Introduction to Natural Language Processing

Chapter 1

Today's Outline

Demo Debriefing

Knowledge of Language

The Ambiguity Problem

Models and Algorithms

Can Machines Think?

History

Administration

Demo Debriefing

- JUPITER
- AnswerBus, Ask Jeeves
- Newsblaster, NewsInEssence
- Babelfish
- Text-To-Speech

HAL

Returning to our HAL example ...

```
Dave: Open the pod bay doors, HAL.
HAL: I'm sorry Dave, I'm afraid I can't do that.
```

The knowledge that HAL needs to take part in this exchange can be broken down into somewhat discrete categories.

Each category can then be studied (and modeled computationally) in isolation. This, of course, leaves open the problem of how the categories interact.

Knowledge of Language

Phonetics and Phonology: speech sounds, their production, and the rule systems that govern their use

Morphology: words and their composition from more basic units

- *cat, cats*
- *friend, friendly*

Syntax: the structuring of words into legal larger phrases and sentences

- *The textbook for the NLP class is great.*
- *Jane met Mary*
- *Mary was met by Jane* (passive)

Semantics: the meaning of words and phrases

- **lexical semantics**: the study of the meanings of words
- **compositional semantics** how to combine words
- *river bank, financial bank*
- *met(Jan,Mary)*

Pragmatics: utterance interpretation in situational context

- *Do you know the time?*

Discourse: utterance interpretation in context of previous utterances

- *Sue took a trip to New York. She had a great time there.*

Deconstructing HAL

Recognizes speech and understands language

Decides how to respond and speaks reply

With personality

Recognizes the user's goals, adopts them, and helps to achieve them

Remembers the conversational history

Customizes interaction to different individuals

Learns from experience

Possesses vast knowledge, and is autonomous

Ambiguity

Why is the following sentence problematic?

- *I saw the woman with the telescope.*

Ambiguity, cont.

Why is the following sentence problematic?

- *I saw the woman with the telescope.*

It is syntactically ambiguous...

- *I saw (NP the woman with the telescope)*
- *I saw (NP the woman) (PP with the telescope)*

The categories of knowledge of language can be thought of as ambiguity-resolving components.

"I made her duck"

How many different interpretations does this sentence have?

- I cooked waterfowl for her
- I cooked waterfowl belonging to her
- I created the (plaster?) duck she owns
- I caused her to quickly lower her head or body
- ...

What if the sentence was spoken? ("I" vs "eye")

How can the lexical, syntactic, and semantic ambiguities be resolved/disambiguated?

- tagging (part of speech) - is "duck" a noun or verb?
- parsing (syntactic structure) - is "her" part of the duck phrase?
- word sense disambiguation (lexical semantics), e.g., does "make" mean create or cook?

Disambiguation

Pick the most likely of n choices.

Tightly coupled processing

- decisions that can be made easily based on one kind of knowledge can inform the others

Loosely coupled processing

- do your best within a single kind of knowledge (maybe several)

Models and Algorithms

Models

- formalisms to represent linguistic knowledge

Algorithms

- used to manipuate the representaitons to produce desired behavior

A small set of standard models and algorithms have proven to be effective

Some Example Models

State machines

- finite state automata, finite state transducers

Formal rule systems

- context free grammars, unification grammars, probabilistic grammars

Logic-based formalism

- first order predicate calculus, higher order logics

Models of uncertainty

- Bayesian probability theory

Use of states and rules has roots in speech, morphology, syntax, while logic has roots in semantics, pragmatics, and discourse. But this is changing! Probability often augments other models.

Algorithms

Almost all of the algorithms we'll study can be viewed as transducers or parsers, i.e., algorithms that accept an input and construct some structure based on that input.

Unfortunately, since language is ambiguous at all levels this is almost never simple. This leads us to employ algorithms that fall into two related categories:

- state space search
- dynamic programming

State Space Search

The states in a search represent a pairing of partial structures with pieces of onput.

The goal is to arrive at the right/best structure after having processed all of the input.

As with most AI problems the spaces are too large and the criteria for "bestness" difficult to encode.

Dynamic Programming

Informally, in the course of searching we often run across or construct structures that must be present in all possible solutions.

Dynamic programming allows us to avoid recomputing such structures.

Have the students talking in the back row of this room ...

Can Machines Think?

The **Turing Test**: language as a test for intelligence

Three participants (2 humans, 1 computer)

- Human interrogator's goal: tell the machine and human apart
- Machine's goal: fool the interrogator into believing that a person is responding
- Other human's goal: help the interrogator reach his goal

Did any of the demos pass the Turing test?

Loebner Prize

"The first formal instantiation of a Turing test"

- http://www.loebner.net/Prizef/loebner-prize.html
- 2001 winner (http://www.alicebot.org/)

July 7, 2002: Approximating Life. The New York Times Magazine. "Each morning, he wakes before dawn and watches conversations stream by on his screen. Thousands of people flock to his Web site every day from all over the world to talk to his creation, a robot called Alice. It is the best artificial-intelligence program on the planet, a program so eerily human that some mistake it for a real person. As [Richard] Wallace listens in, they confess intimate details about their lives, their dreams; they talk to Wallace's computer about God, their jobs, Britney Spears. It is a strange kind of success: Wallace has created an artificial life form that gets along with people better than he does. ...Is she intelligent? If so, how? In 1950, the pioneering British mathematician Alan Turing grappled with this question in the journal Mind, where he first posed the 'Turing Test' – the gold standard for artificial thought. 'Can machines think?' he asked – and immediately noted that the question hinges, of course, on what 'thinking' is."

A Brief NLP History

Four paradigms: stochastic, logic-based, nlu, discourse modelling

Empiricism and finite-state models return

Recent years: strong integration of different techniques, different areas (including speech and information retrieval)

For Next Time

Read handouts (optional)

Read Chapter 2

Send me your paper preferences

Get a CS account (if you haven't already)

Send me email for the class mailing list (if you haven't already)

Fill out survey (if you haven't already)

Survey

Name:

Email:

Department:

Year:

Undergraduate Major:

Relevant Courses:

Relevant Research Experience:

Methodologies (state machines, rule systems, grammars, logic, search, probability, automata, machine learning, etc.):

Programming Languages:

Operating Systems:

Goals: