

Dialogue and Conversational Agents

Ling575
Spoken Dialog Systems
March 31, 2016

Roadmap

- Dialog and Dialog Systems
- Facets of Conversation:
 - Turn-taking
 - Speech Acts
 - Cooperativity
 - Grounding
- Spoken Dialogue Systems:
 - Pipeline Architecture
 - Finite-State, Frame-based, Information State Systems
 - Evaluation

Dialog Example

C₁: ...I need to travel in May.

A₁: And, what day in May did you want to travel?

C₂: OK uh I need to be there for a meeting that's from the 12th to the 15th.

A₂: And you're flying into what city?

C₃: Seattle.

A₃: And what time would you like to leave Pittsburgh?

C₄: Uh hmm I don't think there's many options for non-stop.

A₄: Right. There's three non-stops today.

C₅: What are they?

A₅: The first one departs PGH at 10:00am arrives Seattle at 12:05 their time.
The second flight departs PGH at 5:55pm, arrives Seattle at 8pm. And the
last flight departs PGH at 8:15pm arrives Seattle at 10:28pm.

C₆: OK I'll take the 5ish flight on the night before on the 11th.

A₆: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air
flight 115.

C₇: OK.

Travel Planning: Let's Go!

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.

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.

AT&T's How May I Help You?

S: How may I help you?
U: *can you tell me how much it is to tokyo?*
S: You want to know the cost of a call?
U: *yes that's right*
S: Please hold on for rate information.

S: How may I help you?
U: *what's the area code for chicago?*
S: You want an area code.
U: *(silence)*
S: The area code for Chicago is 312.

Figure 19.2 The call routing/telephone information domain: two actual conversations between a user (U) and the AT&T How May I Help You (HMIHY) system (S) (Gorin et al., 1997).

ItSpoke Tutoring System

S: I didn't understand. What force acts on all objects within close proximity to earth?

U: *uh the force of gravity*

S: Fine. Besides the Earth's gravitational force, are there any other forces acting on the person, keys, and elevator after the release?

U: *no*

S: Yep. Thus, keep in mind that after the release, the only force on the person, keys, and elevator is the force of gravity. If gravity is the only force acting on an object, will it be moving or staying still?

U: *it will be moving*

S: Fine. Not only are the person, keys, and elevator moving, they have only gravitational forces on them. When an object is falling and has only gravitational force on it, it is said to be in what?

U: *freefall*

Figure 19.3 The tutoring domain: part of an actual conversation between a student user (U) and the ITSPoke system (S) of (Litman and Silliman, 2004), based on the Why2-Atlas text-based tutoring system (?).

Dialogue is Different

- Two or more speakers
 - Primary focus on speech
- Issues in multi-party spoken dialogue
 - Turn-taking – who speaks next, when?
 - Collaboration – clarification, feedback,...
 - Disfluencies
 - Adjacency pairs, dialogue acts

Conversations and Conversational Agents

- Conversation:
 - First and often most common form of language use
 - Context of language learning and use
 - Goal:
 - Describe, characterize spoken interaction
 - Enable automatic recognition, understanding
- Conversational agents:
 - Spoken dialog systems, spoken language systems
 - Interact with users through speech
 - Tasks: travel arrangements, call routing, planning

Conversation

- Intricate, joint activity
 - Constructed from consecutive turns
 - Joint activity between speakers, hearer
 - Involves inferences about intended meaning
- SDS: simpler, but hopefully consistent

Turn-Taking

- Multi-party discourse
 - Need to trade off speaker/hearer roles
 - Interpret reference from sequential utterances
- When?
 - End of sentence?
 - No: multi-utterance turns
 - Silence?
 - No: little silence in smooth dialogue: < 250ms
 - Gaps less than actual sentence planning time - anticipate
 - When other starts speaking?
 - No: relatively little overlap face-to-face: ~5%

Turn-taking: Who & How

- At each TRP in each turn (Sacks 1974)
 - If speaker has selected A to speak, A must take floor
 - If speaker has selected no one to speak, anyone can
 - If no one else takes the turn, the speaker can
- Selecting speaker A:
 - By explicit/implicit mention: What about it, Bob?
 - By gaze, function
- Selecting others: questions, greetings, closing
 - (Traum et al., 2003)

Turns and Structure

- Some utterances select others:
 - Adjacency pairs:
 - Greeting – Greeting, Question – Answer,
 - Compliment – Downplayer
 - Silence ‘dispreferred’ within adjacency pair
 - A: Is there something bothering you or not?
 - (1.0)
 - A: Yes or No?
 - (1.5)
 - A: Eh.
 - B: No.

Turn-taking in HCI

- Human turn end:
 - Detected by 250ms (or longer) silence
- System turn end:
 - Signaled by end of speech
 - Indicated by any human sound
 - Barge-in
- Continued attention:
 - No signal
- Design problems create ambiguous silences
 - Problematic for SDS users
 - (Stifelman et al., 1993), (Yankelovich et al, 1995)

Speech Acts

- Utterance:
 - Action performed by the speaker (Austin, 1962)
 - Performatives: *name, second*
 - *I name this ship the Titanic.*
 - *I second that motion.*
 - Extend to all utterances

Utterances as 3 Act Types

- Locutionary act:
 - utterance with some meaning
 - *“You can’t do that!”*
- Illocutionary act:
 - Act of asking, promising, answering, in utterance
 - *Protesting*
- Perlocutionary act:
 - Production of effects on feeling, beliefs of addressee
 - *Intend to prevent doing some action*
- Types: assertives, directives, commissives, expressives, declarations

The 3 levels of act revisited

	Locutionary Force	Illocutionary Force	Perlocutionary Force
Can I have the rest of your sandwich?	Question	Request	Intent: You give me sandwich
I want the rest of your sandwich	Declarative	Request	Intent: You give me sandwich
Give me your sandwich!	Imperative	Request	Intent: You give me sandwich

Collaborative Communication

- Speaker tries to establish and add to
 - “common ground” – “mutual belief”
- Presumed a joint, collaborative activity
 - Make sure “mutually believe” the same thing
- Hearer must ‘ground’ speaker’s utterances
 - Indicate heard and understood

Closure

- Principle of closure:
 - Agents performing an action require evidence of successful performance
 - Also important to indicate failure or understanding
- Non-speech closure:
 - Push elevator button → Light turns on
- Two step process:
 - Presentation (speaker)
 - Acceptance (listener)

Degrees of Grounding

- Weakest to strongest
- Continued attention:
 - Silence implies consent
- Next relevant contribution
- Acknowledgment:
 - Minimal response, continuer: *yeah, uh-huh, okay; great*
- Demonstrate:
 - Indicate understanding by reformulation, completion
- Display:
 - Repeat all or part

Dialog Example

C₁: ...I need to travel in May.

A₁: And, what day in May did you want to travel?

C₂: OK uh I need to be there for a meeting that's from the 12th to the 15th.

A₂: And you're flying into what city?

C₃: Seattle.

A₃: And what time would you like to leave Pittsburgh?

C₄: Uh hmm I don't think there's many options for non-stop.

A₄: Right. There's three non-stops today.

C₅: What are they?

A₅: The first one departs PGH at 10:00am arrives Seattle at 12:05 their time.
The second flight departs PGH at 5:55pm, arrives Seattle at 8pm. And the
last flight departs PGH at 8:15pm arrives Seattle at 10:28pm.

C₆: OK I'll take the 5ish flight on the night before on the 11th.

A₆: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air
flight 115.

C₇: OK.

Grounding

- Display:
 - C: I need to travel in May.
 - A: *And what day* in May did you want to travel?
- Acknowledgment + Next relevant contribution:
 - *And what day* in May did you want to travel?
 - *And you are flying into what city?*
 - *And what time* would you like to leave Pittsburgh?

Travel Planning

S1: Hello. You've reached the [Communicator]. Tell me your full name
U2: *hi i'd like to fly to seattle tuesday morning*
S3: Traveling to Seattle on Tuesday, August eleventh in the morning. Your full name?
U4: *[John Doe]*
S5: Hello, [John], I'm glad to hear from you again. You can now ask me for help.
S6: OK. Let's see I have a United flight departing Pittsburgh at ten forty five a.m. arriving into Seattle at one fifty seven p.m. Is that OK?
U7: *Yes*
S8: Will you return to Pittsburgh from Seattle?
U9: *no i will go to Los Angeles*
S10: Traveling to Los Angeles. on what date
U11: *Sorry can we go back to talk about the previous leg again*
S12: For the first leg, you are on a United flight 1446 departing Pittsburgh at ten forty five A.M. tomorrow, arriving in Seattle at one fifty seven P.M. Do you want to take that?
U13: *actually i want to leave on wednesday*
S14: a flight on wednesday.
S15: I have a flight departing Pittsburgh at ten forty five a.m., arrives Seattle at one fifty seven p.m. Is that OK?
U16: *Yes*

Figure 19.1 The travel domain: a fragment from a successful conversation between a user (U) and the Communicator system (S) of Xu and Rudnicky (2000).