

Referring Expressions

A mention that is about some entity, the *referent*

- ▶ Pronoun
- ▶ Name: Toby
- ▶ Definite description: Samuel's cat
- ▶ Indefinite description: A cat
- ▶ Mentions can be nested: *her son's manager's husband* has
 - ▶ *Her*
 - ▶ *Her son*
 - ▶ *Her son's manager*
 - ▶ *Her son's manager's husband*

Focus on discourse entities \neq real-world entities

Sample of Theories of Reference

About definite descriptions, e.g., with *the*

The man with the wine glass

- ▶ Bertrand Russell
 - ▶ The P
 - ▶ $\iota x : P(x)$ means the unique x such that $P(x)$ holds
 - ▶ Undefined if zero or two or more
- ▶ Keith Donnellan
 - ▶ Suppose the glass contains grape juice
 - ▶ The speaker still meant a specific person
 - ▶ You would still understand whom they meant
- ▶ John Perry
 - ▶ “Essential” Indexical
 - ▶ I vs. a description that refers to me

The man with the wine glass is leaving stains on the new carpet

Coreference

When two expressions have the same referent

- ▶ Coreference is crucial for understanding natural language
 - ▶ Within a sentence
 - ▶ Across sentences by the same speaker or writer, as in a *discourse*
 - ▶ Across sentences by different parties, as in a *dialog*

Anaphora

- ▶ A referent being evoked
 - ▶ First mention of a referent
 - ▶ Natural with indefinite descriptions
 - ▶ Singleton: Referent with single mention
- ▶ A referent being accessed
 - ▶ Subsequent mention of a referent
- ▶ Anaphora
 - ▶ Reference to a referent that has been already introduced into the discourse
 - ▶ Not just pronouns but also proper names (when repeated)
 - ▶ Not just NPs but also VPs—virtually any construct
- ▶ Cataphora: from previous referent to subsequent reference
 - ▶ Works only for pronouns
- ▶ Entity linking: Identify referent in the real world or in an ontology

Anaphora Examples

I saw a man with a wine glass. He was drunk.

I saw a man with a wine glass. Both he and it were foggy.

He is 75 years old but the man behaves like an unruly teenager.

I was to give my friend a ride but my car didn't start so I canceled it.

I was to give my friend a ride but my car didn't start so I canceled it.

I was to take my friend to an appointment but my car didn't start so I canceled it.

I was to take my friend to her final but my car didn't start so I canceled it.

My car didn't start because it was faulty but my friend doesn't believe it.

I spent hours trying to repair my car. It was a tedious job.

Types of Referring Expressions

- ▶ Indefinite noun phrases
 - ▶ Indefinite article
 - ▶ Quantifiers: some, all
 - ▶ Generalized quantifiers: three of seven
 - ▶ Demonstratives: this [unusual reading]:
 - ▶ I came across this struggling actor who works as a barista
- ▶ Definite noun phrases
 - ▶ Definite article
 - ▶ Known and identifiable to the reader or listener
 - ▶ Demonstratives: this, that [common reading]
- ▶ Pronouns in quantified expressions
 - ▶ In *Every mother remembers her child's birthday*
 - ▶ There is no direct referent for her since it is bound within the scope of the *every* over mothers

Indefinite and Definite Noun Phrases

19th century grammar terms: inaccurate but established

▶ Indefinite NPs

- ▶ Primarily introduce a referent into the discourse
- ▶ Don't need to be indefinite: Some are *specific*

I've been through the desert on a horse

I've been through the desert on a horse with no name

I met a man by the name of John Doe

▶ Definite NPs

- ▶ Notionally anaphoric: refer to some referent introduced into the discourse by an indefinite NPs
- ▶ Depending on corpus, often ($\leq 50\%$ for newswire) not anaphoric by fact of being clear

I went to a restaurant. The waiter brought me the menu.

Zero Anaphora

- ▶ Prominent in several languages
 - ▶ Chinese
 - ▶ Italian
- ▶ More apparent in spoken dialog or casual discourse

So the boss calls me in

☐ Says I'm not pullin' my weight

- ▶ Beatles: *A Day in the Life*

Woke up, fell out of bed

Dragged a comb across my head

- ▶ Kenny Rogers: *The Gambler*

You got to know

When to hold 'em

Know when to fold 'em

Know when to walk away

Know when to run

Information Processing View of Discourse

A cognitive model of discourse processing

- ▶ Some NPs introduce entities into the discourse
 - ▶ New to the discourse and new to the hearer (or reader)

I saw a man enter the building

- ▶ New to the discourse but old to the hearer (or reader)

I saw Samuel enter the building

- ▶ Some NPs evoke entities already in the discourse

- ▶ Old to the discourse and old to the hearer (or reader)

I saw a man enter the building.

He was carrying a package.

Salience and Accessibility

- ▶ Present in the hearer's mind
- ▶ Or easy to recall
- ▶ Therefore, requires less linguistic material to refer to
- ▶ Some NPs evoke entities that are readily inferred
 - ▶ New to the discourse and new to the hearer (or reader), but definite

I went to the restaurant. The waiter brought me the menu.

I went to the restaurant. They brought me the menu.

- ▶ Rely upon the applicable frame being selected

Non-Referring Expressions: Noun Phrases

- ▶ Blocked by negation (Karttunen)

Janet doesn't have a car

*It's a Toyota

- ▶ Blocked by nonfactive verbs (Asher)

I doubt Janet has a car

*It's a Toyota

- ▶ Appositives don't refer but provide parenthetical information

United, a unit of UAL, matched the fares

- ▶ But worth linking appositives to the main NP for understanding

- ▶ Predicative: properties of the head noun, not a separate entity

NC State is a university in Raleigh

- ▶ Attributive: also properties

NC State was established as a land-grant institution

Non-Referring Expressions: Exercise

Give examples of such expressions

Non-Referring Expressions: Expletive Pronouns

- ▶ Expletives or pleonastic

It's cold in here

- ▶ Clefts

It was Xerox who invented the mouse-based UI

- ▶ Extraposition

It surprised no one that Russia invaded Crimea

Non-Referring Expressions: Generics

- ▶ Generic nouns: refer to a type rather than an individual or individuals

The lion is the king of the jungle

But he scavenges food more than the lowly hyena

- ▶ Generic: you (Kenny Rogers: *The Gambler*)

You got to know

When to hold 'em

...

You never count your money

When you're sittin' at the table

- ▶ Habitual verb phrases similarly capture types of events

You never count your money

When you're sittin' at the table

Constraints on Coreference: 1

- ▶ Number agreement
 - ▶ Singular: you/she/her/he/him/his/it
 - ▶ Plural: you/we/us/they/them/their
 - ▶ How would you classify y'all? And, youse?
 - ▶ What about they/them/their used for an individual?

- ▶ Noteworthy: singular they

- ▶ Shakespeare's *Comedy of Errors*, circa 1594

There's not a man I meet but doth salute me
As if I were their well-acquainted friend

- ▶ P. G. Wodehouse's *The Inimitable Jeeves*, circa 1923

Personally, if anyone had told me a tie like that suited me, I should have risen and struck them on the mazzard, regardless of their age or sex; . . .

Constraints on Coreference: 2

- ▶ Person agreement
 - ▶ First
 - ▶ Second
 - ▶ Third
- ▶ Gender (and personhood)
 - ▶ Male
 - ▶ Female
 - ▶ Nonpersonal

Constraints on Coreference: 3

- ▶ Binding theory: how mentions relate to an antecedent in the same sentence
 - ▶ Reflexives: himself, herself, themselves
 - ▶ Consider coreference with the subject of the most immediate containing clause of a pronoun

- ▶ Reflexives must: herself = Sanjana

Sanjana bought herself a new lease on life

- ▶ Nonreflexives must not: her \neq Sanjana

Sanjana bought her a new lease on life

- ▶ Recency: prefer more recent utterance or nearer preceding sentence

Grammatical Role and Complications

- ▶ The subject of a sentence is more preferred as an antecedent than its object

Meenakshi worked on a project with Maya. She prepared their joint presentation

- ▶ But gender agreement matters more

Meenakshi worked on a project with Luke. He prepared their joint presentation

- ▶ As do constructs that block singular reference

*Meenakshi and Maya worked together on a project. She prepared their joint presentation

- ▶ Leap frogging?

?Meenakshi interned at IBM. Meenakshi and Maya worked together on a project. She prepared their joint presentation

Verb Semantics

- ▶ Influence of the deep meaning of a verb with respect to salience and causality

John telephoned Bill

He lost the laptop

John criticized Bill

He lost the laptop

Selectional Restrictions

- ▶ It's more natural to cook soup than to cook a bowl

I ate the soup in my new bowl after cooking it for hours

*I ate the soup in my new bowl after cooking it for hours

- ▶ Jurafsky's explanation focuses on *ate* but it seems to me that *cooking* is the verb of interest
- ▶ But if you are into pottery and your interlocutor knows it, both work, especially the second:

I ate the soup in my new bowl after spending hours preparing it

I ate the soup in my new bowl after spending hours preparing it

NP Content

Emmon Bach *Problominalization*, famous 1.5 page article from 1970

- ▶ The full NP matters, including relative clauses [commas added]

My neighbor, who is pregnant, said that she was very happy

*My neighbor, who is pregnant, said that he was very happy

- ▶ Bach takes the above as evidence that the full NP is relevant, so we might show it as

My neighbor, who is pregnant, said that she was very happy

Pronominalization

- ▶ Naïve thinking: we can substitute an NP for the pronoun
- ▶ Consider a so-called Bach-Peters sentence

The man_i who shows he_i deserves it_j will get the prize_j he_i desires

- ▶ This sentence has no finite resolution

The man who shows that the man deserves the prize that the man who shows that the man deserves the prize that the man ... (ad infinitum) will get the prize that the man who shows that the man deserves the prize that the man who shows ... (ad infinitum)

- ▶ Additional examples:

I gave the book_b that he_m wanted to the man_m who asked for it_b

The pilot_i who shot at it_j hit the MiG_j that chased him_i

Karttunen's Examples: 1

- ▶ Specific reading; can refer (Bean Blossom is a town in Indiana)

The director is looking at an innocent blonde
She is from Bean Blossom

- ▶ Nonspecific reading: can't refer

*The director is looking for an innocent blonde
She is from Bean Blossom

- ▶ Nonspecific reading: can refer in a modal or hypothetical context

The director is looking for an innocent blonde
She must be 17 years old

- ▶ We interpret her age as a requirement on such a person

- ▶ Specific reading: works in a modal (epistemic) or hypothetical context

The director is looking at an innocent blonde
She must be 17 years old

- ▶ We interpret her age as a fact about her

Karttunen's Examples: 2

- ▶ This interpretation is correct

I gave each student a cookie

Some of them ate it right away

- ▶ But there is no unique cookie being referred to

Coreference Task: Identify Coreference *Clusters* or *Chains*

Focused on pronominal anaphora

- ▶ Superscript identifies chain
- ▶ Subscript identifies mention

[Victoria Chen]_a¹, CFO of
 [Megabucks Banking]_a², saw
 [[her]_b¹ pay]_a³ jump to \$2.3 million, as
 [the 38-year-old]_c¹ also became
 [the company]_b²'s president.

It is widely known that
 [she]_d¹ came to
 [Megabucks]_c² from rival
 [Lotsabucks]_a⁴ .

Clusters indicated by labeling:

- 1 {Victoria Chen, her, the 38-year-old, She}
- 2 {Megabucks Banking, the company, Megabucks}
- 3 {her pay}
- 4 {Lotsabucks}

Notice the pleonastic **It**
 Cleft or extraposition?

Example Dataset: OntoNotes

Associated with the CoNLL 2012 Shared Task

ACL SIGNLL Conference on Computational Natural Language Learning

- ▶ Chinese and English, 1 million words each from newswire, magazine articles, broadcast news, broadcast conversations, web data and conversational speech
- ▶ Arabic, 300,000 words from newswire sources
- ▶ Includes coreferring NPs as mentions
- ▶ Includes appositive clauses within a mention
- ▶ Doesn't label singletons \Rightarrow simplifies the task by removing a confounder
- ▶ Doesn't label generics and pleonastic pronouns
- ▶ Labels pronominal modifiers only when they are proper nouns
 - ▶ Not *wheat* in *wheat fields*
 - ▶ No *American* in *American policy*
 - ▶ But *UN* in *UN policy*

Challenges

- ▶ Separating anaphoric from nonreferential (expletive) pronouns
- ▶ Confounding due to singleton NPs, of which there are many (typically, 60%–70%)

Mention Detection

- ▶ Spans of text corresponding to mentions
- ▶ Current techniques generally err on the side of recall
- ▶ May involve parsing and named entity recognition
 - ▶ Any NP
 - ▶ Any possessive pronoun
 - ▶ A named entity
- ▶ Newer techniques go further by extracting all n-grams
 - ▶ For $1 \leq n \leq 10$
 - ▶ Most such are not NPs
- ▶ Filtering out the useless hits is this crucial

Rule-Based Filtering

- ▶ Early rule-based approaches for pleonastic *it*
 - ▶ Lists of cognitive verbs: *believe*
 - ▶ Lists of modal adjectives: *necessary, certain*

It is ⟨Modal Adjective⟩ that S

It is ⟨Modal Adjective⟩ (for NP) to VP


It is ⟨Cognitive Verb⟩-ed that S

It seems/appears/means/follows (that) S

- ▶ Supplement rules with classifiers for three subtasks:
 - ▶ Mentions (referentiality—is a referent)
 - ▶ Anaphoricity: is an anaphor
 - ▶ Discourse-new: new mention that may later (in the text) be pointed to by an anaphor
- ▶ Piecemeal is not effective: Modern approaches combine the classifiers into a single model

Methods Customized to Evaluation

Sounds flaky from the scientific standpoint

- ▶ Some ideas developed for a dataset (and task) may generalize
- ▶ A specific process for OntoNotes
 - ▶ Take all NPs, possessive pronouns, and named entities
 - ▶ Remove
 - ▶ Numeric quantities, e.g., *100 dollars*, *8%*—rarely coreferential
 - ▶ Mentions embedded in larger mentions, e.g., $[[her]_b \text{ pay}]_a^3$
 - ▶ Adjectival forms of nations, e.g., *Canadian*
 - ▶ Stop words, e.g., *there*
 - ▶ Regular expressions to identify (and remove) pleonastic 

Selecting Anaphoric and Referential Mentions: Example

Victoria Chen, CFO of Megabucks Banking, saw her pay jump to \$2.3 million, as the 38-year-old also became the company's president. It is widely known that she came to Megabucks from rival Lotsabucks.

Victoria Chen

the company

CFO of Megabucks Banking

the company's president

Megabucks Banking

It

her

she

her pay

Megabucks

\$2.3 million

Lotsabucks

the 38-year-old

Appositive

Predicate nominal

Pleonastic

Embedded in larger mention

Numeric

Anaphoricity Classification

- ▶ Labeled examples
 - ▶ Positive: Any span labeled as an anaphor
 - ▶ Negative: Any span that is not labeled as an anaphor
- ▶ Features (run into the dozens)
 - ▶ Head word
 - ▶ Context words
 - ▶ Definiteness
 - ▶ Length
 - ▶ Position in discourse
 - ▶ Animacy: volitional doer of an action

Pleonastic Pronouns and Nonreferring NPs

Common error: when they are deemed coreferents of something

- ▶ Detecting nonreferring NPs
 - ▶ Generalize over first-occurring NPs, which can't be anaphoric
 - ▶ Frequently occurring head nouns that are never labeled referring
- ▶ Mining web data for identifying anaphoric pronouns

Anaphoric	You can make <u>it</u> in advance
Nonanaphoric	You can make <u>it</u> in Hollywood

- ▶ Anaphoric: ordinary expression
 - ▶ Lots of hits and variety on “Make _____ in advance”
 - ▶ Make pasta in advance
 - ▶ Make them in advance
- ▶ Nonanaphoric: idiomatic expression
 - ▶ Few hits and limited variety on “Make _____ in Hollywood”
 - ▶ * Make pasta in Hollywood
 - ▶ * Make them in Hollywood

Pleonastic Pronouns and Nonreferring NPs

Mining web data for identifying anaphoric pronouns

Anaphoric	You can make <u>it</u> in advance
Nonanaphoric	You can make <u>it</u> in Hollywood

- ▶ Anaphoric: ordinary expression, e.g., “It’s been a problem”
 - ▶ Lots of hits and variety on “Make _____ in advance”
 - ▶ Make pasta in advance
 - ▶ Make them in advance
- ▶ Nonanaphoric: idiomatic expression, e.g., “It’s been ages”
 - ▶ Few hits and limited variety on “Make _____ in Hollywood”
 - ▶ * Make pasta in Hollywood
 - ▶ * Make them in Hollywood
- ▶ Some words are not discriminatory, e.g., *money*

Mention-Pair Task

Considers mentions, not the underlying entities

- ▶ Computes probability of coreference for two mentions: a candidate antecedent and a candidate anaphor
- ▶ Heuristic to create dataset: for each anaphor mention, m_i
 - ▶ One positive instance
 - ▶ (m_i, m_j) where m_j is the closest (correct) antecedent of m_i
 - ▶ she \Rightarrow the 38-year-old
 - ▶ the company \Rightarrow Megabucks Banking
 - ▶ Several negative instances
 - ▶ (m_i, m_k) where m_k occurs between m_i and m_j
 - ▶ the company \Rightarrow the 38-year-old
 - ▶ the company \Rightarrow her pay
- ▶ Closest-first clustering: proceed right to left
 - ▶ Link to first antecedent with probability of coreference > 0.5
- ▶ Best-first clustering: evaluate globally in the discourse
 - ▶ Link to antecedent with highest probability of coreference, if > 0.5

Mention-Rank Task

- ▶ For the i th mention treated as an anaphor
 - ▶ Random variable, $y_i \in \{1, \dots, i-1, \varepsilon\}$ points to its antecedent
 - ▶ Here ε indicates no antecedent
- ▶ Training is nontrivial
- ▶ Heuristics such as
 - ▶ Positive: closest antecedent
 - ▶ Negative: all mentions within two sentences that are not antecedents

Entity-Based Task: Relevant Features

Link to previous discourse entity, i.e., a cluster of mentions

- ▶ Size of a cluster
- ▶ Shape of cluster, indicating sequence of types of the mentions in it
 - ▶ Proper Noun (P)
 - ▶ Definite NP (D)
 - ▶ Indefinite NP (I)
 - ▶ Pronoun (Pr)
 - ▶ Example sequence (in order of occurrence in the text)
 - ▶ ⟨Victoria Chen, her, the 38-year-old⟩
- ▶ Mean of mention-anaphor probability for each pair drawn from two clusters
 - ▶ Indicates closeness of the clusters as a basis for combining them

Clustering has not proved competitive and mentioned-ranking methods are more prevalent

Sample Features

Drawn from long list in the book

Consider an anaphor *she* and an antecedent *Victoria Chen*

- ▶ Attributes of antecedent and of anaphor (in that order)
 - ▶ Number, gender, animacy, person, NER type
Sg-F-A-3-PER / Sg-F-A-3-PER
 - ▶ Mention type: Proper noun (P), Definite, Indefinite, Pronoun (Pr)
P / Pr
- ▶ Attributes of antecedent entity
 - ▶ Entity shape of the cluster (sequence of mentions)
P-Pr-D
- ▶ Features of an anaphor-antecedent pair
 - ▶ Sentence distance: 1
 - ▶ Mention distance (intervening mentions): 4
- ▶ Document features
 - ▶ Genre: Dialog, News, ...
N

Evaluation of Coreference Resolution: F-Measure

MUC: Message Understanding Conference

- ▶ MUC F-measure is based on coreference links (pairs of mentions)
- ▶ H : set of hypothesis clusters, i.e., what the tool finds
- ▶ R : set of reference clusters, i.e., the ground truth

- ▶ Precision

$$\frac{|H \cap R|}{|H|}$$

- ▶ Recall

$$\frac{|H \cap R|}{|R|}$$

- ▶ Somewhat conventional

Evaluation of Coreference Resolution: B³

Based on the presence or absence of a mention relative to entities with which it is confused

- ▶ H_e : hypothesis cluster containing mention e
- ▶ R_e : reference cluster containing mention e
 - ▶ Precision for mention e

$$\frac{|H_e \cap R_e|}{|H_e|}$$

- ▶ Recall for mention e

$$\frac{|H_e \cap R_e|}{|R_e|}$$

- ▶ Overall precision: weighted sum of precisions for all mentions
- ▶ Overall recall: weighted sum of recalls for all mentions
- ▶ Information extraction: Equal weights for all mentions
- ▶ Information retrieval: Equal weights for all clusters

Winograd Schema Problems

Indicates the need for deep world knowledge for successful coreference

- ▶ Pairs of discourses with minor difference that inverts the interpretation

The trophy didn't fit into the suitcase because it was too **large**

versus

The trophy didn't fit into the suitcase because it was too **small**

- ▶ Two relevant entities: e.g., *trophy*, *suitcase*
- ▶ A pronoun that could refer to either entity, e.g., *it*
- ▶ Traditionally phrased as a question answering problem
 - ▶ What was too large?
 - ▶ What was too small?

Bias in Coreference

Automated tools and people can be biased in their interpretations

The nurse didn't meet the surgeon because he was late

The nurse didn't meet the surgeon because she was late

- ▶ Exercise: give an example in the spirit of a Winograd Schema that demonstrates bias in
 - ▶ Gender
 - ▶ Age
 - ▶ Race or ethnicity