Lecture 12: Anaphora Resolution

Lexical Semantics and Discourse Processing
MPhil in Advanced Computer Science

Simone Teufel

Natural Language and Information Processing (NLIP) Group
UNIVERSITY OF CAMBRIDGE
Simone.Teufel@cl.cam.ac.uk

Slides after Advaith Siddharthan

March 2, 2011

1. Referring Expressions
   - Cognitive Status and Givenness Hierarchy
   - Syntactic Constraints
   - Salience
2. Pronoun resolution algorithms
   - Hobbs
   - Lappin and Leass
   - Ge et al.

Reading:

- Jurafsky and Martin, chapter 21.3-21.6

Referring Expressions

From The Aventures of Tom Sawyer by Mark Twain

The old lady pulled her spectacles down and looked over them about the room; then she put them up and looked out under them. She seldom or never looked THROUGH them for so small a thing as a boy; they were her state pair, the pride of her heart, and were built for "style," not service—she could have seen through a pair of stove-lids just as well.

Referring Expressions

Not Mark Twain...

Aunt Polly *pulled Aunt Polly's spectacles down and looked over Aunt Polly's spectacles about the room; then* **Aunt Polly** *put Aunt Polly's spectacles up and looked out under Aunt Polly's spectacles.* **Aunt Polly** *seldom or never looked THROUGH Aunt Polly's spectacles for so small a thing as a boy...*

This one neither (all pronominalised)...

She *pulled* **them** *down and looked over them about* **it***; then she put them up and looked out under them. She seldom or never looked THROUGH them for so small a thing as* **that***; they were her state pair, the pride of* **it***, and were built for "style," not service—She could have seen through* **them** *just as well.*

Appropriate use of referring expressions reduces communication effort for both listener and speaker.

Motivation

- **Machine Translation**: translate from languages with grammatical gender into English (*elle*→ *she?/it?*
- **Information Extraction**: merge information about same referent
- **Text Summarisation**: Identify salient entities and events
- **Question Answering and Information Retrieval**: better question/answer matching

They also. . .

- are frequent
- display a wide range of reference phenomena
- are central to discourse theories

Terminology

- **anaphora:** the phenomenon of referring to an antecedent (metonymically also refers to the referring expression). Subtypes are pronouns and definite NPs.
- **referent:** a real world entity that some piece of text (or speech) refers to.
- **referring expressions:** bits of language used to perform reference by a speaker.
- **coreference:** two references to the same referent
- **antecedent:** the text evoking a referent.
- **cataphora:** the phenomenon where the referring expression precedes the antecedent (metonymically also refers to the referring expression)
  - After **his** class, John will play football.

Anaphora resolution vs. coreference resolution

Anaphora resolution

Task of finding an antecedent for each anaphor (typically, pronoun).

Reference resolution

Task of partitioning the set of all referring expressions into equivalence classes (chains) that refer to one referent.

Types of referring expressions

- **Indefinite Noun Phrase:** introduce new entities into the discourse; e.g., *a pair of stove-lids*
- **Proper Noun:** evoke uniquely identifyable known entity.
- **Definite and Demonstrative Noun Phrase:** refer to entities that are uniquely identifiable by the listener; e.g., *the room*. (Not all definite NPs are referring, e.g. *the fact that the earth is round; the US president*)
- **Personal Pronoun:** refers to entities that have high level of activation in the listener's attentional state; e.g., *her, them*.
- **Demonstrative Pronoun:** can refer to entities and to events (e.g., *I had not expected that*).
- **One-Anaphora:** select one from a set of entities. It can introduce a new entity into the discourse, but this is dependent on an existing representation for the larger set; e.g., *I would like one*.

Types of Reference

- **Coreference**
  - referring expression refers to an entity that has been explicitly evoked
    - John owns a car.
    - **It** is a Ford.
- **Bridging Reference**
  - refer to entities that are inferable from previously evoked entities
    - John's car is very old.
    - **The engine** is noisy and **a door** is dented.
  - can involve *Synonymy*, *Hyponymy*, *Meronymy*
  - or other form of inference, e.g.,
    - I bought an iPad today.
    - **They** are so cool.

Non-referential usage

- Cleft: *It was Frodo who took the ring.*
- Pleonastic: *It was raining.*
- Extraposition: *It was unnecessary to repeat it.*

Cognitive Status Constraints

- Form of referring expression that is appropriate in any given context depends on
  - Attentional State of Listener
  - Shared Knowledge between Speaker and Listener
- Example from Gundel et al. (1993)):
  - *I could not sleep last night.*
    1. *A dog next door kept me awake.* (type identifyable)
    2. *This dog next door kept me awake.* (referential)
    3. *The dog next door kept me awake.* (uniquely referential)
    4. *That dog next door kept me awake.* (familiar)
    5. *That kept me awake.* (activated)
    6. *It kept me awake.* (in focus)

Cognitive Status

- **type identifiable:** Listener is able to access a representation of the object type (in 1, knowing what a dog is).
- **referential:** Listener can either retrieve from memory the specific dog referred to, or construct a new representation for this specific dog.
- **uniquely identifiable:** Listener can uniquely identify the intended referent on basis of the noun phrase alone.
- **familiar:** Listener already has an accessible representation in memory. (4 can be used if the listener knows there is a dog next door.)
- **activated:** Listener has immediate access to the referent, i.e., it is in short-term memory, either through discourse or real world. (5 is acceptable if the listener can hear the dog barking.)
- **in focus:** The referent is the focus in the discourse, not only in short-term memory (compare to 5).

Givenness Hierarchy

focus > *activated* > *familiar* > *unique* > *referential* > *type_identifiable*

	Focus	Activated	Familiar	Unique	Referential	Type Identifiable
English	it	HE, this, that, this N	that N	the N	indef., this N	a N
Chinese	∅, ta (he, she, it)	TA, zhe, nei, zhe N (this, that N)		nei N		vi N (a N), ∅ N
Japanese	∅	kare (he), kore (this), sore (that-medial), are (that-distal), kono N (this N), sono N (that-medial N)	ano N (that-distal N)	∅ N		
Russian	∅, on (he)	ON, eta (this), to (that)	eto N (this N), to N (this N)	∅ N		
Spanish	∅, el (he)	EL, este (this), ese (that-medial), aquel (that-distal), este N (this N)	ese N (that-medial N), aquel N (that-distal N)	el N (the N)	∅ N, un N (a N)	

Agreement Constraints on Coreference

- **number** = singular, plural
- **person** = first, second, third
- **gender** = masculine, feminine, non-personal
- **case** = nominative, accusative, genitive

	First Person		Second Person		Third Person	
	Singular	Plural	Singular	Plural	Singular	Plural
Nominative	*I*	*we*	*you*	*you*	*he, she*	*they*
Accusative	*me*	*us*	*you*	*you*	*him, her*	*them*
Genitive	*my*	*our*	*your*	*your*	*his, her*	*their*

Binding Theory (Chomsky, 1981)

Principle A: Reflexives must have local antecedents:

- *John$_i$ washed himself$_i$*
- **John$_i$ asked Mary to wash himself$_i$*

Principle B: Personal pronouns must not have local antecedents

- *John$_i$ asked Mary to wash him$_i$*
- **John$_i$ washed him$_i$*

Principle C: A referring expression cannot have an antecedent that c-commands it.

- **He$_i$ asked Mary to wash John$_i$.*
- **The car had a trailer$_i$ behind it$_i$.*

(A c-commands B if and only if neither A dominates B nor B dominates A; and every branching node that dominates A, also dominates B.)

Semantic Constraints on Coreference

In general, any shared knowledge between the speaker and the listener can be used to constrain the choice of referring expression. In particular:

- **Selectional Restrictions**
  - *Jerry bought coffee from the store. Henry drank* **it**.
- **Verb semantics and "implicit cause"**
  - *John telephoned Bill. He had lost the laptop.*
  - *John criticised Bill. He had lost the laptop.*
- **Discourse Accessibility**
  - *George didn't buy a Volvo. *It was blue.*

Salience and Preferences

- **Recency:** Entities introduced in recent utterances are more likely to be referred to by a pronoun than entities introduced in utterances further back.
- **Grammatical Role:** Entities introduced in subject position tend to get topicalised, and are more likely to be referred to by a pronoun than entities in object positions.
- **Repetition:** Entities that have already been referred to frequently are more likely to be pronominalised than those that have not.
    - *George needed a new car. His previous car got totaled, and he had recently come into some money. Jerry went with him to the car dealers. He bought a Nexus.*
- **Parallelism:** Pronouns are more likely to refer to those entities that do not violate syntactically parallel constructions.
    - *John took Bill to the zoo; Mary took him to the park.*

Pronoun Resolution

- Many factors influence pronoun resolution
- Many of these factors might contradict each other for specific examples
- No pronoun resolution algorithm successfully accounts for all these factors
- Next: three pronoun resolution algorithms
    - Purely syntax-based (Hobbs)
    - Salience model (Lappin & Leass)
    - Supervised ML (Ge et al.)
- These give a broad overview of the field

Hobbs' (1978) Algorithm

- Simple syntax-based algorithm for 3rd person anaphoric pronouns
- Relies on:
    - syntactic parser (with X-Bar output)
    - morphological number and gender checker
- Searches syntactic trees of current and preceding sentences in breadth-first, left-to-right manner. Stops when it finds matching NP.

Simplified Algorithm

- Right-to-left search in current sentence, starting with first c-commanding NP to the right of the pronoun
- While no antecedent found, left-to-right search in preceding sentence; step back sentence by sentence.
- If still no antecedent found, search current sentence from left-to-right, starting with first NP to the right of pronoun (for cataphora).

Hobbs: An Example where it gets it right

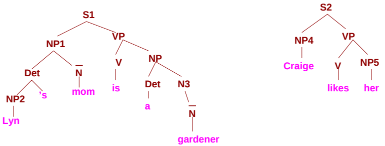

- Start search at NP5 in S2.
- Reject NP4 as no NP node between it and X (S2).
- What would have happened if the subject was *Craige's mom*?
- Move to S1. NP1 is first NP we encounter, so finish.
- Result: *Lyn's mom*

Hobbs: An Example where it gets it wrong

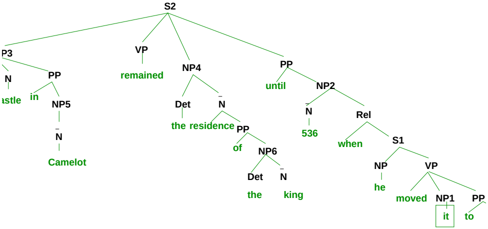

Exercise: What does it do?

Lappin and Leass

Two different operations are performed:

- Maintaining and updating a discourse model consisting of a set of *co-reference classes*:
  - Each co-reference class corresponds to one entity that has been evoked in the discourse
  - Each co-reference class has an updated *salience* value
- Resolving each Pronoun from left to right
  - Collect potential referents from up to 4 sentences back
  - Filter out coreference classes that don't satisfy agreement/syntax constraints
  - Select remaining co-reference class with the highest salience value; add pronoun to class.

Salience

- The salience of a referent is calculated on the basis of recency and grammatical function.

Salience Factor	Example	Weight
Current sentence		100
Subject emphasis	**John** opened the door	80
Existential emphasis	There was **a dog** standing outside	70
Accusative emphasis	John liked **the dog**	50
Indirect object	John gave a biscuit to **the dog**	40
Adverbial emphasis	*Inside **the house**, the cat looked on	50
Head Noun emphasis	**The cat** in the house looked on	80

Salience

- The salience of a referent is the sum of all applicable weights
- The salience of a referent is halved each time a sentence boundary is crossed
  - This, along with the weight for being in the current sentence, makes more recent referents more salient
- Weights are calculated for each member of the salience class
  - Previous mentions can boost the salience of a coreference class
  - This accounts for the repetition effect
- Lappin and Leass report 86% accuracy for their algorithm on a corpus of Computer manuals

A Worked Example

Example

Niall Ferguson is prolific, well-paid and a snappy dresser. Stephen Moss hated ***him****, at least until he spent an hour being charmed in the historian's Oxford study.*

- Discourse Referents:
  - $N_0 = \{$*Niall Ferguson*$\} = 105$
    (subject 80+ head-noun emphasis 80 + non-adverbial 50)/2
  - $S_0 = \{$*Stephen Moss*$\}$ ***does not pass syntax filter***
- New Discourse referents
  - Add *him* to N_0; $N_1 = \{$*Niall Ferguson, him*$\}$

A Worked Example

Example

Niall Ferguson is prolific, well-paid and a snappy dresser. Stephen Moss hated him, at least until ***he*** *spent an hour being charmed in the historian's Oxford study.*

- Discourse Referents:
  - $N_1 = \{$*Niall Ferguson, him*$\} = 405$
    (subject 80+ head-noun emphasis 80 + non-adverbial 50)/2 + direct object 70+ head-noun emphasis 80 + non-adverbial 50 + recency 100
  - $S_1 = \{$*Stephen Moss*$\} = 310$
    subject 80 + head-noun emphasis 80 + non-adverbial 50 + recency 100
- New Discourse referents
  - Add *he* to N_1; $N_2 = \{$*Niall Ferguson, him, he*$\}$

Ge et al. Algorithm

- The algorithm by Ge et al. (1998)
  - does not use an explicit model of discourse
  - collapses the distinction between **hard** constraints and soft preferences
    - Gender information is often noisy (eg: *Clinton*, *Alex* etc)
    - Number agreement not an absolute constraint in all cases

U1. I bought **an iPad** today.
U2. **They** are so cool.

U1. Maybe the key is under **a flowerpot**.
U2. Try looking under **them**.

- They use a Bayesian Approach that incorporates all factors in a machine learning framework.

Ge et al. Algorithm

- Features are derived from agreement values, grammatical roles, recency and repetition
- Calculate the probability $p(a|p, f_1...f_n)$ that a is the antecedent of a pronoun p given the features f_{1-n}.
- Pronoun is resolved by maximising $P(a_i|p, f_{1-n})$ over all potential antecedents a_i.

Bootstrapping Gender Information

Unsupervised approach to learning gender information:

- First run Hobbs' algorithm on the entire Penn Treebank (WSJ)
- Count number of times a noun was labelled as the antecedent of *he/his/him/himself*, *she/her/herself/hers* and *it/its/itself*
- This allows to compute $p(m|w_i)$, $p(f|w_i)$ and $p(n|w_i)$ for every word w_i in Penn Treebank (the probabilities that a word w_i is male, female or inanimate)
- Now use (preliminary) gender information to improve the pronoun resolution algorithm
- This results in recalculation of revised gender probabilities for all words in the Penn Treebank.

Ge et al. results

- Ge et al. report 82.9% of pronouns resoved correctly by their algorithm.
  - removing the syntax features brings the accuracy down to 43%
  - providing perfect gender information improves the accuracy to 89.3%

Summary

- Referring expressions and cognitive status
- Salience Factors:
  - Recency
  - Grammatical position
  - Repetition
  - Parallelism
- Knock-out Criteria:
  - Clashes in Gender, Number
  - Binding Theory
- Three algorithms:
  - Hobbs
  - Lappin and Leass
  - Ge et al

References

- Niyu Ge, John Hale, and Eugene Charniak. 1998. A statistical approach to anaphora resolution. In Proceedings of the Sixth Workshop on Very Large Corpora, pages 161171, Montreal, Canada.
- J.K. Gundel, N. Hedberg, and R. Zacharski. 1993. Cognitive Status and the Form of Referring Expressions in Discourse. Language, 69(2):274307. Emiel Krahmer, Sebastiaan van Erk, and Andr Verleg. 2003. Graph-based generation of referring expressions. Computational Linguistics, 29(1):5372.
- Shalom Lappin and Herbert Leass. 1994. An algorithm for pronominal anaphora resolution. Computational Linguistics, 20(4):535561.