

Natural Language Processing

Info 159/259

Lecture 16: Coreference resolution (March 18, 2024)

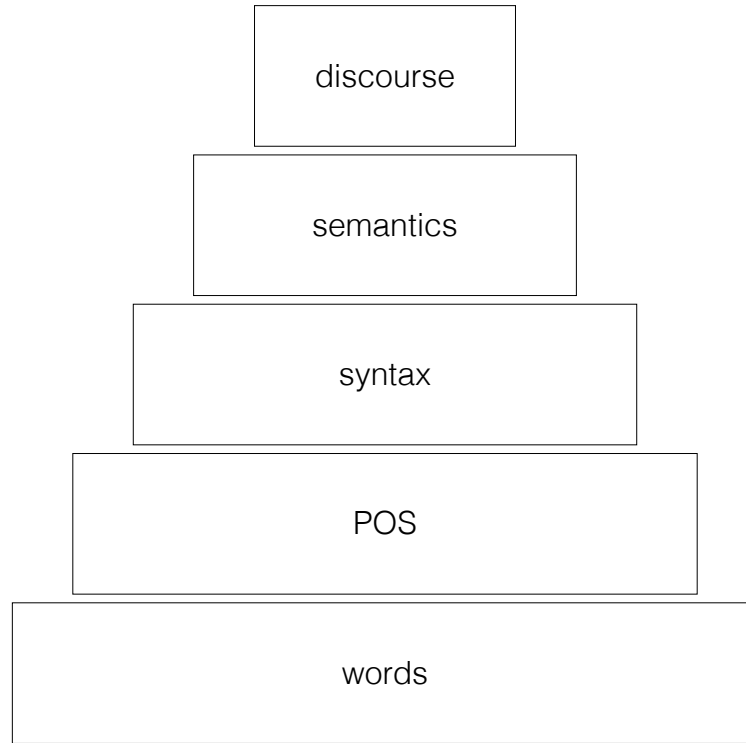
*Many slides & instruction ideas borrowed from:
David Bamman, Kemal Oflazer & Dan Jurafsky*

Logistics

- Quiz 7 is due tonight
- Homework 5 is due this Friday 3/22
- AP2 & 259 Mid-project report
- No quiz this week
 - (Exam 2 on Wednesday)

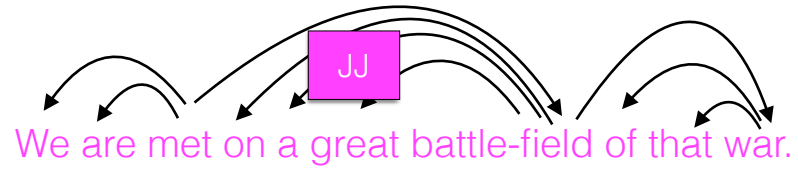
Exam 2

- Wednesday 3/20, 6:40-8:00pm PDT on bCourses.
- Same format as Exam 1
- You can expect more attention to the topics we cover in lectures and in the homeworks, but everything in the lectures and readings is fair game. The exam is cumulative, but expect more attention to topics covered between 2/26 (POS tagging) and 3/18 (coreference resolution), inclusive.
- We drop the lowest scoring grade of the two midterms, so if you don't take midterm 2, your course midterm grade will be your grade for midterm 1.



JJ

great



We are met on a great battle-field of that war.

The diagram features a pink square box containing the letters 'JJ'. From the top of this box, several black curved arrows point downwards to various words in the sentence below. Specifically, the arrows point to 'are', 'met', 'on', 'a', 'great', 'battle-field', and 'of'. The sentence itself is written in a pink color.

We are met on a great battle-field of that war.

Four score and seven years ago our fathers brought forth on this continent, a new nation, conceived in Liberty, and dedicated to the proposition that all men are created equal. Now we are engaged in a great civil war, testing whether that nation, or any nation so conceived and so dedicated, can long endure. **We are met on a great battle-field of that war.** We have come to dedicate a portion of that field, as a final resting place for those who here gave their lives that that nation might live. It is altogether fitting and proper that we should do this. But, in a larger sense, we can not dedicate -- we can not consecrate -- we can not hallow -- this ground. The brave men, living and dead, who struggled here, have consecrated it, far above our poor power to add or detract. The world will little note, nor long remember what we say here, but it can never forget what they did here. It is for us the living, rather, to be dedicated here to the unfinished work which they who fought here have thus far so nobly advanced. It is rather for us to be here dedicated to the great task remaining before us -- that from these honored dead we take increased devotion to that cause for which they gave the last full measure of devotion -- that we here highly resolve that these dead shall not have died in vain -- that this nation, under God, shall have a new birth of freedom -- and that government of the people, by the people, for the people, shall not perish from the earth.

Discourse

- Discourse covers linguistic expression **beyond the boundary of the sentence.**
 - Dialogues: the structure of turns in conversation
 - Monologues: the structure of entire passages, documents



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!



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!

Coreference resolution

Barack Hussein Obama II ( [/bəˈrɑːk huːˈseɪn ɒˈbɑːmə/](#); 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 [he](#) served as president of the [Harvard Law Review](#). [He](#) was a [community organizer](#) in Chicago before earning his [law degree](#). [He](#) worked as a [civil rights attorney](#) and taught [constitutional law](#) at the [University of Chicago Law School](#) from 1992 to 2004. [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.

Coreference resolution

attend graduate school at [Harvard University](#) on a scholarship. Obama's parents divorced in March 1964.^[11] Obama Sr. returned to Kenya in 1964 where he remarried; he visited Barack in Hawaii only once, in 1971.^[12] He died in an automobile accident in 1982 when his son was 21 years old.^[13]

Did Barack Obama die in an automobile accident in 1982?

Coreference resolution

“Victoria Chen, Chief Financial Officer of Megabucks Banking Corp since 2004, saw her pay jump 20%, to \$1.3 million, as the 37-year-old also became the Denver-based financial services company’s president. It has been ten years since she came to Megabucks from rival Lotsabucks.”

Coreference

“Referent”

The entities or individuals in the real world that the text is pointing to.

- VICTORIA CHEN
- MEGABUCKS
- LOTSABUCKS

“Victoria Chen, Chief Financial Officer of Megabucks Banking Corp since 2004, saw her pay jump 20%, to \$1.3 million, as the 37-year-old also became the Denver-based financial services company’s president. It has been ten years since she came to Megabucks from rival Lotsabucks.”

Coreference

“Referring expression”

The text that points to entities.

“Victoria Chen, Chief Financial Officer of Megabucks Banking Corp since 2004, saw her pay jump 20%, to \$1.3 million, as the 37-year-old also became the Denver-based financial services company’s president. It has been ten years since she came to Megabucks from rival Lotsabucks.”

Coreference

“coreference”

The set of text strings that all refer to the same ENTITY.

“Victoria Chen, Chief Financial Officer of Megabucks Banking Corp since 2004, saw her pay jump 20%, to \$1.3 million, as the 37-year-old also became the Denver-based financial services company’s president. It has been ten years since she came to Megabucks from rival Lotsabucks.”

Event coreference

I stubbed my toe on the chair and **it** really hurt.

Sense and reference

Mode of presentation vs.
reference (Frege)

- The morning star/the evening star
- Mark Twain/Samuel Clemens



English constraints

- Number
 - I have a car. They are blue [*they = car]
- Gender
 - My dad is shoveling snow. He's cold. [*he = snow]
- Person
 - We're watching a movie. He likes it [*he = you and I]

English exceptions

- Number

- A: *I have a new friend.*
B: *What's their name?*

- Gender

- “The Nellie, a cruising yawl, swung to **her** anchor without a flutter of the sails, and was at rest.” (Heart of Darkness)

- Person

- “I would prefer not to,” **he** said.

English preferences

- Recency: more recent NPs are preferred
- Grammatical role: subjects are preferred
 - Billy Bones went to the bar with Jim Hawkins. **He** called for a glass of rum.
- Repeated mention: more discourse-salient NPs are preferred (the ones with higher frequency)
- Parallelism
 - Long John Silver went with Jim to the Old Parrot. Billy Bones went with **him** to the Old Anchor inn.
- Verb semantics
- Selectional restrictions

Verb semantics

- John telephoned Bill. He lost the laptop
- John criticized Bill. He lost the laptop.

Winograd challenge

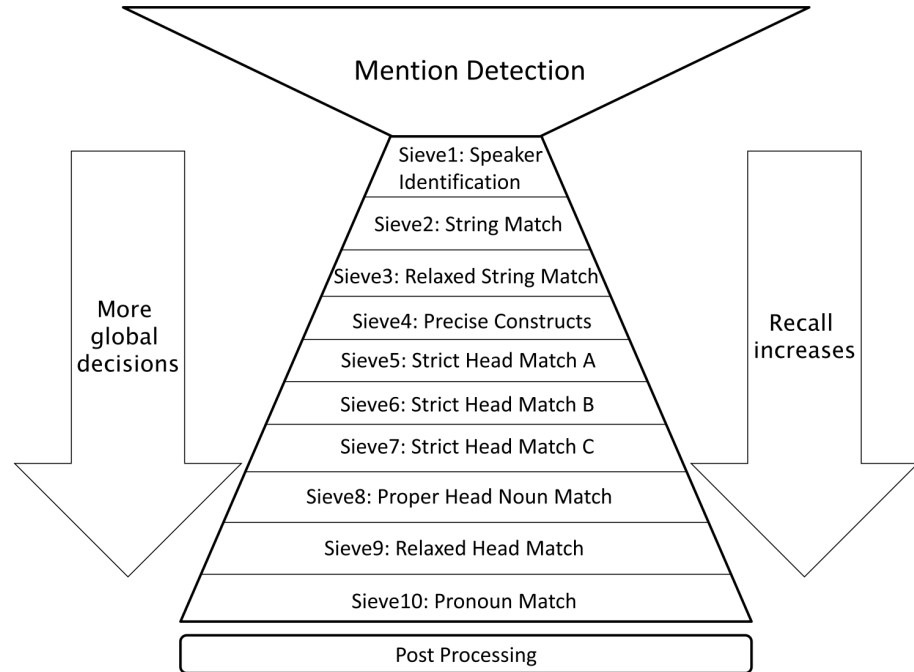
- The trophy would not fit in the brown suitcase because **it** was too big. What was too big?
- The town councilors refused to give the demonstrators a permit because **they** feared violence. Who feared violence?
- The town councilors refused to give the demonstrators a permit because **they** advocated violence. Who advocated violence?

Selectional restrictions

- John parked his car in the garage after driving **it** around for hours.

Stanford “Sieve”

Sequence of pattern matching rules starting at high precision coreference links, progressing to higher recall.



Mention Detection

- All NPs, possessive pronouns, and named entity mentions are **candidate mentions**. Recall is more important than precision.
- Filters to remove candidates:
 - Remove mentions embedded within larger mentions with same headword
 - Remove numeric quantities (100 miles, 9%)
 - Remove existential *there, it*
 - Remove 8 stop words (there, ltd., hmm)
 - ...

Mention-ranking models

- Proceed from the beginning of the document to the end: for each mention m_i in order, make a decision to either:
 - link it to a single *antecedent* mention $\{m_1, \dots, m_{i-1}\}$
 - leave it unlinked (thereby starting a new coreference chain).



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

No candidate antecedents,
so / starts a new
coreference chain



e1

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

Candidate antecedents:

[I]

Do we:

- Link *you* to *I*
- Establish a new entity



e1

e2

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

Candidate antecedents:
[I, you]

Do we:

- Link *you* to *I* (e1)
- Link *you* to *you* (e2)
- Establish a new entity



e1

e2

LUKE

I'll never join you!

e1

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

Candidate antecedents:

[I, you, you]

Do we:

- Link TPOTDS to I (e1)
- Link TPOTDS to you (e2)
- Link TPOTDS to you (e1)
- Establish a new entity



e1

e2

LUKE

I'll never join you!

e1

VADER

e3

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

Candidate antecedents:

[I, you, you, the power of the dark side]

Do we:

- Link *the dark side* to I (e1)
- Link *the dark side* to you (e2)
- Link *the dark side* to you (e1)
- Link *the dark side* to TPOTDS (e3)
- Establish a new entity



e1

e2

LUKE

I'll never join you!

e1

e3

VADER

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

e4

LUKE

He told me enough! It was you who
killed him

Candidate antecedents:

[I, you, you]


Do we:

- Link *the dark side* to *I* (e1)
- Link *the dark side* to *you* (e2)
- Link *the dark side* to *you* (e1)
- Link *the dark side* to *TPOTDS* (e3)
- Establish a new entity

Mention-ranking models

- Mention-ranking models carry out coreference as the result of individual decisions about each mention.
- A coreference chain is the transitive closure of all pairwise links.

Coreference chain 1:
{I, you, you, your, me}



Mention-ranking models

- The core machinery in a mention-ranking model is parameterizing the probability of a link between two mentions m_i and m_j .

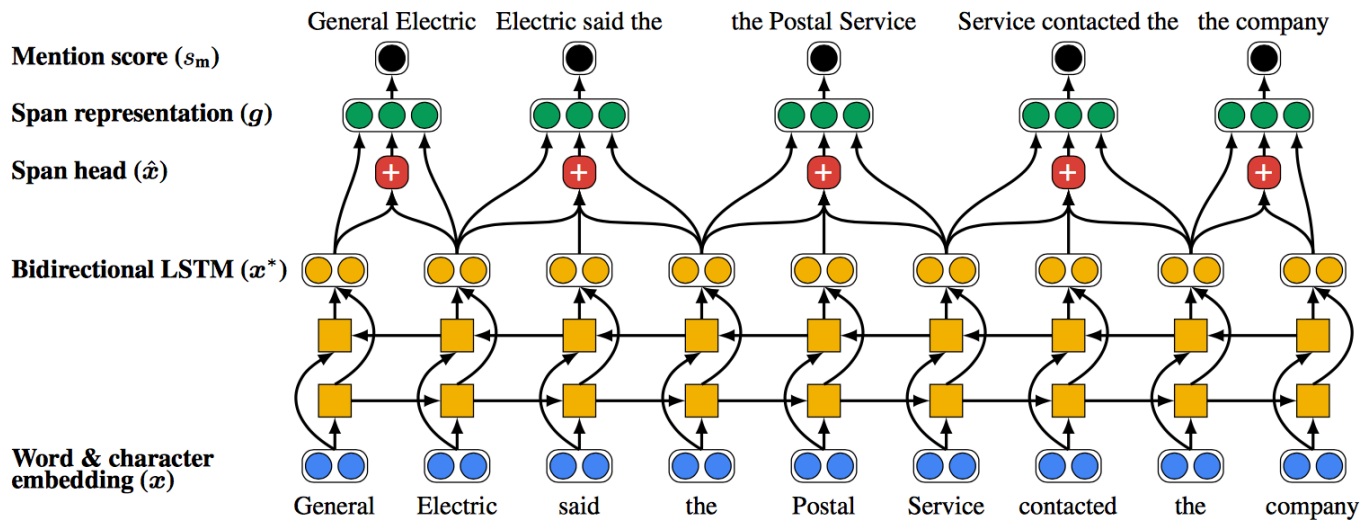
Featurized

- Features use information about the mention type (nominal, proper, pronoun), first/last word of mention, complete mention string, words immediately to left/right of mention, distance between mentions
- Decision to link to antecedent a_i is based on a linear scoring function involving a set of learned weights w and a feature function f .

$$\log P(a_i | x) \propto w^\top f(i, a_i, x)$$

Feature name
Features on the current mention
[ANAPHORIC] + [HEAD WORD]
[ANAPHORIC] + [FIRST WORD]
[ANAPHORIC] + [LAST WORD]
[ANAPHORIC] + [PRECEDING WORD]
[ANAPHORIC] + [FOLLOWING WORD]
[ANAPHORIC] + [LENGTH]
Features on the antecedent
[ANTECEDENT HEAD WORD]
[ANTECEDENT FIRST WORD]
[ANTECEDENT LAST WORD]
[ANTECEDENT PRECEDING WORD]
[ANTECEDENT FOLLOWING WORD]
[ANTECEDENT LENGTH]
Features on the pair
[EXACT STRING MATCH (T/F)]
[HEAD MATCH (T/F)]
[SENTENCE DISTANCE, CAPPED AT 10]
[MENTION DISTANCE, CAPPED AT 10]


Neural coref



Lee et al. (2018), "End-to-end Neural Coreference Resolution"

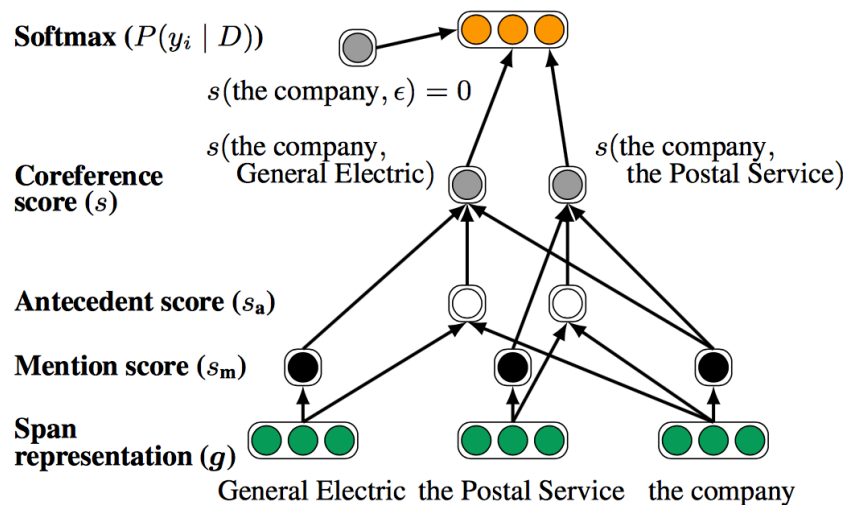
Neural coref

- Representation for mention =
 - BiLSTM output for first token in mention
 - BiLSTM output for last token in mention
 - Attention over BiLSTM output for all tokens in mention
 - Features: size of the mention



Neural coref

- Representation for mention pair (m_i, m_j) :
 - m_i representation g_i
 - m_j representation g_j
 - elementwise product of g_i and g_j
 - Features scoped over pair: distance between m_i and m_j



Neural coref

- Distance embeddings:
 - Discretize the distance between two mentions into a bucket ID

Distance(m_i, m_j)	Bucket
1	0
2-5	1
5-10	2
10-25	3
...	...

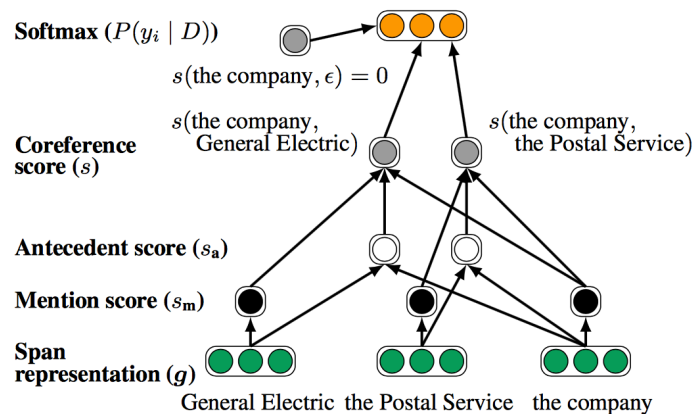
Neural coref

Each distance bucket then has a learnable embedding


0	2	-0.5	1.1	0.3	0.4	-0.5
1	-1.4	0.4	-0.2	-0.9	0.5	0.9
2	-1.1	-0.2	-0.5	0.2	-0.8	0
3	0.7	-0.3	1.5	-0.3	-0.4	0.1
4	-0.8	1.2	1	-0.7	-1	-0.4
5	0	0.3	-0.3	-0.9	0.2	1.4
6	0.8	0.8	-0.4	-1.4	1.2	-0.9
7	1.6	0.4	-1.1	0.7	0.1	1.6
8	1.2	-0.2	1.3	-0.4	0.3	-1.0

Neural coref

- Representation for mention pair (m_i, m_j):
 - m_i representation g_i
 - m_j representation g_j
 - elementwise product of g_i and g_j
 - Features scoped over pair: distance between m_i and m_j



Progress — Coreference resolution



Evaluation

- Evaluating general reference resolution (i.e., all noun phrase entities) is more complicated than straightforward accuracy/precision/recall

$$B_{precision}^3 = \frac{1}{n} \sum_i^n \frac{|Gold_i \cap System_i|}{|System_i|}$$

$$B_{recall}^3 = \frac{1}{n} \sum_i^n \frac{|Gold_i \cap System_i|}{|Gold_i|}$$

3 entities/coreference chains

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!

7 elements
{I, you, you, your, me, your, your, You}

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!

6 elements
{you, your father, you, him, I, your father}

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!

2 elements
{Obi-Wan, He}

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!

Example system output: 4 entities

3 = {I, me, I}

8 = {you, you, you, your, you, your, your, you}

3 = {Obi-Wan, your father, your father}

2 = {He, him}

Evaluation

- Evaluating general reference resolution (i.e., all noun phrase entities) is more complicated than straightforward accuracy/precision/recall

$$B_{precision}^3 = \frac{1}{n} \sum_i^n \frac{|Gold_i \cap System_i|}{|System_i|}$$

$$B_{recall}^3 = \frac{1}{n} \sum_i^n \frac{|Gold_i \cap System_i|}{|Gold_i|}$$

n ranges over all mentions in gold and system output



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!

$ \text{Gold}_i \cap \text{System}_i = 2$	$ \text{Gold}_i = 8$	$ \text{System}_i = 3$
--	-------------------------	---------------------------

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!

| Gold_i ∩ System_i | = 2

| Gold_i | = 6

| System_i | = 8

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!

| Gold_i ∩ System_i | = 6

| Gold_i | = 8

| System_i | = 8

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!

| Gold_i ∩ System_i | = 1

| Gold_i | = 2

| System_i | = 3

LUKE
I'll never join **you!**

VADER
 **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
 **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!

$|\text{Gold}_i \cap \text{System}_i| = 6$

$|\text{Gold}_i| = 8$

$|\text{System}_i| = 8$

Evaluation

- Evaluating general reference resolution (i.e., all noun phrase entities) is more complicated than straightforward accuracy/precision/recall

$$B_{precision}^3 = \frac{1}{n} \sum_i^n \frac{|Gold_i \cap System_i|}{|System_i|}$$

$$B_{recall}^3 = \frac{1}{n} \sum_i^n \frac{|Gold_i \cap System_i|}{|Gold_i|}$$

n ranges over all entities in gold and system output

Solve it

- Ontonotes
 - <http://catalog.ldc.upenn.edu/LDC2013T19>
- MUC 7
 - <http://catalog.ldc.upenn.edu/LDC2001T02>
- ACE 2003
 - <http://catalog.ldc.upenn.edu/LDC2001T02>