Sentence Planning 3: Referring Expression Generation (REG aka GRE)

Lecture 11

March 15, 2013

Reading: *Chapter 5, Reiter and Dale*

R. Dale and E. Reiter (1995)

Pipelined Microplanning

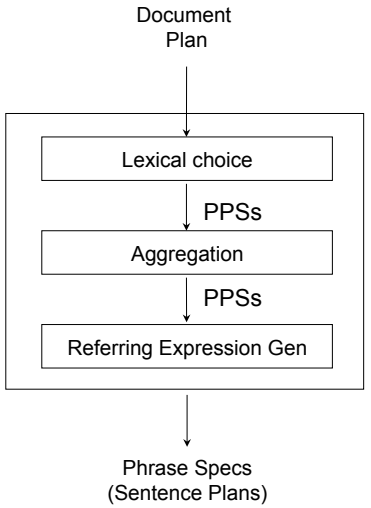

Lexical choice: selects words and syntactic structures to express messages. Result is proto-phrase-spec. PPSs may contain refs to domain entities

Aggregation: combines multiple PPSs into single PPS

REG: takes each PPS and replaces references to domain entities with a phrase spec corresponding to a noun phrase that will uniquely identify that entity to the reader

The GRE task

- Symbolic names of knowledge base entities in PPSs must be replaced by semantic content for referring expressions that are sufficient to identify intended referent(s) to the reader/hearer

Two issues:

- Initial introduction of an object
- Subsequent references to an already salient object

Definite and Indefinite Noun Phrases

Definite NPs:

- Noun phrases marked by definite determiners
  - *The train is about to leave.*
  - *Those trains will leave before ours.*
- Proper names
  - *The Caledonian Express leaves for Aberdeen from track 11*
- Pronouns
  - *It's about to leave*
  - *You'll miss it if you don't hurry*

Indefinite NPs:

- Marked by indefinite determiners
  - *A train is about to leave.*
  - *Some trains already arrived.*

Initial Reference

Introducing an object into the discourse:

- Often use Indefinite NP
  A woman came into the room
- Can also use Definite NP
  - Full Proper Name
    Barack Obama gave a speech on Friday
  - Relate to an object that is already salient
    The lecturer for NLG sent email about the assignment
  - Introduce physical objects by their location
    The train on track 12 is just about to leave
  - When entity is likely to be known or inferable by hearer
    Where is the train station?

Subsequent Reference

- Refer to an entity already introduced into the discourse
- Interpretation dependent on preceding material in the discourse, i.e., ***anaphoric reference***
- To interpret reference, must identify the ***antecedent***
- Typically use Definite NPs
  - Pronouns
    The Prime Minister was accused of bullying. He denied the charges.
  - Definite NPs
    The leader was shocked by the allegations.
  - Proper names, possibly abbreviated
    Mr. Brown said he would never do such a thing.

Generation of Referring Expressions

- Is a microcosm of NLG
  - Content Selection: which properties to express to uniquely describe the intended referent
  - Syntactic Realization: which syntactic configuration to use
  - Lexical Choice: which words to choose
- Today we're going to focus on a sub-problem in GRE
  - Content Selection: choosing the properties of the entities to be included in the referring expression
  - Definite descriptions only

Definite NP Referring Expressions

- Definite NPs that identify a (salient) entity
  - *E.g., identify a particular animal from a visible group of animals*
    - *The poodle*
    - *The black cat*
- Content decisions
  - Attributes: *the cat, the black cat, the large cat,* or *the large black cat*?
  - Specificity: *the dog* or *the poodle*?

Overview of Lecture

- Gricean conversational implicature
- Applying Grice to GRE
  - 3 algorithms
- A better algorithm?
  - Incremental algorithm
- Some recent work on GRE
- GRE as a shared task for NLG

Grice's Cooperative Principle

- Describes how people normally behave in conversation.

  "Make your contribution such as it is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged."
- Speakers (generally) observe the cooperative principle, and listeners (generally) assume that speakers are observing it.
- Allows implicatures, meanings that are not explicitly stated, but which can be inferred, e.g.,

  A: Mark looks exhausted.

  B: He's got a new baby.

Gricean Maxims

Pragmatic directives for speakers:

- Maxim of **Quality**: Truth
  - Do not say what you believe to be false.
  - Do not say that for which you lack adequate evidence.
- Maxim of **Quantity**: Information
  - Make your contribution as informative as is required for the current purposes of the exchange.
  - Do not make your contribution more informative than is required.
- Maxim of Relation: **Relevance**
  - Be relevant.
- Maxim of **Manner**: Clarity
  - Avoid obscurity of expression.
  - Avoid ambiguity.
  - Be brief.
  - Be orderly.

Conversational Implicature

Violation of the maxims leads hearer to make inferences about why the violation occurred.

Miss X produced a series of sounds which corresponded closely with the score of "Home Sweet Home."

implies (to hearer) that speaker could not say

Miss X sang "Home Sweet Home"

(Grice, 1975)

Referring expression example

- In a room full of MacBooks and a single PC, if a speaker says

  Use the PC with the DVD drive.

- The hearer will infer that the speaker could not just say

  Use the PC.

- Implicates to hearer that there is at least one PC (in the room) that does not have a DVD drive
  - *false implicature* if there is no PC that does not have a DVD drive (in the room)
- Saying too much can be dangerous!

Amount of content

- Saying too much can lead to false implicatures
  - *The PC with a DVD drive.*
- Saying too little leads to ambiguity
  - *The MacBook* when there are 10 salient MacBooks
- How much should we say?
  - Applies to many NLG tasks, not just generation of referring expressions

Dale & Reiter Reference Task

- Only examined referring expressions that are
  - definite NPs, e.g., *the black dog*
  - references to physical objects
  - references to "salient" objects
  - intended to distinguish the target object from the set of salient objects
- Simple task enabled in depth analysis of algorithms, impact of Gricean maxims

Distinguishing Descriptions

- A referring expression successfully identifies the target if it is a ***distinguishing description***.
- A distinguishing description matches the target, but none of the ***distractors*** (other salient objects)
- Represented semantically by object type (head noun) and properties (modifiers)

Identify an entity to the hearer

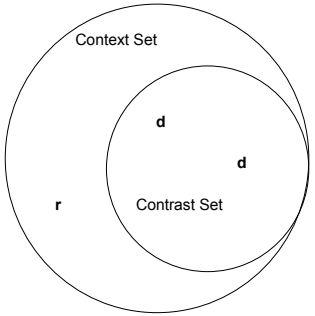

r = intended referent
d = potential distractor
C = Contrast Set (distractors)
L = properties in generated description

Object1: <type,dog> <size,small> <colour,black>
Object2: <type,dog> <size,large> <colour,white>
Object3: <type,cat> <size,small> <colour,black>

r = Object1
C = {Object2, Object3}
L= {<type,dog>,<colour,black>}

Distinguishing descriptions
the black dog
the small dog
the small black dog

Non-distinguishing descriptions
the dog
the small animal

Interpreting the Maxims

- **Quality:** referring expression must be accurate description of target (always true for distinguishing description)
- **Quantity:** no "extra" attributes beyond what's needed to distinguish the target from other salient entities
- **Relevance:** all attributes must rule out at least one distractor
- **Manner:** use the shortest possible NP

(Manner subsumes Quantity and Relevance.)

Full Brevity Algorithm

To obey maxims and avoid false implicatures, never use more than the minimal number of properties required for uniquely identifying the intended referent *(Dale 1989)*

An algorithm:

1. **Check whether 1 property is enough**
2. **Check whether 2 properties are enough**

….

Etc., until

success {minimal description is generated} *or*
failure {no description is possible}

Computational Analysis: Full Brevity Algorithm

- Finding shortest distinguishing description requires exhaustive search
- Worst-case, this algorithm would have to inspect all combinations of properties
  n properties $\Rightarrow 2^n$ combinations
- Equivalent to finding a minimal set cover (where length is number of properties)
- NP-Hard task, therefore computationally intractable!
- Obeying (this interpretation of) the maxims can be very expensive, even for this simple reference task

Solutions

- Ignore complexity issues, hoping this won't be a problem in real applications
- Approximate "shortest possible"
- Look at what people do

Approximation 1: Greedy Heuristic

- Uses a well-known approximation algorithm for minimal set-cover to produce referring expressions
- Chooses property with smallest set of values first

  Intended Referent: Object1 (note: all objects are of type cup)

  Object1: <size,large>, <colour,red>,<material,plastic>
  Object2: <size,small>, <colour,red>,<material,plastic>
  Object3: <size,small >, <colour,red>,<material,paper>
  Object4: <size,medium>, <colour,red>,<material,paper >
  Object5: <size,large>, < colour,green >,<material,paper >
  Object6: <size,large>, < colour,blue >,<material,paper >
  Object7: <size,large>, <colour,blue>,<material,plastic>

- Would first select *plastic*, then *large* or *red*, and then the other of *red* or *large*, whichever was not picked second
- In this case would select: *the large red plastic cup* when the true minimal description is: *the large red cup*

Approximation 2: Local Brevity

- A declarative approximation to a straightforward interpretation

  "it must not be possible to replace 2 or more existing attributes by a single new attribute."

- Requires an initial distinguishing description (perhaps using the Greedy Heuristic), followed by iterative improvement steps
- Preference Rules:

  No Unnecessary components:

  the small black dog vs *the black dog*

  Local Brevity: *the sleeping female dog* vs. *the small dog*

  Lexical Preference: basic-level words should be used

  chihuahua vs. *dog*

  (Reiter, 1990)

Observations from Psycholinguistics

- Humans often include unnecessary modifiers in the referring expressions they generate

  (Levelt, 1989, pp. 129-143)

- Humans can begin uttering a referring expression before they have finished scanning the set of distractors (incremental generation)
  - From eye-tracking studies *(Pechmann, 1989)*

Typical Psycholinguistic Study

- Show a person a picture containing
  - (A) a white bird
  - (B) a black cup
  - (C) a white cup
- Ask person to identify (A)
- Result is often:
  - *the white bird*,

  not minimal
  - *the bird*

What Do People Do?

- D&R guess that people use a simple incremental algorithm
- Start with a null description, and then scan through the set of distractors, adding (speaking) attributes as necessary to rule out distractors
  - Attributes have a preference order
- Attributes that become unnecessary are not removed (unsaid)

Example

- Set of objects:

  {white(b_1), bird(b_1),
  black(b_2), cup(b_2),
  white(b_3), cup(b_3)}
- Target: b_1
- Processing
  - Add "white" to rule out b_2
  - Add "bird" to rule out b_3
- Result is "white bird"
  - "White" is now unnecessary, but is kept

Incremental algorithm

- Basic algorithm:
  - order properties to be used in distinguishing descriptions (preference order):

    $\mathbf{P} = p_1 < p_2 < \ldots < p_n$
  - iterate through ordered list of properties **P**
  - add attribute to description being constructed, if it rules out any distractors that haven't been ruled out yet,
  - terminate when a distinguishing description has been constructed (or no more properties)
- Violates brevity and quantity maxims, but is ***very*** fast.
- If people can "get away" with it, why shouldn't NLG systems?

(Dale and Reiter, 1995)

Evaluation of the Four Algorithms

Interpretation	Theoretical complexity	Typical Run-time	Mimics human behaviour?
Full Brevity	NP-Hard	$\approx n_a^{nl}$	no
Greedy Heuristic	polynomial	$\approx n_a n_d n_l$	yes
Local Brevity	polynomial	$\approx n_a n_d n_l$	no
Incremental Algorithm	polynomial	$\approx n_d n_l$	yes

n_a: the number of properties known to be true of the intended referent
n_d: the number of distractors in the current context
n_l: the number of attributes in the final referring expression

Recap

- D&R tried to formalize the Gricean maxims for a simple reference task.
- The formalization was computationally intractable, so they approximated (weakened) it to make it tractable
- The resulting algorithms were still more complex/ expensive than what people do
- D&R now recommend a simple/fast algorithm that does not adhere to all of the maxims
- Can we safely ignore the maxims?
  - Hypothesis: goal-oriented and sub-language conformant NLG systems will not create false implicatures

On-going debate

- Grice's theory assumes speakers spend lots of effort optimizing what they say for the hearer's benefit, called ***audience design***
  - Many theories and psycholinguistic studies support the idea of audience design
- But other psycholinguistic studies show speakers make non-optimal choices for readers/listeners
  - Speakers are "lazy", they do what's easiest for them
  - Time pressure affects production of descriptions *(Horton & Keysar, 1996)*

Subsequent Work on GRE: van Deemter

- Extensions of the incremental algorithm
  - Overlapping attribute values
    - *bought by Philips* does not rule out *bought by Sony*
  - Reference to sets
    - *the black dogs*
  - Negations, conjunctions, disjunctions
    - *The black dog that is not a poodle*

K van Deemter (2002), Generating Referring Expressions: Boolean Extensions of the Incremental Algorithm. *Computational Linguistics* **28**, pages 37-52

Krahmer et al

- Graph-based model
  - Vertices are objects
  - Edges are attributes or relations
  - Goal: find subgraph that is isomorphic to target but not to any distractor
  - Allows referring expressions to include relations
    - *The black dog next to an oak tree*

E Krahmer *et al* (2003), Graph-Based Generation of Referring Expressions. *Computational Linguistics* **29**.

GRE: A Shared Task for NLG

- First NLG Challenge on Attribute Selection for Generating Referring Expressions (ASGRE), was held in Copenhagen in September 2007 in conjunction with the UCNLG+MT Workshop
- Referring Expression Generation Challenge 2008

  http://www.itri.brighton.ac.uk/research/reg108
- TUNA Challenge 2009

  http://www.itri.brighton.ac.uk/research/genchal09/tuna/
- *Current NLG Challenges*

  http://www.itri.brighton.ac.uk/research/genchal12/

Conclusion

- Def NP reference is good "laboratory"
  - Relatively simple task
  - Can explore deep issues about pragmatics, semantics, communication …
    - Hopefully these findings generalise
- Also important for high-quality text
  - Hard to do in template-based system
  - Advantage of real NLG

References

- R. Dale (1989). Cooking up Referring Expressions. *Proc of 27th Meeting of the Association of Computational Linguistics*, pp. 68-75.
- R. Dale and E. Reiter (1995). "Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions". *Cognitive Science* 19:233-263.
- H.P. Grice (1975). "Logic and Conversation." In R. M. Harnish (ed.), "Basic Topics in the Philosophy of Language", pp. 57-73.
- W.S. Horton and B. Keysar (1996). "When do speakers take into account common ground?" *Cognition 59:* 91–117.
- W. Levelt (1989). *Speaking: From Intention to Articulation.* MIT Press.
- T. Pechmann (1989). "Incremental speech production and referential overspecification." *Linguistics* 27: 89-110.
- E. Reiter (1990). "The Computational Complexity of Avoiding Conversational Implicatures." In *Proc of 28th Meeting of the Association of Computational Linguistics*, pages 97-104. MIT Press.