

Who decides what a text means?

And what the answer implies for computational linguistics

Graeme Hirst
University of Toronto

Acknowledgements

Stephen Regoczei

Jean-Pierre Corriveau

Jane Morris

Susan McRoy

Peter Heeman

David Traum

Philip Edmonds

Nadia Talent



Philosophically naïve

– like CL itself

Text-meaning

Any complete utterance

- Spoken or written
- Interactive or not
- Long or short

Text-meaning

Meaning of whole message,
including implicatures,
inferences, affect, subtext

- Not just word-
or sentence-meaning
- Could be more than, or less
than, sum of sentence-meanings

What is the locus of text-meaning?

- 1 Meaning is in the text itself

The text is all we have

We can't read the writer's mind

The New Criticism:

W.K. Wimsatt, Monroe Beardsley,
Cleanth Brooks, ...

Facts about the writer aren't
relevant to the text's meaning

What is the locus of text-meaning?

- ① Meaning is in the text itself
- ② Meaning is in the writer / speaker

Speaker's intention is essential to the meaning of an utterance

John Austin, Paul Grice, John Searle, ...

- *It's cold in here* → *Turn up the heat*
- *That might be difficult* → *No!*
- *We must do lunch some time*
→ *I never want to see you again*

Take the hint, see the intent

What is the locus of text-meaning?

- ① Meaning is in the text itself
- ② Meaning is in the writer / speaker
- ③ Meaning is in the reader / hearer

The meaning of an utterance is
the reader's response to it

Stanley Fish, Roland Barthes, ...

Different meaning for each
reader?

Interpretive communities give
stability to meaning of a text.

What is the locus of text-meaning?

- ① Meaning is in the text itself
- ② Meaning is in the writer / speaker
- ③ Meaning is in the reader / hearer

Or two of these, or all of these

We can ask the same question for lower levels of linguistic elements

Words

Sentences

Semantic roles

Lexical relations

...

The same three answers are possible

But they don't have to be the same answer at each level

Maybe ...

Effects of individual writer or reader are apparent only at text-meaning level, not below

Or maybe conversely ...

Individual writer or reader's lower-level idiosyncrasies are *dampened* at text-meaning level

Or maybe ...

Views of text-meaning in CL and NLP

- ▶ My view here: Text is always a locus of meaning
- ▶ The issue: Reader and/or writer as additional loci?

Views of text-meaning in CL and NLP

- ▶ Dominance of each view in CL varies with era
- ▶ CL has become *less* sophisticated in its view

The history of
the philosophy of text-meaning
in computational linguistics


1970–2015

1970-1985

1970–1985

- Simple utterances
- All texts are massively ambiguous
All texts are enthymematic (incomplete)
- Use knowledge of world and beliefs to interpret texts.

1970-1985



1970–1985


- Find the interpretation most consistent with what's already known
- Construe input as best match to own prior knowledge
- ③ Meaning is in the reader / hearer

1970-1985

- Example:

*The city councillors denied the demonstrators a permit because **they** were communists.*

Who were the communists?




1985-1995

1985–1995

- Interactive dialogues
- Gricean and pragmatic theories of “real” language use
- Determine real intent of user’s utterances through, e.g., abductive reasoning ...
- ... and hence the user’s goals and plans

1985-1995



1985–1995

- Example:

Talking to domestic robot:

I'd like a beer

→ *Bring me a beer and do it right now*

1985-1995


- A text means whatever the speaker thinks it means or intends it to mean
 - ② Meaning is in the writer / speaker
- The computer's job is to read the user's mind

1995-2015

1995–2015

- Large, non-interactive texts
- Statistical and machine-learning methods
- NLP tasks as meaning-preserving statistical transformations

1995-2015



1995–2015

- Texts regarded as *objets trouvés* (‘found objects’)
- Meaning is “extracted” by “processing” the words and their context
 - ① Meaning is in the text
- “The text is all we have.”

1995–2015

- Examples:

Find articles on raptor migration in Colorado.

Find follow-ups to this news story.

Summarize this report.

Monitor this chat room.

Roles of the linguistic computer

1970-1985: Independent agent

1985-1995: Servant of the user

1995-2015: Reader and transformer of text

Computational linguistics
vacillates between the three
views of locus of text-meaning

But computational linguists
don't notice and don't care

Philosophically naïve

Two types of system


- **Observer:** Reads external text on behalf of a user
- **Conversant:** Actively participates in a dialogue with a user

CL's naïve assumptions about meaning


- User *or* writer is perfect language user
- Meaning is conveyed solely by positives
- No distinction between sentence-meaning and text-meaning or interpretation

CL's naïve assumptions about meaning

- If observer: User's knowledge and agenda are same as the writer's
- If conversant: System's knowledge and agenda are same as user's
 - Same framework or understanding



2016-202X



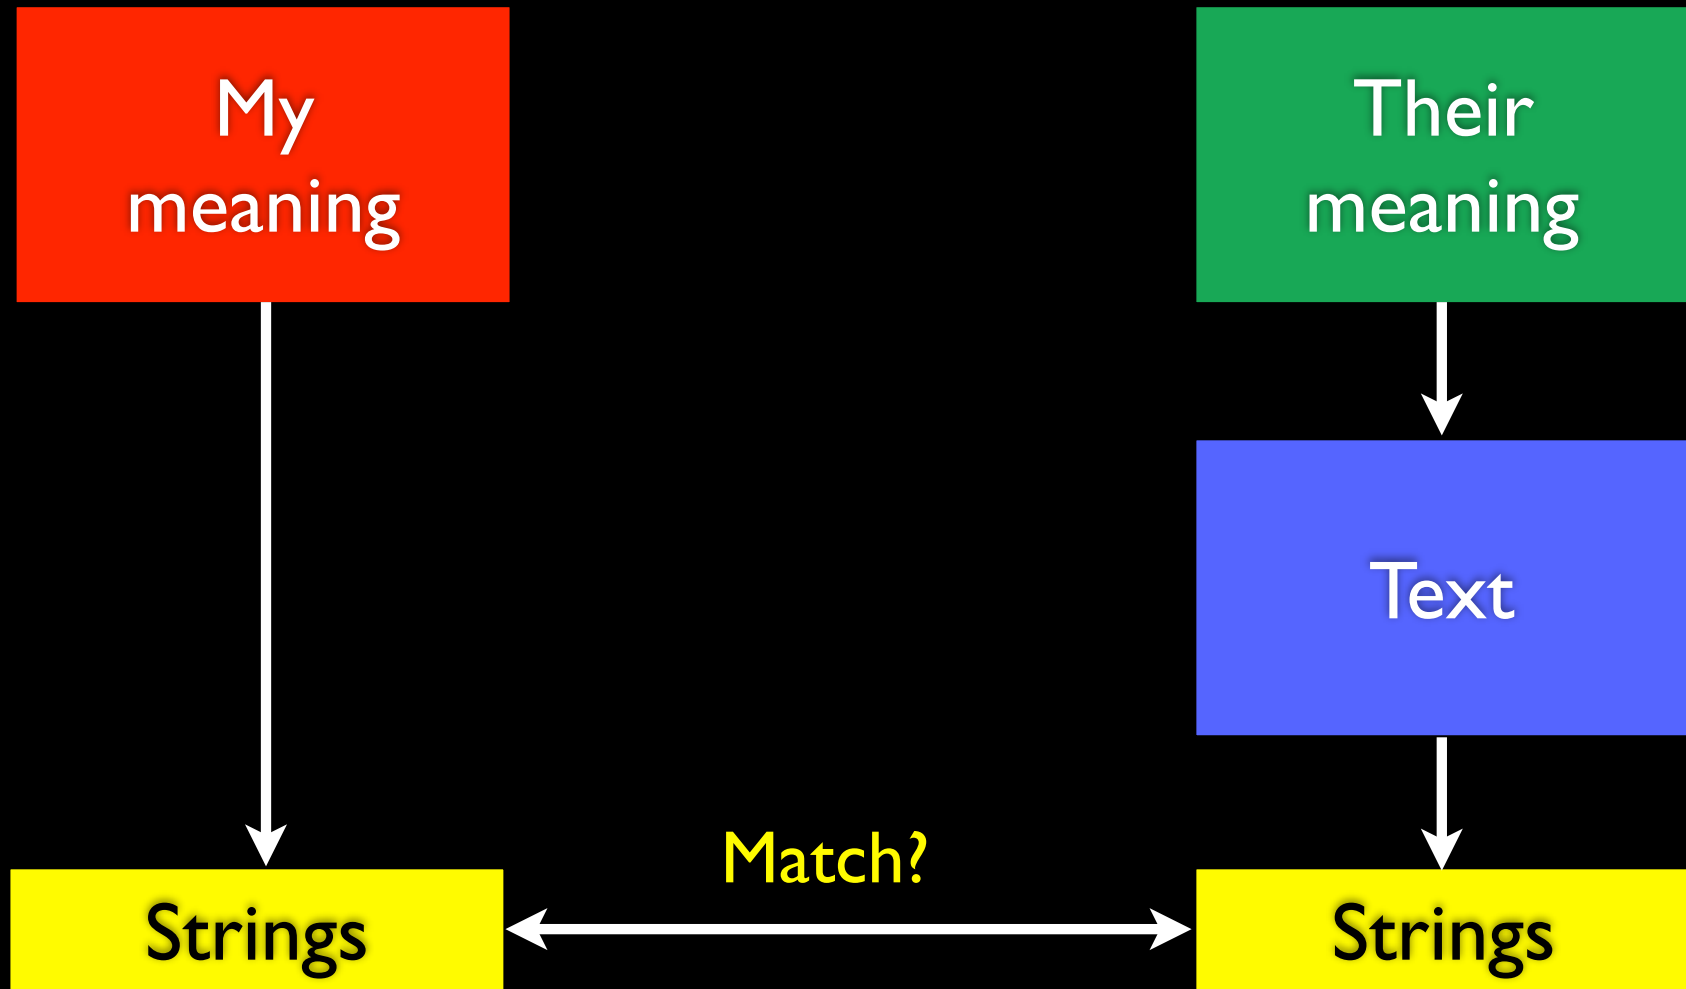
2016–202X

- Elimination of assumption of identical agendas
- Interpretation distinguished from meaning
- User models and consideration of text provenance
- Return of in-reader and in-writer views
- Negotiated meaning


Google
has turned everyone into
researchers

but with only
an **impoverished** view of meaning

≤ 2015



By 202X



① What does this mean for me?

② What are they trying to say?

2016–202X Observers of text

① What does this mean for me?

- *Goal:* Research intermediaries that can interpret from the user's perspective
- To get at reader's meaning, system first needs to understand their **purpose** and **viewpoint**
- Sophisticated user model

2016–202X Observers of text

① What does this mean for me?

- Abstract, wide-ranging, or unusual information needs and query-oriented multi-document summarization
- “Learning by reading”
 - Integrating content of new document into existing knowledge base

~~“The text is all we have.”~~

We also know what the user knows and wants

2016–202X Observers of text

② What are they trying to say?

- *Goal:* Research intermediaries that can interpret text from the writer's perspective
- Published models of writers
- Hermeneutic (interpretive) task
- Intelligence gathering

2016–202X Observers of text

② What are they trying to say?

- Sentiment analysis and classification
- Opinion extraction and ideological analysis
- Detecting verbal deception
- Learning by reading: answering test questions
- Semantic / knowledge-based machine translation

~~“The text is all we have.”~~

We may know the writer and the context

2016–202X Conversants

- 1980s–2000s: User spoke, system acted.
- 2016: User speaks, system might act.
- ③ Negotiated meaning
 - collaboratively constructed by speaker and hearer

Misunderstanding and not-understanding

Computers don't understand very well

ASR and software limitations

People don't understand very well

Inattentiveness, mishearing,
misreading

Difficult material

Poor expression

People don't understand very well

But people succeed anyway

Notice and recover from
problems in understanding

Computers should too

Not-understanding

No unique interpretation
found by hearer

Hearer is aware of the failure

Misunderstanding

Hearer finds interpretation unintended by speaker

Hearer is not immediately aware of the failure

Self-misunderstanding vs other-misunderstanding

Recovery requires negotiation
of meaning

Collaboration on the
construction of a meaning ...

... that works for both

Unifies reader-based view
and writer-based view

Recovering from not-understanding

Example

Collaboration on referring expressions

A: What's that weird creature over there?

B: In the corner?

A: Uh-huh.

B: It's just a fern plant.

A: No, the one to the left of it.

B: That's the television antenna. It pulls out.

Data from Svartvik & Quirk 1980, S.2.4a:1-8

Collaborative repair of non-understanding

- Repair of text-level non-understanding
- Speaker and listener negotiate and refine description of object
- Integrates speaker-based and listener-based views of meaning
- Computational models of this process

Plan-based model of collaboration on referring

- Speaker has goal of referring; plans a description
- Listener tries to infer plan, identify referent
- If unsuccessful, rejects some or all of plan; may suggest an improvement to it

Heeman & Hirst 1995

Plan-based model of collaboration on referring

- Two copies of the model (each with own beliefs and goals) can perform simplified version of fern-plant dialogue

Heeman & Hirst 1995

Example

Collaboration on referring expressions

A: You just stay on 2A until you get to Lowell Street.

B: Is it marked?

A: Yeah, I think there's a street sign there, it's an intersection with lights.

B: Okay.

Data from Psathas 1991

Extending the plan-based model

- Add beliefs about visual salience of attributes
- Add notion of confidence in referring expression

Edmonds 1994

Recovering from misunderstanding

If the present text is unexpected or uninterpretable

then hypothesize a present or earlier misunderstanding

by self or other

Re-interpret or clarify

Example

MOTHER: Do you know who's going to that meeting?

Data from Terasaki 1976

Pragmatic ambiguity of *Do you know...*?

1. Asking for information.
 - *Yes, it's Sara and Nadia.*
2. Asking whether the hearer knows.
 - *Yes, don't worry, I have the full attendance list.*
3. Testing the hearer.
 - *It's Sara and Nadia, right?*
 - *No, you're wrong, it's Martin and Tim.*
4. Making a **pre-announcement** for some surprising information.
 - *Who??*
 - *The Rector, that's who!!*

Example

MOTHER: Do you know who's going to that meeting?

RUSS: Who?

MOTHER: I don't know.

RUSS: Oh. Probably Mrs McOwen and some of the teachers.

Data from Terasaki 1976

Collaborative repair of misunderstanding

- Repair of text-level misunderstanding
- Speaker and listener negotiate and refine meaning of prior utterance
- Integrates speaker-based and listener-based views of meaning
- Abductive-reasoning models of this process (McRoy and Hirst 1995)

There's more ...

- Many more situations for negotiation of meaning
 - All kinds of misunderstanding, misalignment of interpretations, misconceptions
- Elicitation and construction of knowledge (Regoczei and Hirst 1991)

Negotiated meaning wasn't
useful in CL in the 1990s

By 202X it will be necessary

Future roles of the linguistic computer

- Servant of the user
- “Neutral” reader and transformer of text
- Proxy for the world
- Proxy for the user in the world

Future roles of the linguistic computer

- **Mediates** between the user and the world
- Interprets the world to me
- Interprets me to the world

Conclusion

- Three loci of text-meaning
 - in text, in writer, in reader
- CL varies in its view
 - but has lately forgotten the writer and reader
- New applications will bring them back

Conclusion

- Further sophistication in text-meaning
 - Searching for and reconciling **different interpretations** of text
 - **Collaborative construction** of meaning in interaction and elicitation of knowledge

Future role of the linguistic computer

- Mediation and reconciliation
- Worldwide love and peace