An Introduction to Natural Language Generation

Robert Dale
Microsoft Institute of Advanced Software Technology
and
School of Mathematics, Physics, Computing and Electronics
Macquarie University
Sydney
Australia
rdale@microsoft.com

1. An Overview of NLG
2. Linguistic Realization
3. Text Planning
4. Generating Referring Expressions

Topic 1:

An Overview of NLG

Course Objectives

- to give a broad overview of the field of NLG
- to show the state of the art in NLG
- to give an overview of the more prominent NLG systems and approaches
- to highlight the current major issues in NLG research

An Overview of the Course

1. An Overview of NLG
2. Linguistic Realization
3. Text Planning
4. Generating Referring Expressions

Overview

1. What *is* NLG? Definitions and applications
2. A short history of the field
3. What you need for generation
4. NLG vs NLU

What Is Natural Language Generation?

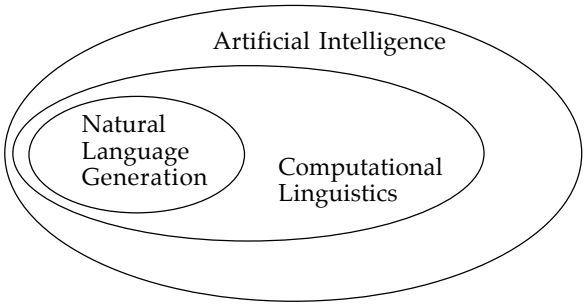

What Is Natural Language Generation?

Generation is
not
the inverse of parsing.

What Is Natural Language Analysis?

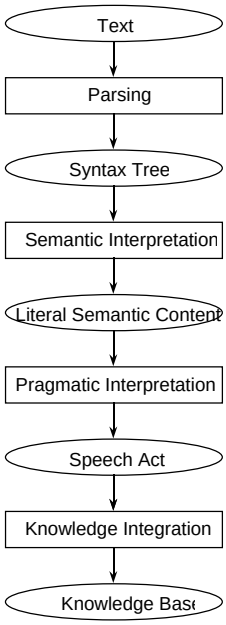

What Is Natural Language Generation?

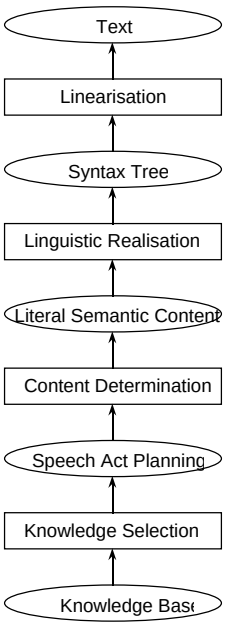

What Is Natural Language Generation?

A definition:

> Natural language generation is the process of deliberately constructing a natural language text in order to meet specified communicative goals.

Key elements:

- making choices
- pursuing goals

Why Bother to Generate Language at All?

Cognitive Science:

- to better understand the human language production facility

Engineering:

- problems with canned output:
  - no guarantee that the text strings will always be consistent with the program's behaviour
  - have to anticipate all the possible questions and answers ahead of time
- as machines become more intelligent, they need more sophisticated means of expression

Principal Uses of NLG Systems

- planning natural language utterances in dialog systems
- planning extended monologic discourses
- psycholinguistic modelling
- composing target language texts in machine translation systems

Other Uses of NLG Systems

Uses we won't consider here:

- evaluation of linguistic grammars
- language and grammar teaching: for example, paraphrasing sentences
- linguistic functions in word processors: for example, composing semi-standard documents

Other caveats:

- we'll focus on generation work where the output is English
- we won't consider systems which produce speech output

Overview

1. What *is* NLG? Definitions and applications
2. A short history of the field
3. What you need for generation
4. NLG vs NLU

The Early 1970s

- Generation from semantic networks: Simmons and Slocum [1972]
- Game descriptions: Davey [1972]
- The formalisation of conversational structure: Power [1974]
- Modelling psychoanalytic discourse: Clippinger [1974]
- Lexical choice in the Conceptual Dependency framework: Goldman [1975]

Simmons and Slocum

Input: a semantic net whose nodes are word-sense meanings and whose arcs are deep case relations

Realisation Mechanism: an augmented transition network representation of a grammar

Effect: produces surface sentences that realize the content of the semantic network

Simmons and Slocum

John saw Mary wrestling with a bottle at the liquor bar.
John went over to help her with it before he drew the cork.

John and Mary together drank the champagne.

Davey's PROTEUS

Sample output:

I started the game by taking the middle of an edge, and you took an end of the opposite one.

I threatened you by taking the square opposite the one I had just taken, but you blocked my line and threatened me.

However, I blocked your diagonal and threatened you.

If you had blocked my edge, you would have forked me, but you took the middle of the one opposite the corner I had just taken and adjacent to mine and so I won by completing my edge.

Davey's Proteus

Sample input:

[[⟨Proteus⟩ 7] [⟨ACD⟩ 5] [⟨Proteus⟩ 3]]

⇓

[[⟨Proteus⟩ ⟨square 7⟩ start ⟨game⟩ take ⟨square 7⟩]
[⟨ACD⟩ ⟨square 5⟩ take ⟨square 5⟩]
[⟨Proteus⟩ ⟨square 3⟩ take ⟨square 3⟩]]

⇓

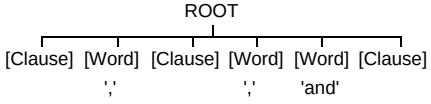

⇓

I started the game by taking a corner. . . .

Power's John and Mary

Mary: John.
John: What's bothering you?
Mary: I want to suggest a goal.
John: Go to hell.
Mary: Will you help me get the door open.
John: No.
Mary: —
John: Mary.
Mary: Yes.
John: I want to propose a goal.
Mary: Go ahead.
John: Will you help me get the door open, even though I was rude.
Mary: By all means.
John: Then let's make a plan.
Mary: John.
John: What.
Mary: May I ask you something.
John: Yes of course.
Mary: Is the door open.
John: No.
. . .

The Late 1970s

- Chester's [1976] program for translating formal proofs into English
- Meehan's story-writing program, TALE-SPIN [1976]
- Thompson's strategy/tactics distinction [1977]
- Cohen's PhD work on the planning of speech acts [1978]
- Moore and Mann's KDS [1979]

Chester's EXPOUND

Sample input:

L1	$\forall x(Nx \rightarrow Hx)$	(PR)
L2	$\forall x((Wx \wedge Hx) \rightarrow Bx)$	(PR)
L3	$\forall x \exists y(Wy \wedge Py, x)$	(KN)
L4	$\forall x \forall y((Py, x \wedge By) \rightarrow Bx)$	(PR)
L5	$\forall y(Py, x \rightarrow Ny)$	(PR)
L6	$\exists y(Wy \wedge Py, x)$	(U1 L3)
L7	$Wc \wedge Pc, x$	(E1 L6)
L8	$Pc, x \rightarrow Nc$	(U1 L5)
L9	$Nc \rightarrow Hc$	(U1 L1)
L10	$(Wc \wedge Hc) \rightarrow Bc$	(U1 L2)
L11	Bc	(TF L7 L8 L9 L10)
L12	$\forall y((Py, x \wedge By) \rightarrow Bx)$	(U1 L4)
L13	$(Pc, x \wedge Bc) \rightarrow Bx$	(U1 L12)
L14	Bx	(TF L7 L11 L13)
L15	$(\forall y(Py, x \rightarrow Ny)) \rightarrow Bx$	(CD L5 L14)
L16	$\forall x((\forall y(Py, x \rightarrow Ny)) \rightarrow Bx)$	(UG L15)

Chester's EXPOUND

Sample output:

THEOREM: Suppose that every native of Ajo has a cephalic index in excess of 96. Suppose furthermore that every woman who has a cephalic index in excess of 96 has Pima blood. Suppose that if any person is parented by any person, say w, and ... Then if every person is a native of Ajo whenever he is a parent of any person, say p, then the person p has Pima blood.

Proof: Suppose that every person who is a parent of person x is a native of Ajo. Since some woman is a parent of y ...

The Early 1980s

- McDonald's MUMBLE [1980]
- The HAM-RPM German dialogue system [von Hahn *et al* 1980]
- Work started on the Penman/Nigel framework at ISI [Matthiessen 1981]
- Explanation generation for expert systems [Swartout 1981]
- McKeown's and Appelt's influential PhD theses appeared in 1982
- Kempen and Hoenkamp's work on incremental sentence generation [1982]
- Appelt's TELEGRAM [1983]: unification grammar in generation
- Kukich's [1983] stock market report generator

McDonald's MUMBLE

Sample input:

```
(discourse-unit
  :head (general-clause
          :head (chase
                  (general-np
                    :head (np-proper-name "Fluffy")
                    :accessories
                       (:number singular
                        :determiner-policy no-determiner))
                  (general-np
                    :head (np-common-noun "mouse")
                    :accessories
                       (:number singular
                        :determiner-policy kind))
                    :further-specifications
                       ((:specification
                           (predication_to-be *self*
                              (adjective "little"))
                         :attachment-function
                             restrictive-modifier)))))
          :accessories (:tense-modal present
                        :progressive
                        :unmarked))))
```

Sample output:

Fluffy is chasing a little mouse.

Kukich's ANA

Some Input Data:

06/24	IND	814.69	821.63	805.56	810.41	DN	2.76
06/24	TRN	317.97	321.57	313.85	317.00	UP	0.30
06/24	UTL	106.13	106.83	105.61	106.83	UP	0.70
06/24	CLOSINGSTOCK		30	INDUS	810.41		
06/24	CLOSINGSTOCK		20	TRANSP	317.00		
06/24	CLOSINGSTOCK		15	UTILS	106.83		
06/24	CLOSINGSTOCK		65	STOCKS	314.90		
06/24	330PM		30	INDUS	812.02		
06/24	330PM		20	TRANSP	317.60		
06/24	330PM		15	UTILS	106.52		
06/24	330PM		65	STOCKS	315.28		
06/24	3PM		30	INDUS	810.12		
06/24	3PM		20	TRANSP	316.47		
06/24	3PM		15	UTILS	106.70		
06/24	3PM		65	STOCKS	314.63		
06/24	230PM		30	INDUS	810.50		
06/24	230PM		20	TRANSP	316.77		
...							
06/24	11AM		15	UTILS	106.22		
06/24	11AM		65	STOCKS	315.46		
06/24	1030AM		30	INDUS	814.12		
06/24	1030AM		20	TRANSP	318.05		
06/24	1030AM		15	UTILS	106.39		
06/24	1030AM		65	STOCKS	315.79		
06/24	ISSUESTRADED			1849			
06/24	ADVANCED			772			
06/24	DECLINED			660			
06/24	UNCHANGED			417			
06/24	NATION			65223320	72820880		
06/24	NYSEVOL			55860000	62710000		
06/23	IND	801.75	815.35	795.76	813.17	UP	13.51
06/23	TRN	308.98	317.37	306.65	316.70	UP	8.92
06/23	UTL	106.13	106.83	105.61	106.13	DN	0.17

Kukich's ANA

Sample output:

Thursday June 24, 1982

wall street's securities markets meandered upward through most of the morning, before being pushed downhill late in the day yesterday. the stock market closed out the day with a small loss and turned in a mixed showing in moderate trading.

the Dow Jones average of 30 industrials declined slightly, finishing the day at 810.41, off 2.76 points. the transportation and utility indicators edged higher.

volume on the big board was 55860000 shares compared with 62710000 shares on Wednesday. advances were ahead by about 8 to 7 at the final bell.

Appelt's KAMP

- KAMP is given the goal of removing the pump from the platform
- KAMP plans for John (the novice) to remove the pump
- KAMP plans to request that John remove the pump
- KAMP expands the request as an imperative utterance with the main verb *remove*
- KAMP knows that John needs to know where the tool is, so he plans to tell John; this goal is subsumed into the realization of the higher level goal.

The output:

Remove the pump with the wrench in the toolbox.

The Mid 1980s

- Mann and Thompson explored and developed the idea of rhetorical structure [1984]
- People started thinking about user modelling: McCoy [1984, 1985], Paris [1985]
- McDonald and Pustejovsky [1985] picked up Joshi's Tree Adjoining Grammar as a suitable formalism for use in NLG

The Late 1980s

- Hovy's work on affect in PAULINE [1986, 1987]; Jameson's work on violating Gricean maxims [1987]
- Work began on the interaction of language and gesture [Kobsa *et al* 1986]
- Increasing interest in bidirectionality [Lancel *et al* 1986; Appelt 1987]
- The generation of referring expressions [Novak 1986, Reithinger 1987, Dale 1988]
- RST was developed further, and first attempts were made at operationalising [Hovy 1988, Moore 1988]

Hovy's RST Text Structurer

Sample input:

```
((SHIP.EMPLOYMENT A105)
 (SHIP.R A105 KNOX)                    (ENROUTE E105)
 (SHIP.COURSE.R A105 195)              (EBEG.R E105 870420)
 (CURRENT.POSITION.R A105 P102)        (EEND.R E105 870424)
 (POSITION P102)                       (DESTINATION.R E105 SASEBO)
 (LONGITUDE.R P102 79)                 (LOAD E107)
 (LATITUDE.R P102 18)                  (EBEG.R E107 870425)
 (READINESS.LEVEL.R A105 C4)           (EEND.R E107 870428)
 (NEXT.MAJOR.EMPLOYMENT.R A105 E107)
 (CURRENT.MAJOR.EMPLOYMENT.R A105 E105)
```

Sample output:

> Knox, which is C4, is en route to Sasebo. It is at 79N 18E heading SSW. It will arrive on 4/24, and will load for four days.

The Late 1980s

- Examination of the relationships between systemic grammar and functional unification grammar, and the formal properties of these formalisms
- The use of well-founded grammatical formalisms for linguistic realization blossomed: LFG, GPSG, TAG
- Mumble-86 stabilised
- First explorations of connectionism in NLG: [Dyer 1988]

The Early 1990s

- Multimodality: WIP, COMET, IDAS
- Another big systemic grammar: Fawcett's GENESYS
- The discourse relations explosion, and the debate on intentions
- DiMarco *et al*s first attempts at a notion of style in generation
- Collocations, and the increasing visibility of MTT
- The input to the generation process: do words come first? A reappraisal of lexical choice
- RST in instructional texts

The WIP Project

Sample input:

> a formal description of the communicative intent of a planned presentation

The modality choice question:

- what should go into text
- what should go into graphics
- what kind of links between verbal and non-verbal fragments are needed

The WIP Project

Sample output:

Lift the lid.

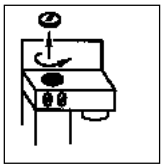

To fill the watercontainer, remove the cover.

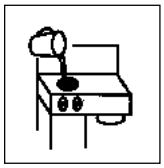

Use cold tapwater.

The State of the Field

Reasons for increased interest in NLG over the last 10 years:

- several major pieces of research
- emergence of work relevant to both NLU and NLG
- increasingly sophisticated underlying programs

Some Provisos

The field is not yet mature:

- work in generation research is not as consolidated as work in, for example, parsing;
- there is little consensus on the nature of the problems;
- there is no common starting point.

Overview

1. What *is* NLG? Definitions and applications
2. A short history of the field
3. What you need for generation
4. NLG vs NLU

Knowledge Sources in NLG

Levels of representation and processing:

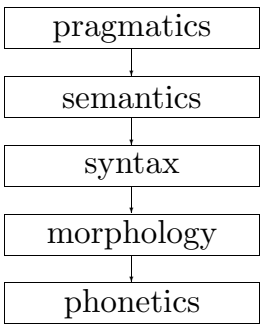

The Stages of the Generation Process

Natural language generation is a goal oriented process with three identifiable stages:

- identifying the goals the utterance is to achieve;
- planning how the goals may be achieved, including evaluating the situation and the available communicative resources; and
- realizing the plans as a text.

Implicated Components

A non-trivial generation system requires the following:

- a non-linguistic reasoning component;
- some representation of the discourse;
- some notion of a model of the audience;
- some representation of the available linguistic resources, typically encoded within both a grammar and a lexicon.

Some Terminology

Linguistic Resource: an item, supplied by the natural language, that can carry information and can therefore be used to realize some element of the utterance content or goals.

Message: a structure at an intermediate representational level that records the information that is selected to go into an utterance.

A Standard Archiecture for NLG

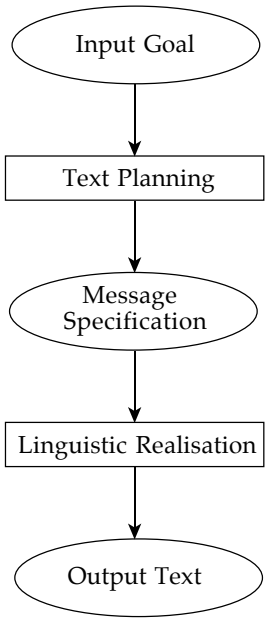

A Standard Archiecture for NLG

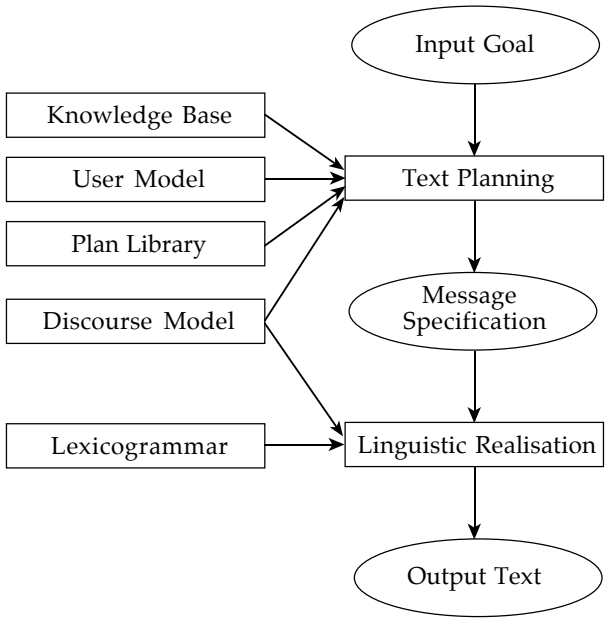

Language Generation as Choice

At least the following decisions have to be made by any complete NLG system:

- text content
- what information should be omitted
- organisation of content into a coherent discourse
- tone or degree of formality
- decomposition into sentences
- choice of syntactic constructions
- how entities should be described
- choice of words

An Example

Consider:

> This course is being taught by Robert Dale. It is an introduction to natural language generation.

This text embodies the following decisions:

- Of all the things known about the course, it states the lecturer's name and the topic of the course.
- It uses two simple sentences rather than one more complex sentence.
- It uses a passive rather than an active sentence for the first piece of information.
- It uses the phrase *being taught* rather than *being given*.
- It uses the pronoun *it* in the second sentence, in preference to a full noun phrase.

Overview

1. What *is* NLG? Definitions and applications
2. A short history of the field
3. What you need for generation
4. NLG vs NLU

Differences Between NLG and NLU

- the balance of research effort
- information flows in opposite directions

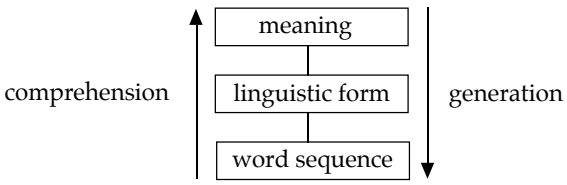

What Generation and Comprehension Share

- same basic notion of a lexicon, using a taxonomy of basic word classes, word senses, and morphology;
- fairly shared notion of grammar as a means of describing the constructions available in a language; and
- descriptions of various discourse phenomena (particularly anaphora) are important in both areas.

One View

The perceived differences between the NLG and NLU tasks are unreal.

- every evident problem has a counterpart which may or may not be evident on the other side of the fence
- underlying claim: if a process is used in generation, it has effects which may be discernible, interpretable, and possibly significant.

The Focus of Research

- In NLU the known is the text, perhaps with intonational information; the unknown is whatever the researcher chooses as a stopping point—typically some form of semantic specification with anaphors resolved.
- In NLG the known is the system's goals and intentions: but at what level do you specify these?

Problems Particular to NLU

- covering all the ways to say things
- goal identification
- vocabulary coverage
- ambiguity

Problems Particular to NLG

- deciding how much to say, and what not to say:
  - maintaining brevity
  - avoiding stating the obvious
- designing text structure:
  - may need to add material to the basic subject matter
  - controlling the effects of the structure and ordering of the material
  - making the text flow smoothly
- problems in carrying out a detailed text plan once built:
  - determining the sentence boundaries and the use of conjunctions
  - deciding when to use anaphora
  - lexical selection
  - use of marked syntactic structures for particular rhetorical effects

A Summary of the Two Processes

Comprehension:

- the known is the wording of the text
- the primary effort is to scan the text, during which its linguistic form and meaning gradually become apparent
- the algorithms are based on hypothesis management
- the major problems are ambiguity and under-specification

A Summary of the Two Processes

Generation:

- the known is the speaker's intentions, and the selected content and perspective
- the primary effort is choosing from alternatives and establishing plans, constructing specifications and then realizing them
- the algorithms are typically organized as planning by progressive refinement
- ambiguity is a non-issue; the process is over-supplied with source information and must decide what to highlight and what to omit

Summary

What we've done so far:

- defined NLG as
  - a goal-oriented process
  - a process of choice
- looked briefly at the history of the field and some examples
- introduced some concepts that are important in NLG
- compared NLG and NLU

What's Coming Next ...

1. An Overview of NLG
2. Linguistic Realization
3. Text Planning
4. Generating Referring Expressions