Coreference Resolution

Bonan Min

bonanmin@gmail.com

Some slides are based on class materials from Ralph Grishman, Thien Huu Nguyen, Luke Zettlemoyer

Reference Resolution: Objective

Identify all phrases which refer to the same real-word entity

- first, within a single document
- later, also across multiple documents

Terminology

referent: real-world object referred to

referring expression [mention]: a phrase referring to that object

Mary was hungry; she ate a banana.

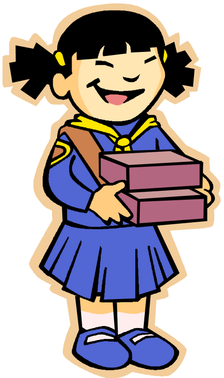

Terminology

coreference: two expressions referring to the same thing

Mary was hungry; she ate a banana.

antecedent
(prior expression)

anaphor
(following expression)

So we also refer to process as anaphora resolution

Coreference Resolution

Barack Hussein Obama II (/bəˈrɑːk huːˈseɪn oʊˈbɑːmə/; born August 4, 1961) is the 44th and current President of the United States, and the first African American to hold the office. Born in Honolulu, Hawaii, Obama is a graduate of Columbia University and Harvard Law School, where $\boxed{\text{he}}$ served as president of the *Harvard Law Review*. $\boxed{\text{He}}$ was a community organizer in Chicago before earning his law degree. $\boxed{\text{He}}$ worked as a civil rights attorney and taught constitutional law at the University of Chicago Law School from 1992 to 2004. $\boxed{\text{He}}$ served three terms representing the 13th District in the Illinois Senate from 1997 to 2004, running unsuccessfully for the United States House of Representatives in 2000.

Types of Referring Expressions

definite pronouns (he, she, it, …)

indefinite pronouns (one)

definite NPs (the car)

indefinite NPs (a car)

names

Referring Expressions: Pronouns

Definite pronouns: he, she, it, …

generally anaphoric

- Mary was hungry; she ate a banana

pleonastic (non-referring) pronouns

- It is raining.
- It is unlikely that he will come.

pronouns can represent bound variables in quantified contexts:

- Every lion finished its meal.

Referring Expressions: Pronouns

Indefinite pronouns (one)

refers to another entity with the same properties as the antecedent

- Mary bought an IPhone6.
- Fred bought one too.
- *Fred bought it too.

can be modified

- Mary bought a new red convertible.
- Fred bought a used one.
  = a used red convertible

  (retain modifiers on antecedent which are compatible with those on anaphor)

Referring Expressions: Pronouns

Reflexive pronouns (himself, herself, itself)

used if antecedent is in same clause

- I saw myself in the mirror.

Referring Expressions: NPs

NPs with definite determiners ("the")

reference to uniquely identifiable entity

generally anaphoric
- I bought a Ford Fiesta. <u>The car</u> is terrific.

but may refer to a uniquely identifiable common noun
- I looked at <u>the moon</u>
- <u>The president</u> announced …

or a functional result
- <u>The sum of 4 and 5</u> is 9.
- <u>The price of gold</u> rose by $4.

Referring Expressions: NPs

NPs with indefinite determiners (“a”)

generally introduces a new ‘discourse entity’

may also be generic:

- A giraffe has a long neck.

Referring Expressions: Names

Subsequent references can use portions of name:

- Fred Frumble and his wife Mary bought a house. Fred put up a hammock.

Complications

Cataphora: pronoun referring to a following mention:

- When she entered the room, Mary looked around.

Bridging anaphora: reference to related object

- Entering the room, Mary looked at the ceiling.

Zero anaphora: many languages allow subject omission, and some allow omission of other arguments (e.g., Japanese)

- these can be treated as zero (implicit) anaphors
  - similar resolution procedures
- some cases of bridging anaphora can be described in terms of PPs with zero anaphors:
  - IBM announced the appointment as Fred as president

Non-NP anaphora: Pronouns can also refer to events or propositions:

- Fred claimed that no one programs in Lisp. That is ridiculous.

Conjunctions and collective reference

Conjunctions and Collective Reference

With a conjoined NP,

… Fred and Mary …

we can refer to an individual ("he", "she") or the conjoined set ("they")

We can even refer to the collective set if not conjoined …

"Fred met Mary after work. <u>They</u> went to the movies."

Resolving Pronoun Reference

Constraints

Preferences

Hobbs Search

Selectional preferences

Combining factors

Pronouns: Constraints

Pronoun must agree with antecedent in:

animacy
- We're watching a movie. He likes it [*he = you and I]

gender
- Mary met Mr. and Mrs. Jones. She was wearing orange pants.
- needs first-name dictionary
- some nouns gender-specific: sister, ballerina

number
- some syntactically singular nouns can be referred to by a plural pronoun: "The platoon … they"

Pronouns: Preferences

Prefer antecedents that are

recent
- at most 3 sentences back

salient
- mentioned several times recently

subjects

Recency and preference for subjects are often captured by Hobbs search order, a particular order for searching the current and preceding parse trees

Hobbs Search Order

Traverse parse tree containing anaphor, starting from anaphor

Then

- traverse trees for preceding sentences, breadth first, left-to-right
  - incorporates subject precedence
- stop at first NP satisfying constraints
  - Gender and number agreement
    - Male-female
    - Singular-plural

Relatively simple strategy, competitive performance: correct 72.7%

- As a competitive baseline
- Provide features for more sophisticated algorithms

Hobbs Search Order: Example

The castle in Camelot remained the residence of the king until 536 when he moved it to London.

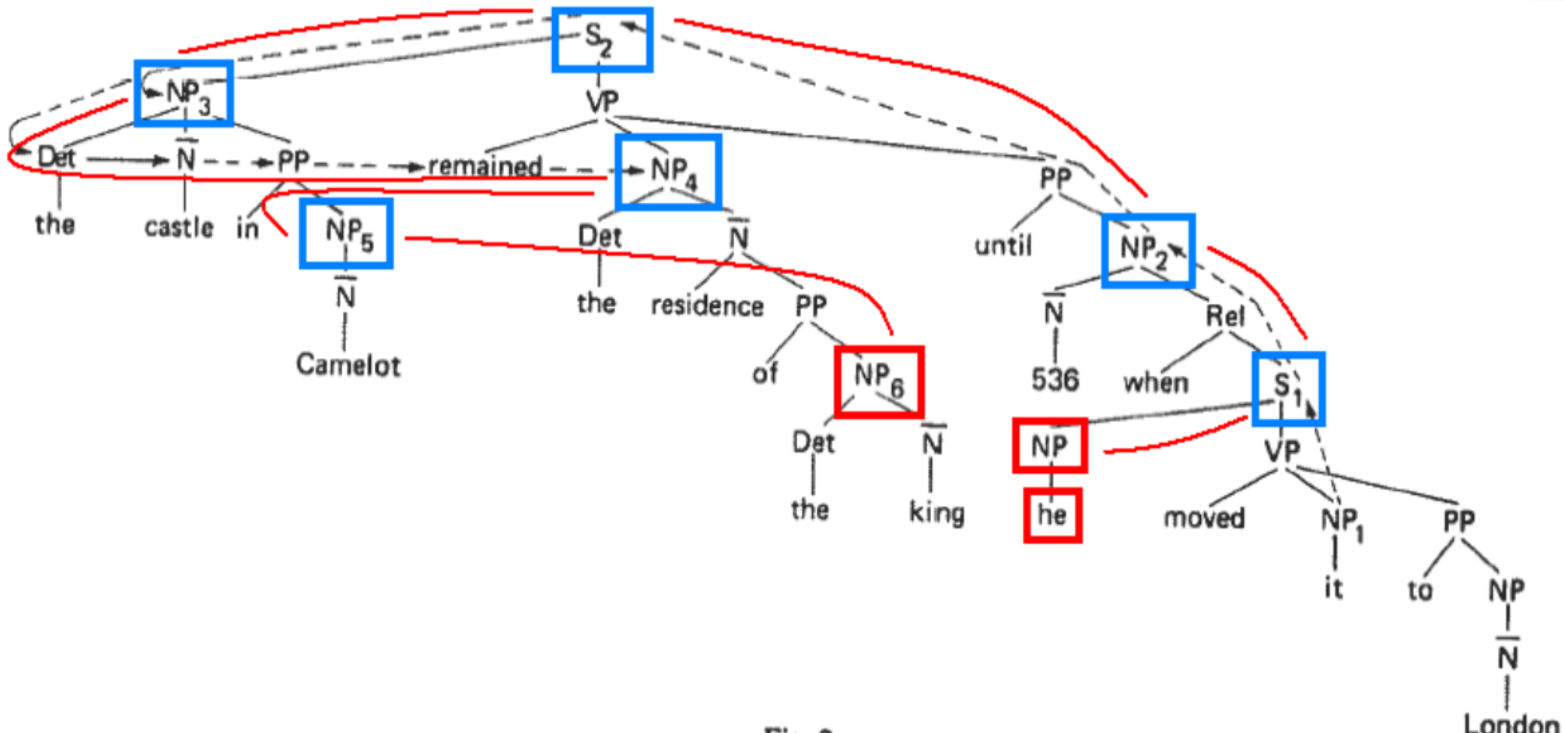

Fig. 2.

Hobbs, Jerry R., 1978, Resolving Pronoun References, Lingua, Vol. 44, pp.311-338.

Pronouns: Selectional Preferences

Prefer antecedent that is more likely to occur in context of pronoun

- Fred got a book and a coffee machine for his birthday. He read it the next day.
- can get probabilities from a large (parsed) corpus

Pronouns: Combining Probabilities

P = P (correct antecedent is at Hobbs distance d) ×

P (pronoun | head of antecedent) ×

P (antecedent | mention count) ×

P (head of antecedent | context of pronoun)

Ge, Hale, and Charniak 1998

83% success

Resolving Names

Generally straightforward: exact match or subsequence of prior name

- some exceptions for locations

Resolving Common Noun Phrases

generally difficult

typical strategies for resolving “the” + N:

- look for prior NP with same head N
- look for prior name including token N
  - “the New York Supreme Court” … the court

more ambitious: learn nouns used to refer to particular entities by searching for “name, N” patterns in a large corpus

- “Lazard Freres, the merchant bank”

Types of Models

mention-pair model

- train binary classifier: are two mentions coreferential?
- to apply model:
  - scan mention in text order
    - link each mention to the closest antecedent classified +
    - link each mention to antecedent most confidently labeled +
  - cluster mentions
- weak model of partially-resolved coreference

entity-mention model

- binary classifier: is a mention part of a partially-formed entity?
- richer model: entity has features from constituent mentions

Diversity of Approaches

Three recent systems show range of approaches:

Stanford [CL 2013]

- hand-coded rules
- 10 passes over complete document, using rules of decreasing certainty

Berkeley [EMNLP 2013]

- classifier trained over large corpus with simple feature set
- single pass

UW [EMNLP 2017]

- Performs mention detection and coreference in a single model
- Use FF over contextualized word embeddings to represent mentions

Systems generally do not work very well on anaphoric NPs

Sieve-based, Hand-coded System (Stanford)

sieve: set of hand-coded rules, applied starting with most precise rule

each rule applied across entire document (total of 10 passes)

rules reflect detailed linguistic analysis

most rules involve nominal anaphors; final pass (pass 10) resolves pronouns using agreement constraints

entity-centric model, uses information from all mentions gathered so far

All NPs, possessive pronouns, and named entity mentions are candidate mentions. Recall is more important than precision.

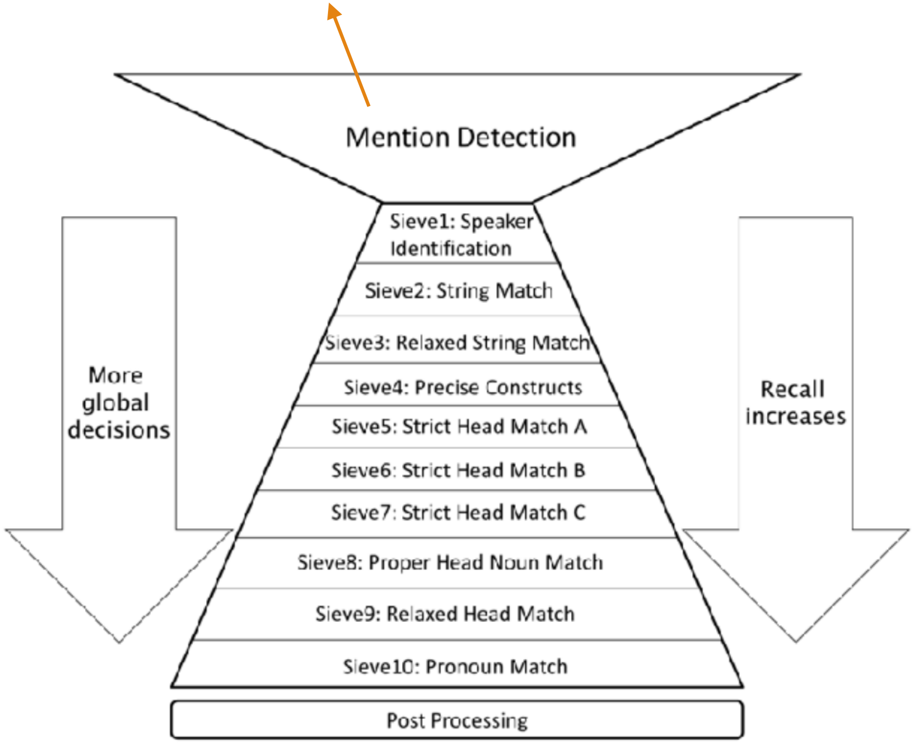

Lee et al., 2013: Deterministic Coreference Resolution Based on Entity-Centric, Precision-Ranked Rules

Shallow Feature Statistical System (Berkeley)

statistical approach based on large annotated corpus (OntoNotes)

mention-synchronous: single pass through document

features make minimal reference to specific linguistic phenomena

- large training corpus enables simple rules to capture most constraints

Anaphoric nominals remain the weak point for all approaches. Durrett and Klein report that when an anaphoric mention is a nominal or name, their system identifies the proper antecedent less than 8% of the time.

Durrett and Klein, 2013: Easy Victories and Uphill Battles in Coreference Resolution (EMNLP)

Features for the mention-pair models

Unary features (valid of a single token)

- Token, lemma, part of speech
- Salience

Binary features (valid of a pair of tokens)

- Number agreement (plural pronoun/plural NP)
- Gender agreement
- Sentence distance
- Hobbs distance
- Syntax: grammatical role
- …

Neural Networks for Coreference Resolution

Aim to learn a conditional probability distribution whose most likely configuration produces the correct clustering

pairwise score for a core
link between span i and
span j in *D*

$$P(y_1, \ldots, y_N \mid D) = \prod_{i=1}^{N} P(y_i \mid D)$$
$$= \prod_{i=1}^{N} \frac{\exp(s(i, y_i))}{\sum_{y' \in \mathcal{Y}(i)} \exp(s(i, y'))}$$

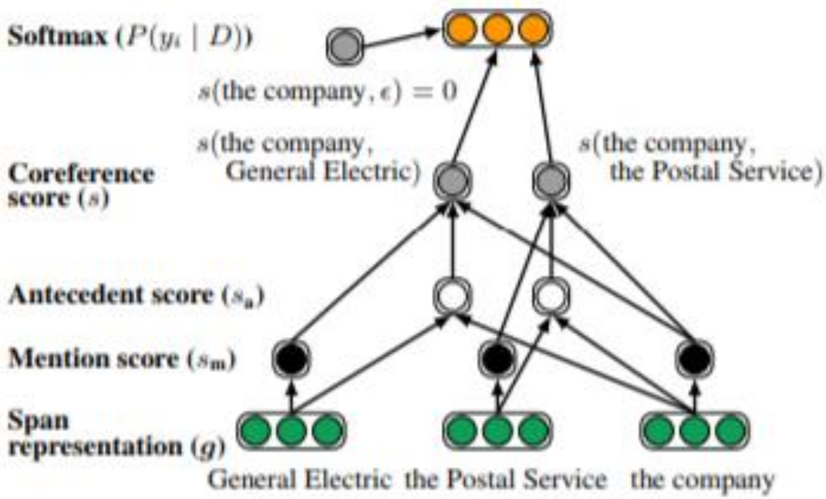

$$s(i, j) = \begin{cases} 0 & j = \epsilon \\ s_m(i) + s_m(j) + s_a(i, j) & j \neq \epsilon \end{cases}$$

$$s_m(i) = \boldsymbol{w}_m \cdot \mathrm{FFNN}_m(\boldsymbol{g}_i)$$
$$s_a(i, j) = \boldsymbol{w}_a \cdot \mathrm{FFNN}_a([\boldsymbol{g}_i, \boldsymbol{g}_j, \boldsymbol{g}_i \circ \boldsymbol{g}_j, \phi(i, j)])$$

Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. End-to-end neural coreference resolution. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017.

Neural Networks for Coreference Resolution con'd

Span representation: ELMO

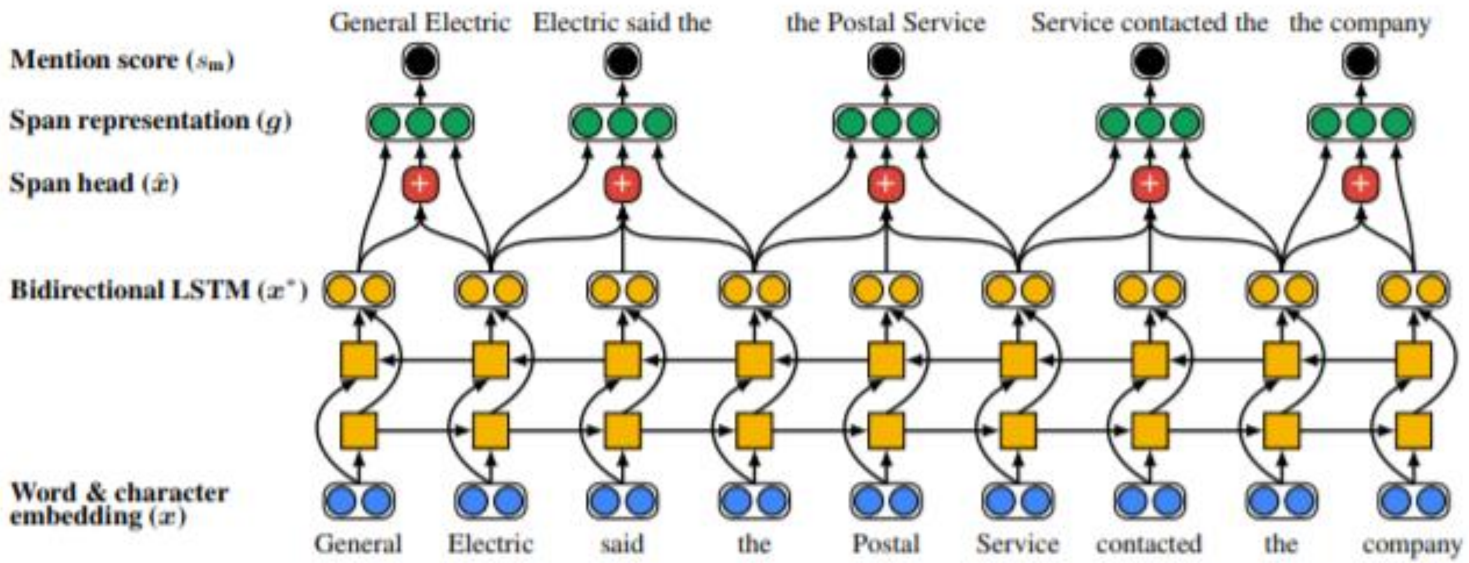

Span representations: BERT

Others?

Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. End-to-end neural coreference resolution. EMNLP 2017.
Mandar Joshi, Omer Levy, Luke Zettlemoyer, Daniel Weld. BERT for Coreference Resolution: Baselines and Analysis. EMNLP-IJCNLP 2019.

Evaluation

Coreference key is a set of links dividing the set of mentions into coreference classes

System response has similar structure

How to score response?

MUC scorer

- based on links …
  recall error = how many links must be added to system response so that all members of a key set are connected by links
- Does not give credit for correct singleton sets

Evaluation

$$B^3_{precision} = \frac{1}{n} \sum_{i}^{n} \frac{|Gold_i \cap System_i|}{|System_i|}$$

B-cubed metric:

- Mention-based

$$B^3_{recall} = \frac{1}{n} \sum_{i}^{n} \frac{|Gold_i \cap System_i|}{|Gold_i|}$$

- For each mention m,
    - r = size of response set containing m
    - k = size of key set containing m
    - i = size of intersection of these sets

  Recall(m) = i / k

  Precision(m) = i / r
- Then compute average of recall, average of precision

Example

Golden: 3 entity/coreference chains

1. {I, you, you, your, me, your, your, You} (8 elements)
2. {you, your father, you, him, I, your father} (6 elements)
3. {Obi-Wan, He} (2 elements)

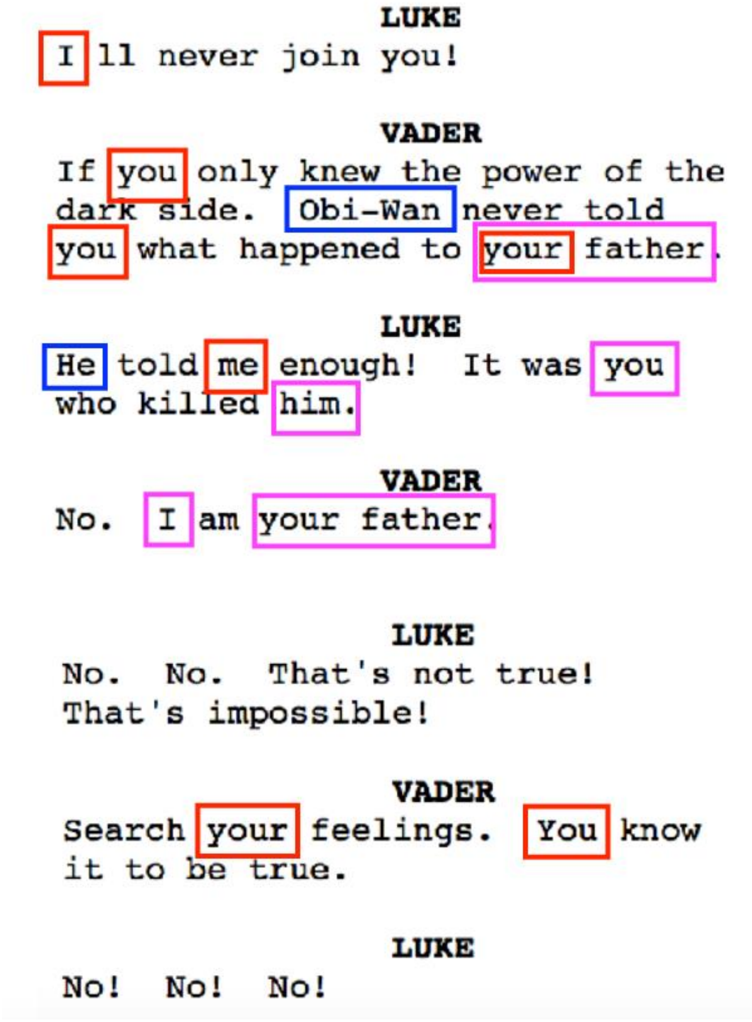
LUKE
I ll never join you!

VADER
If you only knew the power of the dark side. Obi-Wan never told you what happened to your father.

LUKE
He told me enough! It was you who killed him.

VADER
No. I am your father.

LUKE
No. No. That's not true! That's impossible!

VADER
Search your feelings. You know it to be true.

LUKE
No! No! No!

Example

System output: 4 entity/coreference chains

1. {I, me, I} (3 elements)
2. {you, you, you, your, you, your, your, you} (8 elements)
3. {Obi-Wan, your father, your father} (3 elements)
4. {He, him}

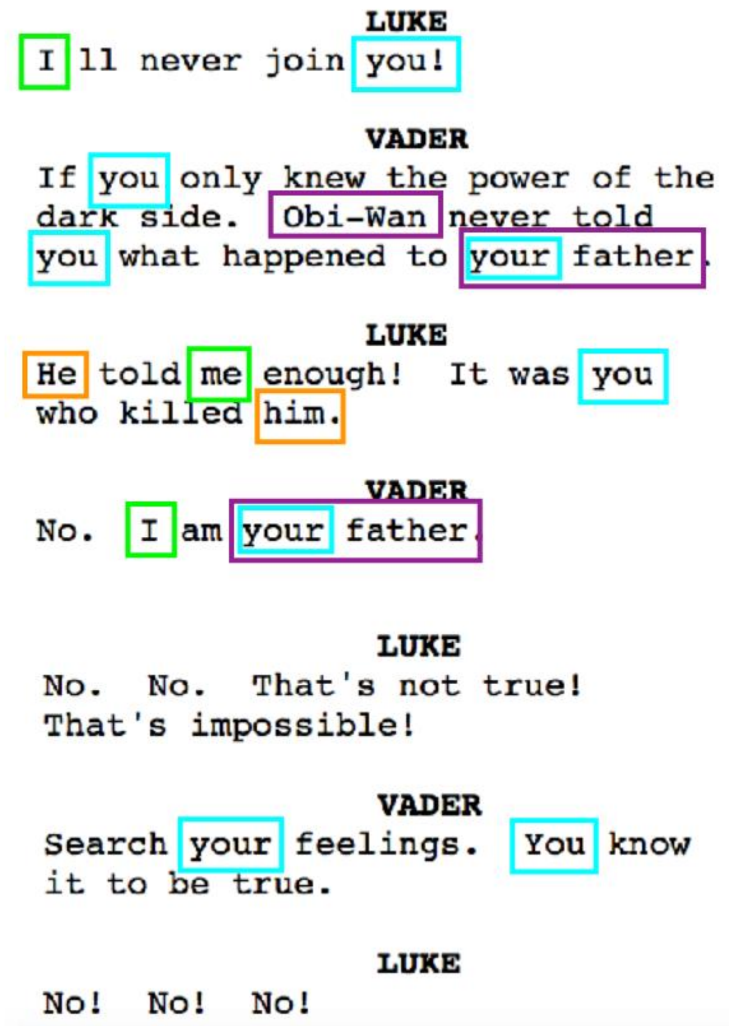

LUKE

I 'll never join you!

VADER

If you only knew the power of the dark side. Obi-Wan never told you what happened to your father.

LUKE

He told me enough! It was you who killed him.

VADER

No. I am your father.

LUKE

No. No. That's not true! That's impossible!

VADER

Search your feelings. You know it to be true.

LUKE

No! No! No!

Example

LUKE
I'll never join you!

VADER
If you only knew the power of the dark side. Obi-Wan never told you what happened to your father.

LUKE
He told me enough! It was you who killed him.

VADER
No. I am your father.

LUKE
No. No. That's not true! That's impossible!

VADER
Search your feelings. You know it to be true.

LUKE
No! No! No!

LUKE
I'll never join you!

VADER
If you only knew the power of the dark side. Obi-Wan never told you what happened to your father.

LUKE
He told me enough! It was you who killed him.

VADER
No. I am your father.

LUKE
No. No. That's not true! That's impossible!

VADER
Search your feelings. You know it to be true.

LUKE
No! No! No!

$\mid Gold_i \cap System_i \mid = 2$	$\mid Gold_i \mid = 8$	$\mid System_i \mid = 3$

Example

LUKE
I'll never join you!

VADER
If you only knew the power of the dark side. Obi-Wan never told you what happened to your father.

LUKE
He told me enough! It was you who killed him.

VADER
No. I am your father.

LUKE
No. No. That's not true! That's impossible!

VADER
Search your feelings. You know it to be true.

LUKE
No! No! No!

LUKE
I'll never join you!

VADER
If you only knew the power of the dark side. Obi-Wan never told you what happened to your father.

LUKE
He told me enough! It was you who killed him.

VADER
No. I am your father.

LUKE
No. No. That's not true! That's impossible!

VADER
Search your feelings. You know it to be true.

LUKE
No! No! No!

$\lvert Gold_i \cap System_i \rvert = 2$	$\lvert Gold_i \rvert = 6$	$\lvert System_i \rvert = 8$

A Coherent Discourse

A text is not a random collection of facts

A text will tell a story, make an argument, …

This is reflected in the structure of the text and the connections between sentences

Most of these connections are implicit, but a text without these connections is *incoherent*

Fred took an NLP course in the Spring.
He got a great job in June.

? Fred took an NLP course in the Spring.
He got a great cat in June.

A Coherent Discourse

Criteria for coherence depend on type of text

Most intensively studied for narratives

- causal connections
- temporal connections
- scripts (conventional sequences)

Coherence and Coreference

Select anaphora resolution more consistent with coherence.

Jack poisoned Sam. He died within a week. *vs.*

Jack poisoned Sam. He was arrested within a week.

How to do this in practice?

- Collect from a large corpus a set of predicate/role pairs, such as:

  subject of poison -- subject of arrest

  object of poison -- subject of die.
- Prefer anaphora resolution consistent with such pairs

Cross-document Coreference

Quite different from within-document coref:

within document (single author or editor)

- a single person will be consistently referred to by the same name
- the same name will consistently refer to the same person

across documents

- the same person may be referred to using different names
- a single name may refer to multiple people (“Michael Collins”)

Limitation

Assume each document separately resolved internally

Only link entities which are named in each document

- general NPs very hard to link
- “Fred’s wife” may refer to different people at different times
- details may change over time:
  - “the dozens of people killed in the bombing”
  - “the 55 people killed in the bombing”

Two Tasks

Entity linking: map each document-level entity to an entry in a standard data base

- e.g., wikification
- entities not in data base are left unlinked

Cross-document coreference

- cluster all document-level entities

Tasks have a lot in common

- often cross-doc coreference begins with entity linking against a large knowledge base or Wikipedia

Features

Features for cross-doc coref:

Internal (name) features

External (context) features

- whole-document features
- local context features
- semantic features

Consistency

Internal (Name) Features

Finding a match:

- exact match suitable for edited text in languages with standard romanization
- use edit distance for informal text
- use edit distance or pronunciation-based measure for other languages (e.g., Arabic)

Estimating probability of coref for exact match:

- for people, use name perplexity, based on
  - number of family names with same given name
  - number of given names with same family name

External (Context) Features

Names are more likely to be coreferential if:

- documents are similar (using tf-idf cosine similarity)
- local contexts are similar
- values of extracted attributes match (birthplace, religion, employer, …)

Conversely, distinct values of some attributes (birthplace, birthdate) are strong indicators of non-coreferentiality

Consistent Wikification

If multiple names are being resolved in a single document, they should preferably be resolved to related entities

- if “New York” and “Boston” are mentioned in the same sentence, prefer that
  - both resolve to cities
  - both resolve to baseball teams
  - both resolve to hockey teams
- in ranking referents, include as a factor the number of links connecting the referents

Consistent Wikification

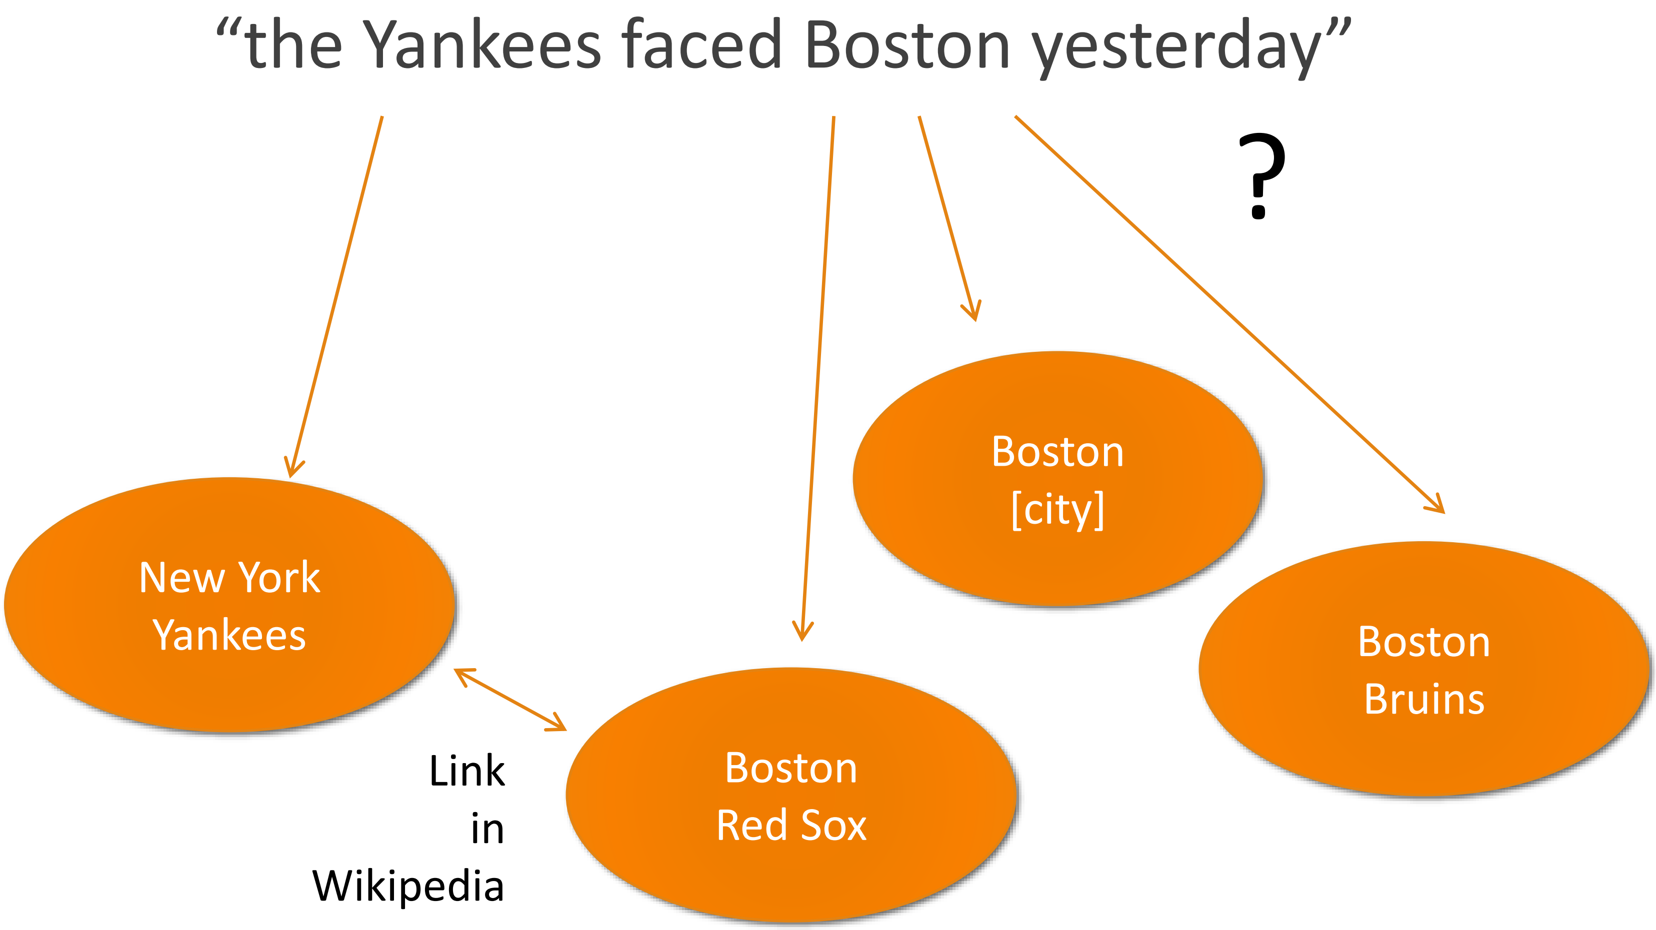

Scaling Up

Potential scale for cross-doc coref much larger

- collection may have 10^7 documents with 10-100 entities each: 10^9 document-level entities
- computing all pairwise similarities infeasible
- use hierarchical approach to divide set
  - analog of entity-mention representation within a document
  - potentially with multiple levels ('sub-entities')