



I L L I N O I S

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

PRODUCTION NOTE

University of Illinois at
Urbana-Champaign Library
Large-scale Digitization Project, 2007.

Technical Report No. 481
SITUATED COGNITION AND
THE CULTURE OF LEARNING

John Seely Brown
Xerox Palo Alto Research Center and
Institute for Research on Learning

Allan Collins
Bolt Beranek and Newman Inc.

Paul Duguid
Institute for Research on Learning

August 1989

Center for the Study of Reading

TECHNICAL REPORTS

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
174 Children's Research Center
51 Gerty Drive
Champaign, Illinois 61820

THE LIBRARY OF THE

SEP 11 1989

UNIVERSITY OF ILLINOIS
URBANA-CHAMPAIGN

CENTER FOR THE STUDY OF READING

**Technical Report No. 481
SITUATED COGNITION AND
THE CULTURE OF LEARNING**

**John Seely Brown
Xerox Palo Alto Research Center and
Institute for Research on Learning**

**Allan Collins
Bolt Beranek and Newman Inc.**

**Paul Duguid
Institute for Research on Learning**

August 1989

**University of Illinois at Urbana-Champaign
51 Gerty Drive
Champaign, Illinois 61820**

The work upon which this publication was based was supported in part by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research under Contract No. N00014-C-85-0026, Contract Authority Identification Number, NR 667-540 and in part by the Office of Educational Research and Improvement under Cooperative Agreement No. OEG 0087-C1001, with the Reading Research and Education Center. The publication does not necessarily reflect the views of the agency supporting the research.

EDITORIAL ADVISORY BOARD
1988-89

Beck, Diana

Commeyras, Michelle

Foertsch, Daniel

Hartman, Doug

Jacobson, Michael

Jehng, Jihn-Chang

Jimenez, Robert

Kerr, Bonnie

Kerr, Paul

Meyer, Jennifer

Moran, Juan

Ohtsuka, Keisuke

Roe, Mary

Schommer, Marlene

Scott, Judy

Stallman, Anne

Wilkinson, Ian

Wolff, Phillip

MANAGING EDITOR
Mary A. Foertsch

MANUSCRIPT PRODUCTION ASSISTANTS
Delores Plowman
Debra Gough

Abstract

Many teaching methods implicitly assume that conceptual knowledge is independent of the situations in which it is learned and used. The authors examine one such method and argue that its lack of success is a direct result of this assumption. Drawing on recent research into learning in everyday activity and not just in the highly specialized conditions of schooling, they claim that knowledge is not independent but, rather, fundamentally "situated," being in part a product of the activity, context, and culture in which it is developed. Teaching, however, often overlooks the central, but restrictive, contribution made by the activities, context, and culture of schools to what is learned there. A theory of situated knowledge, by contrast, calls for learning and teaching methods that take these into account. As an alternative to conventional, didactic methods, therefore, the authors propose teaching through "cognitive apprenticeship" (Collins, Brown, & Newman, 1989). They examine two examples of mathematics teaching that exhibit important features of this approach.

SITUATED COGNITION AND THE CULTURE OF LEARNING

Introduction

Deeply embedded in our society is a profoundly misleading, theoretical separation between *knowing* and *doing*. The folk categories of *know how* and *know what* reflect this, but the division extends well beyond folk categories. Much current cognitive science, philosophy of mind, and artificial intelligence, for instance, assumes that knowing is a process that can be separated from the activities and situations in which it is used. And many of the practices of conventional schooling exhibit the same assumption. Thus G. B. Shaw could make his notorious comment that "He who can, does. He who cannot, teaches." An understanding of the inseparability of knowing and doing, which we argue for in this paper, leads to a distinctly different view of teaching (see, for example, Dreyfus & Dreyfus, 1986; and Haroutunian-Gordon, 1988) and a more enlightened appreciation of good teachers.

In his epigram, Shaw encapsulated, unfairly but concisely, conventional educational assumptions that knowledge can be usefully regarded as self-contained and discrete, and that it can be adequately transferred from teachers to students in the activities of the classroom independent of the activities in which that knowledge might normally be used. Knowing, in this view, is assumed to go on in heads and teaching to go on in classrooms isolated from the complexity of the world outside, from which abstract knowledge can successfully be distilled. A growing body of research that focuses on cognition in everyday activity, however, is beginning to undermine the plausibility of these presuppositions.¹ Knowing (and not just learning), this research suggests, is inextricably *situated* in the physical and social context of its acquisition and use. It cannot be extracted from these without being irretrievably transformed.

If knowledge is situated, then many conventional assumptions must be questioned. In particular, a situated theory of knowledge challenges the widely held belief that the abstraction of knowledge from situations is the key to transferability. An examination of the role of situations in structuring knowledge indicates that abstraction and explication provide an inherently impoverished and often misleading view of knowledge. Knowledge is fundamentally a co-production of the mind and the world, which like wool and warp need each other to produce texture and to complete an otherwise incoherent pattern. It is impossible to capture the densely interwoven nature of conceptual knowledge completely in explicit, abstract accounts.

In the first part of this paper, we explore, in a consciously speculative way, just how cognition is situated.² Knowledge, we claim, is partially embedded in the social and physical world. And these embedding circumstances allow people to share the burden of solving problems efficiently. Learning is the process of constructing robust understanding out of this sort of embedded activity. Such an approach to learning and knowledge offers, among other things, new insights into pedagogical practices. These we address in the second part. In particular, a theory of situated cognition helps explain "cognitive apprenticeship" (Collins, Brown, & Newman, 1989). In methods of cognitive apprenticeship, teachers deliberately deploy the embedding circumstances of knowledge to help students construct robust understanding.

Situated Knowledge and Learning Activity

Learning and Learning Activity

Miller and Gildea's (1987) work on vocabulary teaching shows how the assumption that knowing and doing can be separated leads to a teaching method that ignores the way situations structure cognition. Their work describes how children are taught words from dictionary definitions and a few uprooted exemplary sentences. They compare this method of teaching with the way vocabulary is normally learned outside school.

People generally learn words in the context of ordinary communication. This process is startlingly fast and successful. Miller and Gildea note that by listening, talking, and reading the average 17-year-old has learned vocabulary at a rate of 5,000 words per year (13 per day) for over 16 years. By contrast, learning words from definitions and sentences abstracted from the context of normal communication,

the way vocabulary has often been taught, is slow and unsuccessful. There is barely enough classroom time, Miller and Gildea suggest, to teach more than 100 to 200 words per year. Moreover, much of what is taught turns out to be almost useless in practice. They give the following examples of students using vocabulary acquired this way:

Me and my parents correlate, because without them I wouldn't be here.

I was meticulous about falling off the cliff.

The redress for getting well when you're sick is to stay in bed.

I relegated my pen pal's letter to her house.

Mrs. Morrow stimulated the soup.

The news is very tenet.

Our family erodes a lot.³

Given the method, such mistakes seem unavoidable. Teaching from dictionaries assumes that definitions and exemplary sentences are self-contained "pieces" of knowledge. But words and sentences are not islands, entire unto themselves. They rely on the context in which they are used--and not just the linguistic context--to be understood. Language use would involve an unremitting confrontation with ambiguity, polysemy, nuance, metaphor, etc., were it not for the extralinguistic help that the context of an utterance provides (Nunberg, 1978).

Prominent among the intricacies of language that depend on extralinguistic help are *indexical* words--words like *I*, *here*, *now*, *next*, *tomorrow*, *afterwards*, *this*, which thoroughly ground interpretation in situations. Indexical terms are those that "index" or more plainly point to a part of the situation in which communication is being conducted.⁴ They are not merely context sensitive; they are completely context dependent. They can only be understood in relation to specific situations, from which indexicals are, for the purposes of interpretation, inseparable. A different context for *I* or *now* produced a different meaning. These indexical words rely on a great deal of contextual, extralinguistic information to be understood. And surprisingly, all words can be seen as at least partially indexical (Barwise & Perry, 1983).

Experienced readers implicitly understand that words are situated and therefore ask for "the rest of the sentence" or other contextual information before committing themselves to an interpretation of a word. They go to dictionaries with situated examples of usage in mind. Then the context as well as the dictionary support the interpretation and continually refine a person's understanding of a word. Learning from dictionaries alone, like any method that tries to teach abstract concepts independent of authentic situations, overlooks the way understanding is acquired and developed through continued, situated use. This involves complex social negotiations, which never crystallize into a categorical definition. Because it is dependent on situations and negotiations, the meaning of a word cannot, in principle, be captured by a dictionary definition and a couple of exemplary, but uprooted, sentences alone. Yet the students who produced the sentences above were given only definitions, with no support from the context of normal discourse. Dictionary definitions were assumed to be self-sufficient. The extralinguistic props that structure, constrain, and ultimately allow interpretation in normal communication were ignored by the teaching process, and the students, as a result, had nothing to help limit possible interpretations. And what they learned in the process was therefore either useless or deeply misleading.

All knowledge is, we believe, like language. Its constituent parts index the world and so are inextricably a product of the activity and situations in which they are produced and used. Because new situations, negotiations, and activities inevitably recast them in a new, more densely textured form, concepts continually evolve with each new occasion of use. Thus all concepts--even apparently well-defined, abstract technical concepts--are always under construction. They are never wholly definable, and they

defy the sort of categorical descriptions that are conventionally used in teaching: part of their meaning is always inherited from the contexts in which they are used.

Learning and Tools

To explore the idea that concepts are both situated and progressively developed through activity, we should abandon once and for all any notion that a concept is some sort of abstract, self-contained substance. Instead, it may be more useful to consider conceptual knowledge as in some ways similar to a set of tools.⁵ Tools share several significant features with knowledge; in particular, both can be fully understood in use, and both, if fully understood, change the way their users look at the world.

If knowledge is thought of as tools, we can illustrate Whitehead's (1929) distinction between the mere acquisition of inert concepts and the development of useful, robust knowledge. It is quite possible to acquire a tool but to be unable to use it. Similarly, it is common for students to acquire algorithms, routines, and decontextualized definitions that they cannot use and that therefore lie inert. Unfortunately, this problem is not always apparent. Old-fashioned pocket knives, for example, have a device for removing stones from horses' hooves. People with this device may know what it is for and be able to talk wisely about horses, hooves, and stones. But they may never betray--or even recognize--that they wouldn't begin to know how to use this implement on a horse. Similarly, students can often manipulate with apparent competence algorithms, routines, and definitions they have acquired and yet not reveal--to their teachers or themselves--that they would have no idea what to do if they came upon the domain equivalent of a limping horse.

By contrast, people who use tools in authentic activity actively build an increasingly rich implicit understanding both of the tools themselves and of the world in which they use those tools. Their understanding, initially narrow, is continually broadened through use. Learning and acting are, as a result, interestingly indistinct, learning being a continuous, life-long process resulting from acting in situations.

To develop this comparison between tools and knowledge, between manual tools and conceptual tools, a little further, it can be noted that learning how to use a tool involves far more than can be accounted for in any set of explicit rules. The occasions and conditions for use arise directly out of the context of activities of each community or culture that uses the tool, framed by the way members of that culture see the world. The culture and its cultural viewpoint, quite as much as the tool itself, determine how a tool is used. Thus, carpenters and cabinet makers, for example, both use chisels and lumber, but each profession uses them quite differently. These differences reflect the viewpoint and insights of the members of each particular culture. The culture's viewpoint, a product of its activity, is only gained through entering that culture. Thus, it isn't possible to use a tool fully and appropriately without entering the culture in which it is used, observing practitioners at work, and engaging in authentic cultural activity.

Conceptual tools similarly reflect the cumulative wisdom of the culture in which they are used and the insights and experience of individuals. Their meaning is not invariant but a product of negotiation within a culture and of practice in authentic activity. Again, appropriate use is not simply a function of the abstract concept alone. It is a function of the culture and the activities in which the concept has been developed and used. Just as carpenters and cabinet makers use chisels differently, so physicists and engineers use mathematical formulae differently. Activity, concept, and culture are interdependent. No one can be totally understood without the other two. Learning must therefore involve all three. Teaching methods, however, often try to impart abstracted conceptual tools as fixed, well-defined, independent entities that can be explored in prototypical examples and textbook exercises. But such teaching denies students the access to either the activity or the culture that they need in order to develop an active understanding of a particular concept.

To talk about academic disciplines, professions, or even manual trades as communities or cultures may seem strange. Yet practitioners in each of these areas are connected by far more than their ostensible tasks. They are bound by intricate, socially constructed webs of belief, which make it possible to see them as cultures (Geertz, 1983), and which are essential to interpreting what they do. The activities of most cultures are unfathomable, unless they are viewed from within the culture, for membership of a

culture provides a set of cultural eyeglasses that are the key to understanding and carrying out its activities.

The culture, its belief system and the way it uses its tools--whether they are manual or conceptual--determine the way practitioners see the world. And the resulting way the world appears reciprocally affects the belief system and the activity. Activity, culture, and tools form a complex, interdependent, and inseparable unit. Unfortunately, students are too often presented with only a part of this complex. They are asked to use the conceptual tools of a discipline without being able to look through its cultural eyeglasses. To learn to use tools as practitioners use them, students, like apprentices, must be enabled to enter that community and its culture. Thus, in a significant way, learning is, we believe, a process of enculturation.

Learning and Enculturation

Enculturating may, at first, appear to have little to do with learning. But it is, in fact, what people do in learning to speak, read, and write or to become school children, office workers, researchers and so on. From a very early age and throughout their lives, people, consciously or unconsciously, adopt the behavior and belief systems of new social groups. Given the chance to observe and practice *in situ* the behavior of members of a particular culture, people pick up relevant jargon, imitate behavior, and gradually start to act in accordance with the culture's norms, though these are often recondite and extremely complex. Students, for instance, can quickly get an implicit sense of what is suitable diction, what makes a relevant question, what is legitimate (and also illegitimate) behavior in a particular school activity. The ease and success with which they adopt the school culture (as opposed to the intricacy of describing what it entails) belie the immense significance of the process and obscure the fact that what they pick up is a product of the ambient culture rather than explicit teaching.

The practices of contemporary schooling, however, usually deny students the chance to engage the relevant domain culture, because that culture is not in evidence. Although students are shown the tools of many academic cultures in the course of a school career, the pervasive culture that they observe, in which they participate, and which some enter quite effectively is the pervasive culture of school life itself--a culture that can unintentionally be quite antithetical to useful domain learning. School becomes the dominant cultural framework within which many students assimilate what they learn. But the way schools use dictionaries (or math formulae, or historical analyses) is very different from the way practitioners use them (Schoenfeld, in press). Thus students may pass school-based exams (a distinctive part of the school culture) but still not be able to use a domain's conceptual tools in authentic practice.

Before they can use the tools of a domain in the way a practitioner uses them--either formally or informally--students need to see the activity from the practitioner's cultural standpoint. But this standpoint is the product of domain activity, not explicit teaching. You cannot tell apprentice carpenters how to hit the blunt end of a chisel and assume they can infer the practices of carpentry. And you cannot present students with some of the uprooted conceptual ideas of a domain and typical textbook examples and exercises and assume that they can infer the subtle belief system implicit in the culture's use of those concepts in the world.

This is not to suggest, however, that all students of mathematics or history must be expected to become professional mathematicians or historians. Rather, we claim that in order to learn these subjects (and not just to learn about them) students need much more than abstract concepts and self-contained examples. They need to be exposed to practitioners using these tools in the authentic, ill-defined problems of their world. It is from this sort of activity that the culture's belief system--the key to understanding its behavior--can be inferred and adopted. Thinking mathematically doesn't necessarily include grappling with the unresolved problems of the subject. Mathematicians can think mathematically about apparently trivial issues. As Schoenfeld's teaching shows (see below), even relatively simple problems can be used to tease out the way a mathematician looks at the world and solves emergent problems. The process may appear informal, but it is nonetheless full-blooded, authentic mathematical activity; it can thus be deeply informative for students--in a way that textbook examples and declarative explanations are not.

Activity

In noting the centrality of activity to learning and knowledge, we have leaned heavily on the as yet undefined concept of *authentic* or *appropriate* activity. As a result of our discussion of cultures and enculturation, we are now in a position to clarify and explore these terms a little further.

Authentic activity. Like a tool, activity is only understood with regard to a particular culture and its belief system. It cannot be understood in isolation. Each culture implicitly embraces a set of relevant activities--some of which it may not even recognize explicitly--and constructs meaning and purpose in relation to them; conversely, the common activities, tools, and belief systems are what help to define a particular culture. Authentic activity is simply the ordinary activity of the practitioners in a culture, or activity which is congruent with their ordinary activity. In order to enter a culture, to develop its belief system, and to understand its goals, those are the activities that need to be undertaken.

Again, it is worth stressing that authentic activity is not only done by people who are already experts; it is not necessarily work in the forefront of the field. Apprentice tailors (Lave, 1988a), for instance, begin by ironing finished garments (which tacitly teaches them a lot about cutting and sewing). The activity is simply, valuable, and absolutely authentic. Students of Palincsar and Brown's (1984) reciprocal teaching of reading may read only elementary texts, but they nonetheless use authentic strategies common in some form or other to all readers. The students in Miller and Gildea's study, by contrast, take part in a pedagogical activity that is not even congruent with the normal behavior of practitioners.

School activity tends to be hybrid--implicitly framed by one culture while explicitly being attributed to another. Most classroom activity inevitably takes place within the culture of schooling, but it is attributed by both teachers and students to the cultures of readers, writers, mathematicians, historians, and so on. What students do in school thus tends to be a sort of *ersatz* activity, distorting both what is learned and the culture to which it is attributed.

This sort of school activity is very different from what we have in mind when we talk of authentic domain activity because it is in important ways not congruent with what a domain's authentic practitioners do. Much school work has become a highly specialized, self-confirming activity in a culture of its own. When, for pedagogic purposes authentic domain activities are transferred to the classroom, their context is usually transmuted; they become classroom tasks and part of school culture. Classroom procedures are then applied to what have become classroom tasks. As a result, the system of learning and using (and, of course, testing) can remain hermetically sealed within the self-confirming culture of the school. Consequently, contrary to the aim of schooling, success within this culture often has little bearing on performance elsewhere.

Math word problems, for instance, are generally encoded in a syntax and diction that is common only to other math problems. Thus the word problems of a textbook from 1478 are instantly recognizable today (Lave, 1988c). Many of today's word problems, however, are as foreign to contemporary authentic math practice as they are to the math practice of the fifteenth century. By participating in this *ersatz* activity students risk misunderstanding entirely what practitioners actually do.

Most classroom tasks are inevitably *ersatz* because, in their creation, apparently peripheral features of authentic tasks (like the extralinguistic supports involved in learning vocabulary) are dismissed as "noise" from which salient features can be abstracted for the purpose of teaching. The features of the environment that people actually use to perform mathematical calculations, for instance, are not included in word problems (see next page). Classroom exercises assume that such calculations are performed solely through abstracted mathematical algorithms. In all activity, the context offers an extraordinarily complex network of support for all practitioners, and, as a result, its absence vitiates the learning exercise. Furthermore, the source of such support is often only tacitly understood by practitioners, or even by teachers or designers of simulations. Even well-planned classroom exercises can, therefore, completely fail to provide the supporting contextual features that allow authentic activity. At the same time, students may come to rely on features of the classroom context, in which classroom tasks are inevitably embedded, that are wholly absent in authentic activity. Thus, much of what is learned in school may apply only to the *ersatz* activity, if it was through such activity that it was learned.

Activities of students, practitioners, and JPFs. The idea that most school activity exists in a culture of its own helps explain many of the difficulties of cultivating robust domain learning in school. Jean Lave's ethnographic studies of learning and everyday activity (1988a) reveal how schools can divorce students from the activities and culture that give meaning and purpose to what they learn elsewhere. This is the separation between what Resnick calls "learning in and out of school" (1988). Lave focuses on the behavior of JPFs (just plain folks) and records that the ways they learn are quite distinct from the ways students are expected to learn.

Three categories primarily concern us here: JPFs, students, and the practitioners to whose status both JPFs and students aspire. Put most simply, when JPFs want to learn a particular set of practices, they have two apparent options. First, they can enculturate through apprenticeship. Becoming an apprentice doesn't involve a qualitative change from what JPFs normally do. People enculturate into different communities all the time. The apprentices' behavior and the JPFs' behavior can thus be thought of as pretty much the same.⁶

[Insert Table 1 about here.]

The second and now more conventional option is to enter a school as a student. Schools do seem to demand a qualitative change in behavior, however. What the student is expected to do and what a JPF does are significantly different. The student enters the school culture, while ostensibly being taught something else. And the general strategies for intuitive reasoning, resolving issues, and negotiating meaning that people develop through everyday activity are superseded by the precise, well-defined problems, formal definitions, and symbol manipulation that characterize much school activity.

We try to represent this discontinuity in Table 1, which compares salient features of JPF, practitioner, and putative student behavior:

This table (somewhat like the JPFs' problems) is a little ill defined. But it is intended to help make apparent the important similarity between JPFs' and practitioners' activity. Both have their activities situated in the culture in which they work, within which they negotiate meanings and construct understanding. The issues and problems that they face arise out of, are defined by, and are resolved within the constraints of their belief system and the activity they are pursuing.

Lave's work (1988b) provides a good example of a JPF engaged in authentic activity using the context in which an issue emerged to help find a resolution. The example comes from a study of a Weight Watchers class, whose participants were preparing their carefully regulated meals under instruction:

In this case they were to fix a serving of cottage cheese, supposing the amount laid out for the meal was three-quarters of the two-thirds cup the program allowed. The problem solver in this example began the task muttering that he had taken a calculus course in college. . . . Then after a pause he suddenly announced that he had "got it!" From then on he appeared certain he was correct, even before carrying out the procedure. He filled a measuring cup two-thirds full of cottage cheese, dumped it out on the cutting board, patted it into a circle, marked a cross on it, scooped away one quadrant, and served the rest.

Lave's account nicely brings out the central features of this example:

Thus, "take three-quarters of two-thirds of a cup of cottage cheese" was not just the problem statement but also the solution to the problem and the procedure for solving it. The setting was part of the calculating process and the solution was simply the problem statement, enacted with the setting. At no time did the Weight Watcher check his procedure against a paper and pencil algorithm, which would have produced $3/4 \text{ cup} \times 2/3 \text{ cup} = 1/2 \text{ cup}$. Instead, the coincidence of the problem, setting, and enactment was the means by which checking took place. (p. 165)

The dieter's solution path was extremely expedient and drew on the sort of inventiveness that characterizes the activity of both JPFs and practitioners. It reflected the nature of the activity, the resources available, and the sort of resolution required in a way that problem solving that relies on abstracted knowledge cannot. This inventive problem-solving resolution depended on the dieter seeing

the problem in the context of ongoing activity, which provided him with privileged access to the solution path he chose. (This probably accounts for the certainty he expressed before beginning his calculation.) He was able to see the problem and its resolution in terms of the measuring cup, cutting board, and knife. Activity-tool-culture (cooking-kitchen utensils-dieting) moved absolutely in step throughout this procedure because of the way the problem was seen and the task was performed. The whole micro-routine simply became one more step on the road to a meal.⁷ Knowing and doing were interlocked and inseparable.

This sort of problem solving, carried out in conjunction with the environment, stands quite distinct from the idea of processing solely inside heads to which many teaching practices subscribe. By off-loading part of the cognitive task onto the environment, the dieter automatically used his environment to help solve the problem. In doing this, his actions are not in any way exceptional; they resemble many ordinary JPF and expert practices. Scribner (1984) records, for instance, how complex calculations can be performed by practitioners using their environment directly. In the case she studied, dairy loaders used the configuration of crates they were filling and emptying almost as an embedded abacus. Nor are such problem-solving strategies limited to the physical or social environment. This sort of reliance on situations can be seen in the work of physicists, who see "through" formulae onto the situations of an envisioned world, which then provides support for inferences and approximations (deKleer & Brown, 1984). Hutchins' (in press) study of intricate naval navigation also records the way people engaged in difficult collaborative tasks distribute the burden throughout the environment and the group. The resulting cognitive activity can only be explained in relation to the entire context in which it takes place. "When the context of cognition is ignored," Hutchins observes, "it is impossible to see the contribution of the structure in the environment, in artifacts, and in other people to the organization of mental processes" (p. 31).

Instead of taking problems out of the context of their creation and providing them with an extraneous framework, JPFs and practitioners seem particularly adept at solving them within the framework of the context that produced them. This allows them to share the burden--of both defining and solving the problem--with the task environment as they respond directly to emerging issues. The adequacy of the solution they reach becomes apparent in relation to the role it must play in allowing activity to continue. In the end, the problem, the solution, and the cognition involved in getting between the two cannot be isolated from the context in which they are all fundamentally embedded.

So even though students are expected to behave differently, they inevitably do behave like the JPFs they are and solve most of their problems in their own situated way. Schoenfeld (in press) describes math students using well-known strategies, such as the position of a problem in a particular page or section of a book (where the first questions at the end of chapters are always simple ones, and the last usually demand concepts from earlier chapters) or the occurrence of a particular word in the problem (e.g., "left" signals a subtraction problem), to find solutions quickly and efficiently. Such ploys indicate how thoroughly learners really are situated, and how they always lean on the embedding context for help. Within the culture of schooling this can obviously be very effective. But the school situation is extremely specialized. Viewed from outside, where problems do not come in text books, an unacknowledged dependency on such school-based cues makes the learning extremely fragile.

Conversely, though schooling seeks to encourage problem solving, it disregards most of the inventive heuristics that students bring to the classroom. Instead of deploying such inventiveness to good effect, schools tend to dismiss it out of hand. It thus implicitly devalues not just individual heuristics, which may be fragile, but the whole process of inventively structuring cognition and solving problems. Lave (1988c) describes how some students feel it necessary to disguise effective strategies so that teachers believe the problems have been solved in the school-approved way. In order to suggest how this sort of problem-solving can be deployed (which we attempt later), it will be helpful, first, to have some sort of idea of how cognition is structured by situations.

Structuring Cognition

Authenticity in activity is paramount for learning if conceptual knowledge is not self-contained but, rather, if it is the product of and structured by the activity in which it is developed and deployed; if, in short, not just learning but knowledge itself is situated. It thus needs to be explained in some fashion

how knowledge can be situated. In this section we try to outline, in an admittedly tentative way, the mechanisms whereby situations in the world structure knowledge.

Language, it is now widely recognized, is situated. A proposition, as we pointed out above, can only be correctly interpreted in the context of the situation in which it is used. This, we claim, is also true of knowledge itself. Like an indexical proposition, knowledge "indexes" the situation that produces it, without which it is, in the end, unintelligible. Situations are thus integral components of knowing.

To grasp this idea of indexing, it is perhaps helpful to remember the association of the word *index* with pointing (hence the term *index finger*) and signs. Indexical propositions are situated somewhat in the way that signposts are. Signposts are not universal statements but specifically situated "propositions," whose meaning depends very much on where they stand. To be interpreted correctly, the signpost must be in the right situation. The sign reading Berkeley 5, Oakland 7, San Francisco 15 needs to be quite specifically situated. It relates its component parts by indexing the situation in which it is placed.

Face-to-face conversations enable people to interpret indexical expressions because the indexed features of the situation are immediately available, though people are rarely conscious of their contribution. Their importance becomes apparent, however, in communications between people at a distance. Then indexical expressions become problematic until ways are found to secure their interpretation by situating their reference (see, for instance, Rubin, 1980 on the difference between speech and writing).

Authors of a collaborative work such as this one will recognize the problem if they have ever discussed the paper over the phone. "What you say here" is not a very useful remark. "Here" in this setting needs an elaborate description (such as "page 3, second paragraph, third sentence") and can often lead to long conversations at cross purposes. The problem gets harder in conference calls when "you" becomes as ambiguous as "here" is unclear.

When descriptions replace indexical terms (which is analogous, we hold, to classroom descriptions of activity replacing activity), the nature of discourse changes and understanding becomes much more problematic. Indexical terms are virtually transparent. They, like a signpost, draw little or no attention to themselves. They do not inherently add to the difficulty of understanding a proposition in which they occur, but simply provide essential structure for the discourse. Descriptions, by comparison, are at best translucent and at worst opaque, intruding emphatically between speakers and their subject. (They are less like a signpost and more like a map, which requires extra steps of interpretive work to discover its relationship to the situation.) The audience has first to focus on the descriptions and try to interpret them, then find what they might refer to, and only then can the proposition in which they are embedded be understood. However elaborate, a description does not merely replace the indexical word, just as a map doesn't replace a signpost. The more elaborate the description is in an attempt to be unambiguous, the more opaque it is in danger of becoming. And as Perry argues (1979), in some circumstances, the indexical term simply cannot be replaced.

Knowledge, we suggest, similarly indexes the situations in which it arises and is used, without which it cannot be fully understood. The embedding circumstances efficiently provide essential parts of its structure and meaning. So knowledge, which comes coded by and connected to the activity and environment in which it is developed, is spread across its component parts--some of which are in the mind and some in the world--as the final picture on a jigsaw is spread across its component pieces.

Indexical representations developed through engagement in a task may greatly increase the efficiency with which subsequent tasks can be done, if part of the environment that structures the representations remains invariant. This is suggested by the ability to perform tasks that cannot be described or remembered in the absence of the situation. Recurring features of the environment seem to afford recurrent sequences of actions. Memory and subsequent actions, as knots in handkerchiefs and other *aides memoirs* reveal, are not context-independent. Routines (Agre, 1985) may also be a product of this sort of indexicalization. Thus authentic activity in context becomes a central component of learning.

As Hutchins (in press), Pea (1988), and others point out, the structure of cognition is widely distributed throughout the environment, both social and physical. The environment therefore contributes importantly to indexical representations people form in activity. These representations, in turn,

contribute to the ability to perform future activity efficiently. Recurrent activity--in thoroughly conventional terms, *practice*, but practice in authentic activity--is what ultimately leads from specific, individually situated actions through a process of acting and re-acting, to generalizable knowledge. The generality is not abstract, but situated across multiple contexts.

In language learning, for example, the original frail understanding of a word is developed and extended through subsequent use and social negotiation, though each use is no less situated. Miller and Gildea (1987) describe two stages to this process. The first, in which people learn the word and assign it a semantic category (e.g., the word *olive* is first assigned to the general category of color words), is quickly done. The second, in which distinctions within this semantic category (e.g., between olive and other colors) are explored as the word occurs again and again, is a far more gradual process, which "may never be completely finished" (p. 95). This second phase of word learning corresponds to the development through activity of all conceptual knowledge. The threadbare concepts that initially develop out of activity are gradually given texture as they are deployed in different situations. In the ensuing situations, the appropriateness of the word is tested and, as with all conceptual knowledge, each application has the potential to change an individual's understanding of both the word and the world--the two are interdependent.

Out of this sort of continual activity grows an increasingly more densely textured conceptual knowledge that is enriched by each application in situations. Eagleton (1983) describes reading as just such an enriching process, through which people continually change how they read and this, in turn, changes what they are reading. He begins with the assumption that in order to read, people apply a code to a text. This is analogous to the deployment of a cultural belief system in order to undertake a cultural activity that we described earlier. That belief system is itself constantly changing as a result of the activity it supports:

In applying a code to the text, we may find that it undergoes revision and transformation in the reading process; continuing to read . . . we discover that it now produces a "different" text, which in turn modifies the code by which we are reading it, and so on. This dialectical process is infinite; and . . . it undermines any assumption that once we had identified the proper codes for the text our task was finished. (Eagleton, 1983, p. 125)

This sort of dialectical process is common to all activity, and it similarly undermines any assumption that once conceptual knowledge has been described and applied in exercises, the task of teaching or learning has been finished. Learning is instead a continual process needing above all authentic activity to support its development.

In the following section, we examine the practical ways in which students might be given access to some sort of authentic activity.

Learning Through Cognitive Apprenticeship

Cognitive Apprenticeship

In discussing tools, culture, activity, JPFs, and situated cognition, we have been accumulating a set of characteristics of human learning and reasoning that, we feel, school practices need to honor. Though there are many innovative teachers, schools, and school programs that are exceptions, prevalent school practices still broadly assume that knowledge is individual and self-structured, that schools are neutral with respect to what is learned in them, that concepts are abstract and immutable and independent of the activity in which they are acquired and used, and that JPF behavior is something to be discouraged.

Cognitive apprenticeship (Collins, Brown, & Newman, 1989) methods, by contrast, try to enculturate students into authentic practices through authentic activity and social interaction in a way similar to that which is evident--and evidently successful--in craft apprenticeship. In this section, we examine two examples of mathematics teaching in an attempt to elucidate how some of the characteristics of learning we have discussed can be honored in the classroom.

Schoenfeld's teaching of problem solving--thinking mathematically. Schoenfeld's teaching of problem solving (1985, in press) attempts to show college students how to think mathematically about the world, how to see the world through mathematicians' eyes, and thus how to use the mathematician's tools. By deliberately engaging students in authentic mathematical practice, he reaches well beyond the inculcation of problem-solving skills. His method, much more importantly, provides students with the opportunity to enter the culture of mathematical practice, which enables them to use those skills as practitioners.

As a means to generate spontaneous mathematical thinking and to permit his students to see and engage in mathematics as a sense-making pursuit, Schoenfeld has students bring problems to class that he and they investigate. The approach is distinctive because conventionally, before graduate school, few students get the opportunity to see their teachers engaged in mathematical practice--yet the students are, nonetheless, expected to understand the nature of that practice.

In one case (Schoenfeld, in press), he and his class faced the problem of the magic square (see Figure 1). Though the problem is relatively straightforward, the collaborative work involved in solving it and, importantly, in analyzing the solution helped reveal to the class the way mathematicians look at problems. The class worked collectively through a number of strategies, in which, on reflection, they recognized more general and more powerful mathematical ideas. Thus, from discussing whether 9 can go in the center of the square, they developed the ideas of "focus[ing] on key points that give you leverage," and "exploit[ing] extreme cases." Although Schoenfeld may seem only to be teaching strategy rather than subject matter, he is, more fundamentally, building with his class a mathematical belief system around his own and the class' intuitive responses to the problem.

[Insert Figure 1 about here.]

As an indication that he is working in the culture of mathematics, not in the culture of schooling, he does not stop at what, in culture of school practice, marks the end: an answer.

Are we done? In most mathematics classes the answer is "yes." Early in the semester my students all say "yes," expecting me to go on to another problem. My answer, however, is a resounding "no." In most classes, so-called "problems" are exercises; you are done when you've shown that you've mastered the relevant technique by getting the answer. (Schoenfeld, in press, p. 34)

His class' goal, by contrast, is to understand the mathematical nature of magic squares, and it is in part by doing this that the belief system is exemplified. The class went on to explore other possible magic squares and thereby discovered general principles (e.g., an algebraic form for describing them). This process also led to some further generalizable mathematical strategies that are less commonly seen in classroom practice, such as working forwards from an initial solution, using systematic generating procedures, and having more than one way to solve a problem. Schoenfeld is careful to emphasize that all these strategies are illustrated in action, developed by the class, not declared by the teacher. The belief system is instilled in the only way it can be, through practice in which the students actively take part.

Lampert's Teaching of Multiplication

Lampert (1986) also involves her students in mathematical exploration, which she tries to make continuous with their everyday knowledge. She has devised methods for teaching mathematics to fourth-grade students that lead from students' implicit understanding of the world beyond the classroom, through activity and social construction in the culture, to the sort of robust learning that direct teaching of algorithms usually fails to achieve.

She starts teaching multiplication, for example, in the context of coin problems, because in the community of fourth-grade students, there is usually a strong, implicit, shared understanding of coins; next she has the students create stories for multiplication problems, drawing on their implicit knowledge to delineate different examples of multiplication. Then she helps them toward the abstract algorithm that everyone learns for multidigit multiplication, in the context of the coin problems and stories the

community has created. Thus, Lampert presents the algorithm as one more useful strategy to help them resolve community problems.

[Insert Figure 2 about here.]

The first phase of Lampert's teaching starts with simple coin problems, such as "Using only nickels and pennies, make 82 cents." With these problems, Lampert helps her students explore their implicit knowledge. Then, in the second phase, her students create stories for multiplication problems (see Figure 2). The students perform a series of decompositions and discover that there is no one magically "right" decomposition decreed by authority, just more and less useful decompositions, whose *use* is judged in the context of the problem to be solved and the interests of the problem solvers.

The third phase gradually introduces students to the standard algorithm, now that such an algorithm has a meaning and a purpose in their community. The students start with an extended procedure that parallels the story problem above. Eventually they find ways to shorten the process, and they usually arrive at the standard algorithm, justifying their findings with the stories they had created earlier.

Lampert hopes to develop a composite understanding of four different kinds of mathematical knowledge: (a) *intuitive knowledge*, the kind of short cuts people invent when doing multiplication problems in authentic settings; (b) *computational knowledge*, the basic algorithms that are usually taught; (c) *concrete knowledge*, the kind of concrete models of the algorithm associated with the stories they created; (d) *principled knowledge*, the principles such as associativity and commutativity that underlie the algorithmic manipulations of numbers. She tries to inculcate an inseparable understanding of all of these kinds of knowledge and the connections between them, and thus to bridge the huge gap that emerges from much conventional teaching between conceptual knowledge and problem-solving activity--between, as we characterized them at the beginning, knowing and doing.

Lampert's approach exhibits several qualities of social construction and situated cognition that exemplify cognitive apprenticeship:

by beginning with a task embedded in a familiar activity, it shows the students the legitimacy of their implicit knowledge and its availability as scaffolding in apparently unfamiliar tasks.

by pointing to different decompositions, it stresses that heuristics are not absolute, but assessed with respect to a particular task--and that even algorithms can be assessed in this way.

by allowing students to generate their own solution paths, it helps make them conscious, creative members of the culture of problem-solving mathematicians. And, in enculturating through this activity, they acquire some of the culture's tools--a shared vocabulary and the means to discuss, reflect upon, evaluate, and validate community procedures in a collaborative process.

Schoenfeld's approach, on the other hand, points in particular to the way in which students can be exposed to the authentic ways of thinking of a culture and its conceptual viewpoint, as much as to its subject matter.

In the terms of cognitive apprenticeship, we can represent the progress of the students from embedded activity to general principles of the culture as presented in Figure 3.

[Insert Figure 3 about here.]

In this sequence, apprenticeship and coaching begin by providing modeling *in situ* and scaffolding for students to get started in an authentic activity. As the students gain more self-confidence and control, they move into a more autonomous phase of collaborative learning, where they begin to participate consciously in the culture. The social network within the culture helps them develop its language and the belief systems and promotes the process of enculturation. Collaboration also leads to the articulation or manifestation of strategies, which can then be discussed and reflected on. This in turn fosters generalizing, grounded in the students' situated understanding.

From here, students can, as we have suggested, use their fledgling conceptual knowledge to see their activity in a new light, which in turn leads to the further development of the conceptual knowledge, and so on. This, then, suggests a helical process (Figure 4), rather than the linear process outlined in Figure 3.

[Insert Figure 4 about here.]

Cognitive apprenticeship attempts to develop densely textured concepts out of and through continuing authentic activity. The term is closely allied to our image of knowledge as tool. Cognitive apprenticeship supports learning in a domain by enabling students to acquire, develop, and use conceptual tools in authentic domain activity, just as craft apprenticeship enables apprentices to acquire and develop the tools and skills of their craft through authentic work and membership of their trade. Through this process, apprentices enter the culture of practice. So the term *apprenticeship* helps to emphasize the centrality of activity in learning and knowledge and highlights the inherently context-dependent, situated, and enculturating nature of learning. And *apprenticeship* also suggests the paradigm of situated modeling, coaching, and fading (Collins, Brown, & Newman, 1989), whereby teachers or coaches promote learning first by modeling their strategies for students in authentic activity. Then teachers and colleagues support a student's attempts at doing the task. And finally they empower the student to continue independent of the teacher. The progressive process of learning and enculturation perhaps argues that Burton, Brown, and Fischer's (1984) "increasingly complex microworlds" should be replaced by "increasingly complex enculturating environments."

Cognitive emphasizes that apprenticeship techniques actually reach well beyond the physical skills usually associated with apprenticeship to the kinds of cognitive skills more normally associated with conventional schooling. This extension is not as incompatible with traditional apprenticeship as it may at first seem. The physical skills usually associated with apprenticeship embody important cognitive skills if our argument for the inseparability of knowing and doing is correct. Certainly many professions with generally acknowledged cognitive content--for example, law, medicine, architecture, business, etc.--have nonetheless traditionally been learned through apprenticeship. Moreover, advanced graduate students in the humanities, the social sciences, and the physical sciences, acquire their extremely refined research skills through the apprenticeship they serve with senior researchers. Then they, like all apprentices, must recognize and resolve the ill-defined problems that issue out of authentic activity, in contrast to the well-defined exercises that are typically given to them in text books and on exams throughout their earlier schooling. It is at this stage, in short, that students no longer behave as archetypal students but participate in the activity of practitioners and develop their conceptual understanding through social interaction and collaboration in the culture of the domain, not of the school.

Social interaction and collaboration play such a central role in this sort of learning that before concluding this discussion of cognitive apprenticeship, we should highlight their importance. And we should emphasize that though we have been explicitly discussing school education, features of this form of learning, in particular "legitimate peripheral participation" and collaboration are particularly relevant to workplace training as well as school learning.

Social interaction and the social construction of knowledge. Lave (1988a) and Lave and Wenger (in preparation) point out the importance for apprentices of learning their craft in the appropriate community of practice, towards the center of which they continuously progress through gradual enculturation. The apprentice tailors Lave studied are surrounded both by master tailors and by other apprentices, all engaged with similar tools in similar authentic activity. This allows apprentices to observe and gradually participate fully in the practices of the community. The apprentices learn to use the relevant tools in the context of their use within the belief system that gives purpose and meaning to the tasks they undertake. The knowledge of the community is evidently constructed, acquired, developed, distributed, and validated through intense social interaction.

This sort of social interaction should not be seen as a facet of some distant and exotic community. Resnick has pointed out (1988) that throughout most of their lives people learn and work collaboratively, not individually, as they are asked to do in many schools. Lampert's and Schoenfeld's work described above, Scardamalia, Bereiter, and Steinbach's (1984) teaching of writing, and Palincsar and Brown's

(1984) work with reciprocal teaching of reading all employ some form of social interaction, social construction of knowledge, and collaboration.

Within a culture, ideas are exchanged and modified and belief systems developed and appropriated in part through conversation and narratives, so these must be promoted, not inhibited. Though these are often anathema to traditional schooling, they are an essential component of social interaction and thus of learning. They provide access to much of the distributed knowledge and support of the social matrix (Orr, 1987). So learning environments must allow narratives to circulate and "war stories" to be added to the collective wisdom of the community.

The role of narratives and conversations is perhaps more complex than might at first appear. An intriguing role in learning is played by "legitimate peripheral participation," whereby people who are not directly taking part in a particular activity learn a great deal from their legitimate position on the periphery (Lave & Wenger, in preparation). It is a mistake to think that important discourse in learning is always direct and declarative. This peripheral participation is particularly important for people entering the culture. They need to observe how qualified practitioners behave and talk to get a sense of how expertise is manifest in conversation and other activities.

Cognitive apprenticeship and collaborative learning. If, as we propose, learning is a process of enculturating that is supported in part through social interaction and the circulation of narrative, groups of practitioners are particularly important, for it is only within groups that social interaction and conversation can take place.

Some of the most salient features of group learning include:

- a. *Collective problem solving.* Groups are not just a convenient way to accumulate the individual knowledge of their members. They give rise synergistically to insights and solutions that would not come about without them (Schoenfeld, in preparation).
- b. *Displaying multiple roles.* Successful execution of most individual tasks requires students to understand the many different roles needed for carrying out that task. Getting one person to be able to play all the roles entailed by authentic activity and to reflect productively upon his or her performance is one of the monumental tasks of education. The group, however, permits different roles to be displayed and engenders reflective narratives and discussions about the aptness of those roles.
- c. *Confronting ineffective strategies and misconceptions.* We know from an extensive literature on the subject (diSessa, 1982, 1983, 1986; McCloskey, Caramazza, & Green, 1980; White, 1983) that students have many misconceptions about qualitative phenomena in physics. Teachers rarely have the opportunity to hear enough of what students think to recognize when the information that is offered back by students is only a surface retelling for school purposes (the handing back of an uncomprehended tool, as we described it at the beginning) that may mask deep misconceptions about the physical world and problem-solving strategies. Groups, however, can be an efficient means to draw out, confront, and discuss ineffective strategies.
- d. *Providing collaborative work skills.* Students who are taught individually rather than collaboratively can fail to develop skills needed for collaborative work. In the collaborative conditions of the workplace, knowing how to learn and work collaboratively is increasingly important. If people are going to learn and work in conjunction with others, they must be given the situated opportunity to develop those skills.

In looking at Schoenfeld's and Lampert's teaching, in noting what we believe are particularly important features of their methods, and in stressing in particular social interaction and collaborative learning, we are trying to show how teaching through a form of apprenticeship can accommodate the new view of knowledge and learning we have been outlining. The increasing role of the teacher as a master to apprentices, and the teachers' use of authentic domain activity as a major part of teaching will perhaps

once and for all dismiss Shaw's scurrilous criticism of teachers with which we opened. His comment may then be replaced with Alexander Pope's more admirable wish to

Let such teach others who themselves excel

Conclusion: Toward an Epistemology of Situated Cognition

A great deal of work in investigating situated features of cognition remains to be done. It is, however, already possible to point to areas of education already under serious reappraisal (see, for example, Resnick, 1988; Shaker, 1988) for which continued research in this area may be particularly relevant.

One of the particularly difficult challenges for education (which many exceptional teachers may have independently solved) is how to separate what should be made explicit in teaching from what should be left implicit. A common strategy in trying to overcome difficult pedagogic problems is to make as much as possible explicit. Thus we have ended up with wholly inappropriate methods of teaching. Whatever the domain, explication lifts implicit and possibly even nonconceptual constraints (Cussins, 1988) out of the world and tries to make them explicit or conceptual. These then occupy a place in our ontology and become something more to learn about rather than simply something useful in learning. In contrast, indexical representations seem to gain their efficiency by leaving a great deal of the context underrepresented or implicit. Future work into situated cognition, from which educational practices will benefit, must, among other things, try to frame a convincing account of the relationship between explicit knowledge and implicit understanding.

We have described here only part of the implications of a fully developed theory of situated cognition. The major theoretical work to shift the traditional focus of epistemology remains to be done. For centuries epistemologists have concentrated primarily on conceptual representation and its problematic relation to objects in the world, assuming that representation is cognitively prior to all else. This has led them to battle with the seemingly intractable question of reference--the problematic alignment of conceptual representations of the world and objects in the world. A theory of situated cognition suggests that activity and perception are importantly and epistemologically prior--at a nonconceptual level--to conceptualization and that it is on them that more attention needs to be focused. An epistemology that begins with activity and perception, which are first and foremost embedded in the world, may simply bypass the classical problem of reference--of mediating conceptual representations.

In educational terms alone, however, the unheralded importance of activity and enculturation to learning suggests that much common practice is the product of an inadequate epistemology. This further suggests that a new epistemology might hold the key to a dramatic improvement in learning and completely new insights into education.

References

- Agre, P. (1985). *Routines*. Cambridge, MA: MIT AI Memo 828.
- Barwise, K. J., & Perry, J. (1983). *Situations and attitudes*. Cambridge, MA: MIT Press.
- Burton, R., Brown, J. S., & Fischer, G. (1984). Skiing as a model of instruction. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its development in social context* (pp. 139-150). Cambridge, MA: Harvard University Press.
- Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive apprenticeship: Teaching the craft of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning and instruction: Essays in honor of Robert Glaser*. Hillsdale, NJ: Erlbaum.
- Cussins, A. (1988). *The connectionist construction of concepts* (SSL Research Report). Palo Alto, CA: Xerox Palo Alto Research Center.
- deKleer, J., & Brown, J. S. (1984). A qualitative physics based on confluences. *Artificial Intelligence Journal*, 24, 1-3.
- diSessa, A. (1986). Knowledge in pieces. In G. Forman & P. Pufall (Eds.), *Constructivism in the computer age*. Hillsdale, NJ: Erlbaum.
- diSessa, A. (1983). Phenomenology and the evolution of intuition. In D. Gentner & A. Stevens (Eds.), *Mental models* (pp. 15-33). Hillsdale, NJ: Erlbaum.
- diSessa, A. (1982). Unlearning Aristotelian physics: A study of knowledge-based learning. *Cognitive Science*, 6, 37-75.
- Dreyfus, H., & Dreyfus, S. (1986). *Mind over machine*. New York: The Free Press.
- Eagleton, T. (1983). *Literary theory*. Oxford: Basil Blackwell.
- Engestrom, Y. (1987). *Learning by expanding*. Helsinki: Orienta-Konsultit Oy.
- Fillmore, C. (1974). *Santa Cruz Lectures on Deixis*. Bloomington: Indiana University Linguistics Club.
- Geertz, C. (1983). *Local knowledge*. New York: Basic Books.
- Haroutunian-Gordon, S. (1988). Mind over machine. *Educational Researcher*, 17(3), 50-53.
- Hutchins, E. (in press). Learning to navigate. In S. Chalkin & J. Lave (Eds.), *Situated learning*.
- Lampert, M. (1986). Knowing, doing, and teaching multiplication. *Cognition and Instruction*, 3, 305-342.
- Lave, J. (1977). Tailor-made experiments and evaluating the intellectual consequences of apprenticeship training. *The Quarterly Newsletter of the Institute for Comparative Human Development*, 1, 1-3.
- Lave, J. (1988a). *The culture of acquisition and the practice of understanding* (IRL Report 88-0007). Palo Alto, CA: Institute for Research on Learning.
- Lave, J. (1988b). *Cognition in practice*. Boston, MA: Cambridge University Press.
- Lave, J. (1988c). *Word problems: A microcosm of theories of learning*. Paper presented at AERA annual conference, New Orleans.

- Lave, J. (in preparation). *Tailored learning: Education and everyday practice among craftsmen in West Africa*.
- Lave, J., & Wenger, E. (in preparation). *Situated learning: Legitimate peripheral participation*.
- McCloskey, M., Caramazza, A., & Green, B. (1980). Curvilinear motion in the absence of external forces: Naive beliefs about the motion of objects. *Science*, 210, 1139-1141.
- Miller, G. A., & Gildea, P. M. (1987). How children learn words. *Scientific American*, 257(3), 94-99.
- Nunberg, G. (1978). *The pragmatics of reference*. Bloomington: Indiana University Linguistics Club.
- Orr, J. (1987). *Talking about machines*. SSL Report. Palo Alto, CA: Xerox Palo Alto Research Center.
- Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and monitoring activities. *Cognition and Instruction*, 1, 117-175.
- Pea, R. D. (1988). *Distributed intelligence in learning and reasoning processes*. Paper presented at the meeting of the Cognitive Science Society, Montreal.
- Perkins, D. N. (1986). *Knowledge as design*. Hillsdale, NJ: Erlbaum.
- Perry, J. (1979). The problem of the essential indexical. *Nous*, 13, 3-21.
- Resnick, L. (1988). Learning in school and out. *Educational Researcher*, 16(9), 13-20.
- Rogoff, B., & Lave, J. (Eds.). (1984). *Everyday cognition: Its development in social context*. Cambridge, MA: Harvard University Press.
- Rubin, A. (1980). A theoretical taxonomy of the differences between oral and written language. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 411-438). Hillsdale, NJ: Erlbaum.
- Scardamalia, M., Bereiter, C., & Steinbach, R. (1984). Teachability of reflective processes in written composition. *Cognitive Science*, 8, 173-190.
- Schoenfeld, A. H. (in preparation). *Ideas in the air* (IRL Report 88-0011). Palo Alto, CA: Institute for Research on Learning.
- Schoenfeld, A. H. (in press). On mathematics as sense-making: An informal attack on the unfortunate divorce of formal and informal mathematics. In D. N. Perkins, J. Segal, & J. Voss (Eds.), *Informal reasoning and education*. Hillsdale, NJ: Erlbaum.
- Schoenfeld, A. H. (1984). *Mathematical problem solving*. Orlando, FL: Academic Press.
- Scribner, S. (1984). Studying working intelligence. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its development in social context* (pp. 9-40). Cambridge, MA: Harvard University Press.
- Shanker, A. (1988). Exploring the missing connection. *New York Times*, 6.19.1988, E7.
- White, B. (1983). Sources of difficulty in understanding Newtonian dynamics. *Cognitive Science*, 7, 41-65.
- Whitehead, A. N. (1929). *The aims of education*. Cambridge: Cambridge University Press.

Footnotes

¹All work in this area is to a greater or lesser degree built upon foundational research of activity theorists such as Vygotsky, Leontiev, and others. For examples of recent work upon which we have drawn, see Rogoff and Lave (Eds.), 1984; Scribner, 1984; Hutchins, in press; Engestrom, 1987; Lave and Wenger, in preparation; and in particular Lave, 1977, 1988a, 1988b, 1988c, in preparation. Anyone familiar with Jean Lave's work on learning, apprenticeship, and everyday cognition will realize at once that we are deeply indebted to her groundbreaking work.

²This paper does not attempt to produce a historical survey or to trace the genesis of ideas that we discuss. And our purpose is far from giving a critical analysis. We hope merely to present the flavor of a situated theory of cognition and to suggest what it might imply, particularly for education.

³The excerpted dictionary definitions that led the students to these sentences are as follows: *Correlate*: be related one to the other; *meticulous*: very careful; *redress*: remedy; *relegate*: send away; *stimulate*: stir up; *tenet*: doctrine held as true; *erode*: eat out. As they were given these definitions with little or no contextual help, it would be quite unfair to regard the students as foolish for using the words in this way.

⁴In linguistics literature, accounts of indexicality usually use the term *deixis*. See for example Fillmore (1974).

⁵This image is, of course, not original. For the way it is developed here, though, we are particularly indebted to Richard Burton, who explored it during a symposium organized by the Secretary of Education of Kentucky and to D. N. Perkins' book *Knowledge as Design* (1986).

⁶The JPFs must, of course, have access to a culture and become what Lave and Wenger (in preparation) call a "legitimate peripheral participant" in that culture. And, of course, as an apprentice, a JPF usually has to do a great deal of work. We are not trying to suggest anything magical occurs in the process of enculturation. Medical interns testify to exactly how hard it can be. But the process, we stress, is not qualitatively different from what people do all the time, adopting the behavior and belief systems of their peers.

⁷To get some sense of how foreign this is to school tasks, it might be useful to imagine the impropriety of a student given this as a mathematical problem asking "Does the dieter have a measuring cup, cutting board, and knife at hand?" Though word problems are meant to ground theory in activity, the things that structure activity are denied to the problem solvers. Textbooks ask students to solve supposedly "real-life" questions about people who do very unreal things, such as driving at constant speeds in straight lines or filling leaking troughs with leaking buckets. Students are usually not allowed to indulge in real-life speculation. Their everyday inventiveness is constrained by prescribing and proscribing ways in which the solution must be found. The ubiquitous Mr. Jones might, after all, wisely repair the hole in his bucket or fill the trough with a hose. Sitting down and calculating how many journeys it will take with a leaking bucket is probably the very last thing he would do (see also Lave, 1988c).

Author Notes

Many of the ideas of this paper emerged from group discussions at the Institute for Research on Learning. In particular we are grateful to Jim Greeno, Jean Lave, Susan Newman, Roy Pea, and John Rheinfrank who read earlier drafts and commented upon them with great care. We are also grateful to Richard Burton, Bill Clancey, and Alan Schoenfeld for helpful and insightful contributions. More generally, we would like to acknowledge the influence of Brian Cantwell Smith's pioneering research into a theory of computation and semantics, built on notions of situatedness, embeddedness, and embodiedness; of Susan Stucky's important new idealization of mind in terms of "radical" efficiency rather than rationality; and also of the work on indexicality of Philip Agre and David Chapman.

Table 1**JPF, Practitioner, and Putative Student Learning Activity**

	JPFs	Students	Practitioners
reasoning with:	causal stories	laws	causal models
acting on:	situations	symbols	conceptual situations
resolving:	emergent problems and dilemmas	well-defined problems	ill-defined problems
producing:	negotiable meaning & socially constructed understanding	fixed meaning & immutable concepts	negotiable meaning & socially constructed understanding

Figure Captions

Figure 1. The Magic Square Problem (from Schoenfeld, in press).


Figure 2. Story Problems for Teaching Multiplication (from Lampert, 1986).

Figure 3. The Progress of Students in Cognitive Apprenticeship.

Figure 4. The Helical Process of Learning in Cognitive Apprenticeship.

Can you place the digits 1,2,3,4,5,6,7,8,9 in the box to the right, so that the sum of the digits along each row, each column, and each diagonal is the same? The completed box is called a magic square.

Figure 1



T: Can anyone give me a story that could go with this multiplication... 12×4 ?

S1: There were 12 jars, and each had 4 butterflies in it.

T: And if I did this multiplication and found the answer, what would I know about those jars and butterflies?

S1: You'd know you had that many butterflies altogether.

T: Okay, here are the jars. [Draws a picture to represent the jars of butterflies--see diagram.] The stars in them will stand for butterflies. Now, it will be easier for us to count how many butterflies there are altogether, if we think of the jars in groups. And as usual, the mathematician's favorite number for thinking about groups is?

S2: 10

T: Each of these 10 jars has 4 butterflies in it. [Draws a loop around 10 jars.] ...

T: Suppose I erase my circle and go back to looking at the 12 jars again altogether. Is there any other way I could group them to make it easier for us to count all the butterflies?

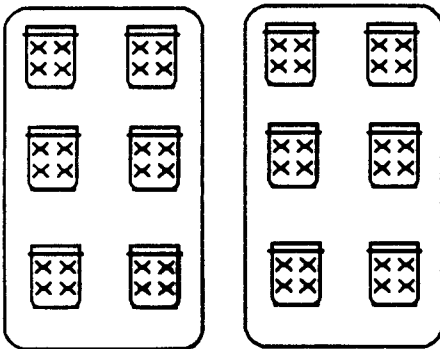
S6: You could do 6 and 6.

T: Now, how many do I have in this group?

S7: 24

T: How did you figure that out?

S7: 8 and 8 and 8. [He puts the 6 jars together into 3 pairs, intuitively finding a grouping that made the figuring easier for him.]



T: That's 3×8 . It's also 6×4 . Now, how many are in this group?

S6: 24. It's the same. They both have 6 jars.

T: And now how many are there altogether?

S8: 24 and 24 is 48.

T: Do we get the same number of butterflies as before? Why?

S8: Yeah, because we have the same number of jars and they still have 4 butterflies in each.

Figure 2




Figure 3




Figure 4

This page is intentionally blank.

