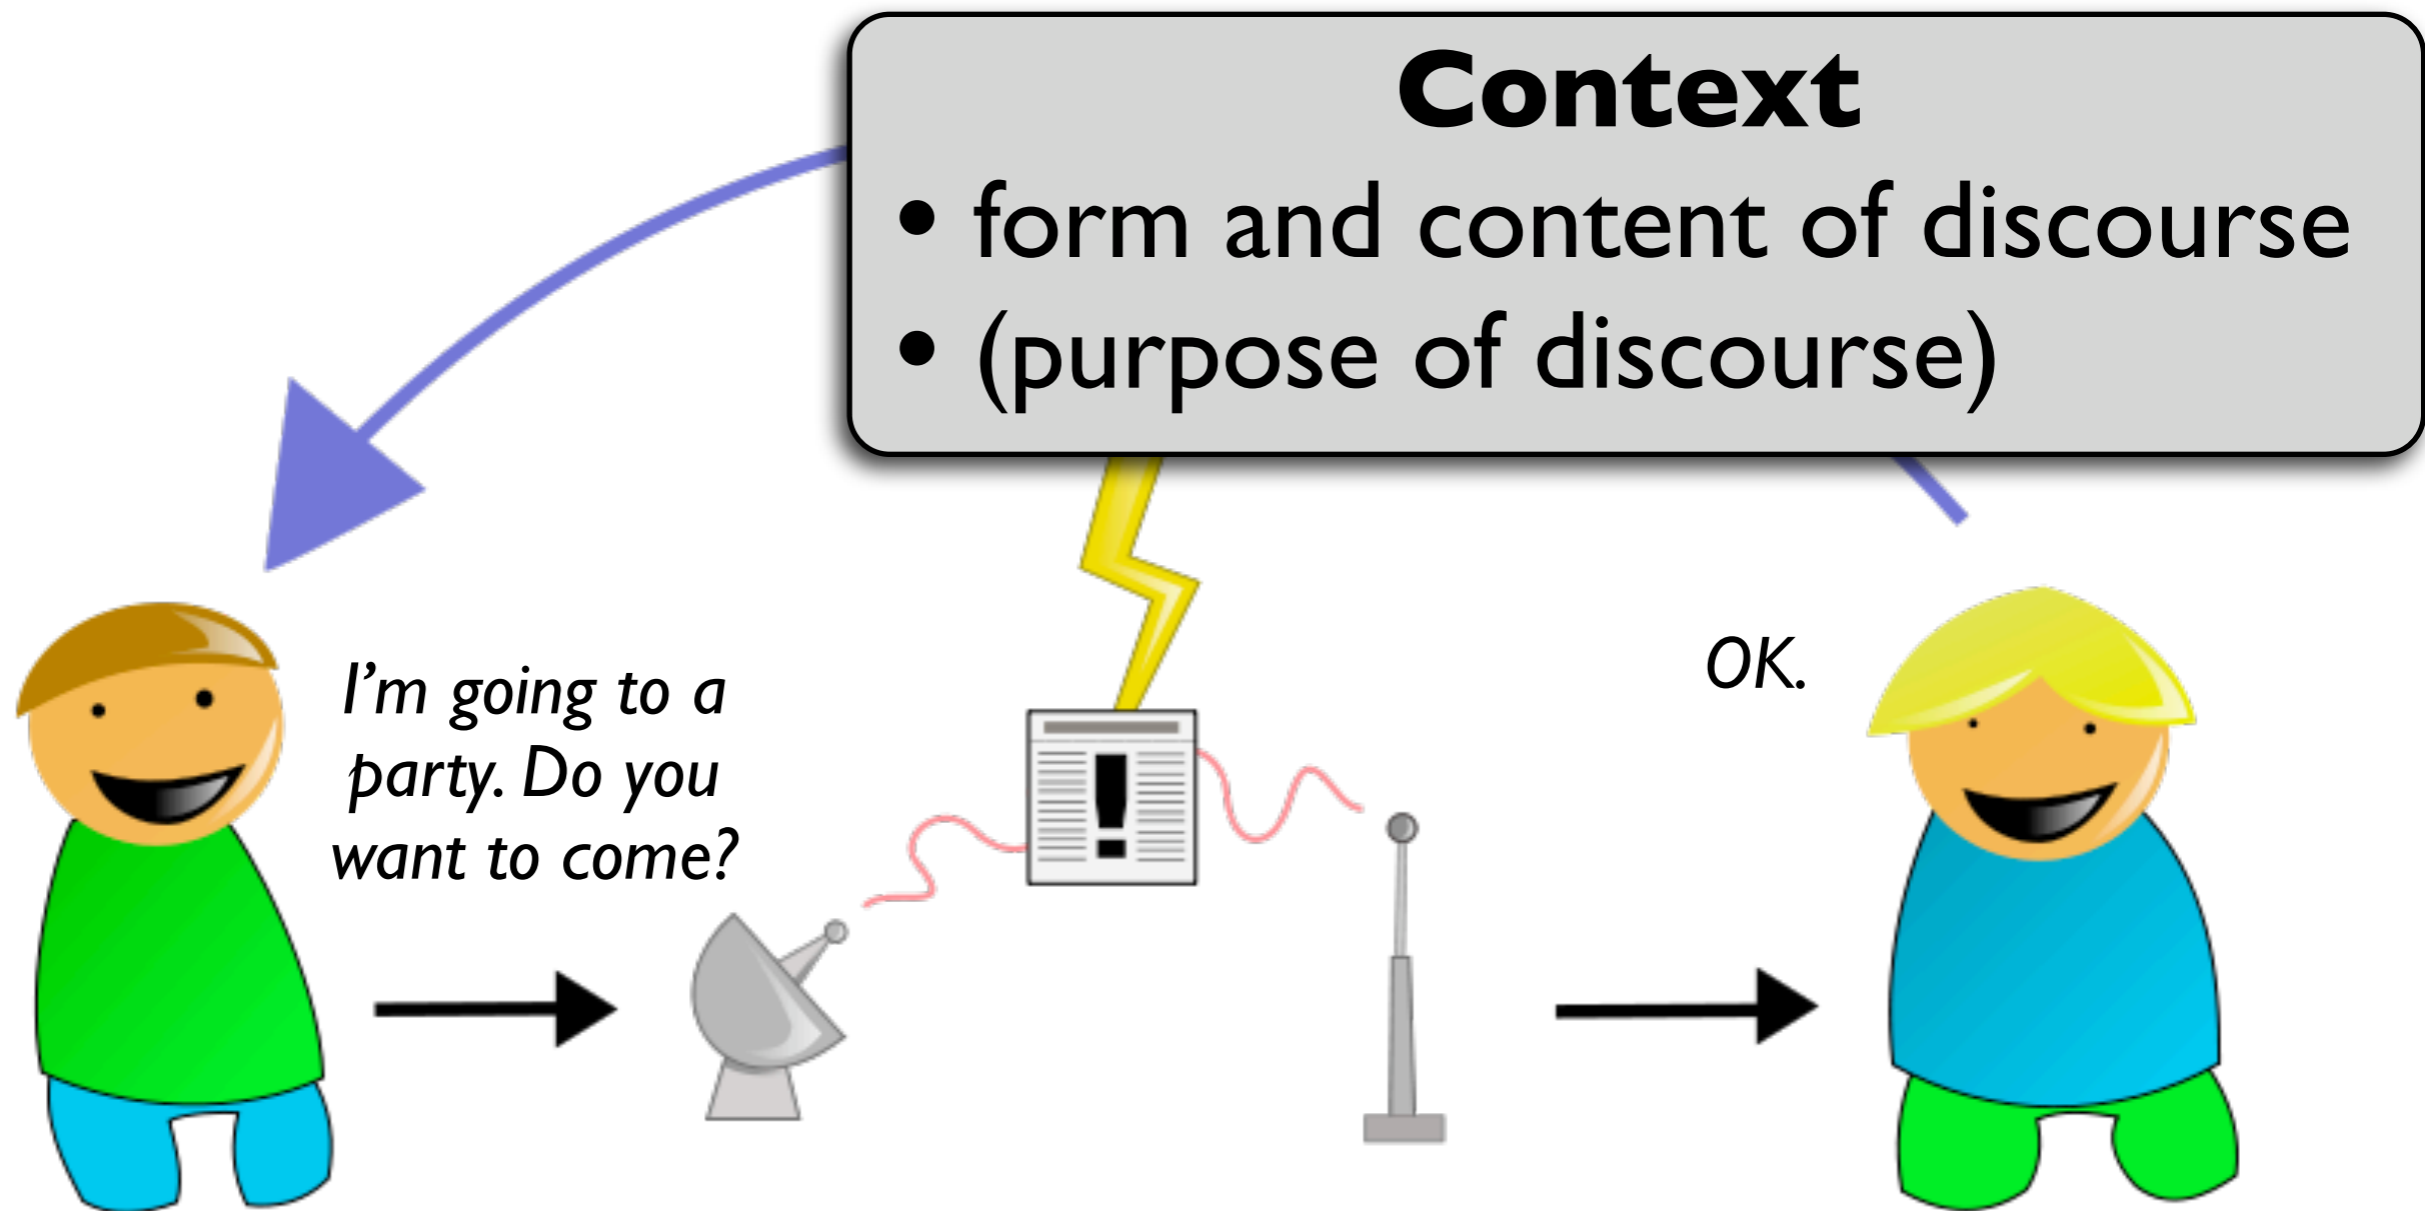


Generating Instructions in Virtual Environments

Session 4: Context in situated communication

Konstantina Garoufi
5 November 2009


Non-situated communication



Non-situated communication

Context

- form and content of discourse
- (purpose of discourse)



Situated communication



Appelt (1982)

Situated communication

Context

- form and content of discourse
- purpose of discourse
- objects of the scene in the visual field
- spatial configuration
- gestures, gaze
- history of interaction
- task at hand



The issues

- Which are the linguistic and extra-linguistic aspects of context that are important for situated communication?
- How do these contextual aspects interact with and guide NLG?
- In what ways can we model this interaction?

Outline

- What is context and why bother with it?
- A small chronicle of approaches to capturing the context in NLG
 - ▶ Classical linguistic context modeling
 - ▶ Multidimensional context modeling
 - ▶ Towards situated context modeling
- Conclusions

Outline


- What is context and why bother with it?
- A small chronicle of approaches to capturing the context in NLG
 - ▶ Classical linguistic context modeling
 - ▶ Multidimensional context modeling
 - ▶ Towards situated context modeling
- Conclusions

What is context?

Context is what constrains a problem solving without intervening in it explicitly.

Brézillon (1999)

IJCAI-95 Workshop on context in NLP



Some of the workshop's topics

- What is the relationship between formalization of context and natural language ideas of context?
- Which phenomena and inferences observed in natural language are context-independent and which ones always depend on context?
- How to automatically identify context-provided constraints resulting in conveying additional or different aspects of information?

Some of the workshop's main conclusions

- Context increases information content of natural language utterances
 - Context provides constraints on reasoning
 - Context facilitates learning
- but...**
- Context in NLP stays rather unexplored
 - Theory-versus-practice gap
 - Context is a big bag of issues that seem to be highly domain-specific
 - Works addressing context are too vague

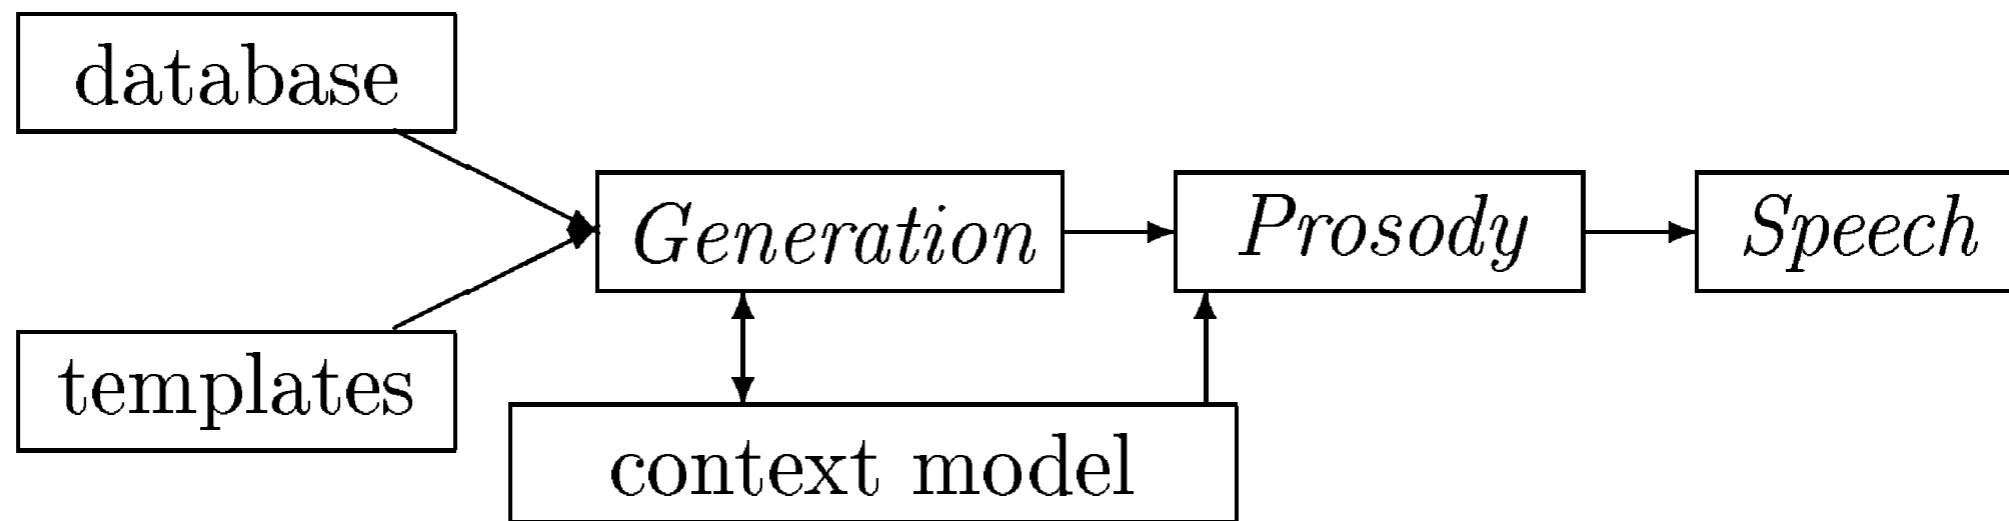
A couple of years later...

- Dial Your Disc (DYD) system
- One of the first NLG systems with a **dedicated** context model
- Generation of spoken monologues about W.A. Mozart's instrumental compositions

*van Deemter & Odijk
(1997)*



How DYD works



*The following+ composition+, from **which** you are going to hear a fragment+ of part one+, was written+ by Mozart in **his** Salzburg+ period, in Salzburg, in Brussels+, and in Paris+. The KV+ number of **this work**+ in c+ is K. six+. **It** is also+ known as the first+ sonata+ for violin+, and harpsichord+. The violin+ is played by Gerard Poulet+, and the harpsichord+ by Blandine Verlet+.*

Context modeling in DYD

Find a level of representation that is both rich and explicit enough to allow a system of rules to exploit the information in there for contextually appropriate utterances

- ▶ Set up a **data structure** and fill it with information
- ▶ Formulate rules that exploit this data structure



context model

DYD's context model

- **Knowledge state:** Which information has been expressed so far, and when?
- **Topic state:** Which topics have already been dealt with, which are still to be considered?
- **Context state:** Which objects have been introduced? How and when?
- **Dialogue state:** What recordings have been selected so far?

So what information does that encompass?

Submodel	Contains
Knowledge state	what, when, how told
Topic state	todotopics, donetopics, attemptedtopics, current topic
Context state	D-model, templates/words used, recently used concepts, location, alwaysgiven
Dialogue state	recording (part) selected; history of previous selections, whattotell

- both syntactic and semantic
- some generally required, some system-specific
- granularity subject to application
(here: speech generation \Rightarrow prosody important)

Two questions arising

1. But, of course, context models have also come up in other settings (AI, DRT, ...). Wouldn't it have been possible to re-use these context models for our purposes?
2. OK, so this looks like a reasonable model of context. Isn't our problem solved now?

Context modeling in AI: *ist* theory

- Goal: Introduce contexts as abstract mathematical entities with properties useful in AI
- Example:
ist(context-of("Sherlock Holmes stories"), Holmes is a detective)
- Axioms and rules for entering, leaving, lifting, transcending contexts

McCarthy (1993)

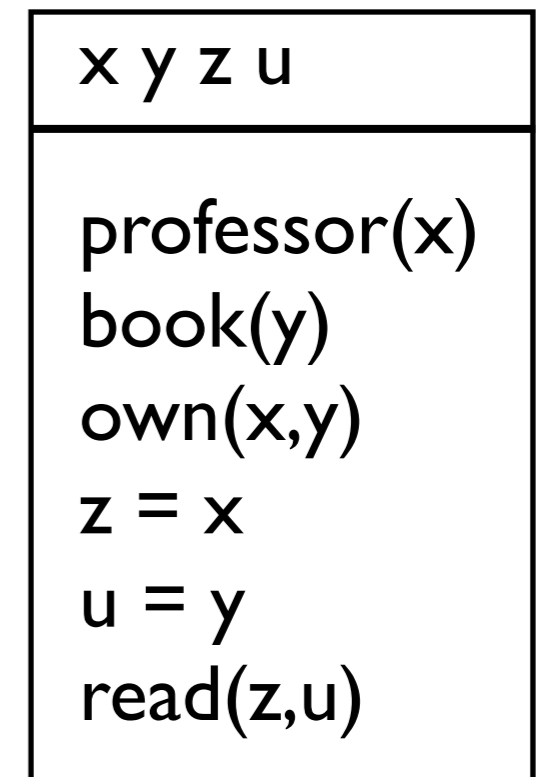

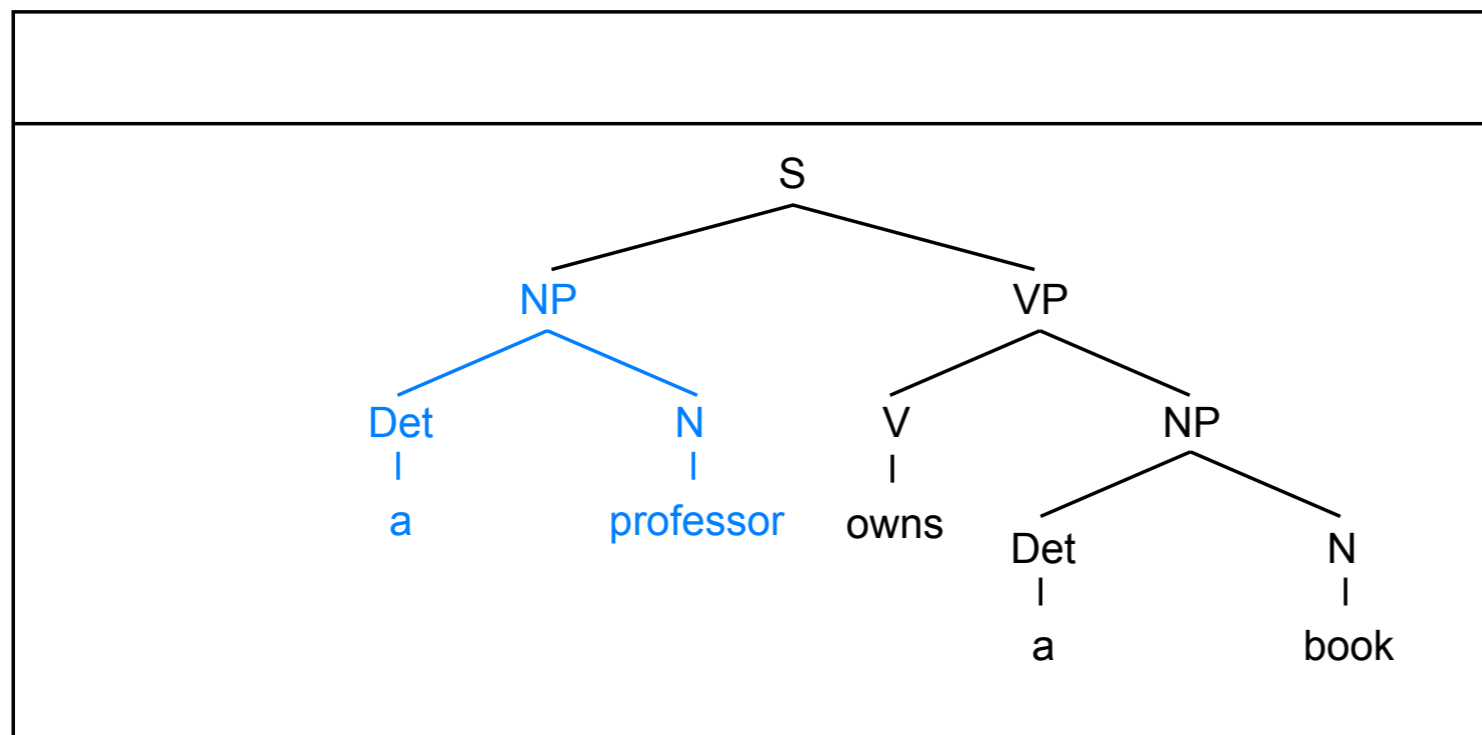
Problems with *ist*

- Context dependence in language is pervasive!
- Linguistic contexts change unexpectedly:
discourse entities added, objects and
expressions move into/out of focus
- We need context update and comparison
mechanisms suitable for linguistic contexts
- Hmm how about Discourse Representation
Theory (DRT) then?

Context modeling in DRT

Kamp & Reyle
(1993)

- *A professor owns a book. He reads it.*



DRS

Problems with DRT

- DRSs do not contain all the information we need. E.g. how do we deduce **topic** from a DRS?
- DRSs contain information irrelevant for generative purposes
- So what we really need (at least for DYD) is a computationally viable variant of DRT

Has DYD's context model solved all our problems?

- Consider the following text:

*M. Walker will give a presentation **later today** in the same room as where the opening session was held. **He** **is** **currently** in the coffee room, **just around the corner** and he might be an **interesting person** for setting up a project on ubiquitous computing.*

- Is DYD's context model sufficient here?


Outline

- What is context and why bother with it?
- **A small chronicle of approaches to capturing the context in NLG**
 - ▶ Classical linguistic context modeling
 - ▶ **Multidimensional context modeling**
 - ▶ Towards situated context modeling
- Conclusions

The Parrot-Talk NLG system

- Human agents in the physical world are supported by software agents
- Text is generated for output on a wearable device (parrot)
- Conference center application: parrots search for information and encounters with other users who share same interests

Geldof (1999)



Context dimensions in Parrot-Talk

- **Linguistic:** How far ahead in the discourse have objects been mentioned?
- Extra-linguistic
 - ▶ **temporal:** date, time
 - ▶ **physical:** how close is target user?
 - ▶ **social** implicature: what is target user doing?
- **User profile:** interest in which topics and persons?

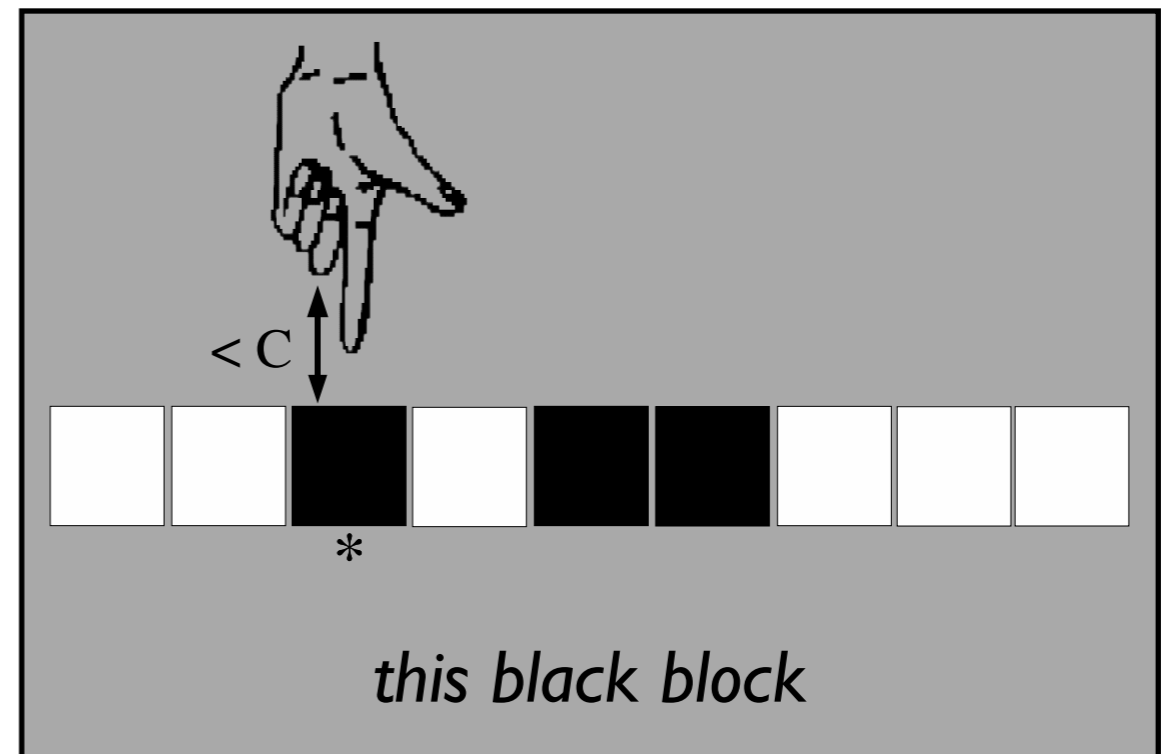
That's nice, but...

- We still seem to be far from a comprehensive situated context model
- What happened in situated NLG in the 21st century?

Multimodal context in GRE

Richer notions of multimodal context, with focus on REs

- ▶ deictic pointing gestures

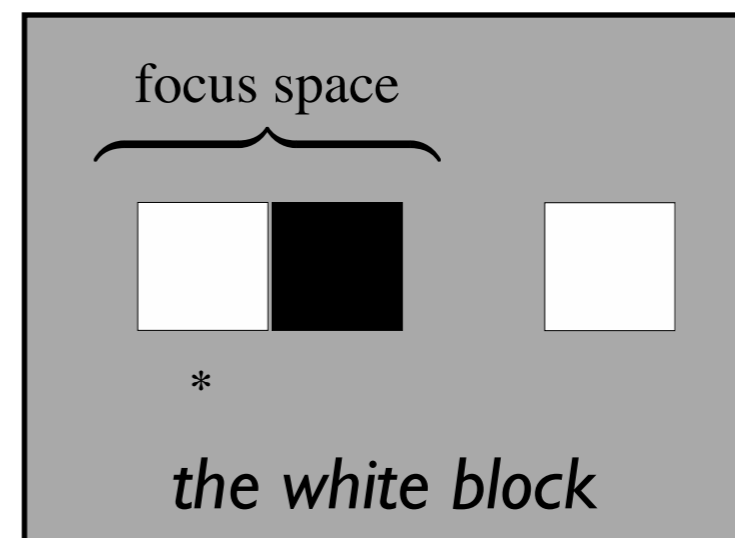



van der Sluis & Krahmer (2001)

Multimodal context in GRE

Richer notions of multimodal context, with focus on REs

- ▶ deictic pointing gestures
- ▶ current focus of attention

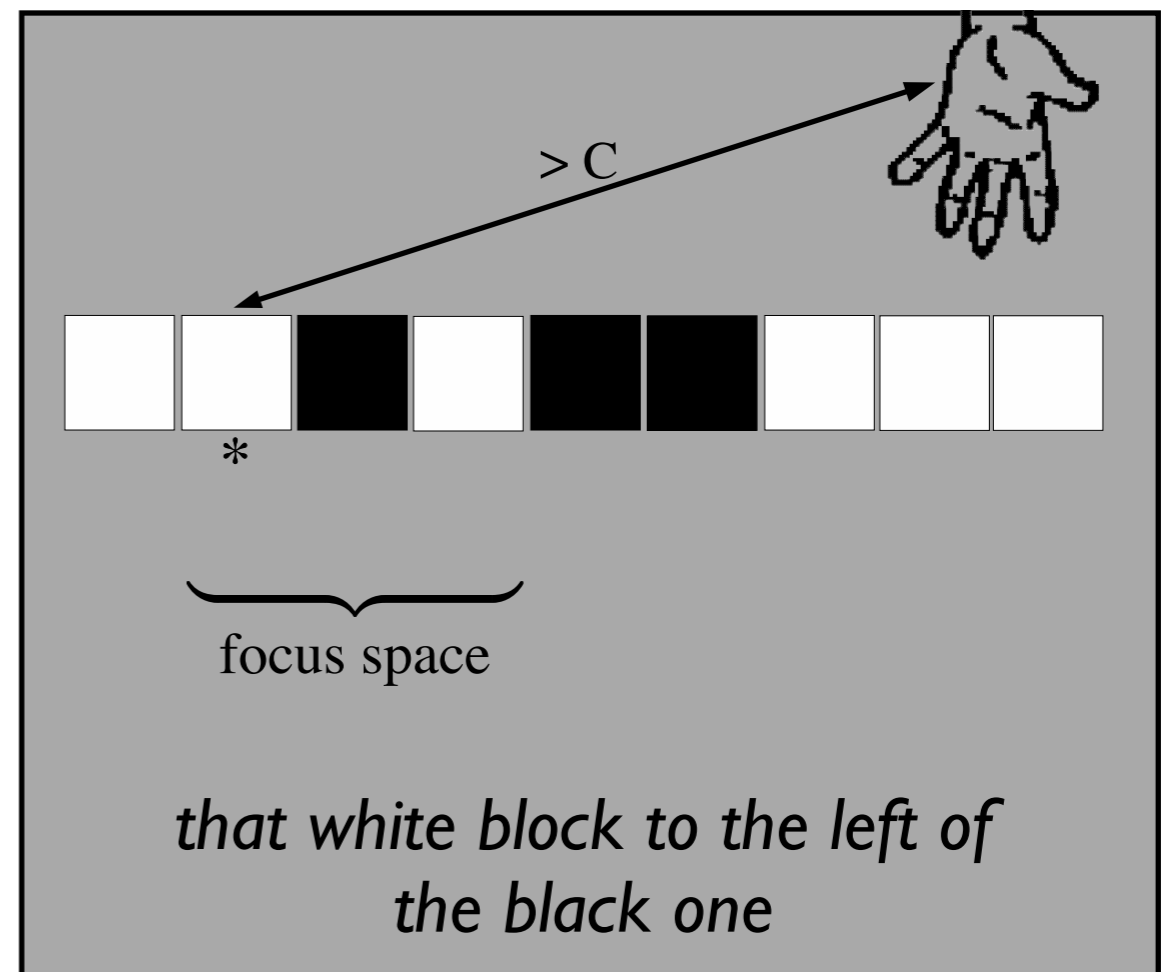


van der Sluis & Krahmer (2001)

Multimodal context in GRE

Richer notions of multimodal context, with focus on REs

- ▶ deictic pointing gestures
- ▶ current focus of attention
- ▶ three-dimensional salience: linguistic, inherent, and focus space salience




van der Sluis & Kraemer (2001)

Outline

- What is context and why bother with it?
- **A small chronicle of approaches to capturing the context in NLG**
 - ▶ Classical linguistic context modeling
 - ▶ Multidimensional context modeling
 - ▶ **Towards situated context modeling**
- Conclusions

A few years later



Byron & Fosler-Lussier (2006)

OSU Quake 2004 corpus of two-party situated problem-solving dialogs

- ▶ deictic and exophoric (i.e. situational) reference
- ▶ language calibrated against spatial arrangement of world
- ▶ perceptual limitations

A sister corpus: SCARE



- 15 spontaneous English dialogue sessions
- Each session records the joint problem-solving of a pair of human partners working through a treasure-hunt style task in a 3D virtual world

Stoia et al. (2008)


The SCARE corpus

- instruction giver (IG) guides instruction follower (IF) through completing tasks



*IF's view of the world,
as displayed on IG's monitor*

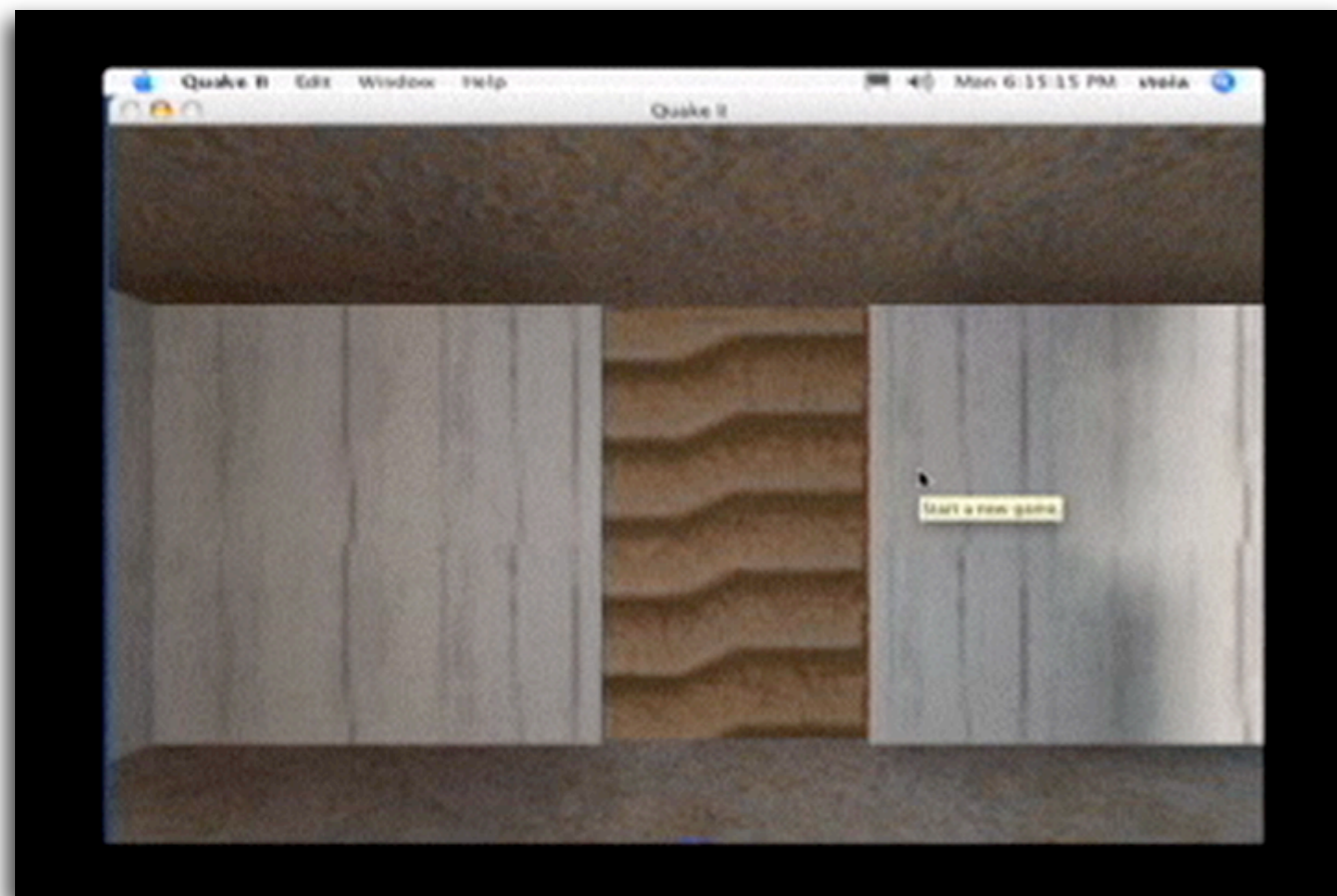
Stoia et al. (2008)



IG's map of the world



Example transliterated



walk forward and go through the first door you see [pause] and then go through the next one right in front of it [pause] yeah that one [pause] ok [disfluency - w] and then turn to your right [pause] and then hit the button in the middle [pause]



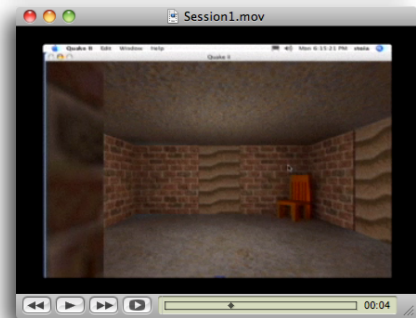
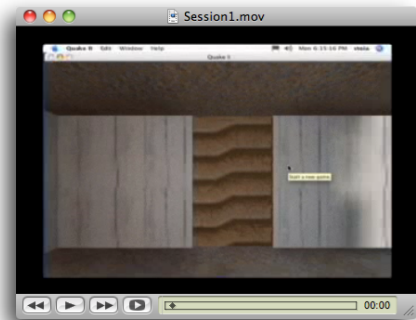
Example step-by-step

walk forward and go through the first door you see

and then go through the next one right in front of it

and then turn to your right

and then hit the button in the middle



} navigation

} referring expression generation

What is happening here?

- More than mere GRE!
- Looks like the IG is manipulating the extra-linguistic context of the discourse in a way that allows him to use a linguistic utterance of lower cognitive complexity
- How can an NLG system model that?

Outline

- What is context and why bother with it?
- A small chronicle of approaches to capturing the context in NLG
 - ▶ Classical linguistic context modeling
 - ▶ Multidimensional context modeling
 - ▶ Towards situated context modeling
- **Conclusions**

Conclusions

- We've come quite a bit of a way since 1995, but there's still a lot to be explored
- Fundamental questions about the nature of context in situated communication remain open
- No unified account of the various notions of situated context exists
- However, with GIVE, we now have an infrastructure for systematically addressing these issues