

CS 533
Natural Language Processing
Lecture 1 - January 27, 2003

Matthew Stone



Department of Computer Science
Center for Cognitive Science
Rutgers University

Outline

Administrative
Language and information
Collaborative conversation

Contact info

My web site

<http://www.cs.rutgers.edu/~mdstone>

Class web site

<http://www.cs.rutgers.edu/~mdstone/class/533>

Office hours: Wednesday 3-5.

Materials

Herbert H. Clark. *Using Language*.
Cambridge, 1996.

Daniel Jurafsky and James H. Martin. *Speech
and Language Processing*. Prentice Hall, 2000.

Additional tutorial and systems readings.

Syllabus

1. Architectures for dialogue.
2. Algorithms for linguistic structures.
3. Applications.

Syllabus

1. Architectures for dialogue.
 1. Jan 27. Collaborative conversation.
 2. Feb 3. Utterances.
 3. Feb 10. Language and the world.
 4. Feb 17. Reference and common ground.
 5. Feb 24. Goals of utterances.
 6. Mar 3. Flow of conversation.

Syllabus

1. Architectures for dialogue.
2. Algorithms for linguistic structures.
 1. Mar 10. Syntactic representations.
 2. Mar 24. Syntactic analysis and interpretation.
 3. Mar 31. Syntactic structure and generation.
 4. Apr 7. Building resources from data.

Syllabus

1. Architectures for dialogue.
2. Algorithms for linguistic structures.
3. Applications.
 1. Apr 14. Speech recognition.
 2. Apr 21. Text processing.
 3. Apr 28. Translation, information extraction.
 4. May 5. Project presentations.

Requirements

Short weekly homework.

Writing
Data analysis
Programming

Final project.

Writing
Data analysis
Programming

Language and Information

Tasks you might know about

Retrieve web pages that match a query (or scientific articles relevant to a research question, or news articles on a given topic).

Decide if an email message is spam (or a job ad, or a talk announcement, or related to class).

Highlight parts of a document that let somebody know if they want to read it (or what the thesis of the document is, or what its point of view is).

Language and information

These are information tasks

Language is written by people for people.
Computer system helps users manage language.

Linguistic knowledge tends not to be important in information tasks.

Language and information

These tasks can often be solved with a vector-space or "bag of words" model of language.

To reason about a text, list of all of the words in it and how often they occur there.

Language and information

Example.

I want a man who knows what love is all about. You are generous, kind, thoughtful. People who are not like you admit to being useless and inferior. You have ruined me for other men. I yearn for you. I have no feelings whatsoever when we're apart. I can be forever happy. Will you let me be yours? - Gloria.

Language and information

Vector space representation

I: 4
want: 1
a: 1
man: 2
who: 2
know: 1
what: 1
love: 1
be: 6
all: 1
about: 1
you: 5
...

Language and information

This model does not capture meaning.

I want a man who knows what love is. All about you are generous, kind, thoughtful people, who are not like you. Admit to being useless and inferior. You have ruined me. For other men, I yearn. For you, I have no feelings whatsoever. When we're apart, I can be forever happy. Will you let me be? Yours, Gloria.

Language and information

Vector space representation

I: 4
want: 1
a: 1
man: 2
who: 2
know: 1
what: 1
love: 1
be: 6
all: 1
about: 1
you: 5
...

Language and information

In information tasks, topic not content is what matters.

Language is written by people for people.

Computer system helps users manage language.

Language and Information

Retrieve web pages that match a query

Reuse of words from the query shows what the web page is about.

More important than content is quality – you want to get a good web page, which means link analysis, anchor text, etc.

Language and information

Decide if an email message is spam

Again, the words are a great clue – one way
(Nigeria, free, xxx,...)
or the other
(Matthew, publication, due,...)

Language and information

Displaying important sentences.

I yearn for you.

For you, I have no feelings whatsoever.

In this class, we will be looking at tasks where meaning does matter

Dialogue

Language is created by people for machines or
by machines for people.

Computer system helps users in some real-
world task.

Collaborative conversation

Sources:

Clark, *Using Language*, Chapter 1.

Stone, *Communicative Intentions*.

II. Language as Joint Activity

Utterances are real-world actions taken
publicly and collaboratively.

Meaning is the *product* of this
collaboration.

[e.g., Grice 57; Searle 69; Lewis 79; Thomason 90; Clark 96]

Understanding real-world action



attributing mental state,
intention or *commitment*,
linking action to context and goals.
[e.g., Pollack 90]

Same for language use




Implementing NLG as coordination on interpretation

Construct
 a grammatical utterance
 that achieves desired effects
that links with context so as to be understood.

Stone & Doran, ACL, 1997
 Stone & Webber, INLG, 1998
 Stone, INLG, 2000
 Cassell, Stone & Yan, INLG, 2000
 Stone, Bleam, Doran & Palmer, TAG+, 2000
 and
 Stone, Doran, Webber, Bleam & Palmer,
 Microplanning with communicative intentions,
 In submission to *Computational Intelligence*

Collaborative (dialogue) agents



This class is about the algorithms and data structures you can use in this architecture

Algorithms:

- Plan recognition
- Deliberation
- Coordinated execution

Data structures:

- Representations of communicative intentions
- Representations of task and discourse context

Task-oriented Dialogue

- Interlocutors talk something through

A: So are we all set?
 B: The vegetables are still too crunchy.
 A: The zucchini there?
 B: Yeah, the zucchini...
 A: OK, I'll take care of it.


Links to Real-world Collaboration

- Reference and attention
 - each subtask involves key objects
 - interlocutors' talk about those objects reflects progress through the task

[Grosz, 1977; Sidner, 1983; Webber, 1983]

Reference and Attention

A: So are we all set?
 B: The vegetables are still too crunchy.
 A: The zucchini there?
 B: Yeah, the zucchini...
 A: OK, I'll take care of it.



Links to Real-world Collaboration

- Dialogue structure
 - each subtask involves key steps of action and coordination
 - subtasks can structure the dialogue itself into segments with a common purpose

[Grosz & Sidner, 1986; Lochbaum, 1992]


Discourse Structure

Proceed with task
Identify next subtask

A: So are we all set?
 B: The vegetables are still too crunchy.
 A: The zucchini there?
 B: Yeah, the zucchini...

Plan next subtask

A: OK, I'll take care of it.



Links to Real-world Collaboration

- Individual utterances make progress
 - achieving the task requires specific steps
 - steps can often be mapped to high-level speech acts achieved by utterances

[Litman & Allen, 1983; Carberry, 1990]

Speech Acts in Discourse


Proceed with task

Identify next subtask

A: YNQ: So are we all set?
 B: YNQ Ans: The vegetables are still too crunchy.
 A: CLQ: The zucchini there?
 B: CLQ Ans: Yeah, the zucchini...

Plan next subtask

A: Commit: OK, I'll take care of it.



A Rich Functional Organization


Proceed with task

Identify next subtask

A: YNQ: So are we all set?
 B: YNQ Ans: The vegetables are still too crunchy.
 A: CLQ: The zucchini there?
 B: CLQ Ans: Yeah, the zucchini...

Plan next subtask

A: Commit: OK, I'll take care of it.



[cf. Grosz & Sidner, 1986; Litman & Allen, 1983; Carberry, 1990; Lochbaum, 1992]


Mediated by Specific Intentions

*Proceed with task
Identify next subtask*

A: YNQ: So are we all set?

here the dialogue is intended to continue with an answer on whether *A+B* are *done-cooking* as of *6:30pm*

Yet Made up of Complex Linguistic Structures



With Very Abstract Meanings

So are we all set?


At the current time does the group containing the speaker have the property they require to be ready for the upcoming event?

Problem of Interpretation

- The hearer observes some actions
 - meaningful elements in the utterance in a dependency tree


- **and asks**

- what could the speaker have been trying to achieve with them?




Semantics as a Constraint

- Speaker makes presuppositions
 - schematic form suggested by grammar




Context as a Constraint

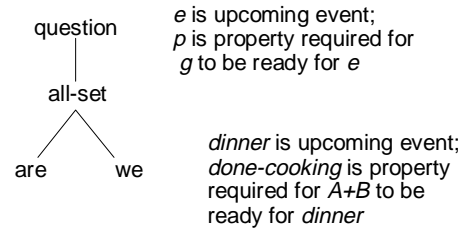
- Speaker makes presuppositions
 - specific instances found in context (or added)



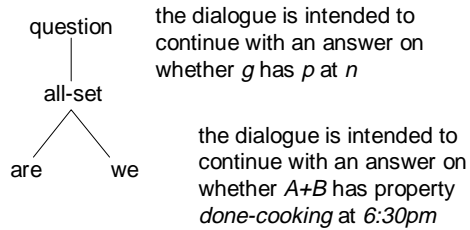
More of These Constraints



More of These Constraints

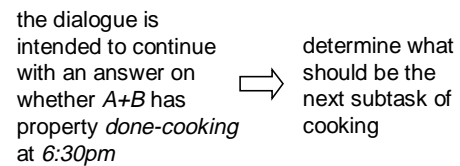


Suggests a Specific Intention



A Further Constraint


- Linking specific intent to broader goals in dialogue



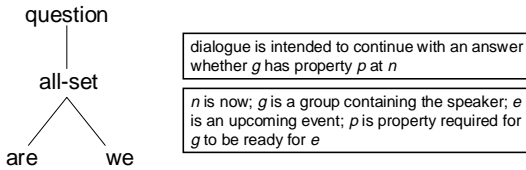
Theoretical Statement

- Hearer has a plan-recognition problem
Cf [Grice, 1957; Thomason, 1990]
- Plan is a representation laying out
 - actions to do
 - assumptions about how actions cause change
 - assumptions about what the world will be like
 - to show how to achieve intended effects.
 Cf [Bratman, 1987; Pollack, 1990]

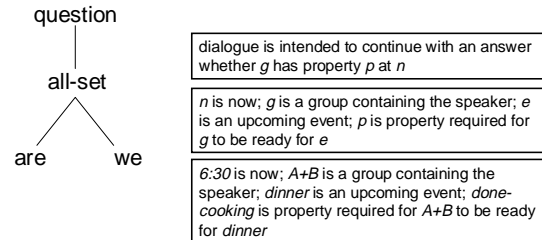
Actions to Do



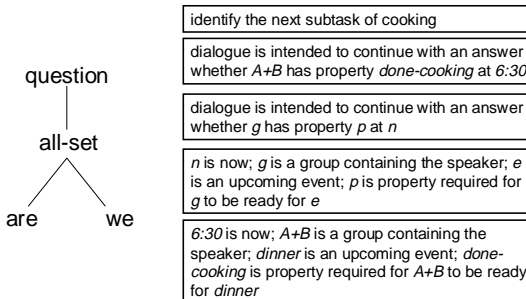
Assumptions About Change



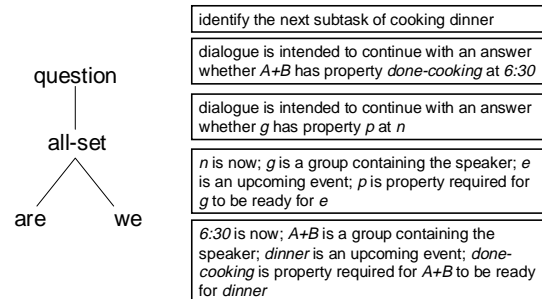
Assumptions About World



Intended Effects



Result of Interpretation



Flexible Framework

- Representation answers key questions
 - what was said
 - what did that utterance mean
 - what did the speaker take for granted
 - what did the speaker intend for the dialogue
- But allows this information to be brought together and used in many ways

Production

- The speaker starts with a desired effect
 - in the ongoing dialogue

identify the next subtask of cooking dinner

- and asks
 - how could I use an utterance in context to contribute to achieving this effect?

Start with Goal

identify the next subtask of cooking dinner

Add Elements One by One

question

identify the next subtask of cooking dinner

dialogue is intended to continue with an answer whether $A+B$ has property *done-cooking* at 6:30

dialogue is intended to continue with an answer whether g has property p at n

Add Elements One by One

question

all-set

identify the next subtask of cooking dinner

dialogue is intended to continue with an answer whether $A+B$ has property *done-cooking* at 6:30

dialogue is intended to continue with an answer whether g has property p at n

e is an upcoming event; p is property required for g to be ready for e

dinner is an upcoming event; *done-cooking* is property required for $A+B$ to be ready for *dinner*

Add Elements One by One

question

all-set

are

identify the next subtask of cooking dinner

dialogue is intended to continue with an answer whether $A+B$ has property *done-cooking* at 6:30

dialogue is intended to continue with an answer whether g has property p at n

n is now; e is an upcoming event; p is property required for g to be ready for e

6:30 is now; *dinner* is an upcoming event; *done-cooking* is property required for $A+B$ to be ready for *dinner*

Add Elements One by One

question

all-set

are

we

identify the next subtask of cooking dinner

dialogue is intended to continue with an answer whether $A+B$ has property *done-cooking* at 6:30

dialogue is intended to continue with an answer whether g has property p at n

n is now; g is a group containing the speaker; e is an upcoming event; p is property required for g to be ready for e

6:30 is now; $A+B$ is a group containing the speaker; *dinner* is an upcoming event; *done-cooking* is property required for $A+B$ to be ready for *dinner*

Until Plan Can Be Recognized

question

all-set

are

we

identify the next subtask of cooking dinner

dialogue is intended to continue with an answer whether $A+B$ has property *done-cooking* at 6:30

dialogue is intended to continue with an answer whether g has property p at n

n is now; g is a group containing the speaker; e is an upcoming event; p is property required for g to be ready for e

6:30 is now; $A+B$ is a group containing the speaker; *dinner* is an upcoming event; *done-cooking* is property required for $A+B$ to be ready for *dinner*