Distributed Tutorial Strategies

Sandra Katz (katz+@pitt.edu)

Learning Research and Development Center, University of Pittsburgh

3939 O'Hara Street, Pittsburgh, PA   15260 USA

Abstract

In one-on-one tutoring sessions, lessons are sometimes distributed across several interaction episodes. We observed recurring patterns in the way that physics tutors parcel dialogue sub-goals between problem-solving discussions and post-solution, reflective discussions. We call these patterns "distributed tutorial strategies." They attempt to achieve two main functions: generalizing from the current problem and building the student's repertoire of methods for achieving particular problem-solving goals. This paper illustrates several distributed strategies that physics tutors use to achieve these instructional goals and presents a framework for describing the goal structure of distributed tutorial strategies.

Introduction

A number of studies have demonstrated the effectiveness of human one-on-one tutoring and its superiority over traditional classroom-based instruction (e.g., Bloom, 1984; Cohen, Kulik & Kulik, 1982)—often called the "2 sigma effect." In an attempt to explain this effect, several researchers have investigated the tactics that tutors implement (e.g., McArthur, Stasz, & Zmuidzinas, 1990; Merrill, Reiser, Raney, & Trafton, 1992). A vast repertoire of tactics have been identified, including giving didactic explanations, scaffolding, hinting, engaging the student in Socratic-style dialogues or "directed lines of reasoning" (Hume, Michael, Rovick, & Evens, 1996), and various forms of questioning (Graesser & Person, 1994).

Although this research has enhanced our understanding of how tutors handle impasses and other situations that occur at a single point during a problem-solving session, it has little to say about lessons that are distributed across several dialogue episodes. The distributed nature of instructional interaction has been acknowledged by several scholars (e.g., Akhras & Self, 2000), but the goal composition and structure of distributed dialogues has not been investigated. This paper stems from our observation of recurring patterns of tutorial dialogue goals that were distributed between physics problem-solving sessions and post-solution, reflective discussions. We refer to these patterns as "distributed tutorial strategies."
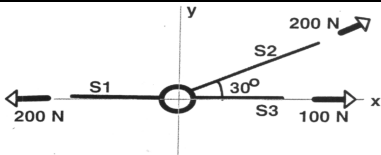
Distributed tutorial strategies consist of a sequence of dialogue sub-goals that collectively attempt to achieve a higher-level instructional goal. Some sub-goals are executed in the initial dialogue episode, while other sub-goals are executed in one or more later episodes. Table 1 contains an example of a distributed tutorial dialogue. (Annotations are in italics.) This dialogue took place via teletyped interaction between a tutor and a student while the latter was working on a problem in an intelligent tutoring system for basic mechanics called Andes (e.g., Gertner & VanLehn, 2000). In the problem-solving excerpt, the student's first equation signals a misconception: that all force problems deal with stationary objects. This misconception is also evident in the student's reply to the tutor's question (problem-solving turn 2: "there's no acceleration"). The overall goal of the distributed dialogue is to resolve this misconception and, correspondingly, to abstract the correct principle (Newton's Second Law) and a schema associated with this principle.[1] During problem solving, the tutor makes the student aware of his error, gets him on track, and lets the solution bring out a contradiction in the student' beliefs about acceleration, but the tutor does not address the misconception directly. During the post-solution discussion, the tutor guides the student in confronting the contradiction brought out by the solution (post-solution turns 1-6) and reifies the principle that explains it—Newton's Second Law—and the schema that maps onto the solution (post-solution turn 7). Because distributed interactions that implement this strategy rely on the solution to contradict the student's reasoning, we dub this strategy *Let the Solution Speak for Itself*.

Previous research has shown that tutorial dialogues that are distributed between problem solving and post-solution reflection, such as the one shown in Table 1, strengthen conceptual understanding and promote near transfer—that is, students' ability to solve problems similar to those discussed during tutoring (Katz, O'Donnell, & Kay, 2000; Katz, Allbritton, & Connelly, in press). Not surprisingly, Katz et al. (in press) found a transfer effect for distributed dialogues whose overall goal was to generalize from the case at hand, or to increase the student's repertoire of methods for achieving particular problem-solving goals. The main motivation for the research discussed in the current paper was to better understand the mechanisms by which distributed dialogues achieve these instructional functions—in particular, to determine whether there are recurring distributed tutorial strategies for abstraction and problem-solving skill-building that can be specified in terms of their sub-goal composition and structure. We identified several such strategies. Following a discussion of the research methodology, we describe and illustrate a framework for specifying distributed tutorial strategies.

---

[1] By "schema," we mean a mental representation of the features that all instances of a type of problem have in common—the concept(s) and principle(s) associated with it, the quantities that are given or need to be derived, and so on.

Table 1. Example of a Distributed Tutorial Dialogue

Andes Problem Statement:	
In the figure at right, each of the three strings exerts a tension force on the ring as marked. Use the labels S1, S2 and S3 to refer to the three strings. Find the components of the net force acting on the ring.	y 200 N S2 S1 30° x 200 N S3 100 N

Problem-Solving Dialogue	**Post-Solution Dialogue**
The student writes the following equation in the Andes equations window: (1) Fnet=0 *The tutor interrupts to flag an error: the student wrongly assumed that there is no acceleration, so he set the net force (Fnet) equal to 0.* 1. **T:** Hold on; How do you know that Fnet = 0? 2. **S:** there's no acceleration 3. **T:** That's not for sure; maybe there _is_ acceleration; look at what the question is asking... *No more discussion takes place during problem solving. The student enters the following main equations (intermediate steps have been deleted), which solve the problem and show that there is acceleration, because there is a net force in both the x and y directions:* (2) Fnet_x= -S1 + S3 + S2cos(30) (3) Fnet_x = -200 + 100 + 200cos(30) = 73.2 (4) Fnet_y = 200sin(30) = 100N	1. **T:** So let me ask you a question; looking at your equations—is there any acceleration? 2. **S:** I guess so, is it because there is a force in the positive x direction? 3. **T:** Not only a force in the positive x-direction...are there any other net forces? 4. **S:** no 5. **T:** really? What is equation 4, then? 6. **S:** oh, the y direction also. 7. **T:** Yes. So you can see that there is a net force in the x-direction, which would result in an acceleration in the x-direction: Fnet_x = m * ax and similarly in the y-direction: Fnet_y = m * ay. This is a very important idea, and is called "Newton's Second Law". Is it familiar to you? 8. **S:** yes. 9. **T:** Good. Are you ready to move on, or do you have any more questions? 10. **S:** no. I am ready to move on. 11. **T:** Ok.

Methods

The distributed dialogues analyzed in this study are a sub-corpus of the corpus used to investigate the roles and effectiveness of post-solution reflective dialogues, as cited previously (Katz et al., in press). The study procedure, corpus, and coding methods are described in detail in that paper.   Here we summarize the methodological features that are most pertinent to the current study.

Fifteen student volunteers taking an introductory physics course at the University of Pittsburgh were randomly assigned to one of seven tutors. Tutors had prior experience teaching physics in a classroom or one-on-one tutoring setting; some had done both.   Student and tutor participants were paid a nominal amount.

The study followed a pretest→intervention→posttest design. The intervention consisted of 24 problems that students worked on in the Andes physics tutoring system. The tutor and student sat in separate rooms and interacted via teletype. Andes' automated coaching was suppressed so that all of the help that students received came from the live tutors. The system automatically logged students' actions and conversations with their tutor. To highlight the roles of post-solution dialogues and their potential impact on student-tutor interaction during problem solving, we presented the problems in one of two formats:

"debrief" and "no debrief." At the start of each problem, the experimenter told participants whether they would be allowed to discuss the problem further after the student solved it (i.e., debrief).   Students worked on twelve problems in each of these within-subject conditions.

The full corpus consists of 315 transcripts, 160 from "debrief" sessions, 155 from "no debrief" sessions. (Due to time constraints and other factors, four students did not complete the 24-problem set.)   We focused on the "debrief" sessions in the previous and current study. The transcripts were segmented into dialogue episodes. Problem-solving episodes typically begin with an initiating query or statement and end with a student action. Post-solution episodes (debriefs) begin at the point where the tutor confirms that the student's answer is correct. We segmented each episode further into sub-dialogues—one sub-dialogue per topic—referred to henceforth simply as "dialogues."

Two raters coded the following dialogue features:

- **Initiator:** who initiated the discussion—the student or the tutor?
- **Information status:** This feature is only coded for post-solution dialogues, but refers to problem-solving discussions. What is the informational relationship between the post-solution dialogue and problem-solving discussions?   Does the post-

Table 2. Abstraction and Skill Repertoire-Building Functions of Distributed Dialogues

Description	Frequency
Schema construction: enable the student to recognize the features that a set of problems have in common, identify the associated schema (e.g., a projectile problem; a static force problem), and see how the current problem maps onto this schema.	13(6,4,1,1,1)
Schema extension: enable the student to see how a familiar schema applies to problems with different surface features than the problems he or she has encountered so far—for example, to realize that the formal representation of Newton's Second Law (Net force = mass * acceleration) can not only be used to solve problems involving stationary objects, but problems with moving objects as well.	5(5)
Conceptual specialization: enable the student to understand the distinction between related concepts—e.g., the difference between instantaneous, average, and constant acceleration.	5(1,1,2,1)
Tactical generalization: encourage and enable the student to use standard "tricks of the trade" for solving quantitative problems, such as breaking complex problems into more manageable sub-goals, reading the problem statement carefully, and expressing relations symbolically before instantiating variables.	11(4,1,3,1,2)
Alternative method—optimal: broaden the student's repertoire of problem-solving skills, by enabling the student to see why a recommended procedure is preferable to the one that he or she applied.	6(4,1,1)
Alternative method—equivalent: broaden the student's repertoire of problem-solving skills, by enabling the student to see why a recommended procedure is equivalent to the one that he or she applied.	2(2)
Total	42

solution dialogue bring a new topic to the table? If it extends a problem-solving discussion about the same topic, does the post-solution dialogue contain new content or summarize the problem-solving dialogue? Based on these distinctions, we code the information status of a post-solution dialogue as "new," "elaboration," or "summary."

Given the current study's focus on identifying the sub-goal structure of distributed dialogues that attempt to achieve abstraction and skill-building functions, the sub-corpus used in this study consists of tutorial lessons with the following features:

- **Tutor-initiation.** Because we are interested in characterizing the strategies used by tutors, problem-solving and post-solution dialogues had to be initiated by the tutor.
- **Distribution:** The lesson contains at least one problem-solving dialogue and one post-solution dialogue, and the latter is a reprise of the former— that is, the information status of the post-solution dialogue was coded as "elaboration" or "summary." Typically there is only one problem-solving dialogue that corresponds to a post-solution reprise, but sometimes there are several, as when a student repeats an error.
- **Abstraction or skill-building as an overall goal.** The problem-solving and post-solution dialogues collectively attempt to achieve one of the functions described in Table 2. Our main interest here is to specify how this inter-dialogue cooperation takes place. We used a similar goal classification scheme in the current study to the one used to describe the roles of post-solution dialogues in Katz et al. (in press). In the few cases where there

was more than one function, the primary function was assigned as the overall goal descriptor.

The sub-corpus consists of 42 distributed dialogues that have these three features. Table 2 shows the frequency of distributed dialogues that attempt to achieve each overall goal. The frequency notation signifies the total number of dialogues in a category and the breakdown by tutor. For example, a frequency of 13(6,4,1,1,1) for schema construction means that there were 13 distributed dialogues of this type, involving five student-tutor dyads; one tutor initiated 6 distributed schema construction dialogues, another tutor initiated 4, and three tutors initiated one each.

Fifteen transcripts were selected at random to test for inter-rater reliability. Thirteen transcripts contained a post-solution dialogue and the two coders' judgments of whether a post-solution discussion occurred were in perfect agreement (100%). Agreement rates were 92% for initiator (kappa = .83), 85% for information status (kappa = .78), and 94% for overall instructional goal (Table 2) (kappa = .77).

Specifying Distributed Strategies

Instructional dialogues whose sub-goals are distributed between problem solving and post-solution reflection can be described in terms of two main features:

- **Sub-goal status.** Is a sub-goal optional or necessary for achieving the overall goal?
- **Sub-goal staging.** Where is each sub-goal executed—during problem solving, during post-solution reflection, or distributed between these two phases?

Table 3: Sub-goal Composition and Structure of Three Goals Described In Table 2

Overall Goal	Problem-solving Dialogue Goals	Post-solution Dialogue Goals
Schema construction	• {*If S has trouble getting started* → S knows which principle(s) to apply to the current problem} • {*If S reaches an impasse in applying a principle* → S knows how to apply principle to current problem} • {*If error signifies a misconception and error correction suffices to get S on track* → S is aware that S made an error, but may not understand the underlying misconception}	• S understands the main features of a problem schema • {S understands how the current solution maps onto this schema}
	• S is aware that the current problem is a member of a class of problems (schema) • {S knows what the schema is called (identification)} • {*If there is evidence of misunderstanding or tutor suspects that S is applying (applied) schema by rote* → S understands the concept(s) and principle(s) associated with the schema}	
Schema extension	• {*If S has trouble getting started* → S knows which principle(s) to apply to the current problem} • {*If S reaches an impasse in applying a principle* → S knows how to apply principle to current problem}	• {S's understanding of how to apply principle(s) to situation represented in the current problem is strengthened} • S understands how to apply the same principle(s) to problem variation(s)
	• {*If there is evidence of misunderstanding or tutor suspects that S is applying (applied) schema by rote* → S understands the concept(s) and principle(s) associated with original problem and variation(s)}	
Alternative method— optimal	• S understands tutor-recommended procedure • S is convinced that S should replace sub-optimal procedure with optimal procedure	• {S's understanding of the superiority of recommended procedure is strengthened}
	• S understands that original procedure is appropriate, though sub-optimal • S understands why suggested procedure is optimal • {*If the two procedures are conceptually related, or one can be derived from the other* → S understands relationship between original and recommended procedures}	

Table 3 presents a formal specification of three of the instructional goals described in Table 2, in terms of these features. The sub-goal composition and structure of these goals was determined by analyzing matching cases in the corpus (Table 2). We refer to the resulting specification as an *instructional goal frame* (referred to henceforth as "goal frame"). Below we illustrate the approach with reference to the goal frame for schema construction shown in Table 3 and the sample distributed schema-construction dialogue shown in Table 1.

As shown in Table 3, dialogue sub-goals can be stated in terms of their perlocutionary effect (Austin, 1962)—the intended effect on the hearer (in this case, the student)—to the extent that this can be inferred from the content of the tutor's utterances and manner of presentation. In tutorial discourse, the tutor typically tries to alter the student's mental state (e.g., increase understanding about a concept, convince the student that a recommended method is preferable to the one that the student has implemented) and/or actions (e.g., enable the student to apply a principle). Student is abbreviated as S in Table 3.

Optional dialogue sub-goals are in brackets in Table 3. For example, the only sub-goals that all schema-construction dialogues share in common are making the student aware that the problem represents a schema and

schema abstraction—that is, enabling the student to understand the core features of the schema (e.g., its slots and associated principle(s)). Schema awareness and abstraction are achieved in post-solution turn 7 of the Table 1 dialogue. Abstraction is typically supported by other optional sub-goals, such as identifying the schema and reifying how the current solution maps onto it. The latter emerges through post-solution turns 1-7 of Table 1.

Sub-goals are optional when the student does not make an error associated with the sub-goal, in response to tutor preferences, and when other conditions apply. To the extent that these conditions can be inferred from representative dialogues, optional sub-goals can be specified as productions (if-then rules). In Table 3, the conditional part of optional sub-goals is shown in italics. For example, in the schema construction goal frame, the third problem-solving sub-goal states that the tutor can defer addressing a misconception if flagging the error suffices to get the student on track. This is precisely what happened in problem-solving turns 1-3 of Table 1.

With respect to staging, some sub-goals must be addressed during problem solving. Other sub-goals must take place during post-solution reflection. Still other sub-goals are "floating"—they can occur during either or both phases. "Floating goals" are shown spanning the

problem-solving and post-solution sub-goal columns in Table 3. For example, as part of schema construction, the tutor will sometimes identify the schema at the start of a problem-solving session, possibly to get the student on course, as when one tutor said, "Hint: Newton's second law, only this time you are concerned with acceleration in x and acceleration in y." At other times, schema identification takes place during post-solution reflection—for example, with respect to a different problem than the one shown in Table 1, a tutor asked, "Did you see that this is work done by a variable force?"

A distributed tutorial strategy is one instantiation of the goal frame for an overall goal. It consists of a subset of the dialogue sub-goals in the goal frame, arranged in a particular order. As discussed previously, the distributed dialogue in Table 1 implements the *Let the Solution Speak for Itself* strategy, which instantiates the goal frame for schema construction. It consists of the sub-goal structure specified in Table 4. Optional sub-goals in the goal frame can be required by the strategy definition. For example, the *Let the Solution Speak for Itself* strategy only takes place when the student has a misconception about the core principle associated with a schema. Error-flagging is an essential (not optional) part of the problem-solving discussion and reification of misconceived concept(s) and principle(s) is an essential part of the post-solution discussion.

Table 4: The *Let the Solution Speak for Itself* Strategy

> **Problem-solving Dialogue Sub-goals:**
> - S is aware that S made an error, but may not understand the underlying misconception
> - {S knows which principle(s) to apply to the current problem}
> - {S knows how to apply principles to current problem}
>
> **Post-solution Dialogue Sub-goals:**
> - S is aware that the current problem is a member of a class of problems (schema)
> - S understands the main features of the schema
> - {S understands how the current solution maps onto this schema}
> - S understands the concept(s) and principle(s) associated with the schema.

There are other ways that schema construction can take place. The most common strategy—implemented in 8 of the 13 schema construction cases—could be called the *Advanced Organizer* strategy, because the tutor identifies the schema or the central principle associated with it at the start of the problem—for example, "I'll get you started…Conservation of energy." The tutor then helps the student, as needed, with applying the relevant concept(s) or principle(s). During post-solution reflection, the tutor reifies the schema and how the current solution maps onto it.

The goal frame for schema extension provides another example of distributed tutoring. Some tutors appear to be aware of the limited time that a typical physics course affords for exposing students to the wide range of physical situations that a given schema can be applied to. Schema extension attempts to address this problem. It typically happens through the *What if…* distributed strategy specified in Table 5—named as such because the tutor poses one (or more) problem scenarios to the student during post-solution reflection which alter the physical situation in the current problem.

Table 5: The *What if…* Strategy

> **Problem-solving Dialogue Sub-goals:**
> - {S knows which principle(s) to apply to the current problem}
> - {S knows how to apply principle(s) to the current problem}
> - {S understands principle(s) associated with current problem}
>
> **Post-solution Dialogue Sub-goals:**
> - {S's understanding of how to apply principle(s) to situation represented in the current problem is strengthened}
> - S understands how to apply the same principle(s) to problem variation(s)
> - {S understands concept(s) and principle(s) associated with original problem and variation(s)}

The primary goal of the problem-solving discussion (if any) is to strengthen the student's ability to apply the schema and its associated principle(s) to the current problem and similar ones that the student has encountered. This is necessary before moving on to "new territory." The problem-solving dialogue that preceded the sample post-solution excerpt shown below centered around helping the student solve a typical static Newton's Second Law problem (to find the tension in a rope holding a suspended bungee jumper). The "what if" scenario posed during post-solution reflection challenges the student to apply the schema to a dynamic situation (T = tutor):

> **T:** What if I now told you that this guy is going to rappel down the cliff and the maximum tension in the rope is 450 N. How could you find the minimum acceleration he would need to not break the rope?

The post-solution discussion focused on enabling the student to see that the formalism for Newton's Second Law (Net force = mass * acceleration) is the core of the schema that maps onto both situations:

> **T:** you need to think of these problems as the same in the sense that when you are dealing with problems where several forces are acting the first equation that should enter your mind is Newton's 2nd law. Think of it as a recipe because although the ingredients may change the base is always that: F_net = m * a.

The sub-goal structure of a distributed strategy is fairly stable. However, the particular tactics that tutors use to implement each sub-goal can vary greatly. For example, in the post-solution dialogue shown in Table 1, the tutor takes an indirect approach to enabling the student to realize the contradiction brought out by his solution. Post-solution turns 1-6 illustrate a Socratic-style tactic called a "directed line of reasoning" (Hume et al., 1996). Alternatively, the tutor could have stated the contradiction didactically, in the same manner as he summarized Newton's Second Law and its general, symbolic representation in post-solution turn 7.

Conclusion

According to constructivist views of learning, learning takes place through interactions between individuals and their environment (e.g., Newman, Griffin & Cole, 1989). Since interactions evolve over time, there has been increasing interest in understanding the relationships between instructional interactions, e.g.:

> As a time-extended process, learning depends on the relations that develop over time between aspects of single interactions in situations. Therefore, the role of a theory of time-extended processes of interaction is to formalize the various ways that interactions relate to one another over time in a course of interaction… (Akhras & Self, 2000, p. 349)

In this paper, we proposed an approach to formally specifying the intentional relationships that hold between problem-solving and post-solution dialogues that collectively attempt to achieve higher-level instructional functions. At times, dialogue sub-goals are aligned in patterns that we call "distributed tutorial strategies."

Because tutors are rational beings, it is tempting to infer that they consciously decide, during problem solving, to parcel a lesson between problem solving and post-solution reflection. There is some evidence of deliberate distribution in this corpus. For example, during a problem-solving discussion, a tutor presaged a post-solution lesson as follows: "ok…let's do it this way...solve for the tension first and then I will ask you a variation to show that the way you are thinking of it is not good." However, we believe that in most cases, reprises in post-solution discussions are not pre-planned; they are simply the result of the tutor seizing an opportunity to extend a previous discussion. Further research is needed to determine the degree to which distributed interactions stem from deliberate planning, and to uncover the constraints that determine which strategies tutors use in particular contexts.

As Chi, Siler, Jeong, Yamauchi, and Hausmann (2001) have noted, descriptive studies of human tutoring need to be followed up by outcome-based research—that is, studies of the effectiveness of identified strategies and tactics. Along these lines, further research is needed to determine which of the distributed strategies described in this paper—and others we identified—support conceptual understanding and transfer, and whether similar effects could be achieved through non-distributed interaction. This research would provide valuable guidance for tutor training and the design of effective automated tutors.

References

Akhras, F. N., & Self, J. A. (2000). System intelligence in constructivist learning. *International Journal of Artificial Intelligence in Education, 11,* 344-376.

Austin, J. L. (1962). *How to do things with words.* Cambridge: Harvard University Press.

Bloom, B. S. (1984). The 2 sigma problem: the search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13,* 4-16.

Chi, M. T. H., Siler, S. A., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from human tutoring. *Cognitive Science, 25,* 471-533.

Cohen, P. A., Kulik, J. A., & Kulik, C. L. C. (1982). Educational outcomes of tutoring: a meta-analysis of findings. *American Educational Research Journal, 19,* 237-248.

Gertner, A., & VanLehn, K. (2000). Andes: A coached problem solving environment for physics. In Gauthier, G., Frasson, C. & VanLehn, K. (Eds.), *Intelligent Tutoring Systems: 5th International Conference, ITS 2000.* Berlin: Springer.

Graesser, A. C., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal*, *31*, 104-137.

Hume, G. D., Michael, J., Rovick, A., & Evens, M. (1996). Hinting as a tactic in one-on-one tutoring. *Journal of the Learning Sciences, 5(1),* 23-47.

Katz, S., O'Donnell, G., & Kay, H. (2000). An approach to analyzing the role and structure of reflective dialogue. *International Journal of Artificial Intelligence and Education, 11*, 320-343.

Katz, S., Allbritton, D, & Connelly, J. (in press). Going beyond the problem given: How human tutors use post-solution discussions to support transfer. *International Journal of Artificial Intelligence and Education,13.*

McArthur, D., Stasz, C., & Zmuidzinas, M. (1990). Tutoring techniques in algebra. *Cognition and Instruction, 7(3),* 197-244.

Merrill, D. C., Reiser, B. J., Raney, M., & Trafton, J. G. (1992). Effective tutoring techniques: a comparison of human tutors and intelligent tutoring systems. *The Journal of the Learning Sciences, 2(3),* 277-306.

Newman, D., Griffin, P., & Cole, M. (1989). *The construction zone: Working for cognitive change in school.* Cambridge, MA: Cambridge University Press.