Digging for Bedrock: A casual inspection of the epistemological foundations of three disciplines

Eckehard Doerry

CIS-TR-94-25
November 1994

Department of Computer and Information Science
University of Oregon

Digging for Bedrock:

A casual inspection of the epistemological foundations of three disciplines

Eckehard Doerry
Dept. of Computer and Information Science
University of Oregon
Eugene, OR. 97403
eck@cs.uoregon.edu

" 'I'm afraid that's wrong', says the quiz master, 'that's not the answer I have on the card'"
— Robin A. Hodgkin, *Playing and Exploring*

Abstract

In the last decade we have seen the boundless enthusiasm for tutoring humans using "intelligent" machines founder, as system after system has failed to live up to expectations. For the most part, such systems can be seen to rest on similar philosophical foundations, conceptions of mind and thought rooted firmly in the Rationalist tradition that forms the cornerstone of modern science. This foundation is two-tiered, centered around an epistemology that characterizes knowledge as internalized and finite symbolic representations of reality and a metaphysics that assumes the existence of absolute truth, not only where physical reality is concerned, but also in terms of semantics[1]. Evidence of the Representationalist mindset can be found in the underlying assumptions of many areas of science, education, and philosophy throughout history, but is especially apparent in Artificial Intelligence (AI), where the status of knowledge and meaning have become the central issues. Examples of such assumptions found in AI, and in tutoring systems in particular, include:

- The notion that meaning can exist outside of context. That meaning is absolute and finite and can therefore be described symbolically.
- That human knowledge exists as a finite, concrete interpretable structure in the brain.
- That "not knowing" is generally a problem of content. Symbolic forms are semantically deterministic and so learning is reduced to the definition of new facts/relationships.
- A metaphor of learning based on knowledge transference; the teacher the font of knowledge and the student a vessel to be filled.

[1] More accurately, unambiguous semantics arise as *an integral part* of an absolute symbolic reality.

Recent work by both sociologists and educators suggests that the assumptions listed above are erroneous and situate the source of meaning firmly in its context of use. Among other things, this implies that learning can not be separated from acting physically and (especially) linguistically.

Introduction

Despite the fact that new tutoring systems are being created every day and that research in Intelligent Tutoring Systems (ITS) continues at a brisk pace, it is increasingly difficult to deny the obvious: machine-based tutoring is at an impasse. Of all the systems and approaches that have been suggested, very few have ever found their way into the classroom, and almost none have ever been rigorously tested in the real world. At best, anecdotal evidence of short-term trials is presented to support the claims of breakthrough by system designers. Looking at the field of ITS, we find it to be well-ordered with research proceeding along well-defined lines, investigating pedagogic strategies, novel representations of student and expert knowledge, modeling explanation and discourse. For those set solidly in this engineering mindset, it is only a matter of finding the right knowledge representation, understanding spoken speech, designing the right interface — in general, a matter of solving just a few more technical problems to bridge the gap between human and machine.

However, computer scientists are only the most recent thinkers to try their hand at unraveling one of the most fascinating conundrums in all of human experience: How do we, as humans, find meaning in the real world, how does that meaning arise, and how do we support the growth of similar conceptions (i.e. teach) in others? Philosophers have tangled with the metaphysical and epistemological issues that form the core of this conundrum for centuries. More recently, sociologists, cultural anthropologists and professional educators have all suggested novel approaches to the problem, each based on the particular perspective of learning and knowing natural to their discipline. Education, in particular, is currently undergoing some fundamental changes due, presumably, to the failure of traditional methods, and thus provides a rich source of new ideas about how humans learn and understand.

In sum, then, human learning and knowing has been the focus of research in a variety of disciplines, all of which have been in existence far longer than computer science, much less the area of Intelligent Tutoring Systems. In an earlier paper (Doerry, 1994), I argued that the underlying source of troubles in ITS lies in the unthinking allegiance to certain philosophical commitments, in particular, the philosophy of Representationalism, which serves as the foundation for much of modern science. The goal of this paper is not so much to suggest solutions as it is to broaden the perspective. Accordingly, I begin by characterizing and comparing the epistemological foundations of Artificial Intelligence, Ethnomethodology, and the education sciences, working to describe the conception of knowledge and the sources of meaning in each discipline. This framework is then elaborated with a characterization of the dynamic learning process under each perspective. With this basis firmly in hand, I move to an examination of the epistemological foundations of several "alternative" approaches to teaching that have been suggested recently. In the final sections I speculate on the roles that might be played by machine-based tutors in the future, based on our revised conception of teaching and knowing.

1.0 Epistemology

Pedagogical approaches come in almost infinite varieties, from Socrates to the solemn lecture halls of Great Britain to the latest advances in intelligent machine tutoring. With such a wide range of techniques in existence, it is difficult to decide on a framework for comparison and evaluation. Which approach holds the most promise for future development? The solution is to adopt a more abstract analysis by observing that all such techniques have one thing in common: they seek to "teach" something to the student. In other words, the goal of all pedagogy is to *change* the student in some way, so that the student has a more robust conception of the subject matter after the session than before. Thus, characterizing the various conceptions of *what* it is that is changed through the teaching process and *how* that change is brought about constitutes the analytic cornerstone of this paper. An implicit claim in the sections that follow is that this epistemological perspective, how one conceives of knowledge and how it is manipulated has a profound effect on how the pedagogical task is approached, where the presumed areas of difficulty lie, how communication breakdown is characterized, and on the perceived limits of what can be directly explained.

Finally, though the philosophies discussed here are associated with various fields, namely, Artificial Intelligence, Ethnomethodology and Education Sciences, this association is prototypical in nature only. Indeed, a major theme of this paper is that it is the underlying epistemologies that are important, regardless of the field. For example we will discover that the philosophy of Representationalism, here associated with artificial intelligence, can also be seen to underlie more traditional approaches in education as well as much of psychology and cognitive science.

1.1 Traditional ITS: Representationalism

The central tenet of Representationalism is that we, as humans, carry inside of our heads a symbolic model, or representation, of the world, which serves as the basis for all reasoning and action that we perform. This notion is illustrated in figure 1.1.

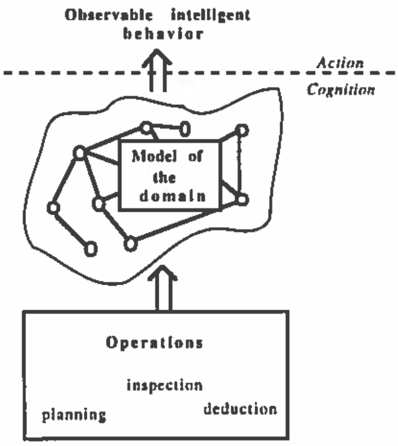

Figure 1.1: Representationalist model of behavior

An interesting feature of this approach is that, since knowledge is packaged into discrete symbolic units, it is trivial to draw dividing lines between knowledge, based on what each chunk of knowledge is "about". For example, it is common in artificial intelligence to distinguish between "domain knowledge", expertise in the area that the system is designed to address, and "common sense knowledge", which is knowledge about all other aspects of reality.

In most cases, the behavior of an intelligent system is influenced (if not driven) by information external to the reasoning system. Under a Representationalist model, this is accounted for by allowing the system to accept any number of external inputs. However, the number of such inputs (though perhaps quite large) must always be finite and pre-defined and, more importantly, the significance ascribed to an input (i.e. how it may influence the outcome of reasoning) is deterministic and pre-defined by how it is integrated in the overall symbolic structure. This observation leads us to the heart of the Representationalist paradigm, the underlying assumption that the meaning of action can be determined succinctly, monotonically, and independent of a context of use.

1.2 Ideas from Education Sciences: Mixed Metaphors

As with many other disciplines, education is primarily a pragmatic field, concerned with discovering approaches that work, rather than the articulation of the epistemological foundations of those theories. Thus it is difficult (and perhaps dangerous) to posit a foundation by inference alone. In this section, I rely mostly on the work of Hodgkin (1985) who, in turn, draws heavily on Polanyi (1958) for her epistemological foundations. As we shall see in section two, recent work in education science can be seen to rest squarely on these same foundations.

Polanyi's conception of human knowledge (Polanyi, 1958) is two-tiered, with an upper tier of "articulate" knowledge resting on a foundational layer of "tacit" knowledge. The arrangement is illustrated in figure 1.2.

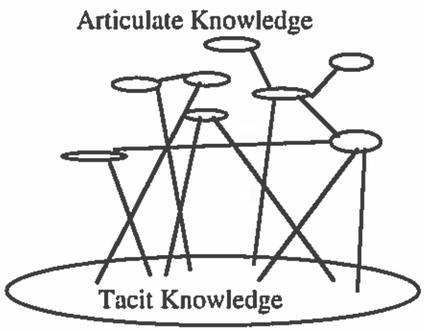

Figure 1.2: Polanyi's two-tiered model.

In the diagram, the unstructured lower layer represents the tacit knowledge, an amorphous amalgam of information that has been *uniquely* patterned by heredity and experience. Polanyi refers it (in the context of performing skilled action) as follows:

> I have pursued the roots of personal knowledge to its most primitive forms which lie behind the operations of a scientific formalism. Tearing away the paper screen of graphs, equations, and computations, I have tried to lay bare the inarticulate manifestations of intelligence by which we know things in a purely personal manner. I have entered on an analysis of the arts of skilful doing and skilful knowing, the exercise of which guides and accredits the use of scientific formulae, and which ranges further afield, unassisted by the formalism, in shaping our fundamental notions of most things which make our world." (Polanyi, 1958, p.64)

As evidenced in the quote, Polanyi feels that tacit knowledge precedes all other forms of knowing, in particular, the articulate knowledge which comprises the upper layer of the model. It is this tacit layer, a compilation of all past experience, that serves as the primary resource for ascribing meaning to the everyday experiences and action in the world. Note that, unlike the representational system explored earlier, the tacit knowledge is entirely amorphous. That is, the experiences are somehow stored in a "raw" form, with no interpretation or structure defined for them[2]. Only in the articulate layer of knowledge does the model allow for "other, more explicit and articulate kind of knowledge which dominates our books and conversations" (Hodgkin, p. 6). Specifically, Polanyi indicates that the articulate layer comes about by essentially compiling and naming common patterns of action.

1.3 Ideas from Ethnomethodology: Situated Action

The emergence of Ethnomethodology as a distinct branch of sociology in the last decades has introduced a new perspective on human knowing and learning, based heavily on the ideas of Martin Heidegger. This has not always been the case. It can be argued that sociology, as a science, was founded by Franz Boas early in this century. As a former mathematician, Boas was a champion of rationalism and, in particular, empiricism, a fundamental axiom of the hard sciences, which states that the only things that can be considered true are those that can be directly observed or that can be formally deduced from such observations. On this view, all information outside of this boundary is considered meaningless speculation. In his investigation of values, Pirsig (1991) quotes Margaret Mead, a pupil of Boas, speaking about Boas: "He feared generalization like the plague, and continually warned us against it. Generalization should be based on the facts and only the facts". As a result of Boas' empiricist mindset and in an effort to attain a semblance of the detached, formal status accorded the hard sciences, sociology and related sciences, such as cognitive science and psychology, have placed a heavy emphasis on formally proving their hypotheses. Unfortunately, formal proof generally involves representation of the domain as a set of the sort of " hard facts" that Boas was concerned with. Consequently, the sciences of the human mind and culture have been led by this thirst for rationalist validity to adopt the same Representationalist foundations that support artificial intelligence.

The establishment of Ethnomethodology, however, marks a breaking away from this rigid rationalist perspective. The premise is that human behavior is not amenable to formal logical analysis and, more importantly, that by shoe-horning human behavior into the Representationalist model, we obscure the true nature of human intelligence. Suchman argues for a dynamic, situated view of human behavior in which the meaning of a situation is constructed as the events unfold. Even then, the meaning of a sequence of the events is continuously reassessed in light of further events and may therefore change in

[2] Viewed in light of insights reached later in the paper, it might be more accurate to say that such experiences retain an *infinite* potential for interpretation.

the future. Pirsig summarizes the objection as follows: "The trouble was that man isn't suited to this kind of [Representationalism-based] study. Objects of scientific study are supposed to hold still. They're supposed to follow the laws of cause and effect in such a way that a given cause will always have a given effect, over and over again. Man doesn't do this. Not even savages." (Pirsig, 1991, p.61)

Suchman's conception of knowledge is much the same as Polanyi's "tacit" knowledge, described earlier: an unstructured record of experience, with no structure or meaning imposed on it at all.

According to Suchman, this record of experience serves as the primary resource for making sense of actions we perform or events we perceive in the world. Another critical resource, on this view, is the experience of others is the society, as communicated through interaction. On this view, symbols (e.g. words) are simply linguistic tools used to rationalize and objectify action retrospectively. That is, symbolic models are applied descriptively, after the action has occurred, as illustrated in figure 1.3. Clearly this perspective reverses the generative process suggested by the Representationalist model depicted in figure 1.1.

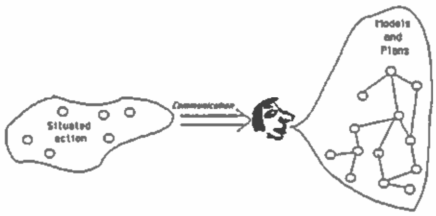

Figure 1.3: Symbolic structures as descriptive tools for objectifying action.

1.4 Discussion: A Focus on Semantics

By working to describe what it means to know under each of the three perspectives, I have focused on relatively mechanical aspects of the three perspectives: *what* knowledge is conceived to be on each of the three views. This descriptive effort obscures a number of subtle pitfalls, especially for those steeped in the Representationalist tradition. First, one might be led to believe that the knowledge models presented here are somehow *more* than illustrative metaphors, that they actually exist in the brain in some form or the other. For instance, this discussion might be construed as an argument for connectionism or other "distributed" forms of knowledge representation. However, the debate here is not about appropriate representations, but rather a critical look at the appropriateness of **all** symbolic representation as a basis for intelligence. A second, even more insidious consequence of this discussion is that the focus on description[3] might be seen to imply that the significance of knowledge — its *meaning* — lies in the structure. In fact, neither Polanyi's tacit knowledge nor Suchman's amorphous pool of experience *have* any sort of structure (in the linguistic or symbolic sense) and, as we shall see, do not tie meaning to static structure at all.

Indeed, the issue of structure appears to be a critical distinguishing feature. On the one hand, we have Representationalism, in which knowledge is synonymous with structure; on the other hand, both Suchman and Polanyi argue for an entirely unstructured, uninterpreted amalgam of experience as the source of intelligent behavior. Unfortunately, this leaves our analysis in a precarious position, with one foot on the concrete structure of symbolic representation, the other in thin air as we work (in vain) to reify the notion of "amorphous experiential knowledge". How can we compare if we can't articulate explicitly? The underlying assumption is, of course, that there *is* something to compare, that knowledge exists independently of a context of use. In other words, the assumption is that knowing can be separated from acting — a separation between the dynamic and the static. In this way, we have arrived at the true partitioning feature of epistemology: static versus dynamic semantics.

Under the philosophy of Representationalism, all knowledge is in a symbolic form: internalized syntactic representations of the world that serve as the basis for all reasoning and action. Since no other forms of knowing exists, the meaning of such structures must be inherent in the symbols themselves which, in turn, leads to the inevitable conclusion that a given symbolic structure has a fixed, finite meaning. In other words, the representational significance of a symbolic structure (in terms of the world it represents) must be unambiguous and deterministic in order for it to be of any use at all [4].

---

[3] This highlights a fascinating paradox in our exploration of situated frameworks. Heidegger and many others (including Suchman) claim that, representationally speaking, objects *do not exist* as mental entities until we use words to communicate about them. Conversely, language must inevitably objectify the experience or concept being communicated. Thus the paradox: *communicating* about situated frameworks is essentially an objectifying (i.e. representational) activity. This accounts for both the attractiveness of the Representationalist view and the difficulty of considering other alternatives.

[4] There is no value of a representation whose "truth value" is uncertain. That is, there's no point in representing anything if the meaning of the representation cannot be established. This observation justifies the obsession with formalized semantics in the knowledge representation community.

In Suchman's view, the nature of meaning is inherently dynamic — there is no such thing as knowing outside of acting. This implies that the meaning of action is constructed as action occurs and, further, that the significance of all past experience is fluid, being reconstructed continuously as events unfold. The situation is analogous to life in a kaleidoscope: no pattern is ever identical to a previous one though, at the same time, each pattern is dependent on the ones that came before. Similarly, Suchman's model suggests that past experiences influence our current perceptions in myriad ways, but that the meaning of each new action is uniquely and dynamically constructed with respect to the context in which it occurs. An obvious consequence of this perspective is that each person, having a unique history of personal experience, will inevitably find a different meaning in a given scenario. Furthermore, a given context may hold essentially infinite meanings for a single person, depending on which features of the context are emphasized as relevant. This motivates Suchman's conception of language, characterized as a mechanism for objectifying the world and constraining meaning in the process of constructing shared meaning.

In light of this discussion, Polanyi's model can be seen to be very similar to the one proposed by Suchman. Polanyi clearly indicates that all knowing begins with and remains rooted in a person's base of tacit knowledge. But how are we to explain the symbolic representations apparent in his "articulate" layer of knowledge? Are we to understand from this that Polanyi's model is somehow an extension of Representationalism, perhaps even a way of conceptualizing the problematic "common sense" knowledge? I feel that, in fact, Polanyi's articulate level *supports* Suchman's view of cognition and, in particular, the role of language in her model. To see this, recall the role of symbolic knowledge under Representationalism: it is used as the *generative* source for all reasoning and action. In contrast, Polanyi feels that the role of the articulate knowledge in his model is primarily *descriptive*, allowing the association of labels (i.e. symbols) with certain experiences in order that we may verbally and mentally make sense out of action. This notion is clearly illustrated in the following quote, in which Polanyi is comparing the development of humans and apes.

> "At the age of 15 to 18 months, the mental development of the chimpanzee is nearing completion; that of the child is only about to start. By responding to people who talk to it, the child soon begins to understand speech, and to speak by itself. By this one single trick in which it surpasses the animal, the child acquires the capacity for sustained thought and enters on the whole cultural heritage of its ancestors." (Polanyi, 1958, p.69)

The event which occurs, according to Polanyi, at 18 months in a child's life is precisely the formation of the second tier, of articulate knowledge — knowledge about how to objectify raw experience, a descriptive act inherently associated with communication. In this sense, Polanyi's articulate layer does not contain symbolic knowledge, but rather knowledge *about* symbols; in particular, how to use symbols to articulate experience. What Polanyi has done is simply to distinguish between language-using experience and all other experience. For the purposes of this paper, then, Polanyi's and Suchman's views are essentially identical: both models imply that meaning arises only in context, which includes the unique experiential history of the observer, and that language is the primary means by which we objectify, rationalize and communicate about action. Henceforth, I refer to them collectively as the *Situated* view.

In sum, the Representational view and the Situated view are fundamentally different with respect to the nature of knowledge an the origins of meaning. The key points of the two perspectives are contrasted as follows:

Representationalism

- Meaning is finite and quantifiable.
- Establishing "shared meaning" is a matter of arranging for both parties to have the same symbolic knowledge structure.
- A given representation has a fixed meaning or set of meanings. There is no need to negotiate over the meaning of a symbolic structure; the structure *is* the meaning. Also, the existence of multiple meanings is ruled out.

Situated Action

- Meaning does not exist independent of context.
- Meaning is locally and uniquely constructed based on past experience and the particulars of the context. A given situation has potentially infinite meanings.
- Language is a tool for emphasizing relevant situational features in an effort to establish shared meaning.

1.5 Discussion: Learning and Teaching

In the previous section, I worked to present and analyze the philosophical foundations underlying three prominent disciplines concerned with knowing and learning, narrowing them down to two distinct epistemologies, the Representationalist view and the Situated view. As discussed in the last section, the two models differ profoundly with respect to the nature of human knowledge. In this section, I extend the analysis, moving to an analysis of the dynamic aspects of the two models. We have seen, so to speak, *what* *knowledge is* (and how meaning arises) according to the two models, but how does knowledge grow and change under the two perspectives? My goal is to show how the static aspects of the epistemology shape the dynamic aspects (i.e. learning) of the model.

1.5.1 Elaboration and Deduction

The ways in which knowledge grows are straightforward under the Representationalist model, due to the rigid, unambiguous way in which meaning is characterized. Since knowledge under this model symbolic, the primary way of acquiring new knowledge is through the transfer of new symbolic structures, as illustrated in figure 1.4.

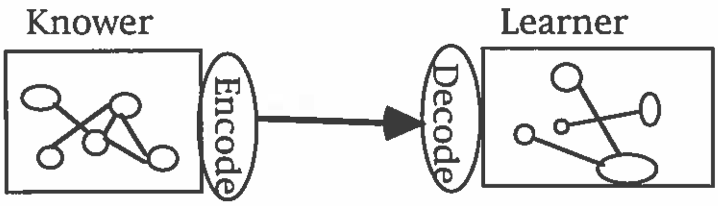

Figure 1.4: Learning in a Representationalist world

In this model, the symbolic knowledge of the knower is (optionally) encoded into a form suited to the particular means of communication, and is then transferred to the learner, who decodes the information to arrive at the same symbolic knowledge. Assuming that there is no error in transfer (see section 2.1), there is no room for confusion under this view, since the symbolic knowledge is assumed to have deterministic and unambiguous meaning.

It is interesting to consider how learning by experience is accounted for in this model. The cornerstone of Representationalism is that anything that can be known **must** be symbolic in nature. Thus, under this model, all perceptions of experience must be recorded as symbolic structures. However, since symbolic structures embody meaning unambiguously, the significance of an experience (in the very act of being symbolically represented) is established for once and for all. To constrain things even further, the attribution of significance, the building of symbolic structure from experience, is itself a symbolic activity, driven by the symbolic structures that already exist. In this way, the significance of experience under the Representationalist view is not only rigidly and unambiguously determined at the time the experience occurs (and is symbolized), the range of *possible* meanings to choose from is finite and predetermined by the symbolic structure of the experience interpretation engine.

A distinction is often drawn by educators between *active* and *passive* learning, reflecting the belief that not all learning involves transfer of new knowledge from a knower to a learner. In passive learning, the learner undergoes an introspective process in which new knowledge is derived from existing knowledge. Under the Representationalist model, this process is taken to be pure symbolic manipulation, in the form of logical deduction and symbolic abstraction (chunking). Existing symbolic structures are analyzed using tightly defined rules to derive new structures (and thereby meaning).

In sum, learning under the Representationalist scheme amounts to symbol manipulation, either in the form of external acquisition or internal derivation. Since meaning resides in the symbolic structures, it arises naturally as such structures are (re)organized.

1.5.2 The Situated view

The Situated view suggests that the source of all knowing lies in the dynamic interaction of some unarticulated, amorphous record of past experience with the particulars of the current context. Thus, active learning in the Situated View simply equates to being exposed to further experience. Words that are spoken and symbols that are presented certainly form a part of the experiential context, but have no inherent power beyond that of the other infinite details of the learning environment. Thus, the process of active learning does not involve any notion of knowledge transfer, focusing instead on the situated construction of meaning in a continuously evolving context.

But how, then, can the Situated model account for passive learning, which seems to occur even in the absence of new experience. Polanyi appears to provide a partial answer in positing a cyclic model of knowledge development in his model, as illustrated in figure 1.5.

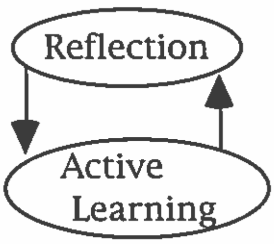

Figure 1.5: Polanyi's cyclic model of development.

In this model, the learner alternates between active learning, in which novel circumstances are experienced, and cautious reflection, during which time the learner works to makes sense of what has been observed. Unfortunately, Polanyi focuses his attention mostly on how this cycle serves to elaborate the "articulate" level of his knowledge model. As discussed earlier, this amounts to developing linguistic tools for describing action, rather than making sense out of action itself. There is no account of how passive reflection expands the tacit experiential knowledge which forms the core of the Situated model. Suchman, too, does not explicitly address the issue of passive learning in her work. However, I feel she would argue that the notions of action and context extend beyond the physical, applying equally well to mental behavior. In this sense, we can "simulate" experience mentally (based, of course, on actual experience) to find new meaning in action that has been experienced or might have been experienced. Indeed, this perspective is supported by Egan (1989) in his analysis of imagination in learning.

As might be expected, the Representational and Situated models vary widely is their conceptions of the learning process. The former view, following its commitment to a rigid, unambiguous notion of meaning, centers on the transfer and manipulation of symbolic knowledge. In the latter view, meaning can not exist outside of context, leaving nothing to "transfer" to the learner. Learning is therefore inseparable from experience. By characterizing the Representationalist and Situated view of the learning process abstractly, I provide a strong framework for investigating the philosophical foundations of both traditional and progressive pedagogical approaches being taken towards education today.

2.0 Education: a look at human learning outside of ITS

With a firm grasp of epistemological foundations, we are ready to move from philosophy to practice, to an examination of how learning has been characterized by cultural anthropologists and those concerned with the practices of effective teaching. One goal is to expose the philosophical foundations of this work; another is to identify features of successful education to inform the design of tutoring systems in the future.

2.1 Traditional approaches in education

It is instructive to consider how confusion, or lack of understanding in general, can even exist under the Representationalist model. After all, if symbolic structures are the unambiguous embodiment of meaning, then all that is required in learning is that pupils store the appropriate symbolic structure. Under the Representationalist model presented in section 1.1, the only way that lectures can fail to transfer knowledge is, *by definition*, through an error in the transfer of those symbolic forms, as illustrated in figure 2.1.

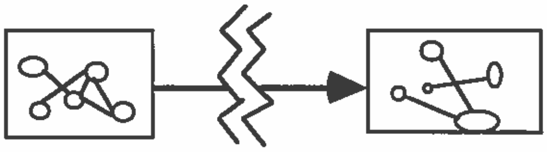

Figure 2.1: Failure of learning defined as interface failure.

It is this philosophy, I believe, that has led to the recent focus on interfaces, with little consideration of underlying philosophical assumptions. For example, the Cardiolab (Douglas and Liu, 1989) had a beautifully crafted direct manipulation interface, but nonetheless presented its explanations as networks of symbols, assuming (erroneously) the meaning of such symbolic structures to be apparent. Similarly, later problems revealed in our work with CVCK are rooted in the inherent ambiguity in constructing the meaning of symbolic information[5].

Outside of machine tutoring, traditional human-human tutoring also can be seen to rely on the same Representationalist foundations. Lecture hall formats in which the teacher presents and the student "absorbs" are clearly an instance of the sort of communication model depicted in figure 1.4. The assumption that meaning resides in symbolic descriptions of the domain, without any connection to the learner's world, is evident in the one-way, non-interactive nature of traditional lecture. A similar argument can be made for the comprehension of knowledge in textbooks. This source of knowledge permits even less interaction (none, to be precise) with the knower than does the lecture format. Thus, the construction of meaning is left entirely to the student. How, then, have countless generations of students learned anything at all? There are several answers to this question. Learners with sufficient background in the domain are able to ground the symbolic information in real experience and thereby construct its meaning. Most often, however, it is through laboratory sessions (experience) and through communication with fellow students that the meaning of the lecture is constructed.

In sum, traditional approaches to learning are heavily biased towards a Representationalist model of knowing. Learning is presumed to be reducible to storing appropriate symbolic forms. Misunderstanding is defined to be a matter of content — of

[5] Note that in the case of the CVCK system, however, two subjects collaborated to find the meaning in the application. In this sense, CVCK can be classified with the "new approaches" discussed in the next section. More significantly, this perspective implies that the use of two cooperating subjects is the main reason anything was learned, instead of just a technique for exposing the hidden mental processes of the individuals, as we thought at the time.

having internalized certain symbolic forms incorrectly or incompletely. This assumption is vividly exposed in tutoring systems like BUGGY (Burton, 1982) and WEST (Burton and Brown, 1979) which explicitly model incorrect and incomplete formation of the rules of the domain.

2.2 New directions in learning

From the Situated perspective, traditional approaches fail (by definition) to account for the highly individualized way in which meaning is constructed from experience, leaving it entirely to the learner to find meaning in the abstract symbolic information that has been presented. In some sense, this amounts to making the process of finding the meaning in symbolic forms "homework", something that is not addressed in class. In this section, I have a brief look at several approaches to education, some old, some very new, which implicitly acknowledge the highly personal and situated nature of learning by relying critically on interaction between the knower and the learner within a specific problem-solving context. In this way, the construction of meaning becomes a collaborative effort, taking place continually as each new impasse is reached and resolved in the course of completing the task.

2.2.1 Apprenticeship learning

The concept of apprenticeship learning is as old as human culture itself. Whether the practice is sanctioned as a formal arrangement within a culture or simply takes place without any explicit social acknowledgment, apprenticeship is defined by an exchange of labor for the opportunity to observe and learn from a master. Recently, there has been renewed interest in apprenticeship learning by thinkers like Lave and Wenger (1991), who feel that this age-old approach must be re-examined with respect to the Situated model of cognition. Not only does the apprentice have interactive access to the knower and the context of the activity, the entire learning activity is actually *embedded* in the context. That is, task learning in apprenticeship is driven by immersion in the environment of task application, coupled with increasing participation of the learner in that environment. Lave and Wenger have coined the term *legitimate peripheral participation* to describe this learning arrangement. In the following sections, I analyze Lave and Wenger's studies in an effort to reveal how various features of an apprenticeship scenario impinge on learning efficacy. The goal is not to compile a prescriptive list for learning success, but rather to develop a framework for characterizing various approaches in apprenticeship learning.

Informal Apprenticeship

In some cases, apprenticeship may never be formalized at all. That is, learning may take place over a period of time in the absence of any social contract between the knower and the learner. Lave and Wenger characterize the apprenticeship of Yucatec midwives as an instance of this form of learning. In this case, the learner, who is almost always a daughter or granddaughter of the master midwife, simply acts as a helper to the master, accompanying her on her rounds. In particular, since the knower-learner relationship is not ever formalized, active instruction (e.g. explanation of some procedure of rational for giving a certain herbal medicine) occurs only rarely, if at all. This may account, in part, for the observation that the period of apprenticeship under this arrangement is lengthy, often stretching over several decades. Furthermore, there is no formal "graduation" from apprenticeship to master — the apprentice simply takes on more and more responsibilities over the years until, in the end, the master simply plays the role of an observer (thereby sanctioning the actions of the apprentice) while the apprentice does all of the real work. In sum, the status of the master-apprentice relationship or, more accurately, the level of curricular *organization* of the apprenticeship, is one dimension of difference in apprenticeship learning.

Organization of Learning

Unlike the example of the Yucatec midwives, most apprenticeships have a more formal status, in which there is some sort of social (or even legal) contract between the master and the learner. In such instances, it is incumbent upon the master to organize the learning experience in some way, ordering the skills that the apprentice is expected to learn; a "curriculum" of sorts must be established. How this is done constitutes another dimension of difference in apprenticeship learning.

In traditional educational theory, certain tacit notions prevail regarding the organization of curriculums including "from simple to complex", "from known to unknown", and "from concrete to abstract" (Egan, 1989). The common thread among these guidelines is their absolute independence from the context of use — they are all organized in a *logical top-down* fashion. In other words, they emphasize task-analytic form, over real world function. In their study of the apprenticeship of Vai and Gola tailors and the apprenticeship of Navy midshipmen, Lave and Wenger find very different curricular organizations.

In West Africa practically all garments, from underwear to formal outfits, are hand-crafted and the tailor's craft is highly respected. The order in which new tailor's apprentices learn skills is interesting. A task-analytic approach would dictate that curricular organization should follow the logical form of the garment: first teach the apprentice how to stitch (on some piece of scrap), then show him how to use scissors and cut pieces (on more scraps), and finally combine these tasks to construct the finished garment. In fact, apprenticed tailors are first put to work putting the finishing stitches on complete pieces, starting with underwear and moving on to hems and liners for formal garments. Next they are taught to sew together undergarments and then formal garments. Only near the end of the apprenticeship are they taught to cut out garment components at all.

In the Navy, midshipmen are responsible for the navigation of the vessel, which is especially critical when navigating into a port. This process involves six midshipmen: four taking sightings and two combining the results on a navigational chart, which involves much trigonometric calculation. Again, a task-analytic curriculum would dictate that learners first be schooled in abstract trigonometry, learning to combine measurements, then be schooled in taking sightings and soundings. In fact, novice midshipmen are first paired with more advanced apprentices at the sighting and sounding stations. Only after rotating to each of the four stations are they brought inside and, again in collaboration with more senior officers, allowed to combine the measurements mathematically on the charts.

Clearly, the curriculums for tailors and sailors are very different. However, they both share a crucial focus on the *context of the task* that separates these curricular approaches from those used in traditional education, which focus on the logical structure of the task. In the case of the tailors, the central aspect of the task is the customer. Novices are first given tasks that don't matter in the eyes of the customer — hidden seams and underwear. As they progress, they are given more "critical" work, things that affect the visible quality of the product. In the Navy, apprentices follow the flow of information as it is produced and consumed in the navigation of the ship starting at the sighting/sounding stations and moving to the charting board. Thus, though the two curricular organizations are different, they share a common focus on the real world aspects of the task.

Failed apprenticeship

What is it about apprenticeship that makes it such an effective approach to learning? Is just the mere act of apprenticeship enough? One way to gain insight into this question is to analyze a context in which apprenticeship has *failed*. Lave and Wenger identify the apprenticeship of American meat cutters as one such instance. Novice butchers first undergo some classroom training to familiarize them with the basic concepts and tools and are then sent, as apprentices, to actual job sites. Lave and Wenger suggest several features of this apprenticeship situation which might be responsible for the observed failure of novices to grow into masters of the trade.

1) The workplace is organized in such a way that work stations, the places at which individual jobs are performed, are separated by walls and doors. Thus, the apprentice tends to work in isolation, with no access to other aspects of the overall task. This leaves the apprentice totally unable to learn through observation as well as obscuring the ways in which other aspects of the task depend on the apprentice's work.
2) The master butchers are not available as a learning resource. Not only are they not in the same space (see above) but rigid social barriers separate them as well. In particular, the masters tend to view the apprentices as (non-learning) laborers, and treat them accordingly.
3) Economic factors emphasize efficiency and division of labor. An apprentice learning a new job is simply not as efficient as a master. Thus, an apprentice may be left doing the same (introductory) job indefinitely. This would be similar to a Gola tailor being assigned to sew underwear for years on end, even after learning the job perfectly[6]. In this way, the learner is never exposed to all tasks in the profession.

In light of the analysis above, in is clear that several central features of apprenticeship learning are responsible for the efficacy of this approach:

- Apprentices do **useful** work at all times. Every task set for the apprentice contributes to the overall effort of the organization. The overall structure of the curriculum is rationalized by the real world context of the work, not by the top-down decomposition of the task.
- Learners have intimate access to knowers and to the overall task. Not only can learners ask questions of the masters but, more importantly, they have continual access to

[6] I feel this point reflects a central underlying assumption of the apprenticeship paradigm: integrity on the part of the master. The whole notion of apprenticeship is based on the master receiving labor for free (or at reduced cost) in exchange for learning. Only the integrity of the master prevents him from abusing the relationship by keeping the apprentice in the same task after that task has been mastered.

other aspects of the task. Even though they may not be able to accomplish these aspects, the relevance of their contribution is evident at all times. In other words, co-presence anchors the subtask in the context of the overall task and thereby *gives meaning* to the chore.

Lave and Wenger's work on apprenticeship contributes to our understanding of experiential learning in several important ways. In the past decade, "learning by experience" has become a popular catch phrase, almost a mantra, to those developing new approaches to education, and especially to designers of tutoring systems based on direct manipulation of microworlds. Clearly, Lave and Wenger's findings validate this approach in many ways, emphasizing the value of exploratory experimentation in the learning process. At the same time, however, Lave and Wenger's work exposes several shortcomings of the "learning by experience" approach:

The role of curriculum.

It is not enough to simply provide an exploratory environment. In most cases, apprenticeship in Lave and Wenger's study was defined by an ordered set of tasks. In fact, the ordering of these tasks, the structure of the curriculum, appears to be both crucial to the success of the apprenticeship and very different from traditional approaches to curricular planning. The cases of the midwives and the butchers highlight two consequences of curricular failure: in the former, the ambiguity of the curriculum (if the term even makes sense in the absence of a formal apprenticeship) led to a very flat learning curve; in the latter, the failure to execute the given curriculum led to curtailment of learning and stagnation.

The interactive nature of learning.

A central aspect of the apprenticeship model of learning is the interaction between knowers and learners to accomplish a common task. The master is continually present, leading not only by explanation, but by example as well. Putting this in terms of the earlier sections of this paper, by face-to-face explanation in a real problem-solving context, the knower actively collaborates with the learner to construct the meaning of action as it occurs.

2.2.2 Collaborative Learning

A slow metamorphosis is taking place in education, gaining momentum only in the last five or ten years. A leading proponent of these new techniques, loosely grouped under the heading of *collaborative learning* describes the changes as follows:

"During the last 90 years over 575 experimental studies comparing the effectiveness of cooperative, competitive, and individual approaches to learning have been published, as well as over 100 correlational studies. Many of these studies focus on college teaching. More is known about the efficacy of cooperative learning than about lecturing or almost any other aspect of education. However, the paradigm of teaching has been slow to shift. Yet, shifting it is." (Jones, 1992)

The "competitive" and "individual" approaches that Jones mentions here refer to the traditional practices in education. By adopting pedagogical techniques (e.g. lecture) and evaluation methods that focus on the individual student, an atmosphere of isolation and competitiveness is promoted. An obvious example of this phenomenon is the practice of "grading on the curve", a practice Jones claims imposes artificial distinctions in performance and discourages cooperation. In contrast, he offers[7] the concept of collaborative learning, in which students cooperate to achieve common goals. Figure 2.2 provides a brief summary of the major distinctions between these two approaches (Egan, 1989; Hamm and Adams, 1992).

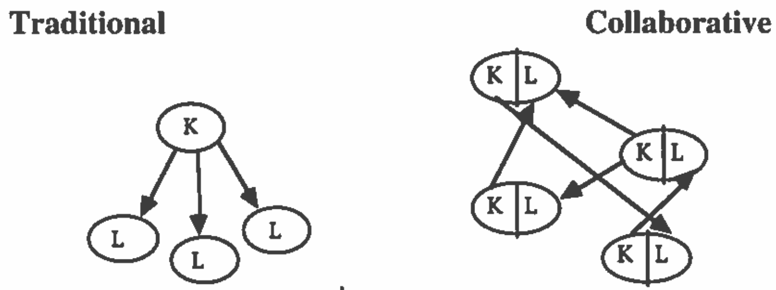

Traditional	Collaborative
Teacher centered	Task-centered
Passive listening	Cooperative problem solving
Ability grouping	Mixed ability groups
Individual success	Group processes
	Projects

Figure 2.2: Traditional vs. Cooperative approaches

As described earlier, the traditional approach focuses on the teacher as the source of all knowledge. In contrast, the collaborative approach views students as multi-faceted entities that act as knowers and learners simultaneously. Each student is empowered to share his or her unique perspective on the subject matter at hand, as the group collaborates to make sense out of the new material. Clearly, there is something missing in the second figure: the source of the new material. Reinforcing a point made in the last

[7] In fact, Slavin (1983) is generally given credit for founding the area of collaborative learning.

section, the presence of a knower to "prime" the discussion and organize the curriculum is universally recognized as necessary by all advocates of the collaborative model[8]. In a sense, then, the collaborative model is more of an extension to the lecture-based paradigm to "expand" the boundaries of the classroom beyond simply the presentation of the material to move the negotiation of its meaning under the auspices of the teacher as well.

While this "expansion of scope" may seem innocuous, I feel that it represents a profound shift in our conception of human learning. Indeed, it is nothing less than a re-alignment of the entire philosophical foundations for education, moving from a Representationalist model of meaning to a Situated one. Collaborative approaches acknowledge the ambiguity and essentially meaningless nature of symbolic information, and implicitly recognize the highly personal nature of meaning, providing an active cooperative context in which students share experiences as they work to negotiate the meaning of a given presentation. According to Hamm and Adams:

> "Small group discussion ... helps ideas take shape and become more alive and personal. Expressing ideas helps the group examine, compare, and affix personal meaning to the concepts or beliefs presented. Placing ideas into a personal and collective experience is a powerful motivator, and information tends to be remembered longer because more meaning is attached to it." (Hamm and Adams, 1992, p.3)

Clearly, collaborative learning is firmly centered around concept formation through interpersonal communication; through the use of language, we learn how to objectify experience, determine relevant features in perceived action, and develop the means to rationalize action. Stated in these terms, the Situated epistemological foundations of collaborative learning are apparent.

2.2.3 Sources of experience

Under the collaborative model, meaning is constructed from the shared experiences of group members. But what if students don't *have* any past experiences in some domain. Mathematics, for example, is an abstract domain in which humans have no innate experience. Electrical engineering, molecular chemistry and higher physics are other good examples. All of these domains are essentially imaginary, symbolic frameworks constructed to objectify and rationalize a seamless flowing reality. Under the Situated model, the only way we can possibly find meaning in these invisible symbolic domains is to somehow connect them with real experiences. There are two ways I think this can be achieved.

---

[8] Hodgkins (1985) and Polanyi (1958) refer to this guidance as "the role of tradition" in their work. This is not to be conflated with the traditional epistemology (i.e. Representationalism).

Direct Modeling of abstract domains

A new approach to teaching children mathematics, known as *Thinking Math* (Gill and Billups, 1992), illustrates one way in which abstract domains can be linked to real experience. Thinking math is clearly a collaborative technique, centered around class discussions and group problem solving. Yet, another crucial aspect of this approach lies not in the collaboration between class members, but in the direct link between abstractions and real objects. It has often been argued that, in order to acquire meaning for learners, abstract domains like mathematics must be applied to a concrete domain like basic motion physics. Who has not seen "story problems" asking the things like "imagine that you are superman and want to jump over a tree; you know you are 50 feet from the base of the tree, the angle to the top is...."? While this story form can certainly be useful (see below), it relies on *imaginary* circumstances, not real ones. Thinking Math is centered around the notion that real circumstances are the most robust source of meaning. For instance, in learning about multiplication or division, students are asked to plan a class picnic, calculate the number of people who will be present, determine how many six-pacs of soda and packages of hot dogs will be required to feed everyone, and how many buses will be needed for transportation. This approach is similar to the use of manipulatives like apples and sticks to reify the abstract quantifiers in basic math. Our previous discussion gives us a framework for rationalizing the efficacy of both: If collaborative learning relies on constructing meaning based on past experience, basing instruction on real circumstances experienced by all class members must dramatically increase the pool of shared perception.

Storytelling and imagination

As desirable as it may be, there will come a point in the curriculum where it is not possible to construct real circumstances to reify the abstract problem. In his work, Egan (1989) argues that imaginary stories (i.e. stories about events that never happened) play a pivotal role in learning and understanding not accounted for in the traditional paradigm. According to Egan, traditional curriculums are founded on a number of *ad hoc* principles, including, for instance, that learning proceeds from concrete to abstract, from known to unknown, and from the simple to the abstract. This leaves no room at all for imagination. Egan writes:

> "Imagination seems to have had no influence at all on the *ad hoc* principles. The sense of the child as an energetic creator of mental images of what may never have been experienced seems, on the face of it at least, to conflict with the sense of the child presented in the *ad hoc* principles that have been so influential in education. Is the child

who manipulates concrete materials derived from everyday experience the same child whose mind is brimming with star-warriors, monsters, and wicked witches? And is it more or less important for future educational development that a child be able to create and mentally manipulate these imaginary creatures than that the child be able to conserve liquid volume?" (Egan, 1987, p.7)

Egan's words emphasize an interesting feature of imaginary stories: not only can they describe circumstances that never happened, they may also describe ones that never *can* happen (e.g. monsters). Furthermore, imaginary stories are apparently used by adults as well as children. A good example is Einstein's use of *Gedankenexperimente* to consider scenarios like faster than light travel, which are both imaginary and are fundamentally impossible.

In this way, imaginary stories are used to essentially *broaden* the base of experience used by humans as a resource in constructing the meaning of action. Note, however, that imagination itself is firmly rooted in real experience — Superman behaves primarily like a man and a person traveling faster than the speed of light still perceives things in a human way.

The goal of this section has been to establish that new pedagogical approaches to learning, collectively labeled collaborative learning, represent the adoption of the more flexible Situated perspective over the traditional Representationalist view. Collaborative learning clearly emphasizes the cooperative, situated construction of meaning from symbolic knowledge. Following Suchman and Polanyi, proponents of the collaborative paradigm identify past experience and local context as the critical resources for this task. A corollary effort of the section has been to explore the sources of such personal experience, arguing that real shared experience creates a particularly rich context and that imagination can be viewed as a powerful mechanism for further enriching a person's pool of experience.

3.0 Thoughts on the future: the role of Intelligent Tutoring

In the previous section, I showed how research in Ethnomethodology and new (and rather successful) approaches in education can be seen to rest on the flexible Situated model of meaning. In particular, the Situated model emphasizes that all past experience, both as an individual and as a human being in general, serves as a critical resource in establishing the significance of action. This highly localized, inherently interactive characterization of meaning makes it difficult to see how machine-based tutors can ever accomplish individualized teaching. By rejecting the philosophy of Representationalism, we appear to have slammed the door on the entire field of Intelligent Tutoring Systems.

In this closing section, I define a new role for ITS, arguing that the future of tutoring systems lies not in trying to become a useful *participant* in a human conversation, but in becoming a useful *subject of* human communication. In other words, we must change our perception of ITS to one in which systems support human activity rather than usurping it. I adopt this new perspective to speculate on interesting new directions for ITS, as well as justifying some old ones:

1) **Simulation over explanation.**
   It appears that humans don't operate on the symbolic Representationalist model and that, for the foreseeable future at least, conversations between humans and machines will remain rigid, crippled by the semantic tunnel vision inherent in symbolic representations. Thus, machines that provide active situated explanation are unlikely. At the same time, the central role of experience in learning has been emphasized over and over. Furthermore, it seems (Egan, 1989) that imaginary experiences constitute an effective and powerful way to elaborate direct experience and extend the learning process. These observations support ongoing work in simulation-based ITS, in which the system becomes more of an exploratory tool than an active tutor. Such systems may still take a *passive* remedial role (Doerry, 1992) by altering the environment (including the tasks that are presented) in various ways to direct the student's exploration.
2) **Collaboration-based.**
   In our work with CVCK, a large amount of time was spent working to create an "intuitive" interface, one in which students would have little trouble operating the application, and even less trouble interpreting the simulation that was presented. This was never accomplished. In light of the points made in this paper, the reasons for this are clear: the central claim of the Situated model is that semantically unambiguous interfaces do not exist. The meaning of all experience (including that in simulated worlds) is constructed locally and uniquely by each user. All evidence from collaborative learning research indicates that interaction with other learners is a strong resource for finding meaning in action. I feel that this is a clear mandate for tutoring systems that support collaboration by encouraging several students to use the machine at once, cooperating to solve tasks within the simulated environment.
3) **The role of teachers.**
   All work in cooperative learning emphasizes the important role of the human knower,

both in helping to resolve local impasses, and in structuring the overall exploratory process. Lave and Wenger's (1991) experience with Yucatec apprenticeship illustrates how the lack of curriculum can severely retard learning efficacy. Most simulation-based tutoring systems tacitly acknowledge this observation by providing students with a pre-defined curriculum, either encoded in the simulation itself, or in the form of a "lab manual". The weakness in these approaches lies in their rigidity — later lessons generally depend on learning (which can be very different from mere successful completion) in earlier exercises. Since the machine (and certainly a lab book) has no access to the meaning negotiated by the learners, there is only the weak evidence of syntactic solution correctness to guide the pace and subject matter of the exploration. Thus, creating tutoring systems designed to support the monitoring and curricular organization by a human knower will be important.

4) Network ITS

Finally, the emphasis on collaboration between human learners to empower the learning process leads to an immediate practical problem for ITS: as computers, most systems rely on the display of information on a screen. While two, or perhaps three, humans might be able to cluster around a single screen, accommodating more is problematic. More importantly, a single machine typically has only a single input device, meaning that only one person in the learning group has direct control over the interaction and the direction of the exploratory experience. One common solution has been to create a "lab" in which students work on individual machines. I feel this destroys much of the collaborative context. What is needed is a way for multiple students to work on individual machines in a *shared* exploratory context. For instance, students might have identical dynamic copies of the environment on individual machines , or perhaps all refer to a single copy projected on a central screen. Either way, each group member is equally empowered to act to explore the simulation.

In sum, the suggestions above all center bringing the collaborative learning model recently promoted in classroom education to bear on the design of intelligent tutoring systems. Under the new model, machines are viewed as dynamic (but mute) learning environments, designed to provide an active (though simulated) experiential context for collaborating groups of students under the guidance and a human knower.

4.0 Conclusions

The nature of knowledge — and thereby how we acquire it — has been called into question. In the traditional view, knowledge of the world is represented symbolically in our minds and serves as the basis for all behavior, including internal reasoning, action and communication. However, in order for such a model to work, there must be an assumption that the semantics of the representation are finite and deterministic, arising from the symbolic structures themselves. In the alternative view, knowledge does not really exist as an independent entity. That is, there is no meaning outside of a particular situation of use — what we bring to bear on a situation cannot be determined beforehand and can only be partially (and abstractly) described post hoc, in the form of a description or justification for action.

The consequences for learning and teaching, both for human teachers and the field of ITS are profound. In human education the shift began 10 years ago, away from lecture-oriented approaches towards more collaborative interactive approaches. The matter is more difficult for ITS designers. The computer is inherently a symbolic device and, therefore, is fundamentally constrained to operate under a Representationalist system. It is not the purpose of this paper to argue that it is impossible for machines to communicate anything to humans. Clearly this is not the case. Nor is it to claim that artificial intelligence is a doomed science. Again, there are clearly areas in which "intelligent" machines have led to great advancements. But I do claim that

a) human conversation is not one of these areas and

b) that tutoring and learning is intimately dependent on situated conversation.

Early in the paper, I worked to establish that the perceived nature of meaning — where it resides, and how it is constructed — is the crux of the entire debate, defining an insurmountable difference between the philosophy of Representationalism and the Situated view. As a result, I concluded that machines can never communicate in the human sense since, essentially, there is no meaning inherent in *any* symbolic representation; that meaning is constructed uniquely and collaboratively in each new situation with reference to the particulars of an evolving context. Machines have no access to this context, especially the experiences of humanhood that we all share, and thus cannot participate effectively.

In conclusion, I want to extend this line of reasoning somewhat to suggest that Intelligent Tutoring Systems have actually degraded the learning environment, rather than improving it. To see this, consider two central claims (roughly stated) of this paper:

Part I: Representationalism is a corrupt foundation for human intelligence. All responsibility for discovering the meaning in symbolic forms is dumped in the lap of the human in human-machine interaction. Learners are the least likely to manage this interpretive task. The situation is analogous to giving an illiterate person written instructions[9].

Part II: Learning is intimately rooted in interaction, both with other learners and with knowers, as learners work to find the significance in action.

Now consider the professed goal of intelligent tutoring systems, the Holy Grail for an entire generation of researchers. It has been stated many times that the promise of ITS lies in providing each student with a "personal tutor", perfectly adapted to that individual student's unique learning style. Seated in their cubicles with such machines, students would undergo efficient and personalized instruction and remediation. Most tutoring systems are designed to tutor a single student at a time, fully embracing this one-on-one metaphor.

In light of the conclusions reached in this paper, this is distinctly counter-productive: we have taken the student *away* from all interpersonal interaction — the truly strong resource for finding meaning in action — and isolated her with a machine that has no capacity for helping the student find the meaning in the machine's behavior.

Despite all of this, things are not as bad as they seem. Humans are incredibly adept at learning, at ascribing meaning to action — some systems have enjoyed moderate success. I argue however, that any learning that occurred did so *in spite of* the explanatory efforts of the system, and not because of them.

Bibliography

ANDERSON, R. C. & NAGY, W. E. (1992). The Vocabulary Conundrum. *American Educator*, 16, (No.4):p.14.

BURTON, R. R. & BROWN, J. S. (1979). An investigation of computer coaching for informal learning activities. *Int Jrnl of Man-Machine Studies*, 11, 5-24.

---

[9] It may in fact be *worse*. At least the illiterate person has no preconceptions regarding the written symbols. Our work in CVCK showed that this is emphatically not true for novices in cardiovascular physics.

BURTON, R. R. (1982). Diagnosing bugs in a simple procedural skill. In SLEEMAN, D. H. & BROWN, J. S., Eds, *Intelligent Tutoring Systems.* London: Academic Press.

DOERRY, E. (1994). Evaluating Foundations: A Critique of Traditional Conceptions of Knowing and Learning in Intelligent Tutoring, Dept. of Computer Science Tech. Report TR-94-24, Uni versity of Oregon, Eugene, OR.

DOUGLAS, S. & LIU, Z. Y. (1989). Generating Causal Explanation from a Cardio-Vascular Simulation. *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, MI. 1989.

EGAN, K. (1989). Teaching as Story Telling. Chicago, ILL.: University of Chicago Press.

GIBBS, J. B. (1987). Tribes: a process for social development and cooperative learning. Santa Rosa, Ca.: Center Source Publications.

GILL, A. J. & BILLUPS, L. H. (1992). The Power of Thinking Mathematics. *American Educator*, 16, (No.4):pp. 4.

GLEICK, J. (1987). CHAOS: Making a New Science. New York: Penguin Books.

HAMM, M. & ADAMS, D. (1992). The Collaborative Dimensions of Learning. Norwood, NJ.: Ablex Publishing Corp.

HERTZ-LAZAROWITZ, R. & MILLER, N. (1992). Interaction in cooperative groups: a theoretical anatomy of group learning. Cambridge: Cambridge University Press.

HODGKIN, R. A. (1985). Playing and Exploring: Education through the discovery of order. New York: Methuen & Co.

JONES, S. (1993). Workshop presentation. University of Oregon, Eugene, OR.

LAVE, J. (1988). Cognition in Practice. Cambridge, MA.: Cambridge University Press.

LAVE, J. & WENGER, E. (1991). Situated Learning: Legitimate Peripheral Participation. Cambridge, MA.: Cambridge University Press.

MCDONALD, H. (1986). The normative basis of culture. Baton Rouge, La.: Louisiana State Univ. Press.

PIRSIG, R. M. (1991). Lila, an inquiry into morals. New York: Bantam Books.

POLANYI, M. (1958). Personal Knowledge. London: Routledge and Kegan Paul Ltd.

ROGOFF, B. & LAVE, J. (1984). Everyday Cognition: It's development in social context. Cambridge, MA.: Harvard University Press.

SCHARFSTEIN, B. (1989). The dilemma of context. New York: New York Univ. Press.

SIMON, H. A. (1969). The Sciences of the Artificial. Cambridge, MA.: MIT Press.

SLAVIN, R. (1983). Cooperative Learning. New York: Longman Publishers.

SUCHMAN, L. A. (1987). Plans and Situated Action. Cambridge University Press: Cambridge, MA.

WENGER, E. (1987). Artificial Intelligence and Tutoring Systems. Los Altos, Ca.: Morgan Kaufman Publishers.