

Meaningful Interaction with Unreal Speakers?

David Schlangen

Computational Linguistics, Department of Linguistics

University of Potsdam

david.schlangen@uni-potsdam.de

Abstract

The last time that the semdial workshop took place in Bielefeld was almost 25 years ago, in 2001. Incidentally, this was also where my own first academic presentation happened (Schlangen et al., 2001). This is too much symbolism to ignore, so I will use the occasion of “coming back” to reflect on what happened in these (almost) 25 years, to my research on “(formal and computational approaches to) the semantics and pragmatics of dialogue”, and the field in general. (Because, oh boy, did something happen.)

The semdial workshop series, at least in my understanding, was founded on the idea that bringing together formal, empirical, and computational approaches to the study of dialogue would be possible in a certain way: Formal studies would be informed by empirical studies and in turn help guide computational modelling attempts, which would not need distinguish very clearly between being models of cognition and being human/computer interfaces, and in either way would be evaluated for how closely they mirror empirical findings. This particular way of setting up the relations between the constituent parts has been strained for a while now, with computational modelling becoming more and more “empirical” and machine learning-guided. But it has been exploded by modelling approaches that not even pretend to pay attention to any prior knowledge on the semantics and pragmatics of dialogue, and still (apparently?) succeed better than anything before in “modelling dialogue”. (Yes, I’m talking about “chat optimized LLMs”.)

In my talk, I will try to pick up the pieces, and hopefully show how they can be reassembled: First, I will show that the kind of analyses that we do are useful to understand the status of these “unreal speakers”. In particular, I will analyse the speech act of “assertion”, and show that LLMs perform an atypical variant of it, that in its consequences and how it relates to “real assertion” is not yet well understood. If this analysis is correct, this gives us an interesting new task, which is to devise a normative pragmatics of how the semantics and pragmatics of dialogue with machines ought to be understood and designed. As a direct consequence of this, the second part will make the claim that now that we see what “human-likeness” of human computer interfaces can lead to, we need to be more explicit about our goals for designing interfaces, and especially about how to separate desirable properties (ease of use) from potentially undesirable (blurring the boundaries between real and unreal speakers). In the final part, I will talk about how in my research group we set up the relation between non-computational models of cognition and computational behavioural models. In particular, I will talk about our “clembench” framework for evaluating LLMs through Dialogue Games, and very recent results on post-training of LLMs in this framework.

References

- David Schlangen, Alex Lascarides, and Ann Copestake. 2001. Resolving underspecification under discourse information. In *Proceedings of BI-DIALOG, the 5th Workshop on the Semantics and Pragmatics of Dialogue*, pages 79–93, Bielefeld, Germany.