L113 Word Meaning and Discourse Understanding

Session 7: Coreference Resolution

Simone Teufel

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

2012/2013

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).

Coreference 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 identifyable)
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:

***John**$_i$ *saw* **John**$_i$.

c-command: the relationships "brother, uncle, great-uncle, great-great-uncle . . ."

C-command

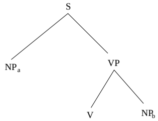

- NP_a c-commands NP_b if and only if neither NP_a dominates NP_b nor NP_b dominates NP_a; and every branching node that dominates NP_a, also dominates NP_b.
- c-command prevents coreference between a c-commanded NP and the commanding NP (unless a reflexive pronoun is used locally)

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.

Salience and Preferences

- **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.

Hobbs' Algorithm

1. Find the lowest node N1 which is an NP or S and contains the NP above pronoun P
2. Check the children of N1 left to right for NPs to the right of P that do not c-command P and do not violate morphological constraints; propose the leftmost of these as **antecedent**.
3. If unsuccessful, repeat step 2 recursively for each child of N1 – Breadth-first search
4. Go up the tree to the lowest NP/S containing N1; call it N2.

Hobbs' Algorithm, continued

5. If N2 is an NP which is not in c-command, propose it as the **antecedent**.
6. Otherwise, apply steps 2-3 to N2.
7. If no antecedent NP is found, continue to apply steps 4 and 5 and then steps 2-3 to progressively higher NP/S nodes.
8. If no antecedent found at highest S of sentence, find the highest S node of the immediately preceding sentence and apply steps 2-3.
9. If still no antecedent found after *n* sentences, search for cataphora in current sentence from left-to-right, starting with first NP to the right of pronoun.

Hobbs: An Example

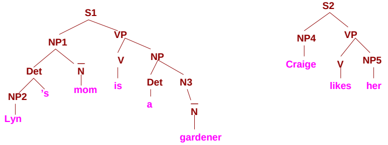

- Start search at NP5 in S2.
- Reject NP4 – c-commands NP5
- Move to S1. NP1 is first NP we encounter, so finish.
- Result: *Lyn's mom*
- What would have happened if the subject of S2 was *Craige's mom*?

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
Non-adverbial emphasis	Inside the house, **the cat** looked on	50
Head Noun emphasis	**The cat** in the house looked on	80

Non-adverbial emphasis penalises nouns in adverbial phrases.
Head-noun emphasis penalises NPs contained in other NPs.

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

The Camelot Example

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

Disc. Referents	Salience		
castle	cur_sent + subj + non-PP + head	100+80+50+80	310
Camelot	cur_sent + subj	100+80	180
residence	cur_sent + dir obj + non-PP + head	100+50+50+80	280
king	cur_sent + non-PP	100+50	150
536	cur_sent + indir obj + head	100+40+80	220

A Longer 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*$\} = 105$
    (subj + head + non-PP 80 + 80 + 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 Longer 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*$\} = 405$
    (subj+ head + non-PP 80 + 80 + 50)/2 + dir obj + head + non-PP + recency 70 + 80+ 50 + 100
  - $S_1 = \{$*Stephen Moss*$\} = 310$
    subj + head + non-PP + recency 80 + 80 + 50 + 100
- New Discourse Referents
  - Add *he* to N_1; $N_2 = \{$*Niall Ferguson, him, he*$\}$

Ge et al.'s 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.

Referring Expressions
Pronoun resolution algorithms
Centering (Grosz et al. 1995)
Hobbs
Lappin and Leass
Ge et al.

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%

Centering Theory (Grosz et al. 1995)

Motivation I: Centering provides a model for judging the coherence aspect of text quality.

Less Coherent Text

John went to his favourite music store to buy a piano. It was a store John had frequented for many years. He was excited that he could finally buy a piano. It was closing just as John arrived.

More Coherent Text

John went to his favourite music store to buy a piano. He had frequented the store for many years. He was excited that he could finally buy a piano. He arrived just as the store was closing for the day.

Referring Expressions
Pronoun resolution algorithms
Centering (Grosz et al. 1995)

Centering Theory (Grosz et al. 1995)

Motivation II: It can also be used for pronoun resolution, by predicting which references would be hard to process by a human.

A bad example

Tony was furious at being woken up so early. He told Terry$_i$ to get lost and hung up. Of course, he$_i$ hadn't intended to upset Tony.

- We want to predict that the use of *he* is inappropriate for referring to *Terry*.

Centering Theory

- A model of the local aspects of attentional state
  - tracks changes in local focus
  - does not provide an account of entities that are globally relevant throughout the discourse.
- The term **center** is used for an entity that links an utterance to other utterances in the same discourse segment
- Hence, the centers introduced by an utterance are also influenced by the surrounding context, not just by the utterance in isolation.

Centering

Every utterance U in a discourse introduces

- a set of forward-looking centers $C_f(U)$ (contains all the discourse entities evoked by the utterance U)
  - $C_f(U)$ is ordered according to the prominence of its member entities in the utterance U.
  - Ordering principle: grammatical function (subjects>objects > everything else).
- exactly one backward-looking center $C_b(U)$.
  - $C_b(U_n)$ of an utterance U_n is defined as the entity with the highest rank in $C_f(U_{n-1})$ that is evoked in U_n.
  - The backward-looking center $C_b(U_n)$ thus serves as a link with the preceding utterance U_{n-1}.

Centering: A model of discourse

- The forward-looking centers $C_f(U_{n-1})$ are a rough model of the listener's attentional state after U_{n-1}
- They predict what the backward-looking center of the utterance U_n will be; in particular, $C_b(U_n) = C_{f,top}(U_{n-1})$
- Need to perform pronoun resolution as you go along, in order to build forward-looking centers (use the same model)
- Abrupt changes in the focus of the discourse are reflected in changes in the backward-looking center.
- Discourse is then modelled by the types of transitions in the backward-looking centers from sentence to sentence.
- A discourse that keeps its center is most coherent, but if changes in topic occur, they should be transitioned smoothly

Four Types of Transitions

Two contributing factors:

- Did C_b change from U_{n-1} to U_n? ([Undefined-to-any-C_b] counts as "no change")
- Was $C_{f,top}$ correctly predicted by C_b?

	Same C_b	**Change in C_b**
$C_{f,top}$ **predicted**	CONTINUE	SMOOTH SHIFT
$C_{f,top}$ **not predicted**	RETAIN	ROUGH SHIFT

CONTINUE: $C_b(U_n) = C_b(U_{n-1}) = C_{f,top}(U_n)$

U$_1$: John went to his favourite music store to buy a piano.

$C_b(U_1)$ = **Undefined;** $C_f(U_1)$ = {**John, store, piano**}

U$_2$: He had frequented the store for many years.

CONTINUE: $C_b(U_2)$ = **John;** $C_f(U_2)$ = {**John, store, years**}

U$_3$: He was excited that he could finally buy a piano.

CONTINUE$C_b(U_3)$ = **John;** $C_f(U_3)$ = {**John, piano**}

In center continuation, the discourse stays focused on the same entity.

Referring Expressions
Pronoun resolution algorithms
Centering (Grosz et al. 1995)

RETAIN: $C_b(U_n) = C_b(U_{n-1})$ but $C_b(U_n) \neq C_{f,top}(U_n)$

U$_1$: John went to his favourite music store to buy a piano.

$C_b(U_1)$ = **Undefined;** $C_f(U_1)$ = {**John, store, piano**}

U$_2$: He had frequented the store for many years.

CONTINUE: $C_b(U_2)$ = **John;** $C_f(U_2)$ = {**John, store, years**}

U$_3$: It was closing just as John arrived.

RETAIN:$C_b(U_3)$ = **John;** $C_f(U_3)$ = {**store, John**}

In center retaining, a connecting sentence which evokes the next focus of discourse. C_b is retained from U_{n-1} to U_n, but it is likely to change in U_{n+1}.

Referring Expressions
Pronoun resolution algorithms
Centering (Grosz et al. 1995)

Smooth Shift: $C_b(U_n) \neq C_b(U_{n-1})$ but $C_b(U_n) = C_{f,top}(U_n)$

U$_1$: John was excited that he could finally buy a piano.

$C_b(U_1)$ = **Undefined;** $C_f(U_1)$ = {**John, piano**}

U$_2$: He went to his favourite music store to buy it.

CONTINUE: $C_b(U_2)$ = **John;** $C_f(U_2)$ = {**John, store, piano**}

U$_3$: It was about to close for the day.

RETAIN: $C_b(U_3)$ = **John;** $C_f(U_3)$ = {**store, day**}

U$_4$: It was his favourite shop in the world.

S-SHIFT: $C_b(U_4)$ = **store;** $C_f(U_4)$ = {**store, John, world**}

Smooth shifts are predictable changes in focus.

Referring Expressions
Pronoun resolution algorithms
Centering (Grosz et al. 1995)

Rough Shift: $C_b(U_n) \neq C_b(U_{n-1}) \neq C_{f,top}(U_n)$

U$_1$: John had always liked going to this store.

$C_b(U_1)$ = **Undefined;** $C_f(U_1)$ = {**John, store**}

U$_2$: It had a wide selection of musical instruments.

RETAIN: $C_b(U_2)$ = **John;** $C_f(U_2)$ = {**store, instruments**}

U$_3$: Mary visited it just as he left.

R-SHIFT: $C_b(U_3)$ = **store;** $C_f(U_3)$ = {**Mary, store, John**}

Rough shifts are unpredictable changes in discourse focus.

Center-Realisation Rules

So far, all pronoun resolution was unambiguous. Now let's move to non-trivial pronoun resolution with this algorithm.
Centering theory postulates two rules that constrain center-realisation:

Rule 1

If any element in $C_f(U_{n-1})$ is realised by a pronoun in U_n, then the center $C_b(U_n)$ must also be realised by a pronoun.

Rule 2

Sequences of center continuation are considered less disruptive than sequences of retaining, which are in turn less disruptive than sequences of shifts (smooth being better than rough).

Centering Algorithm

Goal: Find the referent that causes the smoothest C_b transition according to Rule 2, without violating Rule 1 or any agreement or syntactic constraints.

1. Move through the discourse window from left to right. At each pronoun:
    1. Generate C_f combinations for each possible set of referent assignments; this will create C_bs (top-ranked).
    2. Filter by agreement and syntactic constraints and Rule 1.
    3. Rank remaining referent assignments using Rule 2, i.e., transition orderings

Pronoun Resolution

U$_1$: Tony was furious at being woken up so early.

$$C_b(U_1) = \textbf{Undefined;}\ C_f(U_1) = \{\textbf{Tony}\}$$

U$_2$: He told Terry$_i$ to get lost and hung up.

CONTINUE: $C_b(U_2) = \textbf{Tony;}\ C_f(U_2) = \{\textbf{Tony, Terry}\}$

U$_3$: *Of course, he$_i$ hadn't intended to upset Tony.

$$C_b(U_3) = \textbf{Tony;}\ C_f(U_3) = \{\textbf{Terry, Tony}\}$$

- As Terry is a member of $C_f(U_2)$ that is realised as a pronoun in U_3, Rule 1 says that *Tony*, being $C_b(U_3)$, must also be realised as a pronoun in U_3 (but it isn't).
- Rule 1 filters this interpretation out.

Pronoun Resolution

U$_1$: Brennan drives an Alfa Romeo.

$$C_b(U_1) = \textbf{Undefined;}\ C_f(U_1) = \{\textbf{Brennan, Alfa}\}$$

U$_2$: Friedman races her on Sundays.

RETAIN: $C_b(U_2) = \textbf{Brennan,}\ C_f(U_2) = \{\textbf{Friedman, Brennan}\}$

U$_3$: She often beats her.

$$C_b(U_3) = \textbf{Friedman}$$

- Case 1; She=Brennan, her=Friedman
    - $C_f(U_3)$={**Brennan, Friedman**} → ROUGH SHIFT
- Case 2; She=Friedman, her=Brennan
    - $C_f(U_3)'$={**Friedman, Brennan**} → SMOOTH SHIFT

Therefore: She=Friedman and her=Brennan

Looking at the coherence examples again

U_1: John went to his favourite music store to buy a piano.

$C_b(U_1)$ = **Undefined**; $C_f(U_1)$ = {**John, store, piano**}

U_2: It was a store John had frequented for many years.

RETAIN: $C_b(U_2)$ = **John**; $C_f(U_2)$ = {**store, John, years**}

U_3: He was excited that he could finally buy a piano.

CONTINUE: $C_b(U_3)$ = **John**; $C_f(U_3)$ = {**John, piano**}

U_4: It was closing just as John arrived.

RETAIN: $C_b(U_4)$ = **John**; $C_f(U_4)$ = {**store, John**}

U_5: It would open again tomorrow.

SMOOTH SHIFT: $C_b(U_5)$ = **store**; $C_f(U_5)$ = {**store**}

Looking at the other coherence example

U_1: John went to his favourite music store to buy a piano.

$C_b(U_1)$ = **Undefined**; $C_f(U_1)$ = {**John, store, piano**}

U_2: He had frequented the store for many years.

CONTINUE: $C_b(U_2)$ = **John**; $C_f(U_2)$ = {**John, store, years**}

U_3: He was excited that he could finally buy a piano.

CONTINUE: $C_b(U_3)$ = **John**; $C_f(U_3)$ = {**John, piano**}

U_4: He arrived just as the store was closing for the day.

CONTINUE: $C_b(U_4)$ = **John**; $C_f(U_4)$ = {**John, store, day**}

U_5: It would open again tomorrow.

RETAIN: $C_b(U_5)$ = **John**; $C_f(U_5)$ = {**store**}

Commonalities Centering vs. Lappin/Leass

- Both Lappin & Leass and Centering Approach
  - first identifying possible antecedents
  - then applying a set of filters to rule out some of them
  - and finally applying a decision procedure to select one of the remaining candidates
    - Centering uses Rule 2 (Continuation>Retain>Shift)
    - Lappin & Leass uses Salience Value
- Both algorithms
  - maintain a Discourse Model
  - differentiate between constraints (hard) and preferences (soft)

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
- ...and a Discourse Theory
  - Centering Theory

References

- Jurafsky and Martin, Chapter 21.4
- Ge, Hale, Charniak. 1998. A statistical approach to anaphora resolution. In Proceedings of the Sixth Workshop on Very Large Corpora, pages 161171, Montreal, Canada.
- Gundel, Hedberg, Zacharski. 1993. Cognitive Status and the Form of Referring Expressions in Discourse. Language, 69(2):274307.
- Krahmer, van Erk, Verleg. 2003. Graph-based generation of referring expressions. Computational Linguistics, 29(1):5372.
- Lappin and Leass. 1994. An algorithm for pronominal anaphora resolution. Computational Linguistics, 20(4):535561.
- Grosz, Joshi, Weinstein. 1995. Centering: A framework for modelling the local coherence of discourse. Computational Linguistics, 21(2):203226.