


11. Introduction to natural language processing (NLP)

11.0 Communication

- Several examples of sophisticated communication systems in the animal world exist:
 - Honeybee dance
 - Communicates the direction and distance to the food source.
 - Most advanced is the 'waggle dance.'
 - Vervet monkey alarm calls
 - Different alarm calls for different predators: eagles, leopards, and snakes (Seyfarth & Cheney, 1992)
 - Responds to these alarm calls differ as if the animal had seen the predator itself.
 - Cheating involved.



11.1 What is natural language ?


- Spoken or signed language used by humans; does not include written language.
- Language of native speakers
- Evolved naturally
 - Not constructed
 - Not formal
- Has symbolic reference
- Compositional: meaning of an expression is a function of its constituents.
- Combinatorial: finite number of elements can be combined to infinite number of expressions.
- Open-ended: there is no limit in the (number of) meanings that can be expressed.

11.2 Study of natural languages

- Natural language processing (NLP) divided into study of automated

1. generation (speech production)

2. understanding (comprehension)




3. acquisition (evolution)

- Traditionally studied in isolation from other domains ← prevailing strongly modular theory of human mind.

- In several academic disciplines by

- Linguists
- Psycho/cognitive linguists
- Philosophers
- Computational linguists



Potential of bringing together several sub-fields of AI:

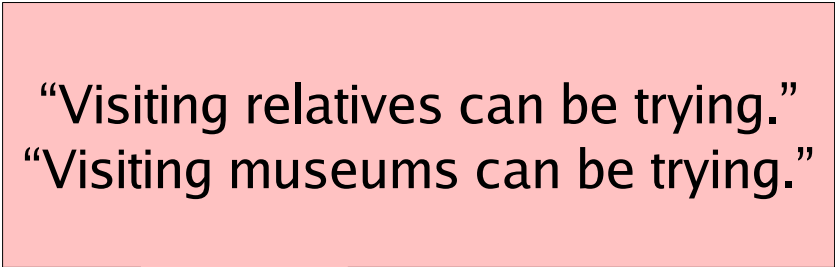
- computer vision
- information retrieval
- machine learning
- planning

11.3 Goals of NLP

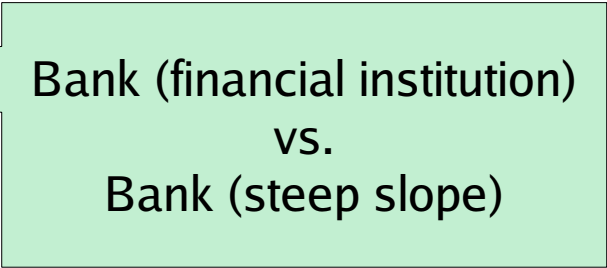
- *Theoretical* goal of computational models of language is to gain understanding in how language works:
 - Structure of language
 - Human language related processes
 - How words get their meaning, and how cognitive constituents, such as beliefs, goals, intentions, relate to language.
- *Practical* goal is to revolutionize technology:
 - Natural language interfaces (e.g., databases, virtual reality, games)
 - Speech control of devices (e.g., cars, doors)
 - Automated customer service (over phone)
 - Tutoring systems

11.4 Other applications of NLP

- Text mining and information retrieval (search engines)
- Automatic summarization or textual report generation
 - Sport summaries
 - Weather forecasts
- Machine translation
- Subgoals:
 - Text and speech segmentation
 - Syntactic disambiguation
 - Word sense disambiguation
 - Understanding of context
 - Dealing with imperfect input
 - Speech acts



“Visiting relatives can be trying.”
“Visiting museums can be trying.”



Bank (financial institution)
vs.
Bank (steep slope)

11.5 Levels of analysis

1. Prosody

- Deals with rhythm and intonation
- Hard to formalize

2. Phonology

- Study of how sounds function in the language; for instance minimal distinctive units (phonemes): /p/ and /b/ in English (pin ≠ bin)
- Important only in speech processing.

3. Morphology

- Study of how words are constructed from meaningful components (=morphemes), for instance friend+ly, drop+ped, John+s.

4. Syntax

- Application of grammar
- Determines the role of words in a sentence.

5. Semantics

- Examination of meaning of words and how these meanings combine to meaning of sentences.
- Context independent meaning not regarding the use.

6. Pragmatics

- How language is used in different situations.
- Conveying meaning not obvious from the semantics: “Do you know what time it is?”- “ Yes.”

7. Discourse

- Temporal aspect: understanding of sequences of expressions.
- How the preceding sentence influences the interpretation of the following sentence, e.g., pronoun reference.

8. World knowledge

- General knowledge about the structure of the world.
- Understanding and knowledge of other language users' beliefs and intentions.

11.6 Syntax vs. semantics vs. pragmatics:

Suggestions for opening words of *Natural Language Understanding* by James Allen, 1995:

1. Language is one of the fundamental aspects of human behavior and is a crucial component of our lives.
2. Green frogs have large noses.
3. Green ideas have large noses.
4. *Large have green ideas nose.

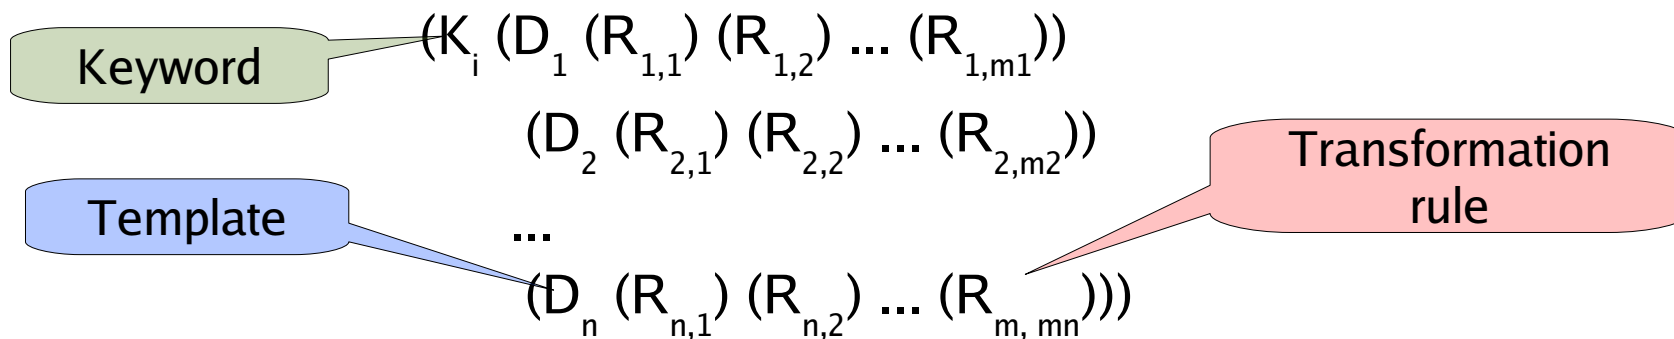
“Colorless green ideas sleep furiously.” (Chomsky, 1957)

11.7 ELIZA

- Joseph Weizenbaum, 1965
- Rogerian therapist
- Holds conversation with a human through keyboard and screen.
- Sample exchange:
 - Person: Men are all alike.
 - Eliza: In what way?
 - Person: They're always bugging us about something specific or other.
 - Eliza: Can you think of a specific example?
 - Person: Well, my boyfriend made me come here.
 - Eliza: Your boyfriend made you come here?
 - Person: He says I'm depressed much of the time.
 - Eliza: I am sorry to hear you are depressed.
 - Person: It's true. I'm unhappy.
 - Eliza: Do you think coming here will help you not to be unhappy?
 - Person: I need some help, that much seems certain.
 - Eliza: What would it mean to you if you got some help?
 - Person: Perhaps I could learn to get along with my mother.
 - Eliza: Tell me more about your family.
- Online demo: <http://www-ai.ijs.si/eliza/eliza.html>

11.8 ELIZA's knowledge

- List of *keywords*
- *Templates* (or decomposition rules) — associated to each keyword — implement a minimal context.
- *Transformation rules* (reassembly rules) associated to the keywords and templates:



- Keywords have ranks or precedences → the most important keyword is chosen.
- Keywords can be tagged, e.g., 'mother' may have a tag for 'family'.

11.9 ELIZA's algorithm

- Given an input sentence S
 - Scan S for potential keywords.
 - Select one with highest rank.
 - Choose the template for that keyword that matches S, e.g., template 'Y are you X' matches sentence "Why are you looking at me?" ('you' is the keyword)
 - Output the response associated to the keyword and the template, e.g., to the input 'I am X', ELIZA replies 'How long have you been X?'
 - If no keywords found, output clarifying questions ("Tell me more," "Go on").
- Makes also unconditional transformations, such as 'I' ↔ 'you', 'yourself' ↔ 'myself'.

11.10 ELIZA's sample data

Keyword	Rank	Template	Transformation rule
alike	10	X	In what way? What resemblance do you see?
are	3 3	X are you Y X are Y	Would you prefer if I weren't Y? What if they were not Y?
always	5	X	Can you think of an example? When? Really, always?
what	2	X	Why do you ask? Does that interest you?

11.11 Why does ELIZA work so well?

- It actually misled quite a few people in 60's into thinking there was a human behind it.
- People attribute meaning to sentences they understand and assume the meaning was intended by the producer.
- In conversational setting easy to sustain illusion of intelligence.
- ELIZA never argues or answers questions, it only asks a series of questions.

11.12 When doesn't ELIZA work?

- It has no world knowledge or colloquial understanding:
 - To the expression “Necessity is the **mother** of invention.” it replies “Tell me more about your **family**.”
- It only understands English keywords:
 - To the input “\$ %%#@ @&!!!” it produces “Can you elaborate on that?”
- No understanding of language structure:
 - To the expression “Green the adzabak are the a ran four” it will reply “What if they were not the a ran four?”
- It does not remember early parts of the exchange → it appears to lose focus as the conversation proceeds.

11.13 Other conversation programs

- Parry
 - Written by Kenneth M. Colby 1971.
 - Simulates paranoid schizophrenic.
 - Is said to have passed the Turing test (by psychiatrists).
 - Received a lot of criticism:
 - Does not use grammar, but input analysis is based on pattern matching.
 - Output based on canned responses.
 - Models paranoid behavior, not the processes producing the behavior.
 - Contains no knowledge.
 - Ad hoc model, lacks generality.

More than iron, more than lead, more than gold I need electricity. I need it more than I need lamb or pork or lettuce or cucumber. I need it for my dreams.

- Racter

- Written tongue-in-cheek by William Chamberlain and Thomas Etter, 1983.
- Artificially insane raconteur that generates English prose at random:
 - plays an active role in conversation
 - jumps from topic to topic
 - makes wild associations
- Written in basic and runs with 64K ram.
- It is claimed to have written a book called “The Policeman's Beard is Half Constructed.”
 - this version of the program never released
- Public version a template based chatterbox.
 - coherence by re-use of text variables

11.14 SHRDLU

- Terry Winograd, 1968-1970.
- First attempt to build a system that understands English language and follows instructions.
- Interacts with the user, answers questions and executes commands in block world.
- What made SHRDLU successful?
 - The world was simple; objects and their locations could be described using about 50 words such as “block”, “pyramid”, “move to”, “big”, “blue”.
 - It has memory to provide context
 - Because of memory, could answer hypothetical questions.
 - Could name object formations and remember (and use) the names.
- How SHRDLU got its name: <http://hci.stanford.edu/~winograd/shrdlu/name.html>

11.15 Transcript

Person: PICK UP A BIG RED BLOCK.

Computer: OK.

Person: GRASP THE PYRAMID.

Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.

Person (changing their mind): FIND A BLOCK WHICH IS TALLER THAN THE ONE YOU ARE HOLDING AND PUT IT INTO THE BOX.

Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.



Computer: OK.

Person: WHAT DOES THE BOX CONTAIN?

Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.

Person: WHAT IS THE PYRAMID SUPPORTED BY?

Computer: THE BOX.




11.16 SHRDLU (cont.)

- Consists of
 - Parser
 - English grammar recognizer
 - Semantic analyzer
 - General problem solver
- Limitations acknowledged soon
 - Generated language not quite as complex (syntactically varying) as language comprehended.
 - The world — block, table, robot, etc. — are internal to the system, but still SHRDLU does not know anything about blocks, or about moving them, dropping them, or which blocks are visible etc.
 - Tremendous success at first, but attempts to model more realistic situations failed.

11.17 Challenges for NLP systems


- Not only about knowing words and grammar. Much more is required.
 - “How much do I owe you?” – “What do you mean by 'do'?”
- World knowledge and disambiguation
 - We gave monkeys the bananas because they were hungry.
 - We gave monkeys the bananas because they were over-ripe.
 - Jane carried the girl with a spade.
 - Jane carried the girl with a banana.
- Context and disambiguation
 - I went into the cave. It was full of bats.
 - I looked into a locker. It was full of bats.



- **Garden-path sentences**
 - The old man the boat.
 - Time flies like an arrow. Fruit flies like a banana.
- **Manuscripts**
 - John entered a restaurant. He ordered a chicken in red wine sauce. He paid the check and left.
 - John entered a park. He asked the dwarf for a mouse. He took the box and left.

11.18 Connecting language to world

- Cross-modal representation



- Tendency to study language in isolation as a formal system.
- Study of language as a *situated* system, and addressing the symbol grounding problem:
 - Translate sensory (not only visual) signals into language.
 - Translate language into physical actions.
 - Language also used to communicate non-physical entities.
- Language learning and language evolution.

11.19 Technical challenges

- Visual representation
 - 'Bag of features' approach in many modern vision systems.
 - Needs to conserve the spatial structure and whole-part relationship of an object. E.g., modifiers in the language; “č up without a handle.”
- Motor control and action representation
 - Needs to conserve the temporal structure of actions.

⇒ Associate words with perceptual and action categories

11.20 Assigning words

- Words are labels to sensory-grounded categories of objects, not individual objects.
- Feature selection
- Word choice depends on context and the sensor data
 - Red hair vs. red car vs. red wine.
 - Non-linguistic information: Green banana
- How specific and detailed description need be?
 - Red vs. crimson
 - Red car with silver trim vs. red car with black tires.

11.21 Verb meanings

- Temporal structure of actions.
- Not only the movement needs to be encoded but the causal aspect: *<someone/thing does something to someone/thing>*
- Challenge in designing planning algorithms for situated language processors:
 - Word choices to resolve possible ambiguities.
 - Planning in communicative and motor actions: use language to coordinate joint physical effort.
 - Mix of action planning and discourse planning.

11.22 Evolution and evaluation

- How is the choice of language to world rules adopted?
 - Learned from examples and feedback.
 - Adopted rules likely differ from system to system.
 - The communication system evolves over time.
- How do we know if the system performs well?
 - Psychological experiments with humans
 - Questionnaires
 - Simulations