Cognitive and Interactional Aspects of Correction in Tutoring

Barbara A. Fox

Technical Report #88-2

Unclassified
SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY 2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
4. PERFORMING ORGANIZATION REPORT NUMBER(S) 1530634	5. MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION University of Colorado 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Personnel and Training Research Programs Office of Naval Research (Code 1142PT)
6c. ADDRESS (City, State, and ZIP Code) Campus Box 345 Boulder, CO 80309-0345	7b. ADDRESS (City, State, and ZIP Code) 800 North Quincy Street Arlington, VA 22217-5000
8a. NAME OF FUNDING / SPONSORING ORGANIZATION 8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DUNS-067-001412
8c. ADDRESS (City, State, and ZIP Code)	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. / PROJECT NO. / TASK NO. / WORK UNIT ACCESSION NO.

11. TITLE (Include Security Classification)

Cognitive and Interactional Aspects of Correction in Tutoring

12. PERSONAL AUTHOR(S)
Barbara A. Fox

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Technical	FROM ____ TO ____	88, 01, 19	35

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

As intelligent tutoring systems have grown in power, there has been an increasing desire for more robust interfaces which can manage such necessary tutoring tasks as correction, intervention, redirection, and the like. Until now, systems which have tried to model these tasks have done so on the basis of researcher intuition, or a few psychological experiments. The research presented here begins to remedy this situation by providing detailed evidence from video-taped human-human tutoring sessions regarding effective strategies for correction, tutor intervention (what I have called assistance), and a variety of other topics in tutorial dialogue. The findings discussed here should be taken as critical areas for exploration in the development of future intelligent tutoring systems.

Continued

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified
22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE (Include Area Code) / 22c. OFFICE SYMBOL

DD Form 1473, JUN 86 Previous editions are obsolete. SECURITY CLASSIFICATION OF THIS PAGE

19. Abstract (continued)

The present paper explores the interactional and cognitive processes involved in correction and tutor assistance in a set of tutoring sessions. As we will see in sections 2 and 3 these two facets--interactional and cognitive--of correction and assistance are really flip sides of one another: tutors and students interact in a way that is cognitively beneficial to the student and the cognitive processes needed are in turn at least partially determined by the interaction between tutor and student. In a sense, interaction and cognitive live together in an ecological system and thus have "evolved" such that their characteristics are intertwined.

The remainder of this section of the paper provides an overview of the project within which this work is situated. Section 2 describes the interactional aspects of correction (including tutor assistance). Section 3 outlines some of the cognitive processes involved in correction, and section 4 presents some ideas for future systems.

Cognitive and Interactional Aspects of Correction in Tutoring

Barbara A. Fox
Dept. of Linguistics and
Institute of Cognitive Science
University of Colorado, Boulder

1 Introduction

As intelligent tutoring systems have grown in power, there has been an increasing desire for more robust interfaces which can manage such necessary tutoring tasks as correction, intervention, redirection, and the like. Until now, systems which have tried to model these tasks have done so on the basis of researcher intuition, or a few psychological experiments. The research presented here begins to remedy this situation by providing detailed evidence from video-taped human-human tutoring sessions regarding effective strategies for correction, tutor intervention (what I have called assistance), and a variety of other topics in tutorial dialogue. The findings discussed here should be taken as critical areas for exploration in the development of future intelligent tutoring systems.

The present paper explores the interactional and cognitive processes involved in correction and tutor assistance in a set of tutoring sessions. As we will see in sections 2 and 3 these two facets--interactional and cognitive--of correction and assistance are really flip sides of one another: tutors and students interact in a way that is cognitively beneficial to the student and the cognitive processes needed are in turn at least partially determined by the interaction between tutor and student. In a sense, interaction and cognition live together in an ecological system and thus have "evolved" such that their characteristics are intertwined.

The remainder of this section of the paper provides an overview of the project within which this work is situated. Section 2 describes the interactional aspects of correction (including tutor assistance). Section 3 outlines some of the cognitive processes involved in correction, and section 4 presents some ideas for future systems.

1.1 Overview of Project

The research within which the prepsent paper is situated seeks to gain an understanding of some of the processes of human tutorial dialogue, including conversational repair, tutor intervention, and linguistic encoding of communicative strategies. It also seeks to characterize the effect on these processes of manipulating several variables, such as mode of communication (i.e., face-to-face vs. terminal-to-terminal), subject matter of instruction, and expertise level of the student. This work is meant to provide an empirical foundation for the enhancement of design principles for tutorial components of intelligent tutorial systems and knowledge-based help systems.

In addition, the research attempts to characterize the differences between human-human tutorial dialogue and human-computer tutorial dialogue, so that a realistic picture of the dialogue requirements on current intelligent tutorial systems can emerge.

As intelligent tutoring systems grow and become more robust, they will need to enter into increasingly extensive natural language dialogues with users. As a result, these systems will need complex natural language interfaces; in particular, they will need sophisticated models of how to carry on natural dialogue.[1]

Because of this need to model human dialogue, the study of human-human instructional dialogue is a critical step in the further development of intelligent tutoring systems, in that such research will provide the theoretical foundation on which working systems can be built. That is, this project is based on the belief that the best tutorial systems will arise out of a fine interplay between what would be theoretically ideal (based on what humans actually do) and what is practically possible (based on what the current technology allows). The focus of this portion of the project is on the former; we examine naturally-occurring tutorial interactions (as one type of instructional interaction which is closely related to the problem at hand and readily observable in a university environment) in the context of furthering our understanding of how people learn and teach through extended dialogue.

The theoretical background for this particular approach to human-computer interaction is captured in the following quote from John Seely Brown:

> A dialogue involves constant conversational repairs between two people. When someone doesn't understand what I've just said, I must try to diagnose not only what, but why, he didn't understand and then accordingly repair what I said. It is basically an adjusting process that goes on between the two of us communicating....We have been misled into thinking that natural language, per se, is so powerful. Instead, I think it is the dialogue process that is so powerful, e.g. the notion of conversational repairs that occur between two people. If we can understand this process and how to capture it in man-machine

---

[1]Hayes (1983) puts forth an interesting perspective on the relationship between human conversation and the kinds of dialogue we should expect between people and computers:

> In brief, we believe that direct simulation of human conversation will not play an important role in user-friendly interfaces until speech processing has made sufficient technical advances to allow spoken language to be used freely and in conjunction with pointing, but that the study of human communication will continue to be relevant to interactive computer interfaces of all kinds. (p.230)

I agree with this statement, and I would emphasize that what we need to model from human conversation is critical communicative processes, not all of human communication or human conversation.

communication, we will have made a major breakthrough on the perceived friendliness of machines.

Brown makes two crucial points in this statement: one, the study of human-human dialogue will form a fundamental part of the development of robust human-computer interfaces; and two, even though human-computer dialogue may differ on the surface in extreme ways from human-human dialogue, at least some (if not most) of the basic underlying processes of human-human dialogue must be captured in a human-computer interface if it is to be user-friendly. Thus, even if users communicate with computers in a way that suggests they believe the computers to be less than adequate communicative partners, some of the processes of natural human-human dialogue will still be apparent. In spite of the potential differences, then, between human-human dialogue and human-computer dialogue, the study of human-human interaction will provide necessary insight into design principles of human-computer interfaces.

In addition to the theoretical support for the study of human-human dialogue offered by Brown's statement, there has been a precedent in AI research for the approach proposed here, in the work of Grosz (1977), Reichman (1981), Hobbs (1976), and others. In these studies, findings from human-human interactions were effectively modeled in computer systems. The widespread recognition of this work throughout the cognitive sciences indicates that it has added significantly to our understanding of human cognitive processes in general, as well as to the robustness of intelligent systems. It is clearly a fruitful avenue of research to pursue.

Nonetheless, it is important to recognize that human-computer instructional dialogue will differ in critical ways from human-human instructional dialogue. The proposed research in part seeks to explore the nature of the differences between human-human instructional dialogue and human-computer instructional dialogue in order to determine what the current dialogue demands on tutorial systems are. By studying both types of dialogue, it will be possible to reach a balance between what is theoretically ideal and what is currently required to achieve an optimal model of dialogue for intelligent tutorial systems.

1.2 Data for the Study

The data for the study have been gathered in the following way. Graduate students from chemistry, physics, computer science and math were located and asked to participate as tutors in the study. They were paid for their work, in keeping with normal tutoring agreements. Students were located by advertising in the student newspaper for people interested in being tutored as part of a research project. Each student-tutor pair met for a total of three hours, spaced out over three one-hour sessions. They were sequestered in a small lab room and video-taped. They were given no instructions concerning the content or structure of their sessions.

The students so far represent a fairly wide spread of expertise, from nearly complete novice to intermediate learner.

Modes of communication: Two "modes" of communication were examined. The first is simple face-to-face, naturally-occurring dialogue between two people. The second is slightly different, in that the two people communicated via terminals (what I'll call the terminal-to-terminal mode), instead of face-to-face. In the terminal-to-terminal mode, only the student was video-taped, and the dialogues were stored electronically by the computer the subjects were using.

Supposed nature of tutor: Within the terminal-to-terminal mode, there are two subconditions. In the first, the student is told (correctly) that s/he will be communicating with another person, who will be using a terminal in another room. In the second, the student is led to believe (incorrectly) that s/he is communicating via the terminal directly with an intelligent tutorial computer system, when in fact a hidden human tutor is responding to the student. In the second case, the human acting as the computer tutor was instructed to produce only a limited set of grammatical constructions and also was instructed to respond with error messages when the student used grammatical constructions outside the set "available" to the tutor. In this way, the student was provided with a situation that was as close to a real human-computer interaction as possible, without requiring a highly sophisticated intelligent tutoring system.

Several hours of each type of tutorial interaction were collected. When possible, the same person served as the tutor in all of the interactions within a domain. The video-taped interactions were transcribed according to the method of Sacks, Schegloff and Jefferson (1974).

Before turning to data, I would like to present and discuss some of the notational conventions used in the transcripts.

All talk is transcribed in standard American English orthography.

A double slash (//) indicates the place at which a speaker's utterance is overlapped by talk from another speaker.

M. No. They're a//ll thin.

C. They're not

Thus with this notation we can see that C's utterance starts after the a in M's all.

An utterance which has more than one double slash in it is overlapped at more than one place, and the utterances which do the overlapping are given in sequential order after the overlapped utterance.

G. they're all Keegans like the ones around Greensprings

they're all kind'v, // bout five five, five si//x,

M. They're all from around Greensprin//gs.

C. Yeh

Here, M's utterance overlaps G's starting at bout, and C's utterance overlaps G's starting with the x in six. Notice that C's utterance also overlaps the very end of M's.

A left-hand bracket at the beginning of two lines indicates that the two utterances begin simultaneously.

M. Yeh.
[
C. Lo:ng time ago it reminds me

M and C begin talking simultaneously.

The equals sign (=) indicates latching, that is, the next speaker begins without the usual "beat" of silence after the current speaker finishes talking. In this case there is an equals sign at the end of the current speaker's utterance and another equals sign at the beginning of the next speaker's utterance. If two speakers simultaneously latch onto a preceding utterance (that is, they begin talking simultaneously), this is indicated in the transcript with a left-hand bracket preceded by an equals sign.

(R) (h)hh (h)uh (h)uh (h)uh!=

(S) hhh(h) H(h)m
=[
K. Which la:mpost?

Here S and K simultaneously latch onto R's laughter.

Numbers given in parentheses indicate elapsed silence, measured in tenths of seconds. Single parentheses with a dot between them represent a silence that is less then a tenth of a second but still longer than the usual beat of silence.

Certain facts about the production of the talk are given through the orthographic symbols used. Punctuation is used to suggest intonation; underlining indicates stress. A colon after a letter means that the sound represented by that letter is somewhat lengthened; a series of colons means that the sound is increasingly lengthened.

The letter h within parentheses indicates "explosive aspiration," and usually means some type of laughter is being produced. A series of h's preceded by a raised dot represents an inbreath (where number of h's is meant to correspond to the length of the inbreath), while the same series preceded by nothing represents exhaling.

Questionable transcriptions are enclosed within single parentheses; the transcribers thereby indicate that the exact form of the utterance is not clear. Speaker's initials given in single parentheses means that there is some question about the speaker's identity. Single parentheses with capitalized words -- e.g.,(CLEARS THROAT)--represent non-transcribed material (i.e., noise which is non-linguistic).

These are the major transcription conventions which will be used in the data fragments in this paper. For a more detailed guide to CA notational conventions see Sacks, Schegloff and Jefferson (1974).[2]

2 Interactional Aspects of Correction

2.1 Organization of Tutoring Interaction: Correction and Tutor Assistance

The main focus of this research to date has been understanding the process of a subcomponent of repair, namely correction, including the process of tutor intervention, or, as we prefer to think of it, tutor assistance. We have chosen to start with this large area because of its significance to a wide range of topics, including the role of the student in creating the learning environment, differences between child and adult learners, the use of interactional and linguistic resources to accomodate a less than normal situation (i.e. an otherwise competent adult in a non-competent role), questions in natural language processing, and issues in interface design.

While modeling the **cognitive** processes involved in correction is clearly a necessary task in the design of more robust user interfaces, especially in the area of ITS, it may at first seem somewhat less critical to model the **interactional** processes involved. While this view is certainly understandable given the vast differences in interactionality between humans and computers, nothing could be further from the truth. First, as we will see in section 3, it is impossible to separate entirely the cognitive processes to be modeled from the interactional ones. And second, we know that human tutors are still more effective than computer tutors (Resnick, p.c.), and so we must assume, at least for now.

[2] For a discussion of notation as theory see Ochs (1979).

that the strategies used by human tutors to accomplish things like correction are effective, and thus worthy of being modeled. Some systems, for example the ACT* system developed by Anderson (Anderson, et al, 1987), already assume a model of the interactional aspects of correction based on intuitions about what is psychologically effective for students. The research presented here aims to provide an empirical foundation for such models, by examining how tutors and students actually do carry out such necessary tasks as correction.

We have found that the best way to understand correction (including tutor assistance) is to examine the overall structure of the tutoring sessions; in other words, correction is best studied within the sequential organization of the interactions themselves. The structural details of each interaction are beyond the scope of the present paper; below I provide only enough description to motivate the analysis of correction.

The sessions divide themselves basically into two groups: the first group is characterized by the student working through problems in a textbook (in one case the student brought in a set of problems from the text that she had been having difficulties solving); the second group is characterized by something more of a lecture format, wherein the tutor talks about domain concepts and demonstrates how to solve problems. This description is an oversimplification of the sessions, but it is sufficient for our purposes here.

When students work through a problem with a tutor, they very often verbalize what they are doing, step-by-step. For example in the following excerpt, the student doesn't just write down the steps as he works through them--he talks about what he is doing as he does it:

S: So:, ah:m (1.0) this, okay, secant, of theta, I know
eqs- equals three:. Now, so but- my equation here is
(0.4) secant squared theta. So what I want to say is
.hh the square root, (0.6) of secant, (0.5) squared

theta, .hh would equal three:. Right?

Why do students produce this on-going commentary? They do it to display to the tutor how they have understood the problem and how they understand what they are currently doing to solve the problem. So every student utterance of this sort is a display of understanding. As with any understanding, the student's understanding of a particular step in a problem can match or not match the interlocutor's understanding, so that any such display of understanding calls for a confirmation (or disconfirmation) from the tutor, by which the tutor displays that the student's displayed understanding matches her own (see Labov and Fanshel, 1977, for a related point). A confirmation agrees with the student's understanding; a disconfirmation disagrees with it.

Now, it is known from past research (Pomerantz, 1975) that agreements come very quickly after the utterance they agree with while disagreements are somewhat delayed. In these tutoring sessions we have found that when the tutor agrees with the student's displayed understanding her signal of confirmation comes quickly after the student's turn, as in:

S: Mkay. .hh And I know it's negative, just to follow your
thought process, because I know that the sine is positive.
T: Mhm

S: And this (draw it out). (0.3) And the double bond goes
away
T: Right

whereas if the tutor disagrees with the student's understanding, the delivery of the disagreement is somewhat delayed (and in some cases what might be described as "hesitant"):

S: And it's going to change when I put this in- there, right?
(1.7)
T: I don't think so.

S: So that triple bond is like ess pee three?
(1.1)
T: Ah:: no:, that- a triple would be an ess pee.

So here is the first point about correction: tutor correction (or indication of a problem with the student's understanding) is delayed with regard to the relevant student utterance.

Furthermore, it follows from this and past work on disagreements (Pomerantz, 1975) that when a student has produced a display of his/her understanding and there is no immediate response from the tutor, the student can anticipate that the tutor is going to disagree with his/her understanding; students in this situation will very often re-phrase their statement as a question, thereby **inviting** correction from the tutor:

S: Okay, just for review for my sake, .hh a cosecant

is .hh uh:m, one over the tangent.
(1.3)
S: Am I correct?

T: N:o.

S: And it's going to change when I put this in- there, right?
(1.7)
T: I don't think so.
[
S: Does the capacitance change?
(0.5)
T: I think the charge changes.

Students can also use this "predisagreement" silence to try by themselves to correct whatever may be wrong with their understanding:

The student and tutor go on this way, displaying and repairing their understandings until the student gets stuck.[3] Here again there are several alternative responses that either could make, assuming that the student is displaying overt signs of "being stuck" (and not, for example, "thinking").

Notice first that the situation of a student being stuck (and showing being stuck) creates a potential conflict. In our everyday interaction, if we see someone having difficulty in some way, it is preferred for us to offer help before that help is requested (Pomerantz, 1975). I do not mean that it is personally or psychologically preferred for us to offer help in this situation; indeed, we may be in a hurry, or not like the person, or have something else we'd rather be doing. Rather, it is preferred socially and structurally, so that if we are not going to offer help we must provide an excuse for not doing so, or pretend we didn't see the trouble, etc. Preference organizations of this sort are independent of the momentary preferences of individual participants.

So it is possible that in tutoring interactions the preferred response to the student's difficulty would be for the tutor to in some way offer help. But, as we have seen, there is a conflicting preference organization which indicates that participants should be allowed to repair their own trouble (Schegloff, Sacks and Jefferson, 1977).

The tutors in this study display an orientation to both kinds of preferences (for a similar finding with regard to compliment responses, see Pomerantz, 1978). Tutors do provide assistance, but they do so in such a way that the student is given the opportunity to unstick themselves, both before the assistance is provided and while the assistance is emerging. Consider the following exchange:

---

[3]In some cases the student avoids getting stuck by indicating before the difficult step that this next step was exactly the reason s/he needed help with this problem:

S: And kay is the: constant that, I know that (0.5) and

um (1.2) eff, that's what I had a problem with, was eff

This strategy only works if the student has worked the problem before coming to see the tutor.

S: and they want to know what the tangent is.So,

I have one over cosine of theta equals three.

(0.8) And I have the sine of theta over cosine

of theta (1.0) hmm:. (0.8) .hh Okay, so I guess

I somehow have to: (0.8) tangent of theta is going

to be: (0.4) sine of theta over cosine of theta.

(2.0) One over cosine of theta,// so (0.3) three.

T: Mkay. Now,

S: Okay

T: ts looking up here, ju//st at what=

S: Aha

T: =they've done. (0.4) .hh cause I can tell, we're

headed in the wrong direction.

S: Ye:ah, they used to con-they use // one of the

pythagoreans.

T: One plus (0.6) tangent squared e-quals the secant squared

[

S: secant squared.

The student in this case is going around in circles--he has repeated by now several times that tangent equals sine over cosine--without finding a new way to look at the problem. The tutor thus intervenes, but not without giving the student a fair opportunity to figure out the answer for himself. Furthermore, the assistance is produced in such a way that the student can collaborate in the redirection, as he in fact does with his last lines. Here again, correction/asssistance is momentarily withheld to give the student a chance to fix the problem himself. The withholding time is not long, however, and a student who wishes to be given a longer opportunity to work the problem out for him/herself must specifically request such an "extention":

S: Now, .hh let's see, when we said tangent of theta

was less than zero .hh u::hm the tangent was

(0.7)

S: give me a second. The tangent was sine over cosine.

T: Mhm

Correction/assistance of this sort thus is slightly delayed but is still offered without being "overtly" requested.

Tutors regularly provide assistance for the student if the student has produced one step in a chain of reasoning but apparently does not see the inference(s) which should be drawn for the next step:

S: Ey plus cee equals zero.

(0.3)

T: Right, so that tells // you

S: Ey equals cee.

(0.4)

T: Minus cee.

S: So it's got to be: in our fourth quadrant.

T: Right.

(1.4)

S: Aha. (0.8) Aha, o//kay.

T: Which means?

S: Then I have to come back down here and I'm- you're

asking me to choose a sign, righ//t?

T: Right

S: .hh Okay.

T: For your cosine.

The conflict between tutor providing help and student working through the trouble him/herself is overtly displayed in the following passage. The fragment starts with the student trying to determine the quadrant for the tangent given in the problem; he first gets in to trouble by giving the wrong formula for cotangent. The tutor provides help, and then the two appear to play a very tame kind of tug of war to see who is going to do the next steps of the reasoning. They subtly try to wrestle a few turns from one another, culminating in the student finally saying "let me see if I can figure that out."

S: I have to place it in a quadrant, is what you're telling
me, right?

T: Mhm

.
.
.

S: I would say: (.) uh, a cotangent, in terms of
ex wai, (0.3) is let's see, one over the uh-
cotangent is one over the uhm, (0.7) hold on
a sec, (LAUGH) uh: cotangent is one over the sine.
(0.4)

T: N:o

.
.

T: Cotangent is one over the tangent.

.
.

S: Now, if I'm thinking in terms of ex and wai, though,
(0.8) fo:r (0.8) the sake of the quadra//nts>

T: It would be cosine over sine.
(0.8)

S: Right=

T: =Which is ex over wai.

S: Okay

.
.

T: And the cosecant//is-

S: Co- cosecant, it's that's the one over the sin//e right

T: One over sine.

.

.

T: Which means that your ex value is positive.

(0.2)

S: Right.

T: . Which puts you in:

(0.5)

S: 'kay, .hh let's see if I can figure that out.

T: Okay.

Quite often, especially in more conceptual domains (such as physics), tutor assistance is provided in the form of a question whose answer will serve as a resource for getting the student unstuck if the student sees how the answer is a resource. This strategy thus has two parts: the first part requires that the student be able to answer the question, and the second requires that the student see how that answer is a resource for continuing the problem. Since both of these processes may end up involving correction, and since correction is dispreferred, this strategy is undertaken very cautiously and with a heavy degree of support from the tutor. Consider the following.

S: eff, that's what I had a problem with, was eff, they said

(3.1) if (0.5) if the electric force between them is

equal to the weight>

(0.4)

T: ts Okay.

S: So: I tried to look at the wei:ght,

(0.9)

T: And all's they give you is the ma:ss.

(0.2)

S: and it- yeah: (0.7) Oh, that's what it was, it

was the mass.

(0.8)

T: Yeah:.

(1.0)

S: Oh:, I s//ee, I want weight.

T: You wrote down mass.

(0.3)

T: Yeah, what's the difference between weight and mass.

S: I used to know this let's see
(1.7)
T: I think, (0.8) I think what it is (0.8) is that,

(0.5) n- what is uh, when you do (0.6) uh:,

gravity problems.=

S: =Right.//It's that-

T: What do you always do?
(1.2)
S: You have to multiply it by the-=

T: =by gee.=

S: =Gee. Right.

.
.
.

S: So I need to multiply this time (0.8) gee.

T: Right.

There are two tutor assistance questions in this passage, marked by arrows in the margin. The second question is of course meant as a resource for answering the first question, which the student shows difficulty answering. Notice that the tutor gives the student a fairly long space in which to answer the question, and when she shows that she is not being able to answer, the tutor then does not directly provide the answer but rather asks a further question, whose answer and import will enable the student to answer the first question. The student and tutor work together to produce the answer to the second question, at which point it is relevant for the student to indicate that she sees the import of the question for the problem they are trying to solve. In fact, the student eventually sees how the answer is a resource and goes on to the next step of the problem. It is critical to note in this passage that the tutor always provides a safety net around the student, so that if she shows signs of not being able to answer the question, the tutor offers a resource for answering; if the student shows signs of not seeing the import of a question for the problem at hand, then the tutor steers the student towards seeing the connection. All of this is of course kept in balance with not correcting or redirecting the student before she has had the opportunity to do those things herself.

Tutors also ask questions before the student has gotten stuck, often to help frame the problem, and the solution.[4] In the following passage, for example, the tutor checks to see if the student understands the kinds of units--and therefore the appropriate formulae--the problem involves:

T: (STARTS READING THE PROBLEM) A one ohm wire.

Okay. what's ohms
(1.4)
S: It's resistance

T: Okay. Is drawn out to three times its original length.

What is the resistance now.

The response to these questions are carefully monitored, as with all tutor questions, to make sure that the chances of the student producing an appropriate answer are maximized. If the student has difficulty answering the question, the tutor will provide clarification, hints, etc, as in the following example:

T: What is the speed of a three hundred and fifty

ee vee electron.Okay.
(1.2)
T: So the main thing he:re, I mean, when you look at that,

what is electron volts, what kind of a

S: It'//s ahm,

T: what are we talking about.
(1.5)
S: Isn't it- the charge of an electron times?
(0.9)
T: Right,// but what is it-what is that. Is it-=

---

[4]This framing is reminiscent of Vygotsky's concept of scaffolding (see Cole, 1978), by which the teacher in an apprenticeship situation structures a task such that it is always within the learner's abilities; as the learner's abilities increase, the teacher gradually removes the "props" which have made the task approachable, until finally the learner can perform the task without assistance.

S: The voltage?

S: =It's s:maller

(0.2)
S: I//t's
T: No- okay, I'm n- I'm a I'm a I'm not asking

a specific enough question. U:hm, (1.1) Uh

(0.3) is this units of length
?
(0.9)
S: Oh, no it's uhm
(1.9)
S: It's voltage, isn't it?

[the tutor goes on to redirect the student until she sees that the answer is energy]

For the sesions in which the tutor does most of the talking (both about general procedures and concepts and about solving problems for the student), the student can display his/her understanding by finishing the tutor's utterances, as in:

T: So this is really the integral of ex (0.8) the who- e-

ex goes in the whole time=
[
S: Right
T: =.hh plus the remainder one over ex.
[
S: one over ex. Yeah

T: In order for this to become basic,
(1.2)
S: It'll have to lose a=

T: =it'l l have to lose a hydro//gen

S: hydrogen

In fact, the tutors often capitalize on collaborative completion as a means of finding out what the student understands by starting an utterance, with a slightly rising intonation, cueing the student to finish the utterance appropriately:

T: So I need to: do what.
(2.0)
T: Multiply on the inside by
(0.1)
S: half

T: one half to get rid of that. (0.8) And so on the outside

I'm going to be:
(1.3)
T: multiplying by>
(1.7)
S: a factor of two, yeah
[
T: two

This strategy, which uses statement syntax to elicit information from the student, neatly avoids the problem of correcting a wrong answer from the student--if the student provides an inappropriate completion for the utterance, the tutor can provide the "correct" answer as if she were merely finishing her own sentence.

T: Integral of ex gives us ex squared over two> One

over ex gives
(0.4)
S: One ex
(0.5)
T: Natural log
(0.4)
S: Absolute value of ex

The tutor can also incrementally add clues to the partial utterance if the student fails to finish the utterance at the first opportunity:

T: .hh Second thing> i:s
(1.0)
T: to>
(1.7)
T: fa//ctor>

S: factor the denominator.

In this passage, the student is unable to complete the tutor's utterance until he hears the first syllable of factor: as soon as he is equiped with that clue, he is able to produce the appropriate completion. This strategy provides the tutor with a kind of metric for judging the student's understanding, inasmuch as each opportunity to complete the utterance that is passed is some indication of how well the student is keeping up (see Fox, 1987 for a detailed discussion of this phenomenon).

In some rare cases, the tutor can even end up completing her own utterance entirely, if the student fails to provide a completion:

T: the main thing with the exponential function,
(1.8) you have (0.2) ee to the ex (1.4) and its
derivative is
(2.3)
T: ee to the ex ( )

The utterance-completion strategy tends to be used mostly for working on a step of a problem (usually a problem being solved out loud by the tutor in front of the student). For more conceptual issues, or more tactical issues, the tutor will often ask a question, with the syntactic form of an interrogative:

T: Is there something we can put in here, to wipe this
part out?=
S: =Well that one will go to zero, yeah.
T: Yeah (//)
S: The ex equals zero.

Here, as with the tutor questions discussed above, the issue of correction is prominent. This format is overtly a question, which of course makes an answer from the student socially appropriate. Both parties work to avert an incorrect answer from the student. They manage this in an intricate way: the silence that grows after the tutor's question is carefully monitored for signs that the student either will or will not be able to answer the question. The student participates by displaying signs of working on a calculation, for example, or by displaying signs of confusion or lack of comprehension; the tutor participates by looking for these signs and responding appropriately. For example, if the tutor has asked a conceptual problem and the student responds by "staring blankly" at the textbook, the tutor is likely to provide assistance before the silence has grown beyond 2 seconds; on the other hand, if the tutor asks the student what the outcome of a particular calculation is, and the student responds by displaying signs of "working on it," the tutor may allow the student a fairly long silence, usually lasting until the student either answers the question or displays further signs of being stuck. Eye movement, facial expression, body posture, position of pencil, and non-linguistic verbal cues (such as sighs, inbreaths, clicks) are all used and monitored during this particularly sensitive time.

There are three central outcomes of this interaction: the student answers the question (usually correctly); the student asks a question about some portion of the problem or question, the tutor answers and the student then answers the original question; or the tutor provides assistance--in the form of clarification, hint, a more leading question--and the process starts again, most often until the student answers. In a few rare cases, the tutor provides the answer itself, but this occurs only after the student has passed several opportunities to answer. The most striking outcome of this whole process, of course, is the low rate of incorrect student answers--I found only 12 incorrect answers produced after tutor questions, out of a total of 97 possible answer slots. This low rate is seen as an achieved outcome of the processes described above, rather than a natural fact of, say, the student's IQ or knowledge of the subject matter.

Summary of Correction Strategies

Above we saw that tutors withhold correction and even assistance until the student has had an opportunity to initiate correction on his/her own. In some cases this strategy is effective in getting the student to produce a correction, but in other cases the student is unable or unwilling to produce a correction; in these cases, the tutor takes on, in some instances with the overt collaboration of the student, the task of correction or assistance. How the trouble is handled depends on where in its sequence it is produced.

There are four main positions in which the tutor engages in correction, or initiation of correction. In the first case, the student has produced a display of understanding which is in some way incorrect; the tutor withholds correction and in this space the student, anticipating disagreement from the tutor, invites correction. The tutor responds with a correction.

S: Because secant squared of theta is square root of
(0.8)
S: Can I do it that way?

S: S- can I say three minus one?
[
T: Mm::

T: No, you want to say three squared. Because the secant is three.

The second position is the case of a wrong answer produced after a tutor question. The tutor in this situation regularly initiates correction and the student attempts self-correction.

T: Did the area change?
(1.0)
S: Wouldn't the area be the same?
(0.9)
T: We only have the same amount of copper.

S: Yeah.

T: Well think of taking silly putty.

T: Like a block of silly putty like this> and you

pulled it out? What would ha//ppen?

S: It's the same.
(0.2)

T: It would get long and skinny, though.

Notice that in this case, when the correction was not invited by the student, that the tutor does not overtly correct the student; rather the tutor tries to redirect the student's thinking. The behavior of the tutor is thus clearly sensitive to the context of utterance of the problem.

The third position follows an utterance or set of utterances by the student which usually exhibit(s) that the student is stuck. In this position, the tutor regularly initiates correction and allows the student the opportunity to actually accomplish the final correction:

S: tangent of theta is going to be: (0.4) sine of theta
over cosine of theta. (2.0) One over cosine of
theta. // so, (0.3) three.

T: Mkay. Now,

S: Okay

T: ts looking up here, ju//st at what=

S: Aha

T: =they've done, (0.4) .hh Cause I can tell, we're
headed in the wrong direction.

S: Ye:ah, they used to con- they use one of the pythagoreans.

((Student working problem: 10.8))
T: .hh Where did this minus sign come from.

S: .hh It's minus ex. This minus ex shouldn't be here.

The fourth position is the rarest, and involves the student producing an utterance--usually in conjunction with working a step of a problem--that completes a tutor prompt. In this case, the tutor initiates and accomplishes correction by simply producing the rest of her original utterance, with the correct piece of completing material serving as the correction (which gives something like an embedded correction). These instances all involve low-level calculations which the student has produced while working a problem.

T: One over ex gives>
(0.4)
S: One ex.
(0.5)
T: Natural log.

We can see from these examples that where an error arises affects very much how it is handled by the tutor and the student. Thus in order to model the kinds of strategies employed in tutoring, we need models of conversational structure in tutoring, such as those presented above.

3 Cognitive Aspects of Correction

In section 2 we saw the interactional considerations relevant to correction and tutor assistance. In this section we examine some of the cognitive processes which must underly correction, to see if we can shed some light on the kinds of processes an intelligent system must be capable of to manage correction.

For the purposes of this discussion, let us use the following passage as a testbed. In this passage, the tutor asks a question, gets an incorrect answer, points out that it is incorrect, and then initiates the process for finding the correct answer.

S: I use this equation, and what I need to find the speed?
(0.4)
T: Right. And velocity (0.4) Do you know the difference
between velocity and speed?
(2.4)
S: Velocity is the distance times time
(1.7)
S: Speed is uhm
(2.9)
T: No, distance is- velocity isn't distance times//time.

S: It's not?

T: uh-uh
(0.7)
T: Think of it=

S: =Huh

T: ..hh You go: six hundred miles per hour.

S: Oh, that's uhm (0.4) that's the speed?
(0.8)
T: Yeah. .hh And velocity and speed are the same thing.
(0.3)
T: Except the thing is, is that- velocity is a vector.

And speed is just a number.

.

.

.

T: That equation's still wrong.

S: Is it?

T: Yeah, think of miles per hour.
(1.5)
T: So what is that.

S: Over time=

T: =Right.
(1.1)
S: Okay.
(0.9)
T: The one you might be thinking of is- distance is velocity

times time.
(1.3)
S: Oh, that's got to be it.He did that in class today.
(0.2)
T: Yeah, velocity times time.

.

.

T: But that's really easy just to (carry) to the side,

all of a sudden

There are two issues being dealt with in this passage: the relationship between velocity and speed, and the correct equation for velocity. When the student starts to answer the initial question, she shows that (a) she has gotten the equation for velocity wrong, and (b) she thinks there is a difference in equations between velocity and speed. She is thus wrong on two counts, and the tutor must guide her out of both errors.

The tutor begins by trying to correct the equation for velocity. She starts off by saying that the equation for velocity is not what the student has written down and then gives an example (of velocity) in vernacular formulation from which the student can deduce the correct formula for velocity (distance over time). The student misinterprets this utterance; she hears "six hundred miles per hour" as a vernacular formulation of **speed**--which in fact it is--and therefore not as a resource for correcting her equation for velocity (since she clearly still thinks that velocity and speed are different in some fundamental way). The tutor then repairs this misconception by informing the student that velocity and speed are the same. This information should allow the student to go back, using the formulation of speed, to correct for herself her equation for velocity. She fails to do this, so the tutor comes in (after giving the student a chance to fix the problem herself) with "that equation's still wrong." The tutor once again offers an English version of the formula for velocity, and this time, with a bit more prompting, the student is able to correct her equation. Notice that the last lines in the passage are the tutor indicating to the student how her error arose; as the tutor describes it, the student followed a very natural kind of process--taking the same equation and switching two of the terms in a reasonable but mathematically incorrect way:

$d=vt$ $\quad\quad$ $v=dt$ (should be $v=d/t$)

What processes enable the tutor to carry out the correction as she does? In even the most superficial analysis, at least four components must be acknowledged:

1. The tutor must be able to interpret the utterance, or set of utterances, requiring correction.
2. The tutor must be able to determine that the interpretation displayed by that set of utterances requires correction.

3. The tutor tries to understand the path by which the student was led to the error or misconception.

4. The tutor must be able to articulate a new direction for the student so that the "error" can be discovered and corrected by the student. This new direction must be designed such that it takes into account the exact error made (and its underlying misconceptions), and, if possible, the path which led to the error.

In the passage under consideration, for example, the tutor recognizes that "velocity is the distance times time" is meant as a formula for velocity, which is part of the answer to the tutor's first question. Furthermore, the tutor recognizes that this formula is incorrect and that its being incorrect represents a stumbling block for solving the problem at hand. And finally, and most interestingly, the tutor formulates a redirection, based on manoevering the student away from the particular error she made--she gives an example velocity which does not conform to the formula given by the student. She also tries to articulate the path which led to the error.

Modeling these four processes, which are only the most obvious aspects of correction, remains in my view an extremely difficult task, one whose seriousness has not been adequately acknowledged in the ITS literature. While there is an abundance of work on finding the "bugs" in students' domain knowledge, there has been very little work done on how tutors determine whether something is (a) incorrect for the purposes of the on-going problem solving, and (b) deserving of correction during the tutoring. Furthermore, given that most intelligent systems operate with canned text, it is difficult to see how the fourth process could be modeled at all, since it requires careful design of the redirecting utterances to take into account the specific error made by the student. Since the range of student errors cannot be determined in advance, it is of course impossible for a system using canned text to tailor its "utterances" to steer a student away from every possible errorful conception.

The difficulties grow considerably if we also consider correction of communicative problems, in addition to correction of domain knowledge. Consider the following passage.

T: [READING THE PROBLEM] What is the speed of a three

hundred and fifty ee vee electron.Okay.
(1.4)
T: So the main thing he:re, I mean, when you look
at that, what is electron volts, what kind of a
what are we talking about.
(1.5)
S: Isn't it- the charge of an electron times?
(0.9)
T: Right,//but what is it-what is that. Is it-=
S: The voltage?
S: =It's smaller
(0.2)
T: No- okay, I'm n- I'm a I'm a I'm not asking a specific
enough question. U:hm, (1.1) Uh (0.3) is this units
of length?
(0.9)
S: Oh, not it's uhm
(1.9)
S: It's voltage, isn't it?

The tutor in this passage must see what kind of an answer "electron volts is the charge of an electron times the voltage" is. It is, of course, a formula for how to calculate electron volts. It arises from the formulation of the question "What is electron volts, what kind of a what are we talking about" by the possible interpretation of is as 'equals', a common interpretation in physics problems. The tutor recognizes that as not the kind of answer she was looking for and sees that it will not serve as a resource for helping the student solve the problem at hand. She rephrases the question as "but what is it- what is that" with the stress of what, indicating that the first answer was inappropriate because it stressed is--an equivalence rather than a description. The student's answer to this new formulation responds with a description, "it's smaller," which is not the kind of description the tutor is looking for. The tutor reformulates the question once more, this time not in the form which gave rise to the errorful answers. This time she has designed her question to specifically rule out the various interpretations of what is X; she accomplishes this by naming a member of the class from which the right answer should be chosen but which is itself obviously not the right answer (the obvious incorrectness of her choice is critical--otherwise the question would have been heard as asking for the correct answer rather than guiding the student to the right level of analysis). The student finally produces an answer at the right level (although it is not the answer the tutor was getting at).

In order to perform these kinds of corrections (or repairs), the tutor must engage in even more troublesome (for computer systems) behavior:

1. The tutor must be able to interpret the utterance, or set of utterances, requiring correction (or repair).

2. The tutor must be able to recognize that the interpretation displayed by that set of utterances requires repair.

3. The tutor must be able to see what in the preceding discourse could have led to the interpretation requiring correction and how it could have led to it.

4. The tutor must be able to re-articulate the source of trouble such that the new formulation takes into account the path which led to the "errorful" interpretation (and eliminates it as a viable interpretation).

I think it is safe to say that no system currently in use could manage these four processes. In particular, the last two steps--tracing the path which led to the misinterpretation and reformulating the trouble source--seem beyond the capabilities of current systems. These processes involve not an abstract domain of knowledge--such as physics or chemistry--but rather a history of the preceding discourse, the ability to find alternative interpretations, based on another's utterances, of one's own verbal behavior and the ability to re-design one's own utterances to rule out undesired interpretations and guide the hearer towards the desired interpretation. These are extremely sophisticated processes, which demand an extensive familiarity with possible relationships between syntactic/semantic structure and discourse structure--natural language understanding in the very deepest sense of the phrase--and an entirely different kind of cognitive flexibility than computers have displayed up to now. The issue of interpretation, and multiple interpretations, is crucial here, more specifically what kinds of utterances in what kinds of contexts lead to what kinds of interpretations. We need to begin addressing these issues seriously if are to capture the power of the process of correction.

It is extremely important to point out here that as I have formulated the processes above, they are at once cognitive and interactional. That is, while the processes themselves must take place "in the mind," the goals motivating these processes are clearly interactional, and their product--for example, a correcting utterance--is a piece of interaction, fitted to the context of utterance. There is thus no way to formulate cognitive processes, at least of the kind having to do with human discourse, which could be independent of interaction.

4 Conclusion

In addition to the implications for modeling correction and tutor assistance, the work presented here brings a new emphasis to the role of the student in the learning process (although see Woolf, 1984; Oberem, 1987; and Miyake, 1982 for suggestions in roughly the same direction). In particular, it should be clear from this work that, given the opportunity, learners play a critical role in the structure and substance of a tutoring session. It is essential, then, that ICAI systems allow for maximum learner participation, in particular by instituting mechanisms for enabling learners to signal states--confusion, lack of understanding, etc.-- to invite correction when appropriate, and of course to ask questions when necessary. These requirements of course, put pressure on systems to be able to interpret student behavior, including such potentially ambiguous behavior as silence. This kind of interpretation can only be done if the system can create a model of the structure of the dialogue so far.

In keeping with this view towards the student as well as the tutor, it is important to stress a view of tutoring as an accomplishment achieved by both parties, a cooperative endeavor. Any outcome of the tutoring process (e.g., low rate of student error, or more tutor guided in some places and more student guided in others) should be seen and analyzed as produced by both parties.

This work also brings a slightly different perspective to errorful bungling as an inherent part of tutoring. According to most of the literature in AI (e.g. see most of the papers in Sleeman and

Brown, 1982), mistakes are good for students--students should be allowed to play with a problem and bungle along unless they get far off the track. I--and most of the tutors and students in the project--would agree that mistakes and bungling have their place in the learning process. But I think we would want to disagree with the "bop til you drop" philosophy that is common in the literature. Tutors and students work together closely to produce a secure but not overly constraining safety net around the student, so that student errors, and in particular tutor correction of student errors, are kept to a minimum. These two different philosophies of learning may simply represent two stages in the process--at first it may be helpful for the student to work mainly alone, floundering and bungling at will, until they know what kinds of problems they are having trouble with, or where they often get stuck, and can bring these problems in for more guided instruction. I think they may both need to be accomodated by robust systems. But it is important to stress the degree to which the tutor provides assistance, after as little as a few seconds difficulty. In fact, only one tutor-student pair we observed accomodated any great degree of bungling, and this occurred only because the student stated at the beginning of the session (and also at several points during the session) that he liked to "flail" around for a while and the tutor agreed that that was an acceptable strategy for their particular interaction. In other words, tutors will only let a student bungle along if the student has requested "bungling" as a strategy, and, of course, if the tutor agrees to it. Bungling and flailing are not the default values in the tutoring we observed.

(Notice that even though the kind of learning I have observed in this project is guided, it does not constrict the autonomy of the student--far from it. The student is always given the opportunity to repair his or her own errors and to check on his or her own understanding before the tutor steps in. Granted, the opportunity may consist of 0.4 of the tutor withholding correction or initiation of correction, but given the pace of face-to-face interaction, even this apparently slight bit of silence is sufficient to allow both parties to see whether the student is going to perform correction or initiate repair him/herself.)

REFERENCES

Anderson, J., Boyle, C.F., Corbett, A. and M. Lewis. 1987. "Cognitive modelling and intelligent tutoring." MS.

Cole, M. (ed.) 1978. Mind in Society. Cambridge: Harvard University Press.

Fox,: B. 1987. "Interaction as a source of diagnostic information in human-human tutoring." MS.

Grosz, B. 1977. The representation and use of focus in dialogue understanding. SRI Technical Notes, 151.

Hayes, P. 1983. Introduction. International Journal of Man-Machine Studies 19, 229-230.

Hobbs. 1976. A computational approach to discourse analysis. City University of New York.

Miyake, N. 1982. "Constructive interaction." Report 113, Center for Human Information Processing, University of California, San Diego.

Oberem, G. 1987. "ALBERT: a physics problem-solvinng monitor and coach." Paper presented at Third International Artificial Intelligence and Education.

Ochs, E. 1979. "Transcription as theory." In E. Ochs and B. Schieffelin (eds.) Developmental Pragmatics. New York: Academic Press.

Pomerantz, A. 1975. "Second assessments: a study of some features of agreements/disagreements." Ph.D. dissertation, University of California, Irvine.

----------. 1978. "Compliment responses: Notes on the cooperation of multiple constraints." In J. Schenkein (ed.) Studies in the Organization of Conversational Interaction. New York: Academic Press.

Reichman, R. 1981. Plain speaking: a theory and grammar of spontaneous discourse. BBN, report no. 4681.

Sacks, H., Schegloff, E. and G. Jefferson. 1974. A simplest systematics for the organization of turn-taking. Language 50, 696-735.

Schegloff, E., Jefferson, G., and H. Sacks. 1977. The preference for self-correction in the organization of repair in conversation. Language 53, 361-382.

Sleeman, D. and Brown, J.S. (eds.). 1982. Intelligent Tutoring Systems. New York: Academic Press.

Woolf, B. 1984. "Context-dependent planning in a machine tutor." Ph.D. dissertation, University of Massachusettts, Amherst.

I. Notational Conventions in Transcripts

The following notational conventions are used in the transcripts.

//	point at which current utterance is overlapped by the next utterance produced by another speaker.
(0.0)	Numbers enclosed in parentheses indicate length of silence.
Underlining	indicates stressed syllables
:	lengthened syllable
-	Glottal stop cutting off a word
=	indicates a relationship between two utterances in which there is not the usual beat of silence between them.
?	rising intonation
(( ))	non-linguistic action
( )	unintelligible stretch
hh	audible outbreath
'hh	audible inbreath
(hh)	laughter within a word