

JOURNAL OF

Educational Computing Research

Editor: Dr. Robert H. Seidman

Volume 16, Number 2—1997

Artificial Tutoring Systems: What Computers Can and Can't Know

Theodore W. Frick

Baywood Publishing Company, Inc.
26 Austin Avenue, Amityville, NY 11701

Call (516) 691-1270 Fax (516) 691-1770
Orderline — call toll-free (800) 638-7819
email: baywood@baywood.com web site: <http://baywood.com>

ARTIFICIAL TUTORING SYSTEMS: WHAT COMPUTERS CAN AND CAN'T KNOW

THEODORE W. FRICK

Indiana University Bloomington

ABSTRACT

After more than four decades, development of artificially intelligent tutoring systems has been constrained by two interrelated problems: knowledge representation and natural language understanding. G. S. Maccia's epistemology of intelligent natural systems implies that computer systems will need to develop qualitative intelligence before these problems can be solved. Recent research on how human nervous systems develop provides evidence for the significance of qualitative intelligence. Qualitative intelligence is required for understanding of culturally bound meanings of signs used in communication among intelligent natural systems. S. I. Greenspan provides neurological and clinical evidence that emotion and sensation are vital to the growth of mind—capabilities that computer systems do not currently possess. Therefore, we must view computers in education as *media* through which a multitude of teachers can convey their messages. This does not mean that the role of classroom teachers is diminished. Teachers and students can be empowered by these additional learning resources.

INTRODUCTION

The electronic information age is now here. Computers were invented about half a century ago. Until the late 1970s, computers were large, expensive, and typically found only in big corporations, universities, and governmental institutions. Invention of inexpensive microchips containing very large-scale integrated circuits has changed things dramatically. The personal computer, as it has come to be called, began to appear on our desk tops around 1980. Television, computers, and telecommunications are becoming more and more commonplace around the world [1].

The global electronic village that Marshall McLuhan envisioned in the 1960s is now here [2]. We take it for granted, for example, that we can watch on

TV the Olympics taking place half-way around the world. As another example, many of us were horrified by the drama of the unfolding War in the Persian Gulf in early 1991, brought to us live by *CNN World News* as Baghdad was being bombed.

For thousands of years humankind has built and improved routes and vehicles for transporting goods and people. We have built ships, highways, automobiles, trucks, railways, trains, airplanes, and even dirigibles to transport things and people from one place to another. As we enter the twenty-first century, the most important kind of new "highways" emerging are globally interconnected digital communications networks (e.g., the Internet) [3].

What is transported through these networks are bits of information. These bit collections can represent whatever we want them to, whether they stand for linguistic symbols such as letters of an alphabet, icons or ideographs, for moving pictures and sound, for computer programs, stock market prices, or bank account balances.

Computers, televisions, telephones, stereo systems, and radios all process encoded information. Whether information is digitally or analogically encoded is less important than the trend that these information technologies are merging. Will these globally interconnected, multimedia computer-television-telephone-stereo-radio systems largely replace teachers? Will these multimedia tutoring systems be intelligent enough to do so? Imagine, for example, intelligent tutoring systems on the Web available to anyone in the world.

The notion of intelligent tutoring systems is seductive. The professional development of conventional computer-mediated learning products is very expensive—typically 200-300 hours of development time for one hour of learning time [4]. If somehow we could put knowledge into computers which were smart enough to teach students, this would be a solution to the labor-intensive methods that we currently use.

INTELLIGENT TUTORING SYSTEMS

Intelligent tutoring systems are typically conceived as having a knowledge base, a set of pedagogical rules, a model of the student, and a natural language interface [5, 6]. The artificial intelligence (AI) community has run into serious difficulties in both knowledge representation and natural language understanding [7, 8]. While there were some notable early successes, such as Terry Winograd's blocks world [9], the knowledge representation and natural language problems have proven to be largely intractable [6, 10].

Nonetheless, two areas of AI have met with some success. *Expert systems* and *neural networks* have proven themselves to be quite useful for certain kinds of tasks [11-13].

Expert Systems

Expert systems consist of rules, usually in an "If . . . , then . . ." form. These rules are created by knowledge engineers, who try to capture the reasoning processes of experts who solve problems in some domain. These rules constitute a database for a computer program called an "inference engine." When a person uses such an expert system, he or she interactively answers questions that were previously programmed into the database by the knowledge engineers. The inference engine uses as further data the answers provided, in order to follow a particular path of reasoning (by doing forward or backward chaining through the rule set). Through such interaction with a user, the expert system reaches a conclusion or decision and displays the result.

While expert systems can reason deductively quite well within a narrow domain, they are helpless and useless outside that domain. Moreover, such systems cannot be said to understand the meaning of the terms used in the reasoning process [10]. Expert systems follow the rules for reasoning that are established by the knowledge engineers. If such rules are faulty, incorrect inferences will be drawn by the expert system.

Neural Networks

Neural networks are similar to expert systems in that a given set of inputs will produce the same output. The major difference is that neural networks create the "rules" themselves. This learning occurs through practice and feedback—a sort of Pavlovian conditioning. Dreyfus points out that we humans do not know what the networks are necessarily learning [8, 10]. He referred to one example, where such a network appeared to have learned to identify military tanks from aerial photos. When it was tested on a new set of stimuli, it made numerous mistakes. On hindsight, researchers discovered that during training, the photos which had tanks in them were taken on cloudy days, and ones without tanks were taken on sunny days. The network had learned to discriminate cloudy and sunny days, not whether tanks were present. In another example, a neural network apparently had learned to drive a van, where the primary input came from a video camera focused on the road ahead [10]. The net appeared to have learned to keep the van on its side of the road when driving at slow speeds. However, when it was tested on an interstate highway, it had a tendency to take exits to the right. In this case, the network had probably learned to follow the white line on the right side of the road.

At this time, it is not clear how far we can go with neural networks. The primary obstacle is that the number of trials necessary during training increases exponentially with the number of components in the network.

Neural networks and expert systems have had some success in voice and handwriting recognition. Again, however, such systems can only deal with

limited numbers of spoken words or hand strokes with moderate reliability. These systems do not understand the meanings of words but instead discriminate sound patterns or two-dimensional spatial configurations. Their discrimination is aided by data on typical grammatical patterns and normally occurs within a narrowly prescribed context—e.g., purchasing an airline ticket—in which a limited vocabulary is used with a finite set of choices and tasks to be performed.

Ongoing Projects

Two current projects are attempting to “educate” computers. Douglas Lenat and colleagues at the University of Texas have been attempting to put factual knowledge and rules of human conduct into computers [14]. This ambitious Cyc project has been under way for a number of years, and according to some, is the last great hope for the AI community [8, 10]. Time will tell if it succeeds, and the discussion below may shed some light on the prospects for its success. Rodney Brooks and colleagues at the Massachusetts Institute of Technology have been designing computers capable of learning from experience [15]. According to Greenspan, “both approaches have failed to reach the levels projected for them, and in creative reasoning they can be outdistanced by a young child” [16, p. 126].

THE CRUX OF THE PROBLEM

Feigenbaum and Dreyfus conclude that the problems of natural language understanding and knowledge representation are not likely to be solved until machines have perceptual-motor capabilities that would, for example, allow them to move physically across a room without running into things [10]. This kind of sensorial experience with the world around us appears to be necessary for grounding the *meaning* of natural language that is used in communication among humans. On hindsight, this should not be surprising. John Dewey, for one, poignantly discussed the vital relationship between experience and thinking as being central to the learning process [17]. Maria Montessori, for another, designed her system of education around this relationship [18].

Recent biological and neurological research on how human brains develop is shedding new light on just how vital sensory experience is [16]. Human genes and their environment interact in ways that resemble a “dance” between nature and nurture. Evidence is mounting that human sensory and emotional experiences literally build neural connections in the brain [19-21]. This research provides empirical evidence which supports the significance of qualitative intelligence in creating meaning of natural language.

In this article, I first provide an overview of George Maccia’s epistemology which distinguishes qualitative, quantitative, and performative intelligence. Since his epistemology is relatively unknown, details beyond this article are provided on the Web (<http://education.indiana.edu/ist/faculty/episcomp.html>). Next,

I discuss implications of this epistemology for artificially intelligent tutoring systems. Finally, I discuss research done by Greenspan and others which complements this epistemological view.

OVERVIEW OF MACCIA'S EPISTEMOLOGY

George Maccia—a scientist, mathematician, general systems theorist, and philosopher of education during his long career—was interested in epistemology and its relationship to educational objectives and curriculum. After several decades of research and writing, he published a seminal article in *Systems Research*, "Genetic Epistemology of Intelligent Systems" [22]. Just before retirement, he further elaborated his epistemology in a presentation to the Fourth International Conference on Systems Research, Cybernetics and Informatics in 1988 [23]. Maccia's research was motivated in part by concerns about inadequacies of the taxonomy of educational objectives developed by Benjamin Bloom et al., which has had considerable impact on conceptions of curriculum in U.S. public schools during the last forty years [24].

Figure 1 illustrates the basic kinds of cognition: "knowing that," "knowing that one," and "knowing how." Earlier epistemologists such as Ryle [25] and Scheffler [26] had made distinctions between "knowing that" and "knowing how." Maccia's contribution was the addition and distinction of "knowing that one" from "knowing that." Moreover, his contention was that intelligence is distributed among systems of various kinds—not just human beings—but other forms of living and non-living systems. Such non-living systems could include computers.

Quantitative Intelligence: "Knowing That"

"Knowing that" includes understanding of concepts. To identify some four-legged, hairy creatures as dogs is an example of "knowing that" or what Maccia called "quantitative intelligence" (see Figure 1). Quantitative intelligence is "usually associated with mental acts that employ abstractive inference (i.e., modes of generalization or instantiation)" [22, p. 215]. Steiner adds, "Theoretical structures allow one to shape and group instances; they are universals and so are generals that are independent of time and place. Although 'quantitative' in a common sense pertains to numbers, in its technical sense it involves extension. Generals independent of time and place are universal classes and so have range. 'All' is a quantifier" [27, p. 18].

Maccia distinguishes three kinds of "knowing that": instantiation, theoretical knowing, and criterial knowing (see Table 1). For example, to categorize Earth, Mercury, Venus, Mars, etc. as instances of the concept "planet" is quantitative *instantiation*. To provide an evidentiary argument to justify the relationship between matter, energy and light, as expressed in Einstein's famous equation,



Figure 1. Illustration of three fundamental kinds of cognition based on Macchia's epistemology (drawings by Elizabeth Boling).

Table 1. Types of "Knowing That" from Maccia's Epistemology

Types of "Knowing That"	Tutorial Conditions: S: Student, Q: Object of Knowing, T: Teacher
Instential	<p>S <i>instantiates</i> Q if and only if</p> <ol style="list-style-type: none"> 1. S believes that Q. 2. S identifies Q as an instance of a kind. 3. S correctly believes Q. 4. Q is a state of affairs. 5. T knows that the above conditions hold in order that S identify Q.
Theoretical	<p>S <i>knows the theory</i> of that Q if and only if</p> <ol style="list-style-type: none"> 1. S believes that Q. 2. S is in a position to know that Q. 3. S correctly believes that Q. 4. S presents an evidentiary argument that completely justifies S's belief that Q. 5. S explicates the relevance and fruitfulness of the theory of that Q. 6. Q is a state of affairs. 7. T knows that the above conditions hold in order that S knows the theory of that Q.
Criterial	<p>S <i>knows the criteria</i> of that Q if and only if</p> <ol style="list-style-type: none"> 1. S believes that Q. 2. S is in a position to know that Q. 3. S correctly believes that Q. 4. S presents a justificatory argument to establish the credibility of criteria of that Q. 5. S demonstrates the relevancy and fruitfulness of criteria of that Q. 6. Q is a state of affairs. 7. T knows that the above conditions hold in order that S knows the criteria of that Q.

$E = mc^2$, is an example of *theoretical knowing*. To be able to justify and establish the credibility of the value "freedom of speech" is an example of *criterial knowing*. Further examples of "knowing that" are also provided on the World Wide Web at URL: <http://education.indiana.edu/ist/faculty/episcomp.html> [28].

Qualitative Intelligence: "Knowing That One"

On the other hand, "knowing that one" is ability to index "none-other." It is to discern the specific features that make some entity unique—what sets it apart from all else. For example, in Figure 1, the man recognizes his dog, Rover. In this

case, he is not seeing his dog as an instance of dog (quantitative instantiation), but instead as a specific unique dog—namely his dog, Rover.

Maccia concludes that four conditions are necessary for “knowing that one”: “1) the state of affairs which is the object of cognition must be epistemically present; 2) the presence may be perceptual or imagined, but the image must be accurate and complete; 3) the association between the knower and the known must be intimate or heightened; and 4) the object of cognition must be a unified whole with its own identity having characteristics that are discrete” [22, p. 217].

Steiner further asserts that “. . . qualitative structures, if adequate, allow one to be sensitive to the immediacy of the given, to the uniques; they are pervasive qualities. Uniques cannot be members of classes, and so no extension is involved; each is what it is. It cannot even be said of an unique that it is one of a kind” [27, p. 18].

As can be seen in Table 2, there are three kinds of “knowing that one”: recognition, acquaintance, and appreciation. An example of *recognition* is to identify our planet Earth from a wide-angle photograph taken from a satellite, distinguishing it from other bodies such as our moon, Jupiter, Venus, etc. If *acquainted* with the Washington Monument one could describe its unique qualities, such as its being an obelisk several hundred feet tall, dedicated as a memorial to the first U.S. president, which is located in the center of a cross-shaped mall, respectively ended by the Lincoln Memorial, the Jefferson Memorial, the White House, and the Capitol Building in Washington, D.C. As an example of *appreciation*, an Olympic diving judge observes and evaluates certain features of a given dive such as height of take-off, form during tucks and twists, verticality of entry into the water, and lack of splash.

Maccia refers to qualitative intelligence as ability to come to “know that one.” Further examples of “knowing that one” are also provided on the Web [28].

Roger Shank has conducted research on the primary role of *story telling and comprehension* as a fundamental element of human cognition and intelligence [29]. Stories are comprised of linguistic symbols (signs) which normally represent unique states of affairs. To be able to tell and comprehend stories requires acquaintance and perhaps appreciation in Maccia’s epistemology. Knowing a story is an example of “knowing that one.”

The reader should note that—like a double-edge sword—the same object of knowing can be known quantitatively and qualitatively. The man in Figure 1 can cognize that Rover is an instance of dog as well as his own unique pet. Furthermore, that same object may be relevant to “knowing how”—for example, knowing how to bathe the dog.

Performative Intelligence: “Knowing How”

The third basic kind of knowing is “knowing how”—to use means to achieve ends. Know-how requires cognition of particular circumstances, making

Table 2. Types of "Knowing That One" from Maccia's Epistemology

Types of "Knowing That One"	Tutorial Conditions: S: Student, Q: Object of Knowing, T: Teacher
Recognitive	<p>S <i>recognizes</i> Q if and only if</p> <ol style="list-style-type: none"> 1. S believes that Q. 2. S is completely justified in believing that Q. 3. No other statement or belief defeats S's belief that Q. 4. S selects Q from not Q and not Q from Q. 5. Q is a state of affairs. 6. T knows that the above conditions hold in order that S recognize Q.
Acquaintive	<p>S is <i>acquainted with</i> Q if and only if</p> <ol style="list-style-type: none"> 1. S recognizes Q. 2. S selects elements $[q_i \dots q_j]$ determinate of Q; and relations $[r_1 \dots r_j]$ determinate of Q. 3. Q is a state of affairs. 4. T knows that the above conditions hold in order that S be acquainted with Q.
Appreciative	<p>S <i>appreciates</i> that Q if and only if</p> <ol style="list-style-type: none"> 1. S is acquainted with Q. 2. S selects elements $[q_i \dots q_j]$ appropriate of Q; and relations $[r_1 \dots r_j]$ appropriate of Q. 3. Q is a state of affairs. 4. T knows that the above conditions hold in order that S appreciate that Q.

judgments based on the conditions, and choosing appropriate courses of action when warranted—all in order to achieve some desired outcome. In short, know-how is purposeful. "Performative structures are enactions. They allow one to act" [27, p. 18]. Table 3 provides definitions of six kinds of "knowing how," or what Maccia referred to as performative intelligence [23].

Maccia distinguishes among procedures, performances, innovation, and creation. *Procedural knowing* is exemplified by knowing a recipe for baking chocolate chip cookies, whereas *performative knowing* is illustrated by successfully making a batch of such cookies. To build a better mousetrap is an instance of *innovation*. To write this article is an example of *creation*. Further examples are given on the Web [28].

Tutorial Conditions of Knowing and AI

The reader should note that Maccia was concerned about tutorial conditions of knowing, as can be seen in Tables 1 to 3. In other words, how could a teacher tell if a student had achieved one of these kinds of knowing? As indicated above, he

Table 3. Types of "Knowing How" from Maccia's Epistemology

Types of "Knowing How"	Tutorial Conditions: S: Student, P: Means to Achieve End, T: Teacher
Procedural- Protocolic	<p><i>S knows the protocol of P</i> if and only if</p> <ol style="list-style-type: none"> 1. S iterates the constituents and succession of movements in executing the protocol. 2. The protocol is the way of performing P. 3. P is a single-pathed doing. 4. T knows that the above conditions hold in order that S knows the protocol for doing P.
Procedural- Conventional	<p><i>S knows the convention of P</i> if and only if</p> <ol style="list-style-type: none"> 1. S iterates the preferred constituents and succession of movements in executing P. 2. The convention is a way of performing P. 3. P is a multi-pathed doing. 4. T knows that the above conditions hold in order that S knows the convention for doing P.
Performative- Protocolic	<p><i>S knows how to do the protocol P</i> if and only if</p> <ol style="list-style-type: none"> 1. S has the capacity for doing P. 2. S has the facility for doing P. 3. S smoothly executes P. 4. P is a single-pathed doing. 5. T knows that the above conditions hold for doing P.
Performative- Conventional	<p><i>S knows how to do the convention P</i> if and only if</p> <ol style="list-style-type: none"> 1. S has the capacity for doing P. 2. S has the facility for doing P. 3. S smoothly executes P. 4. P is a multi-pathed doing. 5. T knows that the above conditions hold for doing P.
Innovative	<p><i>S knows how to innovate the doing of P</i> if and only if</p> <ol style="list-style-type: none"> 1. S has the capacity for doing P. 2. S has the facility for doing P. 3. S smoothly executes constituents and succession of movements into some performance P_n when P includes P_n, and P_n is not equivalent to P. 4. P is a doing.
Creative	<p><i>S knows how to create the doing P</i> if and only if</p> <ol style="list-style-type: none"> 1. S has the capacity for doing P. 2. S has the facility for doing P. 3. S smoothly executes constituents and succession of movements of $P(1, 2, \dots, n)$ into P_2 where $P(1, 2, \dots, n)$ are elements of P and P_2 is not included in P. 4. P is a doing.

was interested in characterizing different kinds of cognitive learning objectives that went beyond Bloom's taxonomy. These tutorial conditions could also be used to determine cognitive accomplishments in artificial intelligence as well, although that was not Maccia's original intention.

IMPLICATIONS OF MACCIA'S EPISTEMOLOGY FOR TUTORING SYSTEMS

Without qualitative intelligence, the other kinds of knowing are *not grounded in experience with the world around us*. This in turn implies that tutoring systems, whether natural or artificial, must possess qualitative intelligence as well as quantitative and performative; otherwise they blindly reason and follow procedures that manipulate symbol systems, images, sounds, icons, and the like with no cognition of their meaning. Inventive intelligence (innovation and creation) is also a desirable property of a tutoring system, so it can realize new ways of doing and develop deeper understandings through disciplined inquiry [28].

Computer Qualitative and Quantitative Knowing

Cognitive scientists have begun to realize the epistemological significance of "knowing that-one." Jerome Bruner reviews the evolution in cognitive psychology before discussing the current revolution, which he refers to as cultural psychology [30]. Bruner views "mind" as a creator of meanings, a special interaction through which it both constitutes and is constituted by culture. This is an entirely different paradigm than the view of mind as information processor, a view that was previously dominant in cognitive psychology. Bruner believes that "... we shall be able to interpret meanings and meaning-making in a principled manner only in the degree to which we are able to specify the structure and coherence of the larger contexts in which specific meanings are created and transmitted. . . . It simply will not do to reject the theoretical centrality of meaning for psychology on the grounds that it is 'vague.' Its vagueness was in the eye of yesterday's formalistic logician. We are beyond that now" [30, pp. 64-65].

The larger context to which Bruner refers is culture. When we communicate we use "signs." "Signs" are the "stuff" of communication which are embedded in our culture. "Signs" are not limited to spoken and written words (linguistic symbols), but also include icons, gestures, pictorial representations, non-linguistic sound patterns, smells, touch, facial expressions, expression of emotion, demonstration by enactment, etc. "Signs" take on meaning for humans through use during situated action.

It is true that computers can process "signs" such as key presses and mouse clicks in the sense of decoding, storing, transmitting, and encoding them. However, present-day computer systems cannot be said to understand the "signs" they manipulate. The only "signs" computers do appear to know are the bit patterns

which represent machine language instructions. These instructions (patterns) are literally hardwired into a computer. Each instruction is associated with a primitive action the computer can do, such as move a collection of bits from one place to another, and perform arithmetic and logical (Boolean) operations on those bit collections.

While a collection of bits may be an encoded representation of external "signs," a computer is literally isolated from the culture in which those "signs" are given meaning. Computer systems are blind and deaf to the symbol systems, images, icons, and sounds that they process. Indeed, John Searle concludes, "... the symbols have no meaning; they have no semantic content; they are not about anything. They have to be specified purely in terms of their formal or syntactical structure. The zeroes and ones, for example, are just numerals; they don't even stand for numbers" [31, p. 31].

Searle provides a thought experiment—the Chinese Room example—to illustrate this point. In the imaginary Chinese room, someone who understands English but *no* Chinese follows a rule book (in English) for manipulating these symbols. When certain Chinese symbols are passed into the room from the outside, the person follows the rules and passes some other Chinese symbols out of the room. From outside the room, it might appear that the person (or a computer program executing these same rules) might have passed the Turing Test—i.e., the answers to questions posed to the Chinese room were indistinguishable from a native Chinese speaker. However, the person inside does not understand a word of Chinese but is simply following the rules for manipulating the symbols. Searle concludes that "[u]nderstanding a language, or indeed, having mental states at all, involves more than just having a bunch of formal symbols. It involves having an interpretation, or a meaning attached to those symbols" [31, p. 33].

From another perspective, Roger Shank concludes that "[i]ntelligence, for machines as well as people, is the telling of the right story at the right time in the right way. Thus, the key problem in artificial intelligence is the indexing problem" [29, p. 242]. Maccia also views indexing as that which distinguishes qualitative from quantitative intelligence: "... recognition, acquaintance, and appreciation employ the ostensive operator 'none-other'. '*None-other*' is an operator which is *indexical not instantial*. It does not negate the 'other' as 'non-other'. It separates one from all others" [italics added, 22, p. 215].

Indexing, then, is something quite different from abstraction (see Tables 1 and 2). When we index something, we discern those features which identify its uniqueness. For example, when a victim of an assault describes the appearance of the assailant to a police sketch artist, he or she is indexing those features such as a scar below the left eye, thick eyebrows, receding hairline, etc., which might allow others to recognize the assailant. As another example, several years ago I absent-mindedly attempted to unlock the door of someone else's old car in a campus parking lot of the same model, color, and year of my own. When my key did not

work, I realized that this car did not have the right worn spots in the upholstery, and the rust spots on the body were in the wrong places. Those unique features which indexed my car were missing.

If Maccia, Shank, and Searle are right, then current efforts to design artificially intelligent tutoring systems capable of understanding human languages will not succeed unless they can satisfactorily address the grounding problem—grounding that occurs through qualitative cognition. That is indeed a very significant implication. From Bruner's perspective, such systems will never learn the meaning of the 'signs' they manipulate unless they become truly interactive with the culture in which the signs are embedded. In other words, computer systems will somehow need to "live in" and experience the culture with us. This appears to be the very same grounding that concerned Maccia in his discussion of genetic epistemology of natural intelligence [22, 23].

Computer Performative Knowing

Computers are very good at following protocols such as instructions for computation and logical operations. When such instructions are properly sequenced, a computer can do complex tasks such as flying airplanes, multivariate statistical analysis, medical diagnosis, buying and selling of stock, diagnosing faults in equipment, anti-lock braking to stop a vehicle, etc. To date, computer systems have evidenced more ability in the domain of performative intelligence, than in any of the others. This should not be surprising, since computation is an example of know-how. It is also not surprising that those artificially intelligent tutoring systems which have been successful usually teach know-how (e.g., solving algebraic equations, diagnosing equipment malfunctions, etc. [6]). See Table 3, especially protocolic and conventional performance.

Such "applied intelligence" is one of the areas where computers can literally extend the capacity and facility of human know-how. Expert systems and neural networks are two of the best examples to date. In effect, computers can learn how to do tasks—including deductive reasoning—that humans perform. Once learned, the execution of these tasks is normally much faster and more reliable than we ourselves can do them. Indeed, electronic computers were invented because human computers were too slow and error prone. Prior to the 1940s "computer" was the name of a human occupation. Since then the methods of teaching electronic computers how-to-do have advanced from low- and high-level programming languages to extant methodologies such as knowledge engineering for expert systems and training of neural networks via practice and feedback.

Because computer systems exhibit performative intelligence, we can teach them to do tasks. It is this very capability that makes it possible to use computers as an interactive medium for instruction and learning. It is *interaction* which sets computers systems apart from other media such as books, television, and film. However, present-day computers literally do not understand the culturally bound

meanings of the messages which they manipulate during these interactions because such computers lack qualitative intelligence.

Computer-mediated learning products—whether they are guided practice exercises, tutorials, simulations, games, hypertext, multimedia, or Web documents—merely carry out the directives (rules, programs, scripts, mark-ups) from a human tutor who originally designed the particular instructional or informational system and its subject matter. A computer system has no idea of the meanings of the messages (i.e., groups of “signs”) being sent back and forth between the tutor and students. A computer system is a medium which *conveys* those human messages.

Thus, it is clear that computers can learn or be taught at least some kinds of know-how. Further discussion of limitations of computer ability to innovate and create, which are additional kinds of performative intelligence in Maccia’s epistemology (see Table 3), is available on the Web [28].

IMPLICATIONS FROM NEUROLOGICAL RESEARCH ON GROWTH OF MIND¹

The discussion to this point has represented views from epistemology and cognitive psychology. Recent neurological research on how human brains develop sheds further light on human cognition and the challenges faced by researchers in artificial intelligence. Human genes and their environment interact in ways that resemble a “dance” between nature and nurture. Evidence is mounting that human sensory and emotional experience—especially during the first ten years of life—literally build neural connections in the brain [16, 19-21].

Neuroscientists have found that electrical activity in the brain actually changes its physical structure. At birth a human brain contains approximately 100 billion neurons. As a baby’s brain develops it overproduces neurons and quadrillions of connections among them. Then experience with the world begins to prune the emerging structure. In essence, a child’s behavior and experience influence the physical structure and formation of her brain [19, 20].

Moreover, since each individual’s experience is unique, we each evolve different brain structures. Sensitive periods are critical during early years. If a child’s brain is deprived of a stimulating environment, brain growth is retarded [21]. It is no longer a question of nature or nurture, but a “dance” between our genes and our experience [16]. Starting around the age of ten, the brain begins to destroy the weakest connections; thus, the remaining structures constitute a unique brain that is grown out of experience.

¹ This discussion is not intended to imply that artificially intelligent systems will need to resemble human nervous systems. As Minsky has noted [10], there is no reason that such systems will need to be constructed isomorphic to biological systems, any more than airplanes need to fly in the way birds do.

Stanley Greenspan provides clinical evidence that "emotion organizes experience and behavior" [16, p. 23]. He believes that the Western tradition of treating cognition and emotion separately—going back to the ancient Greeks—has blinded us to the role of affect in organizing experience. "We are able to get at our stored experience so rapidly and reliably because our affective capacity organizes information in an especially functional and meaningful manner" [16, p. 29]. He suggests that

if . . . information is dual-coded according to its affective and sensory qualities, then we have a structure or circuitry set up in our minds that enables us to retrieve it readily. . . . Affects enable us to identify phenomena and objects and to comprehend their function and meaning. Over time they allow us to form abstract notions of interrelations. . . . Affect, behavior, and thought must be seen as inextricable components of intelligence. For action or thought to have meaning, it must be guided by intent or desire (i.e., affect). Without affect, both behavior and symbols have no meaning [16, pp. 30-37].

Maria Montessori realized this important connection between affect and cognition about 100 years ago when she began to design learning environments for young children. Central to her methods was that a child's natural curiosity should motivate learning. Therefore, adults should structure the environment in ways which allow that curiosity to be satisfied through a child's interaction with the environment. The child in effect is deciding what to learn, when, and for how long [18]. Montessori was acutely aware of the sensitive periods of childhood which modern neuroscience is documenting [e.g., 16, 19-21], and developed her curriculum to match those changing early needs.

It would appear that the mounting evidence from neuroscience and clinical psychology has significant implications for both education of human beings as well as for research in artificial intelligence. Greenspan concludes:

Computers may be able to perform certain cognitive operations, even more effectively, and certainly faster, than humans. But unless they acquire the ability to experience and react to emotion, silicon chips will be unable to exercise intelligent discrimination. . . . What separates human intelligence from that of computers, robots, androids, and any other cyber-creatures we can imagine, is the fact that we possess a nervous system capable of—indeed, specifically designed for—generating and evaluating affect. . . . Unless and until we solve the problem of creating living cellular reactivity and affects, as well as the capacity to abstract patterns of affects, in an artificial form, no machine will think in a truly human way [16, pp. 126-127].

CONCLUSION

In this article I have discussed the limitations of intelligence achievable by computer systems in 1997. Clearly, such systems exhibit some kinds of know-how, referred to as performative intelligence in Maccia's epistemology. It is also patent that such systems can manipulate symbol systems including codified language. I am using a computer to write symbols at this very moment, although it does not understand the meaning of the signs (words and images) that are being displayed on the monitor or printed with a laser device. It does not possess any significant degree of qualitative intelligence or quantitative intelligence itself, in part because it does not have the sensory or emotional capacities through which we humans derive meaning of the signs used in communication among ourselves.

On March 12, 1997, the birthday of the fictional computer, HAL, was in actuality celebrated in Urbana, Illinois [32]. Arthur C. Clarke imagined HAL when writing his novel, *2001: A Space Odyssey*, some thirty years ago [33]. Can a real HAL be created? I make no predictions. The tutorial conditions of "knowing that," "knowing that one," and "knowing how" (listed in Tables 1 to 3) indicate some of the problems that the artificial intelligence community will likely need to address.

Artificially intelligent tutoring systems are beyond our reach at this juncture except for those which teach certain kinds of know-how. Therefore, we should view computers in education largely as *media* through which a multitude of teachers can convey their messages. This does not mean that the role of classroom teachers is diminished. Teachers and students can be empowered by these additional learning resources.

REFERENCES

1. P. McKeown, *Living with Computers* (4th Edition), Dryden Press, Fort Worth, Texas, 1993.
2. M. McLuhan, *Understanding Media: The Extensions of Man*, McGraw-Hill, New York, 1964.
3. W. Gates with N. Myhrvold and P. Rinearson, *The Road Ahead*, Viking, New York, 1995.
4. D. Merrill, *Automating Instructional Design: Recent Advances in Instructional Transaction Theory and ID Expert*, paper presented at the conference of the Association for Educational Communications and Technology, Nashville, Tennessee, February 1994.
5. G. Kearsley, *Artificial Intelligence and Instruction*, Addison-Wesley, Reading, Massachusetts, 1987.
6. D. Sleeman and J. Brown (eds.), *Intelligent Tutoring Systems*, Academic Press, New York, 1982.
7. H. Dreyfus, *What Computers Can't Do: The Limits of Artificial Intelligence* (Revised Edition), Harper and Row, New York, 1979.

8. H. Dreyfus, *What Computers Still Can't Do: A Critique of Artificial Reason*, MIT Press, Cambridge, Massachusetts, 1992.
9. M. Arbib, *Computers and the Cybernetic Society*, Academic Press, New York, 1977.
10. See *The Machine that Changed the World*, Films for the Humanities and Sciences, Princeton, New Jersey, Video Series, 1993. Part 4, "The Thinking Machine," provides an excellent discussion of the challenges of knowledge representation and natural language understanding. Numerous scholars are interviewed on this PBS documentary including Marvin Minsky, Edward Feigenbaum, and Hubert Dreyfus. The Turing Test is also well illustrated.
11. T. Frick, Computerized Adaptive Mastery Tests as Expert Systems, *Journal of Educational Computing Research*, 8:2, pp. 187-213, 1992.
12. R. Neapolitan, *Probabilistic Reasoning in Expert Systems*, Wiley, New York, 1990.
13. P. Winston and M. Prendergast (eds.), *The Commercial Uses of Artificial Intelligence*, MIT Press, Cambridge, Massachusetts, 1984.
14. D. Lenat and R. Guha, *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*, Addison-Wesley, Reading, Massachusetts, 1990.
15. E. Seels and R. Brooks (eds.), *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*, L. Erlbaum Associates, Hillsdale, New Jersey, 1995.
16. S. Greenspan with B. Benderly, *The Growth of the Mind and the Endangered Origins of Intelligence*, Addison-Wesley, Reading, Massachusetts, 1997.
17. J. Dewey, *Democracy and Education: An Introduction to the Philosophy of Education*, The Free Press, New York, Chapter 11, pp. 139-151, 1916.
18. M. Montessori, *Spontaneous Activity in Education*, Schocken Books, New York, 1965.
19. W. Greenough and J. Black, Induction of Brain Structure by Experience: Substrates for Cognitive Development, *Developmental Behavioral Neuroscience*, 24, pp. 155-299, 1992.
20. I. Weiler, N. Hawrylak, and W. Greenough, Morphogenesis in Memory Formation: Synaptic and Cellular Mechanisms, *Behavioural Brain Research*, 66, pp. 1-6, 1995.
21. S. Schanberg and T. Field, Sensory Deprivation Stress and Supplemental Stimulation in the Rat Pup and Preterm Human Neonate, *Child Development*, 58, pp. 1431-1437, 1987.
22. G. Maccia, Genetic Epistemology of Intelligent Natural Systems, *Systems Research*, 4:1, pp. 213-218, 1987.
23. G. Maccia, *Genetic Epistemology of Intelligent Systems: Propositional, Procedural and Performative Intelligence*, paper presented at the Fourth International Conference on Systems Research, Cybernetics and Informatics, Baden-Baden, West Germany, August 1988.
24. B. Bloom et al., (eds.), *Taxonomy of Educational Objectives, Volume I, Cognitive Domain*, Longmans, Green, New York, 1956.
25. G. Ryle, *The Concept of Mind*, Barnes and Noble, New York, 1959.
26. I. Scheffler, *Conditions of Knowledge*, Scott-Foresman, Glenview, Illinois, 1965.
27. E. Steiner, *Methodology of Theory Building*, Educology Research Associates, Sydney, Australia, 1988.

28. T. Frick, *Implications of Maccia's Epistemology for Use of Computers in Education*, Indiana University School of Education Web Document, Bloomington, Indiana, URL <http://education.indiana.edu/ist/faculty/episcomp.html>, 1997.
29. R. Shank, *Tell Me a Story*, Charles Scribner's Sons, New York, 1990.
30. J. Bruner, *Acts of Meaning*, Harvard University Press, Cambridge, Massachusetts, 1990.
31. J. Searle, *Minds, Brains and Science*, Harvard University Press, Cambridge, Massachusetts, 1984.
32. A. Clarke, *3001: The Final Odyssey*, Ballantine Books, New York, p. 254 (acknowledgments), 1997.
33. A. Clark, *2001: A Space Odyssey*, The New American Library, New York, 1968.

Direct reprint requests to:

Dr. Theodore W. Frick
Associate Professor
School of Education, Rm. 2218
Indiana University
Bloomington, IN 47405-1006