

**CS 672 — Meaning Machines**  
Guide for breakout sessions  
Oct 27 2004

Consider two dialogue fragments that Brennan and Clark present in detail in their paper:

Director: a docksider

Matcher: a what?

Director: um

Matcher: is that a kind of dog?

Director: no, it's a kind of um leather shoe, kinda preppy pennyloafer

Matcher: okay, okay, got it

Director: a shoe

Matcher: uh, which shoe?

Director: ooh, forgot about that, um it's going right, it's red

Matcher: okay

In these dialogues, the Director and Matcher pursue meaningful interactions that are both PERCEPTUALLY and SOCIALLY GROUNDED.

The speakers' descriptions draw on their abilities to categorize objects, to recognize the differences among objects, and to express the distinctive properties in words. The director, for example, categorizes the first shoe as a docksider, a shoe, and a pennyloafer, and categorizes the second as a red shoe going right. Similarly, the matcher's perceptual abilities are implicit in the suggestion that the director might be referring to a dog in the first dialogue or the suggestion that the director might be referring to any of a set of shoes in the second dialogue. These capacities, in turn, implicate the use of symbolic mental representations whose content is causally determined in part by perceptual connections—perceptually grounded representations.

But these interactions ALSO reflect the social abilities and relationships of the interlocutors. The participants must understand their own and others' actions, they must act collaboratively, and they must understand their commitments and responsibilities to one another in doing so. Concretely, in asking for clarification, the matcher holds the director responsible for using *docksider* to refer to a property; the director follows up by describing what a docksider is. Likewise, when the matcher points out that *a shoe* is ambiguous, the director not only provides an alternative description, the director APOLOGIZES for failing to respect the socially grounded context. While the causal grounding of mental representations is a prerequisite for explaining actions semantically in terms of the agent's information and values, the social grounding of conversational action seems to be a prerequisite for attributing non-natural meaning to speakers.

The objective of our breakout sessions on Wednesday will be to organize our understanding of the bases for our attribution of meaning to one another, in light of these examples, our intuitions and expertise from our home disciplines, and the shared reading we've done so far this semester. I propose to proceed from the specific to the general, in three steps.

1. We begin by explaining PRECISELY what's going on in Brennan and Clark's two dialogues. What perceptual information do the interlocutors have to have? What do they have to represent in the context? How does this information come together in the interpretations of their

utterances in dialogue? How does the context evolve as the conversation proceeds? How might the interlocutors choose their contributions? How might they understand them? How might they coordinate their evolving representations of the context?

2. We then turn to a more GENERAL perspective. What general kinds of information do agents need to keep track of? What kinds of representations should agents generally maintain if they are to participate in these dialogues? And what should they do with these representations? What general capabilities are implicated in the explanations you have given? How are these capabilities related? In other words, what kind of cognitive architecture is required to support interactive, situated conversations such as these?
3. Finally, we can think CRITICALLY about the literature in semantics, cognitive science and artificial intelligence, in light of the intuitions and models of conversation we develop. How might philosophical perspectives — including everything from Winograd & Flores (1986) to Fodor's (2001) — be informed by a specific model of our own meaningful communication? What's missing from models of linguistic meaning — including everything from Montagovian model-theoretic semantics (1974) to the state-of-the-art dynamic semantics that grows out of Stalnaker's (1978) work on assertion? How can the challenges that we face in computational semantics be framed more clearly and more insightfully by a more nuanced understanding of our intuitions about meaning and a broader account of the role of meaning in conversation?

So, each group should think about formulating a specific analysis of Brennan and Clark's dialogue, an abstract, high-level specification for an agent that could participate in them, and an interdisciplinary research agenda outlining key difficulties in realizing this architecture.

To prepare for the meeting, take a look at the dialogues, and make sure you have something precise to say about an aspect of the dialogue that you're familiar with — perception, interaction, language, computation. Make sure you're prepared to describe what is happening in the dialogues and discuss models of how this process could be realized.

This isn't a quiz: I expect the understanding we arrive at to be new—to build on our common reading this semester but also to reflect the broader and more disparate background we bring to that reading from our own interests. Nevertheless, I've obviously thought about these issues a lot, and you can probably find echoes of the kind of analysis I would have expected coming into the class in the *Cognitive Science* paper I distributed on the first day. It's now appeared, so before class Wednesday you might enjoy returning to the print version:

<http://dx.doi.org/10.1016/j.cogsci.2004.05.001>

I think you'll find it a lot more approachable now that we have explored the key background on which the paper builds.