Computational Semantics and Pragmatics

Raquel Fernández

Institute for Logic, Language & Computation
University of Amsterdam

Autumn 2016

Overview of topics

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

Linguistic coordination

Speakers in dialogue tend to *adapt to each other* at different levels:

- phonetic production (Babel 2012, Kim et al., 2011)
- lexical choice (Brennan and Clark, 1996)
- syntactic constructions (Pickering and Ferreira, 2008)
- gestures (Furuyama et al.,2005) postural sway (Shockley eat al., 2007)

[Terminology: alignment, entrainment, coordination, convergence, adaptation]

Our interest here is in *linguistic alignment*: adaptation to aspects of our conversational partner's language

- Alteration in likelihood of particular language behaviour
- May be dynamic adjustment to partner's most recent contribution
- or gradual alignment during (and beyond) interaction
- Found in both experimental and natural interactions of many kinds, in many languages

Linguistic coordination

- Empirical *evidence* of alignment / coordination
- What *causes* this adaptation is a matter of debate:
  - the need for mutual understanding (Clark, 1996)
  - priming (Pickering & Garrod, 2004)
  - negotiating social distance (Giles, 2008)

Alignment at different linguistic levels

Phonology/phonetics: speech rate, response latencies, vocal intensity, pronunciation, pausing patterns

Lexicon (word choice): *shoe* vs. *pennyloafer*

Syntax: If your partner uses a syntactic structure, you are more likely to use it too.

The nun is giving a book to the clown (V NP PP) vs.
The nun is giving the clown a book (V NP NP)

The cowboy is giving the banana to the burglar vs.
The cowboy is giving the burglar the banana

Alignment at different linguistic levels

Semantics: dialogue partners converge on semantic conceptualisations

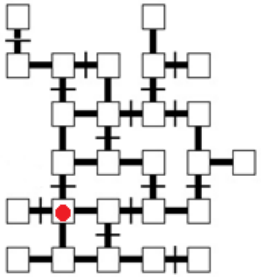

Description schemas:

I'm at B5 vs.
I'm at second column, second row from the bottom

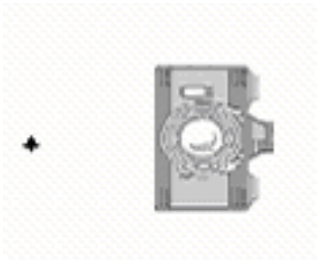

Reference frames:

The dot is below the camera vs.
The dot is to the left of the camera

Alignment in human-computer interaction

Humans also align with artificial dialogue partners.

- Alignment of lexical choice in route-finding task (Koulouri et al, 2014)

> Robot: I am at the junction by the bridge, facing the **bendy road**.
>
> User: Go into the **bendy road**.

- Children modify their speech in response to animated characters (Coulston et al. 2002)
  - greater amplitude with louder 'extrovert' character
  - smaller with quieter 'introvert' character

Exploiting alignment in HCI

User's alignment with the system: Alignment reduces the space of possible user behaviours. This can help HCI by:

- implicitly shaping the user's input in a way that the system can understand: eliciting specific behaviour (word choice, grammatical structures, speech rate, amplitude...)
- predicting user input

System's alignment with the user: generating more naturalistic output

- Users expect that the conversational partner will align
- Increasing user satisfaction

Why do people align language?

So, there is evidence of alignment, but. . . what triggers this type of coordination?

Three different approaches to explaining alignment:

- driven by *communicative* goals and the need for mutual understanding
- consequence of our *cognitive* architecture, triggered by priming mechanisms
- driven by *social* goals, to negotiate social distance

Alignment is driven by communicative goals

Speakers align to maximise *mutual understanding*.

- Appeal to common ground (joint action model by Clark et al.)
- Audience design: *what is my interlocutor likely to understand?*
  - driven by the desire to be understood, to reach mutual understanding
  - leads to more successful communication

Alignment is goal-directed. Goal: *communicative success*

- it requires a model of the dialogue partner as communicative agent

Evidence

- Partner-specific conceptual pacts
- Referential task (lexical choice)

$< 15\%$ chance to use *'seat'* in null context

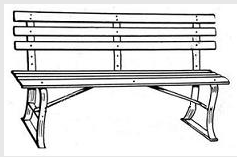

If partner uses *'seat'*:

– 83% alignment when thinking partner is a computer
– 44% alignment when thinking partner is a human

– 80% alignment when thinking partner is an basic computer
– 42% alignment when thinking partner is an advanced computer

More lexical alignment with 'less capable' partner (Branigan et al. 2011)

Communicative beliefs affect lexical alignment.

Evidence

Grounding problems affect alignment.

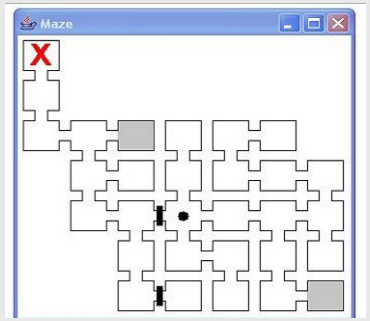

Pattern of semantic shift:

0 mins: The piece of the maze sticking out
2 mins: The left hand corner of the maze
5 mins: The northenmost box
10 mins: Leftmost square of the row on top
15 mins: 3rd column middle square
20 mins: 3rd column first square
25 mins: 6th row longest column
30 mins: 6th row 1st column
40 mins: 6 r, 1 c
45 mins: 6, 1

Reversion to figurative model after clarification:

A: I'm in the 4th row 5th square.
B: Where's that?
A: The end bit.
B: I'm on the end bit right at the top.

Participants systematically favour Figural and Path descriptions when encountering problematic dialogue

Garrod and Doherty (1994) Conversation, co-ordination and convention: an empirical investigation of how groups establish linguistic conventions. *Cognition*, 53:181-215.

Mills and Healey (2008) Semantic negotiation in dialogue: mechanisms of alignment, in *Proceedings of SIGdial*.

Alignment is due to our cognitive architecture

Alignment is a natural consequence of the architecture of *our cognitive system*.

- Interactive alignment model (Pickering & Garrod 2004)
  - alignment driven by activated linguistic representations – priming (stimulus, response)
  - leads to reduction of cognitive load, and indirectly to successful communication

Pickering & Garrod, Toward a mechanistic psychology of dialogue, *Behavioral and Brain Sciences*, 27(2):169–190, 2004.

Pickering & Garrod, The interactive-alignment model: Developments and refinements, *Behavioral and Brain Sciences*, 27(2):212–225, 2004.

Interactive alignment model

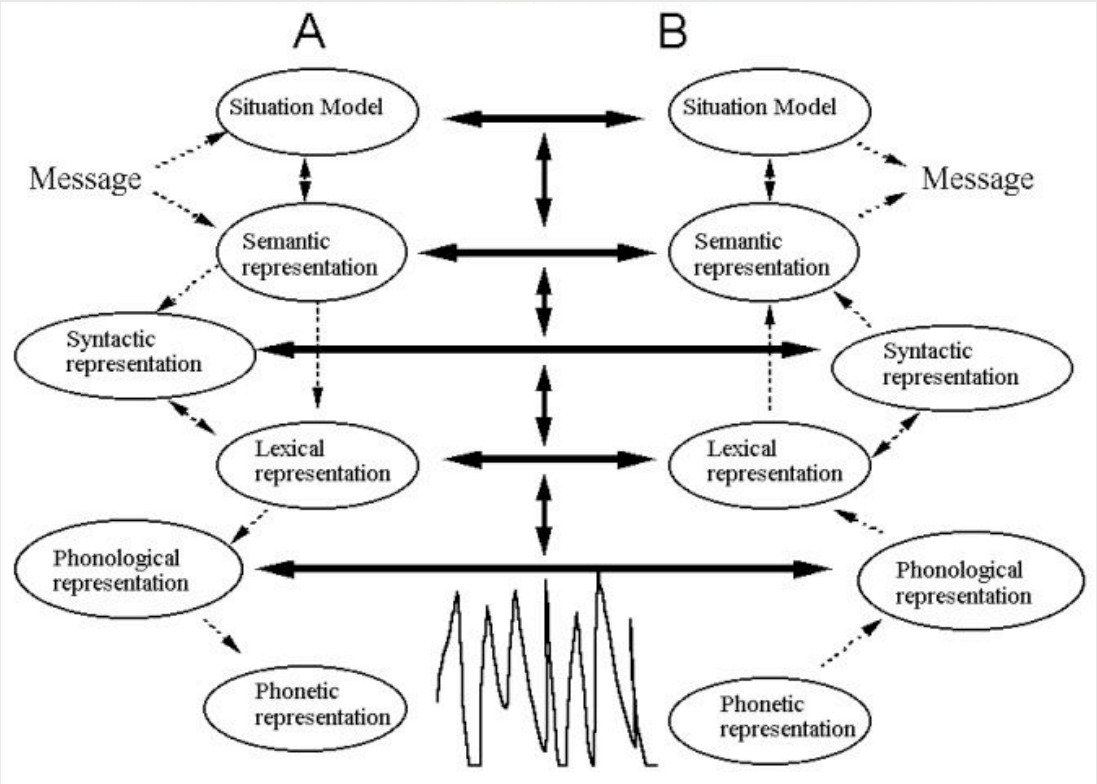

(Pickering & Garrod 2004)

- Priming operates on representations at every level
- Alignment at one level enhances alignment at other levels
  e.g., syntactic alignment is enhanced by lexical / semantic overlap
- Alignment of situation models leads to successful communication

Alignment is due to our cognitive architecture

Alignment is a natural consequence of the architecture of *our cognitive system*.

- Interactive alignment model (Pickering & Garrod 2004)
  - alignment driven by activated linguistic representations – priming (stimulus, response)
  - leads to reduction of cognitive load, and indirectly to successful communication

Alignment is *not goal directed*.

- implicit and automatic (triggered by linguistic features)
- no representation of partner required

Evidence

- Syntactic alignment
- Syntactic alignment with lexical boost

nun **giving** a book to a clown (V NP PP rather than "nun giving a clown a book")
→ "sailor **showing** a hat to a girl"; more priming with "sailor **giving** a hat to the girl"

the **sheep** that's red (Relative Clause rather than "the red sheep")
→ "the **book** that's red"; more priming with "the **goat** that's red"

- Same level of syntactic alignment under differing beliefs – believing partner is human (66%) vs computer (64%)

Bergmann, K., Branigan, H., & Kopp, S. (2015). Exploring the alignment space: lexical and gestural alignment with real and virtual humans. *Frontiers in ICT*, 2(7), 1–11

Mirror Neurons

So called *mirror neurons* fire during both action and perceiving an action (Di Pellegrino et al. 1992).

New Pickering & Garrod model:

- Production and comprehension are tightly interwoven – this underlies people's ability to predict themselves and each other.
- Based on *covert imitation* and *forward modelling*: recreating behaviour and predicting the perceptual outcomes of an action

M. Pickering & S. Garrod (2013) An integrated theory of language production and comprehension. *Behavioural and Brain Sciences*.

Audience design vs. priming

- A lot of evidence is consistent with the two models.
- No single account explains the full range of evidence.
  - different linguistic levels sensible to different mechanisms?
- Most research does not seek to contrast accounts: different tasks, different contexts, different partner behaviour.

Some evidence that speakers fail to adapt to partners in the early moments of processing (Keysar, Barr, and Horton, 1998)

- early processing is *egocentric*
- maintaining and updating a model of the partner is computationally expensive, so is done only when necessary (Pickering & Garrod, 2004)

But this has been countered by Brennan & Hanna (2009):

"early moments of language processing can be flexible, nimble, and responsive to such attributions, rather than reflexive, egocentric, and 'dumb'."

Brennan, S. E. & Hanna, J. E. (2009). Partner-specific adaptation in dialogue. *Topics in Cognitive Science.*

Alignment is driven by social goals

Speakers align to socially index and achieve rapport with conversational partners.

- Communication accommodation theory (Giles et al.)

Alignment:

- driven by affiliation, desired to be liked, need for social approval
- leads to more likeable perception, more acceptance/compliance

Goal: *enhancement of social relations*

- it requires a model of the dialogue partner as social agent

Evidence

- Speech rate alignment implicitly increases compliance with requests (Buller & Aune 1992)
- Repetition increases waiters' tips (Van Baaren et al. 2003)
- Matching of functions words predicts relationship initiation and stability in speed dating conversations (Ireland et al., 2011)
- More alignment towards high-powered partners

Coordination and status-based power

C. Danescu-Niculescu-Mizil, L. Lee, B. Pang and J. Kleinberg (2012). Echoes of power: Language effects and power differences in social interaction, *Proceedings of WWW*.

We need a reasonably *large corpus* with *social asymmetries* amongst interacting agents

⇝ Turn to *online communities*

- community of Wikipedia editors
- some of them are *administrators*
- they interact via "talk pages"

User talk:Mackensen

From Wikipedia, the free encyclopedia

Canadian folk singer talk pages [edit]

....are being recreated. Would you mind deleting them again and salting them? Thank you, JNW (talk) 01:00, 14 June 2014 (UTC)

. Done. I've left the IP a friendly note. Mackensen **(talk)** 01:13, 14 June 2014 (UTC)

. Much appreciated. I noticed some of those talk pages had been deleted a half dozen times since 2012. Maybe a sneaky way of reintroducing deleted articles? JNW (talk) 01:16, 14 June 2014 (UTC)

Style Coordination

How things are said as opposed to *what* is said

$\rightsquigarrow$ *function words* are topic-independent (Pennebaker et al, 2007)

pronouns, articles, quantifiers, prepositions, conjunctions, . . .

Editor$_a$: Corrected. Please check. **Any** more outstanding problems?

Editor$_b$: **Everything** is fine. Thanks a lot.

Coordination of b towards a for a class of function words m, for all pairs of utterances (u_a, u_b) where b directly replies to a:

$$C^m(b, a) = P(u_b \text{ uses } m \mid u_a \text{ used } m) - P(u_b \text{ uses } m)$$

Summary

Coordination / adaptation of style (broadly understood) in dialogue. Three perspectives:

- driven by communicative, partner-specific goals
- mechanistic consequence of our cognitive architecture
- driven by social goals

To read for discussion on Tuesday:

C. Danescu-Niculescu-Mizil, L. Lee, B. Pang and J. Kleinberg (2012). Echoes of power: Language effects and power differences in social interaction, *Proceedings of WWW*.

D. Reitter and J. Moore (2007). Predicting Success in Dialogue, *Proc. ACL*.

↪ More up-to-date longer version: Reitter & Moore (2014)
Alignment and task success in spoken dialogue, *Journal of Memory and Language*

Final projects: start to think about it!

- you are strongly encouraged to collaborate in pairs
- project proposal due on Monday 10 Oct.