Computational Semantics and Pragmatics

Autumn 2014



Raquel Fernández
Institute for Logic, Language & Computation
University of Amsterdam

Practical Matters

- **Lecturer:** Raquel Fernández (raquel.fernandez@uva.nl)
  Science Park 107, room F1.07
- **TA:** Julian Schlöder (julian.shloeder@gmail.com)
- **Website:** Slides, references, and other important information will be posted on the course's website:
  http://www.illc.uva.nl/~raquel/teaching/cosp/cosp2014/
- **Timetable:**
  - Tuesdays & Thursday 11-13h, in different rooms.
- **Relevant seminars at the ILLC:**
  - Computational Linguistics Seminar (CLS)
  - DIP (discourse processing) Colloquium

  Check the ILLC Events webpage for details.

What is this course about?

Obviously about the semantics and pragmatics of natural language: about meaning and interpretation in context, and about language use in interaction.

The precise contents of the course change every year. This year the course is dedicated to dialogue modelling: models of language as it is used in actual conversation (more later).

The course is also about using computational and empirical methods to explore semantic/pragmatic phenomena.

Prerequisites

No formal prerequisites are required to follow the course. However, some basic things are expected from you:

- I'll assume some basic knowledge of semantics and pragmatics
- an empirical orientation: an interest in the empirical evidence behind theoretical claims; and in working with existing data, or in collecting data via experiments.
- a computational inclination: an interest in computational methods of enquiry and evaluation
  - if you don't know any programming, you are still welcome
  - you are encouraged to learn some basic programming skills
  - if you have programming knowledge, you'll have the chance to use it

Dialogue Modelling

Research on dialogue deals with the study of language as it is used in conversation $\rightsquigarrow$ natural setting to study language.

- spontaneous and online: disfluent, fragmentary, elliptical
  - what is grammatical? what are the units of study?
  - multi-modality
- multi-agent phenomenon: coordination
  - content coordination
  - coordination of the communicative process: turn-taking, feedback

A transcript fragment from the Switchboard corpus:

```
B.52 utt1:  Yeah, /
B.52 utt2:  [it's,+ it's] fun getting together with immediate family.  /
B.52 utt3:  A lot of my cousins are real close /
B.52 utt4:  {C and} we always get together during holidays and
             weddings and stuff like that, /
A.53 utt1:  {F Uh, } those are the ones that are in Texas?  /
B.54 utt1:  # {F Uh, } no, # /
A.55 utt1:  # {C Or } you # go to Indiana on that?  /
B.56 utt1:  the ones in Indiana, /
B.56 utt2:  uh-huh.  /
A.57 utt1:  Uh-huh, /
A.57 utt2:  where in Indiana?  /
B.58 utt1:  Lafayette.  /
A.59 utt1:  Lafayette, I don't know where, /
A.59 utt2:  I used to live in Indianapolis.  /
B.60 utt1:  Yeah, /
B.60 utt2:  it's a little north of Indianapolis, about an hour.  /
```

Some key units of analysis

- Turns: stretches of speech by one speaker bounded by that speaker's silence – that is, bounded either by a pause in the dialogue or by speech by someone else.
- Utterances: units of speech delimited by prosodic boundaries (such as boundary tones or pauses) that form *intentional units* – that is, that can be analysed as an action performed with the intention of achieving something.
- Dialogue acts: intuitively, conversations are made up of sequences of actions such as *questioning, acknowledging*,... a notion rooted in *speech act theory*.

Main topics we will cover

- Turn taking
- Dialogue acts
- Grounding
- Dynamic semantics for dialogue
- Disfluencies
- Convergence and alignment

Some references where you can find overviews of the field:

R. Fernández (2014) Dialogue. To appear in *Oxford Handbook of Computational Linguistics*, 2nd edition.

Schlangen (2005) Modelling dialogue: Challenges and approaches. Künstliche Intelligenz, 3:23-28.

D. Jurafsky & J. Martin (2009) *Speech & Language Processing*, chapter 24 "Dialogue and Conversational Agents".

M. McTear (2002) Spoken Dialogue Technology: Enabling the Conversational User Interface, *ACM Computing Surverys*, 34(1).

K. Jokinen & M. McTear (2010) *Spoken Dialogue Systems*, Synthesis Lectures on Human Language Technology.

Course evaluation

- Coursework: 30%
  - some graded exercises
  - some non-graded exercises
    (failing to submit means loosing 0.5 on total coursework grade)
  - reading and discussion of relevant reseach papers
    ($\approx$ 10% at my discretion)
- Final project (paper + presentation): 70%
  - on-topic philosophical/theoretical essays are in principle OK, but
  - ideally, your project should include an empirical/computational component, e.g. analysis of real data or some sort of implementation

Learning outcomes

To succeed in the course, you should demonstrate an understanding of the topics covered by being able to:

- Analyse and critique the research questions and the methodology used to address them in existing relevant literature.
- Formulate your own research questions within the scope of the course.
- Apply appropriate (empirical/computational) techniques to address your research questions.
- Write about the work of others and your own work in proper scientific style.
- Present the work of others and your own work to an audience in a clear and engaging way.

This is a research-oriented course (slightly more appropriate for 2nd-year master's students, but 1st-year's are welcome too if committed!).

Final projects

Any topic related to the themes covered in the course. A few ideas on possible types of projects (abstracting over particular topics):

- a quantitative corpus study of some interesting phenomenon
- a machine learning experiment using an existing corpus
- an analysis of data collected by yourself in an experiment
- an implementation of an interesting problem
- an analysis and small extension of a paper from the literature
- an analysis of interesting connections between different approaches
- ...

Some options in this list may seem unfeasible to you, but they may be perfectly possible – don't abandon an interesting idea before discussing it with me!

It might be possible to work in pairs – to be decided on a case by case basis.

Timeline

- You should propose a project topic to me no later than 9 Oct.
- The last week of the course will be dedicated to working on your project with supervision from my part (no lectures, individual meetings).
- Project presentations on 23 Oct.
- Deadline for final paper: Sunday 26 Oct (no extensions).

You can find these and more details on the course website.

End of Introduction

Turn Taking

Dialogue participants do not only need to make decisions about what to say, but also about *when to say it* $\rightsquigarrow$ timing

- Turn-taking is one of the fundamental organisational principles of conversation.
- Learned early: within the first 2 years of life
- Strong universal patterns: tendency to minimize both *overlap* and *gaps* between turns

How do we do it?

Empirical facts: what we know about how it works

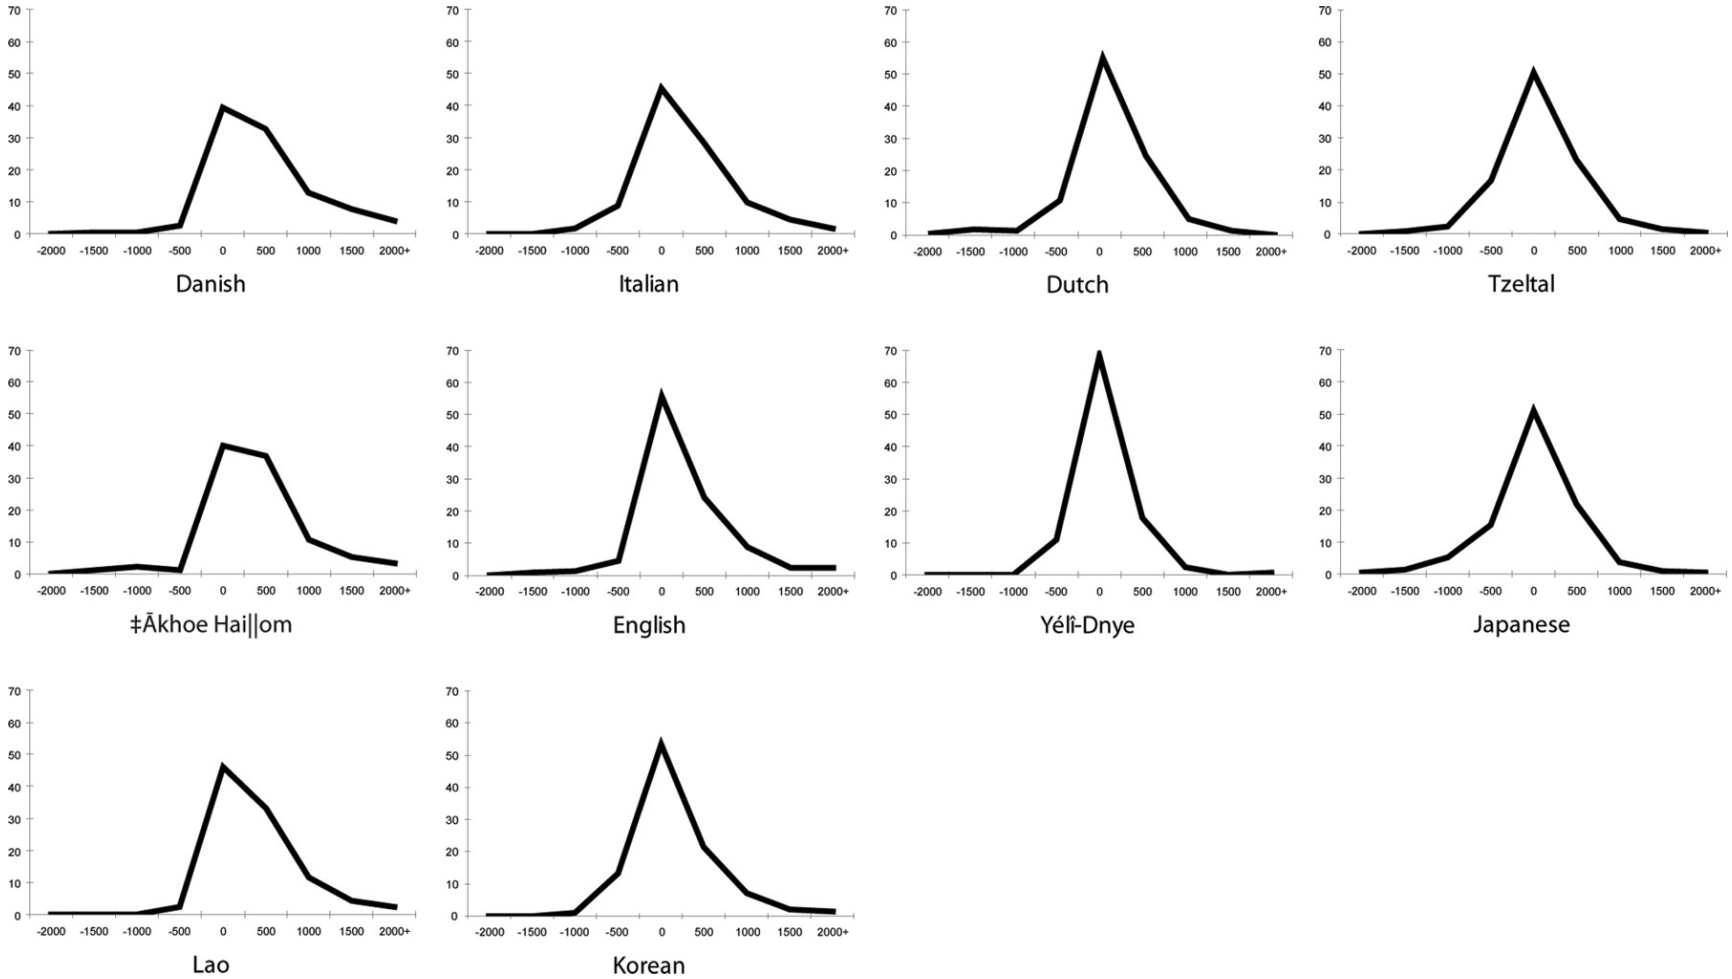
Turn-taking happens very smoothly:

- Overlaps are rare: on average, less than 5% of speech.
- Inter-turn pauses are very short: $\sim$ 200ms (less than 500ms.)

This has long been known for English, but recent large-scale studies show it holds cross-linguisticaly and cross-culturally.

Empirical facts: what we know about how it works

Distribution of turn transition length in milliseconds in 10 languages:



Stivers et al. (2009) Universals and cultural variation in turn-taking in conversation, *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*.

Empirical facts: what we know about how it works

Turn-taking happens very smoothly:

- Overlaps are rare: on average, less than 5% of speech.
- Inter-turn pauses are very short: $\sim$ 200ms (less than 500ms.)
  - even shorter than some intra-turn pauses
  - shorter than the motor-planning needed to produce the next utterance

Turn-taking can't be *reactive* (we do not react to silence as most artificial conversational agents do) — turn taking is predictive: we *project* turn endings and turn transitions.

Conversation Analysis Model

The seminal model of turn taking was put forward by sociologists within the framework of Conversation Analysis (Sacks et al. 1974)

Sacks, Schegloff, & Jefferson (1974) A simplest systematics for the organization of turn-taking in conversation.

- Turns consist of turn constructional units (TCUs) with projectable points that can be predicted beforehand.
- Such projectable points act as transition relevance places (TRPs) where turn transitions are relevant.

Conversation Analysis Model

Three rules govern the expected behaviour at TRPs:

1. if devices to select a next speaker (e.g. questions, gaze, naming) are used, the current speaker stops and the selected speaker takes the turn;
2. else, any other speaker may take the turn (may *self-select*),
3. if no other party takes the turn, then the current speaker may continue.

Predictions:

- Minimal overlap
  - only one speaker may generally be speaking at any time (speakers wait to TRPs)
  - overlap may occur at wrongly projected TRPs (TCUs are varied)
  - overlap may occur when there are competing next speakers
- No silence gaps as the norm
  - to the extent that TRPs are projectable.
  - if silence occurs, it is typically for a reason $\rightsquigarrow$ silence is informative.

Turn-Taking Models

Some subsequent research has focused on how to make more precise the notions of TCU and TRP.

- How can TRPs be predicted? humans are able to predict whether an utterance will continue and for how many words.
  - syntactic closure plus acoustic information (rising/falling intonation; faster speaking rate);
  - syntactic completion is context-dependent - pragmatic completion;
  - certain prosodic patterns signal that the speaker plans to hold the turn beyond syntactic completion;
  - word fragments and filled pauses are indicative of turn-hold.

To do

For next Tuesday 9 Sept:

- read one research paper on turn taking (see website for references); we'll discuss this in class
- if you have time, read some LanguageLog posts on turn taking

For next Thursday 11 Sept:

- Transcribe a fragment of a conversation (preferably between two participants)
- pay attention to: turn segmentation, overlap, pauses, turn holding, turn yielding, and turn taking signals

You will find more details on these exercises on the course website later today.