A Theory of the Multiple Roles of Diagnosis in Collaborative Problem Solving Discourse

Cynthia S. Gadd
School of Business Administration
College of William and Mary
Williamsburg, VA 23185
csgadd@mail.wm.edu

Abstract

A better understanding of the nature of consultations between professionals engaging in the collaborative process of solving complex problems — expertise *in use* — offers the potential to reshape our ideas about how to design computer systems that can engage in collaborative problem solving with their human cohorts. The research reported here has sought to account for key behaviors contributing to successful consultation, as identified by a cognitive task assessment of human-human consultation discourse in the medical teaching rounds setting. We have come to view the communication acts of the presenter/investigator as evidence of his deliberate intention to indirectly construct a particular model of the patient's case — *his model* — in the expert's mind, resulting in two separate but related diagnostic tasks for the expert: one at the patient level and one at the presenter/investigator level. This dual-diagnostic theory of expert understanding of the presenter/investigator's communication actions is partially implemented in the RUMINATE program. The theory provides insights into the expert's capacity to model aspects of the presenter/investigator's competence — insights that contribute to our understanding of expertise embedded in the context of collaborative problem solving discourse.

Introduction

In (Evans & Gadd, 1989; Gadd, 1995; Gadd & Pople, 1988; Gadd & Pople, 1990) we describe a cognitive task analysis of medical rounds discourse, a form of consultation in which one of the partners, the presenter/investigator, may not be a domain expert but possesses a competence in the domain and has been responsible for gathering and performing the initial analyses of the data available about the problem to be solved. The second partner, the expert or consultant, has the benefit of years of experience solving problems in the domain, but in this instance he is reliant upon the presenter/investigator for access to much of the case data. We claim that such consultations require a complex model of expertise, in which skillful execution of the roles of teacher, critic, and communicator leads to superior collaborative problem solving.

Our analysis of a large corpus of teaching rounds discourse data[1] led us to assert that much of the expert behavior observed in rounds could possibly be explained by detecting the *presence* of local incoherencies in the discourse relative to the expert's model of the problem state — a single-model approach which we suggest is sufficient for readily recognizable divergences from the expert's model (see Evans & Gadd, 1989). Our analysis also indicates that there were instances when the expert focused upon the *source*, in addition to the presence, of incoherency, requiring a multiple-model (expert/non-expert) approach in which the expert seeks to uncover the reasoning processes by which the presenter/investigator has come to see the problem. In (Evans & Gadd, 1989) we suggested the possibility of generating an alternative agent model from the expert model, for those situations when the single-model approach was inadequate; and, in (Gadd & Pople, 1988; Gadd & Pople, 1990), described an early version of a multiple-model approach to diagnostic reasoning in collaborative discourse settings. In this paper, we refine this theory of the expert's reasoning regarding the presenter/investigator's intended model of patient-specific reasoning in a clinical consultation and introduce RUMINATE, a computer-based simulation intended to demonstrate the plausibility of the theory. The contributions of this work are addressed in the introduction to the following section and in the conclusions.

A Theory of the Multiple Roles of Diagnosis in Consultation

Stated in general terms, diagnosis is the formulation and evaluation of competing hypotheses that could account for some observed behavior or evidence. In this model we are proposing that the expert uses two very similar versions of this generic diagnostic process to perform inference in two distinctly different hypothesis spaces. These hypothesis spaces correspond to the expert's intertwined goals of *a)* understanding the patient's medical problem, and *b)* evaluating the expertise of the presenter/investigator (the person presenting the patient data, whom we assume to also be the physician working up the case). In medical teaching

[1] The data set consists of eleven complete internal medicine teaching rounds cases recorded at hospitals affiliated with the University of Pittsburgh School of Medicine and Mercy Hospital in Pittsburgh, PA. We also analyzed transcripts of specialty consultations, including those conducted face-to-face and by telephone. The cognitive task analyses of the more than twenty hours of transcripts utilized techniques adapted from linguistics and discourse understanding research.

rounds, case-related observable behavior available to the expert arises from several sources — including the presenter/investigator's communication acts (usually in the form of reports of findings or manifestations of disease that are present or absent in the patient; see Figure 1 for an example), a physical exam performed by the expert at the bedside, and review of imaging output and laboratory data.

Presenter/investigator: ...
his complaints are those of um dysphagia
which is ah marked by sticking in
the low retrosternal area
and some pains
a dull pain in that area
ah the pain is dull
it's not of a burning quality
and it doesn't radiate anywhere ...

Figure 1: Example of presenter/investigator's reports
Case 4 (lines 30-46 in the original transcript)

This research has focused upon the presenter/investigator's communication acts as the source of evidence for both diagnostic processes. The first diagnostic process focuses on an explanation of the patient's observed (or at least reported) medical condition, in terms of diseases that could have caused or contributed to a specific set of findings. The second diagnostic process, and the part of the proposed model demonstrated by the RUMINATE implementation, focuses on a very different explanation task: how the presenter/investigator's communication acts, in the form of reports about the patient's medical problems, can be explained in terms of plausible patterns of medical diagnostic reasoning. We view the communication acts of the presenter/investigator as evidence of his deliberate intention to indirectly construct a particular model of the patient's case — *his model* — in the expert's mind. There are several key phrases in this description of the presenter/investigator-level of diagnosis, each of which will be discussed in turn.

Deliberate intention. First, by referring to the presenter/investigator's *deliberate intention* we emphasize the relationship between this research and the basic idea of the early work of Allen, Cohen, and Perrault (Allen & Perrault, 1980; Cohen & Perrault, 1979): that acts of communication were planned in order to achieve certain goals.[2] The presenter/investigator's presentation is the result of his deliberate intention to influence the expert's beliefs by having the expert construct a model of the patient's case that is congruent with his own, thereby demonstrating his diagnostic and communicative competence. For his part, the expert is a willing constructor of such a model, since it serves his goal of evaluating the presenter/investigator's diagnostic and communicative competencies. Therefore, as he uses the presenter/investigator's communication acts to put this model together, he also judges it for completeness, lack of ambiguity, and accuracy.

Indirect construction. Second, as we discussed in our analysis of the transcript data (see Evans & Gadd, 1989), the *process* of diagnosis is not explicitly featured in the presenter/investigator's communication acts. The institutional norms of the Standard Order of Presentation (SOP) and the local topic structure of rounds discourse stipulate the terms of a contractual exchange between the expert and the presenter/investigator that can be characterized as a dialogue in which the presenter/investigator relies on *indirect construction* to convey his intended model. There are several reasons for these institutional norms, including *a)* the difficulty of reconstructing the extensive diagnostic activity that was performed by the presenter/investigator prior to the presentation and *b)* the appropriateness of examining a specific diagnosis from multiple perspectives (e.g., different physiological systems).[3] In adhering to the terms of this contractual exchange, the presenter/investigator uses the content and timing of the communication acts that comprise his presentation — his choices of which patient findings to present and when to present them — to indirectly construct his intended model.

Evidence interpretation. Finally, our reference to *evidence* of the presenter/investigator's deliberate intention to indirectly construct *a particular model of the patient's case ... in the mind of the expert* corresponds directly to the second level of diagnostic reasoning described in the opening paragraph of this section: how the presenter/investigator's communication acts can be explained by the expert in terms of plausible patterns of patient-specific diagnostic reasoning. Note that in this second diagnostic process, it is the communication act of reporting a finding, rather than the content value of the finding, that serves as observed behavior or evidence in the diagnostic process. Although the source of evidence for both diagnostic processes is a presenter/investigator's utterance, such as "He has no jaundice," the actual evidence interpreted is different for each. The expert's patient diagnostic process uses the *content value of the finding reported*, in this case negative for jaundice, to downgrade the causal or contributory role of liver involvement. The expert's second diagnostic process, aimed at understanding the reasoning processes of the presenter/investigator, uses the *act of reporting the finding* as possible confirming evidence that the presenter/investigator has considered liver involvement as he attempted to perform his diagnosis of the patient. In the first case, the utterance containing the finding report impacts the expert's current understanding of the patient. In the second case the utterance impacts the expert's current understanding of the how the presenter/investigator intends for the expert to believe he has reasoned in his attempt to understand the patient.

The distinction between these two diagnostic processes may seem subtle, but we would argue it is essential to

[2] Many researchers have explored the relationship between beliefs, intention, and discourse plans (Carberry, 1983; Grosz & Sidner, 1986; Litman, 1985; Moore & Swartout, 1989; Pollack, 1986).

[3] Cf. (Cicourel, 1990) discussion of the social complexity of collaborative medical diagnosis.

constructing a complete, multidimensional model of medical expertise. The first, diagnosing the patient, is crucial to the physician's long recognized role of medical problem solver, bringing to bear his knowledge of disease, physiology, etc., in a novel patient setting. The second, diagnosing the presenter/investigator, is equally crucial to the physician's role of consultant to less experienced or peer physicians, in which communication of medical knowledge (as both speaker and hearer) and evaluation of the other person's expertise are important tasks.

Contribution. The dual-diagnostic model provides important insights into the relationship between the presenter/investigator's communication actions and the expert's capacity to model aspects of the presenter/investigator's competence regarding *a)* the presentation of information that is consistent with expectations of an institutionalized setting; and *b)* the appropriateness of the patient-specific reasoning he has intended to communicate by his presentation. Specifically, the model identifies those situations in the rounds discourse in which the expert can no longer bring his understanding of the problem into concordance with the presenter/investigator's reports, and offers plausible alternative reasoning patterns as potential sources of the discordance. Our analyses of rounds discourse data indicate that these types of situations often result in expert interventions (i.e., interruptions of the presenter/investigator's presentation in which the expert asks a relationship-seeking question in prelude to an explanation or Socratic-style dialogue), and as such, are key to understanding expertise embedded in the context of collaborative discourse. The following description of these two diagnostic processes, and the specific nature of their respective hypothesis spaces, contrasts the meaning and use of observed behavior as evidence in each.

Patient-level Diagnosis

The focus of the expert's first diagnostic process is the patient's medical problem. In this diagnostic process, a hypothesis is formed for each disease that could be considered a candidate for explanation of the medical problem. The hypothesis space is made up of all diseases that should be considered in an attempt to account for observed behavior or evidence, here defined as one or more medical findings and their associated values (e.g., usually positive, negative, or a numeric value.) The diagnostic process consists of forming and evaluating hypotheses about which disease(s) would best account for the medical findings as they are provided as evidence.[4] Since there are often competing hypotheses[5], there are situations in which quandaries arise as to which hypothesis best accounts for the finding(s). A classic quandary in clinical medicine is that between assimilating a new finding with an existing hypothesis versus seeking an independent cause for it. An example of a patient-level quandary, derived in this case from the introduction of a positive finding of retrosternal pain, is shown in Figure 2.

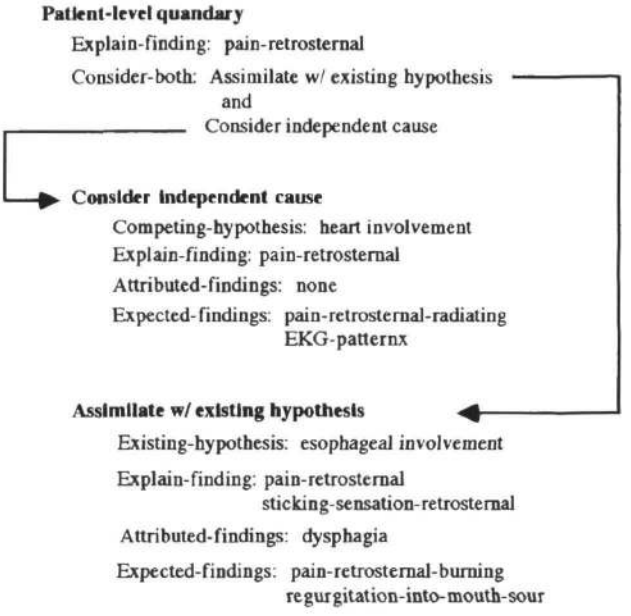

Figure 2: Patient-level quandary

The process of diagnosing the patient's medical problem includes determining what additional evidence, in the form of findings, would be useful to the expert in understanding the patient's problem. In the rounds setting, these additional findings (also shown in Figure 2) form a set of expectations for findings to be included in future reports by the presenter/investigator. Each new finding is interpreted with respect to its ability to increase or decrease the support for diseases in the existing hypothesis space or to change the formulation of the hypothesis space more dramatically, such as when existing patient-level quandaries are resolved or new quandaries are created. The diagnostic process continues in this fashion, with the ultimate goal of converging on a consistent explanation of the patient's set of findings, in the form of one or more diseases.

Presenter/investigator-level Diagnosis

The focus of the expert's second diagnostic process is the construction and evaluation of the presenter/investigator's intended model of patient-specific reasoning. In this diagnostic process, a hypothesis is a pattern of reasoning about the patient-level diagnosis that could be attributed to the presenter/investigator; and the hypothesis space is the set of these reasoning patterns that could plausibly account for

[4] Although this research relies upon the hypothetico-deductive method most frequently associated with descriptions of medical problem solving, other research suggests that physicians' selection of data-gathering and clinical reasoning strategies "is dependent on many factors, including knowledge, the discourse content, and the completeness and precision of information provided by the patient." (Patel, Evans & Kaufman, 1989, p. 308).

[5] This set of competing hypotheses is often referred to in clinical medicine as the differential diagnosis.

the presenter/investigator's observed communication behavior — usually a report of a medical finding. Contrasting this form of evidence with that used at the patient diagnosis level, it is the *act of reporting* the finding, not the content value of the finding (e.g., positive, negative), that appears to be relevant to understanding the reasoning of the presenter/investigator.[6]

In his task of constructing the model of patient-specific reasoning intended by the presenter/investigator's communication actions, our analyses suggest that the expert uses his own reasoning about the patient's problem as the nominal trajectory through the space of possible reasoning patterns.[7] In other words, in the absence of indications to the contrary, the expert assumes the presenter/investigator's reasoning about the patient's problem should approximate his own. Like all default values, the expert's use of his own patient-level diagnostic process as a hypothesis with respect to how the presenter/investigator has reasoned about the patient's medical problem has the potential to be disproven by subsequent evidence. Therefore, the expert's hypothesis space is also populated by plausible alternatives to the expert's preferred reasoning pattern. These alternatives may not be correct, in that they may lead to incomplete or even wrong patient-level diagnoses, but they do provide the expert with some directions to explore if his nominal trajectory is unable to account for the presenter/investigator's subsequent utterances.

Referring again to the point in our example when the new finding of retrosternal pain is reported, the expert's patient-level diagnostic process formed a quandary between "assimilate with the esophageal problem", or "pursue one or more independent causes, such as heart involvement." However, it is also plausible that the presenter/investigator may not have considered the heart involvement at all — preferring to assimilate the pain finding into the existing esophageal problem without consideration of any independent causes. The ultimate resolution of this issue is dependent upon evidence to be found in the presenter/investigator's subsequent utterances (either voluntary or perhaps in response to indirect or direct questioning by the expert). Therefore the expert forms a second type of quandary, this one at the level of understanding the presenter/investigator's reasoning. This *presenter/investigator-level quandary* is between two plausible reasoning patterns, either of which could account for the presenter/investigator's observed behavior:

- Is he following the expert's nominal reasoning pattern (called the nominal path), in which a patient-level quandary is created between assimilating the pain finding with the esophageal problem or seeking an independent cause? or
- Is he following an alternative pattern of reasoning (called the alternative path) in which the pain is assimilated within the esophageal problem without consideration of heart or other independent causes?

In our example, this quandary forms the current hypothesis space in the presenter/investigator-level diagnosis, shown in Figure 3.

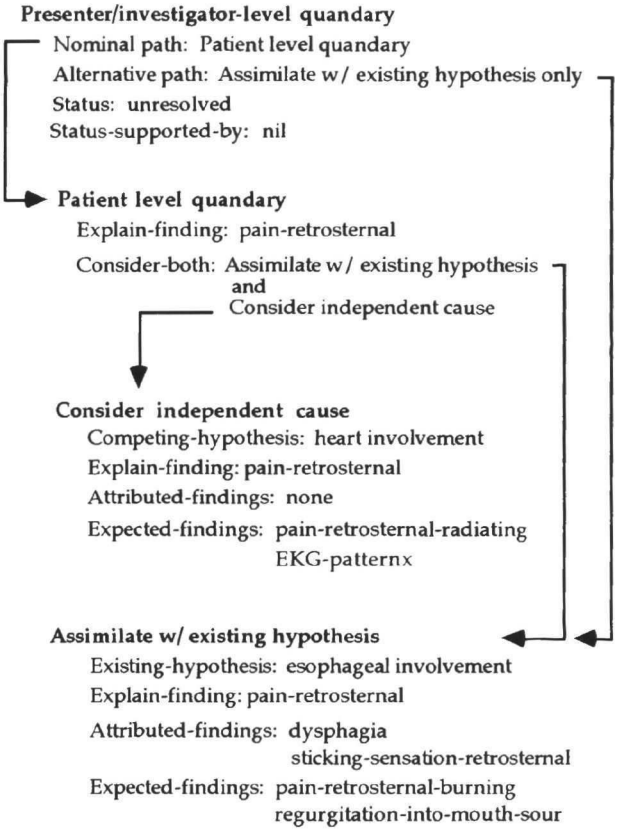

Figure 3: Presenter/investigator-level quandary

The presenter/investigator-level diagnostic process includes determining what additional evidence, in the form of reports of findings, would be useful in furthering the expert's understanding of the presenter/investigator's reasoning. In this example, the expert can use the presence of heart finding reports at appropriate times within the organization of the rounds presentation as evidence against the "assimilate with existing problem only" leg of his quandary about the presenter/investigator's reasoning (and the absence of such reports as evidence for the "assimilate-only" leg).

Each new report by the presenter/investigator is interpreted with respect to its ability to increase or decrease the support for one of the competing reasoning patterns in the existing hypothesis space at this level of diagnosis. A new report may also have the effect of changing the form of the hypothesis space, by resolving a quandary at this level or causing a new one to be created. Quandary resolution is demonstrated in our example when the finding of "no radiating quality associated with the retrosternal pain" is

[6] Evidence may also arise from the failure to report an expected finding or even the act of reporting the ordering of a test without reporting its presently unavailable result. These sources of evidence are discussed fully in (Gadd, 1995).

[7] The expert's nominal trajectory refers to the reasoning that he has used to arrive at his current understanding of the patient, i.e., his patient-level diagnostic process. The nominal trajectory does not necessarily imply correctness in the absolute sense; instead it serves as a default hypothesis regarding the presenter/investigator's reasoning.

reported. This report of a finding, which could indicate a heart problem such as angina but is not linked to any of the esophageal diseases, is likely to be interpreted by the expert as resolving his quandary about whether the presenter/investigator was considering heart involvement as well as esophageal involvement as a possible cause for the pain. In other words, the expert is able to assume, with more certainty as a result of this latest report, that the presenter/investigator's reasoning is congruent with his own.

If, in this situation, the presenter/investigator does not report about the radiating quality of the pain before moving on to some other topic, the expert could interpret this "failed expectation" as resolving the existing quandary about the presenter/investigator's reasoning in favor of the "assimilate-only" path — or at least as supporting this alternative over his own preferred reasoning pattern. The expert has reached an *evaluation point* due to a perceived lack of congruence between his reasoning and that of the presenter/investigator. At such an evaluation point, the expert wants to explain the source of the apparent incongruence. He may decide to intervene with a direct or indirect question to the presenter/investigator regarding consideration of heart involvement in this case. Or the expert may decide to defer action until some later point, continuing to gather evidence that may (or may not) provide insight into the source of the apparent incongruence. In the rounds discourses we analyzed, evaluation points such as these appear to be an important source of the expert's discourse interventions, i.e., questions or explanations.

Finally, consider the situation if the presenter/investigator had reported ordering a complete heart work-up instead of reporting the negative finding for radiating pain. Such a report would be of no value in the diagnosis of the patient, since it neither increases or decreases the support for any of the diseases hypothesized. However, it serves nicely as evidence to downgrade or eliminate the possibility that the presenter/investigator has not considered the potential causality between the retrosternal pain finding and heart involvement in his reasoning about the patient. In this situation, the report of ordering a diagnostic-specific procedure suggests that the presenter/investigator's reasoning and the expert's reasoning are congruent.[8]

These examples of how three quite different resolutions to the original presenter/investigator-level quandary could occur serve to reiterate the points made earlier about what constitutes evidence at the presenter/investigator level of diagnosis. First, it is not the content value of the finding but the act of reporting it that matters since either a positive or negative value for the finding of radiating pain would have served equally well as evidence in favor of the presenter/investigator having considered heart disease. In our analyses of the teaching rounds transcripts (Evans & Gadd, 1989) we observed that many so-called "significant negatives" seem to be attributable to this type of signaling on the part of the presenter/investigator. Second, evidence about the presenter/investigator's reasoning can also be found in his failure to report a finding at an appropriate time in the rounds presentation — as demonstrated by the variation on our example in which the he does not mention the radiating quality of the pain before moving on to another topic. Our theory illustrates the interdependence of discourse structure and domain-specific reasoning by using transitions between segments in the SOP and between (sub)topics in the local topic structure as "forcing points" for the consideration of unreported but expected findings. Third, the presenter/investigator level of diagnosis may usefully interpret evidence that does not inform the patient level diagnostic process, such as the heart work-up order.

The interpretation of each new finding report may cause modifications to the presenter/investigator-level hypothesis space, which could occur as a result of the interpretation itself or as a result of pursuing the types of evaluation points described above. Following this interpretation, expectations for reports to be subsequently included in the rounds presentation are also revised. So the diagnostic process at this level continues, with the goal of achieving and maintaining congruence between the expert's and presenter/investigator's reasoning about the patient; or in other words, to facilitate the development of a shared understanding of the patient's medical problem.

The RUMINATE Program

The purpose of the RUMINATE program was to take what has been a fairly theoretical discussion of how consultations work and enable the level of discourse to be brought to a much more concrete level. The RUMINATE program requires several knowledge sources: knowledge about how clinicians perform patient-specific diagnostic reasoning, including a representation language for individual reasoning actions and the permissible relationships between them; knowledge about the types of uncertainty (called quandaries) that an expert clinician may have about the reasoning of a presenter/investigator; and knowledge about what constitutes evidence that will contribute to the resolution of quandaries that an expert clinician may have about the reasoning of a presenter/investigator. The reasoning processes performed by RUMINATE closely follow the presenter/investigator-level diagnostic reasoning process described in the previous section. First, RUMINATE models how an expert clinician uses both the content and timing of a presenter/investigator's communication actions to gain insight into the presenter/investigator's patient-specific diagnostic reasoning. This is accomplished through plan induction, in which a default reasoning path, its alternatives, and uncertainties with respect to the presenter/investigator's adherence to one or more of these paths — are inferred. Second, RUMINATE models how the expert clinician seeks to reduce uncertainty with respect to his insight into the presenter/investigator's patient-specific diagnostic reasoning, in the following situations:

[8] There are still some open issues here, such as does the pain have a radiating quality, is it appropriate to report a heart work-up at this point in the rounds presentation (without reporting clinical evidence for doing so), and what exactly is meant by heart work-up (e.g., a cardiologist consult, a specific set of tests)? These issues may trigger evaluation points similar to the one described above.

- when the expert doesn't know how the presenter/investigator is reasoning, by assuming the presenter/investigator is following the default (is "with me") until evidence to the contrary presents itself[9];
- when the expert knows/suspects the presenter/investigator's reasoning differs from his own, by identifying the alternative reasoning pattern and establishing an explanation task to find out its source; and
- when the expert believes uncertainty has been reduced/eliminated regarding his understanding of presenter/investigator's reasoning, by establishing expectations for subsequent presenter/investigator reports that are consistent with this understanding and creating explanation tasks when they fail.

The RUMINATE program is described fully in (Gadd, 1995), including examples of its behavior and evaluation of the coverage it provides for a range of rounds discourse exemplars.

Conclusions

The primary goal of this research has been to gain a better understanding of the nature of consultations between professionals engaging in the collaborative process of solving complex problems. Our focus on expertise *in use*, rather than isolated problem solving, has led us to consider two important aspects of communication actions in consultation settings: their intentional role within the institutional forms of discourse that characterize consultations and their specific relationship to the expert's capacity to develop a competence model of the presenter/investigator. Here we have presented a theory to account for a wide range of behaviors that were observed in the actual teaching rounds discourses. The theory was subsequently described in the concrete terms of a simulation, thereby allowing assessment of its strengths and limitations, and providing a base from which to explore unresolved issues.

Communication settings characterized by uneven knowledge distribution, mentorship roles, institutional discourse norms, and a multiplicity of participant objectives are prevalent in professions other that medicine, such as law and business. In public accounting, for example, junior audit staff routinely report their field findings to more senior auditors in consultation-like sessions in which critique, advice, and assessment are active goals. The model proposed and explored by this research provides a basis for facilitating a supportive environment for human-computer problem solving discourse in such settings. Both computer-based consultants and learning environments could benefit from the opportunity to gain a better shared view of the problem to be solved, through exploration of *the process* of domain reasoning afforded by this model.

Acknowledgments

The research reported here is part of my dissertation, done at the Decision Systems Laboratory and the Joseph M. Katz Graduate School of Business at the University of Pittsburgh. It was funded in part by the Josiah Macy, Jr. Foundation. Harry E. Pople, Jr., William E. Spangler, and Karen Locke provided helpful comments and suggestions on an earlier draft of this paper.

[9] Cf., discussion of "fit-deficit triggers" in (Evans & Gadd, 1989, pp. 243-245).

References

Allen, J. F., & Perrault, C. R. (1980). Analyzing Intention in Utterances. *Artificial Intelligence*, 15(3), 143-178.

Carberry, S. (1983). Tracking User Goals in an Information-Seeking Environment. *In Proceedings of the Conference of the American Association for Artificial Intelligence* (pp. 200-206). Washington, D.C.

Cicourel, A. V. (1990). The Integration of Distributed Knowledge in Collaborative Medical Diagnosis. In J. Galegher, R. E. Kraut, & C. Edigo (Eds.), *Intellectual teamwork: social and technological foundations of cooperative work* (pp. 221-242). Hillsdale, NJ: L. Erlbaum Assoc.

Cohen, R., & Perrault, C. R. (1979). Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science*, 3(3), 177-212.

Evans, D. A., & Gadd, C. S. (1989). Managing Coherence and Context in Medical Discourse. In D. Evans & V Patel (Eds.), *Cognitive Science in Medicine*, (pp. 211-255). Cambridge, MA: MIT Press.

Gadd, C. S. (1995). RUMINATE: A Model of the Multiple Roles of Diagnosis in the Communication and Evaluation of Expertise. Doctoral dissertation. Pittsburgh: University of Pittsburgh.

Gadd, C. S., & Pople, H. E. (1988). An Interpretation Synthesis Model of Medical Teaching-Rounds Discourse: Implications for Expert System Interaction. *International Journal of Education Research*, 1(1), 81-102.

Gadd, C. S., & Pople, H. E. (1990). Evidence from Internal Medicine Teaching Rounds of the Multiple Roles of Diagnosis in the Transmission and Testing of Expertise. In N. Fredricksen (Ed.), *Diagnostic Monitoring of Skill and Knowledge Acquisition*, (pp. 89-112). Hillsdale, NJ: Lawrence Erlbaum Associates.

Grosz, B. J., & Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3), 175-204.

Litman, D. J. (1985). Plan Recognition and Discourse Analysis: An Integrated Approach for Understanding Dialogues. Doctoral dissertation. Rochester: University of Rochester.

Moore, J., & Swartout, W. (1989). A reactive approach to explanation. In *Proceedings of the Eleventh Annual International Joint Conference on Artificial Intelligence* (pp. 1504-1510). Detroit, MI.

Patel, V L., Evans, D. A., & Kaufman, D. R. (1989). A Cognitive Framework for Doctor-Patient Interaction. In D. Evans & V Patel (Eds.), *Cognitive Science in Medicine*, (pp. 257-312). Cambridge, MA: MIT Press.

Pollack, M. (1986). Inferring Domain Plans in Question-Answering. Doctoral dissertation. Philadelphia: University of Pennsylvania.