Hammers and Nails in Linguistic Interaction

Raquel Fernández

Institute for Logic, Language & Computation
University of Amsterdam

ESSLLI 2015, Barcelona

Outline

1. Agreement and Disagreement
2. Power Asymmetries in Interaction
3. Child-Adult Dialogue

Part 1:
Agreement and Disagreement

Keeping Track of the Common Ground

- To maintain coherence in dialogue, speakers must keep track of their *common ground*.
- What is agreed upon? what is still under discussion?

Sometimes, this is easy:

```
(1) A: That slogan is quite obvious.
    B: I agree.  / That's not true.
```

Sometimes it is not at all trivial:

```
(2) A: I never did care for him, in the James Bond movies.
    B: I was never into those movies, either.
(3) A: This is a very interesting design.
    B: It's just the same as normal.
```

[All examples from AMI and Switchboard corpora]

Sometimes it *seems* easy, but it is not . . .

> (4) A: But it's uh yeah it's an original idea.
>   B: Yes, it is. ⤳ **acceptance.**
>
> (5) A: A banana is not it's not really handy.
>   B: Yes, it is. ⤳ **rejection.**
>
> (6) A: It's not very well advertised.
>   B: No, it's not. ⤳ **acceptance.**

Although the responding utterance seems trivial, determining its
dialogue function – *acceptance* vs. *rejection* – is not.

What's at stake is how the *polarities* of *proposal & response* interact.

Polarity

Logical polarity has not been explored in *computational approaches*
Galley et al. (2004), Germesin & Wilson (2009), Misra & M. Walker (2013)

It has seen renewed interest in *formal semantics* regarding polarity
particles and negation Farkas & Roelofsen (2013), Cooper & Ginzburg (2013)

```
(7) A: Sue failed the exam.
    B: Yes she did. / No she didn't.
(8) A: Sue did not pass the exam.
    B: Yes she did.  / No she didn't.
```

In classic semantics, A's assertions have the same propositional content
⤳ include polarity to account for different 'meaning' of *yes* / *no*.

Our aim: determine the *accepting* or *rejecting force* of a response.

J. Schlöder and R. Fernández. The role of polarity in inferring acceptance and rejection in dialogue. *SIGdial* 2014.

Relative (dis)agreement: Formal Model

Basic model: assume a proposal P is on the table. The next move R *accepts* P iff $P \wedge R$ is consistent.

Assign a polarity (pos/neg) to proposal and response, respectively:

- aligned polarities $\rightsquigarrow$ *accepting force*
- misaligned polarities $\rightsquigarrow$ *rejecting force*

  R: **relative agreement**
  - P positive $\rightsquigarrow$ default case (positive-positive)
  - P negative $\rightsquigarrow$ reverse case (negative-negative)

  R: **relative disagreement**
  - P positive $\rightsquigarrow$ default case (positive-negative)
  - P negative $\rightsquigarrow$ reverse case (negative-positive)

  R: **absolute agreement / disagreement**

Empirical Study

How widespread is relative polarity in *actual dialogue*?
Can our formal model be operationalised and have *practical value*?

Computational experiment:

- ~ 1300 P-R pairs from two dialogue corpora (AMI & Switchboard) of which only 12% are rejections
- Task: identification of rejections
- Naive Bayes classifier with several standard features
- Use of surface-form heuristics for polarity assignment
- **Relative polarity boosts results substantially**
  (F-score increased from .52 to .60 in AMI and from .33 to .58 in SWB)

Open Issues . . .

Some *logically consistent* responses may act as *rejections*:

> (9) A: We are all mad, aren't we?
>   B: Well, some of us. ↝ *not (necessarily) all of us?*

J. Schlöder and R. Fernández. Pragmatic rejection. *IWCS* 2015.

Many exchanges are not clearcut acceptances or rejections . . .

Crowdsourcing experiment: beyond gold-standard corpus annotations and our intuitions, *what does the crowd think?*

Please indicate which of the following options best captures what speaker B meant:
  • definitely / • possibly *agrees* with A   • definitely / • possibly *disagrees* with A

(10) A: All drug dealers can be   (11) A: Let's start with Dim Sum.
         sentenced to the death sentence.   B: Or have some vegetables.
     B: Convicted drug dealers.
     ↝ **25% disagreement category**   ↝ **95% disagreement category**

**Part II:
Power Asymmetries 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)

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)

Focus today: *social factors behind linguistic adaptation*

What kind of data?

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

⤳ *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:

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

Overall *coordination towards* a: average across all editors b who
address a (adapted from Danescu-Niculescu-Mizil et al. 2012)

Status-based power : Wikipedia editors coordinate more towards admins (Danescu-Niculescu-Mizil et al., 2012. *Echoes of Power.*)

What about other more implicit forms of social power, such as how *'central'* you are within the *social network* – do they impact linguistic style matching?

B. Noble and R. Fernández (2015). Centre Stage: How Social Network Position Shapes Linguistic Coordination. *Workshop on Cognitive Modeling and Computational Linguistics, NAACL 2015.*
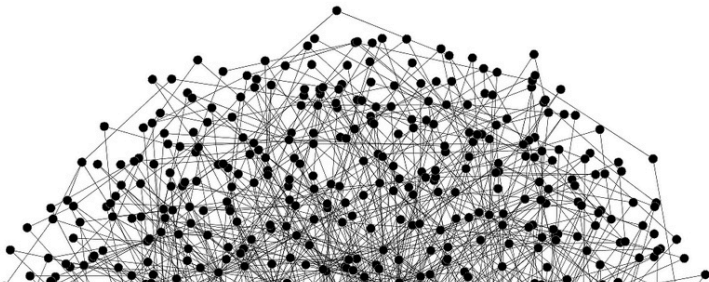
The Wikipedia Social Network

We want to construct a *social network* that reflects the *linguistic interactions* between the Wikipedia editors:

- nodes represent individuals in a community – Wikipedia editors
- edges give some measure of social connectivity between individuals – weighted according to the number of direct replies

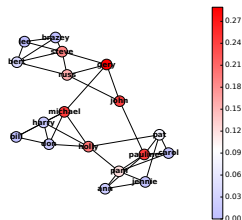Corpus: 342,800 posts, 26,397 editors (1,825 of whom are admins)
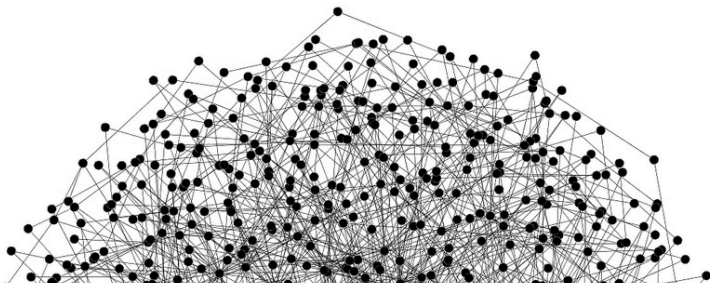
Centrality Measures

Betweenness centrality: How important
are you to community connectivity?

$$\mathrm{BC}(n^*) = \sum_{n \neq m \in N} \frac{|\{\sigma \in \mathrm{Path}(m, n) \mid n^* \in \sigma\}|}{|\mathrm{Path}(m, n)|}$$
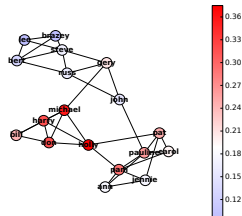
where $\mathrm{Path}(m, n)$ is the set of shortest paths
between m and n

Eigenvector centrality: How important
are your neighbours?

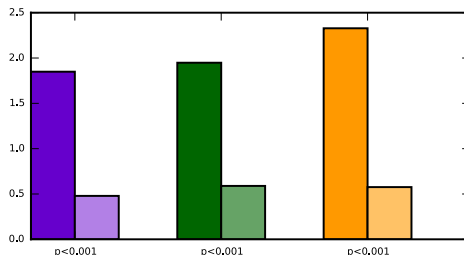$$\mathrm{EC}(n^*) = \frac{1}{\lambda} \sum_{n \in M(n^*)} \mathrm{EC}(n)$$

where $M(n)$ is the neighbourhood of n and λ is
the largest *eigenvalue*

Highly central editors: over one standard deviation above mean score.

Results

- More style coordination towards administrators.
- More style coordination towards editors in central social positions.

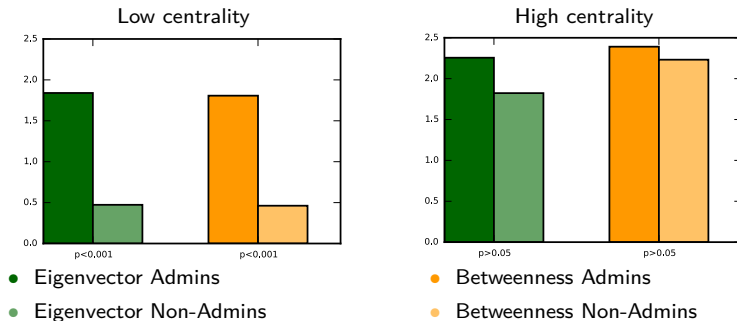- Admins
- Non-Admins
- High Eigenvector
- Low Eigenvector
- High Betweenness
- Low Betweenness

- On average, admins occupy more central positions, but the impact of adminship and centrality turn out to be largely independent ...

Results

- Low-centrality editors receive more coordination if they are admins.
- But adminship is less important for high-centrality users.

Low centrality

High centrality

- Eigenvector Admins
- Eigenvector Non-Admins
- Betweenness Admins
- Betweenness Non-Admins

⤳ *social network centrality* sometimes eclipses status-based power in triggering linguistic style adaptation.

Open Issues . . .

Is adaptation to central users (rather than admins) more important for *social acceptance*?

- how does this happen even though centrality is more *implicit*?
- do highly central users exhibit speech *more typical* of the community?

Several *practical applications* within computational social science:

- automatic discovery of social relations,
- tracking evolution of relations over time, . . .

Part III: Child-Adult Dialogue

Coordination in Child-Adult Dialogue

child $\rightarrow$ adult *language learning*
child $\leftarrow$ adult *child-directed speech*

---

input *vs.* interaction

sensitivity to statistical regularities
in the input ignoring interaction

sensitivity to when & how the
input if offered in interaction

```
Adult: Help me put your toys away, darling.
Child: I'm going to Colin's and I need some toys.
Adult: You don't need a lot of toys.
Child: Only a little bit toys.
Adult: You only need a few.
Child: Yes, a few toys.
```

Focus here: ways of investigating how speakers pick up on each other's language (*coordinate*) at different degrees of locality.
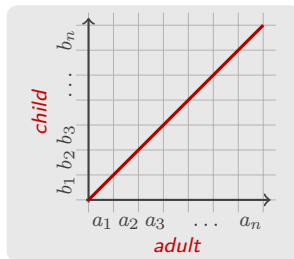
Turn-based Cross-Recurrence Plots

Two-party dialogue transcript

Cross-recurrence plot: each cell corresponds to a pair of turns (i, j)

```
A₁: which one do you want first
B₁: that one
A₂: you like this one
B₂: yeah, give me
    .
    .
    .
Aₙ: ...
Bₙ: ...
```

$\Rightarrow$

Recurrence (coordination) score for each (i, j)

- *global recurrence*: average coordination over all turn pairs
- *local recurrence*: recurrence in (semi-)adjacent turns, separated by at most distance $d < n$ (diagonal line of incidence)
- *upper recurrence*: child's turn comes after adult's *adult ← child*
- *lower recurrence*: adult's turn comes after child's *child ← adult*

Turn-based Cross-Recurrence Plots

CRP of a dialogue with Abe (2.5 years old):

order of turns shuffled   original dialogue
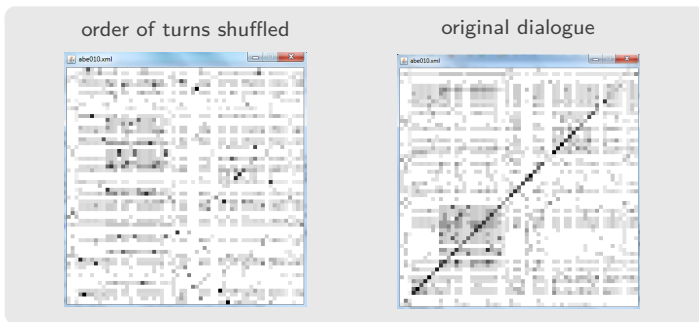
Same *global* recurrence but very different *local* recurrence

⇝ global: chance recurrence regardless of temporal development of interaction

Measuring Recurrence

Many measures are possible: lexical, conceptual, syntactic,. . .

R. Fernández & R. Grimm. Quantifying Categorical and Conceptual Convergence in Child-Adult Dialogue, *36th Annual Conference of the Cognitive Science Society*. 2014.

Syntactic coordination: number of shared part-of-speech bigrams factoring out lexical identity, normalised by length of longest turn.

> Adult: you are pressing a button and what happens ?
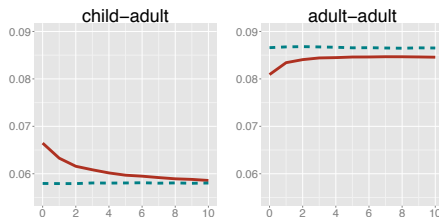>   PRO|you AUX|be PART|press DET|a N|buttton CJ|and PRO|what V|happen
>
> Child: what happens the horse tail
>   PRO|what V|happen DET|the N|horse N|tail

Results

Data: 380 dialogues from 3 children over a period of ∼3 years.
For comparison: ∼1000 adult-adult dialogues from Switchboard.

- *local vs. global*: significantly more local coordination.

- *directionality*: both coordinate more at local levels, but the adult recurs with the child significantly more.
- *difference with adult dialogue*: very different coordination patterns, with adults showing syntactic *divergence* at adjacent turns
  ⤳ less recurrence than expected by chance.

Open Issues . . .

Contrast with previous evidence of syntactic alignment in adult-adult dialogue (e.g., Pickering & Ferreira 2008), but not surprising

↝ advancing a conversation requires *different dialogue acts* with distinct syntactic patterns.

Why is there syntactic recurrence in child-adult dialogue?

- *feedback mechanism* to ratify linguistic constructions?
- possibly related to *corrective feedback*

```
Child:   you're good to sharing.
Mother:   I'm good at sharing?
```

S. Hiller & R. Fernández. Towards the Automatic Extraction of Corrective Feedback in Child-Adult Dialogue. *SemDial* 2015.

Ultimate question: to what extent does *interaction* contribute to *language acquisition*?

Recap

1. Agreement and Disagreement
2. Power Asymmetries in Interaction
3. Child-Adult Dialogue

logic, machine learning, crowdsourcing,
corpus studies, social network analysis,
recurrence quantification analysis, . . .

Thanks!

Julian Schlöder

Bill Noble

Sarah Hiller

Robert Grimm

I'm hiring: 1 PhD & 1 postdoc

`raquel.fernandez@uva.nl`

NWO
Netherlands Organisation
for Scientific Research

essence
Marie Curie
Initial Training Network

Take a look at SIGdial conference, SemDial workshop, *Dialogue & Discourse* journal