

The Computer and Natural Language (Ling 445/515) Dialogue Systems

Markus Dickinson

Department of Linguistics, Indiana University
Autumn 2011

What are dialogue systems good for?

Anytime we have a straightforward task, dialogue systems seem like a good idea:

- ▶ We can book airline tickets over the phone without dealing with error-prone humans.
- ▶ We can interact with our computer while keeping our hands free.
- ▶ We can talk to a system that won't have the same prejudices a human might (... depending on the programmer)

Before we look at dialogue systems, though, we need to understand something about human dialogue

Introduction

Why dialogue?

Human dialogue

Basic facts
Dialogue moves
Speech acts
Rules

Early dialogue systems

ELIZA
Other chatterbots

Modern dialogue systems

Conversations

The speaker's perspective

Conversations work because speakers & listeners understand a lot about dialogue

1. **Why to speak:** both assume a common goal
2. **What to say:** both can understand direct & indirect statements, requests, questions, etc.
3. **How to say it:** both properly handle intonational changes, politeness conventions, etc.

Introduction

Why dialogue?

Human dialogue

- Basic facts
- Dialogue moves
- Speech acts
- Rules

Early dialogue systems

- ELIZA
- Other chatterbots

Modern dialogue systems

From the the listener's perspective, they must make sense of what is being said by the speaker, by:

- 1. Accurately noticing what the speaker said.**
 - ▶ This involves hearing what was said, then breaking up the speech signal into words and phrases
- 2. Working out the literal meaning of what was said.**
- 3. Working out the reason that the speaker did what they did.**

In order to do this, the listener has to reason about the speaker's intent, on the basis of the observed evidence

- ▶ A first step toward building a computer that can do the same is to precisely describe how this process works in human-human dialog

Introduction

Why dialogue?

Human dialogue

- Basic facts
- Dialogue moves
- Speech acts
- Rules

Early dialogue systems

- ELIZA
- Other chatterbots

Modern dialogue systems

Communicating is a collaborative event

The speaker and the listener must manage the following processes, in real-time:

1. Speaking and listening in an orderly way, without talking over each other more than is acceptable.
2. Checking that the other person is understanding.
3. Fixing confusions and misunderstandings before they cause the conversation to break down.

Human beings are skilled in this activity long before they enter pre-school

- ▶ Trying to involve a computer in dialog makes the difficulties more evident

Introduction

Why dialogue?

Human dialogue

Basic facts
Dialogue moves
Speech acts
Rules

Early dialogue systems

ELIZA
Other chatterbots

Modern dialogue systems

Why dialogues happen?

One of the main reasons for using language is in order to collaborate with others on getting things done.

- ▶ If people have everything they need, or can get it without asking for help, there's no real reason to speak
- ▶ More often, people and computer systems need to work with others in order to achieve their goals

Dialog is full of rules and conventions that have evolved over the years in order to make communication run smoothly

- ▶ A question is usually followed by an answer
- ▶ A request is often made *indirectly* because outright demands feel rude
- ▶ A refusal is often heavily disguised so as to soften the impact of the unexpected and unwelcome response

A sample dialog from the Let's Go system

(<http://www.speech.cs.cmu.edu/letsgo/example.html>)

- s: Thank you for calling the CMU Let's Go! Bus Information System. This conversation will be recorded for research purposes. If you do not want to be recorded, please hang up now. What can I do for you?
- u: I would like to go to the airport tomorrow morning.
- s: To the airport. When do you want to arrive?
- u: I'd like to arrive at 10:30.
- s: Arriving at around 10:30 AM. Where do you want to leave from?
- u: I'd like to leave from Carnegie Mellon.
- s: From Carnegie Mellon. There is a 28X leaving Forbes Avenue at Morewood at 9:20 AM. It will arrive at Greater Pittsburgh Airport at Departure Stops at 10:15 AM.

Automating dialogue (2)

- u: When's the next bus?
- s: There is a 28X leaving Forbes Avenue at Morewood at 9:40 AM. It will arrive at Greater Pittsburgh Airport at Departure Stops at 10:35 AM.
- u: I'd like to go from the airport back to CMU on Friday.
- s: From the airport to CMU. When would you like to leave?
- u: At 6:15 PM.
- s: Leaving at around 6:15 PM. There's a 28X leaving Greater Pittsburgh Airport at Departure Stops at 6:26 PM. It will arrive at Forbes Avenue at Morewood at 7:20 PM.
- u: Thank you very much.

How does dialogue, or **discourse** (= conversation), work?

We will look at a few aspects of human-human interactions:

- ▶ Basic facts about dialogues
- ▶ Organization of conversation: Dialogue moves
- ▶ What utterances do: Speech acts
- ▶ Rules of conversation: Gricean maxims

Basic facts about dialogues

Utterances

People do not necessarily speak in sentences

- ▶ **Utterances:** basic unit of conversation
 - ▶ may span over several turns
 - ▶ may have several utterances within one turn

AGENT: Yeah yeah the um let me see here we've got you on American flight nine thirty eight

CUSTOMER: Yep.

AGENT: leaving on the twentieth of June out of Orange County John Wayne Airport at seven thirty p.m.

CUSTOMER: Seven thirty.

AGENT: and into uh San Francisco at eight fifty seven.

Basic facts about dialogues

Taking turns

How humans communicate is relatively systematic

- ▶ **Turn-taking:** must know when it's the right time to contribute your turn
 - ▶ Comes naturally to humans: overlaps and long pauses are actually somewhat rare
- ▶ **Adjacency pairs:** two-part conversational structures where turn-taking is usually quite clear
 - ▶ question/answer: question is looking for an answer from dialogue partner
 - ▶ greeting/greeting
 - ▶ request/grant
 - ▶ etc.

Basic facts about dialogues

Common ground

Common ground: the set of things which both speakers believe to be true of the conversation.

- ▶ Part of classroom common ground is that we are at Indiana in a linguistics class. Not part of the common ground is what I had for lunch.
- ▶ People assess what others know before making a new contribution

In order to establish common ground, speakers do various things:

- ▶ **Backchannels** = short utterance which indicates the utterance was heard and that the speaker should continue.

A: *That's a nice shirt.*

B: *Mm-hmm.*

- ▶ Acknowledgment of utterance: continued attention, completing speaker's utterance, and so on
- ▶ Asking clarifying questions

Can view dialogue as a series of moves, as in a game

1. As an exchange has to start somewhere, we can group together a set of **initiating moves**
 - ▶ Making an assertion (“you have a friendly face”)
 - ▶ Issuing a command or making a request (“let’s go and have a bite to eat”)
 - ▶ Asking a question (“are you doing anything tonight?”)

2. Some moves are **responses** to the previous move

- ▶ Saying “yes” or “no” to a question where the person was seeking information
- ▶ Giving an answer to an information-seeking question that needs more than “yes” or “no”
- ▶ Answering the question, but then providing more information than was strictly asked for
- ▶ Agreeing to do something (e.g., saying “ok” to “let’s go and have a bite to eat”)
- ▶ Refusing to do something (e.g., saying “no” to “let’s go and have a bite to eat”)
- ▶ Maybe agreeing to do something (e.g., saying “maybe”)
- ▶ Partially agreeing to do something (e.g., saying “ok, if we’re quick”)
- ▶ Partially refusing to do something (e.g., saying “no, unless you pay”)

3. Others are responses, more or less, but divert the conversation from its expected path—**dialog management moves**
- ▶ Saying “huh?” when you didn’t hear.
 - ▶ Saying “I don’t think I understand” in response to something you didn’t understand
 - ▶ Saying “I can’t believe you said that” when you did hear what the other person said, find it offensive or unacceptable, and want them to retract it
 - ▶ Saying “I take it back” when you want the other person to pretend that you haven’t said what you just did

Dialogue moves

Discourse structure

Discourse, or dialogue, has some structure to it—not just a bunch of random, loosely-connected statements.

- ▶ Every discourse has a **discourse purpose** = a reason for talking (Grosz and Sidner 1986)
- ▶ Purpose is clearer in **task-based dialogues** = speakers are trying to achieve some goal together
e.g., giving directions or fixing a car

If we figure out the discourse structure automatically, the computer can keep track of what's going on.

- ▶ Know what topics are appropriate
- ▶ Better able to know what pronouns and definite noun phrases (e.g., *the dog*) refer to.

What does each utterance do? How does it function?

- ▶ Utterances often equated with **actions** (Austin 1962)

For example, there are **performative verbs**; by saying them, you actually do what you say

- ▶ I (hereby) christen this ship *The Swarthy*.
- ▶ I pronounce you man and wife.
- ▶ I second that motion.
- ▶ I bet you five dollars the Hoosiers will win the NCAA championship this year.

Speech acts (2)

Speech acts help us to understand what is really going on when people say, e.g., “could you pass the salt?”

- ▶ Syntactic form of a question
- ▶ In many contexts, the speech act that it implements is a polite command

For dialogue systems, importance is to recognize that utterances may not be exactly what they seem

- ▶ Some standard patterns:
 - ▶ use of questions to stand for commands
 - ▶ use of objectively unnecessary, informative statements to stand for requests (e.g., using “you are blocking my view” vs. “get out of my eyeline”)

Common speech acts: informing, reminding, inviting, complaining, refusing, accepting, requesting, commanding

Speech acts (3)

Preconditions & effects

Speech acts can be broken down into **preconditions** and **effects**

e.g., informing someone that a concert is on Friday

- ▶ Preconditions:
 - ▶ actually know that the concert is on Friday
 - ▶ the other person does not know when the concert is
- ▶ Effect: the other person knows that the concert is on Friday

Automated systems often explicitly represent speech acts in terms of preconditions & effects

Speech acts can be **direct** or **indirect**

- ▶ Direct: Form matches meaning.

A: *Please take out the garbage.*

imperative structure & imperative meaning

- ▶ Indirect: The form and the meaning are different.

A: *The garbage isn't out yet.*

declarative structure but imperative/directive meaning

A: *Could you take out the garbage?*

interrogative structure but imperative/request meaning

⇒ How do we encode such distinctions in a computer?

People have marked up text with different labels for speech acts.

The **Dialogue Act Markup in Several Layers (DAMSL)** is a scheme of speech acts people use to annotate texts

- ▶ **forward-looking functions:** look for a response
 - ▶ question, request, etc.
- ▶ **backward-looking functions:** provide a response
 - ▶ answer, agreement, etc.

Note the resemblance to adjacency pairs.

Introduction

Why dialogue?

Human dialogue

Basic facts

Dialogue moves

Speech acts

Rules

Early dialogue
systems

ELIZA

Other chatterbots

Modern dialogue
systems

Using DAMSL in a computer

Using the DAMSL scheme,

- ▶ People annotate a text, or a **corpus**
- ▶ People can train a system on it and attempt to mark a new text with these labels.
- ▶ People can use the functions to tell a dialogue system what to do next
 - ▶ The system keeps track of the previous forward-looking function to know how to respond.
 - ▶ If a dialogue system needs an answer, it knows what kind of forward-looking function to use to get the human to provide one

But if a dialogue system wants to know what to do next, it needs more than forward and backward-looking functions ...

Introduction

Why dialogue?

Human dialogue

Basic facts

Dialogue moves

Speech acts

Rules

Early dialogue
systems

ELIZA

Other chatterbots

Modern dialogue
systems

Rules of Conversation: Gricean Maxims

When people speak “indirectly,” it’s usually clear what they mean

⇒ They obey what Grice (1975) referred to as the four **maxims** (rules) or conversation

- ▶ All based on the **cooperative principle**
 - ▶ speakers are both trying to contribute to the purposes of the conversation.
- ▶ We use these to infer what a speaker is really saying

Quantity: Be exactly as informative as is required.

- ▶ Make your contribution as informative as is required (for the current purposes of the exchange)

A: *How many pencils do you have?*

B: *Two.*

→ means exactly two (or would have said more)

- ▶ Do not make your contribution more informative than is required.

Violation:

(passing by somebody on the way to class)

A: *How's it going?*

B *It's complicated. Yesterday I was sick ...*

Quality: Try to make your contribution one that is true.

- ▶ Do not say what you believe to be false
- ▶ Do not say that for which you lack adequate evidence

Violation:

A: *Do you know how to drive a stick-shift?*

B: *Yes, I do. I've seen my dad do it many times.*

Gricean maxim 3: Relevance

Relevance: Be relevant.

A: *Is Gail dating anyone these days?*

B: *Well, she goes to Indianapolis every weekend.*

We make an inference that Gail is dating somebody in Indianapolis

- ▶ Otherwise, B's statement doesn't make much sense.

Gricean maxim 4: Manner

Manner: Be perspicuous (easy to understand).

- ▶ Avoid obscurity of expression
 - ▶ In other words, avoid jargon when it's not a part of the common ground, e.g., words like *prolixity*
- ▶ Avoid ambiguity
- ▶ Be brief (avoid unnecessary prolixity)
 - ▶ Similar to maxim of quantity, but think here of long-winded conversations that provide the same amount of information as a shorter one
- ▶ Be orderly

We'll start our tour of early dialogue systems with a quick tour of **ELIZA**

- ▶ Written by Joseph Weizenbaum in the mid-1960s
- ▶ Simulation of a (Rogerian) psychotherapist.
 - ▶ Rogerian therapist asks questions and gives backchannels more than providing advice
- ▶ User types questions, and ELIZA responds.

For some, the test of such a system is how human-like it is:

- ▶ If I can tell I am interacting with a computer, it has failed the **Turing Test**

Example conversation

Taken from Weizenbaum's 1966 paper:

USER: Men are all alike.

ELIZA: IN WHAT WAY?

USER: They're always bugging us about something or other.

ELIZA: CAN YOU THINK OF A SPECIFIC EXAMPLE?

USER: Well, my boyfriend made me come here.

ELIZA: YOUR BOYFRIEND MADE YOU COME HERE

USER: He says I'm depressed much of the time.

ELIZA: I AM SORRY TO HEAR YOU ARE DEPRESSED

USER: It's true. I am unhappy

ELIZA: DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY

Example conversation (cont.)

USER: I need some help, that much seems certain.

ELIZA: WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP

USER: Perhaps I could learn to get along with my mother.

ELIZA: TELL ME MORE ABOUT YOUR FAMILY

USER: My mother takes care of me.

ELIZA: WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU

USER: My father.

ELIZA: YOUR FATHER

USER: You are like my father in some ways.

ELIZA: WHAT RESEMBLANCE DO YOU SEE

Example conversation (cont.)

USER: You are not very aggressive but I think you don't want me to notice that.

ELIZA: WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE

USER: You don't argue with me.

ELIZA: WHY DO YOU THINK I DON'T ARGUE WITH YOU

USER: You are afraid of me.

ELIZA: DOES IT PLEASE YOU TO BELIEVE I AM AFRAID OF YOU

USER: My father is afraid of everybody.

ELIZA: WHAT ELSE COMES TO YOUR MIND WHEN YOU THINK OF YOUR FATHER

USER: Bullies.

ELIZA: DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR BOYFRIEND MADE YOU COME HERE

How ELIZA works

There are many different **implementations** of ELIZA out there, but they all work in the same general way:

1. Transform input to change the point of view (*I* and *you*)
2. Perform **pattern recognition** (pattern matching)
3. Substitute keywords into **canned phrases**, or submit a general canned phrase

They can also more or less include some history of the conversation

How can we tell that ELIZA isn't a human?

A number of clues tell us that ELIZA is a computer:

- ▶ Lack of initiative
- ▶ Lack of knowledge
ELIZA knows nothing about the real world
- ▶ Lack of common sense
- ▶ Lack of true grammar/**syntax** = structure of sentences
 - ▶ Say *you peabrain* to it and you might get ME PEABRAIN back.
 - ▶ Need two separate patterns for *Computers worry me very little.* and *I'm not worried much by computers.*
- ▶ Doesn't have (much of) a memory

ELIZA also started a whole phenomenon of **chatterbots**

PARRY (Colby et al 1971)

- ▶ More verbal than ELIZA in some ways, i.e., talks for longer durations
- ▶ keeps track of “global” emotional state
 - ▶ when the *anger* variable is high, for example, PARRY will choose from a set of “hostile” options.

“With Parry, Dr. Colby established that a computer chip could be programmed to imitate a paranoid schizophrenic.”

<http://www.edu-cyberpg.com/Linguistics/Parry.html>

ALICE is a more advanced chatterbot, which won the 2004 Loebner Prize

- ▶ Loebner Prize = instantiation of the Turing Test:
\$100,000 for the first computer to pass the Turing Test
(hasn't happened yet)
- ▶ Each year, the most human-like computer wins \$2000

You can chat with ALICE at: <http://alice.pandorabots.com/>

All of these chatterbots are just pattern-matchers, albeit of varying complexity

- ▶ There has been a push to add linguistic and real-world knowledge to dialogue systems
- ▶ Which is why we spent so much time talking about how real dialogue works

Modern dialogue systems

We'll look at Jason Baldridge's slides from here on out ...