Coreference Resolution

COMP-599

Nov 3, 2015

Midterm

Tuesday, Nov 10, in class, 80 min.

Closed book

Question types:

- M/C
- Short answer
- Problem solving

Expectations

You'll be expected to know:

- Definitions (including formulas)
- How to run algorithms by hand
- Simple derivations
- Simple probability and calculus (up to, say, chain rule)
- Basic linguistic terminology and theory
- Lambda calculus

You won't be expected to remember*:

- Lagrange multipliers
- What the derivative of, say, e^x is.
- Specific lexical rules for translating items into FOL

*But could still be asked to apply or use, given relevant facts

How To Study

Review and redo the exercises we did in class

Learn the definitions of all the technical terms (most of them are in bold or are the titles of slides)

Review your assignments

Do practice problems

List of practice problems from the textbook are on the course website. Note that they are not meant to be comprehensive, but are supplementary.

Outline

Coreference and anaphora

Hobb's algorithm

Machine learning for coreference resolution

Coreference resolution tasks

Reference and Coreference

that cat
Whiskers
something furry
it

Referring expressions **Referent**

"*That cat*", "*Whiskers*", "*it*", and any other expression that point to the same referent are said to **corefer**.

Anaphora and Antecedents

In a passage:

> **Tardar Sauce** (born April 4, 2012), better known by her Internet name "Grumpy Cat", is a cat and Internet celebrity known for her grumpy facial expression.[1][3][4][5] Her owner, Tabatha Bundesen, says that her permanently grumpy-looking face is due to an underbite and feline dwarfism.
>
> https://en.wikipedia.org/wiki/Grumpy_Cat

An **anaphor** points to a *previous* linguistic expression, which is its **antecedent**.

Cataphora

Despite their name, cataphors are not anaphors that point to cats.

When he's grumpy, ***Whiskers*** *refuses to eat.*

Instead, a **cataphor** points to an antecedent that *follows* it.

Types of Referring Expressions

Proper names

McGill University

Whiskers

Montreal

Pronouns

I

you

it

their

ours

herself

More Types of Refering Expressions

Noun phrases

- **Indefinite**
  - *Some water*
  - *A deer*
  - *This random dude* (Note that *this* is ambiguous)
- **Definite**
  - *The cat*
  - *The election*

Demonstratives (They point to something)

- *That hotdog*
- *These problems*

Cross-linguistically Speaking

Zero anaphora

- Many languages omit pronouns in certain contexts.
- Often called **pro-drop** languages
- Computational task: *detect* and *resolve* them

Sometimes, you can tell what pronoun is missing:

- e.g., Spanish:

  No *habl-o* *español.*

  NOT Speak-1Sg Spanish

  (I) don't speak Spanish.

  Languages like this: Spanish, Italian, Russian, and many others

Omitting Pronouns

Other times, you really have to tell from the surrounding context.

- e.g., Japanese:

  ai shi -te- -ru.

  Love PROG PRES

  (I) love (you).

  But could also be *(He) loves (her). (They) love (me).*

  Languages like this: Japanese, Chinese varieties, Korean

This occasionally happens in informal English, usually in the first person.

Went to a dope COMP-599 class today.

Other Kinds of Reference

Bridging reference

Reference to entities that can be inferred from a previously mentioned entity

I like my office. The windows are large and the table is made of mahogany.

You should get a cactus. They are easy to care for.

Non-Referential Pronouns

Pleonastic pronouns

It is raining.

Snap out of it!

Clefting

It is COMP-599 which is giving me headaches.

- Used to put the focus on some point
- Seems marginally referential?

Pronominal Anaphora Resolution

What is some relevant information?

- Gender, number (*it* 3SG-inanimate vs. *we* 1PL)
- Syntactic information (grammatical role information)
- Recency

Binding Theory (Chomsky, 1981)

Chomsky defined syntactic constraints for determining when an antecedent can **bind** a referring expression.

The students taught themselves. *[themselves = the students]*

The students taught them. *[them ≠ the students]*

Reflexives must be bound by a subject in a certain syntactic relationship called **c-command**. Personal pronouns *must not* be bound in this way.

Hobb's Algorithm (1978)

A traversal algorithm which requires:

- Constituent parse tree
- Morphological analysis of number and gender

Overall steps:

1. Search the current sentence right-to-left, starting at the pronoun
2. If no antecedent found, search previous sentence(s) left-to-right

Steps in Hobb's Algorithm

1. Begin at the NP node immediately dominating the pronoun.
2. Go up to the first NP or S above it. Call this node X and the path to it p.
3. Do a left-to-right breadth-first traversal of all branches below X to the left of p . Propose as antecedent any NP node encountered that has an NP or S between it and X.
4. If X is the highest S in the sentence, consider the parse trees of previous sentences in recency order, and traverse each in turn in left-to-right breadth-first order. When an NP is encountered, propose it as an antecedent. If X is not the highest S, continue to step 5

Steps in Hobb's Algorithm

5. From X, go up to the first NP or S above it. Call this new node X and the path to it p.
6. If X is an NP and p doesn't pass through the Nominal that X immediately dominates, propose X as an antecedent.
7. Do a left-to-right breadth-first traversal of all branches below X to the left of p. Propose any NP encountered as the antecedent.
8. If X is an S, do a left-to-right breadth-first traversal of all branches below X to the right of p, but don't go below any NP or S encountered. Propose any NP encountered as the antecedent.
9. Go to step 4.

Example of Hobb's Algorithm

Alice saw a beautiful cupcake in the patisserie window.

She showed it to Bob.

She devoured it.

Assume a standard parse of the sentences of the type we have been drawing in this class.

Assume a perfect gender/entity type checker.

Exercise: Run Hobb's Algorithm

Parse the following sentences and run Hobb's algorithm on the pronouns in it.

Bob opened up a new dealership last week.

John took a look at the Acuras in his lot.

He ended up buying one.

Other Heuristic Approaches

Lappin and Leass (1994)

- Assigns weights to the various factors that we have discussed by hand.

Centering Theory (Grosz et a., 1995)

- A theory of entity transitions in a discourse, looking at their syntactic positions in sentences.
- Brennan et al. (1987) used this for pronominal anaphor resolution. Antecedent is selected in order to yield a series of entity transitions that are preferred, according to Centering Theory.

Coreference Resolution by ML

Soon et al. (2001) defined 12 features for NP coreference resolution (not just pronominal):

Feature Type	Feature	Description
Lexical	SOON_STR	C if, after discarding determiners, the string denoting NP_i matches that of NP_j; else I.
Grammatical	PRONOUN_1*	Y if NP_i is a pronoun; else N.
	PRONOUN_2*	Y if NP_j is a pronoun; else N.
	DEFINITE_2	Y if NP_j starts with the word "the;" else N.
	DEMONSTRATIVE_2	Y if NP_j starts with a demonstrative such as "this," "that," "these," or "those;" else N.
	NUMBER*	C if the NP pair agree in number; I if they disagree; NA if number information for one or both NPs cannot be determined.
	GENDER*	C if the NP pair agree in gender; I if they disagree; NA if gender information for one or both NPs cannot be determined.
	BOTH_PROPER_NOUNS*	C if both NPs are proper names; NA if exactly one NP is a proper name; else I.
	APPOSITIVE*	C if the NPs are in an appositive relationship; else I.
Semantic	WNCLASS*	C if the NPs have the same WordNet semantic class; I if they don't; NA if the semantic class information for one or both NPs cannot be determined.
	ALIAS*	C if one NP is an alias of the other; else I.
Positional	SENTNUM*	Distance between the NPs in terms of the number of sentences.

Table from (Ng and Cardie, 2002)

Soon et al., 2001

They trained a supervised decision tree classifier using these features.

Results on MUC-6 data set:

58.6/67.3/62.6 in terms of R/P/F1

Ng and Cardie, (2002) extended the feature set.

62.4/73.5/67.5

Durrett and Klein (2013) incorporated many features into a log-linear model (~3M).

Word-level features + simple recency, syntax, and gender/number features actually work very well.

Other Types of Reference Resolution

Event coreference resolution

Anaphoric shell nouns

Cross-document coreference resolution

Event Coreference Resolution

s1: Hewlett-Packard is negotiating to **[buy]** technology services provider Electronic Data Systems.

s8: With a market value of about $115 billion, HP could easily use its own stock to finance the **[purchase]**.

s9: If the **[deal]** is completed, it would be HP's biggest **[acquisition]** since it *[bought]* Compaq Computer Corp. for $19 billion in 2002.

(Bejan and Harabagiu, 2010)

Event Coreference Resolution

What does it mean for events to corefer?

- Same causes and effects (Davidson, 1969)
- **Happen in same time and place** (Quine, 1985)

Cues for event coreference (Bejan and Harabagiu, 2010):

- Share same event properties
- Share same participants

"*This*"-anaphora

Anaphoric shell nouns provide nominal shells for complex chunks of information (Kolhatkar et al., 2013).

Despite decades of education and widespread course offerings, the survival rate for out-of-hospital cardiac arrest remains a dismal 6 percent or less worldwide.

This fact prompted the American Heart Association last November to simplify the steps of CPR to make it easier for lay people to remember and to encourage even those who have not been formally trained to try it when needed.

Cross-document Alignment

Even trickier: multiple documents discussing an overlapping set of events and entities

View as an alignment problem (Wolfe et al., 2013; Roth and Frank, 2012)

(2) a. The regulator ruled on September 27 that Nasdaq too was qualified to bid for OMX [...][3]

b. The authority [...] had already approved a similar application by Nasdaq.[4]

Sure alignment

(3) a. Myanmar's military government said earlier this year it has released some 220 political prisoners [...][5]

b. The government has been regularly releasing members of Suu Kyi's National League for Democracy party [...][6]

Possible alignment

(Roth and Frank, 2012)

References (Others in J&M)

- Bejan and Harabagiu. 2010. Unsupervised event coreference resolution with rich linguistic features. *ACL*.
- Durrett and Klein. 2013. Easy Victories and Uphill Battles in Coreference Resolution. *EMNLP*.
- Kolhatkar, Zinsmeister, and Hirst. 2013. Annotating Anaphoric Shell Nouns with their Antecedents. *LAWS*.
- Ng and Cardie. 2002. Improving Machine Learning Approaches to Coreference Resolution. *ACL*.
- Roth and Frank. 2012. Aligning Predicate Argument Structures in Monolingual Comparable Texts: A New Corpus for a New Task. *Joint Conference on Lexical and Computational Semantics.*
- Soon, Ng and Lim. 2001. A machine learning approach to coreference resolution of noun phrases. *Computational Linguistics*.
- Wolfe et al. 2013. PARMA: A Predicate Argument Aligner. *ACL*.