Modelling Variations in Goal-Directed Dialogue

Jean Carletta*
University of Edinburgh
Department of Artificial Intelligence
jcc@aipna.ed.ac.uk

1 Introduction

This research investigates human dialogue variations by having simulated agents converse about a simple map navigation task using a computational theory of human dialogue strategies. The theory to be tested is that agents have a variety of strategies which they use in goal-directed dialogue for deciding what information to include, how to present it, how to construct references, and how to handle interactions. Agents choose strategies for each of these aspects of the dialogue depending on the complexity of the current task, responses from the partner, and the task's importance; in general, agents try to minimize the collaborative effort spent on the task but still complete it satisfactorily. The software allows simulated agents to be constructed using any combination of strategies, showing how the strategies interact and allowing the decisions made in human dialogues about the same task to be modelled. Currently, the system works on a subset of the strategies found in Shadbolt [6], but a corpus of human dialogues is being studied to construct an improved theory of dialogue strategies, and these strategies will be incorporated into later versions of the system.

2 The Dialogue Domain

The task domain for the dialogues involves navigating around a simple map containing approximately fifteen landmarks. Two participants are given maps which differ slightly; each map may contain some features omitted on the other, and features in the same location on the two maps may have different labels. The first participant also has a route from the labelled start point to one map feature. The second participant must draw the

route on his or her map. In this task, the participants must cooperate because neither of them knows enough about the other's map to be able to construct accurate descriptions. At the same time, small changes in the map test how participants handle referential ambiguities, how information is carried from one part of the dialogue to the next, and when agents decide to replan rather than repair an existing plan. Despite the possibilities for referential difficulties, this task minimizes the use of real world knowledge as long as all participants understand how to navigate. The task is simple enough to be completed in computer-simulated dialogue, but admits the dialogue variations to be tested in the research.

3 The Central Idea

The central idea behind the research is that agents need multiple strategies for engaging in goal-directed dialogue because they do not necessarily know the best way to communicate with a given partner. Self [5] shows that dialogue is crucial where neither agent has all of the relevant domain knowledge. Dialogue is also necessary for any explanations where agents don't have accurate models of their partners [3]. Even if agents have all of the relevant domain knowledge, they may not know how best to present that knowledge, especially since explanations are about exactly that part of the task which is not mutually known to the dialogue partners [1]. Shadbolt [6] presents evidence that humans handle uncertainties about what information to give and how to present it by having a set of strategies for each aspect of the dialogue. Then the agent can tailor explanations to a particular partner by using the strategy that best fits the situation. For instance, human agents who believe that much domain information will have to be communicated structure their presentation carefully and often elicit feedback from the partner, like partici-

*Supported by a scholarship from the Marshall Aid Commemoration Commission. Thanks to Chris Mellish for supervising me and Alison Cawsey for helpful comments.

pant A of Shadbolt's [6] example 6.16:

A: have you got wee palm trees aye?
B: uhu
A: right go just + a wee bit along to them have
you got a swamp?
B: er
A: right well just go + have you got a water-
fall?

Agents who believe that most domain information is known to their partner are more likely to rely on interruptions from the partner and replanning, as in example 6.11:

A: and then + go up about and over the bridge
B: I've not got a bridge I've got a lion's den
and a wood
A: have you got a river?

One way to make computer generated explanations look more natural is to plan them using strategies modelled on the human ones. Although strategies like these could be built into the way a system plans an explanation, making strategy choices explicit allows the strategies themselves to be investigated, providing a way to test out how variations affect the ensuing dialogue. The goal of the present research is both to show how using dialogue strategies can improve the "naturalness" of computer-generated task explanations and to provide insight into the dialogue strategies which humans use and how they interact.

4 The Project

The project involves creating a theory of human dialogue strategies and modelling it using two computer processes that converse. Communication for the computer agents, based on the model in Power [4] and Houghton [2], is simplified in a number of ways. A convener wakes the agents in turn and interactions are made by placing messages in mailboxes, leaving out the complications of turn-taking and interruption. Rather than reason from "visual" images of the maps, agents begin with sets of beliefs about the positional relationships among objects and share knowledge about both dialogue conventions, expressed as *interaction frames* [2], and navigational concepts like "toward", "between",

and "around". The agents converse in an artificial language resembling their shared planning language, but substituting referring expressions for internal feature identifiers. Under these constraints, the agents use dialogue strategies to decide on the content and form of the dialogue. The existing system is a prototype designed to show that incorporating such strategies can explain some variations in human dialogue and make agents more flexible. An improved set of strategies is being extracted from the corpus of human dialogues. The end result of the project will be a theory of how communicative strategies control variations in dialogue, and software in which computer-simulated agents use these strategies to complete the navigation task.

5 The Program

The current version of the software uses dialogue strategies adapted from Shadbolt [6]. He lists seven different aspects of dialogue along which strategies may be developed. Agents may vary strategies for *feedback* (how they handle the partner's utterances), *specification* (how they construct and resolve referring expressions), *ontology* (how they decide from what features are available how to construct route descriptions), *focus* (the amount of explicit focus information given), *difference* (the effort spent determining what the partner's utterances mean), *decentering* (whether information is presented using the agent's or the partner's names for features), and *hypothesis formation* (the effort spent making hypotheses about the partner's knowledge). Agents choose strategies for each of these aspects depending on how explicit they want to be, which in turn depends on how likely the partner is to misunderstand each aspect of the dialogue. Some of Shadbolt's aspects are interrelated; for instance, agents that provide explicit information about the current focus do not need to construct referring expressions as carefully as agents who provide no focusing information at all. Our own work divides the strategies slightly differently so that they can be divided into sets depending on whether they affect planning the dialogue interaction, planning the content, planning the presentation, or realizing references; the goal is to make the strategies as modular as possible so that they can be modelled simply. Each simulated agent takes on a set of strategies for the duration of a dialogue. Currently, the prototype varies how much information about the structure of the dialogue is explicitly given, which features are included in a route description depending on a model of the partner's be-

liefs, how often an agent allows interactions from the partner, and how much repair an agent is willing to do rather than replan a description. The agents also use heuristics to prefer plans where the partner already understands the plan's prerequisites. The output of the program is a simulated dialogue where each agent keeps the same strategies for the course of the dialogue; an obvious future step is to allow agents to adapt to a particular partner or part of the task by varying their strategies within a dialogue.

6 Examples

The system currently has several strategies which affect how much structuring information is given in a dialogue and how often feedback is elicited from the partner. The following dialogue, an English gloss of two simulated agents conversing, shows how agent A might act if it believed that the maps had many differences:

- A: I'm going to tell you how to get to the buried treasure. I'm going to tell you how to navigate the first part of the route. Do you have a palm beach?
 B: Yes.
 A: Do you have a swamp?
 B: No.
 A: Do you have a waterfall?
 B: Yes.
 A: The swamp is between the palm beach and the waterfall. OK?
 B: Yes.
 A: The route goes to the left of the palm beach and around the swamp. OK?
 B: Yes.

If agent A believes that there will be few misunderstandings, or that B will understand enough to say what it misunderstood, it might choose to give information first and repair afterwards:

- A: The route goes to the left of the palm beach and around the swamp. OK?
 B: Where's the swamp?
 A: The swamp is between the palm beach and the waterfall. OK?
 B: Yes.

7 Conclusion

The theory and program are designed to show how variations that occur in human dialogue can be explained in terms of deciding among communicative strategies governing the form of interaction, the content and presentation of information, and the construction of referring expressions. The strategies found by examining a human corpus of goal-directed dialogues are implemented in a dialogue system where two computer processes using an artificial language and simplified turn-taking complete the task. This approach is useful in itself for determining what makes human dialogues seem natural, but it also has implications for human-computer interaction, since it is one step towards making computer dialogues with humans operate flexibly to fit in with human expectations.

References

- [1] S. Garrod. Explanations in dialogues as attempts to overcome problems in coordination. In *Proceedings of the Third Alvey Explanation Workshop*, September 1987.
- [2] G. Houghton. *The Production of Language in Dialogue: A Computational Model*. PhD thesis, University of Sussex, April 1986.
- [3] J. Moore and W. Swartout. A reactive approach to explanation. In *Proceedings of the Fourth International Workshop on Natural Language Generation*, 1988.
- [4] R. J. D. Power. *A Computer Model of Conversation*. PhD thesis, University of Edinburgh, 1974.
- [5] J. Self. Bypassing the intractable problem of student modelling. In *Intelligent Tutoring Systems*, 1988.
- [6] N. R. Shadbolt. *Constituting reference in natural language: the problem of referential opacity*. PhD thesis, Edinburgh University, 1984.