Wisdom is not the product of schooling but the lifelong attempt to acquire it.
- Albert Einstein

Natural Language Processing

Gerhard Fischer

AI Course, Fall 1996, Lecture, Nov 11

The Larger Picture

beyond the isolation assumption of AI (Bobrow article)

chapter 22 in Russell/Norvig: Communication and Collaboration

joint human-computer systems
- computer-supported cooperative work (CSCW) and computer-supported cooperative learning (CSCL)
- cooperative human-computer systems (critiquing)
- distributed artificial intelligence

formal languages versus natural languages
- formal languages: Lisp, first-order languages
- natural languages: English, Chinese
- question: what is an e-mail message?

Tools of Formal Language Theory
(e.g. Backus-Naur Form (BNF))

phrase structure

noun phrases (NP)
verb phrase (VP)

nonterminal symbols

rewrite rules: S ----> NP VP

grammatical versus ungrammatical categories:

the wumpus is dead
wumpus the dead is

deep structure and surface structure (Chomsky):

a man bites a dog
a dog is bitten by a man

Steps of Communication Processes

communication episode: speaker S, hearer H, proposition P, using words W

processes of the speaker:

intention:	S want H to believe P (where S typically believes P)
generation:	S chooses the words W (because they express P)
synthesis:	S utters the words W (addressing them to H)

processes of the hearer:

perception:	H perceives W' (ideally W'=W, but misperception possible)
analysis:	H infers that W' has possible meanings P1, ...Pn
disambiguation:	H infers that S intended to convey Pi (ideally Pi = P)
incorporation:	H decides to believe Pi (or rejects it)

References for Natural Language Processing

chapter 23 in Russell/Norvig: “Practical Natural Language Processing”

James Allen: “Natural Language Understanding”, 2nd edition, 1995, Benjamin/Cummings

Eugene Charniak: “Statistical Language Learning”, 1993, MIT Press

M.J. Bates, R.J. Bobrow, “Natural Language Interfaces: What's Here, What's Coming, and Who Needs It”, in W. Reitman (ed): “Artificial Intelligence Applications for Business”, Ablex Publishing, chapter 10, pp 179-194, 1984

D.G. Bobrow, R.M. Kaplan, M. Kay, D.A. Norman, H. Thompson, T. Winograd, “GUS, A Frame-Driven Dialog System”, Artificial Intelligence Journal, Vol 8, pp 155-173, 1977

M. Kay: “The Proper Place of Men and Machines in Language Translation”, Xerox Palo Alto Research Center, CSL-80-11, October 1980

Why is Natural Language Processing Difficult?

one has to have a lot of specific knowledge about the words and the grammar rules of a language

one must be able to integrate this knowledge with other knowledge about the world

there are other agents in the world who have their own beliefs, goals, and plans

encoded message model of communication: a speaker encodes a representation of a proposition into language, and the hearer then decodes the message to uncover the proposition

situated language model: the meaning of a message is a function both of the message and the situation in which occurs

Different Levels of Language Analysis

phonological knowledge — relates words and sounds; crucial for speech-based systems

morphological knowledge — relates morphemes and words (friend ---> friendly)

syntactic knowledge — words to sentences; structural roles of words in sentences

semantic knowledge — meaning of words and sentences; context-independent

pragmatic knowledge — how do situations affect the interpretation

discourse knowledge — beyond the sentence boundary (e.g., pronouns

world knowledge — user's belief and goals

Practical Applications of Natural Language Processing

machine translation (see example from SYSTRAN Software HTML Translation)

database access

natural language interfaces

natural language based help systems (e.g., UNIX Consultant)

information retrieval

text categorization

extracting data from text (“data mining”, CACM, Vol 39, No 11, Nov 96)

Examples of Understanding

I saw the Grand Canyon flying to new York.

I saw the man in the park with the telescope.

Question: What does the "**they**" refer to?

1. The city councilmen refused to give the women a permit for a demonstration, because **they** feared violence.

2. The city councilmen refused to give the women a permit for a demonstration, because **they** advocated violence.

Version 1 of a Dialogue

P: **Men are all alike.**
C: In what way?

P = Person
C = Computer Program

P: **They are always bugging us about something or other.**
C: Can you think of a specific example?

P: **Well, my boyfriend made we come here.**
C: Your boyfriend made you come here?

P: **He says, I am depressed much of the time.**
C: I am sorry to hear that you are depressed.

P: **It is true, I am unhappy.**
C: Do you think coming here will help you not to be unhappy?

P: **I need some help, that much seems to be certain.**
C: What would it mean to you if you got some help?

P: **Perhaps I could learn to get along with my mother.**
C: Tell me more about your family!

Version 2 of a Dialogue with the same System

P: **My mumble bzzz grumble grumble.**

C: That's very interesting.

P: **Ugh, you izzle flerp me.**

C: What makes you think that I izzle flerp you?

P: **Gorch.**

C: Does that have anything to do with the fact that your mumble bzzz grumble grumble?

GUS: A Travel Planning Assistant

Computer: When do you want to leave?

User: I must be in LA at 10:00am

Computer: Continental Flight #342 arrives at 9:30am?

User: What is the next flight?

Syntax, Semantics and Pragmatics

S_1: Language is one of the fundamental aspects of human behavior and is a crucial component of our lives.

S_2: Green frogs have large noses.

S_3: Green ideas have large noses.

S_4: Large have green ideas nose.

Research Issues

indirect speech acts:

"Do you have Tabasco?"
"Do you know what time it is?"

Unix Consultant:

Question by User: "How can I get more disk space?"
Unix Consultant: "Type rm '*' "

speaker versus listener role

command systems (e.g., UNIX)
menu-systems (e.g., GUI; but notice: short-cuts in Mac Interface)
menu-based natural language systems

limitation of language

visualization
"a picture is worth a thousand words"

Natural Language Interfaces

when is natural language the most appropriate interface language? — not for:

driving a car or playing a video game

object identification through pointing (rather than descriptions)

how can we assess natural language interfaces:

coverage = characterization of the linguistic competence of a system (lexical, syntactic, semantic)

habitability = measures how quickly and comfortably can a user recognize and adapt to the system's limitations

Human-Human versus Human-Computer Interaction

claim by Lucy Suchman (in “Plans and Situated Action”, p 185): *“because of the asymmetry of user and machine, interface design is less a project of simulating human communication than of engineering alternatives to interaction’s situated properties”* ----> mimicking versus complementing

human conversational partners

share a lot of information
model one another’s knowledge and capabilities
process huge amounts of (conflicting) information
update all of these structures as the conversation progresses
manage trouble / communication breakdowns

hypotheses:

technology for useful, cost-effective, natural language interfaces is available
but: these interface will not behave like human conversational partners, so users must carefully examine such systems to understand their capabilities and limitations

Some Theoretical Issues

Chomsky/Fodor: “a language module exists only in the human brain”

do animals have models of the world? — e.g. can an animal respond to: “Bring the red ball over to the blue ball”

language and thought reinforce each other, but

did humans evolve to use language well because they are smart?

are they smart because they use language well?

from artificial intelligence to statistical natural language processing (Charniak):

language comprehension from an AI point of view requires real world knowledge / common sense knowledge

knowledge representation (theoretical <----> CYC; focus on mechanisms and formalisms rather than on very large content)