Computational Linguistics: From Language Tasks to Language Games

David Schlangen
Universität Potsdam
http://clp.ling.uni-potsdam.de
david.schlangen@uni-potsdam.de

Overview

I. What CL could be

II. What CL actually is (and how it does its thing)

III. How it could become what it could be

Overview

I. What CL could be

II. What CL actually is (and how it does its thing)

III. How it could become what it could be

using computers to do linguistics

learn sthg. about language

computational linguistics

using computers to do what (non-academic) “linguists” do

solve this one problem

using computers to do linguistics

computational linguistics

using computers to do what (non-academic) “linguists” do

learn sthg. about language

solve this one problem

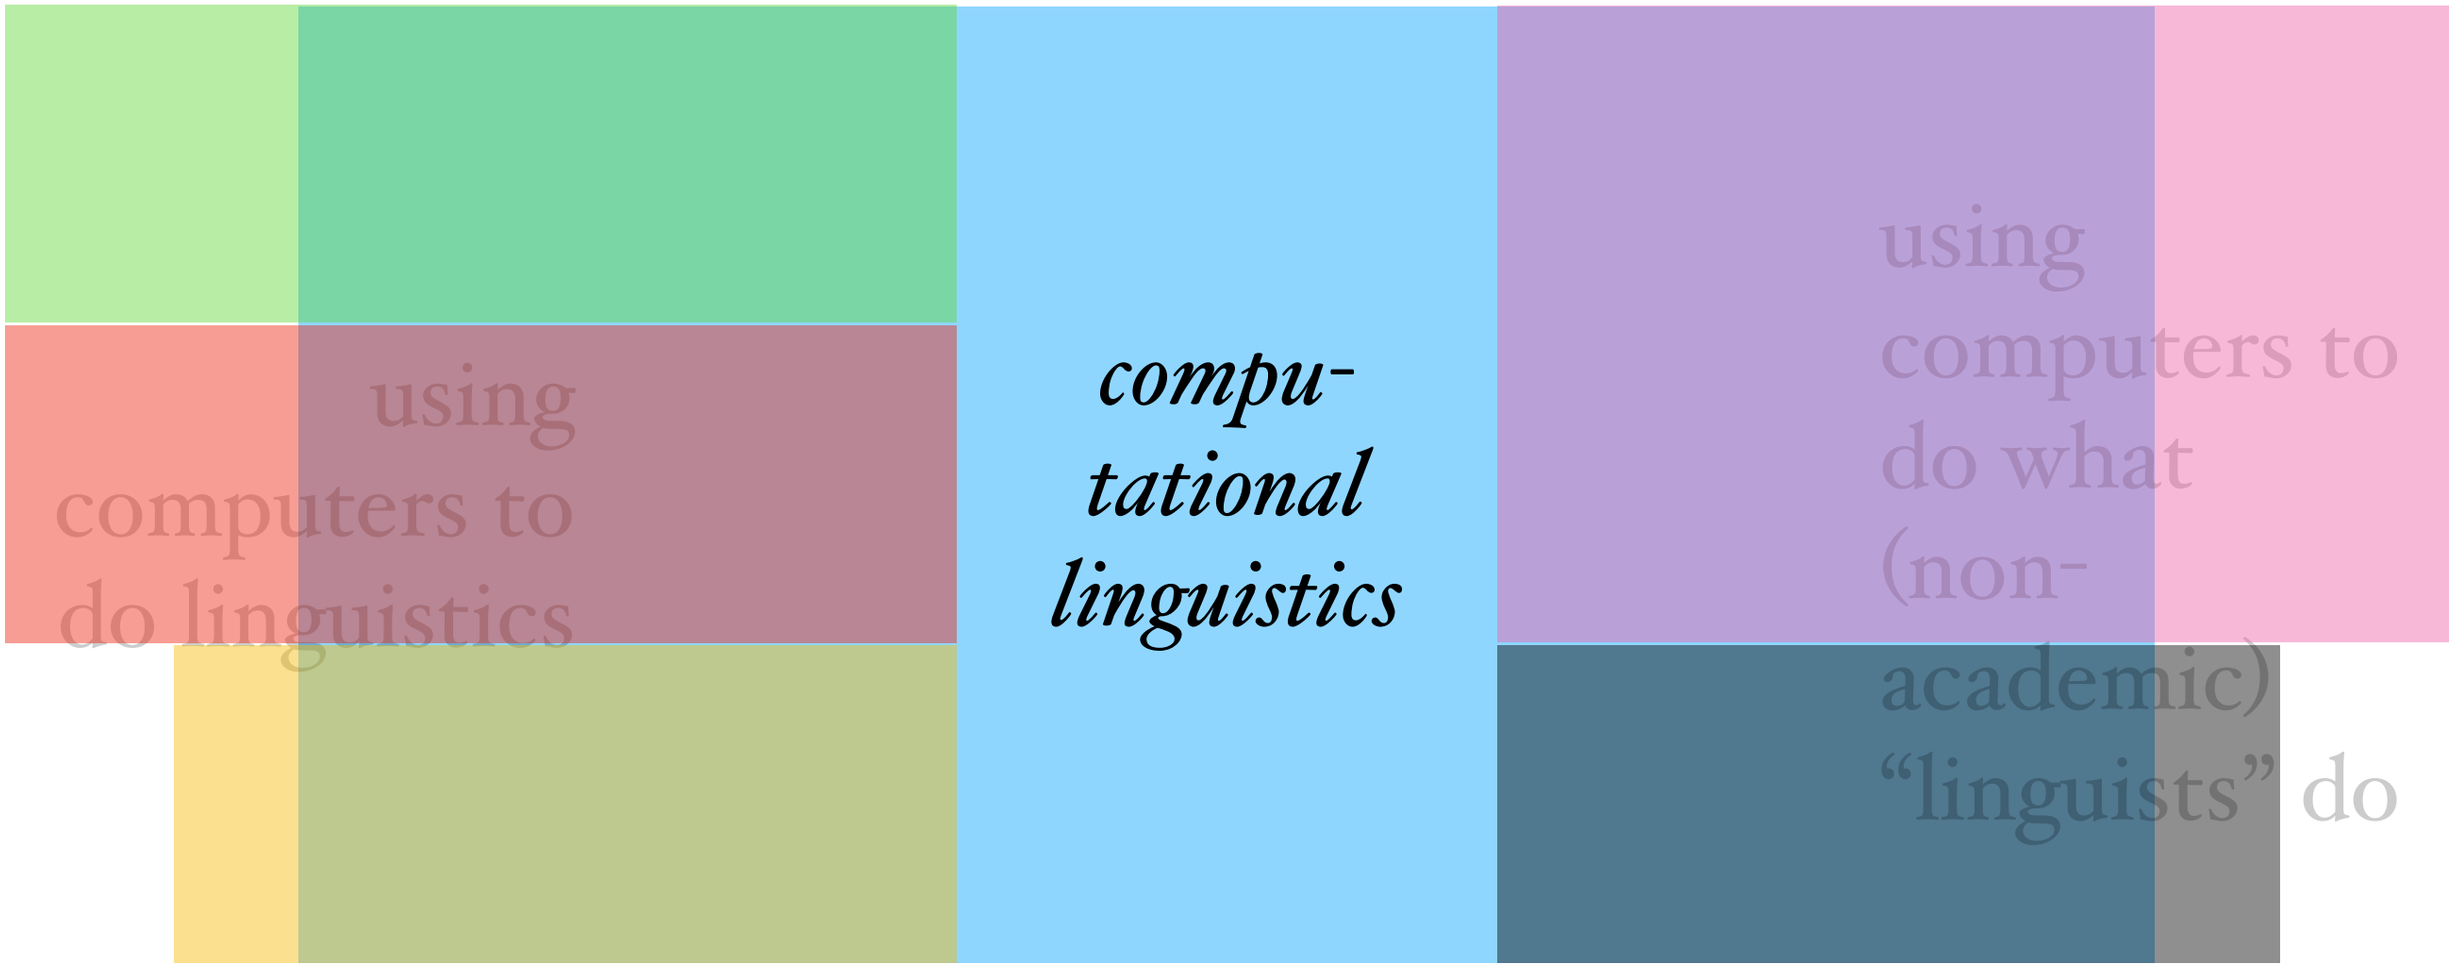

learn sthg. about language

learn about language use / users

solve this one problem

theoretical linguistics
(computability of structure)

psycholinguistics
(human computation of structure)

corpus linguistics
(manage ling.
data)

*compu-
tational
linguistics*

natural language processing
(get language stuff done)

digital humanities
(manage ling.
data)

learn sthg.
about
language

learn about
language
use / users

solve this one
problem

using computers to do linguistics

compu-tational linguistics

using computers to do what (non-academic) “linguists” do

learn about language use / users

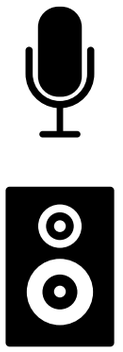

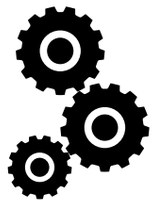

using computers to do linguistics

compu-tational linguistics

using computers to do what (non-academic) "linguists" do

learn about language use / users

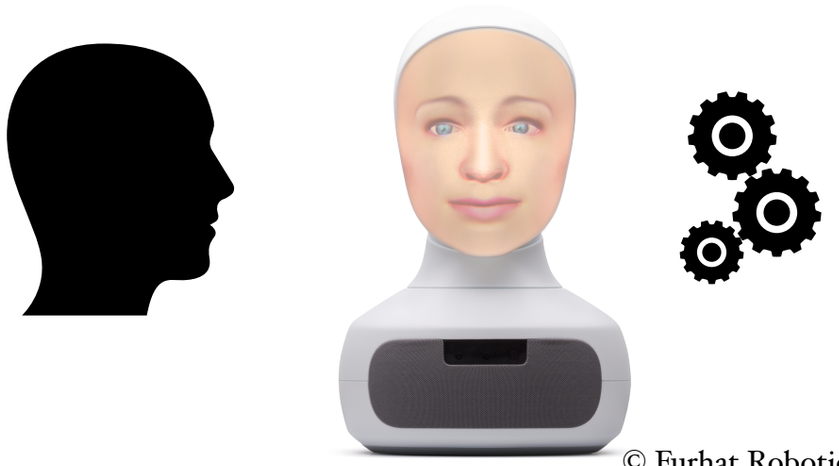

© Furhat Robotics

Overview

I. What CL could be

The field that studies language use
through the simulation of language users

II. What CL actually is (and how it does its thing)

III. How it could become what it could be

Overview

I. What CL could be

II. What CL actually is (and how it does its thing)

III. How it could become what it could be

The "Language Task Approach"

- CL (just like much of linguistics, and all of NLP) approaches its object of study in the form of small, well-defined *tasks*.

1. Translate the following text into German:
   How much wood would a
   woodchuck chuck […]

2. Translate the following text into German:
   Hold the newsreader's nose squarely,
   waiter, or friendly milk will countermand
   my trousers.

3. Fill in the gaps so that the result is a grammatical sentence:
   Colourless ____ ideas sleep _____.

4. Translate the following sentence into First Order Logic:
   Every fish owns a bicycle.

Modelling language tasks

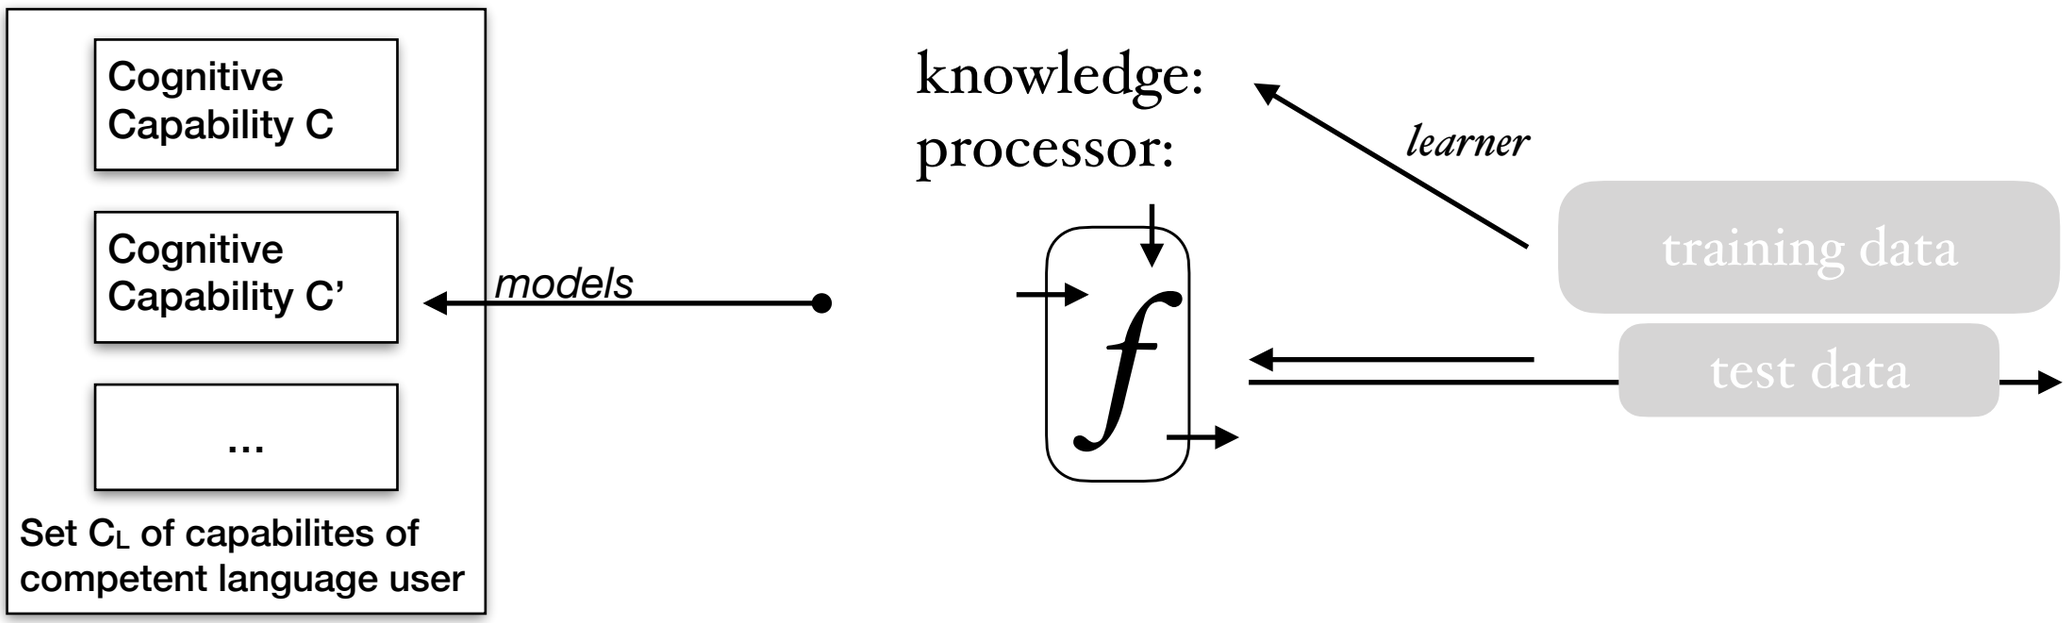

intensional task description: E.g., “translate the sentence”, or “determine the grammaticality of the sentence.”

extensional task description: Do what has been done here.

Example: Syntactic Analysis

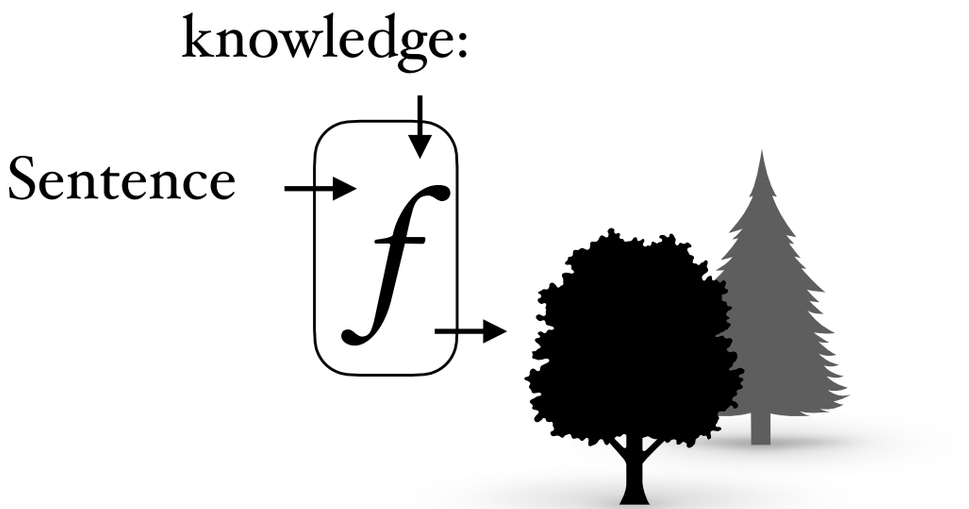

Example: Translation

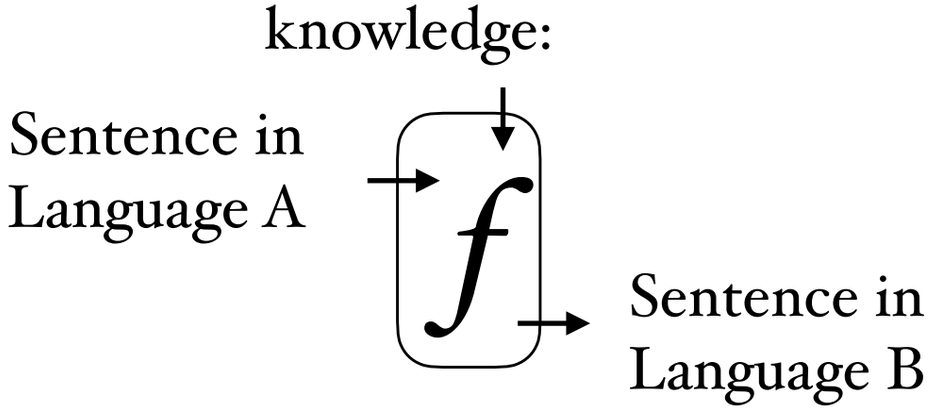

Example: Captioning

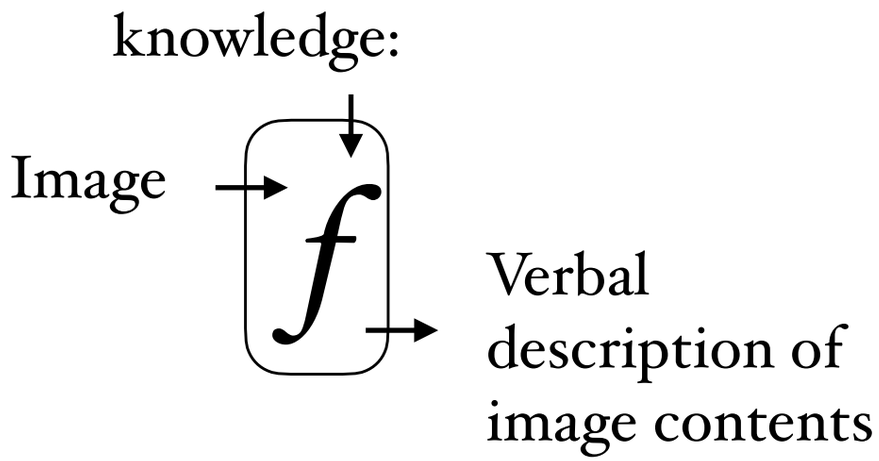

kitteh drinks out of coffee cup while glenn beck rages

MSCOCO corpus

Example: “Machine Reading Comprehension”

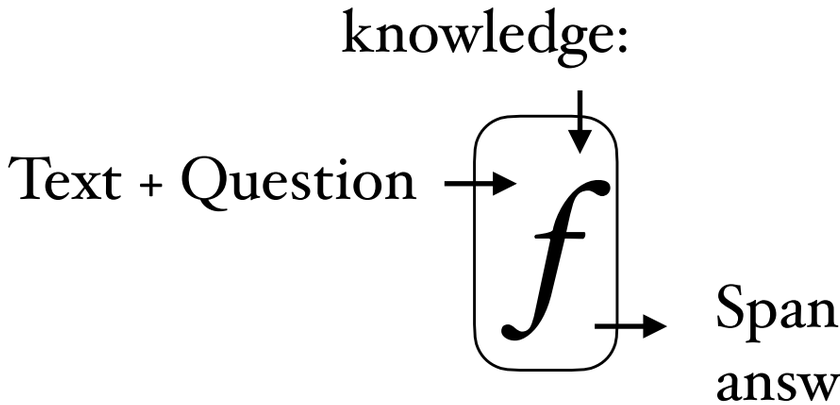

Linguistics is the scientific study of language.[1] It involves an analysis of language form, language meaning, and language in context[2], as well as that of the social, cultural, historical, and political factors that influence language.[3] Linguists traditionally analyse human language by observing an interplay between sound and meaning.[4] Historical and evolutionary linguistics focus on how languages change and grow, particularly over an extended period of time.

Which factors influence language?

Example: “Machine Reading Comprehension”

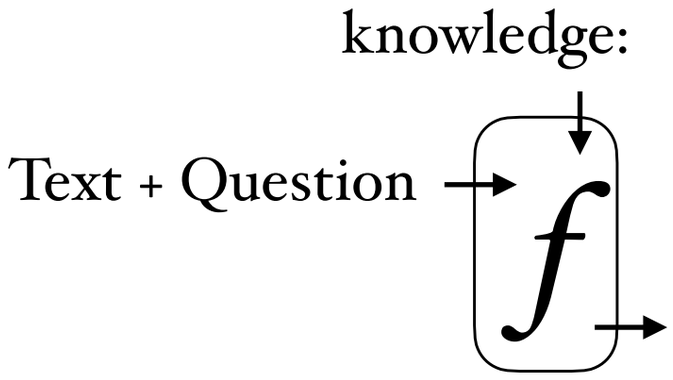

All these tasks are *transducers* that explicate information present in the input, by applying knowledge about how to do that.

Linguistics is
language in context[2], as well as that of the social, cultural, historical, and political factors that influence language.[3]
Linguists traditionally analyse human language by observing an interplay between sound and meaning.[4] Historical
and evolutiona

Can this scale up to modelling unrestricted / less restricted language use?

Which factors influence language?

Modelling language tasks

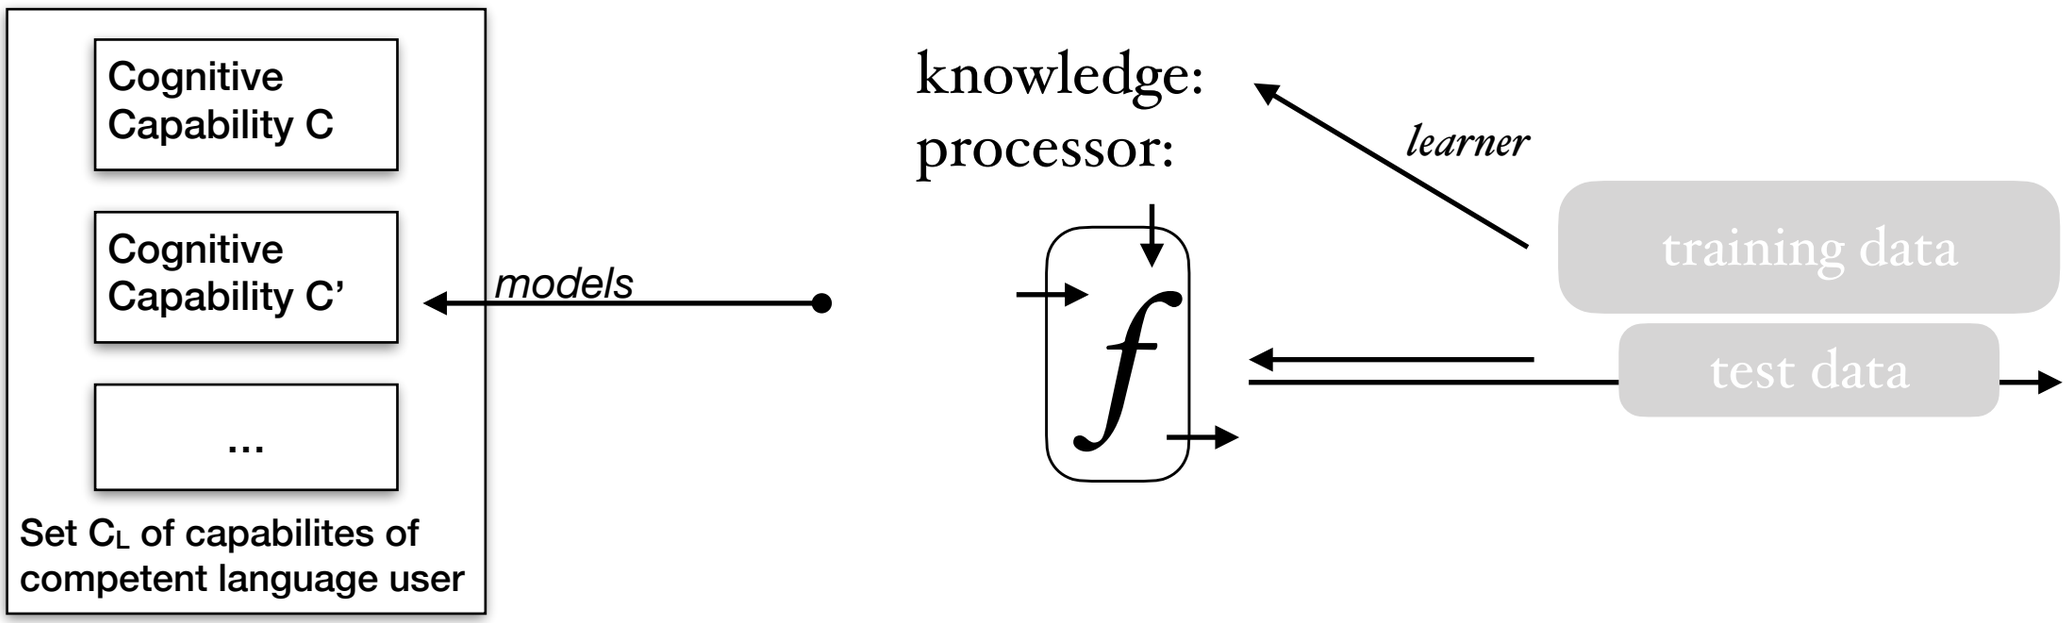

"Traditional" CL has a simple scaling up story: Model phonology, morphology, syntax, semantics, pragmatics … and you're done!

or "determine the grammaticality of the sentence."

It's much less clear where scaling up end-to-end tasks gets you.

Overview

I. What CL could be

II. What CL actually is (and how it does its thing)

III. How it could become what it could be

Overview

I. What CL could be

II. What CL actually is (and how it does its thing)

III. How it could become what it could be

Is Dialogue a Language Task?

A: Hey! B: Hello! A: I have a question. B: Sure, what's up?

Is Dialogue a Language Task?

A: Hey! B: Hello! A: I have a question. B: Sure, what's up?

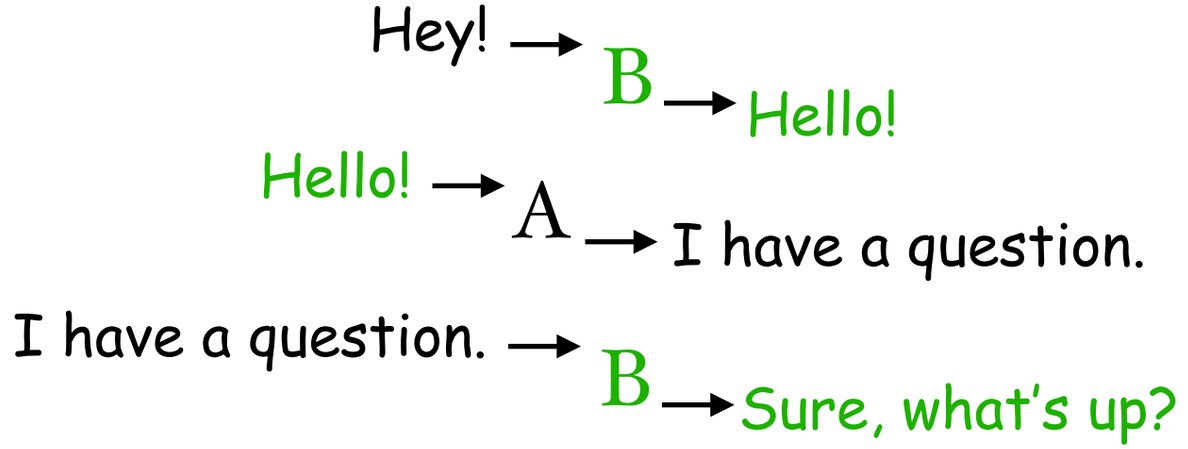

This mapping (*hello* to ..*question*..) is probably not a great idea in the general case. You need more context.

Is Dialogue a Language Task?

A: Hey! B: Hello! A: I have a question. B: Sure, what's up?

Hey! → B → Hello!

Hey! Hello! → A → I have a question.

Hey! Hello! I have a question. → B → Sure, what's up?

Is that still i.i.d.?

Still not enough context.
Different agents can react differently to same input sequences, due to *goals*, *personality*…

Is Dialogue a Language Task?

A: Hey! B: Hello! A: I have a question. B: Sure, what's up?

`cheerful; helpful` Hey! → B → Hello!

`ask_time` Hey! Hello! → A → I have a question.

`cheerful; helpful` Hey! Hello! I have a question. → B → Sure, what's up?

Just add personality, goals, etc. to context!

That doesn't seem right...
Can that make the right kinds of generalisations?

Is Dialogue a Language Task?

A: Hey! B: Hello! A: I have a question. B: Sure, what's up?

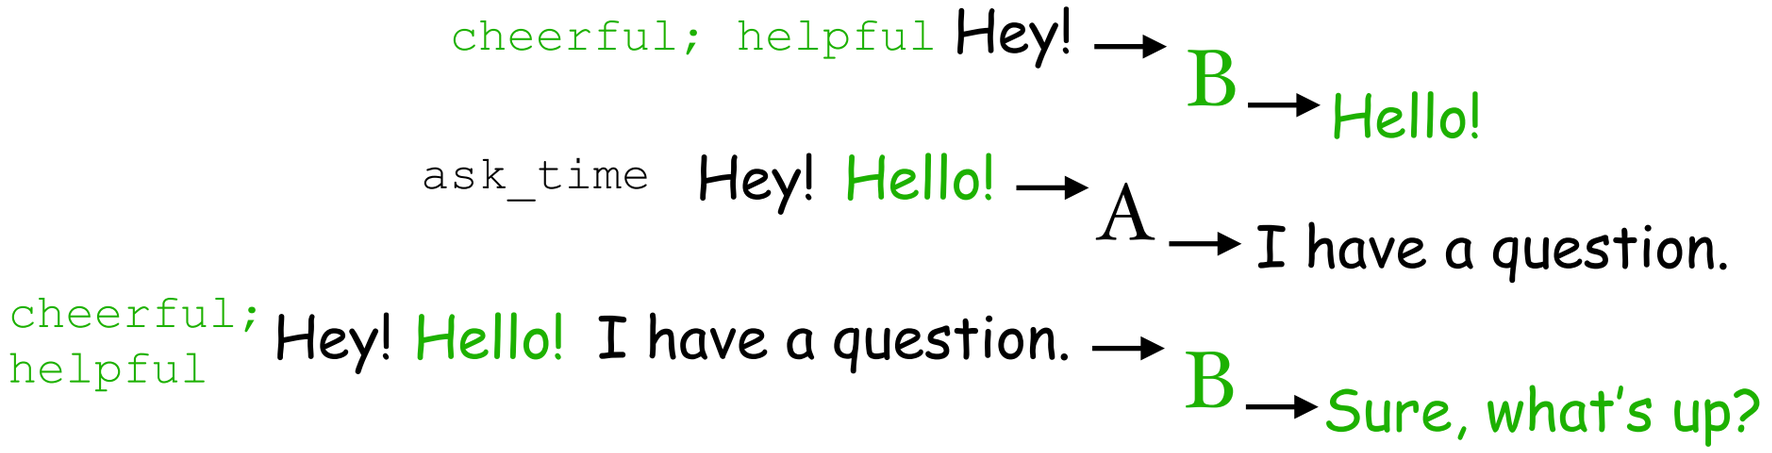

First observation: Whereas previously discussed language tasks were about *extracting* information from input, dialogue is about *adding* information (in the widest sense).

Second observation: This is just a *re-agent*, not an *agent*. We want action, not reaction.

Is Dialogue a Language Task?

A: Hey! B: Hello! A: I have a question. B: Sure, what's up?

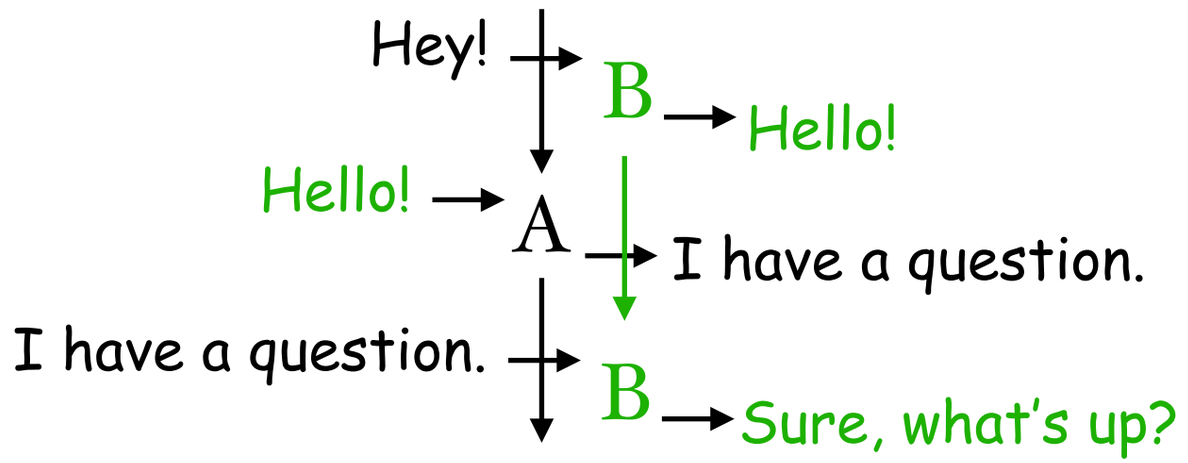

Seems more plausible to assume that some kind of state is kept.

In general, stimulus / response model seems inadequate. Agents make *decisions*, based on input and internal state.

The right level of abstraction seems to be to learn how to make these decisions...

And to learn not by generalising from observed examples, but from experienced interactions.

Is Dialogue a Language Task?

A: Hey! B: Hello! A: I have a question. B: Sure, what's up?

`cheerful; helpful` Hey! → B → Hello!

`ask_time` Hey! Hello! → A → I have a question.

`cheerful; helpful` Hey! Hello! I have a question. → B → Sure, what's up?

Then what is the role of data in learning dialogue behaviour?

Humans may partially learn from imitiation, but they also learn from just trying, and from getting help in reaching their goals.

(They also spend at least 10 years at this, with mostly fairly generous experts around to help them...)

How can we model interaction?

- Two challenges:
  - Can we retain some of the control that the task framing gives over the problem?
  - Can we avoid spending human language acquisition time on training interaction agents?

Language Games

- Dialogues / interaction episodes happen in concrete situations. These can be distinguished according to these factors:
  - The *environment* in which they are happening, insofar as it is important for the interaction.
  - The *setting* in which it takes place: Face to face, over the phone, etc.
  - The *game* that the participants are "playing", what they are trying to achieve.
- (Why separate environment & game? Think about board games: different games can be played on checkers board.)

Environments

AI Habitat, Facebook AI, https://aihabitat.org

Environments

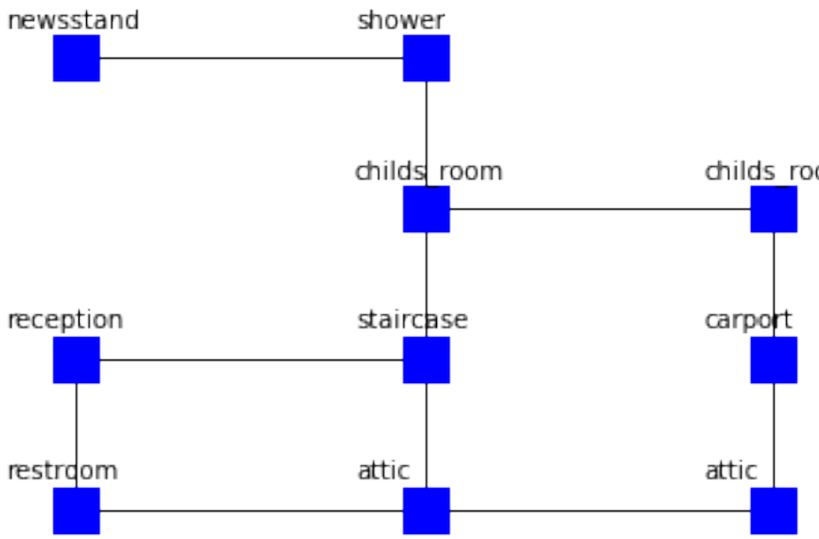

49	01:01	*GM (to A): You can go: [/e]ast [/w]est*
50	01:11	**A: I am in the basement**
51	01:11	**B: I'm in a basement.**
52	01:23	**B: Mine has a white staircase**
53	01:28	**A: no**
54	01:37	**A: mine has wooden stair case**
55	01:55	**B: Okay. Should I try to move towards you?**
56	02:09	**A: Sure**
57	02:11	**B: Wooden? What else?**
58	02:16	B (privately): e

MapWorld / MeetUp game (Ilinykh *et al.* 2019)
https://github.com/clp-research/meetup

How can we model interaction?

- Two challenges:
  - Can we retain some of the control that the task framing gives over the problem?
  - Can we avoid spending human language acquisition time on training interaction agents?

Strategy / Workflow

1. Define language game; be clear about what it is meant to target. Restrict it as much as possible, but not more.

2. Collect human / human data of game playing.

3. Model whatever you can as language task, where possible leveraging existing task da
m

4. Synthesise this into a system, potentially using explicit mechanism for deciding what the current state is and for keeping global state. Use human / human data on game to guide this (which doesn't have to mean to use it as training data).

5. Evaluate system with humans. Improve.

6. Treat system as "user simulator" and train next gen / less modular system using reinforcement learning.

using computers to do linguistics

compu-tational linguistics

using computers to do what (non-academic) "linguists" do

learn about language use / users

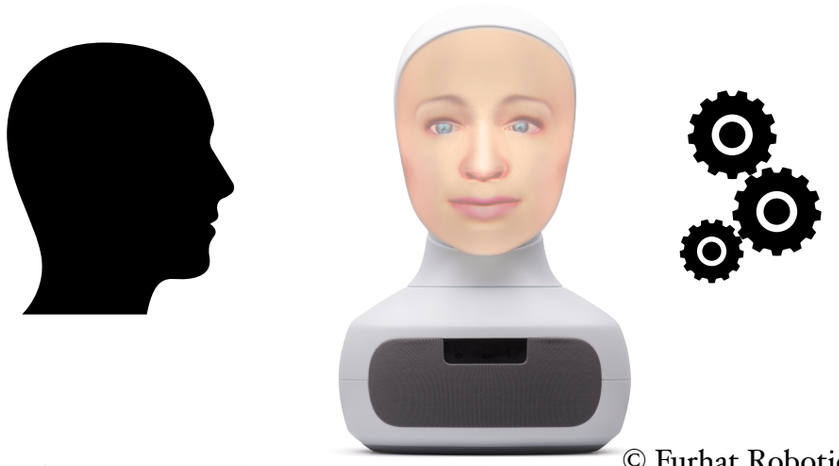

© Furhat Robotics

Conclusions

I. What CL could be

The study of language games

II. What CL actually is (and how it does its thing)

The (fascinating, and v. useful) study of language tasks

III. How it could become what it could be

By setting up interesting games, & by overcoming supervised learning on aggregate data (other agents' experience)

Thank you!

http://clp.ling.uni-potsdam.de

david.schlangen@uni-potsdam.de

Thanks also to my Phd students, Postdocs, & other collaborators for helping me develop and clarify this perspective.
(Timo Baumann, Okko Buß, Gabriel Skantze, Casey Kennington, Ting Han, Sina Zarrieß, Soledad López, Julian Hough, Nikolai Illinykh, Nazia Attari, Anne Beyer, Brielen Madureira, Robin Rojowiec, Fabian Galetzka, Sharid Loáiciga, Jana Götze, Philipp Sadler)

Gratefully acknowledged: Funding from DFG (Inpro, DUEL, RECOLAGE), CITEC / Excellence Initiative, SFBs Alignment in Communication, Limits of Variability.