CONCEPTUALIZING GROUP LEARNING

Paul S. Goodman

Brandi Martin Pearce

Jeanne M. Wilson

Graduate School of Industrial Administration

Carnegie Mellon University

412-268-2288
412-268-7357 (fax)
pg14@andrew.cmu.edu

I. INTRODUCTION

This paper is about the concept of group learning. The specific focus is on three questions:

- How can we sharpen our definitions and conceptualizations of group learning?
- How does context change our models of group learning?
- How can we better understand the mechanisms of group learning?

We approach these questions by first delineating some current research on group learning, then indicating gaps in current thinking about group learning and finally by suggesting possible new approaches.

There are three main reasons for exploring these questions. First, groups have become an important unit of analysis in organizational research and practice. Over the past 20 years, groups have emerged as an important building block of organizational effectiveness. Therefore, understanding more about groups should improve organizational functioning.

Second, the focus on group processes, such as group learning, is important because most of the research on groups has an "outcome orientation." A scan of many reviews shows that most empirical studies focus on how variables such as size, composition, and identity affect group effectiveness indicators (e.g., productivity, satisfaction). Although outcome-oriented research is important, we also need to understand basic group processes in their own right. What does group level vs. individual level learning mean? What mechanisms drive group learning? In addition to understanding processes in their own right, we also need to understand them because research on processes can inform outcome-oriented research.

Third, over the last five years a number of models and studies about group learning have emerged. The research has occurred in different settings (e.g., hospitals, students in a laboratory, the military) and has used different theoretical orientations and methodological approaches. Given the emergence of a mass of studies on group learning, it is timely to look at the current direction of this research and suggest future directions.

This working paper also explores how context affects models of group learning. One position we develop is that many models of group learning are built around traditional groups, and as the context of groups changes, the models also need to change. Lastly, we explore some of the mechanisms that drive learning. As with any developing area of research, in many cases these mechanisms are not clearly specified.

II. GAPS IN CURRENT DEFINITIONS / CONCEPTUALIZATIONS

Although the empirical research on group learning has grown tremendously (Argote, Gruenfeld & Naquin, 2001), the definition of the construct has varied considerably across studies. The goal of this section is to critically examine the construct of group learning. Our strategy will be to analyze three representative studies, focusing on four basic questions:

1. How have the researchers conceptually defined learning?
2. How have the researchers operationalized the construct?

3. What level of analysis is used?
4. How is learning differentiated from related concepts?

Ultimately, our objective in examining these four questions is to expand our current understanding of group learning and to specify the features that are necessary in a definition of group learning. Several existing studies could be considered in this review; however, in the interest of space, we are only going to focus on three (Gruenfeld, Matorana & Fan, 2000; Edmundson, Bohmer & Pisano, 2000; Liang, Moreland & Argote, 1995). We selected these because they are exemplar studies.

Moreland, Argote, and their colleagues have contributed significantly to our understanding of group learning through their work on transactive memory systems (Moreland, Argote & Krishnan, 1996). Conceptually, this line of research examines how knowledge possessed by individual group members is combined to develop a shared awareness of who knows what. The group's objective is to learn who knows what about a complex routine task, who can be trusted, and how the group coordinates its work along with learning about the features of the task. This information is ultimately stored in the individual memories of participating members. The researchers suggest that training may serve as one mechanism by which groups develop transactive memory systems, because it provides a forum for members to develop a shared experience through observation. Transactive memory systems may also develop informally, as groups spend time together. Transactive memory systems are operationally defined by indicators of memory differentiation: group members specializing in different aspects of the task, task coordination and task credibility, and members' trust of one another's knowledge. These factors have been measured at several levels. Memory differentiation and task credibility are individual measures of behavior within the group whereas task coordination examines the group's behavior as a whole. Ultimately, learning who-knows-what is assessed by evaluating the group's performance on three criteria: its ability to collectively recall the information learned from training, its speed of task completion, and the number of errors the group makes.

Another line of research that has contributed to our understanding of group learning is the work of Edmonson (1996). Her research focuses our attention on the antecedents of group learning, such as the degree of psychological safety and other organizational factors. In one recent study, she and her colleagues discuss the conditions in which operating room teams that use well-learned routines and have clearly identifiable roles adapt to new ways of working together (Edmonson, Bohmer & Pisano. 2000). In this domain, "learning" involves both group-level learning as members work together over time and organizational-level learning, in which routines and each member's role are highly prescribed allowing hospital workers to be easily substituted for one another. Edmonson et al. (2000) identify two stages of learning: early learning and new learning. Conceptually, early learning refers to a group's ability to learn a specific process or procedure, which involves the learning of new tasks, roles, and relationships. In their operating room study, early learning occurs through a training program attended by all members of a surgical team in order to learn the specifics of the procedure. New learning occurs when groups discover discrepancies that lead to changes reflecting their innovations. These discrepancies may involve a gap between what one knows and what one is able to do, and what one <u>wants</u> to know or be able to do. A perceived error may

serve as another type of discrepancy that invokes the learning process. Operationally, learning is facilitated through communication regarding necessary changes, which are then tested providing additional feedback, which may lead to further discrepancies, initiating the cycle again.

At the organizational level, early learnings are stored in roles and procedures. It appears that subsequent new learnings at both the group and organizational levels are stored in the memory banks of each individual, but it is unclear whether other facilities for storage, such as changes to procedure manuals, are used.

An additional stream of research on group learning has focused on the generation and transfer of new information. Gruenfeld and her colleagues recently examined how temporary changes in group membership affect what groups learn (Gruenfeld, Martorania & Fan, 2000). Conceptually, the researchers suggest that group learning is the transfer of knowledge from an individual to the group through direct or indirect influences. They also suggest that a change in membership can initiate the learning process by providing individuals with access to new information. Operationally, learning occurs when there is a change in the group output based on a change in membership. In this case, the object of learning was the generation and utilization of new ideas, which eventually appeared at the group level of analysis through the joint output of a group essay.

What else needs to be considered? Let's look at these three conceptualizations in the light of the questions: How is learning conceptualized? How is it operationalized? What is the level of analysis? And how is learning differentiated from related concepts? Our goal is to identify new theoretical areas for development rather than to evaluate particular studies.

Conceptual Definitions

Group learning is defined in a variety of ways. From all of the research programs discussed above, we can extract a general definition of group learning (i.e., knowing who- knows-what in the transactive memory studies, or the transfer of knowledge from the newcomer to the group in the Gruenfeld studies). However, it is not clear whether these definitions represent the whole construct of group learning or some subset of the construct space. We think it would be helpful to consider the basic elements of a definition of group learning.

Let's take the feature of "sharing," implicit in most conceptualizations of group learning, as an example. In the Gruenfeld study (2000), the inclusion of a unique idea into a group essay does not necessarily ensure that all the members of the group agree or share the idea. It is possible that the individual author has decided to include the idea without the agreement of the others, or that a strong leader is exerting influence over what information is included. This raises some important questions about group learning. Is it enough for members of the group to have similar information if they have no shared understanding that they hold similar information? How do we characterize a scenario where group members hold diverse perspectives but have a shared awareness that they hold different views? How many individuals need to hold a shared awareness of others' information before we can say there is group sharing? Does the degree of shared awareness affect the potential strength of group learning? How does sharing evolve over

time, and what effect does this have on group learning? Figure 1 represents three very different pictures of sharing.

Although sharing is an essential feature of group learning, we suggest that three additional characteristics have not yet been adequately discussed: The group's ability to store, access, and revise information shared through member interactions. Psychologists suggest that memory is an important component of individual learning because it facilitates recall and revision. We believe the same is true for group learning, where storage is facilitated through mechanisms other than personal memory. For example, at the group level, information can be stored formally via roles, procedures, protocols, computer systems, email or other text-based communication vehicles. Another important feature of learning is the ability to access stored information. Although information may be stored for the group, its ability to act on the information is limited if only one individual and not all of the members can access it. In addition to sharing, storing, and accessing information, it is also important that the group is able to dynamically revise the data that are available to them, thus allowing learning to continue over time.

Operational Definitions

Although formal definitions of group learning have been appropriately broad, the operationalization of group learning has sampled a relatively limited sphere of learning behavior. In all three studies, learning focuses on how individuals interact with one another to acquire, share, and combine information. The studies tend to focus on simple, mechanistic forms of learning (Mowrer & Klein, 2001), facilitated through trial and error, training, communication, observation and changes in membership. For example, the object of learning in the transactive memory studies and the studies of operating room teams and military crews seems to focus on improving coordination through the learning of roles, tasks, and expertise. Of course, this only represents one form of learning. Other objects of learning can be competitive strategies (in the case of top management teams), new discoveries (which might be particularly important to research and development teams), cause-and-effect relationships, or even "double-loop" learning in which teams learn to question their fundamental assumptions and processes (Argyris, 1982). Current operationalizations also exclude latent learning, which may occur from prior exposure, as well as insight-based learning. How groups develop these more complex types of learning is likely to differ considerably from the processes they use to learn how to coordinate assigned physical tasks. To fully understand the domain of group learning, we need to sample some of these other forms of learning in groups.

Level of Analysis and Measurement

We must be able to clearly distinguish individual learning from group learning, and group learning from organizational-level learning on a conceptual and operational basis. Figure 1 provides a visual representation of group learning and its levels of analysis. The first case is an example of individual-level learning. In the final case, all members have an understanding of X and they know others have an understanding of X. This reflects one attribute of learning at the group level. We have observed the middle case in our own research on exocentric teams. Members individually understand something, but there is not shared understanding. This case represents a necessary

condition, but not sufficient condition for group-level learning. Group learning means that there is a change in the group's repertoire, not just in individuals. If member A simply transfers information to member B, that represents individual but not group-level learning.

We also need to be clearer in linking the operational and conceptual levels of analysis. For example, in the studies of transactive memory systems, two of the three factors used to measure a shared awareness of who knows what have been operationalized at the individual level. Similarly, task credibility evaluates how much members trust one another's knowledge. The less individuals criticize each other's work, the more task credibility exists within the group. However, individual-level criticism may only reflect a dyadic interaction and may not encompass the degree to which all members of the group trust the group's task credibility.

If we are going to consider the concept of shared knowledge at the group level, we need to examine it more directly and ensure that our measures reflect a group process or outcome. Techniques employed in research on shared mental models, which measure the extent to which knowledge is actually "shared" by members of a group, would provide more direct group-level indices of shared learning.

Related Concepts

It is important to distinguish group learning from related concepts. Tolman suggested, in the early part of the last century, that learning is not the same thing as performance (1932). In all three of the studies we reviewed, the outcome of learning was measured using performance indicators. For example, in the Edmonson operating room study (Edmonson et al, 2000), new learning was measured by a negative performance indicator: reversion to a conventional procedure during the operation. In the transactive memory studies, learning is assessed based on the group's collective recall, how quickly the group completes the task, and the number of errors it commits. Finally, in the Gruenfeld (2000) study, learning is measured as the number of unique ideas generated.

However, learning may occur when there is no change in a group's overall performance, and conversely, performance can change without learning actually taking place. There are numerous factors exogenous to groups that may lead to performance improvements or declines but are not related to learning. For organizational groups, this might include a change in the economy that opens new markets and opportunities for revenue generation that are not reflective of changes in the group's internal behaviors. Secondly, performance reflects task completion, which may occur significantly after the time that the learning actually occurs. It is possible that a group may learn from a prior experience, but might not have the chance to apply its knowledge until a future demand presents itself. For example, in the economic downturn of the early 1980s, Intel's top management team decided to stop investing in production plants and R&D. This slowed the pace of their innovation in DROM technology, in which they had been the clear leaders. At the same time, Japan's top management teams continued their pace of R&D investment. When the global economy recovered, Japan emerged as the dominant leader in this business sector, forcing Intel out of the market. The top management team at Intel learned through this experience, but did not have the opportunity to apply the learning

right away. Most importantly, we will most likely observe the performance outcomes from their actions for some time to come.

Much of the work on group learning also assumes that learning results in positive changes in group performance. For example, learning occurs in the Gruenfeld (2000) study if unique ideas are generated and utilized. In the transactive memory studies, improvements in the time of assembly and the reduction of errors are reflections of group members learning who knows what. In the Edmonson (2000) article, reversion to old procedures indicates non-learning. However, group learning can also result in negative changes in the range of potential behavior. Studies of group learning need to account for the possibility of superstitious learning (where the group learns a false connection between their actions and some outcome) or unlearning (where previously learned responses are extinguished). Ultimately, measuring group learning using performance-based metrics may not adequately reflect important changes in groups' repertoires.

Besides performance, information sharing, decision-making, and socialization also have been equated with group learning. Although these processes involve the exchange of knowledge, they do not guarantee that the knowledge will be stored by the group and subsequently applied. In the case of information sharing, it is possible that the information will be disbelieved or disregarded in some cases. Thus, the group can maintain its current repertoire without changes in its interdependent behavior. Decision-making involves the discussion of trade-offs and priorities with the goal of arriving at a final decision. However, once a decision is reached, it could confirm or reinforce prior behavior and result in no change to the group's repertoire. Finally, most research on socialization involves the processes by which individuals acquire knowledge about appropriate attitudes and behaviors required to participate in a group. However, as an individual learns about appropriate behavior within the group, this does not necessarily suggest that the group's behavior will change.

In conclusion, our position is that we need to be more explicit and detailed in our conceptualization of group learning. We have defined group learning as a change in the range of potential group behavior. The features that should be in this conceptualization include the storage, accessibility, sharing, and revisability of shared information.

Figure 1 about here.

III. GAPS IN CONTEXT

To date, most of the work on group learning has been done in the context of co-located groups, with members who belong to the same organization and who remain relatively stable over time. Other organizational researchers have argued that the context in which you study phenomena constrains your thinking about the key research variables and affects the way in which you model (Rousseau & Fried, 2001). Focusing on a narrow range of the phenomenon you are studying can result in a sort of conceptual range restriction. For instance, if you study the relationship between cognitive ability and job performance by looking at samples of graduates from top-tier universities, you would come to the (erroneous) conclusion that there is no relationship between cognitive ability and job performance because most everyone studied would have high cognitive skill. In this section, we consider how context affects our assumptions, our level of analysis, and the mechanisms we associate with group learning.

Working with groups in a particular context shapes the assumptions we hold about groups. In turn, these assumptions limit the variables we study, the models we develop, and even our approach to research. For example, when research centers on stable, co-located groups performing a physical task (product assembly, surgery), it leads to a focus on how groups learn to coordinate actions through the most readily available mechanism: observing other members' performance. When you expand or change the group context, it forces you to re-examine the assumptions, the mechanisms for learning, and even the key variables. For instance, if the members are not co-located, the process and the importance of developing a shared understanding changes dramatically. Examining some of the assumptions associated with observing group learning in particular contexts can indicate ways to expand the study of group learning in general.

Framework for thinking about group context

Our own thinking about group learning has been influenced by our study of "exocentric" teams: groups whose primary work activities, interactions, and relationships are external in nature (Goodman & Wilson, 2000). At their most extreme, exocentric teams have members from multiple organizations, working in a geographically distributed manner, who form and re-form over time. For the purposes of this research we focus on computer emergency response groups. These groups form spontaneously to respond to Internet-based threats or attacks, such as the recent "Love Letters" incident. Members of the group typically represent multiple organizations, often located around the world. The members interact intensely for the life of the attack (usually a period of a few hours or days) before the group disbands. The next event may not occur for several weeks or months, and the membership of the next response team may only partially overlap with previous response groups.

In our research (Goodman & Wilson, 2000) we have developed a framework for describing different forms of groups (see Table 2). The three major categories -- effectiveness, membership characteristics, and group structure -- generate a set of dimensions that characterizes different forms of groups. For illustration purposes, we note that traditional groups, which have been the focus of much group research, have internal and relatively clear effectiveness indicators (e.g., units produced over a period of

time), are composed of members from the same organization devoting their working hours to the group, and include structural features like having formalized roles and being co-located in space and time. The exocentric teams we have investigated have some of the opposite characteristics – membership from multiple organizations, a relatively short life, and roles that are emergent rather than prescribed. If we were to discuss cockpit crews or surgical operating teams, the profile would of course be different.

These varying profiles or group contexts matter because they operate on different theoretical assumptions, which rarely are studied explicitly. These assumptions are important because they permit certain types of learning mechanisms and not others.

What's missing?

When research does not make explicit context and assumptions related to that context, it limits the conceptual development of the group learning construct and its related processes. Let's illustrate this point by first examining the implications of making one's assumptions explicit in terms of group learning mechanisms and selecting the appropriate unit of analysis.

Effect of context on mechanisms for group learning

In the studies of group learning reviewed in this paper, one of the most common mechanisms for learning is training. In studies of transactive memory, one of the ways in which group members learn who knows what is through group training. In the Edmundson et al. (2000) studies of operating room teams, the early learning of routines and coordination occurred in training provided to the intact team. In contrast, joint training is not a realistic possibility in the computer emergency response teams we study. Members come from multiple organizations and work across great distances to get things accomplished in a relatively short time frame. Learning is more likely to occur when members have conflictive assumptions, agendas, or standard operating procedures. There is nothing like working with members from different organizations to cause people to re-examine their taken-for-granted assumptions about how work gets done.

Another learning mechanism in the transactive memory studies, and to some extent in the Edmundson's (2000) research, is that group members can learn by informally watching each other work. In these more traditional group settings, observation is likely to be a powerful mechanism for learning to coordinate activity. However, observation and other mechanisms, such as trial and error, assume some stability in the group. In contexts where members are not devoted full-time to the group and/or the group itself has a relatively short life, these mechanisms will likely be less prominent.

Another example in which context affects the mechanisms for group learning comes from assumptions about the structure of the group. Most studies of group learning assume that the members all share the same experiences. In studies of transactive memory, all the members of the group are present to observe the assembly of a radio. In the studies of operating room teams, all the members observe the same operation. Co-location is one of the key factors that contributes to shared interpretations or common ground. But what happens when the members are experiencing the same event from different locations and from within different organizations? Distance and differing

organizational contexts cause members to learn different things about a shared event. Observing this phenomenon in the computer emergency response teams we study has led us to posit that the larger the percentage of a group that experiences a problem, the more likely learning will occur. In traditional group structures the implicit assumption is that all members will perceive the same (or similar) problems since they have experienced the event together. However, even in traditional contexts members are unlikely to have the same perceptions of events (e.g., in the case of operating room teams, members' perceptions would be influenced by their roles). But researchers have neglected the question of how different perceptions or learnings get combined in groups because of the implicit assumption about shared experience.

Effect of context on level of analysis

Working with non-traditional groups has also helped us to reconsider the appropriate unit of analysis. In this case, it is important to analyze context because it makes us think explicitly about the appropriate unit of analysis. Because much of the group learning literature is (1) based on traditional groups and (2) oriented toward effectiveness studies, the group itself has been the unit of analysis. However, in the computer emergency response teams we study, the exact membership of the group changes from attack to attack, and the members come from multiple organizations. This leads us to conclude that a more appropriate unit of analysis in these contexts might the event and learning that occurs from event to event.

An event is a central incident that occurs over a limited period of time. Events involve a set of activities and interactions that are usually initiated by some exogenous force and have a definable beginning and end (Goodman, 2000). Although events are regarded as the fundamental units of analysis in other fields, such as relativity theory, they have received comparatively little attention in the organizational literature. When events are considered in the organizational literature, there are multiple treatments of events, including using events as variables in the analysis rather than as the unit of analysis. In contrast, we use events as the units of analysis and consider how learning occurs in groups from event to event. For computer emergency response groups, events are:

- unpredictable
- of short duration, and
- discontinuous, with varying time lags between events

Events are typically critical situations that require rapid response, so there is little time for reflection, planning, or structured communication until afterwards. Learning in this context presents a number of interesting challenges. Because the membership changes from event to event, it is difficult for groups to rely on individuals for memory or knowledge transfer. Other groups with changing membership (operating teams, flight crews) respond to this constraint by storing information in clearly delineated roles. This is not possible in computer emergency response teams because the roles emerge to fit the unpredictable nature of each crisis. The use of events as the unit of analysis encourages us to consider other storage mechanisms beyond individual members (who may change from event to event) and beyond formalized roles (which are not present in these new forms of groups).

Although we have suggested that events are an appropriate unit of analysis for group learning in exocentric contexts, events may also be useful in studying how groups learn in more traditional contexts as well. In a wide range of organizational groups (from top management teams to product development teams), one of the most interesting questions is how groups process learning acquired in one critical incident and transfer it to the next incident (Weick & Roberts, 1993). How does the poor performance of a recent acquisition affect a top management team's range of behavior? How do the learnings from one product development experience get transferred to the next development effort when it is likely that some, but not all, of the prior members are involved? In addition to helping us answer interesting questions about the transfer of learnings in groups, using events as the unit of analysis may help us avoid some of the problems associated with traditional outcome-based research (Aldrich, 1999).

We have argued that we need to make our assumptions about groups more explicit, and being more explicit about context may help us identify boundary conditions for group learning models and results. This may be increasingly important as work within and between organizations is increasingly carried out by groups that are formed to respond to a particular set of circumstances and are then disbanded, groups whose members work in different locations, or groups that are composed of members from multiple organizations (Miles, Snow, Matthews & Coleman, 1997). As organizations move from fixed structures to more fluid forms of organizing, they increase the risk that learning will not be transferred across events, across distance or across organizational boundaries. Understanding processes for group learning across events becomes more salient as traditional mechanisms for learning (individual retention, prescribed roles) become less common in new forms of organizing.

Table 1 about here

IV. MECHANISMS THAT CREATE LEARNING

In this section, we consider models of group learning. The basic question is what mechanisms or processes result in learning. We explore the models researchers use to explain group learning and what might be missing from these models. One task is to identify features that need to be considered in models of group learning.

To fully develop models of group learning, researchers need to address the following questions:

Learning Object. What is it that individuals or groups are supposed to learn? In an algebra class, a student team may be learning the basic rule that you can subtract the same amount from both sides of the equation without changing its basic meaning. Alternatively, the learning object could be developing a shared routine for monitoring patients in an intensive care unit (ICU).

Learning Task. Simon (2001) has argued that understanding the task is critical to understanding learning. There are of course many kinds of tasks, which range from solving an algebra problem to writing a paper about the group perception of its own

activities to analyzing the cause of an error in a hospital. How do task characteristics constrain or facilitate learning?

In addition, all group tasks unfold in some organizational environment. It may be in a course on groups, or in an ICU in a hospital. Both the task and its environment are critical issues in understanding learning processes.

Mechanisms. What are the critical mechanisms that lead to the development of the learning object? To illustrate, let's say a group has acquired a new routine to do its task more effectively, and we want to learn how this acquisition came about. At one level, we know the learners must select out some stimuli vs. others (e.g., a new routine.) Also, some of the information about the routine needs to be retained or stored, and some indexing system must exist to permit the recall of this information.

At another level, we need to explain why all of this occurs. What is the nature of the set of problems to be solved? Has the group experienced some error or visible discrepancy in its performance that has spurred it to see change? (Was there an opportunity to observe other groups who experienced similar discrepancies?)

Figure 2 about here

Figure 2 shows the basic elements in specifying a model of group learning. Our job is not to enumerate all possible objects, tasks, or mechanisms; rather we argue that a discussion of group learning should be explicit about:

- These three concepts
- The interrelationship among these concepts. Tasks provide constraints around different possible learning objects. Different mechanisms are more likely to be associated with different objects.
- How the intersection of objects, tasks, and mechanisms relate to the definition of group learning. For example, shared understanding is a critical feature of many definitions of group learning. We need a clear explanation on how a learning object becomes shared.

First, our analysis asks how researchers represent mechanisms of group learning. Next, we examine what else might be involved in specifying the mechanisms - or how learning occurs. To accomplish this we return to the three exemplar studies. Since the nature of this discussion is somewhat detailed, we will focus on one study and note the generalizability of our approach to other studies.

Gruenfeld, et al. (2000) asked what groups learn from their "worldly" members? The study design has members working in a group doing projects when one member (the "itinerant) moves to a new group and eventually returns to his or her group of origin. The researchers hypothesize that the itinerant member's unique knowledge and experience can be transferred from the group where it originated to another group engaged in the same activities. The mechanisms underlying this hypothesis seem to be that the itinerant member has a unique knowledge that other members take into consideration. The itinerant can have direct influence in convincing others to accept their idea or indirect influence by changing how others think.

This experiment unfolds in a lab integrated into a course in organizational communications. In the lab, groups of students completed a content task (e.g., examining leadership styles) and a process task (e.g., describing what happened in their group during the workshop). In the process task, each student had 15 minutes to write an individual essay; then the group had another 15 minutes to write a corrective essay on its process.

Our goal is to use this research to illustrate the ideas of mechanism specification in group learning. The focus is not to do an in-depth analysis of the paper's strengths and weaknesses. (This is a well-designed and presented study.) Let's return to our three constructs.

Examples of Learning Objects. In the Gruenfeld study (19XX), one indicator of learning was whether the itinerants' ideas would appear in the essays written by the temporary groups rather than the essays written by the original groups, or whether the itinerants' ideas would more likely appear in their original groups after they returned than before their departure.

If we assume these operational procedures are appropriate, it is still important to explore what else might be considered. What does "idea" mean? It could be, "Bill did not participate much," or "We did not plan our time well enough and had to rush the essay." These are two very different kinds of ideas. One is about an individual's behavior. The other is about a group-level activity. The implication in the latter example is that planning will help the group in its next task. At another level, we do not know whether either of these example ideas are shared by the other members, a possible attribute of group learning. Given that they have only 15 minutes, it is unlikely that the members of each group could present, discuss, and test for consensus across all of the ideas in their essay.

We do not argue that group learning cannot happen in this environment. Consider the following dialogues:

Itinerant: We did not plan our time well enough.
Member A: I agree.
Member B: We did not plan well in our last lab.
Member C: Let's assign Denise to be our time manager in the next lab.
All members: Yeah!

This discussion represents a process of group learning. An idea is presented about an aspect of the group, there is concurrence, and a routine is agreed upon for future activity (i.e., assigning a timekeeper). The problem and its solution are stored in people's heads, the essay, and in a new group role.

The basic point in this brief analysis is to be clear about the group learning object. Simply counting an idea from a member does not necessarily distinguish individual- or group-level learnings. In our discussion, we pointed out that the nature of the ideas (i.e., we didn't plan our time well), and the degree to which they are shared are both important.

Examples of Tasks. The nature of the tasks in this research were: The group solved a content-based problem, each member wrote an individual essay about group process, the group produced an essay about its processes. If we focus on the third task, the challenge is to combine a set of individually established positions in order to generate a group product and, as pointed out above, all of this must occur in a relatively short period of

time. In simply describing this task, we can draw on a large body of group literature that deals with process loss (Steiner, 1972). This task has inherent features that may or may not facilitate learning. Also, this group task is not isolated. The group's experience in solving the content problem as well as its experience with prior tasks should be important in understanding its behavior. In this study, the nature of the task is well described, but there is not much discussion about the relationship between the inherent nature of the task and the learning, and the relationship between the content group task and/or prior tasks in this learning episode.

Examples of mechanisms. The basic mechanism for group learning in our example is that the itinerant, who has unique knowledge, will directly influence the other members to accept his or her ideas or indirectly influence how the group thinks. The salience of a new member also may affect which of the itinerant member's comments are selected and retained.

This raises some questions about the implied mechanisms for group learning. First, is influence equivalent to learning? Let's assume that because of salience or interaction patterns, the ideas of one or several itinerants are incorporated in the essay. This is evidence that some ideas were selected out over others. But is it evidence of group-level learning? Although there clearly is a group product, the task involves multiple people proposing ideas, and the essay must be produced in a relatively short time. Are there truly clear, shared understandings or commitments to ideas in the report? If one itinerant worker gets the idea that "Bill did not participate much" into the report, will all of the people in the group retain that observation, share it, and recall it in the next lab? Only in some very special circumstances that deal with learning vs. influence processes will the idea be shared.

In any research, it is critical to state the processes or mechanism by which ideas, in this case from the itinerant member, become shared. One mechanism might be the nature of the communication pattern, whereas another might be that the itinerant's comments evoked some latent learnings that group members had from prior experiences. These mechanisms are very different.

The discussion of learning mechanisms needs to deal not only with the selection of certain stimuli, but also with where they are stored (e.g., within people, in reports) and how they are indexed for recall. In the Gruenfeld et al. (2000) paper, the storage mechanism appears to be the group report. Some reasonable questions to ask include: To what extent is the information in the report stored and accessible in the individual members' heads? How many members would remember what is in the report after it is submitted? Would the report ever be accessed by members as a memory source?

Another question deals with the alternative mechanisms in play as the essay is written. A key mechanism is the type and frequency of process behaviors generated during the group discussion. Clarification, building, and summarizing behaviors will shape the final report. How these behaviors lead other members to interact with the itinerant members will be important. The order effect of information, who writes the essay, and the latent understandings of the members (e.g., they may agree that Bill doesn't participate but not say so) will determine what particular ideas get incorporated in

the report and whether there is some new shared understanding of group processes that may be recalled and used in future group activities.

Specifying Mechanisms of Group Learning. Our overall goals are to examine how we specify the mechanisms of group learning and to indicate features that might be included. We selected the Gruenfeld et al. (2000) study as an example. This same approach could have been used in analyzing the transactive memory studies or Edmunson et al's research on learning in the hospital setting. We think the early empirical work on group learning helps us identify important features for more fully specified models of group learning.

First, we should be more explicit about the learning object. This construct was discussed in the first section of the paper. In specifying models of group learning, we need a good categorization of what people are learning. Is it a fact, opinion, contingency (if we do plan better, our group product will be of higher quality), or routine? Is the object a group-level phenomenon? Is it accessible and shared by group members? How does the nature of the object affect where it is stored? Can it be recalled and used in a future action? Can it be revised?

A second need in the group learning literature is to be more specific about tasks. A group writing an essay is different from a group producing a radio or designing a new product. Some tasks are played out in environments where people's lives are on the line; others involve essays worth X percent of a grade in a course. What are the inherent features of the task, and how do they affect learning processes. McGrath and others (1983) have argued that task features affect group processes. Giving a group 15 minutes to reach some consensus about the elements to include in a paper is a different setting from giving a group one year to create a new product. The processes for learning in these two scenarios are likely to be equally important and possibly very different.

Third, a better understanding of interactions among the objects, tasks and learning mechanisms may sharpen our understanding about group learning. If the task is a problem-solving activity aimed at improving group processes, this creates some boundaries around the possible mechanisms that might be effective. In this case, the object is likely to be some explicit process behavior. The key will be the interactions that lead the group to identify possible behaviors, build shared understandings, create possible storage sources, and so on. If the task is discovering a new way to design a product, then the mechanisms are likely to be very different. Finding the answer may require a tacit-to-explicit transformation coupled with other mechanisms to create shared understandings,, storage methods, and ways to revise and modify the new design. The intersection between objects and tasks will lead to a different set of learning mechanisms.

Fourth, the learning mechanisms need to be specified in more detail. Over the past two decades, research on groups has continued to use very general models (Goodman, Ravlin & Argote, 1986). Our real knowledge of the complexity of group-level learning is rarely revealed. We need to be more specific about how learning occurs.

How might one do this? First, frame a set of questions that the models should address. We have suggested some of these questions: Why do groups select certain ideas and not others? What contributes to whether the information is shared, known by other members but not shared, or not shared at all? Where does information get stored, how, and why? What are the implications for accessing stored information? Since learning is a

dynamic process, how do groups learn, unlearn, and re-learn new information? We want to examine these questions in the light of specific task and object intersections.

V. DISCUSSION

There are four basic themes in this paper. First, there is a new, growing interest in group-level learning. Second, our conceptual and definitional approaches to group learning to date have been limited. We indicated the need to specify the construct space of group learning to conceptualize what group learning is and what it is not. Four features were proposed in a conceptualization of group learning – sharing, storage, accessibility, and revisability.

Third, there is growing awareness that context is important in our research. We tried to be more provocative by saying variation in context is correlated with variation in assumptions, and these assumptions affect the kinds of learning processes that are possible. The framework in Table 1 provides a way to think about context variations and learning.

Lastly, while the initial studies of group learning have been invaluable, we argued that most researchers are under-specifying how group learning takes place. Our suggestion is to specify the task and learning object and then postulate the mechanisms that would intersect with the features of tasks and learning object.

If we are motivated to build richer models of group learning, then there are a number of intellectual challenges.

Sharing is one of our proposed features of group learning. Figure 1 presented one way to think about sharing, but further reflection will show that the meaning of sharing is quite complicated. What are the mechanisms that lead to a common idea that has not been communicated vs. a shared idea that has been discussed? Can we conceptualize and represent varying degrees of sharing (e.g., strength)? What are the implications of learning objects that are shared, but in different strengths? Most sharing is about similar objects. What are the implications for group learning when it is the diversity of objects that is shared?

We could explore other features of learning as storage. In our analysis, we saw storage in individuals, group products, and organizational routines. Is there a more analytic way to categorize storage in memory systems such as people, computers, routines, etc. (Olivera, 1999)? Do features inherent to these storage systems affect the accessibility of information to group members? Are some learning objects more suitable to certain types of storage or memory system? An explanation of these and other features provides a marvelous opportunity for new work.

In the discussion of context, the basic point is that varying dimensions such as the location of members or lifespan of a group have a major impact on what mechanisms could be evoked for group learning. The real challenge, however, is not simply to understand the implications of one dimension at a time. The unexplored challenge is understanding the interactive effect of bundles of dimensions. For example, a team with a very short life, composed of different organizational members with emergent roles and diffuse criteria, is different from a group with a similar short life and diverse members,

but with formal roles and operational criteria. We need to understand how these different combinations of factors affect group learning.

Our proposal to use events as a unit of analysis rather than the group itself poses interesting intellectual and operational questions. One simple question is how to define an event, including specifying when it begins and ends. In our research, we generally know when an attack on the Internet starts, but the end of an event is sometimes more difficult to specify. Also, there may be events within events. For example, in our research on computer response teams, there are attacks (such as Love Letters - which is a form of event), and within the attack there are multiple events (such as producing the advisory). There is a growing acknowledgment of event-based research in the literature, and the definition and characterization of events is attractive because events represent a different approach to studying group processes and can apply to both traditional and exocentric groups.

In current empirical research on group learning, it often is hard to understand the specific learning mechanisms at work. One issue may be the way we graphically represent learning mechanisms. Much of the group literature uses boxes and arrows to represent different phenomena. The problem with this mode of representation is that it portrays the mechanisms at a pretty general level. As an alternative, group theorists might want to expose themselves to the “systems dynamic” research, which provides a stylized and more detailed representation of phenomena. Systems representations, particularly of positive and negative feedback loops might provide an informative way to show how learning does or does not occur in groups.

Another interesting challenge in better understanding the mechanisms of group learning is to factor in the role of alternative processes. Group learning involves a complicated set of processes that connect human skills and energies with task requirements. As we mentioned in the analysis of Gruenfeld et al. (2000), many of the processes have nothing to do with learning. It would be interesting to map how learning processes complement or compete with other group processes.

Finally, more exploration is needed into the generalizability of mechanisms of group learning. A fairly straightforward extrapolation from our discussion of context is that mechanisms of learning will vary by context. If this is true, are we left with context-specific models, or are there properties that remain constant across various contexts?

Figure 1. **Degree of Sharing/Level of Analysis**

Individual Information

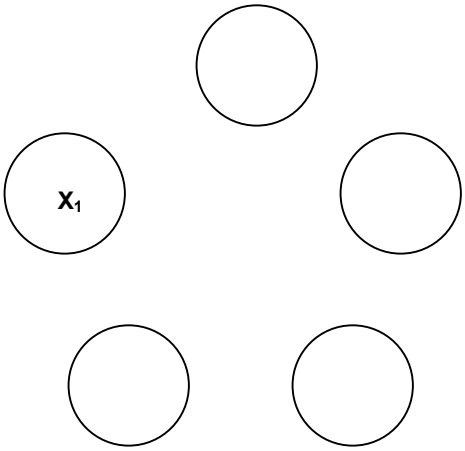

Similar Information, No Shared Awareness – Latent Group Sharing

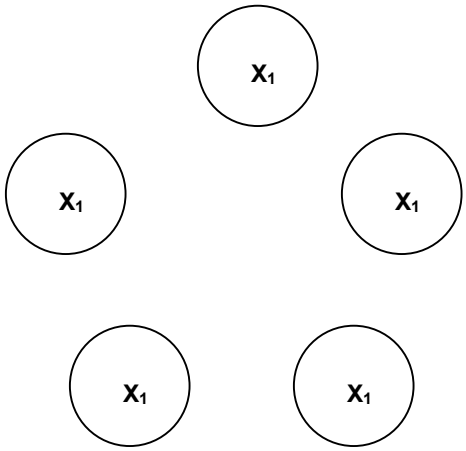

Similar Information, Shared Awareness – Group Sharing

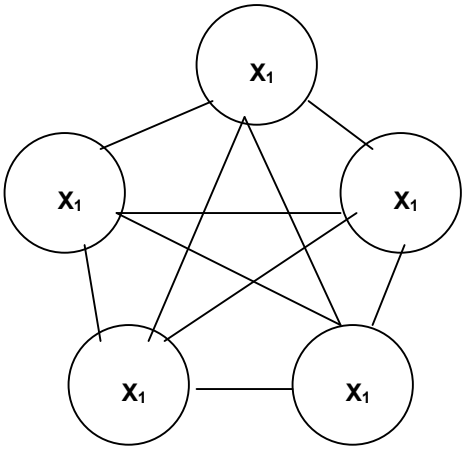

Strength of Sharing

Latent

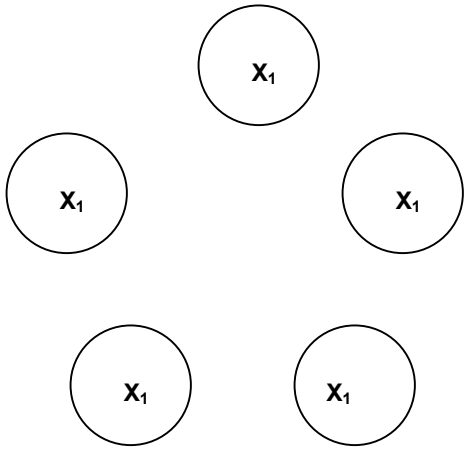

Weak

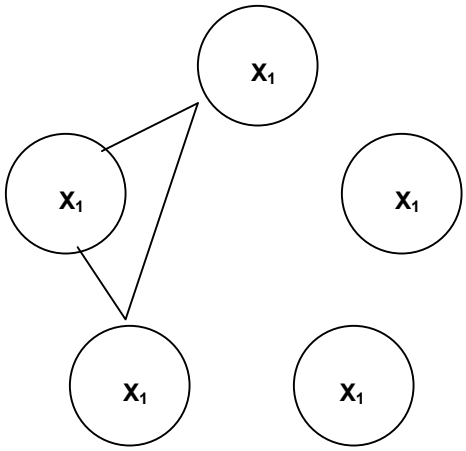

Strong

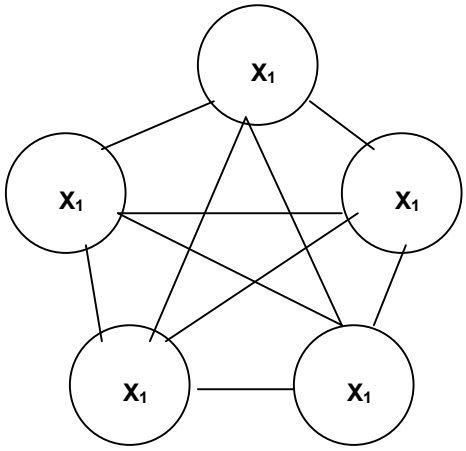

Degree of Diversity

Diverse Information, No Shared Awareness

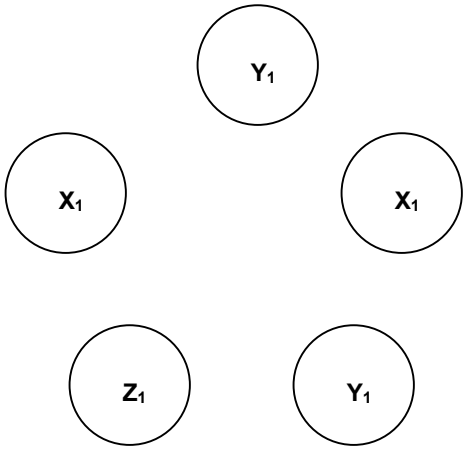

Diverse Information, Shared Awareness

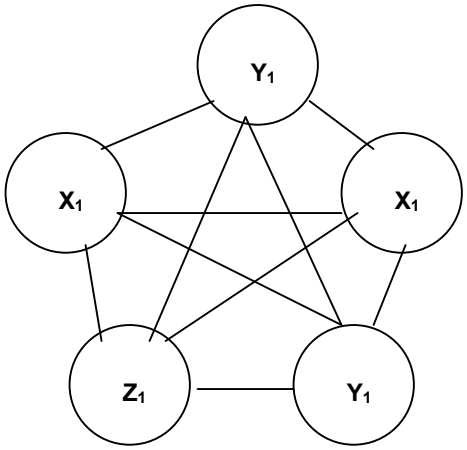

Sharing Information Over Time

Time 1

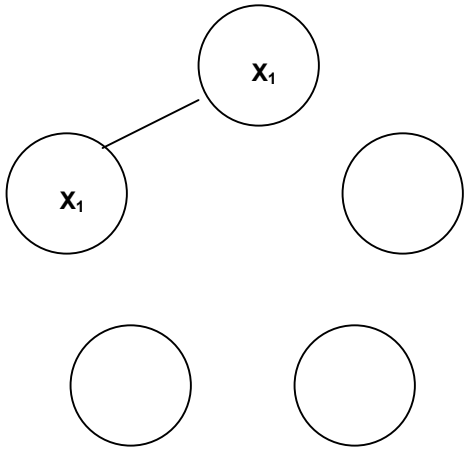

Time 2

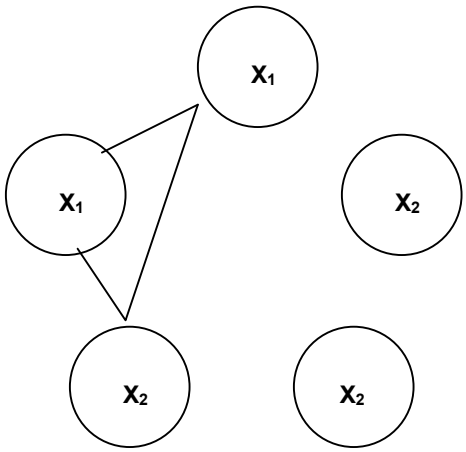

Time 3

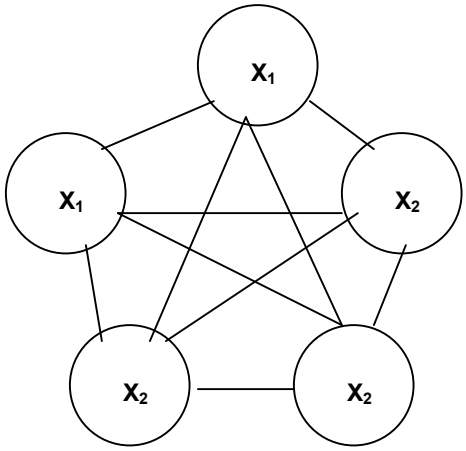

Figure 2: Components of models of learning

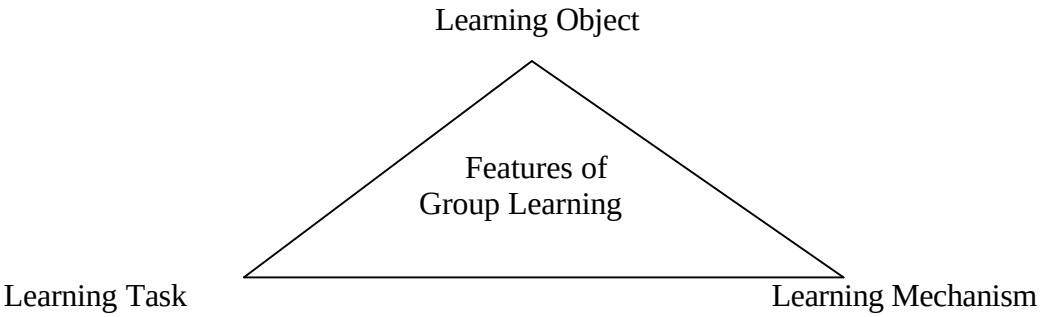

Table 1: Dimensions of Exocentric Teams

	Traditional	***Exocentric***
Effectiveness	• internal • operational	• external • diffuse
Membership	• same organization • 100% of time in group • stable	• different organizations • <100% of time in group • dynamic
Structure	• formalized roles/procedures • co-located in space and time • internally interdependent • temporal interdependence – fixed	• emergent roles/procedures • distributed in space and time • externally interdependent • temporal interdependence – variable

References

Aldrich, H. (1999) Organizations evolving. Thousand Oaks: Sage.

Argote, L. Gruenfeld, D.H. Naquin, C. (2001). Group Learning in Organizations. In M.E. Turner (Ed.), Groups at work: Advances in theory and research. New York: Lawrence Erlbaum.

Argyris, C. (1982). The executive mind and double loop learning. Organizational Dynamics, 11, 5-23.

Edmondson, A.C. (1996). Learning from mistakes is easier said than done: Group and organizational influences on the detection and correction of human error. Journal of Applied Behavioral Science, 32, 5-32.

Edmundson, A.C. (1999). Psychological safety and learning behavior in work teams. Administrative Science Quarterly, 44, 350-383.

Edmondson, A.C., Bohmer, R.M., and Pisano, G.P. (2000). Learning New Technical and Interpersonal Routines in Operating Room Teams. The Case of Minimally Invasive Cardiac Surgery. Research on Managing Groups and Teams. (pp. 29-51). Vol. 3: JAI Press Inc.

Goodman, P.S., Ravlin, E.C. & Argote, L. (1986). Current thinking about groups: Setting the stage for new ideas. In P.S. Goodman (Ed.), Designing Effective Workgroups. (pp. 1-33). San Francisco: Jossey Bass.

Goodman, P.S. & Wilson, J.M. (2000). Substitutes for socialization and exocentric teams. In M.A. Neale, E.A. Mannix & T.L. Griffith (Eds.). Research on managing groups and teams. (pp. 53-78) Stamford, CT: JAI Press.

Gruenfeld, D.H., Martorana, P.V., Fan, E.T. (2000). What do groups learn from their worldliest members? Direct and indirect influence in dynamic teams. Organizational Behavior and Human Decision Processes, 82, 45-49.

Liang, D., Moreland, R. & Argote, L. (1995). Group versus individual training and group performance: The mediating effects of transactive memory. Personality and Social Psychology Bulletin, 21, 384-393.

McGrath, J.E. (1983). Groups: Interaction and performance. New York: Prentice Hall.

Miles, R.E., Snow, C.C., Matthews, J.A., Miles, G. & Coleman, H.J. (1997). Organizing in the knowledge age: Anticipating the cellular form. Academy of Management Executive, 11, 7-20.

Moreland, R.L., Argote, L. Krishnan, R. (1996). Socially Shared Cognition at Work. Transactive Memory and Group Performance. Research on Socially Shared Cognition in Small Groups. (pp. 57-84). Thousand Oaks, CA: SAGE Publications.

Mowrer, R.R. & Klein, S.B. (2001) The transitive nature of contemporary learning theory. In R.R. Mowrer & S.B. Klein (Eds.) Handbook of contemporary learning theories. (pp. 1-21). Mahwah, N.J.: Lawrence Erlbaum Associates.

Olivera, F. (1999). Memory systems in organizations. Unpublished doctoral dissertation, Carnegie Mellon University.

Rousseau, D.M. & Fried, Y. (2001). Location, location, location: Contexts for organizational behavior. Journal of Organizational Behavior, 23, 1-13.

Steiner, I.D. (1972). Group process and productivity. San Diego, CA: Academic Press.

Tolman, E.C. (1932). Purposive behavior in animals and man. New York: Appleton-Century-Crofts.

Weick, K.E. & Roberts, K.H. (1993). Collective mind in organizations: Heedful interrelating on flight decks. Administrative Science Quarterly, 38, 357-381.