

Review Article

Does Conversation Analysis Have a Role in Computational Linguistics?

Computers and Conversation

Paul Luff, Nigel Gilbert, and David Frohlich (editors)
(University of Surrey)

London: Academic Press, 1990, 284 pp.
(Computers and People Series)
Hardbound, ISBN 0-12-459560-X,
\$45.00, £23.95

Reviewed by
Graeme Hirst
University of Toronto

1. Conversation Analysis

Computational linguists are a vicious warrior tribe. Unconstrained by traditional disciplinary boundaries, they invade and plunder neighboring disciplines for attractive theories to incorporate into their own. The latest victim of these attacks is sociology—in particular, a branch of ethnomethodological sociolinguistics called Conversation Analysis (hereafter, *CA*).

The book reviewed here, *Computers and Conversation*, is the first significant result of this. It is based on a symposium held at the University of Surrey in September 1989. Its purpose is to show what *CA* has to offer to the study of human-computer interaction (hereafter, *HCI*), especially interaction in natural language. In this review article, I want to use the book to explore at some length the question of whether *CA* has indeed anything to say to computational linguistics.

Up to now, almost all research in computational linguistics on dialogs between humans and machines has taken the approach of Discourse Analysis. Discourse Analysis concerns itself with talk at the level of linguistic constituents (Levinson 1983, pp. 286ff; Stubbs 1983, pp. 1, 9–10). The disfluencies of spoken language are ignored. Utterances are treated as meaningful units of language, and are analyzed with regard to their semantic content. So the human and the machine are both seen as rational agents whose utterances are always perfectly formed and fully in accordance with Gricean rules. These interlocutors like to spend their time seeking information from each other, testing each other's abilities in plan inference. Typical interchanges are the following:

1. **Human:** The 8:50 train to Montreal?

Computer: The railroad station is three blocks west on Main Street. But it's closed on weekends. In any case, the last train to Montreal ran in 1988. You should call Air Canada reservations at 1-800-AIR-LINE. Thank you for using Mr Askme. Have a nice day!

2. **Computer:** Attach the flageolet to the pump shaft by means of the compressor module.
Human: It doesn't fit inside it.
Computer: No, the blue one.

Conversation Analysis, however, is a contrasting tradition. From its name, Conversation Analysis sounds like it should be the same thing as Discourse Analysis. In practice, however, the two are distinct. CA tends to view talk as a human noise-making behavior. It concerns itself with talk at the sound-stream level, without constituent analysis, and it worries about such questions as how people take turns in conversation, where they look when they speak, and when and how they interrupt each other (Goodwin, 1981, pp. 6–9). Boden (1990b, p. 248) remarks that “a better name for the field would, in fact, probably be something like ‘interactional analysis,’ as everything in the interaction, from a quiet in-breath to the entire spatial and temporal organization of the scene, may be subject to analysis.” In Conversation Analysis, typical interchanges look more like these:

3. **A:** How much did you . in fact . see of the English department upstairs
C: Very little
A: Very little
C: Very little¹
4. **A:** Richard hallo . I've just *[s]* set out
B: *thank you*
A: ⟨⟨2 syllables⟩⟩ make some
B: thanks
A: I've just boiled some water for having coffee cos I haven't had time for tea — *would* you like some
B: *yes* yes²
5. **Patient:** An ah cun take thi:s hand an lif:t it and it duh'unt hurt
Doctor: Yuh mean when yuh *lif* the ar:m up using yer other
han:d=
Patient: =Uh-huh³
6. **A:** How many tubes would you like sir
B: U:hm . what's the price now eh with V.A.T. do you know eh
A: Er I'll just work that out for you=
B: =Thanks
(Pause of 10.0 seconds)
A: Three pounds nineteen a tube sir
B: Three nineteen is it=
A: =Yeah⁴

1 From Svartvik and Quirk (1980, S.1.5: 836–848). I have omitted their extensive prosodic markings.

2 From Svartvik and Quirk (1980, S.1.4: 1–11), simplified. The asterisks mark simultaneous speech; e.g., B's *thank you* is said at the same time as A's [s].

3 From West (1984, p. 133). The colon indicates a lengthening of a sound, the equal-sign that the two utterances are linked without a discernible pause.

4 Levinson 1983, p. 305.

7. (Silence)

A: Huh?

B: I didn't say anything.⁵

The difference between the two research traditions is perhaps best illustrated by the difference in the way they approach the matter of indirect speech acts. In Discourse Analysis and traditional computational linguistics, indirect speech acts are seen as an important phenomenon in discourse, and dealing with them requires considerable effort in plan recognition on the part of the comprehender (e.g., Perrault and Allen 1980). On the other hand, "on the CA view, the problem doesn't even exist" (Levinson 1983, p. 356); rather, the analysis is in terms of *pre-requests* and *dispreferred responses* (for details, see Levinson 1983, pp. 356ff).

Perhaps the most important points that CA makes are that there is much more to human conversation than is studied in Discourse Analysis or theoretical linguistics, and that empirical research methods are required. But it's not immediately obvious that this has any consequences for HCI. After all, human conversation is spoken, whereas HCI, to the extent that it is linguistic at all, is generally textual, research in speech recognition notwithstanding. But on the other hand, both speech and HCI are highly interactive forms of communication. Perhaps, as Frohlich and Luff argue in their chapter of the book, that's what really matters.

For example, one area of particular concern to CA is turn-taking and interruptions. CA studies how it is that speakers signal to others either that they have finished their turn or that they wish to 'retain the floor' for another *turn construction unit*; and it studies the circumstances under which speakers will interrupt one another. In present-day human-machine dialogues, such matters are simply not an issue. Users signal the end of their turn by means such as hitting the RETURN key; the machine signals by displaying some prompt to the user, usually a character such as '*' or '>'. Neither side can interrupt the other. The suggestion of this book is that this is not how things should be. Since elements like turn-taking cues and interruptions are so important in human conversation, if we want to build natural, flexible interfaces, these matters should also be an issue in HCI.

Similarly, Pirkko Raudaskoski, in her chapter, contrasts conventional research in HCI on recovery from misunderstandings with that in CA. The HCI research emphasizes the analysis and classification of the failure (cf. Ringle and Bruce 1982, Goodman 1986), while CA emphasizes the conversational patterns (e.g., Sacks et al. 1974) and the active role of both interlocutors in negotiating a meaning (cf. Clark and Wilkes-Gibbs 1986, Heeman 1991). Raudaskoski sees value in a synthesis of the two approaches. And Hugh Robinson points out in a short chapter that HCI would get more respect within the field of software engineering if it used CA to make itself a little less airy-fairy.

2. The Intellectual Background of Conversation Analysis

Conversation Analysis follows the tenets of *ethnomethodology*, a radical movement that developed in sociology in the early 1960s under the leadership of Harold Garfinkel of the University of California, Los Angeles. In reaction to conventional sociology, ethnomethodology stresses the questions of what the foundations of social interaction

⁵ Schegloff, quoted by West (1984).

are, and how people interpret their everyday social experience:

My purposes ... are to demonstrate the essential relevance, to sociological enquiries, of a concern for common sense activities, as a topic of enquiry in its own right ...: the socially standardized and standardizing, "seen but unnoticed," expected background features of everyday scenes. (Garfinkel 1967, p. 36)

Ethnomethodology thus places great emphasis on its being descriptive, empirical work, detailed and (purportedly) atheoretical: a researcher's task is to gather data and avoid theoretically loaded interpretations.

One "seen but unnoticed" object of study is the mechanism of ordinary conversation: not formal discourse or written text, but common phatic chit-chat and utilitarian exchanges. Conversation Analysis arose in the 1960s from the work of Harvey Sacks, and is associated with Emanuel Schegloff and Gail Jefferson in the U.S. and with John Heritage and J. Maxwell Atkinson in the U.K. For a balanced, easy-to-read introduction to CA that emphasizes the linguistic issues, see Levinson (1983, ch. 6); for an introduction that emphasizes the ethnomethodological roots, see Heritage (1984, ch. 8) or Boden (1990a).

Sacks believed deeply that the world is actually a very orderly place if one just looks closely enough. As Robin Wooffitt says in his introduction to CA in *Computers and Conversation*:

A fundamental assumption informing CA is that ordinary talk is a highly organized, orderly phenomenon. The goal of analysis is to reveal the basis of this phenomenon. ... [Sacks] did not treat an utterance as something that just happened to be said that way. Rather, he worked on the assumption that the things people said were methodically designed. (pp. 10–12)

[In CA,] words used in talk are not studied as semantic units, but as products or objects which are designed and used in respect of the interactions being negotiated through the talk: requests, proposals, accusations, complaints, and so on. ... The analytic objective is to explicate the procedures on which participants rely to produce utterances, and by which they make sense of other people's talk: to describe the 'technology of conversation' (Sacks 1984, p. 413). (p. 10)

Conversation Analysts thus spend their time tape-recording or videotaping conversations like (3)–(7) above, and seeing what sort of structures they contain. Many conversational exchanges are found to be various kinds of *adjacency pairs*: a greeting followed by a greeting in response, a question followed by an answer, and so on. But variations are possible. For example, the response to a question might be another question, thereby initiating a *side sequence* (Jefferson 1972), as seen in dialog (6) above. If there is some sort of trouble with the initial utterance, the response might be a *next turn repair initiator*, as in (5) above.

Throughout all, researchers working with everyday conversational materials have uncovered a veritable gold mine of ... precise and patterned procedures for producing talk that reveal, in their instantiation, the sort of fine-grained order in the social world that so amazed early naturalists in the nineteenth century as they began to systematically observe the natural environment. (Boden 1990b, pp. 252–253)

Wooffitt says (p. 27) that such conversational structures are not fixed and hard-wired cognitive phenomena, but rather are normative and socially organized. There is some dispute on this point. Boden (1990b, p. 250) claims, for example, that the turn-taking model of Sacks, Schegloff, and Jefferson (1974) "has held up across a range of languages

and cultures [and] it now seems quite reasonable to claim that this core machinery for talk transcends both language and culture." Conversely, Levinson (1983, p. 301) uses the claim that the model does *not* hold in Burundi to defend the model against the charge of being trivial or content-free! Regardless, conversational structures are so fundamental and become so deeply internalized that, as David Good points out in his chapter of the book, they are present even in the disordered conversation of schizophrenics. Good suggests that just as studying aphasics can help us learn about syntactic mechanisms, so studying schizophrenics can help us learn about conversation. He exhibits an extract of a conversation in which the patient maintains normal repair and functional cooperation even while the content is seriously awry.

There is a tension in CA between the desire to find neat, universal, conversational structures and the recognition that conversation, like all social interaction, is *situated* and *indexical*—that is, it depends upon, and, *a fortiori*, helps to *create*, the situation or context in which it occurs:

The organizational features of conversation are treated as structures in their own right and are taken to operate—like other social structural factors— independently of specific actors, psychological dispositions, or attributions of particular individuals. This is not to say, of course, that there isn't variation across individuals, but rather that these conversational structures are "context free." Secondly and simultaneously, the structures of talk are assumed to be "context sensitive" in the sense that their instantiation at particular moments and in particular contexts, as well as at specific points in interactional time, constitutes that moment and shapes that interaction. (Boden 1990b, pp. 250–251)

This research is not without its critics of course. Goldthorpe, in his review of early ethnomethodology (1973), finds little if anything to prefer over conventional sociology. He complains of ethnomethodology's impenetrable jargon, and this can certainly be seen in the CA literature—e.g., *transition relevance place* instead of *potential end of utterance*. Presumably, such locutions are intended to distance one from theoretically loaded interpretations of the data.

More recently, O'Connell et al. (1990) have presented a critique of CA research on turn-taking in conversations, a critique so vehement and unrelenting that one is left with the impression that the authors believe that no Conversation Analyst has ever uttered a true word on that topic or any other. The criticisms are largely aimed at the assumptions and methodology of the research. Perhaps the most important, however, is that CA completely ignores the goals and purposes of speakers in a conversation. Even if this is true, it's not clear to me that it matters, at least for CA. A speaker's intents are, after all, not open to direct observation. CA can only describe the external behavior, the structures that conversants have available to them. To expect otherwise is to misunderstand the goals and purposes of CA. For HCI and computational linguistics, however, the point is more serious. If we are in the business of building purposeful machines that respond to purposeful users, the notion of communicative intent seems to be central to the enterprise (cf. Cohen et al. 1990). This casts doubt on the idea that CA could play a useful role in HCI. I'll return to this point at greater length in Section 4 below.

3. Applications of Conversation Analysis in Human–Computer Interaction

Computers and Conversation is implicitly divided into two parts. Chapters in the first part debate the question of what exactly CA is and whether it has any application at all to the study of human–machine interaction. Those in the second part report

research that indeed purports to apply CA to HCI. I'll take the two parts in reverse order, as the applications described in the second are used as evidence in the debate in the first.

3.1 The Advice System

The chapter that is perhaps most convincingly pro-CA is the one by David Frohlich and Paul Luff that describes a CA-inspired interface to an advice system. The advice system, which is named Advice System, gives information and assistance on British government welfare benefits, and is intended for use directly by clients.

The interface is primarily mouse-driven, with the user clicking on 'buttons' on the screen, or picking up phrases from menus to assemble natural language sentences, rather like the Texas Instruments NLMenu system (Tennant et al. 1983a, 1983b). Unfortunately, the authors are fuzzy about many of the details, and it's not clear whether there is also direct keyboard input. (Page 190 and Figure 9.1 suggest that there isn't, yet page 207 and Tables 9.5 and 9.7 suggest that the user might make spelling errors.) It's also unclear whether the input menus are dynamically generated from a true grammar and lexicon, as in NLMenu, or whether they are just a static, pre-determined selection of useful fragments. Likewise, the authors don't say whether the Advice System's responses are just pre-stored text or dynamically generated. (It's characterized as an expert system, so one expects something closer to the latter, but Frohlich and Luff provide no details.)

Frohlich and Luff's explicit claim (p. 188) is that they have tried to build the principles of CA into the system "in such a way as to generate and support orderly sequences of talk." In this design, advice or explanation is not something the system dumps out in a single turn, but is an "emergent property" of a number of turns (p. 219).

For example, CA is applied in formulating *dialogue control policies*, in which the authors try to "recreate some of the dynamics of ordinary conversation" (p. 201) by using CA rules for *turn-taking* in the conversation and for allowable *adjacency pairs* (Sacks et al. 1974). Specifically, the types of utterances possible are questions, statements, and answers, both by the user and by the system. That means that there are $(3 \times 2)^2 = 36$ possible kinds of adjacency pairs (p. 193). For example, a user's question can be followed by an answer by the system (such a pair is denoted UQ-SA) or by a question from the system (UQ-SQ). But each 'speaker' can also 'retain the floor' after an utterance, perhaps adding a statement after an answer, or asking a question (e.g., UA-US, UA-UQ, SA-SS, SA-SQ).

But not all 36 kinds of pairs are allowed in the system (despite the assertion to the contrary in the caption of Table 9.1); 10 or 11 kinds are ruled out by CA's constraints on turn-taking and turn construction (the exact count depends on whether you prefer to believe Table 9.4 or the text on page 205). For example, UQ-US and SQ-SS are eliminated by the rule that a question may not be followed by a statement. In fact, a question always causes a turn transition. This strikes me as too restrictive. A UQ-US sequence seems quite natural, e.g., "Can I get disability benefits? I broke my arm at work last week." But an SQ-UQ sequence is among those allowed; that is, the user can respond to answer a question with another question—although the menu structure is said to make this somewhat difficult to do (p. 204).

The key point that arises is that, unlike ordinary human-computer dialogs, the user and system can interrupt each other at the end of any *turn construction unit*, even if the turn is incomplete. For the user, selecting the key marked "." (period) after constructing an utterance indicates that they wish to end the turn, while selecting "..." (ellipsis) indicates that they wish to say more, that is, to retain the floor. Nevertheless, the system is allowed to interrupt at such points—but only if it requires clarification

of what's been said so far in the turn, that is, if it wants to elicit a *self-repair* by the user (see below). Similarly, the user can interrupt the system. After each turn construction unit uttered by the system, there is a pause during which a button labeled "LET ME SPEAK" becomes available to the user for a short time (p. 204).

Another aspect of conversation studied by CA and incorporated in the Advice System is *repair*, that is, clarifications, corrections of errors, and the like. Schegloff et al. (1977) have shown that there is a strong preference in conversation for *self-repair*—for people to correct their own errors. As already mentioned, the system can use next turn repair initiators; that is, it can interrupt the user and ask "What?". (It's not smart enough to attempt the repair itself: "You mean . . .?") In addition, it enables the user to do the same: it makes *checking moves*, asking "OK?" every so often. If the user answers "No", the system will try to repair the problem. (Similarly, users can ask the system "OK?" if they are not sure that they are being understood.)

But it is unclear to me why the Advice System should ever need to initiate repairs by the user at all. After all, menu-based interfaces are supposed to prevent the user from uttering anything that is beyond the system's ability—a well-structured menu system *precludes* problematic choices. Extending that principle to natural language input was one motivation of the original NLM menu system. But then repair would never be necessary, and Frohlich and Luff's goal of having a system that performs humanlike repair initiations would be undermined! (Frohlich and Luff merely say that the use of menus ensures "high levels of input recognition" (p. 191).)

The closing of the dialogue also follows principles derived by CA. The user can't just walk away from the machine, but must go through an official *pre-closing* (Schegloff and Sacks 1973) in which "each party [must] decline at least one opportunity to continue talking" (p. 214) before the closing and terminal exchange (a "Bye"—"Bye" adjacency pair) is reached. The user can shorten this sequence by clicking on a button labeled "I HAVE TO GO". The pre-closing sequence starts after the system has fully dealt with all questions that the user has asked. In this sequence, the system can volunteer further information and encourage the user to continue to ask questions.

The system also has a deliberate opening sequence intended to help familiarize the user with the system (p. 213). This seems to be a pretty rigid sequence, and it's not clear that the invocation of CA results (such as Schegloff's (1968) summons-answer sequences) is useful or an improvement upon conventional systems.

The rules that control all this are expressed in a Prolog-like notation. Some examples (from the authors' Table 9.3) are shown in Figure 1. It is tempting to view these rules as the rules of a grammar for conversation, a point that I will return to below.

There are some very interesting ideas in this work, but I am not sure what to make of the result. First, the authors are so vague at times, even self-contradictory, that one cannot get an adequate impression of the system as a whole. Are the interface and the underlying advice system genuinely intelligent, or is it nothing more than 'Eliza with a mouse'?⁶

Second, whether the interesting ideas work in practice is an empirical question. The system had not been tried out on real users at the time the chapter was written, though field tests were planned. This is an important point. Many systems that have been designed to be "helpful" have been found to cause more problems and misconceptions than simple, straightforward systems. Examples include the photocopier studied by Suchman (1987) and Raudaskoski's message system, to be described below. Moreover, testing the system is likely to be difficult, because there is so much in it that is novel

⁶ Thanks to George Ferguson for this apt phrase. The allusion is to Weizenbaum (1966).

```

conversation <- opening
                body

body <- if (OR (user-questions) (system-questions))
          then adjacency-pair
          else preclosing1

adjacency-pair <- user-open-floor
                system-floor
                body

user-open-floor <- if (system-questions)
                    then user-answer-turn
                    else user-statement-turn

```

Figure 1

Examples of rules in the Advice System.

that unless it is a runaway success, it might be difficult to determine just which ideas worked and which didn't.

3.2 Repair in Human-Computer Dialog

The second of the two chapters that best show the potential applicability of CA to HCI is that by Pirkko Raudaskoski on repair sequences. Humans not infrequently misunderstand one another and have to initiate repairs or clarifications. Given the imperfect nature of present NLU and speech recognition techniques, the need for a computer to be able to initiate repairs is even greater than it is for people. And similarly, the user should be able to initiate a repair by the system if necessary.

Raudaskoski carried out an interesting experiment with a system that could use a number of different repair-initiation strategies whenever it failed to comprehend an input. The domain was a speech interface to a telephone message store-and-forward system in which users could ask to leave messages for others, or ask to hear their own messages. To finesse the problem of speech recognition, the interface was simulated. Users spoke over the phone to the experimenter, who typed in what they said and read back the system's responses. It's not clear how this might have affected the results of the experiment.

The system had five kinds of repair initiators:

- Parroting back, indicating the source of trouble: e.g., "Leave a message for who?"
- Implying not hearing or understanding: e.g., "Sorry?"
- Explicitly guessing: e.g., "Do you mean ...?"
- Narrowing the user's choices: e.g., "You can either ... or ..."
- Checking: e.g., "You want to ..., right?"

The first two of these were also available to users to initiate a repair by the system.

The experiment was pretty much a failure—but it failed in an interesting way. Raudaskoski found that subjects never availed themselves of repair initiators; and the system didn't go through its whole repertoire either. In one interesting segment

(p. 168), reprinted as (8) below, the system initiates a repair (line 8.1), but the user's response (line 8.2) is not understood either and it tries again (line 8.3). But the user misinterprets this as a prompt for the message to be left (line 8.4). The system keeps trying (line 8.5), but is again misunderstood: what the user repairs is her previous utterance (line 8.6). More patient than any human would be, the system tries yet again (line 8.7). The user takes the fact that the system is ready to start a task as confirmation that the previous task is complete, and starts something new (line 8.8). So the user ends up thinking that she has achieved her goal of leaving a message for Eve, whereas the system has understood nothing at all!

8. (8.1) **System:** So you want to leave a message for Eve—is that right?
 (8.2) **User:** A message for Eve, that's correct.
 (8.3) **System:** What is it that you want to do then?
 (8.4) **User:** Can you say that Judy is at home?
 (8.5) **System:** Do what?
 (8.6) **User:** Judy is at home.
 (8.7) **System:** You can leave messages or listen to messages—select either.
 (8.8) **User:** I want to leave a message for Fred.

What Raudaskoski found was that in speech mode, humans tend to be verbose; they often don't just say "yes" or "no." The system was unable to channel their responses narrowly enough that it could be sure of understanding them. It needed to initiate repairs to repairs, but it couldn't.

I suspect that a similar problem might be found with Frohlich and Luff's Advice System. As we saw earlier, its only way of initiating repairs is to ask "What?" But this is probably a very bad initiator for a computer system that aspires to natural conversation. Valian and Wales (1976) have shown that people's responses to "What?" are generally to offer a syntactically simpler form of the sentence uttered, if it is structurally complex, or to repeat essentially the same utterance, if it is already simple. So when the Advice System says "What?", it is likely to be seen as initiating a syntactic repair, whereas the problem is presumably more likely to be semantic.

Raudaskoski concludes her chapter by suggesting that a computer's next turn repair initiators should *not* be like those used by people, but rather should emphasize exactly where the misunderstanding is. That would certainly be good advice for the Advice System.

3.3 Other Applications

Two other papers in the book also try to apply CA to human-computer interaction. Alison Cawsey is interested in tutorial systems in which an expert (the system) teaches a novice about a domain (here, electronics) and asks test questions. She focuses on explanations that the system gives to the user. While acknowledging that such explanations need to be planned (p. 223), she claims that it is also necessary to allow for interaction within the explanation. She does this with CA-based rules not unlike those of Frohlich and Luff, above. The planning is incremental, and based on each exchange. It occurs at both the high level (discourse structure; e.g., do a pre-closing) and the lower level (e.g., adding a pre-closing to another utterance). In extensions to the work, Cawsey hopes to allow the user to initiate repairs, in effect saying "Huh?" to the system (pp. 232–3). Moore (1989) has described a system with a similar ability, but Cawsey claims that Moore's system, because it's not based on CA's observations

of natural conversations, is limited in two ways: the user can't interrupt in the middle of a turn, and the system doesn't ask the user to confirm its guess as to what is wrong. Whether these limitations make the system less effective than Cawsey's is, of course, an empirical question that is yet to be tested.

Nigel Gilbert, Robin Wooffitt, and Norman Fraser describe a project in progress: Sundial, a speech recognition-based system for giving, for example, flight information over the phone. Utterances are not to be processed singly and in isolation, as in present-day speech recognition systems, but rather as part of a CA-informed discourse. The authors argue that speech is the best HCI medium (p. 236). Of course, the system is still a couple of years off.

A rather different use of CA, in the interpretation of software specifications, is suggested by Anthony Finkelstein and Hugo Fuks. One task of the systems analyst is to take vague, waffly, and often contradictory statements from people as to what is required in a system, and turn them into clear, unambiguous, formal specifications (cf. Regoczei and Hirst 1989). Finkelstein and Fuks suggest that explicit use of CA might ease the analyst's task, allowing one to sort it all out, including the retractions, the iterative clarifications, and so on, to resolve different perspectives on the problem.

They use rules based in part on CA; for example, all challenged statements must be affirmed, denied, or withdrawn. This is as much Discourse Analysis as Conversation Analysis, and the authors claim that (contra some other suggestions, e.g., Levinson 1983, p. 294) the two are quite compatible. The examples given by the authors are made up, not actual dialogue, and it's unclear as to whether they are intended to represent real speech or merely formalizations thereof. (If it's the former, then the authors are advised not to seek careers as playwrights.)

This looks promising at first glance, and the authors hope that it could lead to automated support. But it is unclear that the rules could really do that. They describe a rigid discourse, saying, for example, that all questions must be answered. Real conversations needn't be like that. Nor is it immediately clear how these rules would be used by a human analyst in practice.

4. Is This All Really Possible?

I now turn to the chapters that debate whether CA in fact has useful things to say at all concerning human-computer conversations, and whether the research described above should count as evidence that it does.

Paul McIlvenny invokes Suchman's (1987) notion of the computer as an *interactive artifact*, to emphasize the user's view of it as a reactive, language-using entity that produces seemingly purposeful, social actions (p. 96). CA is thus relevant to HCI "wherever the second person metaphor is present" (p. 99).

The case for the practical usefulness of informing HCI design with CA is put, somewhat unconvincingly, by Michael Norman and Peter Thomas. Norman and Thomas say that CA is important both for its methodology and for its findings. It tells us about normative formats and expectations in conversations, and we can use that, for example, in analyzing situations where users have problems with an interface. CA can "sensitize the designer to matters of the interactional process" (p. 61); it can tell us, for example, not just that feedback is needed, "an often heard and particularly vague piece of design guidance," but that "the next action ... may in fact be ... a *repair* of the user's previous action," and thus "provide us with a view of the normative ways in which repairs are constructed" (ibid.). The chapter closes with some vague pieces of design guidance for HCI from CA, the authors admitting that determining how or whether they can be applied will require the development of a much better

understanding of human–computer interaction.

The opposing case is put by Graham Button, in a rather long and repetitious paper. He says that one might at best be able to produce a “simulacrum” (p. 67) of a naturally occurring conversation. For Button (in the ethnomethodological tradition), talk is an “achievement” (p. 79) or an accomplishment (cf. Garfinkel 1967, p. 1) that *uses* the rules of CA, but is not *caused* by them.

His argument is that the rules of CA are not like the rules of a computer model: CA rules for sequencing, turn-taking, and the like might look like causal rules at first glance, but actually they are “resources” (p. 84) that people “display an orientation to” (p. 78) in situated instances, namely social interactions. Button’s argument can be understood in two slightly different ways, both of which are, I think, valid in CA. The first way is this: A conversant, desiring a particular effect, can choose to invoke one of the rules that will have that effect in the current situation. In this way, conversants can make what they want of the contributions of others. For example, a speaker who says “Can you tell me how to get to your place by bus?” could be heard to be asking for bus information or requesting to be picked up by car; the listener *makes* it one or the other by the way she or he replies.⁷ Button emphasizes the idea, due originally to Sacks, that it is the response that provides for recognition of the first utterance, and sees this as inherently incompatible with rule-based behavior:

If we formulated a rule that upon the production of a question, the next utterance should be an answer, we could not have the case where a second *transforms* a prior utterance into a first and thereby achieves coherence with the first. (p. 85, emphasis added)

Indeed, any aspect of an interaction can be “retrospectively reread to find out in the light of present practical circumstances what [it] ‘really’ consisted of ‘in the first place’ and ‘all along’” (Garfinkel 1967, p. 74).⁸

The second interpretation is this: A conversant, desiring a particular situation, can quite literally *create* the desired situation by invoking a rule that would be appropriate only therein. For example,

Suppose there is a rule for greetings which runs to the effect: do not initiate greetings except with persons who are acquaintances. And suppose we subsequently see a man greeting another whom we know is not an acquaintance. We can either conclude that he broke the rule or we can infer that, *via the use of the rule*, he was seeking to treat the other *as* an acquaintance. The second interpretation is more likely when, for example, our man is greeting a new colleague at the office. (Heritage 1984, p. 126, emphasis supplied)

(Obviously, there are limits to this; you can’t make a 17-year-old eligible to vote simply by singing *Happy Birthday* to them, nor can you kill someone just by holding a memorial service for them.)

Now, from this, Button tries to argue that CA’s rules of conversation are not

⁷ The example is from Blum-Kulka and Weizman (1988).

⁸ For example, in one of Garfinkel’s experiments (1967, pp. 79ff), students who sought advice from a counselor were given completely random answers to their yes/no questions. “Over the course of the exchange, subjects sometimes started with the [random] reply as an answer and altered the previous sense of their question to accommodate . . . the reply as the answer to the retrospectively revised question” (p. 90). (This kind of revision is covert, and is not to be confused with the overt interactional reconstruction of an utterance that can follow a misunderstanding in conversation (Fox 1987, McRoy 1989, Cawsey forthcoming)).

amenable to direct use by computer, that they are not “codifiable” or “reduc[ible] to an algorithm” (p. 82). His first argument is that rules in a computer are *causal*, meaning that they are designed into the hardware itself (p. 77), or into the program it executes. This is, I gather, intended to contrast with the use of rules as resources, as described above. However, I am at a loss to see the contradiction. Any backward-chaining Horn clause theorem prover or planner with appropriate meta-rules can be viewed as using its rules as resources in exactly the ways I described. Indeed, the ability of a second utterance to transform a prior utterance surely *requires* having rules such as “that upon the production of a question, the next utterance should be an answer”! Button seems to be saying nothing more than that CA rules must be represented declaratively, not procedurally, so that they can be deliberately and explicitly invoked with a view to their effect. But far from precluding their use by a computer, this argument suggests that they fit very nicely into present-day AI reasoning systems!

Button’s second argument is that CA rules are not complete; for example, the rules for turn-taking don’t say how to identify a possible transition relevance place:

The fact that these ... places are *possible* transition relevance places means that it is a contextually decidable matter whether or not actual transition takes place. Thus, because the rules do not, so to speak, determine their own application (using the rules is not determined by the rules themselves), there have to be other features involved in the application of the rules. This is the orientation to possible transition places, and in [Sacks et al. (1974)] there is no description of these places being formulated in terms of the rules for their production. Possible transition places are the contextual achievement of the unfolding structuring of the turn in progress. Possible completion places are not, then, provided for in advance by codifiable rules, they are situatedly achieved in and for a particular turn. (pp. 80–81, emphasis supplied)

Button’s intent here is not, as it might seem, to claim that CA chooses to remain fully ignorant about certain elements of its domain; rather, it is, presumably, to tacitly invoke what’s called in ethnomethodology *the et cetera principle* (Garfinkel 1967, pp. 73–75): that the situations in which a rule is and isn’t applicable can never be finitely enumerated. In effect, every rule has an implicit *et cetera* at the end of its list of triggering conditions, and the agent must judge each time whether the current situation is one in which the rule may be applied, or, possibly, broken.⁹ Heritage (1984, pp. 121–122) makes an analogy with the legal system. No matter how carefully a law is written, a complex judicial decision will sometimes be required to determine whether or how the law applies in some particular new situation; and then another judicial decision might be required to determine whether a second situation is covered by the precedent established by the first.

But again, this is hardly an in-principle argument against computer use of CA rules. Case-based reasoning, while far from being a completed science, is nevertheless a well-established subfield of artificial intelligence, and seems to hold considerable promise for being able to do what Button says can’t ever be done.

Underneath, this discussion seems to be little more than a bad version of the old, familiar anti-AI arguments about Turing tests and Chinese rooms, about the limitations of physical symbol systems, about whether one needs to distinguish between a system that is explicitly following a certain rule from one that merely acts as if it is. It seems

⁹ Button seems to be making the stronger point that CA rules contain no triggering conditions at all. But I don’t think that this can be his intent, for as we saw above, it is only through the presence of such conditions that the use of rules as resources becomes possible; a second can transform a first only to the extent that the transformed first is a usual condition for the second.

to me that for HCI the question is whether we can build computer systems whose conversations are usefully similar to those of humans, and it doesn't matter whether or not Button wants to call them merely simulacra.

For Button, the bad guys are Frohlich and Luff, and Gilbert, Wooffitt, and Fraser, whose work I described above. (Cawsey and Raudaskoski are apparently exempted.) Button says that they are wrong in their claims to be "using" CA in their systems; the systems might work, but only because they are *not* truly based on CA or directly using CA rules. Norman and Thomas, incidentally, say that they agree with Button on this point, but think the criticism to be irrelevant—that one can still gain inspiration and help from CA without claiming to be using CA rules directly in one's system (p. 58). Frohlich and Luff, defending their work, seem to make the stronger claim, however: that aspects of CA are indeed directly built into their system (p. 188). The applicability of CA to HCI is, at least, a hypothesis to be refuted, and is the best bet presently available to us (*ibid.*).

Gilbert, Wooffitt, and Fraser also seek to counter Button's arguments, but are not convincing. They give a lengthy argument showing the obvious: that the constraints on adjacency pairs, with nesting and interruption, can be written as a phrase structure grammar (like the rules of Frohlich and Luff that we saw earlier), and such a grammar can be parsed by a chart-parser—provided, of course, that the individual constituents can be recognized! And because the parser is nondeterministic, it can recognize an ambiguity, and either maintain it (p. 251) or 'choose' from among the different possible interpretations (p. 254). But it's not at all clear to me what this proves. It's one thing to use a grammar of CA rules to perform post hoc recognition of the high-level structure of a conversation, or to generate an example of (both sides of) a conversation. It's quite another to use such rules to actually *participate* in a conversation without making a mess of things—especially if one takes the *et cetera* principle seriously. It's the other papers in the book that show that this might be possible.

Gilbert, Wooffitt, and Fraser conclude that "the idea that the findings of Conversation Analysis are not susceptible to computer implementation depends on a rather simple-minded notion of what a computer model can do" (p. 256). And indeed, this is the feeling one gets from Button's paper. Button retorts that "rather, the idea that Conversation Analysis lends itself to computer modelling depends upon a simple-minded notion of what conversation is, and of what order of description of human action Conversation Analysis provides" (p. 70). And he's right, too. Conversation is a complex, highly context-dependent, behavioral phenomenon. The systems described in the book don't come close to it. But that's *not* to say that they are trivial. Button accuses the systems of simply responding with "set pieces" (p. 82), by which he apparently means canned responses. As I mentioned before, it's unclear whether the Advice System's responses are just pre-stored text or dynamically generated; but even if the responses are canned, if it turns out that informants judge the system to be easy or 'natural' to use, if the response selected is always judged to be suitable, even in a variety of complex conversational situations, then something interesting is going on.

While Button's criticism is, in effect, that CA is too rich and fine for HCI, the criticisms of CA by O'Connell et al. (1990) imply the converse: that by ignoring communicative intent, CA is too *impoverished* for use in HCI. On this view, CA is nothing more than a kind of linguistic behaviorism. Unless its descriptions of external behavior can be related to the intents and purposes of speakers, CA can be used neither for understanding intelligent agents nor for constructing them. I see several replies to this. The first is to grant all this, but to point out that these enterprises do, nonetheless, require these descriptions of external behavior, and CA has them available. The challenge is then to use them.

The second reply is to deny the premise that intelligence and understanding require explicit analysis and representation of the intent of other agents. Rather, meaning can be seen as no more than the reader's response to a text, or, more generally, the perceiver's response to perception. This sits well with the notion mentioned above of a response being able to transform a first utterance without regard to the intent of its speaker, and is supported by deconstructionist and reader-response theories in literary studies (e.g., Iser 1974; see also Holub 1984). These theories are highly controversial, and are a current Hot Topic in literary studies (Matthews 1991). But, as Corriveau's reader-response model of text understanding shows (Corriveau 1991), they are not necessarily incompatible with artificial intelligence.

The third reply is to deny the other premise, and claim that intent is indeed accounted for in CA—in fact, is pervasive in CA. When CA uses terms such as *repair* or *closing* or *summons*, or speaks of conversants using rules as resources, it is, in fact, explicitly ascribing intent! When analysts say that a particular utterance is, say, a next turn repair initiator, they are using their own knowledge of the language of the conversation to interpret that utterance in context, and hence they are indeed determining the speaker's intent.

Perhaps the strongest use of intent in CA occurs in the notion of *accountability*: in interaction, agents hold each other accountable, or responsible, for their choice of actions and inactions. For example (Heritage 1984, pp. 106ff), consider again the rules about greetings. If two acquaintances approach each other in the hallway and the first greets the second, the second can choose either to return the greeting, and thereby be "friendly" or "polite," or not to, and thereby be "impolite" or "snub" the first. But he or she cannot just opt out and avoid making the choice, or "*whatever* they do will be intelligible and accountable as a sustaining of, or a development or violation, etc., of, some order of activity. . . . [Their actions] are *condemned to be meaningful*" (Heritage 1984, p. 110, second emphasis added).

One must be careful not to overstate this position, however. CA should not be viewed as just some kind of teenage mutant theory of semantics. Rather, it should be seen as describing the framework that agents use to communicate their meanings and intents. So to study such a framework is necessarily to study aspects of meaning and intent.

I believe that this third reply is correct; but it is problematic for CA for two reasons: it weakens CA's claim to being purely empirical and atheoretical, and it undermines the listener-response orientation of CA. On the first point, CA's defense is to distinguish interpretation from theory, and say that analysts can use their knowledge of the language of the conversation without requiring any theoretical framework. The second point is trickier. Personally, I don't see any incompatibility in the notions of intent and conversational flexibility; we just need to note that the exercise of flexibility by the second speaker can subvert the intent of the first speaker. But I doubt that Button would accept this.

5. Conclusion

Before I finish, a few remarks should be made about the organization and production of the book under review. The chapters that debate whether CA can be applied in HCI seem to have been written iteratively, with each author seeing drafts of the others' chapters in order that their arguments and rebuttals mesh well. This is good, but it

has the awkward consequence that each side's chapters require the other side's as pre-reading. And the chapters are widely separated in the book.

Many of the chapters have cutesy titles that lack the keywords necessary to describe the content adequately (e.g., Frohlich and Luff, "Applying the technology of conversation to the technology for conversation"; Button, "Going up a blind alley"). Some of the chapters seem to be not so much position papers or project reports as memory dumps. David Good, for example, goes out of his way to tell us that his wife isn't always logically omniscient (pp. 142-3),¹⁰ and that before ennui set in, he was pretty good at cryptic crosswords (p. 144).

The typography and layout are mostly fine, except in Chapter 9, where tables are placed badly with respect to where they are cited in the text; parts of pages are left blank for no apparent reason; and one table is split, quite unnecessarily, so that the column headings are at the bottom of a recto page, and most of the body of the table is on the following verso! Conversations are presented in different formats in different chapters. In Chapter 9, they are presented in many different formats within the same chapter; and one fragment, illustrating pre-closing sequences, is given a very large heading, "Our conversation." A number of references cited in the text are missing from the bibliography (e.g., "Schegloff 1981," p. 154; "Hamblin 1971," p. 177; "Heritage 1987," p. 183).

While I have been very critical of much of *Computers and Conversation*, it is clear, I hope, from the length of my comments that I take the work very seriously. I believe that the role of CA in computational linguistics and HCI is definitely worth exploring, and the core of my criticisms is that the faltering first steps in this direction often falter more than they reasonably ought to. There is a sense in which it is clear that CA *must* have a role in NLU, because there is a sense in which ethnomethodology is just a small subfield of artificial intelligence (although that might come as a surprise to the ethnomethodologists), as its goal is simply to explicate and represent some areas of commonsense knowledge, rather like naïve physics and similar projects (cf. Hayes 1985; Hobbs and Moore 1985; Hobbs et al. 1985).

Obviously, it is too early to say that the results of Conversation Analysis research do have a place in the design of human-computer interaction. Button may yet be proved right. But the work described in this book is reason to be optimistic that a valuable synthesis will emerge. And if it doesn't, determining why not would be equally valuable.

Acknowledgments

I am grateful to Deirdre Boden, Richard Frankel, and Tom Smith for discussions that shaped my understanding of ethnomethodology and Conversation Analysis, and for pointers to relevant literature. I am especially grateful to Deirdre for her comments on an earlier draft of this review. My thoughts on the nature of rules were developed in discussions with David Traum. I have also

benefitted from advice from Jeff Pelletier, Susan McRoy, Chrysanne DiMarco, Diane Horton, James Allen, and Stephen Regoczei. The opinions that I express herein, however, are not necessarily anyone's but my own. This review was written while I was a visitor at the Department of Computer Science, University of Rochester, using facilities that the department kindly made available to me.

¹⁰ One of the premises that she was supposed to reason from is the self-contradictory proposition that she knows that he knows that she doesn't know that he knows that one of her acquaintances has a twin, where *know* is to be construed factively. Since everything follows from a contradiction, perhaps her inference engine just became overloaded.

References

- Blum-Kulka, Shoshana and Weizman, Elda (1988). "The inevitability of misunderstanding: Discourse ambiguities." *Text*, 8(3), 219–241.
- Boden, Deirdre (1990a). "The world as it happens: Ethnomethodology and conversation analysis." In: Ritzer, George (ed.) *Frontiers of social theory*. New York: Columbia University Press, 185–213.
- Boden, Deirdre (1990b). "People are talking: Conversation analysis and symbolic interaction." In: Becker, Howard S. and McCall, Michal (eds.) *Symbolic interaction and cultural studies*. Chicago: The University of Chicago Press, 244–273.
- Cawsey, Alison J. (in press). "A belief revision model of repair sequences in dialogue." In: Costa, Ernesto (ed.) *New directions in intelligent tutoring systems*. Springer-Verlag.
- Clark, Herbert H. and Wilkes-Gibbs, Deanna (1986). "Referring as a collaborative process." *Cognition*, 22(1), 1–39. Reprinted in: Cohen, Morgan, and Pollack (1990), 463–493.
- Cohen, Philip R.; Morgan, Jerry; and Pollack, Martha E. (eds.) (1990). *Intentions in communication*. Cambridge, MA: The MIT Press (System Development Foundation benchmark series).
- Corriveau, Jean-Pierre (1991). *Time-constrained memory for reader-based text comprehension*. Doctoral dissertation, Department of Computer Science, University of Toronto (published as technical report CSRI-246).
- Fox, Barbara (1987). "Interactional reconstruction in real-time language processing." *Cognitive Science*, 11(3), 365–387.
- Garfinkel, Harold (1967). *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice Hall. (Reprinted: Cambridge, England: Polity Press, in association with Basil Blackwell, 1984.)
- Goldthorpe, John H. (1973). "A revolution in sociology?" *Sociology*, 7(3), 449–462.
- Goodman, Bradley A. (1986). "Reference identification and reference identification failures." *Computational Linguistics*, 12(4), 273–305.
- Goodwin, Charles (1981). *Conversational organization: Interaction between speakers and hearers*. New York: Academic Press (Language, thought, and culture: Advances in the study of cognition).
- Hayes, Patrick J. (1985). "The second naive physics manifesto." In Hobbs and Moore (1985), 1–36. Reprinted in Brachman, Ronald J. and Levesque, Hector J. (eds.) *Readings in knowledge representation*. Los Altos, CA: Morgan Kaufmann, 468–485.
- Heeman, Peter (1991). *A system that collaborates on referring expressions*. Thesis, Department of Computer Science, University of Toronto (published as technical report CSRI-251).
- Heritage, John (1984). *Garfinkel and ethnomethodology*. Cambridge, England: Polity Press.
- Hobbs, Jerry R. and Moore, Robert C. (eds.) (1985). *Formal theories of the commonsense world*. Norwood, NJ: Ablex.
- Hobbs, Jerry R.; Blenko, Tom; Croft, Bill; Hager, Greg; Kautz, Henry A.; Kube, Paul; and Shoham, Yoav (1985). "Commonsense summer: Final report." Report 85–35, Center for the Study of Language and Information, Stanford University.
- Holub, Robert (1984). *Reception theory: A critical introduction*. New York: Methuen.
- Iser, Wolfgang (1974). *The implied reader*. Baltimore: The Johns Hopkins University Press.
- Jefferson, Gail (1972). "Side sequences." In: Sudnow, David (ed.) *Studies in social interaction*. New York: Free Press, 294–338.
- Levinson, Stephen C. (1983). *Pragmatics*. Cambridge, England: Cambridge University Press (Cambridge textbooks in linguistics).
- Matthews, Anne (1991). "Deciphering Victorian underwear and other seminars." *New York Times Magazine*, 10 February 1991, p. 42 et seq.
- McRoy, Susan Weber (1989). "Nonmonotonic reasoning in natural language." MS, Department of Computer Science, University of Toronto.
- McTear, Michael (1985). "Breakdown and repair in naturally occurring conversation and human-computer dialogue." In: Gilbert, G. Nigel and Heath, Christian (eds.) *Social action and artificial intelligence*. Aldershot: Gower (Surrey conferences on sociological theory and method 3), 104–123.
- Moore, Johanna D. (1989). "Responding to 'huh?': Answering vaguely articulated follow-up questions." *Proceedings of the Conference on Human Factors in Computing Systems*, Austin.
- O'Connell, Daniel C.; Kowal, Sabine; and Kaltenbacher, Erika (1990). "Turn-taking: A critical analysis of the research tradition." *Journal of Psycholinguistic Research*, 19(6), 345–373.
- Perrault, C. Raymond and Allen, James F. (1980). "A plan-based analysis of indirect speech acts." *American Journal of*

- Computational Linguistics*, 6(3–4), 167–182.
- Regoczei, Stephen and Hirst, Graeme (1989). "On 'Extracting knowledge from text': Modelling the architecture of language users." *Proceedings, Third European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Paris, 196–211. Also published as technical report CSRI-225, Computer Systems Research Institute, University of Toronto.
- Ringle, Martin H. and Bruce, Bertram C. (1982). "Conversation failure." In: Lehnert, Wendy G. and Ringle, Martin H. (eds.) *Strategies for natural language processing*. Hillsdale, NJ: Lawrence Erlbaum Associates, 203–221.
- Sacks, Harvey (1984). "On doing 'being ordinary'." In: Atkinson, J. Maxwell and Heritage, John C. (eds.) *Structures of social action: Studies in conversation analysis*. Cambridge, England: Cambridge University Press (Studies in emotion and social interaction).
- Sacks, Harvey; Schegloff, Emanuel A.; and Jefferson, Gail (1974). "A simplest systematics for the organization of turn-taking for conversation." *Language*, 50, 1974, 696–735.
- Schegloff, Emanuel A. (1968). "Sequencing in conversational openings." *American Anthropologist*, 70, 1075–1095.
- Schegloff, Emanuel A. and Sacks, Harvey (1973). "Opening up closings." *Semiotica*, 7, 289–327.
- Schegloff, Emanuel A.; Jefferson, Gail; and Sacks, Harvey (1977). "The preference for self-correction in the organization of repair in conversation." *Language*, 53(2), 361–382.
- Stubbs, Michael (1983). *Discourse analysis: The sociolinguistic analysis of natural language*. Chicago: The University of Chicago Press.
- Suchman, Lucy A. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge, England: Cambridge University Press.
- Svartvik, Jan and Quirk, Randolph (eds.) (1980). *A corpus of English conversation*. Lund: C W K Gleerup (Lund studies in English 56).
- Tennant, Harry R.; Ross, Kenneth M.; and Thompson, Craig W. (1983a). "Usable natural language interfaces through menu-based natural language understanding." *Proceedings of the Conference on Human Factors in Computing Systems*, 1983, 154–160.
- Tennant, Harry R.; Ross, Kenneth M.; Saenz, Richard M.; Thompson, Craig W.; and Miller, James R. (1983b). "Menu-based natural language understanding." *Proceedings, 21st Annual Meeting of the Association for Computational Linguistics*, Cambridge, MA, 151–158.
- Valian, Virginia and Wales, Roger (1976). "What's what: Talkers help listeners hear and understand by clarifying sentential relations." *Cognition*, 4, 155–176.
- Weizenbaum, Joseph (1966). "ELIZA: A computer program for the study of natural language communication between man and machine." *Communications of the ACM*, 9(1), 36–45.
- West, Candace (1984). "Medical misfires: Misgivings, mishearings, and misunderstandings in physician-patient dialogues." *Discourse Processes*, 7(2), 107–134. A revised version appears in: West, Candace. *Routine complications: Troubles with talk between doctors and patient*. Bloomington, IN: Indiana University Press, 1984.

Graeme Hirst is the author of *Semantic interpretation and the resolution of ambiguity* (Cambridge University Press, 1987). His current research concerns repair and reconstruction in human-machine conversations. Hirst's address is: Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4; e-mail: gh@cs.toronto.edu

