U.S. president ← *president*



Error Analysis

Anaphoric pronouns

Obama ← he

72.0%

Referring: head match

the U.S. president ← president

82.7%



Error Analysis

Anaphoric pronouns

Obama ← he

72.0%

Referring: head match

the U.S. president ← *president*


82.7%

Referring: no head match

David Cameron ← *prime minister*



Error Analysis





Compatibility



Compatibility

[David Cameron]₁ ... [The prime minister]₁ ...



Compatibility

[David Cameron]₁ ... [The prime minister]₁ ...

}
} Soon et al. (2001)
inter alia



Compatibility

[David Cameron]₁ ... [The prime minister]₁ ...

▶ Number, gender



Soon et al. (2001)
inter alia



Compatibility

[David Cameron]₁ ... [The prime minister]₁ ...

- ▶ Number, gender
- ▶ Named entity type



Soon et al. (2001)
inter alia



Compatibility

[David Cameron]₁ ... [The prime minister]₁ ...

- ▶ Number, gender
- ▶ Named entity type
- ▶ Unsupervised clustering labels



Soon et al. (2001)
inter alia




Compatibility

[David Cameron]₁ ... [The prime minister]₁ ...

- ▶ Number, gender
 - ▶ Named entity type
 - ▶ Unsupervised clustering labels
 - ▶ WordNet hypernymy / synonymy
- } Soon et al. (2001)
inter alia

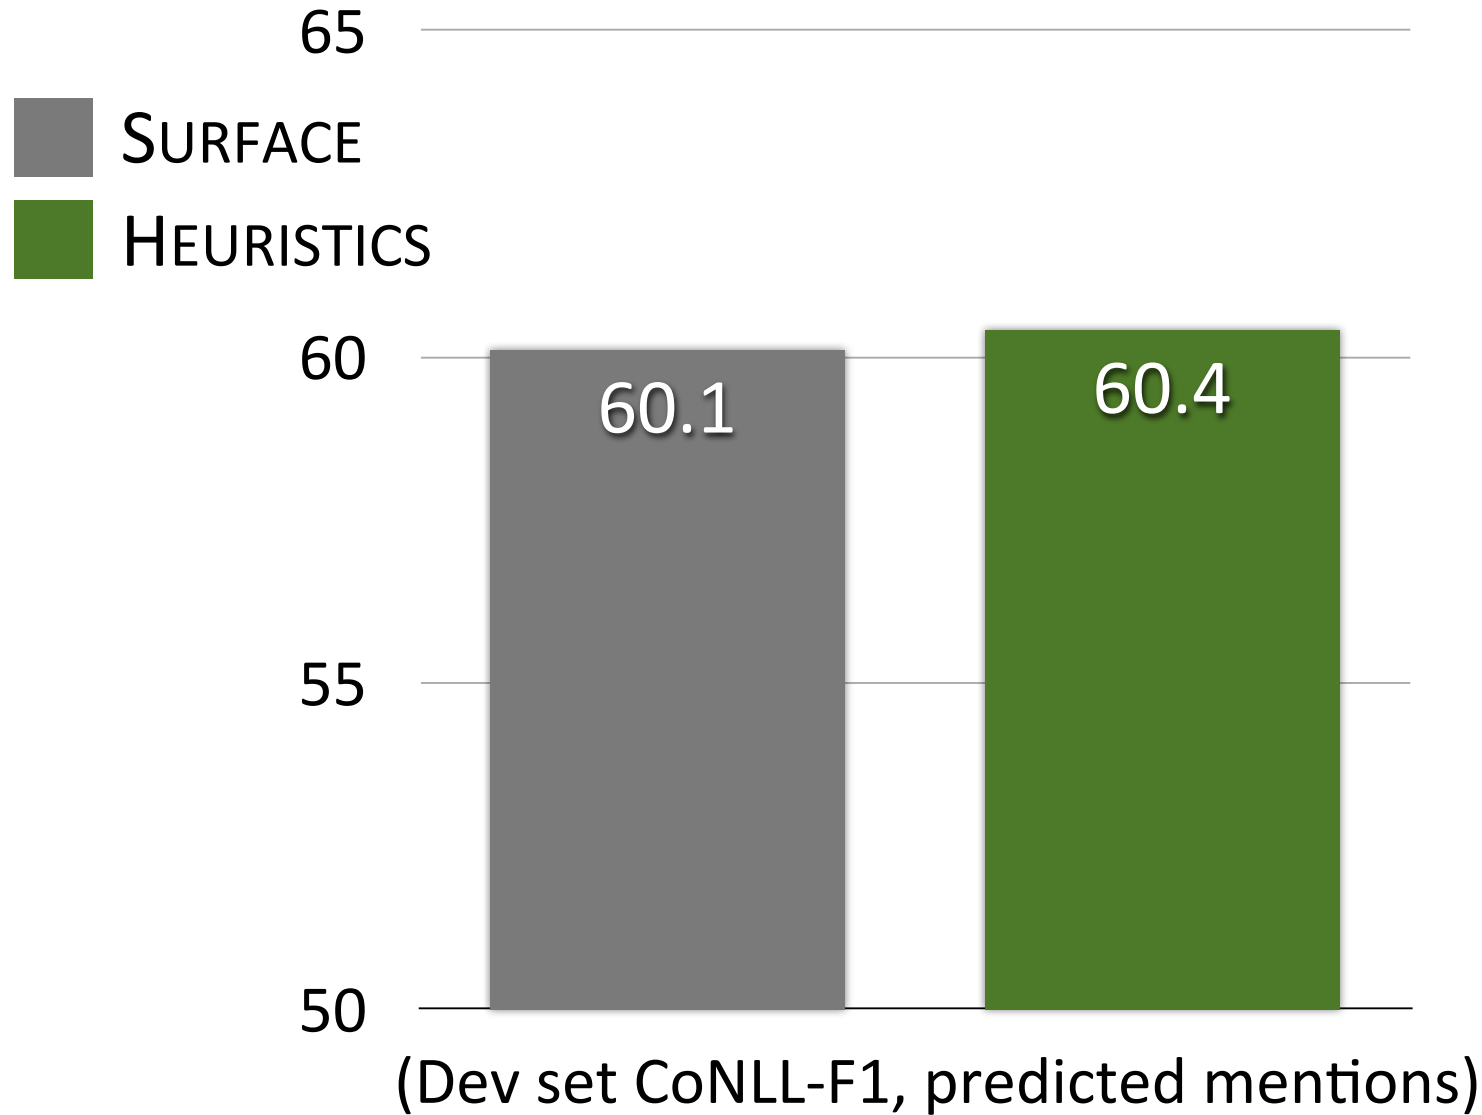


SURFACE Accuracy





SURFACE Accuracy





What works?



What works?

- ▶ Importing external information with sophisticated heuristics {
 - Ponzetto and Strube (2006)
 - Rahman and Ng (2011)
 - Bansal and Klein (2012)




What works?

- ▶ Importing external information with sophisticated heuristics $\left\{ \begin{array}{l} \text{Ponzetto and Strube (2006)} \\ \text{Rahman and Ng (2011)} \\ \text{Bansal and Klein (2012)} \end{array} \right.$
- ▶ We can support additional heuristic features, including number and gender information (Bergsma and Lin, 2006)



Test Set Accuracy

 STANFORD (Lee et al., 2011)




(Test set CoNLL-F1, predicted mentions)



Test Set Accuracy

- STANFORD (Lee et al., 2011)
- IMS (Björkelund and Farkas, 2012)




(Test set CoNLL-F1, predicted mentions)



Test Set Accuracy


- STANFORD (Lee et al., 2011)
- IMS (Björkelund and Farkas, 2012)
- SURFACE (This work)



(Test set CoNLL-F1, predicted mentions)



Test Set Accuracy



(Test set CoNLL-F1, predicted mentions)



Conclusion



Conclusion

- ▶ Surface lexical features capture a wide range of linguistic phenomena in a unified way



Conclusion

- ▶ Surface lexical features capture a wide range of linguistic phenomena in a unified way
- ▶ Semantic errors require heavy-duty information from other knowledge sources



Conclusion

- ▶ Surface lexical features capture a wide range of linguistic phenomena in a unified way
- ▶ Semantic errors require heavy-duty information from other knowledge sources
- ▶ Extensible system that achieves state-of-the-art performance



Conclusion

- ▶ The Berkeley Coreference Resolution System:
<http://nlp.cs.berkeley.edu>
- ▶ Full end-to-end system (accepts raw text as input)




Conclusion

- ▶ The Berkeley Coreference Resolution System:
<http://nlp.cs.berkeley.edu>
- ▶ Full end-to-end system (accepts raw text as input)

Thank you!





Stacking






SURFACE Accuracy



Predicted mentions




SURFACE Accuracy






SURFACE Accuracy





SURFACE Accuracy





Gold vs. System Mentions

New



[the president]



Gold vs. System Mentions

New




[the president]

PERSON




Gold vs. System Mentions






Gold vs. System Mentions





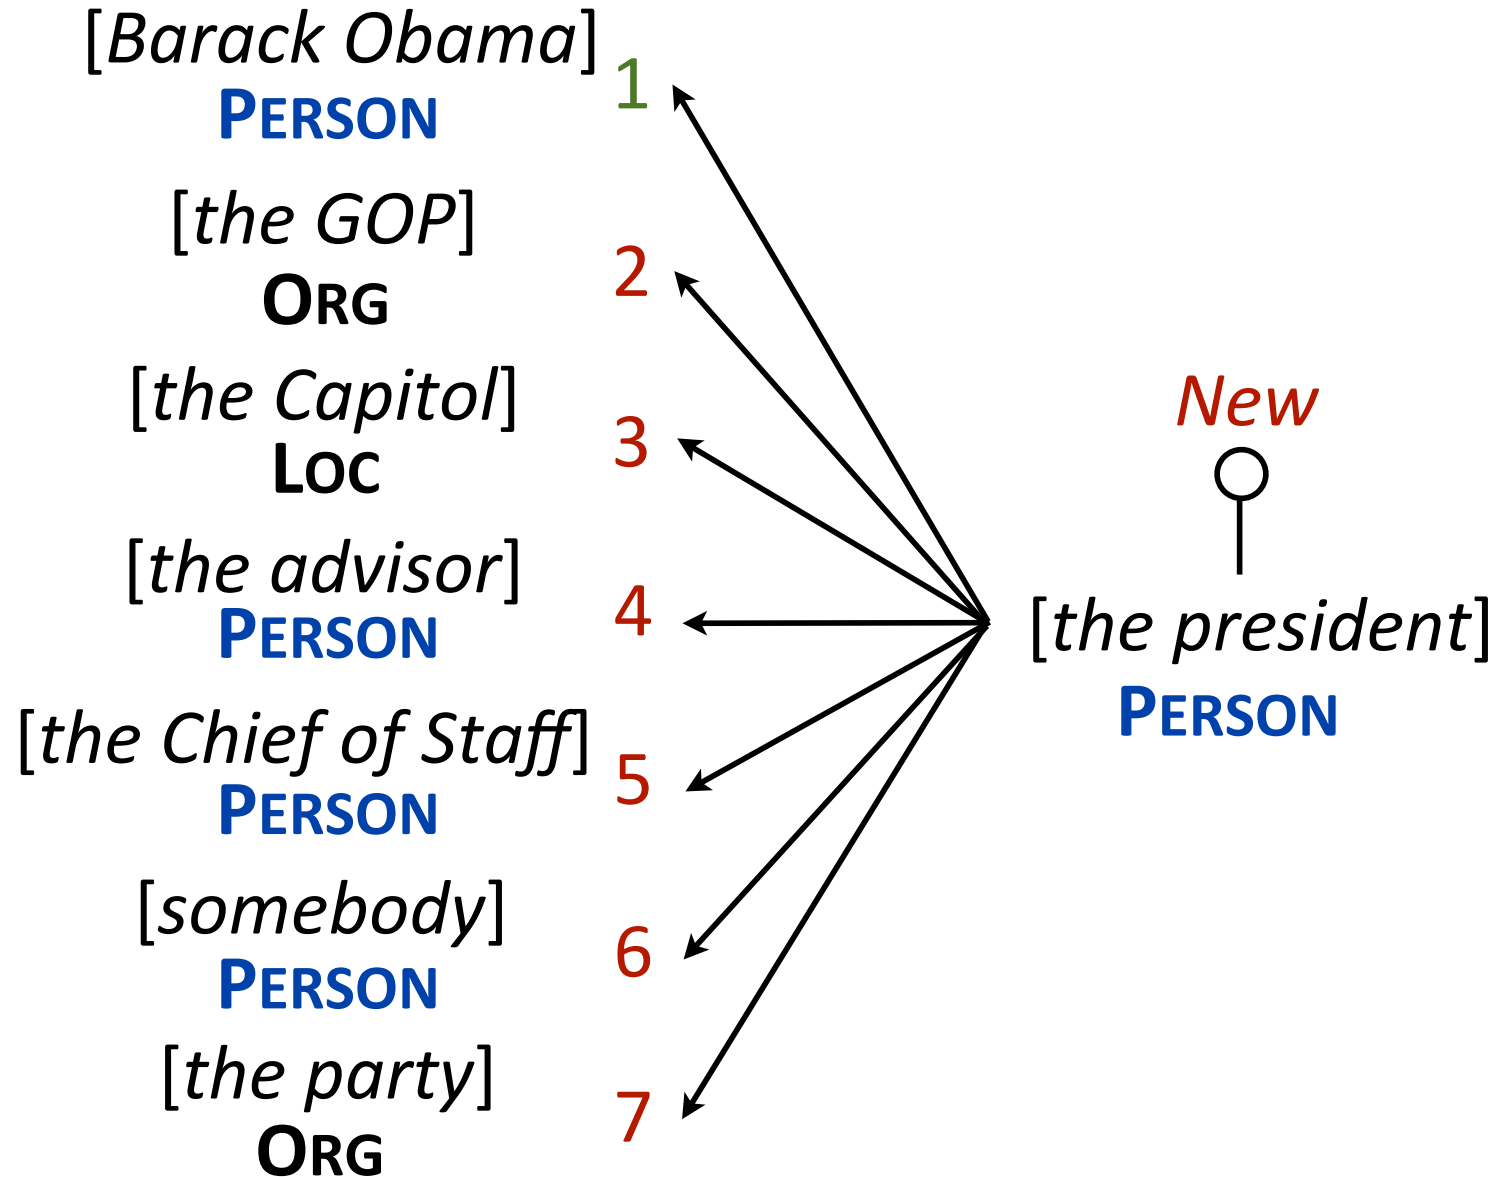
Gold vs. System Mentions



- ▶ 3x more mentions when singletons are included



Gold vs. System Mentions





Gold vs. System Mentions

