

The Construction of Shared Knowledge in Collaborative Problem Solving

Jeremy Roschelle¹ and Stephanie D. Teasley²

¹ Division of Math, Science, and Technology, School of Education, University of California, Berkeley, USA

² Learning Research and Development Center, University of Pittsburgh, USA

Abstract: This paper focuses on the processes involved in collaboration using a microanalysis of one dyad's work with a computer-based environment (the Envisioning Machine). The interaction between participants is analysed with respect to a 'Joint Problem Space', which comprises an emergent, socially-negotiated set of knowledge elements, such as goals, problem state descriptions and problem solving actions. Our analysis shows how this shared conceptual space is constructed through the external mediational framework of shared language, situation and activity. This approach has particular implications for understanding how the benefits of collaboration are realised and serves to clarify the possible roles of the computers in supporting collaborative learning.

Keywords: problem solving, physics learning, collaboration, computer support

1. Introduction

Collaborative work between peers provides a particularly rich environment for studying learning. The social situation maintains student motivation and naturally elicits verbal communication. Furthermore, several prominent theorists [e.g., 4, 8, 9, 19] have argued that learning is fundamentally a social activity. Like several other researchers in this volume, we have begun to investigate systematically the social and cognitive dimensions of collaborative problem solving. Specifically, we have been concerned with collaborative problem solving involving a computer simulation of concepts in physics. This work differs from other chapters in this volume as it focuses specifically on the *process* of collaboration by using a microanalysis of one dyad's work. We believe that a focus on process is necessary in order to understand the value of learning in collaboration with peers.

Studies of children's collaborations have found that successful collaboration involves a large degree of mutual engagement, joint decision making, and discussion [10]. While children's behaviour during collaborative activity has been generally described, most studies focus on outcome rather than process. That is, numerous studies have documented that collaboration can benefit learning and others have described the conditions for successful collaboration (see [1, 18] for reviews). However, little is known about how the collaboration itself actually proceeds.

Before we can begin to analyse the process of collaboration, it is useful to be specific about which phenomena we seek to understand. As the diverse chapters in this book indicate, 'collaboration' is a broadly used term which serves to describe a wide variety of behaviours. In the most general sense, collaboration is said to have occurred when more than one person works on a single task. For our purposes, however, it is helpful and in fact necessary, to draw some specific parameters around what we refer to as collaboration. The following definition delineates the kind of behaviour we have focused on:

Collaboration is a coordinated, synchronous activity that is the result of a continued attempt to construct and maintain a shared conception of a problem.

We make a distinction between 'collaborative' versus 'cooperative' problem solving. Cooperative work is accomplished by the division of labour among participants, as an activity where each person is responsible for a portion of the problem solving. We focus on collaboration as the mutual engagement of participants in a coordinated effort to solve the problem together. We further distinguish between synchronous and asynchronous activity. Although we do not propose that collaboration cannot occur in asynchronous activity, we focus on face-to-face interactions, which can only occur as a synchronous activity.

The notion of 'a shared conception of the problem' is central to our work. We propose that social interactions in the context of problem solving activity occur in relation to a *Joint Problem Space* (JPS). The JPS is a shared knowledge structure that supports problem solving activity by integrating:

- (a) goals;
- (b) descriptions of the current problem state;
- (c) awareness of available problem solving actions;
- (d) associations that relate goals, features of the current problem state, and available actions.

As the microanalysis described below will make clear, we propose that the fundamental activity in collaborative problem solving occurs via engagement with an emergent, socially-negotiated set of knowledge elements that constitute a Joint Problem Space.

Thus, in contrast to traditional cognitive psychology, we argue that collaborative problem solving takes place in a negotiated and shared conceptual space, constructed through the external mediational framework of shared language, situation, and activity — not merely

inside the cognitive contents of each individual's head. This orientation leads to consequences for the design and analysis of computer-supported collaborative learning situations. We see potential applicability of our analysis of the microstructure of collaboration as a methodology for understanding how the benefits of collaboration are realised. Moreover, we see an emphasis on mediational resources as a possible clarification of the role of computer support in collaborative learning.

We present a detailed analysis of collaborative problem solving focusing on data from one of the most collaborative dyads we have studied. The chapter will begin with a description of the computer task the students used and a description of the physics knowledge involved. The next section will be a brief introduction to the framework we use for analysing the collaborative process. The bulk of the chapter will explicate our analysis of the dyads' interactions, focusing on their discourse and activity as they work on one of the problems presented in the microworld. We will focus on the data from this dyad, subjects 'Gary' and 'Sam'. (Names have been changed.)

2. The Task: The Envisioning Machine

Gary and Sam were 15 year old males who were taking a summer course in statistics at the University of California, Berkeley. They were comfortable working together as they had been collaborating on a computer project in the statistics course. They did not have any formal physics training. (Gary, however, had done some reading about physics on his own.) The subjects were asked to work together on an activity involving a computer simulation called "The Envisioning Machine" [11, 12].

The Envisioning Machine (EM) is a 'direct manipulation graphical simulation of the concepts of velocity and acceleration. Figure 1 illustrates the screen of the EM. There are two windows, the "Observable World" and the "Newtonian World". The Observable World displays a simulation of a ball moving across the screen. This represents the goal motion. The Newtonian World displays a particle with velocity and acceleration vectors (the thin and thick arrows, respectively). Using the mouse, the user can manipulate the settings of these vectors. When the simulation is run, the particle in Newtonian World moves with the initial velocity indicated by the velocity vector and the acceleration indicated by the acceleration vector. In both worlds, the moving objects leave a trace of dots behind them as they move. Because the dots are dropped at a uniform time interval, the dot spacing represents speed. All the motions displayed by the EM are constant velocity or constant acceleration motions.




Figure 1. The Envisioning Machine

The specific EM activity used in this study involved matching the goal motion displayed in the Observable World by adjusting velocity and acceleration vectors on the particle displayed in the Newtonian World. This activity was called a 'challenge'. Typically, solving a challenge requires a series of trials in which the students watch the motions in the Observable and Newtonian Worlds, adjust the vectors of a particle in the Newtonian World, run the simulation, and evaluate whether the two motions were the same. Since the students had not previously studied velocity and acceleration, they needed to experiment with the simulation in

order to learn how to adjust the vectors to produce motions that matched motions in the Observable World. Moreover, since the computer did not give explicit feedback on the correctness of a solution, students needed to develop their own criteria for determining whether two motions were 'the same'.

The subjects worked on the EM activity in three sessions, each about 45 minutes long. We will discuss only one challenge from the first session of Gary and Sam's work. The session had the following format: in the beginning of this session, Gary and Sam were instructed on how to use the mouse to do the EM activity. During these instructions, the vectors were given neutral names, "the thin arrow" and "the thick arrow", rather than the more theory-laden terms 'velocity' and 'acceleration'. The task was described as "making the motions the same", though the meaning of "the same" was not specified. Thus the instructions left the meaning of the task substantially underdetermined. After the instructions, Gary and Sam were asked to "work together" on a series of ten challenges. Each challenge consisted of matching a different Observable World motion by adjusting the arrows in the Newtonian World. When the subjects finished the challenges, about 45 minutes later, they were interviewed about what they had learned.

2.1. What is Envisioning Machine Knowledge?

Since our goal is to examine the construction of *shared* knowledge in collaborative problem solving, it is necessary to discuss the nature of the physics knowledge involved in solving the EM activity. Roschelle [11] presents a competence model of EM problem solving and shows that it accounts for key aspects of students' problem solving performances. In this model, the EM activity is seen as a form of difference-reduction: students try to reduce the differences between the motion they control in the Newtonian World and the goal motion in the Observable World. This difference-reduction takes place in two stages: first, students set the directions of the vectors to match the overall shape of a motion. Second, students set the lengths of the vectors to match the speed at which the particle moves along the shape. The types of knowledge corresponding to these two stages are knowledge of *configurations* and knowledge of *qualitative proportionalities*.

Configurations relate the direction of the vectors to the shape of the motion produced. The velocity vector always points in the direction with which the motion begins. Depending on the angle between the acceleration vector and the velocity vector, motions with qualitatively different characteristics are produced. For example, when the velocity and acceleration vectors are colinear and opposed, the motion will go out and come back along a straight line. The Envisioning Machine motions can be categorised into four shapes with four corresponding configurations, as in Table 1 below.

Shape	Direction of vectors
Straight, constant speed	No acceleration
Straight, speeding up	Acceleration in same direction as velocity
Straight, turns around	Acceleration in opposite direction to velocity
Curve	Acceleration at an angle to (not co-linear with) velocity

Table 1. Configurations of Envisioning Machine motions

Within each configuration, it is still necessary to determine the proper length of each arrow. Clues for the correct lengths can be found by comparing the Newtonian and Observable World motions. For example, if the dot spacing in the Newtonian World is greater than the dot spacing in the Observable World, then the velocity vector in the Newtonian World is too long. We call relationships of this form 'qualitative proportionalities', after similar representations developed by computer scientists investigating qualitative reasoning (see [2]). A qualitative proportionality is a relationship between two variables that states that a increase in one variable will result in an increase in the other. The relationship between dot spacing and length of the velocity vector could be stated as "the dot spacing is qualitatively proportional to the length of the velocity vector". Table 2 lists some qualitative proportionalities that students use to solve the EM task.

Property of motion	Length of vector
Initial speed, dot spacing	Proportional to velocity
Time to reach apex	Proportional to velocity and inversely proportional to acceleration
Height of apex	Proportional to velocity and inversely proportional to acceleration

Table 2. Some qualitative proportionalities

Roschelle (in preparation) argues that EM knowledge, as described above, is a valuable form of physics knowledge. Although it is outside scope of this chapter to argue for this view, the line of reasoning is as follows: EM knowledge, as described above, encodes qualitative regularities in the behaviour of the EM. The EM's behaviour, in turn, is based on the mathe-

matical definitions of velocity and acceleration⁵. Thus, the qualitative regularities of the EM are also qualitative regularities in the concepts under their formal definitions. Learning the EM is therefore also learning specific qualitative descriptions of velocity and acceleration. Such qualitative descriptions are particularly important for understanding how physics laws apply to everyday situations.

3. Framework for Analysing Collaboration

An examination of students' discourse and activity as they work together allows us to understand how the social interaction affects the course of learning. This necessitates a microanalysis of not only of the content of students' talk, but also of how the pragmatic structure of the conversations can result in shared knowledge. In particular, it requires understanding how students use coordinated language and action to establish shared knowledge, to recognise any divergences from shared knowledge as they arise, and to rectify misunderstandings that impede joint work. To accomplish this aim, we draw on ideas from pragmatics [e.g., 7], conversation analysis [e.g., 14], and protocol analysis [5] to describe how the communicative exchanges function to construct and maintain a Joint Problem Space. In coordination with an analysis of the development of students' physics knowledge, we are able to identify how social interaction promotes or inhibits learning in key segments of the problem solving process.

Recent work on the coordination of meaning in conversations has stressed that mutual intelligibility is the result of local, interactional work of the participants. Conversants establish shared meaning via the construction and accumulation of a common ground, a body of shared knowledge [3]. Meaning can be coordinated and mutual intelligibility achieved because conversants provide constant evidence, positive and negative, that each utterance has been understood, and engage in repairs when it has not [13].

In our analysis of collaborative learning, we take the point of view that students' work is based on a shared conception of the task. We enlarge the notion of common ground, which has origins in the study of ad hoc conversations, and apply it to the study of a socially-organised task-oriented activity: collaborative problem solving. In doing so, we synthesise the construct of common ground with a cognitive analysis of problem solving activity. Specifically, we hold that collaborative problem solving consists of two concurrent activities: solving the problem together and building a JPS. These activities necessarily co-exist. Conversation in the context of problem solving activity is the process by which collaborators

⁵ The definitions read, "Velocity is the derivative of position with respect to time and acceleration is the derivative of velocity with respect to time".

construct and maintain a JPS. Simultaneously, the JPS is the structure that enables meaningful conversation about problem solving to occur.

The JPS is a pragmatic, rather than an ideal structure. The overlap of meaning in the collaborator's common conception of the problem is not necessarily complete or absolutely certain. Rather this overlap is sufficient to gradually accumulate shared concepts and allow convergence on certainty of meaning.

Thus, to build a JPS, collaborators must have ways of:

- introducing and accepting knowledge into the JPS;
- monitoring on-going activity for evidence of divergences in meaning;
- repairing divergences that impede the progress of the collaboration.

There are a number of structured discourse forms that conversants use in everyday speech to achieve similar goals in the service of mutual intelligibility. These forms utilise language, physical action, and combinations of words and actions. Our analysis will show that students can use the structure of conversation to continually build, monitor and repair a JPS. Below we discuss some of the categories of discourse events that have proved useful for our analysis. A complete review of discourse analysis is outside the scope of the paper (see [7] for a review). Of the categories we discuss, turn-taking is the most pervasive and general. Specific turn-taking forms contribute to various aspects of joint problem solving activity. Socially distributed productions (SDPs) provide means for introducing and accepting problem solving knowledge into the JPS. Narrations and question-answer pairs enable students to monitor each other's interpretations. Repairs offer a means to rectify divergent interpretations. Coordinations of language and action also prove important for introducing, monitoring, and repairing knowledge in the JPS.

3.1. Turn-Taking

Communication between individuals follows a well-specified form of turn-taking that has been extensively described by linguists and sociologists [15]. Discourse units such as questions, acceptances, disagreements, and repairs represent various specific discourse forms available for taking a conversational turn. The flow, content, and structure of turns is used as a measure of whether the participants in a conversation understand each other [3]. Similarly, in our analysis of student's collaborations, we propose that the structure of turn-taking sequences is an indication of the degree to which students share common problem representations. In analysing collaborative work, we look for dialogues in which turn transitions are smooth, and the sequence of talk follows a cooperative pattern. In periods of successful collaborative activity, students' conversational turns build upon each other and the content contributes to the joint problem solving activity.

In addition to joint work, collaborative problem solving includes periods in which partners are not fully engaged with each other. Partners occasionally withdraw from the active interaction with their partner to work on ideas that are too ill-formed or complicated to be introduced into the shared work. These periods are marked in the interaction by periods of significant next-turn deviations such as non-acceptances, disagreements, and empty turns. In a successful collaboration, such periods of withdrawal are usually followed by periods of intense interaction which serve to incorporate the individual insight into the shared problem solving knowledge.

3.2. Socially Distributed Productions

One type of turn-taking structure particularly useful in understanding the production of shared problem solving knowledge is the 'collaborative completion'. As described in the work of Lerner [6] and Wilkes-Gibbs [20], a collaborative completion distributes a compound sentence over discourse partners. That is, one partner's turn begins a sentence or an idea, and the other partner uses their next turn to complete it. One especially relevant type of compound sentence has IF-THEN form. In an IF-THEN collaborative completion, the antecedent and consequent are produced on separate turns. The distribution of the IF-THEN across turns provides an opportunity for partners to accept or repair conditional knowledge.

We call an IF-THEN collaborative completion a 'socially distributed production' because its content consists of a production rule, while its form is socially distributed across turns. We will also include in this category IF-THEN sentences that are delivered in instalments, with the conversational partner producing acceptances in subsequent turns. An SDP may be a particularly effective means for constructing shared knowledge because it spreads the interrelated goals, features, and actions of a knowledge element across conversational turns. This provides multiple opportunities for partners to contribute to the construction and verification of the new piece of shared knowledge.

3.3. Repairs

Since the collaboration process involves periods of individual activity, collaborative activity also produces periods of conflict in which individual ideas are negotiated with respect to the shared work. These periods of conflict usually signify a breakdown in mutual intelligibility, rather than the collaboration per se. In fact, the attempts to reduce conflict by resolving misunderstandings are evidence of the dyad's preference for a working style in which a shared conception of the problem is maintained. Often these attempts take the form of 'repairs'.

Repairs are the method by which participants in talk can deal with problems or troubles in speaking, hearing, or comprehension of dialogue [15]. According to Schegloff [13], repairs are a major means for the achievement and consolidation of understanding and thereby the management of the mutual intelligibility of collaborative problem solving activity.

Without successful repairs, breakdowns in mutual intelligibility continue longer. Both partners use justifications, counter-suggestions, assertions and elaborations in their attempt to get their partner coordinated. Occasionally, failures to re-establish mutual intelligibility (unsuccessful repairs) lead to the students abandoning the current problem. This can be seen when partners give up on a particular challenge or give up a particular aspect of the challenge. In the course of the session, students may return to the particular challenge or problem area, and may resolve the impasse in the shared understanding or continue by working around the impasse.

3.4. Narrations

Narrations are a verbal strategy that enable partners to monitor each other's actions and interpretations. In the EM activity, only one partner can carry out actions with the mouse at a time. These actions may be difficult for the other partner to interpret, because every action can correspond to a number of possible intentions. Narration informs one's partner of the intentions corresponding to actions. This enhances the partner's opportunities to recognise differences in the shared understanding. Continued attention to narrations and accompanying action can signal acceptances and shared understandings [3]. Interruptions to narrations create an immediate opportunity to rectify misunderstandings. Narrations are also useful for the participants to signal that an action is not intended to contribute to the current shared goal; a statement like "I just want to see what this does" signals that the actor is no longer working on the task at hand, but rather is exploring a novel situation.

3.5. Language and Action

Although there are many examples of narratives in collaborative activity, students are not wholly dependent on language to maintain shared understanding. In fact, one major role of the computer in supporting collaborative learning is providing a context for the production of action and gesture. Action and gesture can both serve as presentations and acceptances. An action or gesture can serve as an acceptance when one partner interprets the other partner's utterances by performing an action. Since most of the utterances contain indexical, ambiguous references, the production of the appropriate action both accepts and confirms a shared

understanding of the task. Actions and gestures can likewise serve as presentations of new ideas. Partners often use their hand or the computer mouse to demonstrate an idea. In this case, the partner's ability to interpret the action successfully through an utterance is an indication of mutual intelligibility and acceptability of the idea. The simultaneous production of matching language and action by separate partners can also produce an effective division of labour: while one partner concentrates on carrying out actions, the other concentrates on producing utterances that make the intentions behind the actions available for commentary and repair.

4. Challenge Six

Our goal in the remainder of this chapter is to exemplify the analysis of the process of collaborative problem solving from the point of view of the JPS. Our approach will be to look in detail at the Gary and Sam's construction of a shared conception of the task in 'challenge six', which was the sixth motion that they worked on. These students began challenge six about eleven minutes into their session and finished it about eight minutes later.




Figure 2. The Goal Motion in Challenge 6

In challenge six, the motion of the ball in the Observable World is analogous to that of a ball tossed straight up in the air; it starts upward, slows down, instantaneously pauses at the top, then accelerates downward (Figure 2). To construct this motion in the Newtonian World, subjects must set the direction of the velocity vector upward and the acceleration vector downward. In addition, to copy the Observable World motion exactly, the subjects must adjust the lengths of both vectors appropriately. Two earlier challenges which these students

investigated also required an acceleration vector, but challenge six is the first challenge that they have seen in which the acceleration opposes the velocity.

The detailed analysis that follows shows that the students progressed through challenge six in two main stages. First, they established the correct directions for the Newtonian World vectors. Second, they determined the correct lengths for the vectors. The settings of velocity and acceleration at key moments during the challenge are illustrated in Figure 3. The setting in Figure 3b, which the participants achieved at the end of episode 6-2, shows correct setting of the directions of the vectors, in contrast to 3a. Later, the participants adjusted the lengths of the vectors for a better match to the Observable World motion. Figures 3c and 3d, respectively, illustrate a close approximation and the exact setting of the lengths.

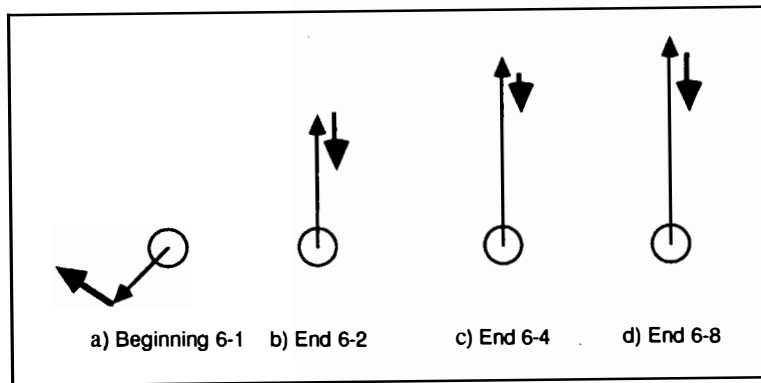


Figure 3. Settings of the vectors during key episodes

We focus on challenge six for several reasons:

1. Challenge six began with strong individual contributions and ended with strong shared contributions. In setting first the directions and later the lengths of velocity and acceleration, the participants began with different, although not necessarily incompatible ideas. By the end of the challenge, their ideas converged. Tracing the course of this interaction enables us to examine processes that the participants use to resolve impasses in shared knowledge.
2. Important aspects of the concepts of velocity and acceleration were learned during this challenge. In particular, the students learned a new configuration of vectors and three new qualitative proportionalities. In addition, the students increased their descriptive capability to include the instantaneous speed of a motion. This was a feature that they had not previously described or used. These accomplishments enable us to investigate processes that students use to construct new elements of shared knowledge.

3. The collaboration during this challenge was particularly successful, both from our point of view as analysts and from the students' own point of view. As analysts, we consider the collaboration successful not just because the participants solved the problem of making the motions the same, but also because they constructed shared knowledge which (a) built effectively on previous shared knowledge, (b) persisted into later challenges and (c) was explained convincingly by both participants in the interview that followed the session. It was successful from the students' point of view as evidenced by their comments both at the end of the challenge and during the interview.

1	G: Challenge six.	<i>S reaches for mouse. G gives it up. S now has mouse for remainder of challenge.</i>
2	S: OK.	<i>Runs simulation. Observable World shows challenge 6 — ball toss, straight up. Newtonian World shows random curve.</i>
3	G: This one's gonna be a curve. No maybe not. (1.5) Oh::: ((falling intonation)) (2.0)	
4	S: It's the acceleration is in the opposite direction to start with. [<i>Stop simulation.</i>
5	G: How can they make a, (,) a double acceleration?	
6	S: See it for a second.	
7	G: Oh they make it go a::ht	<i>Motions up then down with finger tip.</i>
8	S: (()) the first one. [<i>Gestures up.</i>
9	G: Let's see that again. Let's see that again. (6.0)	<i>S resets, runs, NW still random curve.</i>

Table 3. Episode 6-1: In the opening moments of the challenge, the subjects watched motion in the observable world. Both partners simultaneously tried to make sense of a kind of motion they had not yet encountered during any of the previous challenges in the session. Although both partners were engaged, their discourse signalled that they did not yet share the same conception of the challenge.

The transcript⁶ of language and action during challenge six will be presented in a series of 'episodes'. The boundaries for each episode were chosen to be consistent with events in the collaboration, although their exact size was determined for ease of exposition. The contents of the episodes are roughly as follows: in episodes 6-1 and 6-2, the students constructed knowledge of the configuration of vectors required to produce the shape of the challenge six motion. In the remaining episodes, they focused on setting the lengths to match the motion more closely. The problem of setting the lengths began with a considerable difference of opinions in episode 6-3, which was resolved in episode 6-4. In contrast, episodes 6-5 and 6-7 show relatively smooth elaborations of ideas. Episodes 6-4, 6-6, and 6-8 are also interesting because of the new ideas introduced there. In episode 6-8, the participants negotiate a closure to the problem solving activity. The analysis of these episodes focuses on the means by which collaborators introduce and accept ideas into the JPS, monitor emerging interpretations, and maintain the JPS by repairing divergences in understanding.

At the beginning of this challenge (see Table 3), Gary predicted what kind of motion he thought the ball would take (line 3). Upon actually seeing the model run, Sam correctly identified the correct relationship between the directions of the velocity and acceleration vectors: that they should be opposite (line 4). Although his statement is somewhat ambiguous (he could have the direction of the velocity wrong, but the relationship between velocity and acceleration right), his later actions clarified his intent as he set the velocity pointing straight up and the acceleration straight down.

Gary's next utterance occurs as an overlap of Sam's. By this interruption, he did not directly accept Sam's idea that the vectors should be opposed, rather he stated his own conception of the problem (line 5). His question about "double acceleration" suggests that he had a different idea. Sam's statement (line 6), instead of answering Gary's question, directed Gary's attention back to the computer simulation. By re-stating Sam's utterance, Gary agreed to watch the simulation a second time (line 9).

After they watched the simulation again (see Table 4), Gary said "Oh, ok" (line 10) which Sam interpreted as an acceptance of his original idea that the velocity and acceleration should be set in opposition. Sam then began using the mouse to carry out this idea. Although Gary was not contributing to the discourse during this time, he was watching Sam's actions on the screen. Sam's narration as he worked (lines 11-12) allowed Gary the opportunity to comment at any point should he have disagreed. At line 19, after Sam had run the simulation with the arrows opposite, he announced "we almost got it". This utterance reflected Sam's general satisfaction with the shape of the motion. Gary's next utterance (line 20) treated Sam's statement as an invitation to refine the standard for success; he suggested further changes that fit within the framework of setting the arrows opposite via the terms "initial speed" and

10	G: Oh ok. (2.0)	<i>Points to screen and gestures up.</i>
11	S: Yeah, I know what they're doing. Ok. First,	<i>Resets.</i>
12	G: just need to get everything. (2.5)	
13	S: Right about there: =	<i>Sets initial position bottom centre, the correct location.</i>
14	G: =That looks good. (2.0)	
15	S: Ok, so we wanna be cool. (3.0) Just for now is (2.0)	
16	S: Oh. (4.0)	<i>Points velocity vector upwards.</i>
17	S: Ok. (2.0)	<i>Sets acceleration vector downwards.</i>
18	S: So: that one's pretty long. And this one right here is going (.) ba:ck (.5) like that. (3.5)	<i>Makes velocity shorter. Runs simulation, NW goes straight-up/down but speeds don't match.</i>
19	S: Ooh we almost got it=	
20	G: =Initial acceleration is too:: slow and maybe real (.5) our initial speed is too slow and maybe acceleration isn't good either.	<i>S stops simulation.</i>
21	S: Ok we, we, we got the general::	<i>S resets simulation</i>
22	G: Hi (2.0)	<i>G waves at camera. Laughs.</i>
23	S: Hi mom. (4.0)	<i>S waves at camera, laughs.</i>

Table 4. Episode 6-2: The opening episode was followed by a period in which Gary and Sam coordinated their conception of the problem as being one in which the velocity and acceleration arrows are in opposite directions.

⁶ Transcripts are presented using notation found in Suchman [16]. Appendix 1 describes the notation.

24	S: Oh, if we could get rid of one of those things. (2.0) I'm just doing that so I can see what's in there. (2.0)	<i>Makes velocity longer.</i>
25	S: Ok, this one (9.0)	<i>Makes velocity shorter, runs.</i>
26	S: Ok this one's too great. Yeah. (1.5)	<i>Stops, resets simulation.</i>
27	G: Initial speed isn't good enough. Aw [
28	S: Ok (?)	
29	G: You can't tell. [
30	S: Now that looks fine. that's I think=	
31	G: =Initial speed is fine?	
32	S: Just make this one smaller.	<i>Makes acceleration shorter.</i>
33	G: Wait you know we could use particle one to (.) test the initial speed an::	<i>G points to screen, looks at S, gestures up.</i>
34	S: now this [
35	G: Oh we couldn't convert it to a particle two.(1.5)	
36	S: Now ok. That's slightly too much. It should come back.	<i>Runs.</i>
37	G: hhh. eventually.	
38	S: Ok (.) so: (1.0) reset maybe make this one small [<i>Stops, resets.</i>
39	G: It's hard to tell the difference between (2.0) initial speed an::d acceleration. (4.0)	<i>S clicks on a vector but makes no change, runs.</i>
40	S: Aw:: I don't think that did anything.	<i>Stops, resets.</i>

Table 5. Episode 6-3: This episode marks the beginning of a divergence of the partners' ideas about how to set the length of each arrow in order to achieve a closer match between the model and the particle. Throughout this episode, the students were working on different conceptual problems and they were talking more to themselves than to each other.

“acceleration”. Because this utterance pushed for more detail within the context of what had already been accomplished, it is both an acceptance of past work and a presentation of a proposal for future refinement. In this statement, Gary produces a ‘self repair’ [15] in which he corrects the initial part of the utterance, “initial acceleration” to be “initial speed”. Sam’s response (line 21) deflected Gary’s suggestion by referring back to what they had already accomplished (getting the shape right). The episode ended with a brief period of off-task activity, in which Gary and Sam waved at the camera (lines 22 and 23). Off-task behaviour was fairly rare in Gary and Sam’s sessions.

These first two episodes set the context for the rest of the challenge. On one hand, the partners had agreed on the basic configuration of the arrows to match the Observable World motion (see Figure 3b). On the other hand, they had not agreed on how to adjust the lengths of the arrows to achieve a closer match with the target motion.

In this period of divergence, Sam continued to search for the correct solution by experimenting with different lengths for the vectors (see Table 5). During this search he occasionally reported on his work. Meanwhile, Gary’s utterances show that he was working on a conceptual problem: distinguishing initial speed from acceleration. Both participants already understood many differences in the two concepts — they know that each maps onto a different arrow, that the arrows do not have symmetrical effects, and that in this particular challenge the velocity arrow is up, while the acceleration is down. Thus, it seems fairly certain that Gary was trying to make a particular conceptual advance. That is, he wanted to find a principle that would determine the correct length of each vector. This is reflected in Gary’s statements, “you can’t tell” (line 28) and “it’s hard to tell the difference between initial speed and acceleration” (line 39). His first idea of how to do this using the particle without acceleration was a good one (line 33). One could use this particle to match the *initial* motions of the Observable World and Newtonian World, because the initial motion is determined by the velocity vector alone. Then having fixed the velocity, one could focus on the acceleration. This would be a more systematic approach than that which Sam was following. Gary used a question (line 31), an interjection (line 33), and a comment (line 39) to try to get Sam to think about the different effects of velocity and acceleration arrow lengths. Sam’s verbalisations, however, were not responses to any of Gary’s comments during this time. Instead, he used his turns to report on his actions.

The lack of smooth turn-taking in this episode shows that each participant was talking out loud to himself more than to the other. The divergence between the subjects’ work continued and culminated in a breakdown in the interaction. Gary eventually disengaged himself from the task (as can be seen by his verbal unresponsiveness) and began to play with the microphone (see Table 6).

In this section, Gary and Sam renewed a higher level of collaborative engagement. It is interesting to examine the structure of this successful interchange. Gary started by asking a

question (line 44) that directed focus to a part of the motion using shared terminology. Sam's response (line 45), an answer to Gary's question, indicated a willingness to share the activity.

41	S: Ok so we're gonna try another one hhhere Oops wrong one (.) this down? (.) This one up. (4.5)	<i>Makes acceleration bigger.</i> <i>Makes velocity smaller,</i> <i>acceleration smaller.</i>
42	S: Maybe it's slightly (9.0)	<i>Clicks on vector but no change.</i> <i>Runs.</i> <i>G moving microphone around .</i>
43	S: That's cool.	<i>Stops, resets.</i>
44	G: hhh. (.5) Acceleration:: should be increased and (.) Is it going up at a good rate?	
45	S: We could change it.	<i>Makes velocity longer.</i>
46	G: Is it going up too slow or too fast?	<i>G turns to S.</i>
47	S: Too slow.	
48	G: Ok so increase the initial speed [<i>G gestures up.</i> <i>Runs.</i>
49	S: I did=	
50	G: =and::=	
51	S: =OK, now maybe this one we might get it. [
52	G: uh now you need to increase the acceleration too.	
53	S: Yeah, ok. (1.0) Uh. (1.5)	<i>Glances at G, stops, resets.</i>

Table 6. Episode 6-4: Although Gary was disengaged with the task while he played with the microphone, his focus eventually returned to the screen and to Sam's running commentary. In this episode, Gary and Sam became re-engaged in sharing ideas through discourse and action. Through coordinated presentations and acceptances, they began to converge on a shared conception of the properties of the lengths of the vectors.

This willingness contrasts sharply with his previous lack of attention to Gary's utterances. In the next two exchanges, Gary orchestrated the construction of a shared concept corresponding to initial speed. The form of this discourse is interesting. First, Gary specified a particular attribute of the motion. He again framed this statement as a question (line 46). Sam again

responded appropriately and provided the value of the attribute (line 47). Gary's statement in the next turn (line 48) had two overlapping effects: it named an action to be taken (increase velocity) and named the object of that action "initial speed". In his next turn (line 49), Sam confirmed the interchange both by his verbal response and his subsequent actions with the mouse.

This discourse event has the structure of a socially distributed production. This particular SDP was presented in instalments by G, with acceptances in intervening turns from S. The production could be paraphrased:

IF the Goal is to adjust the initial speed,
and the speed "going up" is too slow,
THEN make the velocity vector bigger.

The content of this SDP is a qualitative proportionality between the initial speed and the length of the velocity vector. This understanding is a breakthrough for the collaborators because it connected the length of the velocity vector to a local part of the motion: before this SDP, Gary and Sam had used the term "initial speed", but they consistently used it only as a name for the velocity vector. This use is distinct from the use of "initial speed" to refer to the speed at the beginning of a motion. By connecting the name of the arrow to the speed at which the motion begins ("going up"), Gary and Sam connected the length of the vector to a property of motion. They then adjusted the vector to a close approximation of the correct length (see Figure 3c).

While Gary is the first to give verbal expression to this idea (line 44), it is not clear who originated the idea. In the time period directly preceding his utterance, Sam had been engaged in extensive experiments with the lengths of the vectors. Gary could have been giving verbal expression to an idea that originated in the Sam's experiments with the computer. This interpretation is supported by the fact that Sam was already adjusting the vector appropriately even as Gary was completing the SDP. However, given the nature of the data, we cannot draw definite conclusions about the originator of the idea. Regardless of originator, this episode did mark a convergence in the partners' understanding of the meaning of the length of the velocity vector. This convergence persisted throughout the remainder of the challenge, the session, and into the interview that followed.

The shared nature of their work during episode 6-5 (Table 7) is evidenced in the data in many ways. For example, most of the conversational turns following the statement of a new idea included an acceptance. The students used questions to elicit the consent and involvement of their partner in shared decision-making. The acceptances were sometimes explicit (e.g., lines 54, 57, 59) and sometimes implicit but clearly marked by the discourse structure. One such implicit acceptance was Gary's restatement of Sam's previous utterance (line 56). Furthermore, even though Sam was still in control of the mouse, the control of the activity was shared. This is nicely illustrated in the part of the dialogue that begins with "Here bring it

54	S: Ok:: so (.5) this one goes down right there. (2.0)	<i>Clicks on acceleration but makes no change, runs.</i>
55	S: That might not be enough. [
56	G: It might take a little more than that. (.5) yeah.	
57	S: Ok:: (4.0)	<i>Stop, reset.</i>
58	S: What about there? (1.0)	<i>Makes acceleration longer.</i>
59	G: OK	
60	S: Well that's too much. Definitely.(3.0)	<i>Makes acceleration shorter.</i>
61	G: hhh. here bring it down there we'll see how it is. (4.0)	
62	S: Is it going down?	<i>Clicks twice, but no change, clicks again makes acceleration longer.</i>
63	G: Yeah it is.	
64	S: Is that good?	
65	G: Yeah try it. [
66	S: I:: I'm pretty sure it's gonna be way too much for (2.0)	<i>Runs.</i>
67	S: Yeah, ok. (1.0)	<i>Stops motion only part way up.</i>
68	G: Oh, way too much. (16.0)	<i>Resets, makes acceleration smaller, runs.</i>
69	G: Oh, so close. (1.0)	<i>Stops.</i>
70	S: OK (1.0) I have [<i>Resets.</i>
71	G: How come it took so fa::r (.) uh so long to get back down (.) maybe (.) acceleration:: (.) (looks at S) :::up speed down. (1.0)	
72	G: Doesn't do anything does it? (2.5)	<i>Clicks but no changes, runs.</i>
73	G: Turn off the record hhhh (5.0)	<i>Stops, resets.</i>
74	S: Ok. (.) Top one make it go down. (.5) Top one. (2.0)	<i>Clicks but no changes.</i>

Table 7a. Episode 6-5

75	S: Okay::: Start. (11.0)	<i>Runs. Stop at top of screen on way up, makes velocity, smaller, then larger, velocity now unchanged from last run.</i>
76	G: That's good yeah that's good.	
77	S: Good? Just gonna try it. (3.0)	<i>Runs.</i>
78	S: Mmm::	<i>Stops, resets.</i>
79	G: Maybe more acceleration?	
80	S: Ok, well let's reset (mess with?) this first (3.0)	<i>Makes velocity shorter, runs.</i>
81	S: OK::: ((very low rumble)) (6.0)	

Table 7b. Episode 6-5 (continued): Following the successful re-engagement in the shared conception of the problem, the partners continued a period of mutually shared activity that extended to the end of the challenge. This episode is different from the previous ones in that the students are more mutually engaged in the task. Further, they are working out simple procedural details rather than new concepts. Specifically, they worked on a shared interpretation for the length of the acceleration vector. This reflects both partners' satisfaction with the length of the velocity vector determined in the previous episode.

down there" (line 61). Upon hearing Gary's utterance, Sam began to move the tip of the acceleration vector downwards. While doing so, he involved Gary in the hand-eye feedback loop ("Is it going down?") enabling the pair to co-determine the setting of acceleration (S: "Is that good?" G: "Yeah try it.").

Another difference between this episode and the preceding ones is that the content of the conversation no longer reflects differences in interpretation — the participants were now working out procedural details. This is not to say that Sam and Gary *completely* share a common understanding of the task. As evidence to the contrary, note that Sam's acceptances lack the kind of paraphrasing and elaboration that are often used to signal that participants fully comprehend each other (see [3]). It seems that at this point, Sam was just beginning to appreciate Gary's point of view. Nonetheless, as the activity unfolded, the understanding became sufficient for the partners to make two additional advances in their physics knowledge.

The discourse and the sequence of actions in episode 6-6 (Table 8) provide an interesting example of how partners get new ideas introduced and accepted into an established course of action. Gary marked his new idea by entreating Sam to pay attention ("wait"), and asking, "You know what we can do?" (line 82). This question was a signal to Sam that something new was to follow. However, Sam did not respond to the question. Gary then asked him

82	G: Wait you know (.) you know what we c::can do we can uh stop stop. (1.5)	<i>G points to screen then reaches over and clicks mouse, stops.</i>
83	S: Ok, now I just	
84	G: Now=	
85	S: =wanna see how much we're off	<i>S points to screen.</i>
86	G: we should compare those. (1.0) Now make it go up and compare those and:: (.) um:: hhhh. (.) see what the set off rate is=	<i>Points to dots in OW. Gestures up. Slides finger to NW, gestures up. Slides finger back to OW, gestures up, brings hand to his lap.</i>
87	S: =Oh yeah ok. (1.0) So this is greater. (.) This is way greater. Just look at that. Good= (.)	<i>Moves mouse so mouse-cursor is in OW and shakes it. ((Interpretation: comparing dot spacing, wider in OW than NW.))</i>
88	G: =So we need a higher speed?	
89	S: Yeah. (1.0)	
90	S: So maybe like:: (.) tha::t.	<i>Makes velocity longer.</i>
91	S: And make=	
92	G: =Maybe::=	
93	S: =this one like:: (.) tha::t=	<i>Makes acceleration longer.</i>
94	G: =Now drop it? (1.0)	<i>Voice drops off to inaudible.</i>
95	S: Is that good? (.) Yeah	
96	G: try it.	
97	S: maybe. (1.0)	<i>Runs.</i>

Table 8. Episode 6-6: In this episode, Gary interrupted the current activity to suggest a refinement to the shared understanding of the length of the velocity vector. His refinement connected the length of the velocity vector to the spacing of the dots in the beginning of the motion. This was the first time that the spacing of dots was given a local interpretation.

directly to “stop, stop” and resorted to clicking the mouse (which had up to this point been completely in Sam’s physical control). At this point, Sam was still engaged in the previous course of action and began to justify what he was doing (lines 83 and 85). Gary interrupted (line 86) to offer the new idea. When he said, “We should compare those”, he referred to the trace dots that the particles leave behind as they move across the screen. As the spacing of the dots is an indication of speed, comparing dot spacing at the beginning of the motion is one method for determining the correct value of the velocity vector.

The idea of basing decisions on a comparison of dots was introduced by Sam in an earlier challenge, but until before this time the spacing of dots had not been given a local interpretation — Gary and Sam had compared *all the dots* in the Observable World to *all the dots* in the Newtonian World. The structure of Gary’s utterances shows that he had considerable difficulty expressing his idea, as he paused and interrupted himself several times. The gestures that accompany the idea, on the other hand, were quite clear: he pointed to the first few dots in the Observable World and gestured up, then pointed to the first few dots in the Newtonian World and gestured up, and pointed again to the dots in the Observable World. While his verbal expression “compare those and see what the set of rate is” (line 86) was possibly difficult to interpret, the combination of verbal expression and gesture were enough for Sam to make the correct inference. This is indicated by Sam’s subsequent acceptance (“Oh yeah OK”) and elaboration of the appropriate feature, the spacing between the initial dots (“this is way greater”). This elaboration leads into a SDP which expressed the qualitative proportionality between the dot spacing and the length of the velocity vector (lines 87 and 88). This SDP can be stated as:

IF the goal is to adjust the initial speed and,
the initial dot spacing is greater,
THEN make the velocity vector longer.

Another important point about the interaction during this episode is that both participants were using the association between the length of an arrow and its effect transparently. For example, when Gary suggested making the initial speed *faster* (line 88), Sam took the action of making the vector *longer* (line 90). As the discourse proceeded, reference to the length of vectors dropped out in favour of references to the effect of changing the length.

The end result of Gary’s interruption in this episode was the construction of a important new piece of shared knowledge: the qualitative proportionality between the local dot spacing and the local (instantaneous) speed. This shared interpretation of local dot spacing was confirmed in the interview that followed the task. The process of the construction is also important because it illustrates how participants utilise a combination of linguistic, gestural, and physical resources (the computer screen and the mouse), in order to introduce new ideas into the collaboration.

98	S:	Ok see how the acc(.) initial (1.5) ooh closer and see how much we're off=	
99	G:	=Our acceleration might be too high. (3.0)	<i>Stops, resets.</i>
100	S:	Ok. This is gonna be a li::bit less.	<i>Clicks but no change.</i>
101	G:	Didn't do anything. (14.0)	<i>S makes acceleration shorter, runs, stops, resets.</i>
102	S:	Wants to be cool=(?)	
103	G:	=it's higher? (13.0)	<i>Clicks four times, no changes.</i>
104	S:	Oh. (13.0)	<i>Shakes head. Runs, stops, resets.</i>
105	S:	Ok. (5.0)	<i>Moves face right up against screen.</i>
106	G:	hhh. (3.0)	
107	S:	There right about:?. (1.0)	<i>Makes acceleration longer.</i>
108	G:	Drop it there.	
109	S:	There?	
110	G:	Yeah.	
111	S:	No, up. I can you can see the degrees if you look closely enough. (.5) One more degree. There. (4.0)	<i>Clicks but no change. Runs.</i>
112	S:	uh? Ok. (2.5)	<i>Stops, resets.</i>
113	G:	Maybe higher initial speed?	
114	S:	I'm gonna change the degree one an (.) makes the initial (.)	
115	G:	Aah::=	<i>Falling intonation.</i>
116	S:	=speed go up by one. (.5) Wait. (6.0)	<i>Clicks but no changes. Runs.</i>

Table 9. Episode 6-7: In this episode the partners continued to work together to refine the match between the Newtonian and Observable World motions. In addition, Sam introduced the idea that the vectors could only be moved in discrete amounts. This helped Gary and Sam regulate their search for the correct lengths of the vectors.

Like episode 6-5, episode 6-7 (Table 9) was marked by fluid turn taking and mutual engagement in the task and decision making. They conducted 5 runs, each time making slight changes in the vectors. At one point (line 105) Sam moved physically closer to the screen to be better able to perceive the precision of the vector movements. While doing so, he made an important contribution to the task — he introduced the idea that the vectors are adjustable only in discrete units (“see the degrees”, line 111) and that they could therefore adjust the vectors one unit at a time. This strategy gave the participants better control over trial and error problem solving.

117	S:	hh Ok=	<i>Stops.</i>
118	G:	=It looks pretty good just(.)didn't go up high enough.(.5)	
119	S:	Yeah ok so::=	<i>Resets.</i>
120	G:	=maybe increase the speed.	
121	S:	Ok. (5.0)	<i>Makes velocity longer.</i>
122	S:	There:: (6.0)	<i>Falling intonation, runs.</i>
123	S:	Almost::=	<i>Falling intonation.</i>
124	G:	=Close enough.	<i>OW and NW motion match except NW initial position is one square too high.</i>
125	S:	Wait here I want to do one thing. (2.5)	
126	S:	Reset. (4.0)	<i>Resets, clicks but no change.</i>
127	G:	((inaudible)) (.) There. Saw it move. (7.0)	
128	S:	Yes. We got it perfect.	<i>Looks at G, moves face back from screen.</i>
129	G:	Perfect.	<i>G puts hand on top of S's hand on mouse. They go on to the next challenge.</i>

Table 10. Episode 6-8: This episode shows the negotiation of the challenge ending. This was typical of all challenges during this session; the successful completion of a challenge was jointly determined by the participants. Also during this episode the participants construct a qualitative proportionality between height of a trajectory and the length of the velocity vector.

Episode 6-8 (Table 10) began with Gary making an evaluation of the current state of the problem and what remained yet to be solved (line 118). This proposal took the form of a SDP that expressed the proportionality between the height of the ball's path and the initial speed. It is likely that Sam already recognised this relationship, because he had adopted a procedure of

stopping the simulation as soon as the particle went too high. Once again, while the originator of the idea is uncertain, Gary was the first to verbalise it.

Toward the end of this episode, Sam noted (line 123) that they “almost” had a solution, and Gary responded (line 124) that he was satisfied with the degree of success that they had attained. In fact, the two motions had exactly identical velocities and accelerations, but the initial position in the Newtonian World was about a centimetre too high (see Figure 3d). Sam announced that he wanted to, “do one more thing” and, receiving no objection from G, he proceeded to make one further adjustment. Although, he moved the particle downwards, he did not move it enough for the simulation to recognise the change. Nonetheless, after this change, both partners agreed that their solution was “perfect”.

The challenge closed on this note of mutual satisfaction. The participants’ pride in their performance during this challenge was also apparent in the interview that followed the session: Gary started out the interview by saying, “Do you wanna see number six? Probably one of our most famous ones”.

5. Conclusions

Our perspective has characterised collaboration as a process of constructing and maintaining a Joint Problem Space. Our analysis of challenge six illustrated how coordinated production of talk and action by two participants enabled this construction and maintenance to succeed. The students used language and action to overcome impasses in shared understanding and to coordinate their activity for mutually satisfactory results. But as this analysis made clear, the process of collaborative learning is not homogeneous or predictable, and does not necessarily occur simply by putting two students together. Students’ engagement with the activity sometimes diverged and later converged. Shared understanding was sometimes unproblematic and but oftentimes troublesome. The introduction of successful ideas was sometimes asymmetric, although it succeeded only through coordinated action. These results point to the conclusion that collaboration does not just happen because individuals are co-present; individuals must make a conscious, continued effort to coordinate their language and activity with respect to shared knowledge.

The inherent fragility of the collaborative learning process has led us to consider the resources collaborators employ to surmount difficulties that arise in the course of working together. As our analysis has shown, the most important resource for collaboration is talk. Collaborators use the overall turn-taking structure of talk, as well as specific discourse forms such as narration, questions, socially-distributed productions, and repairs in service of their mutual understanding. These discourse forms allowed the students to produce shared knowledge, to recognise divergent understandings, and to rectify problems that impeded joint

work. Language, however, does not occur in a vacuum. Dewey [4] put it succinctly: “Language would not be the efficacious instrument it is, were it not that it takes place upon a background of coarser and more tangible background physical means to accomplish results.”

We see the ‘computer-supported’ contribution to collaborative learning as contributing a resource that mediates collaboration. In ordinary circumstances, one cannot imagine two 15 year olds sitting down for 45 minutes to construct a rich shared understanding of velocity and acceleration. But in the context of the support provided by the Envisioning Machine activity, our students were successful in doing just that. This leads one to ask: how do resources provided by the computer support collaboration?

Our data suggests several possible answers.⁷ First, we observed the use of the computer as a means for disambiguating language. Gary and Sam do not have a precise, technical vocabulary for talking about motion so they used the objects in their physical situation to support their talk. For example, in the introduction of the idea of comparing dots in episode 6-6, the students used the computer display as a means for establishing shared references. In addition, their maintenance of a shared focus of attention on the computer screen enabled efficient, but ambiguous expressions such as “make it more” to be correctly interpreted. The computer interface also provided an alternate means for producing conversational turns: actions with the mouse could be interpreted as non-linguistic presentations and acceptances of ideas. Second, we observed the use of the computer activity as means of resolving impasses. When students had differing opinions, as in the beginning of challenge six, they resolved their differences by trying out the ideas and seeing what worked. When students had insufficient ideas to progress, as in episode 6-3, they could resort to experimentation with the computer as a means for generating new ideas. Third, we saw that the computer was a device that invited and constrained students’ interpretations. The EM display was carefully designed to suggest appropriate interpretations. An instance of suggestion occurs in the beginning of challenge six, when Sam saw the new motion and leapt to the idea that the arrows should be opposite. But all the interpretations suggested by the EM representation are not necessarily appropriate. Because the simulation behaves according to Newtonian physics, the computer also constrained interpretations towards the underlying scientific model.

Detailed investigations of collaboration such as this can contribute to future investigations of computer-supported collaborative learning in at least two ways. First, this type of analysis can be viewed as a methodology for coming to a deeper understanding of how the benefits of collaboration are realised. Second, this type of analysis can lead to better development of the kinds of supporting resources that computers can provide for collaborative learning. Clearer understanding of the collaboration as a process of constructing and maintaining a shared

⁷ These observations were first reported in Singer, Behrend and Roschelle [17]. We thank Janice Singer for her contributions to this work.

conception of the task can be beneficial for future designs of collaborative learning environments.

Acknowledgements

This work was partially supported by an internship programme at the Institute for Research on Learning, Palo Alto, California.

References

1. Azmitia, M. & Perlmutter, M.: Social influences on children's cognition: State of the art and future directions. In: *Advances in Child Development and Behavior* (H. Reese, ed.). New York: Academic Press 1989
2. Bobrow, D.G. (Ed.): *Qualitative Reasoning about Physical Systems*. Cambridge, MA: MIT Press 1986
3. Clark, H. & Schaefer, E.: Contributing to discourse. *Cognitive Science*, 13, 259-294 (1989)
4. Dewey, J.: *Democracy and Education*. New York: The Macmillan Co 1932
5. Ericsson, K.A. & Simon, H.A.: *Protocol analysis: Verbal Reports as Data*. Cambridge, MA: MIT Press 1984
6. Lerner, G.H.: *Collaborative turn sequences: Sentence construction and social action*. Unpublished Ph.D. dissertation, University of California, Irvine, CA 1987
7. Levinson, S.: *Pragmatics*. Cambridge: Cambridge University Press 1983
8. Mead, G.H.: *Mind, Self and Society*. Chicago: University of Chicago Press 1934
9. Piaget, J.: *The Moral Judgement of the Child*. Glencoe, IL: Free Press 1932
10. Rogoff, B.: *Apprenticeship in Thinking: Cognitive Development in Social Context*. Oxford: Oxford University Press 1990
11. Roschelle, J.: *Students' construction of qualitative physics knowledge: Learning about velocity and acceleration in a computer microworld*. Unpublished doctoral dissertation. Berkeley, CA: University of California 1991
12. Roschelle, J.: *The Envisioning Machine: Facilitating students' reconceptualization of motion*. Unpublished manuscript, Xerox Palo Alto Research Center, CA 1986
13. Schegloff, E.A.: *Conversation analysis and socially shared cognition*. In: *Socially Shared Cognition* (L.B. Resnick, J. Levine, & S.D. Teasley, eds.). Washington, DC: American Psychological Association 1991
14. Schegloff, E.A.: *Discourse as an interactional achievement: Some uses of "uh huh" and other things that come between sentences*. In: *Analyzing Discourse: Text and Talk* (D. Tannen, ed.). Washington, DC: Georgetown University Press 1981
15. Schegloff, E.A., Jefferson, G., & Sachs, H.: *The preference for self-correction in the organization of repair in conversation*. *Language*, 53, 361-382 (1977)
16. Suchman, L.A.: *Plans and situated action*. New York: Cambridge University Press 1987
17. Singer, J., Behrend, S.B. & Roschelle, J.: *Children's Collaborative Use of a Computer Microworld*. *Proceedings of the Conference on Computer-Supported Collaborative Work* (pp. 271-281). New York: Association for Computing Machinery 1988
18. Tudge, J. & Rogoff, B.: *Peer influences on cognitive development: Piagetian and Vygotskian perspectives*. In: *Interaction in Human Development* (M. Bornstein & J. Bruner, eds.). Hillsdale, NJ: Lawrence Erlbaum Associates 1989
19. Vygotsky, L.S.: *Mind in Society*. Cambridge, MA: Harvard University Press 1978
20. Wilkes-Gibbs, D.: *Collaborative processes of language use in conversations*. Unpublished doctoral dissertation, Stanford University, CA 1986

Appendix

Notation

[Bracket indicates a point at which a current speaker's talk is overlapped by the talk of another, with overlapping talk directly beneath.
:	Colons indicate a lengthened syllable, the number of colons corresponding to the extent of lengthening.
?	Question intonation.
.	Full stop with falling intonation.
=	Equals sign indicates no interval between the end of a prior and the start of a next piece of talk.
.hh	Audible breath. Dot before indicates in-breath. No dot indicates out-breath.
()	Words enclosed in parenthesis indicate either non-linguistic action, or transcriber is uncertain of the verbatim.
(())	Double parenthesis indicates features of the audio other than verbalisation, or note from the transcriber.
(0.0)	Numbers in parenthesis indicate elapsed time in tenths of a second.
(.)	Untimed pause.
OW	Abbreviation for "Observable World".
NW	Abbreviation for "Newtonian World".