Foundations of Language Interaction

HANDOUT SEVEN

July 19, 2001

M. Stone
mdstone@cs.rutgers.edu

1 Introduction

Today we look at describing CONVERSATIONAL DYNAMICS from the point of view of modern computer science and artificial intelligence research. Recall the kind of framework we're using:

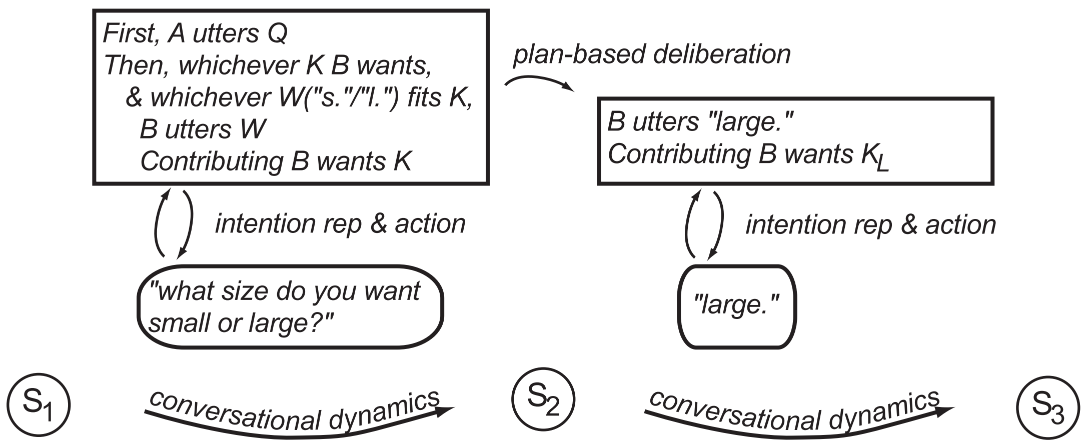

For the moment, we'll abstract away from the communicative intention representations and plan-based deliberation we've been investigating in detail: it's summarized at the level of conversational dynamics. How do we describe conversational dynamics itself, using empirical methods?

2 Some mathematical ideas

First principles might suggest modeling conversational dynamics as a PARTIALLY-OBSERVABLE STOCHASTIC GAME.

(1) A GAME is simply a mathematic formalization of a problem of coordination or competition among agents, where the success of any one agent depends not only on what it does but also on what other agents do.

(2) This game is STOCHASTIC, or governed by probabilistic rather than deterministic laws, because it takes place in a noisy environment in which the effects of actions depend on unreliable components like speech recognizers.

(3) The same noise makes the state of the environment PARTIALLY-OBSERVABLE. Agents must choose their actions based on incomplete information about the world.

There has been some limited work on representing computational problems directly in these terms, such as [Koller and Pfeffer, 1997]. A solution is an equilibrium—a strategy for each participant where the play of each can't be improved if the other plays the strategy. (A wrinkle is that coordination problems typically have multiple equilibria, so players must know what equilibrium to play.) However, practical dialogue research has backed off from these models, which require a lot of parameters and quickly become intractable.

(4) A standard assumption for dialogue research is that the whole community "plays" more or less the same policy, defined by conventions of English. Finding your equilibrium means matching the strategy.

Assumption (4) reduces the problem to a PARTIALLY-OBSERVABLE MARKOV DECISION PROCESS (POMDP), in which you don't have to anticipate other agents' strategic decision-making, and you can analyze dialogue using plain decision theory [Kaelbling et al., 1998]. But POMDPs are intractable too, and further simplification is required.

(5) Walker uses simplified, structured state-representations to finesse the problem of partial observability [Walker, 2000]. This MDP approach emphasizes emphasizes strategy and unpredictability.

Other simplifications are possible.

(6) Horvitz and Paek reason about the value of information in a way that's more faithful to the partial observability of dialogue state, but they accomplish this by giving up on solving strategic problems and using local decision-theoretic heuristics [Horvitz and Paek, 1999, Horvitz and Paek, 2001].

Anyway, from the perspective of computational linguistics, the important thing is not the mathematical formalism you use to solve the problem but the assumptions that you make about linguistic representations in building your model and the techniques that you use to bring the model into correspondence with human behavior.

3 An experiment we're planning for the fall

People answer questions in the terms in which they are posed.

(7) a Q: What size do you want: small or large?
b A: Large.

An example of entrainment. See [Brennan and Clark, 1996].
Dialogue sytems can use this fact to get better performance. The smaller the grammar or the more constrained the vocabulary, the better they are at recognizing what has been said. And good recognition is the most important thing for user satisfaction.
On the other hand, it can be equally frustrating to have to listen to a long list of options when you know how to say what you want—and pretty unnatural to always have to "barge in" to get your task done quickly with a dialogue system. So the system should try to make its questions short whenever possible.

(8) There is a tradeoff of speed and naturalness vs. recognition accuracy between asking short questions (*what size do you want*) and asking verbose questions (*what size do you want: small or large*).

It's an interesting problem because entrainment means you don't always have to be explicit in a context to get the answer you expect.

Here's a simple mathematical model.

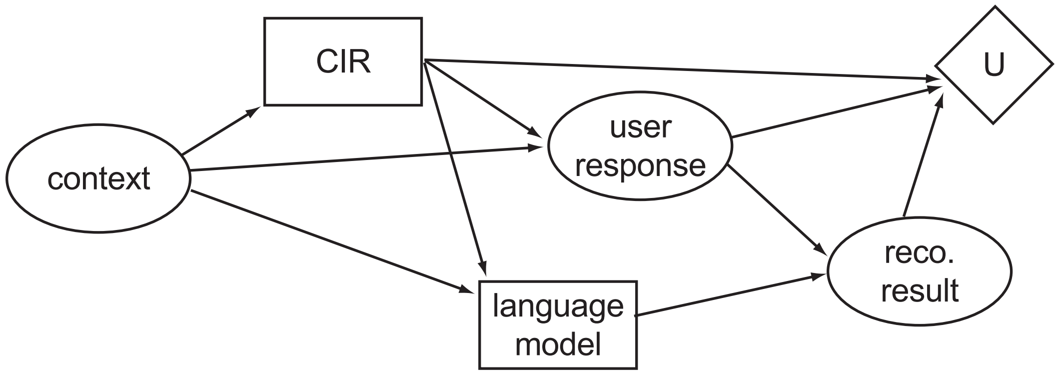

The model is represented as an influence diagram [Shachter, 1986]. Circles represent random variables. Squares represents choices that we can make. (So we're matching strategies here, not finding an equilibrium.) The diamond gives the overall payoff associated with the event. Arrows into circles and diamond indicate probabilistic dependencies. Arrows into squares indicate information avialable for a decision. Here we have

(9) a Context: the value of the dialogue state relevant to entrainment, for example, how have we identified the potential answers previously in the conversation.

b CIR: the communicative intention representation we choose for the question (short form or long form). We get to choose this based on the context, and in general we can specify the dependence nondeterministically in the terms of linguistic theory and plan-based deliberation that we've been developing so far.

c Language model: the grammar or description of possible utterances that we provide to the speech recognizer to describe the ambiguities it should consider in recognizing the user's utterance. For our toolkit we have to pick a symbolic grammar, and we have a few to choose from that we can select based on the context and the question we ask.

d User response: what the user actually says (realized, we might suppose, as a communicative intention representation).

e Recognition result: what we understand, though in practice what suffices is probably what the user meant, something the user didn't mean, or total recognition failure.

f Utility: depends on the difficulty the user has understanding the system, the difficulty the system has understanding the user, and the overall ability of the user to solve her task more easily with the system (as in Walker's PARADISE framework).

What we'd like to do is fit the parameters for this model based on what people actually say and use the model to design a strategy for asking questions in a dialogue system.

(10) With four conditions for running subjects—entrainment from context or no entrainment from context, long question or short question—we can collect the distribution of user responses. Of course we have to code this by hand! Otherwise dialogue systems would work already.

(11) We can reuse the corpus of user responses to understand the error rate of the speech recognizer under different grammars (and, effectively, under different distributions of user response).

(12) As Walker does, we can ask the users what they thought of the system and build a correlation model to assess what's important to them.

That gives a full model. Solving it is then actually pretty easy. Hopefully we'll get the solution we expect – but we'll get it through a quantitative model that says how good it is and that we can validate by running the final system. If so, we'll have learned something about dialogue as well as making a better system.

References

[Brennan and Clark, 1996] Brennan, S. E. and Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 22(6):1482–1493.

[Horvitz and Paek, 1999] Horvitz, E. and Paek, T. (1999). A computational architecture for conversation. In *User Modeling Conference*, pages 201–210.

[Horvitz and Paek, 2001] Horvitz, E. and Paek, T. (2001). Harnessing models of users' goals to mediate clarification dialog in spoken language systems. In *User Modeling Conference*.

[Kaelbling et al., 1998] Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101.

[Koller and Pfeffer, 1997] Koller, D. and Pfeffer, A. (1997). Representations and solutions for game-theoretic problems. *Artificial Intelligence*, 94(1):167–215.

[Shachter, 1986] Shachter, R. D. (1986). Evaluating influence diagrams. *Operations Research*, 34:871–882.

[Walker, 2000] Walker, M. A. (2000). An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *Journal of Artificial Intelligence Research*, 12:387–416.