The Productive Agency that Drives Collaborative Learning

Daniel L. Schwartz

Running Head: Productive Agency in Collaboration

pp 197 - 218 in

Dillenbourg, P. (Ed.). Collaborative learning: Cognitive and computational approaches.

NY: Elsevier Science/Permagon.

In over 60 years of research, there have been very few demonstrations that working in a small collaborative group yields cognitive outcomes that cannot be matched or exceeded by the most competent member of the group (Barron, 1992; Hastie, 1983, Hill, 1982). This finding extends to tasks as diverse as brainstorming (Dunnette, Campbell, & Jaastad, 1963), writing (Fox & Lorge, 1962), problem solving (Kelly & Thibaut, 1969), rope pulling, and rule induction (Laughlin & Futoran, 1985). The research suggests that there is nothing particularly special about working in small groups, at least with regards to cognitive outcomes like learning and intellectual artifacts. Yet, this assertion seems to fly in the face of common sense.

For example, I recently had a discussion with an ethologist, Dick Porter, who studies whether prenatal exposures to odors affect behaviors after birth. He described studies in which chicken eggs are exposed to a particular scent (e.g., Turro, Porter, & Picard, 1994). After hatching, groups of three chicks are placed in cages with a dish of scented grain and a dish of plain grain. He then observes their behavior towards the two dishes. What is of relevance here is why he uses three chicks per cage instead of one. Although chickens are cheap, they are not that cheap; and using three chickens as the unit of analysis increased the number of chickens he needs three-fold. He explained that the reason for using three chicks is that if you put a single chick in a cage it stands still. It loses all productive agency. It will not move towards either dish of grain in the attempt to learn the location of the preferred food. Chickens, like humans, are very social creatures.

In the case of the chickens, it is quite clear that two heads (or three heads) are better than one -- better in ways that have profound effects on their abilities to learn. Yet, in the case of humans, there is little research that shows that 2+ heads are much better than one. What are we to conclude? That chickens are different than humans or that our small group research has been missing something? Well, clearly the former is true, but so is the latter. This latter point is where the papers in this volume come into play. As a collective, the chapters take on the very important and difficult task of trying to develop new research approaches that can help identify what is special about collaboration. The authors of the chapters draw on state of the art research techniques involving robots and humans, and they tackle state of the art problems such as creating computer supported collaborations. Bringing computers into the cooperative equation is a promising new approach to this area of study (see also, Robertson, Zachary, & Black, 1990).

OVERVIEW

In the following discussion of the chapters, I begin by describing what I think is especially fresh and exciting about the use of computers to understand cooperation. Next, I propose a picture of human nature that extends some of the ideas from the papers and helps to counter-balance others. My picture will be distinctly psychological. I hope to show the fundamental importance of considering productive agency in our research on collaboration and collaborative learning. Just as the socially-situated perspective has made it clear to cognitive psychologists that culture is not simply a background variable, I would like to suggest that research into collaboration should make it clear that agency is also not "simply a background variable."

One reason I have chosen to emphasize individuals when the topic is collaborative groups comes from an observation by Weiss and Dillenbourg (this volume). In their discussion of collaborative computer agents, they suggest that it may be difficult to pre-specify all the rules or social conventions needed for collaborative behavior. Collaborative situations are often too complex. Consequently, the agents need to be able to self-improve and organize their collaborations. So, rather than exclusively describing the in-place conventions that regulate on-going group behavior, one may focus on the properties of individuals that make collaborative behaviors emerge. I believe the way to focus this inquiry starts with an understanding of productive agency.

My goal is to step back from the specific models of cooperation that have been generated by the authors and to explore some basic issues involving assumptions about the psychology of human nature. These issues -- individual agency, opportunities to be productive, and constructive learning -- were triggered and informed by the authors, but my discussion sometimes crosscuts the issues the authors considered. One place where I would like to step back involves the authors' frequent wrestling with what counts as cooperative or collaborative learning. I do not wish to add to the clutter of definitions that Dillenbourg (this volume) heroically organizes in his introductory chapter. Missing, however, is the notion of agency which I believe is intimately intertwined with collaborations. Collaboration is not constituted by people following social laws as though they were physical laws. People need to choose whether and when to collaborate and whether to go beyond the minimum necessary to meet the rules of collaboration. I will attempt to argue that the very definition of collaboration involves the notions of agency and an individual's ability to represent other people's agency.

I also think that we need to consider the motivations that lead people to collaborate for the purposes of learning. There are many instrumental motivations that are well-documented in the literature (Slavin, 1983); for example, joining to beat a common enemy. But I would like to focus on motivations that involve learning per se. Many of the current discussions view collaborative learning as the appropriation of ideas from others (implicitly, at least, this seems to be the motive). I wonder whether this is a satisfactory view of human collaboration. I want to propose another view; one that emphasizes the important role of being a contributor rather than a borrower. People appropriate knowledge when they are given opportunities to produce knowledge.

Finally, I want to explore issues of learning and collaboration, and the role of knowledge construction and language in this process. It is tempting to view some forms of collaborative learning as due to the direct communication of linguistic ideas from one agent to another. Computer models of collaboration, for example, often rely on the communication of propositions from one agent to another. But, one does not really learn a proposition, one learns from a proposition (Bransford & Nitsch, 1978). People use propositions to help them construct new knowledge. Linguistic representations can only play a partial role in learning from a partner. I would like to explore the nature of this partial role to help pinpoint what knowledge is most likely to be generated in linguistically mediated collaboration.

All told, I want to paint a picture of humans in which productive agency plays a central role in the characterization of collaboration. One way this agency is expressed is by the decision to collaborate and the effort to make a collaboration work when social rules are not sufficient for successful collaboration. Another way this agency is expressed is by the motivation to produce and contribute. Finally, this productive agency appears in the very way we learn -- we construct knowledge. In the following sections I develop these ideas more fully and consider some of the ways that these characterizations may be addressed empirically. At the end of the chapter, I consider some of the ways these observations could be incorporated into computer models of collaboration and collaborative learning. But first, I next suggest what is so special about computers in the context of understanding collaboration.

THE POTENTIAL OF COMPUTERS FOR UNDERSTANDING COLLABORATION

The chapters offer excellent examples of how computers provide new energy for understanding collaboration. These examples include practical problems and potentials made possible by computer technologies like tele-communication (Hansen, Lewis, Dirckinck-Holmfeld, & Rugelj, this volume) and robotics (Joiner, Issroff, & Demiris, this volume). These real technologies demand new methods of analysis (Baker, Hansen, Joiner, & Traum, this volume; Littleton & Hakkinen, this volume). The answers to the unique challenges of integrating computer technologies into human endeavors may have important, general implications for understanding collaboration.

A different example of how computers provide new direction to collaborative learning research comes from work in artificial intelligence and distributed artificial intelligence (Hoppe & Ploetzner, this volume; Mephu-Nguifo, Baker, & Dillenbourg, this volume; Weiss & Dillenbourg, this volume). This work highlights the great potential of using the computer as a research tool in the tradition of cognitive science -- as a way to model and elucidate the nature of cooperation among humans. In my discussion I will focus on this possibility, because I believe the use of computers to model collaboration is a potent methodology. Regardless of one's theoretical disposition towards information processing, there are two important possibilities. Computer modelling in general may provide a way to handle the awkward size of small groups, and artificial intelligence techniques in particular may help to clarify the importance of representing the thoughts of others during collaboration.

The Awkward Size of Small Groups

Recall the earlier observation that it may be impossible to pre-specify the group rules and social conventions that organize people during their various complex interactions and settings. While this problem may have something to do with human nature, it may also have something to do with the methodologically awkward size of small groups. This point may be understood in the context of a conversation I had with Joe Hamilton, a nuclear physicist at Vanderbilt. In relation to another project, I asked him why he thought mathematics has served modern science so well. His response was surprising. He said that in his domain, traditional mathematics was not working so well. He pointed out that we have excellent mathematical tools, like statistics, for handling large numbers of entities as might be the case when talking about the behaviors of the many atoms and molecules in a gas. He also pointed out that we have excellent mathematical tools, like calculus, for handling small numbers of entities as might be the case when predicting the behavior of a projectile in space. However, we do not have a traditional mathematics for handling middle ranges of interacting entities like the number of protons and neutrons in a nucleus. As a result, nuclear physicists rely on computational models more than other sub-disciplines in physics.

This seems to be most relevant to research involving small group collaborations -- collaborations that often involve several interacting participants (Dillenbourg, this volume). Unlike sociological data, small group behaviors are not smoothed by averages over the thousands of people involved in a social movement. And unlike individual psychological analyses, the multiple interacting agents of a group cannot be turned into the simple equation of a single individual interacting with a single stimulus. It strikes me that computer models of cooperation may be powerful scientific tools that can help in the way that traditional mathematics has helped other physical and social sciences. They can provide a formalism that helps define regularities and that can model complex interactions among medium numbers of entities.

Modelling the Thoughts of Others

The second reason computers may make a powerful research ally, the one that I will emphasize here, is that they permit a type of modelling that is ideally suited to characterizing individuals within collaborations. Specifically, artificial intelligence techniques allow us to model people's thoughts in the folk terms that they experience; terms like, "I agree," and, "This is my goal." This modelling ability may be extremely important to studies of collaboration, even more important than for studies of individuals. This is because a critical component of collaboration involves the representations that individuals have of one another's thoughts. As I describe next, this may be at the heart of understanding collaboration.

A common claim made throughout the papers is that for collaboration to occur, it is necessary for the collaborators to have a model of one another's thoughts, and ideally for the collaborators to have a shared set of models. Weiss and Dillenbourg speak of this necessity with regards to learning. They state, "The 'deep secret' of collaborative learning seems to lie in the cognitive processes through which humans progressively build a shared understanding" (p. xx, this volume). If this secret is in fact critical to collaborative learning, then it seems that the capability of computers to model thought, even if poorly, is particularly relevant for understanding collaboration. This is because computers can model the way we construe other people's thoughts. In fact, this is exactly what computer models of cognition are; they <u>are</u> models of other people's thoughts. In this light, they can give us a chance to explore how models of other people's thoughts might affect collaborative activity.

Learning through collaboration does not always require the representation of the mental states of one's partner. As Joiner, Issroff, and Dmiris (this volume) point out, there are different types of joint activity and learning. In some situations I do not need to know your thoughts, I only need to know the consequences of your actions. For example, in a foraging task, robot A can coordinate with robot B by simply noting where robot A is and has been. More generally, there are many situations, for example traffic lights, where people coordinate and learn behaviors by relying on social conventions. But, as I will try to later, these types of joint activities are not really what we have in mind when we think about collaboration.

Moreover, it is clear that we do learn to use other people's mental states to further our collaborative endeavors. Krauss and Glucksberg (1969), for example, developed a screen referenced task in which participants talking from either side of an opaque screen try to put their respective sets of nonsense shapes into identical orders. Although each partner has the same set of shapes, the shapes are in different orders and do not have conventional names. Therefore, the participants need to infer and represent the shape that their partner has in mind in order to complete the task. Krauss and Glucksberg found that there is steady improvement in collaborative performance with age: five-year old children show little success, six-year old children show some gains, eight-year old children show more success, and so forth. So, even if modeling another person's mental states is not a necessary pre-requisite for learning in groups, it is clear that it is something upon which humans rely. The computer, with its ability to model the way that humans represent other people's thoughts seems ideally suited to sharpening our thinking about collaboration and collaborative learning. At the end of the paper I suggest some possible explorations. But first, it is important to decide what types of traits a computer agent would need in order to model human collaboration.

AGENCY: FROM COMPLIANCE TO CHOICE

To begin developing of the notion of productive agency, I would like to argue that the very notion of collaboration depends on individual agency or intention. Many of the authors in this volume are concerned with a definition of "collaboration" and realize that it is important to begin to understand what does and does not fit this category. In this section I provide some contrasting cases in an effort to tease out some of the features that seem important to understanding what is essential about collaboration.

One contrast comes from Mephu-Nguifo, Baker, and Dillenbourg (this volume) who raise the issue of global versus local control in a cooperative interaction. Computationally, we might ask whether it is better for robot agents to follow globally specified procedures, or is it better to let cooperative behaviors emerge from procedures local to the individuals. In human terms, one might think of this as the issue of whether individuals or social conventions dominate collaborative interactions. Naturally, it is both, and the balance should change depending on the situation. Regardless, I think the essence of collaboration revolves around local control.

There are two types of local control. Local automata, like termites, have simple rules that, when coupled with other termites, can lead to complex structures. The global structure of this behavior arises from locally determined rules. The computer models used by the nuclear physicist might fit this category because the local properties of protons, neutrons, and so forth determine the global patterns they make when interacting. This situation is different from situations involving agents with local autonomy who choose to take on a particular behavior. It is this latter notion of autonomy that is particularly important for collaboration. If I am bound to a social role or predetermined sets of local rules, then it is difficult to say I am collaborative; instead, I am complying without agency, something like termites in a hive. To collaborate, individuals have to enter into relationships, they have to produce ideas, they have to choose whether to communicate, and they have to choose whether to compromise their goals.

In the next sections, I will develop the idea that collaborative learning is not constituted by people simply complying to a role. Collaborative learning takes agency and productive effort precisely because people must develop shared meaning across the differences in their roles and knowledge. To further this point, I will argue that the very definition of collaboration involves the idea of intent and non-compliance. First, I begin with the notion of "the effort after shared meaning," then I move to the idea of non-compliance. Hopefully, this discussion can help anchor the slippery notions of collaboration and collaborating agents. For example, it may provide a solution to the question of whether we should consider the neurons in the brain as collaborative.

Identity versus The Effort After Shared Meaning

A second contrast I would like to draw distinguishes between what is unique about research on collaboration relative to other social topics. In many different social-cultural traditions a common question is how an individual fits into his or her cultural milieu through structures like roles, rules, and social practices. In terms of psychological issues, a key construct here is the notion of identity. An individual gains his or her identity by taking on a role offered by a culture or social practice. Identity is a bridge between the local individual and the global cultural milieu. It is the desire to gain an identity within a cultural milieu that causes an individual to appropriate and come into compliance with the practices and roles of a culture.

While identity and compliance with cultural practices are surely central aspects of humanity, there are others. To borrow Bartlett's (1932) phrase, "the effort after meaning" is also important. Piaget, as well as most of cognitive science, has investigated people's pursuit of meaning and its effect on learning. For example, when people understand the meaning of a text passage they remember more (Bransford & Johnson, 1972). The importance of people's effort after meaning also appears in social settings where there is "the effort after <u>shared</u> meaning." When I talk to my family, friends, or colleagues, I want them to understand me, and I want to understand them. The desire to understand and be understood -- to share meaning -- is a strong motivator of human behavior and worthy of the status of a basic psychological construct. In terms of cognitive science, one might say that individuals want accurate representations of other people's thoughts, and individuals want other people to have accurate representations of their own thoughts. Baker, Hansen, Joiner, and Traum (this volume) point out that we can never actually reach this idealized shared meaning in any absolute sense. But, they explain that it is the effort after shared meaning that helps explain why we learn when we collaborate

Many of the questions for research on collaboration should be about how individuals construct local interactions among themselves to understand one another. Such questions are about how individuals interact with individuals, not how individuals interact with "culture." This is not to say that we should ignore the milieu in which these interactions take place, or the social rules that make exchange possible. Nor is this to say that desires for identity within a group do not play a strong role in group behaviors (e.g., the desire for the status of being team captain). But, there are times when it is worthwhile to minimize the emphasis on the cultural environment in which we swim, and overemphasize people's efforts for achieving shared meaning. Hansen, Lewis, Dirckinck-Holmfeld, and Rugelj (this volume), for example, describe a situation in which people with very diverse and distant "base communities" temporarily come together electronically to work on developing joint projects. The individuals do not have a solid common culture to regulate their interactions. Conceivably, the community differences serve as important forces that cause people to learn through collaboration. Several authors in this volume, for example, propose that misunderstanding <u>is</u> the progenitor of collaborative learning. When people cannot rely on common ground, they need to make the effort after shared meaning, and this often entails learning about and negotiating understanding with one another. For example, one base community needs to learn the merit behind another community's goals.

Collaborative learning often occurs exactly when people cannot assume the common ground of similar thoughts, roles, and goals. Conversely, collaborative learning often fails to occur when people overly rely on social roles. A nice example of this was recently brought to light by Kathleen Jones who studies parent-teacher interactions (e.g., Hoover-Dempsey & Jones, 1996). She explained that in parent-teacher meetings at school, parents and teachers often remain in their roles and maintain very different goals for the child (e.g., to make it through the year without further disruption versus to support a happy childhood and nurture a healthy adult). Because of the roles the teachers and parents maintain in these meetings, they often do not attempt to develop a shared understanding that can help them learn about the child's needs and behaviors from one another. The parent, for example, may passively listen to the teacher's evaluation and prescription. Clearly, an important question broached by nearly all the authors in this volume is when and how different types of affordances, like discourse pragmatics and visual representations, can induce people to make the effort after shared meaning.

One of the theoretical problems of exclusively emphasizing the social conventions and rules that support collaboration is that we can slip into the study of local automata without autonomy. In such theories of collaboration, we do not factor in the agency involved in collaboration. For example, think of the slaves who built the great pyramids. One would say they complied with their masters' rules and the prevailing social conventions (on pain of death), but one would not say they collaborated. Of course, the slaves may have collaborated with one another, for example, by helping to carry a load when they did not have to. But, simply following a social role does not make behavior collaborative, it simply makes it compliant.

We can carry this point further by contrasting the "effort after shared meaning" with "having a shared meaning." Consider a definition of collaborative activity that many of the current authors adopted and that is superficially consonant with the notion of the effort after shared meaning: Collaboration is "a coordinated, synchronous activity that is the result of a continued attempt to construct and maintain a shared conception of a problem" (p. 70, Roschelle & Teasley, 1995). This description strikes me as problematic. It does not capture the essence of collaboration as it appears to the individuals involved; it provides no room for the agency with which people choose to collaborate or not, or choose to make the effort to understand one another or not. It does not capture the sense of compromise and choice that is the hallmark of any collaboration. By this definition, for example, we might have to say that the slaves were collaborating with their bosses as they built the pyramid. After all, they did maintain some shared understanding of the problem -- to build a pyramid.

A Turning Test for Collaboration

I would like to develop another argument for why we may want to factor in an individual's agency in our models of collaboration. I begin by pointing to a second problem with the above definition of collaborative activity. In that definition, collaboration is primarily defined by its outcome, "coordinated and synchronous activity." It is not clear to me what constitutes coordinated activity. Joiner, Issroff, and Demiris state the issue quite nicely, "there are no current universally accepted notions of what constitutes effective or efficient human-human collaboration" (p. xx, this volume). People have their own reasons for collaboration and their own definitions of coordination, and these reasons and definitions often change during the course of an interaction. Unless the researcher is willing to impose particular social norms as to what constitutes a successful collaboration, it seems that a general definition of collaboration should be defined more with an eye towards the view of those people involved rather than the view of the prevailing culture as identified by the researcher. As I argue next, taking the viewpoint of the people involved leads to a definition of collaboration that is grounded in the notion of individual agency and intent.

If we view collaboration as something that involves individuals representing one another's thoughts, then perhaps the most relevant definition of collaboration would be from the eye of the beholder, or representor as it were. So, how would I, as the beholder, decide whether I was interacting with a collaborator? One approach to this question is to consider the traditional artificial intelligence Turing Test to see if we can develop a "Collaborative Turing Test." The original Turning Test was designed to answer the question of how we can decide whether a computer is intelligent. The test stated that if we cannot discriminate between a computer and a person (for example, when interacting over a teletype), then the computer should be considered intelligent. Weizenbaum (1966) made fun of this test with a computer program called ELIZA. ELIZA faked being a therapist. It had a number of generic sentence frames that it would use in response to anything a "patient" typed. For example, if the patient typed, "I'm very unhappy about my mother." ELIZA would respond, "<u>How long have you been</u> 'unhappy about your mother'?" It could just as easily say, "<u>Tell me more about</u> 'your mother'". People were fooled by ELIZA and thought a human was on the other side of the teletype. Weizenbaum concluded that the Turing test was absurd.

Since the original demonstration, there is now a software program based on ELIZA that is doing quite well in the stores. Even though people know that ELIZA is not a person (and not "really" smart), they still choose to use it. Evidently, people are willing to view it as a "smart tool" (Norman, 1993) even if they do not view it as smart. What is the difference between these two -- being a smart tool and being smart? One answer is that people do not view a smart tool as a smart agent that is trying to fulfill its own goals; they simply view it as a clever device.

Perhaps the Turing test should not be whether people can discriminate between computers and humans; one can always manipulate people's short term success at this task anyway. Rather, the test should be whether people are willing to interact with something as an intelligent agent rather than a "smart tool". To borrow from Dennett (1987), maybe the criterion of intelligence is whether people take an 'intentional stance' towards the computer; in their interactions, do they "endow" the computer with an intelligence that they believe the computer uses to fulfill its intentions.

If we map this idea into a Collaborative Turing Test, then a piece of the criterion for identifying a collaborative partner is whether people are willing to view their partner as a collaborative <u>agent</u> during their interactions with it. Do they, for example, assume the computer has collaborative goals that it intends to fill?

People's intentional stance is an important component of collaboration; people need to perceive their partner as a collaborative agent rather than a supportive tool that simply follows or embodies a set of behavioral rules. But still, this definition of collaboration seems empirically toothless as it stands. It provides little insight into collaboration except that it involves the perception of intentional agents with goals. So, I would like to try a slightly different approach. I again begin with the question of intelligence. By some accounts (Eco, 1994), symbolic intelligence is characterized by the ability to lie. Consider work with chimpanzees. Researchers release one chimp into a field where there is a stash of bananas. After finding the bananas, the chimp returns to the band and points to the location of the food for the other chimps to share. So far, there is nothing particularly "intelligent" about this behavior; bees do the same sort of thing. One time, however, legend has it that a chimp came back and pointed in the opposite direction of the food. While the other chimps followed his misdirection, he went back to his now personal stash of food. This misdirection lie depended on the original chimp knowing something about the "thoughts" of the other chimps; namely, they would follow his pointing. This seems distinctly intelligent. The relationship between intelligence and lying makes a good bridge into issues of collaboration and the effort after shared meaning. This is because intentionally lying depends on the ability to construct a representation of another person's (or chimp's) thoughts.

Now let us combine the observation about lying with the importance of the intentional stance to come up with a Collaborative Turing Test: A partner is collaborative if you believe it is possible that the partner could be non-collaborative. In other words, you believe the partner has enough knowledge of your mental states and enough personal agency that it could intentionally thwart you or choose to disengage.

Even though this definition of collaboration is circular and surely not sufficient, I think it does buy us something. Collaboratitve behaviors probably depend a great deal on the extent to which people trust one another's cooperativeness. Will you tell your deep secrets to your therapist, Eliza, if you think she will tell other people against your wishes? Will you be willing to collaborate with another scholar before you have had a chance to assess their intentions and beliefs about intellectual property rights and ownership of joint products? If collaboration depends on representing another person's mental state, then surely one of the mental states we appraise is the other person's goals and likelihood of collaborating at any moment.

Consider the problem of remote collaborations recounted earlier (Hansen et al., this volume). People from different communities need to cooperate electronically on a new project. The authors state that at the beginning of this cooperation, it is important to have a high-band width for exchanging information about constructing joint goals and administrative procedures. An interesting question is whether some of the difficulty people have in early telematic collaborations comes from their difficulty in developing mutual trust of one another's' intentions. I wonder whether early interventions that led to trust among the participants could alleviate some of the early need for a high-band width of information flow. It is amazing how effective a relaxed dinner conversation can be for facilitating future collaborations.

The Collaborative Turing Test also provides some useful analytic distinctions among ideas that frequently become conflated. First, it points out that communication and collaboration are distinct. Imagine that you are in perfect communication with your word processing machine. The computer records every keystroke perfectly. Does that mean the computer is collaborating? Not really, unless you believe it could intentionally put the wrong letter on the screen every now and then. Or, take the converse, imagine that I speak Hindu and you speak Japanese. Even though we would have miserable communications, this does not mean we are not collaborating. I know, for example, that if you get frustrated enough, you might choose to quit. Although communication and the effectiveness of a collaboration are empirically correlated, they are still different things.

Second, the definition helps to clarify that outcomes are not sufficient criteria for identifying collaboration. If basketball team A continually loses to basketball team B, this does not necessarily mean that team A is less collaborative. In fact, one might suspect that team A needs to be much more attentive to collaborative issues because there is a high risk of people becoming uncooperative in a failing group. Finally, I think the Collaborative Turning Test helps to eliminate some of the slippery slope that occurs when we consider embedded systems like individuals in a group, termites in a hive, and neurons in a brain. To put it bluntly, I see no sense in saying that the termites or neurons are collaborating. They have no choice in what they do.

MOTIVATION: FROM APPROPRIATION TO PRODUCTION

The ability to express agency plays an important role in people's motivation and benefit from collaborative learning. First, I think people need to have the intent to learn while interacting in a collaborative group; learning is not automatic (e.g., Bereiter & Scardamalia, 1989). Although there is a place for incidental learning outside of one's intent, I suspect that in many cases it is not enough for people to simply collaborate; they also need to have the intent to learn about the situation over which they are collaborating. There are many examples of where people have scripted cooperation among children, and the children end up producing the script rather than generating the knowledge they are supposed to learn through cooperation (Vye et al., in press). Second, and more profoundly, I think people are motivated to collaborate to the degree they get to exert their agency through productive behavior. First, I will discuss the intent to learn, then I will take up the issue of production and motivation.

The Intent to Learn

A wonderful example of the importance of the intent to learn comes from two different stories involving Eskimos. The first story involves the question of why an ancient European settlement on Iceland disappeared. One hypothesis is that, for whatever reason, the Europeans did not appropriate the knowledge of the Icelandic Eskimos and consequently the Europeans did not survive. This shows that the failure to appropriate can be a big mistake. The second story involves Eskimos and Athabascan Indians in Alaska. In a location called the Middle Yukon, the Indians are separated from the Eskimos by a thin mountain range. In times past, the Indians purportedly crossed this range and sneaked into the Eskimo village to kidnap women (McClellan, 1971). They did not appropriate the Eskimo women because there were not enough Indian women. Rather, they took the women because the Eskimo's had developed excellent technologies. The Indians appropriated the women to gain access to their technological know-how. Interestingly, the Indians had to do this often. When a woman died, they would kidnap another. This is because they never bothered to learn what the women knew; they simply appropriated the technology not the understanding.

I like these examples because they highlight that there is a difference between a culture of appropriation and a culture of learning. I think it complements the observation made by Baker, Hansen, Joiner, and Traum that "learning seems more likely to occur to the extent that agents expend greater cognitive effort towards mutual understanding than that which would be minimally required for communication" (p. xx, this volume). Evidently, the Indians were happy to communicate with their Eskimo women and did not put in the intentional effort needed to learn.

Putting Production Back into the Dialectic with Appropriation

What is it that causes people to put forth that extra cognitive effort towards mutual understanding? There are many things ranging from potential rewards to the desire for a friendly chat to the environmental affordances that support communication. In terms of basic motivations I would like to illuminate the importance of production and original contribution. This appears to have been ignored lately, perhaps because of our overemphasis on the appropriation of cultural practices. Appropriation is clearly an important idea, but it is only half of the story. Marx (1939/1973) spoke of two great forces that help constitute a person. One was appropriation -- we become what we are by appropriating the ideas and artifacts of those around us. Alienation, his construct of psycho-social malaise, can occur when we are not allowed to appropriate the contributions of others. But, Marx did not consider appropriation the "key" quality of what it means to be human. Instead, he felt that humans are quintessentially builders. We want to produce and create ourselves in the world through our ideas and our material products. This way we may put our "element" in the social matrix, and other people may appropriate our ideas. At the same time, we may "reappropriate" our creations as they have been realized in the world, culture, and others. This serves as feedback about ourselves, our learning, and our environment. Without production, there is no feedback.

For Marx, the key to a complete person was not simply access to the material and intellectual wealth of a society, but also access to the means of production. Marx did not advocate a welfare state in which people only had access to appropriation. He advocated a productive state where people could contribute and impress themselves upon the world. Individuals are builders of their society, not simply recipients. For Marx, the key issue was always who had the means of production.

The emphasis on individual production is an aspect of Marx that appears to have faded somewhat in the communitarian tradition that comes through Vygotsky and Engles (e.g., Illyenkov, 1977). In Vygotsky, we read about the movement from external to internal, but we less often hear of the movement from internal to external. Although communitarian scholars speak of the importance of "activity," it still seems that the individual's productivity falls into the background. The emphasis on appropriation at the expense of production strikes me as particularly problematic when we consider collaborative learning.

Consider the World Wide Web. What is it that makes for a successful website? One thing is the opportunity for people to contribute. Hagel and Armstrong (1997) argue that successful websites are those that enable virtual communities. Whereas Prodigy emphasizes news broadcasting and information delivery, virtual communities support contributions by those individuals who use the website. Like the old adage, "The best conversationalist is a great listener," the most inviting website is the one where people can contribute. *Amazon.com* is a very successful on-line bookstore, in part, because it allows individuals to submit their own reviews of books. People contribute and they look forward to the responses. This is one way that we come to learn, by creating ourselves and reappropriating the feedback from our creation.

As another example, consider what makes you most content in a conversation. Is it when you have been told something and understand it; when you have appropriated someone else's idea? Or, is it when you have contributed substantially to the conversation, when you have produced ideas that move other people and that help contribute to the direction of the interaction? Stated less rhetorically, perhaps the most irritatingly uncooperative agent is really the one who denies you agency within the group.

I think there are some wonderful research opportunities here. For example, one might use a simple version of Csikszentmihalyi's (1990) method of measuring flow. Every now and then, interrupt people during a collaborative episode and ask them to rate their "intellectual energy for learning" on a scale from one to ten (Schwartz & Bransford, 1996). I predict that on average people will give much higher ratings when they are contributing their ideas to a receptive group than when they are listening to somebody else's ideas. Moreover, I suspect that when people feel they are contributors, they will be much more willing to go beyond "the minimum necessary" to communicate and complete a task.

The dialectic between production and appropriation suggests that the opportunity to produce should influence people's willingness to appropriate from those with whom they are collaborating. One nice example comes from a study by Kay Burgess (personal communication). Burgess worked with 5th-grade students who learned how to solve a complex problem about rescuing an eagle. Later those students came in as consultants to help college students solve the same problem. The 5th-grade students were empowered and highly engaged as they explained some of the intricacies to the college students. And, importantly, they became aware of and started to appropriate the college students' behaviors such as their diligence. Similarly, my colleagues at the Learning Technology Center at Vanderbilt have found that teachers become more involved in learning about new instructional techniques after they have had an opportunity to present their versions of those techniques to a new group of teachers. A study that formally examined the dialectical benefits of production on subsequent appropriation would make a profound contribution to the literature on collaborative learning.

LEARNING: FROM ASSIMILATION TO CONSTRUCTION

One of the reasons that production is so important to collaborative learning is that learning itself is productive. People construct their knowledge through generative mental and physical activities. They do not simply assimilate someone else's knowledge or practices; they actively produce their understanding. The constructive nature of learning has implications for how people learn, how they come to understand one another, and what they are likely to learn in groups. In particular, there are a set of implications surrounding language that are especially relevant to collaborative learning because collaborative learning typically involves linguistically mediated communication.

Some Background on the Relationship between Language and Learning

As a starting point for understanding the implications of linguistic communication, we can begin with the chapter by Mephu-Nguifo, Baker, and Dillenbourg (this volume). These authors explicitly compare machine learning operators and dialog operators. They point out that there is a similarity between the two because both require the assumption of a common language. By common language, they mean that there is no ambiguity in reference. If you are told 'John is a cow', and you know the referents of the terms 'John,' and 'cow,' then you can assimilate 'John is a cow' into your knowledge base. The work by Hoppe and Ploetzner (this volume) helps clarify the point further. They describe a computer simulation in which one agent has a qualitative representation of a physics problem and another agent has a quantitative representation. Although the two computer models have different conceptual representations, they are able to communicate because their references to the problem at hand are unambiguous.

Weiss and Dillenbourg (this volume) point out that systems of distributed artificial intelligence depend on perfect symbol-referent mapping. They then argue that the necessary absence of referential ambiguity in these systems makes them incapable of modelling an important dynamic of collaborative learning; the systems cannot negotiate meaning. In other words, when the assumption of a common language (perfect referent-symbol mapping) is violated, computers cannot recover and learn from one another in the process. On the one hand, this appears to be a limitation of symbolic computer models; they require a common language to collaborate and exchange information. On the other hand, this also appears to be a strong limitation in humans as well. Baker, Hansen, Joiner, and Traum (this volume) point out that when people cannot assume that their expressions are understood by their partners, there needs to be an effort dedicated to re-constructing a basis for mutual knowledge. Often times, this effort falls short, just like in computers. But, as both sets of authors point out, it is in the attempt to rectify misunderstandings where a great deal of learning can occur.

Language plays an important role in the attempt to establish understanding among partners. To explore what role language plays, particularly with respect to learning, consider the discussion started by Ploetzner, Dillenbourg, Preier, and Traum (this volume). In their chapter, they compare the effects of self-explanations to the effects of receiving and giving explanations to others. By looking at this clean contrast, they hoped to identify essential differences between working alone and together. They do not find many differences, although they smartly speculate that feedback should be a key difference between self-explaining and other explaining (despite the fact that studies have not investigated this possibility). In addition to the lack of feedback, a likely explanation of these null results may have to do with the heavy use of language in both settings. Language and linguistic representations play a particular role in generating new knowledge and may be a primary mediator of learning effects, whether alone or in collaboration. To develop the point, I first explore the sorts of things language can communicate when people try to learn with or about one another.

The Implications of Constructivism for Reaching Shared Meaning

When someone says something to me, I do not simply assimilate or copy that expression into my mental network. People do not learn a text, they learn from a text (Bransford & Nitsch, 1978). The words that you express can serve as a starting point for me to construct or generate my own knowledge. Even if we come from the same culture, there is a distinct possibility that I will come up with something very different than what you had in mind when you uttered your expression. Consider, for example, the following newspaper headlines that were collected over the past year:

- Drunk Gets Nine Months in Violin Case
- Survivor of Siamese Twins Joins Parents
- Iraqi Head Seeks Arms
- New Study of Obesity Looks for Larger Test Group
- Kids Make Nutritious Snacks
- Miners Refuse to Work after Death

I doubt the newspaper editors had the amusing alternatives in mind. The sentences show that people do not simply assimilate the language of others. They actively generate meanings, meanings that may be quite different from what the speaker had in mind.

Although language is notorious in this regard, it is important to remember that "perceptual things" suffer the same fate. They are not simply assimilated like photographs in the head. People actively generate understanding using the input of the physical world just like they do from the verbal world. If you and I look at the same thing, there is no guarantee that we will see the same thing. The mere presence of the physical world does not ensure a common ground between two people who both have access to that world.

Consider, for example, the squiggles in Figure 1. Imagine that your task is to memorize them for an upcoming recognition test. Also imagine that the artist who drew the squiggles is at your side. The artist would probably see something very different than you do. Fortunately, breakdowns in shared meaning are not irreparable. In the current case, to improve your memory, the artist could probably help you see what he sees with a little bit of supportive language. For example, "Turn the figure clockwise 90° and match each of the following labels to its respective squiggle: James Dean, Babyface, St. Nick, Baseball Bob." Let me reiterate this point more telegraphically. Hoppe and Ploetzner (this volume) describe a system that uses a joint blackboard as a source of common ground for conversants. They describe the system as WYSIWIS -- What You See Is What I See. Although not as easily said, I think the system is more appropriately called, WYASIWIAS -- What You Are Shown Is What I Am Shown.

Figure 1 about here -- squiggles

Understanding is generated and constructed. The implication of this is that physical reference suffers the same fate as words; neither guarantees common meaning and learning across individuals. Even so, despite this similarity, there are very real differences between words and objects. Language helps in constructing a particular kind of knowledge. Language and other symbolic representations like mathematics are very good at helping people to build and evaluate an articulable structure. Moreover, there are criteria of understanding that come with symbolic explanations. Ploetzner, Dillenbourg, Preier, and Traum state, "Explanation is a social criterion for operationalizing what is accepted as understanding" (p. xx, this volume). I would like to amend this just a little to state, "what is accepted as <u>structured</u> understanding." Most people, for example, have a very good understanding of how to ride a bicycle although they cannot explain it.

To demonstrate the power of language with regards to structure, consider Figure 2a. What do you see? People see many different things including a plane flying sideways and two olives on a toothpick. But let us assume that it is two men riding a tandem bicycle wearing large, round hats. You are looking down at them. Next, consider Figure 2b. What do you see? If you are like most people, you see two pairs of men riding bicycles. Now, notice how language helps re-structure your thought. It is not really men on bicycles; it is a bear cub clinging to the back side of a tree. The circles are its paws. Suddenly, new structures in the referent become important, like the distance between the two lines that portray the tree.

Figure 2 about here -- Circles and Lines

The Type of Knowledge People Often Generate in Collaborative Learning

Collaboration typically involves heavy doses of language. This language should lead cooperative groups towards structural descriptions in their learning. Moreover, groups may move towards abstractions as the members try to find a safe place to communicate where their idiosyncratic differences of interpretation will not get in the way. It is this pull towards abstraction and structure in the verbal communication of groups that strikes me as the sort of place to find a special effect of collaboration on cognitive outcomes.

Consider the following study (Schwartz, 1995). Seventh grade students were given several descriptions of fictitious fish and their habitat requirements. For example, "the Spotted Frolling lives in lakes with weeds," "the One-Finned Halluck needs weeds and a sandy bottom." Their task was to construct a visualization of the various relationships. Students worked alone or in pairs. The students who worked alone drew pictures of lakes with fish in them. Only 6% created visualizations that were abstract in the sense that they did not actually look like fish and lakes. In contrast, 67% of the pairs constructed an abstract representation like a matrix or chart. This percentage is well-above the probability that a pair would have included at least one member who would have constructed an abstract representation working alone. In other words, the collaboration among the pair members led them to generate something new that was not found in otherwise similar individuals.

Another study, completed for this chapter with Doug Holton, demonstrates the point again, only in a negative fashion. The basic task involves glasses filled with imaginary water. Imagine that there are two glasses of the same height that are filled to equivalent levels of imaginary water. The only difference is that one glass is thin and one is wide. Would they start pouring at the same or different angles? Figure 3 provides a two-dimensional version of the problem. In prior research, Schwartz and Hegarty

(1996) found that only 20% of individuals make the correct explicit judgment that the thin glass needs to be tilted further than the wide one. When solving the problem explicitly, many people abstractly compare quantities like glass width, and this abstract approach leads them to faulty conclusions. However, when otherwise identical people close their eyes and tilt each empty glass in turn until they "see" the imagined water reach the rim, 100% correctly tilt a narrow glass further than a wide one. The question addressed by the following study was what would happen when people worked together to solve the problem. One possibility is that one member of the pair would tilt each glass, and the other member would observe the different angles of tilt. Another possibility is that, because of the pull towards abstraction in group communications, the pairs would rely on more discrete and explicit reasoning and would therefore get the problem wrong.

Figure 3 about here -- two glasses side by side

Twelve pairs tried to solve the task. They were given the wide and thin glasses and were asked to figure out whether the glasses would start pouring (imaginary) water at the same or different angles. Interestingly, 0% of the pairs correctly answered the problem. The pairs did not imagine the behavior of the water which is how people usually reach the correct answer. Instead, they thought and spoke about the problem in terms of static structures and discrete features like width to height ratios. These features do not easily offer a solution and led the participants to the wrong answer. Evidently, there are times when the verbal exchanges in a group lead to productive outcomes, and there are times when they do not.

By considering the type of understanding that people generate with the aid of language, it should be possible to prescribe particularly appropriate times to share linguistic representations during knowledge growth. For example, John Bransford and I have examined whether there is a "time for telling" (Schwartz & Bransford, in press). Is there a way to prepare people to be told something; are there times when the structured knowledge of verbal communication is particularly beneficial?

In one experiment, we tried to teach four target concepts from cognitive psychology (e.g., people tend to remember stereotypical events). Students were separated into three instructional treatments. In the "double telling" treatment students wrote a three page summary of a brief book chapter that described the target concepts and the experiments that exemplified them. Five days later, they heard a lecture on the concepts and experiments. In the "double discovery" treatment, students analyzed and looked for patterns in simplified data sets from experiments that exemplified the four concepts. Table 1 provides one sample of what they had to analyze. Five days later they completed the analysis task again. In the final condition, "discovery and telling," the students analyzed the raw data and then five days later heard the same lecture as the "double telling" students. A week after completing the instruction, all the students tried to predict the outcomes from an hypothetical study. All the target concepts were applicable to the hypothetical study. The relevant experimental question is which treatment made the most correct predictions.

Table 1 -- one example of data the students analyzed

The results were definitive. Students in the "discovery and telling" condition made over twice as many predictions as the students in both the "double discovery" and "double telling" conditions. One interpretation of this result is that the discovery activity helped students discern specific features that differentiated classes of psychological phenomena, much as a botanist can distinguish sub-species of a given flower. For example, by analyzing Table 1, they differentiated events that have high and low frequencies in people's recounting of a doctor visit. After noticing this and other distinctions, the subsequent lecture provided the coherent and abstract structure that enabled these students to construct an understanding of why these differences are significant. For example, the lecture explained how schema theory predicts these types of stereotypy effects in memory.

The results of the preceding study shows that individuals needed both forms of knowledge. Without the "discovery," the "telling" simply provided a set of facts to be memorized. And, without the "telling," the differentiations were simply observations. This result provides an important lesson for those who believe that direct teaching (e.g., a lecture) is contrary to constructivist ideals. As argued above, people construct their knowledge regardless of whether the input comes from the physical or linguistic world. The current study points out that there is a place for texts and lectures in a classroom; namely, when students have sufficiently differentiated domain knowledge to use the expository materials in a <u>constructive</u> manner.

In combination, the three studies indicate that it is useful to entertain which types of understanding people are most likely to construct under different learning conditions and how those types of understanding may complement one another. In particular, it appears that linguistic communications, whether within a collaborative group or a classroom lecture, lead to abstract and structured understandings. How well this hypothesis will hold up under further empirical scrutiny is an open question. Nonetheless, it seems important for researchers to seriously consider the type of knowledge that is most likely for people to construct in group interaction.

COMPUTER MODELLING REVISITED

Ploetzner, Dillenbourg, Preier, and Traum state, “Even if an agent models the other sufficiently to continue the dialogue, this might still be merely at a shallow level of understanding, not enough to actually learn” (p. xx, this volume). The question then is how do we get people to move from shallow interactions to the deep interactions that can yield the special learning benefits of working in a group. As pointed out in the beginning of the chapter, little empirical research has yielded much head way on this problem. Thus, I have been trying to reconsider what is unique about the psychology of collaborative learning. To accomplish this task, I have worked before the backdrop of computer models of collaborative agents. Now, I would like to bring this backdrop into the foreground.

In the preceding sections, I have been laying out an agenda for computer models of collaborative learning by constructing a picture of what is essential to collaboration. I have argued that the computers must have the <u>agency</u> to choose to collaborate or not. Once they have that agency, they need <u>motivations</u> that determine whether they interact or not. Finally, they need to generate <u>learning</u> products that are likely to occur once individuals choose to cooperate. In the next three paragraphs I briefly suggest some of the ways these ideas might be examined computationally. I am way beyond my depth here, so I may suggest things which have already been examined. I apologize in advance.

The way I have framed the issue of agency is by pitting choice versus social compliance. I think it would be very interesting to conduct simulations where one changes the balance or force of social rules and the individual agents’ goals. One might program the social rules as highly rigid as in the case of physical laws. Or, one might make them fairly soft, perhaps allowing agents to rebel against “society.” At the same time, one might manipulate an agent’s desire to maintain personal goals and to trust other agents. I wonder how changes in these various parameters would affect the interactions between the desire to comply, the desire to collaborate, and the desire to reach personal goals in a collaborative task. Artificial intelligence work in game playing seems particularly useful in this regard because it has worked out methods for representing short-term and long-term personal goals, the goals of others, and the rules of the game. Perhaps it would be possible to construct a game playing program where there are multiple agents playing multiple games each trying to attain some form of points. The agents get to decide whether to “break the rules” in any given game, and they get to decide which agents they choose to play against and which agents they team up with. This would be a nice context for exploring the interplay among people’s representations of other people’s collaborativeness and their behavior with respect to social rules.

I framed the issue of motivation in terms of appropriation versus production. Moving this to the computational arena, one might create interesting simulations by developing “I will continue to collaborate” (ICC) heuristics. The ICC heuristics might take into account the degree to which the agent is appropriating another agent’s knowledge, the degree to which the agent is producing knowledge, and the degree to which the agent perceives its productions as having been appropriated by another agent. By changing weights within the heuristics, it would be interesting to see if we could model outcomes from human interactions. Another possibility is that one might model what happens when one agent recognizes its ideas in another agent (i.e., when its ideas have been appropriated). For example, if agent A recognizes its “knowledge” in agent B, it may improve communication because agent A can rely on that knowledge as a common ground. Moreover, with some additional ICC heuristics agent A might be more likely to communicate with agent B because it shows evidence of trying to make a shared meaning.

Finally, I framed the issue of learning in terms of assimilation versus construction. This is a more difficult problem for computer models, because they often rely on assimilation in which the propositions of one agent are directly inserted into the knowledge base of another agent. This problem could be avoided, however, by using inductive algorithms that treat propositional input as data from which to construct an understanding. More interesting in the current context, however, is the possibility of given computer models different options for how they resolve breakdowns in communication. For example, will they try to find a physical referent so they can point to it, or will they generate a more abstract representation so they can communicate the structure of the idea they are having trouble communicating. More generally, I think Hoppe and Ploetzner (this volume) are on the right track when they look at multiple representations in the context of a qualitative and a quantitative agent that try to communicate their knowledge. I wonder what would happen if they added a few more agents into the mix and gave them some ICC heuristics for choosing collaborators. Would quantitative agents end up associating with quantitative agents, and qualitative agents with qualitative agents?

One reason these considerations about collaboration are important, even as half formed as they are, is that they raise the issue about the theoretical level at which we plan to operate. One approach would be to create computer models (or theories) that operate at the level of collaborative rules or roles. For example, one might manipulate the “collaborative rule sets” that different robots use and see what happens. These rules might take the form of Gricean maxims, “only communicate as much information as is necessary.” Or, they might take a more directly functional form, “Exchange information if it will optimize the gathering of food.” This would make for an interesting set of simulations, but they would be simulations about social conventions not about individuals in collaboration. Gricean maxims are only interesting because they can be broken. Moreover, as pointed out at the outset, Weiss and Dillenbourg (this volume) suggest that it may be difficult to pre-specify all the needed rules because collaborative situations are often too complex; the agents need to be able to self-improve in their collaborative behaviors. This presumably requires a lower order layer that can generate cooperative behaviors. I have been trying to propose some psychological elements of this lower layer. Ideally, what we want is a theory that explains how collaboration and collaborative learning emerge, not simply how they look once they appear. I believe analysis at this lower layer may best illuminate what is unique about people's effort to achieve shared meaning and deep understanding in small groups.

REFERENCES

Barron, B. J. (1992). Collaborative problem solving: Is team performance greater than what is expected from the most competent member? Unpublished dissertation, Vanderbilt University, Nashville, TN.

Bereiter, C., & Scardamalia, M. (1989). Intentional learning as a goal of instruction. In L. B. Resnick (Eds.), Knowing learning, and instruction: Essays in honor of Robert Glaser (pp. 361-392). Hillsdale, NJ: LEA.

Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. Journal of Verbal Learning and Verbal Behavior, 11, 717-726.

Bransford, J. D., & Nitsch, K. E. (1978). Coming to understand things we could not previously understand. In J. F. Kavanagh and W. Strange (Eds.), Speech and language in the laboratory, school and clinic. Cambridge, MA: MIT Press.

Bartlett, F. C. (1932). Remembering: A study in experimental psychology. Cambridge: Cambridge University Press.

Csikszentmihalyi, M. (1990). Flow: The psychology of optimal experience. NY, NY: Harper Row.

Dennett, D. C. (1987). The intentional stance. Cambridge, MA: MIT Press.

Dunnette, M. D., Campbell, J. & Jaastad, K. (1963). The effect of group participation on brain-storming effectiveness for two industrial samples. Journal of Applied Psychology, 47, 30-37.

Eco, U. (1984). Semiotics and the philosophy of language. Bloomington: Indiana University Press.

Fox, D. J., & Lorge, I. (1962). The relative quality of decisions written by individuals and by groups as the available time for problem solving is increased. Journal of Social Psychology, 57, 227-242.

Gibson, E. (1969). Principles of perceptual learning and development. NY: Meredith Corporation.

Hagel, J., & Armstrong, A. G. (1997). Net gain: Expanding markets through virtual communities. Boston, MA: Harvard Business School Press.

Hastie, R. (1983). Experimental evidence on group accuracy. In B. Grofman & G. Owen (Eds.), Information pooling and group decision making (pp. 129-157). Greenwich, CT: JAI Press.

Hill, G. W. (1982). Group versus individual performance: Are N + 1 heads better than one? Psychological Bulletin, 91, 517-539.

Hoover-Dempsey, K. V., & Jones, K. P.(April, 1996). Parental perceptions of appropriate roles and role activities in children's education. Paper presented at the International Roundtable on Families, Schools, Communities and Children's Learning, New York.

Ilyenkov, E. V. (1977). The concept of the ideal. In Philosophy in the USSR: Problems of dialectical materialism. Moscow: Progress.

Kelly, H. H. & Thibaut, J. W. (1969). Group problem solving. In G. Lindzey & E. Aronson (Eds.), The handbook of social psychology, Second edition (pp. 1- 101). Reading, MA: Addison-Wesley.

Krauss, R. M., & Glucksberg, S. (1969). The development of communication: Competence as a function of age. Child Development, 40, 255-266.

Laughlin, P. R. & Futoran, G. C. (1985). Collective induction: Social combination and sequential transition. Journal of Personality and Social Psychology, 48, 608-613.

Marx, K. (1939/1973). Grundrisse. Martin Nicolaus, Trns. New York, NY: Random House.

McClellan, C. (1971). Feuding and warfare among Northwestern Athapaskans. In A. McFadyen Clark (Ed.), Proceedings: Northern Athapaskan Conference. National Museum of Man Mercury Series Canadian Ethnology Service, Paper No. 27, Vol. 1.

Norman, D. A. (1993). Things that make us smart: Defending human attributes in the age of the machine. New York: Addison-Wesley.

Robertson, S. P., Zachary, W., & Black, J. B. (Eds.). (1990). Cognition, computing, and cooperation. Norwood, NJ: Ablex.

Roschelle, J. & Teasley, S. D. (1995). The construction of shared knowledge in collaborative problem solving. In C. E. O'Malley (Ed.), Computer-supported collaborative learning (pp. 69-97). Berlin: Springer-Verlag.

Schwartz, D. L. (1995). The emergence of abstract representations in dyad problem solving. Journal of the Learning Sciences, 4, 321-354.

Schwartz, D. L. & Bransford, J. D. (in press). A time for telling. Cognition & Instruction.

Schwartz, D. L. & Bransford, J. D. (1996) Experiences that energize. Unpublished manuscript. Nashville, TN: Vanderbilt University.

Schwartz, D. L., & Hegarty, M. (1996). Coordinating multiple representations for reasoning about mechanical devices. In P. Olivier (Ed.), Cognitive and Computational Models of Spatial Representation: AAAI-96 Spring Symposium Series (pp. 101-109). Palo Alto, CA: AAAI Press.

Slavin, R. E. (1983). Cooperative learning. New York: Longman.

Turro, I., Porter, R. H., & Picard, M. (1994). Olfactory cues mediate food selection by young chicks. Physiology & Behavior, 55, 761-767.

Weizenbaum, J. (1966). "ELIZA -- A computer program for the study of natural language communication between man and machine," Communications of the ACM, 9, 36-44.

Vye, N. J., Schwartz, D. L., Bransford, J. D., Barron, B. J., Zech, L. and Cognition and Technology Group at Vanderbilt. (in press). SMART environments that support monitoring, reflection, and revision. In D. Hacker, J. Dunlosky, & A. Graesser (Eds.), Metacognition in Educational Theory and Practice. Mahwah, NJ: Erlbaum.

ACKNOWLEDGMENTS

I would like to thank John Bransford for the many insights he contributed directly and indirectly. I would also like to thank Pierre Dillenbourg for the opportunity to participate and Trefor Davies for his clever art work in Figure 1.

FIGURES

<u>Figure 1.</u> Do you see the same thing as the original artist? (Adapted from Gibson, 1969.)

<u>Figure 2.</u> What do you see in the two figures?

<u>Figure 3.</u> The shaded area represents water in the two glasses. If they are tilted, do they start pouring at the same or different angles?

Table 1. An example of the tasks that students completed to discover patterns in people's memory (adapted from Schwartz & Bransford, in press).

Researchers asked five people to write down the events that occur when they visit the doctor. The results are shown below. Analyze the data to discover the important patterns. Make a graph that shows those patterns.

Person 1: Enter office. Check in with receptionist. Sit down. Wait. Name called. Enter exam room. Sit on table. Doctor examines. Doctor asks questions. Make another appointment. Leave office.

Person 2: Check in with receptionist. Read magazine. Look at other people. Name called. Sit on table. Nurse tests. Doctor examines. Leave office.

Person 3: Check in with receptionist. Sit down. Read magazine. Talk to nurse. Nurse tests. Talk to doctor about problem. Leave office.

Person 4: Enter office. Sit down. Read magazine. Enter exam room. Undress. Sit on table. Nurse tests. Doctor examines. Get dressed. Leave office.

Person 5: Enter office. Check in with receptionist. Sit down. Read magazine. Name called. Follow nurse. Enter exam room. Nurse tests. Doctor enters. Doctor examines.

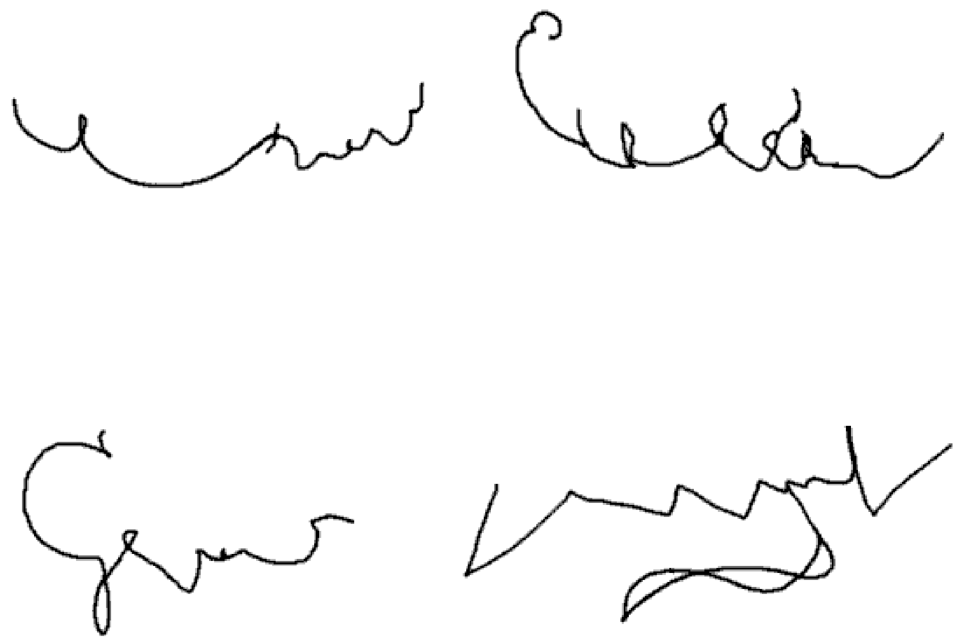

Figure 1.

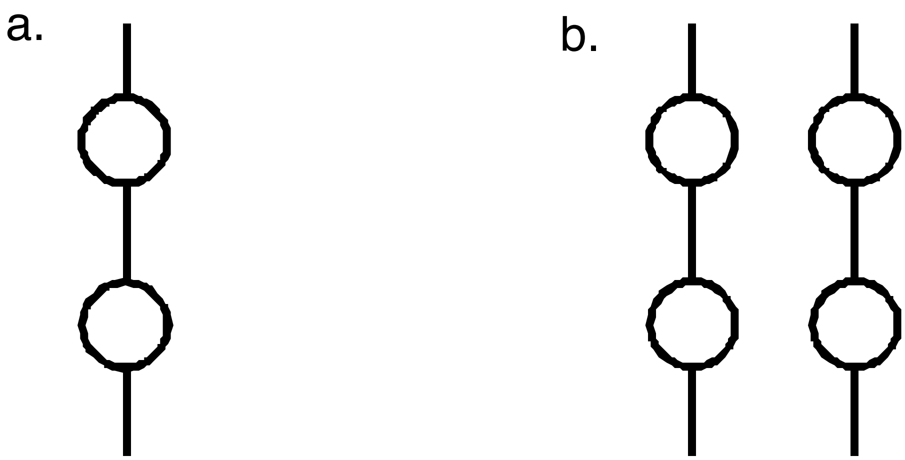

Figure 2.

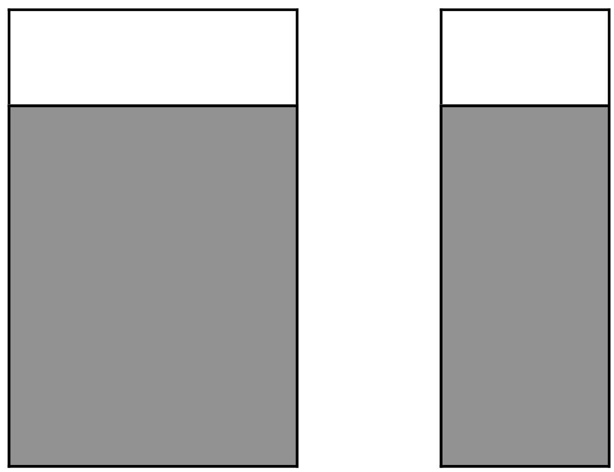

Figure 3.