Discourse: Coreference

Deep Processing Techniques for NLP
Ling 571
March 2, 2016

Roadmap

- Coreference
  - Referring expressions

  - Syntactic & semantic constraints
  - Syntactic & semantic preferences

  - Reference resolution:
    - Hobbs Algorithm: Baseline
    - Machine learning approaches
    - Sieve models

  - Challenges

Reference

- Queen Elizabeth set about transforming her husband, King George VI, into a viable monarch. Logue, a renowned speech therapist, was summoned to help the King overcome his speech impediment...

Referring expression: (refexp)

Linguistic form that picks out entity in some model

Reference

- Queen Elizabeth set about transforming her husband, King George VI, into <u>a viable monarch</u>. Logue, a renowned speech therapist, was summoned to help the King overcome his <u>speech impediment</u>...

Referring expression: (refexp)

- Linguistic form that picks out entity in some model
- That entity is the "referent"
  - When introduces entity, "evokes" it
  - Set up later reference, "antecedent"
- 2 refexps with same referent "co-refer"

Reference (terminology)

- Queen Elizabeth set about transforming her husband, King George VI, into <u>a viable monarch</u>. Logue, a renowned speech therapist, was summoned to help the King overcome his <u>speech impediment</u>...

- Anaphor:
  - Abbreviated linguistic form interpreted in context
    - Her, his, the King
  - Refers to previously introduced item (“accesses”)
    - Referring expression is then anaphoric

Referring Expressions

- Many alternatives:
  - Queen Elizabeth, she, her, the Queen, etc
  - Possible correct forms depend on discourse context
    - E.g. she, her presume prior mention, or presence in world
- Interpretation (and generation) requires:
  - Discourse Model with representations of:
    - Entities referred to in the discourse
    - Relationships of these entities
  - Need way to construct, update model
  - Need way to map refexp to hearer's beliefs

Reference and Model

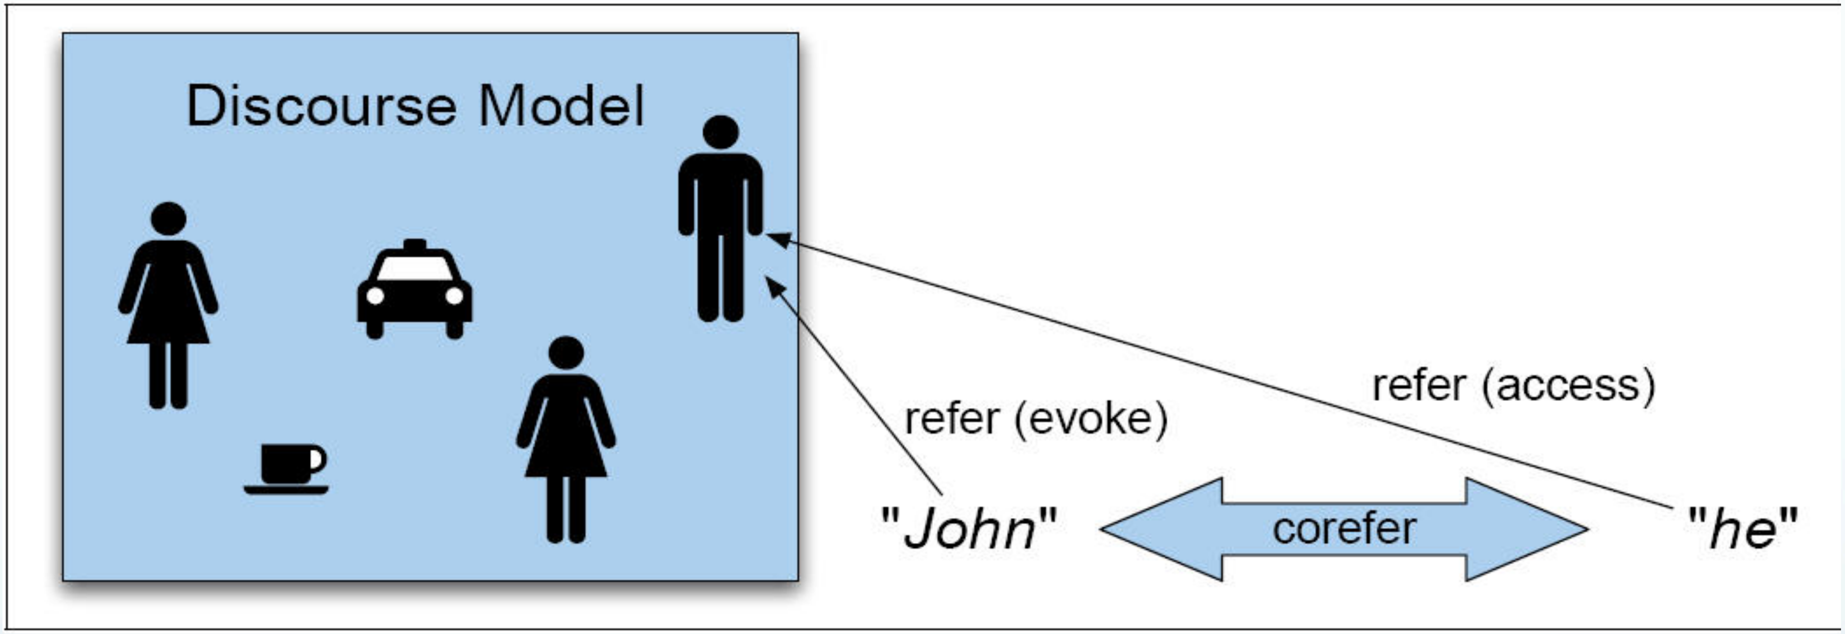

Reference Resolution

- Queen Elizabeth set about transforming her husband, King George VI, into a viable monarch. Logue, a renowned speech therapist, was summoned to help the King overcome his speech impediment...

Coreference resolution:

Find all expressions referring to same entity, 'corefer'

Colors indicate coreferent sets

Pronominal anaphora resolution:

Find antecedent for given pronoun

Referring Expressions

- Indefinite noun phrases (NPs): e.g. “a cat”
  - Introduces new item to discourse context
- Definite NPs: e.g. “the cat”
  - Refers to item identifiable by hearer in context
    - By verbal, pointing, or environment availability; implicit
- Pronouns: e.g. “he”,”she”, “it”
  - Refers to item, must be “salient”
- Demonstratives: e.g. “this”, “that”
  - Refers to item, sense of distance (literal/figurative)
- Names: e.g. “Miss Woodhouse”,”IBM”
  - New or old entities

Information Status

- Some expressions (e.g. indef NPs) introduce **new** info
- Others refer to old referents (e.g. pronouns)

- Theories link form of refexp to given/new status

The givenness hierarchy:

in focus >	activated >	familiar >	uniquely identifiable >	referential >	type identifiable
{it}	{*that*, *this*, *this* N}	{that N}	{the N}	{indef. *this* N}	{*a* N}

- Accessibility:
  - More salient elements easier to call up, can be shorter
    Correlates with length: more accessible, shorter refexp

Complicating Factors

- Inferrables:
  - Refexp refers to inferentially related entity
    - *I bought a car today, but the door had a dent, and the engine was noisy.*
    - E.g. car → door, engine
- Generics:
  - *I want to buy a Mac. They are very stylish.*
  - General group evoked by instance.
- Non-referential cases:
  - *It's raining.*

Syntactic Constraints for Reference Resolution

- Some fairly rigid rules constrain possible referents
- Agreement:
  - Number: Singular/Plural
  - Person: 1st: I,we; 2nd: you; 3rd: he, she, it, they
  - Gender: he vs she vs it

Syntactic & Semantic Constraints

- Binding constraints:
  - Reflexive (x-self): corefers with subject of clause
  - Pronoun/Def. NP: can't corefer with subject of clause
- "Selectional restrictions":
  - "animate": The cows eat grass.
  - "human": The author wrote the book.
  - More general: drive: John drives a car....

Syntactic & Semantic Preferences

- Recency: Closer entities are more salient
  - The doctor found an old map in the chest. Jim found an even older map on the shelf. It described an island.

- Grammatical role: Saliency hierarchy of roles
  - e.g. Subj > Object > I. Obj. > Oblique > AdvP
    - Billy Bones went to the bar with Jim Hawkins. He called for a glass of rum. [he = Billy]
    - Jim Hawkins went to the bar with Billy Bones. He called for a glass of rum. [he = Jim]

Syntactic & Semantic Preferences

- Repeated reference: Pronouns more salient
  - Once focused, likely to continue to be focused
    - Billy Bones had been thinking of a glass of rum. He hobbled over to the bar. Jim Hawkins went with him. He called for a glass of rum. [he=Billy]
- Parallelism: Prefer entity in same role
    - Silver went with Jim to the bar. Billy Bones went with him to the inn. [him = Jim]
    - Overrides grammatical role
- Verb roles: "implicit causality", thematic role match,...
    - John telephoned Bill. He lost the laptop. [He=John]
    - John criticized Bill. He lost the laptop. [He=Bill]

Reference Resolution Approaches

- Common features
  - "Discourse Model"
    - Referents evoked in discourse, available for reference
    - Structure indicating relative salience
  - Syntactic & Semantic Constraints
  - Syntactic & Semantic Preferences
- Differences:
  - Which constraints/preferences? How combine? Rank?

Hobbs' Resolution Algorithm

- Requires:
  - Syntactic parser
  - Gender and number checker
- Input:
  - Pronoun
  - Parse of current and previous sentences
- Captures:
  - Preferences: Recency, grammatical role
  - Constraints: binding theory, gender, person, number

Hobbs Algorithm

- Intuition:
  - Start with target pronoun
  - Climb parse tree to S root
  - For each NP or S
    - Do breadth-first, left-to-right search of children
      - Restricted to left of target
    - For each NP, check agreement with target
  - Repeat on earlier sentences until matching NP found

Hobbs Algorithm Detail

- Begin at NP immediately dominating pronoun
- Climb tree to NP or S: X=node, p = path
- Traverse branches below X, and left of p: BF, LR
  - If find NP, propose as antecedent
    - If separated from X by NP or S
- Loop: If X highest S in sentence, try previous sentences.
- If X not highest S, climb to next NP or S: X = node
- If X is NP, and p not through X's nominal, propose X
- Traverse branches below X, left of p: BF,LR
  - Propose any NP
- If X is S, traverse branches of X, right of p: BF, LR
    - Do not traverse NP or S; Propose any NP
  - Go to Loop

Hobbs Example

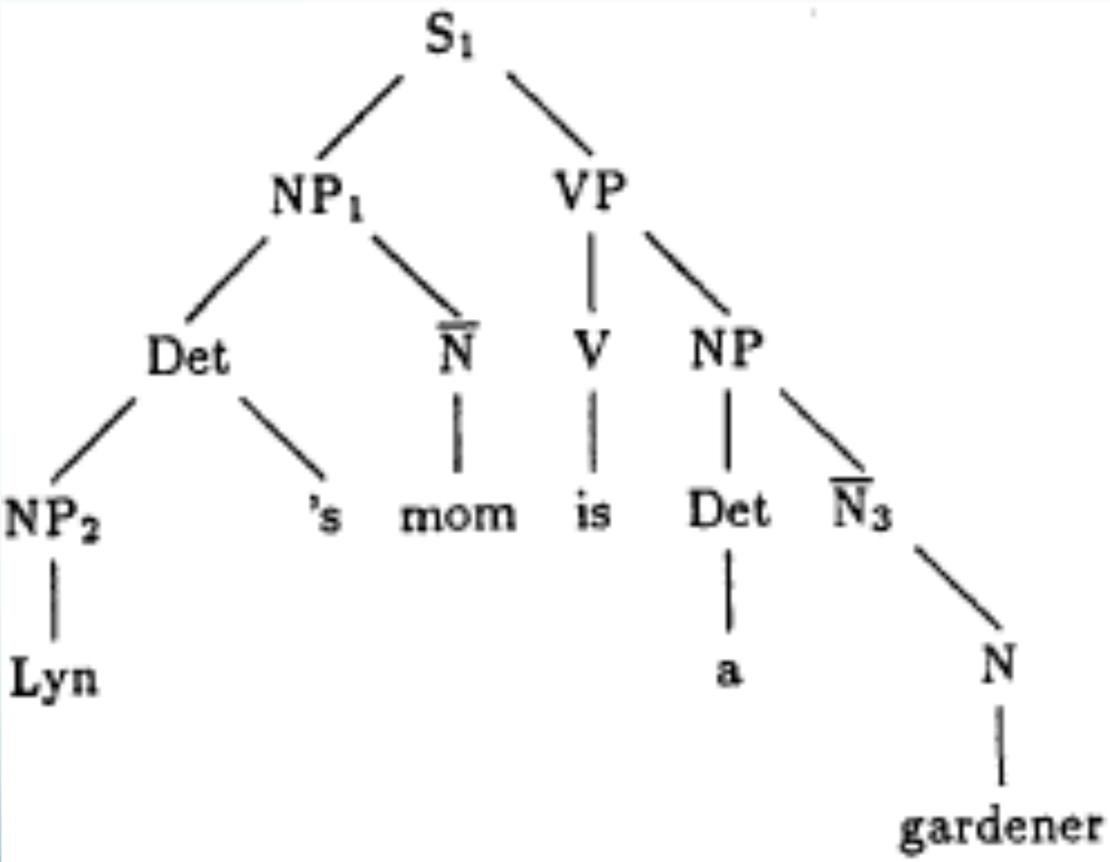

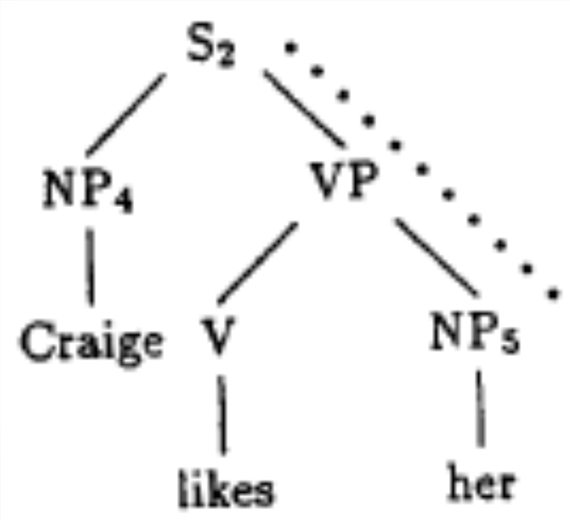

Lyn's mom is a gardener. Craige likes her.

Another Hobbs Example

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

- What is **it**?
  - residence

Another Hobbs Example

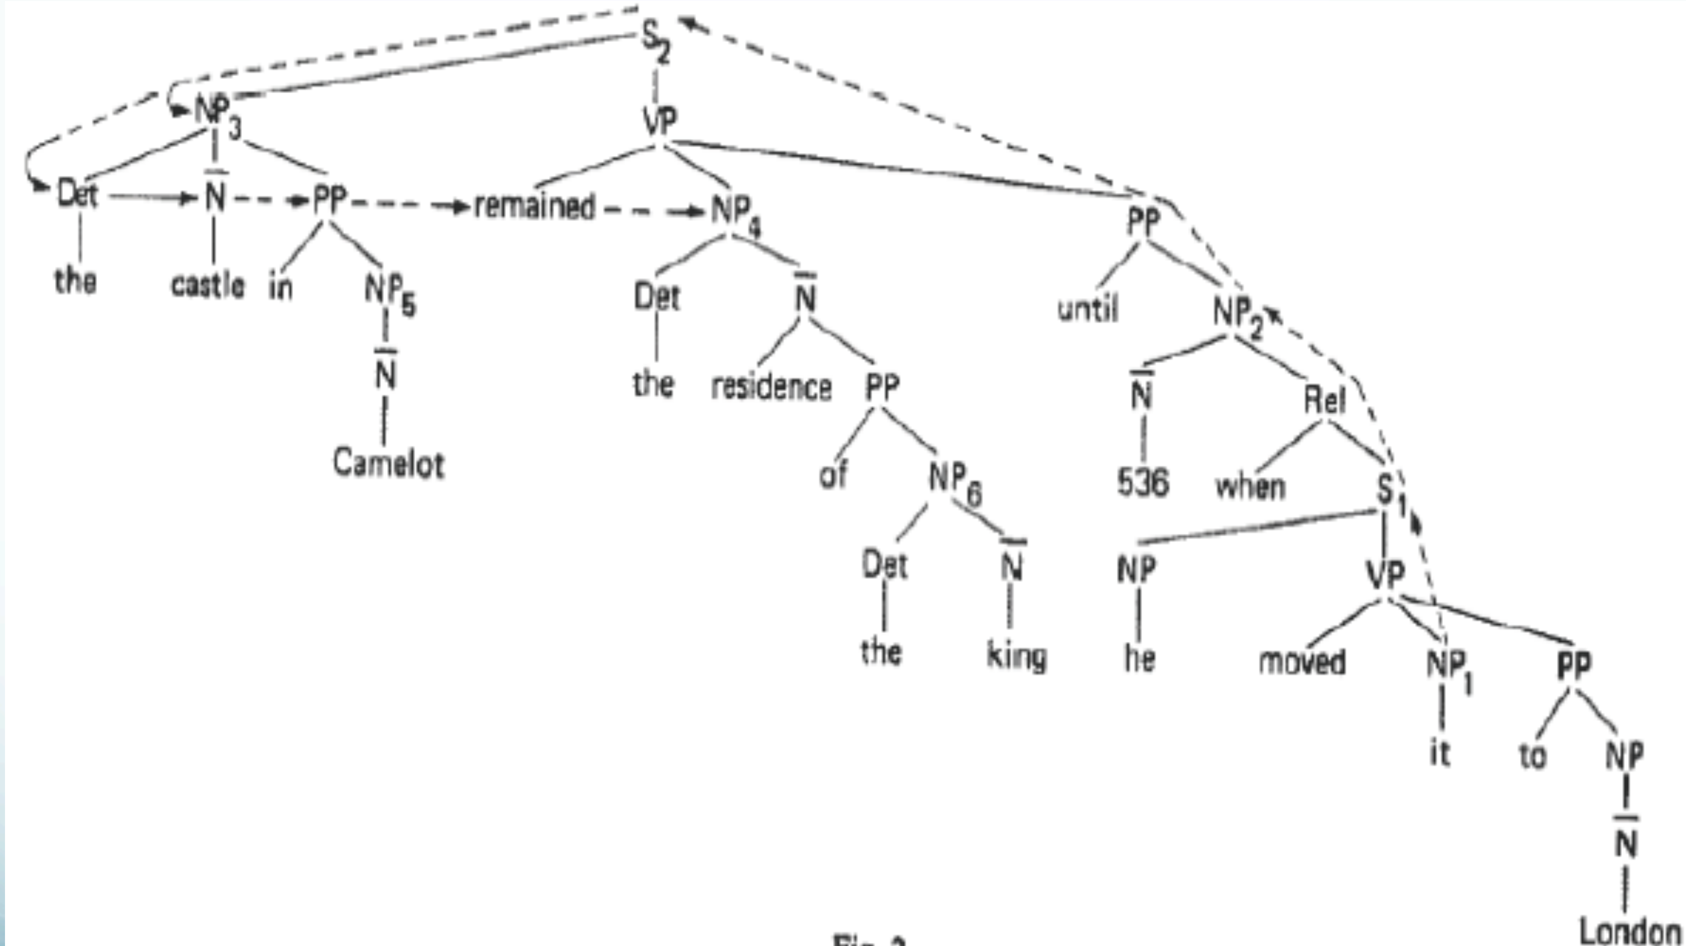

Hobbs, 1978

Hobbs Algorithm

- Results: 88% accuracy ; 90+% intrasentential
  - On perfect, manually parsed sentences
- Useful baseline for evaluating pronominal anaphora
- Issues:
  - Parsing:
    - Not all languages have parsers
    - Parsers are not always accurate
  - Constraints/Preferences:
    - Captures: Binding theory, grammatical role, recency
    - But not: parallelism, repetition, verb semantics, selection

Data-driven Reference Resolution

- Prior approaches: Knowledge-based, hand-crafted
- Data-driven machine learning approach
  - Coreference as classification, clustering, ranking problem
    - Mention-pair model:
      - For each pair NPi,NPj, do they corefer?
      - Cluster to form equivalence classes
    - Entity-mention model
      - For each pair NP_k and cluster $C_{j,,}$ should the NP be in the cluster?
    - Ranking models
      - For each NP_k, and all candidate antecedents, which highest?

NP Coreference Examples

- Link all NPs refer to same entity

Queen Elizabeth set about transforming her husband,
King George VI, into a viable monarch. Logue,
a renowned speech therapist, was summoned to help
the King overcome his speech impediment...

Example from Cardie&Ng 2004

Annotated Corpora

- Available shared task corpora
  - MUC-6, MUC-7 (Message Understanding Conference)
    - 60 documents each, newswire, English
  - ACE (Automatic Content Extraction)
    - Originally English newswite
    - Later include Chinese, Arabic; blog, CTS, usenet, etc
- Treebanks
  - English Penn Treebank (Ontonotes)
  - German, Czech, Japanese, Spanish, Catalan, Medline

Feature Engineering

- Other coreference (not pronominal) features
  - String-matching features:
    - Mrs. Clinton <->Clinton
  - Semantic features:
    - Can candidate appear in same role w/same verb?
    - WordNet similarity
    - Wikipedia: broader coverage
  - Lexico-syntactic patterns:
    - E.g. X is a Y

Typical Feature Set

- 25 features per instance: 2NPs, features, class
  - lexical (3)
    - string matching for pronouns, proper names, common nouns
  - grammatical (18)
    - pronoun_1, pronoun_2, demonstrative_2, indefinite_2, ...
    - number, gender, animacy
    - appositive, predicate nominative
    - binding constraints, simple contra-indexing constraints, ...
    - span, maximalnp, ...
  - semantic (2)
    - same WordNet class
    - alias
  - positional (1)
    - distance between the NPs in terms of # of sentences
  - knowledge-based (1)
    - naïve pronoun resolution algorithm

Coreference Evaluation

- Key issues:
  - Which NPs are evaluated?
    - Gold standard tagged or
    - Automatically extracted
  - How good is the partition?
    - Any cluster-based evaluation could be used (e.g. Kappa)
    - MUC scorer:
      - Link-based: ignores singletons; penalizes large clusters
      - Other measures compensate

Clustering by Classification

- Mention-pair style system:

  - For each pair of NPs, classify +/- coreferent
    - Any classifier

  - Linked pairs form coreferential chains
    - Process candidate pairs from End to Start
    - All mentions of an entity appear in single chain

  - F-measure: MUC-6: 62-66%; MUC-7: 60-61%
    - Soon et. al, Cardie and Ng (2002)

Multi-pass Sieve Approach

  - Raghunathan et al., 2010

- Key Issues:
  - Limitations of mention-pair classifier approach
    - Local decisions over large number of features
      - Not really transitive
      - Can't exploit global constraints
      - Low precision features may overwhelm less frequent, high precision ones