Computational Semantics and Pragmatics

Raquel Fernández

Institute for Logic, Language & Computation
University of Amsterdam

Autumn 2016

Overview

- timing coordination – turn taking
- *meaning* coordination – dialogue acts and grounding
- *style* coordination - alignment and adaptation
- language *acquisition* in interaction

Communication in Dialogue

Two views of communication:

- Shannon (1948) - Information theory: information encoded by the sender, transmitted, and decoded by the recipient.
- Grice (1957) - human communication is characterised by the process of *intention recognition*
  - speech acts / dialogue acts / moves encapsulate intention
  - intention is not fully determined by linguistic form

Goals and intentions beyond language

We have a strong tendency to ascribe *goals* and *intentions* to agents.
Related to

- theory of mind: ability to model internal mental state of agents
- attribution of causation

F. Heider & M. Simmel, (1944) An experimental study in apparent behavior. *The American Journal of Psychology*, 57.

▸ original video ▸ newer rendering

A. Michotte. (1962) The perception of causality. Methuen, Andover, MA. ▸ video

Sensing actions by others triggers attribution of intentions, goals, causes.

Speech act theory: conversations are made up of *linguistic actions*.

Speech Act Theory

Initiated by Austin (*How to do things with words*) and developed by Searle in the 60s-70s within philosophy of language.

Speech act theory grows out of the following observations:

- Typically, the meaning of a sentence is taken to be its truth value.
- There are utterances for which it doesn't makes sense to say whether they are true or false, e.g., (2)-(5):

(1) The director bought a new car this year.
(2) I apologize for being late.
(3) I promise to come to your talk tomorrow afternoon.
(4) Put the car in the garage, please.
(5) Is she a vegetarian?

- These (and generally all) utterances serve to *perform actions*.
- This is an aspect of meaning that cannot be captured in terms of truth-conditional semantics ($\rightsquigarrow$ *felicity conditions*).

Types of Acts

Austin identifies three types of acts that are performed simultaneously:

- *locutionary act*: basic act of speaking, of uttering a linguistic expression with a particular phonetics/phonology, morphology, syntax, and semantics.
- *illocutionary act*: the kind of action the speaker intends to accomplish, e.g. *blaming, asking, thanking, joking,...*
  - these functions are commonly referred to as the illocutionary force of an utterance $\rightsquigarrow$ its *speech act*.
- *perlocutionary act*: the act(s) that derive from the locution and illocution of an utterance (effects produced on the audience)

John Austin (1962), *How to do things with words*, Oxford: Clarendon Press.

Types of Illocutionary Acts

Searle distinguished between five basic types of speech acts:

- *Representatives*: the speaker is committed to the truth of the expressed proposition (assert, inform)
- *Directives*: the speaker intends to ellicit a particular action from the hearer (request, order, advice)
- *Commissives*: the speaker is committed to some future action (promise, oaths, vows)
- *Expressives*: the speaker expresses an attitude or emotion towards the proposition (congratulations, excuses, thanks)
- *Declarations*: the speaker changes the reality in accord with the proposition of the declaration (provided certain conventions hold), e.g. baptisms, pronouncing someone guilty.

John Searle (1975), *The Classification of Illocutionary Acts*, Language in Society.

From speech acts to dialogue moves

Dialogue acts (term introduced by Bunt, 1994):
Coherence and cohesion:

- inspired by dynamic semantics: moves as context-change actions (several semantic/pragmatic formal frameworks: QUD, SDRT, ...)
- structure: forward-looking and backward-looking acts

```
Waitress: What'll ya have girls?
Customer: What's the soup of the day?
Waitress: Clam chowder.
Customer: I'll have a bowl of clam chowder.
```

- *adjacency pairs*: not strict adjacency but expectation.
  - given the first part of a pair, the second part is immediately relevant and expected (*preferred* and *dispreferred* second parts)
  - intervening turns perceived as *insertion sequence* or *sub-dialogue*

Meta-communication: [more on this in the next lecture]

Bunt, H. (1994), Context and dialogue control, *Think Quarterly*, 3:19–31.
Schegloff (1972), Sequencing in conversational openings, in *Directions in Sociolinguistics*.

Dialogue Act Taxonomies: DAMSL

DA taxonomies aim to be effective as tagsets for annotating dialogue corpora.

One of the most influential DA taxonomies is the *DAMSL* schema (Dialogue Act Markup in Several Layers) by Core & Allen (1997).

- Communicative Status
- Information Level
- Forward-looking Function
- Backward-looking Function

▸ DAMSL annotation manual

The taxonomy is meant to be general but not totally domain independent $\rightsquigarrow$ it has been adapted to several types of dialogue.

DA Taxonomies: SWBD DAMSL

The SWBD DAMSL schema is a version of DAMSL created to annotate the Switchboard corpus. Here are the 18 most frequent DA in the corpus:

Tag	Example	Count	%
Statement	*Me, I'm in the legal department.*	72,824	36%
Continuer	*Uh-huh.*	37,096	19%
Opinion	*I think it's great*	25,197	13%
Agree/Accept	*That's exactly it.*	10,820	5%
Abandoned/Turn-Exit	*So, -/*	10,569	5%
Appreciation	*I can imagine.*	4,633	2%
Yes-No-Question	*Do you have to have any special training*	4,624	2%
Non-verbal	*<Laughter>,<Throat_clearing>*	3,548	2%
Yes answers	*Yes.*	2,934	1%
Conventional-closing	*Well, it's been nice talking to you.*	2,486	1%
Uninterpretable	*But, uh, yeah*	2,158	1%
Wh-Question	*Well, how old are you?*	1,911	1%
No answers	*No.*	1,340	1%
Response Ack	*Oh, okay.*	1,277	1%
Hedge	*I don't know if I'm making any sense*	1,182	1%
Declarative Question	*So you can afford to get a house?*	1,174	1%
Other	*Well give me a break, you know.*	1,074	1%
Backchannel-Question	*Is that right?*	1,019	1%

The average conversation consists of 144 turns, 271 utterances, and took 28 min. to annotate. The inter-annotator agreement was 84% (κ=.80).

▸ SWBD annotation manual

Indeterminacy

On the Gricean view, it is possible for the same signal to correspond to different intentions:

The gun is loaded $\rightsquigarrow$ *threatening? warning? explaining?*

Conversely, the same intention can be realised by different signals:

Requesting:

- A day return ticket to Utrecht, please.
- Can you please give me a day return ticket to Utrecht?
- I would like a day return ticket to Utrecht.

$\rightsquigarrow$ *How do we map from utterances to dialogue acts?*

DA Recognition

Two computational models of the interpretation of dialogue acts:

- *Symbolic models*: based on epistemic logic (beliefs, desires, and intentions - BDI); use of logical inference to reason about the speaker's intentions.
- *Probabilistic models*: the surface form of the sentence is seen as a set of cues to the speaker's intentions; use of probabilistic machine learning models.

Both models use a kind of inference: the hearer infers something that was not contained directly in the semantics of the utterance.

Daniel Jurafsky (2004) Pragmatics and Computational Linguistics. *Handbook of Pragmatics*. Oxford: Blackwell.

Symbolic Models

Classic symbolic models of dialogue acts aim to explain *indirect speech acts*

Can you pass me the salt?
⇝ Literal speech act [literal force hypothesis]: *yes-no question*
⇝ Indirect speech act after an inference chain: *request* (pass me the salt)

- S is cooperative, thus U has some aim
- S already knows the answer to the explicit question
- thus S must intend something other than asking
- ability to do something is a pre-condition for requesting
- therefore, given the context, S is probably *requesting* me to pass her the salt.

The *BDI approach* is meant to be a general model of rational action that can be applied to conversation:

- what motivates our actions
- how to understand actions by others

Symbolic Models

BDI approaches have been used as the basis to implement conversational agents in the TRAINS/TRIPS projects.

- see the project's website for access to a dialogue corpus collected to develop the system, movies of the system in action, and links to publications. http://www.cs.rochester.edu/research/trains/

Allen et al. (2001) Towards Conversational Human-Computer Interaction, *AI Magazine*.
Allen et al. (2001) An architecture for more realistic conversational systems, in *Proc. of Intelligent User Interfaces*.

Probabilistic Models

Intuition behind probabilistic models: the listener uses cues in the input to infer a particular interpretation.

Probabilistic models are typically trained on dialogue corpora annotated with dialogue acts (like Switchboard).

Given the observed cues c, the goal is to find the DA d^* that has the maximum posterior probability $P(d|c)$ given those cues.

$$d^* = \underset{d}{\operatorname{argmax}}\, P(d|c) = \underset{d}{\operatorname{argmax}}\, P(d)P(c|d)$$

We need to choose the DA that maximises the product of two probabilities: the prior probability of a DA $P(d)$ and the likelihood $P(c|d)$ of observing a particular combination of features when a particular DA is present.

Daniel Jurafsky (2004) Pragmatics and Computational Linguistics. *Handbook of Pragmatics*. Oxford: Blackwell.

Probabilistic Models

Use of several sources of knowledge:

- *Lexical and Syntactic Cues*: words/phrases that occur more often in particular DAs. presence of particular words, such as *'please'* (requests), word order (questions), tag particle *'right?'* in final position (declarative questions or checks)
- *Prosodic Cues*: final pitch rise (polar questions and declarative questions); loudness or stress can help distinguish *'yeah'* agreement from backchannel.
- *Conversational Structure Cues*: *'No it isn't'* is an agreement after *'It isn't raining'* and a disagreement after *'It is raining'*. *'yeah'* is more likely to be an agreement after a proposal. ($\rightsquigarrow$ adjacency pairs)

Stolcke et al. (2000) Dialogue Act Modeling for Automatic Tagging and Recognition of Conversational Speech, *Computational Linguistics*, 26(3).

Some recent probabilistic models try to bypass feature engineering:

Nal Kalchbrenner & Phil Blunsom, Recurrent Convolutional Neural Networks for Discourse Compositionality, *CVSC Workshop at ACL*, Sofia, Bulgaria, 2013.

Probabilistic Models

Papers to read: come prepared to explain and discuss.

- Friday 23 Sept:

> Dmitrijs Milajevs & Matthew Purver. Investigating the Contribution of Distributional Semantic Information for Dialogue Act Classification. *CVSC Workshop at EACL*, Gothenburg, Sweden, 2014.

- Tuesday 27 Sept:

> Manuvinakurike et al., Toward Incremental Dialogue Act Segmentation in Fast-Paced Interactive Dialogue Systems. *SIGdial*, Los Angeles, 2016.

It is OK to not understand everything in a paper.

- Consider goal, motivation, methods, results, implications, limitations
- Pay attention to style and structure.

Summary

Today:

- conversation as intention recognition
- from speech acts to dialogue acts
- dialogue act taxonomies
- dialogue act recognition

Friday:

- discussion of Milajevs & Purver (2014) on DA recognition
- dialogue as joint action, the grounding process

Assignment 2:

- available today and due on Tuesday 27 Sept at 23:00