NEW PERSPECTIVES ON COGNITION AND INSTRUCTIONAL TECHNOLOGY

William J. Clancey

Institute for Research on Learning
2550 Hanover Street
Palo Alto, CA 94304
USA

In E.Costa (Ed.), New Directions for Intelligent Tutoring Systems: Proceedings of the NATO Advanced Research Workshop on New Directions for Intelligent Tutoring Systems, Sintra, Portugal, 6–10 October, 1990 (pp. 3–14). Berlin: Springer-Verlag, 1992.

New Perspectives on Cognition and Instructional Technology

William J. Clancey

Institute for Research on Learning, 2550 Hanover Street, Palo Alto, CA 94304, USA

Abstract: Diverse perspectives from anthropology, philosophy, and linguistics lead us to view human knowledge as constructed moment-by-moment in interaction between people and their environment. The dynamics of human behavior is central, embracing all levels from *perception* (by which information is defined by the observer, not passively received), *interpretation* (by which representations are commented upon and thereby given meaning, not stored and retrieved from memory and simply "applied"), and *communication* (by which knowledge emerges through group interactions, not transmitted as predefined packets). This new conception leads us to view computer models in a new way.

Keywords: computer-aided instruction / learning / cognitive modeling / knowledge-based tutoring / intelligent systems / interactionalism / social knowledge / expert systems / student models

Introduction

Over the past five years, diverse perspectives from anthropology [10, 11, 17], philosophy [8], and linguistics [13, 18, 19] have become increasingly central to cognitive science research, radically changing our understanding of human cognition. We are led to view human knowledge as something constructed moment-by-moment as an interaction between people and their environment, in order to cope with the situations at hand. Knowledge is not something stored and retrieved unchanged from memory and then simply "applied" [4, 6, 5, 9]. This new conception leads us to view computer models (e.g., "expert systems," "student models," and "intelligent tutoring systems") in a new way [14, 19]. In considering designs for learning environments in particular, we now pay special attention to how knowledge is constantly being created and given meaning through social interactions, as opposed to processes that take place exclusively inside of the heads of individuals. We consider the dynamics of human behavior to be central, embracing all levels from *perception* (by which information is defined by the observer, not passively received), *interpretation* (by which representations are commented upon and thereby given meaning, not retrieved and mechanically applied), and *communication* (by which knowledge emerges through group interactions, not transmitted as predefined packets of meaning). This paper introduces these ideas and considers the implications for how we use intelligent tutoring system technology.

What is learning?

We remain cognizant that, of course, it is individual people who articulate theories and act. However, what individuals say and do is not merely *conditioned* on what other people are saying and doing: Individual behaviors can only be described and understood as part of a larger whole, the goals and emergent activities of a social system. By analogy, the study of knowledge and learning is like thermodynamics, crucially built on an understanding of the properties of atoms and molecules, but fundamentally not about individuals in isolation. Rather, to understand what people believe and how they come to believe and do different things, we must broaden our perspective to describe the encompassing social system, the emergent effects of interactions, what people know about these effects, and how they seek to change them by their behaviors.

One way of summarizing this "interactional" point of view is that learning always involves adapting to the social and physical environment. In contrast, cognitive science has tended to view learning as something that takes place exclusively in the brain. But individual learning cannot be separated from the developing interactions of the group. How do people get their work done as a team and how does successful learning occur on the job? What *resources* does the environment make available as cues for structuring and supporting work [10, 17, 20]? Resources include reference materials, the physical layout of the office or shop floor, as well as team members with different specialities who provide direction and information.

We are particularly concerned about how computer models, no matter how realistic, artificially define the nature of a task by eliminating perceptual and problematic aspects of the actual setting [16]. In part, this awareness can help us develop more appropriate simulations (or justify choices that seemed arbitrary or a matter of opinion before, such as simulating sounds or simulating other agents). With the aim of promoting learning, we focus on the design of tools for accomplishing collaborative work, not isolated "tutoring systems" or even "consultation systems," which are often described in a decontextualized way, as if they could be understood or meaningfully used on any desk in any room.

To summarize some basic hypotheses that are radically changing our research on learning:

- **Representations**--what we say, write, draw, visualize--are constructed as *perceivable things*, they are not stored and retrieved in the brain. Human memory consists of the capacity to replay, organize, and recombine sequences of behavior (phrases) [1, 15]; this is *a memory of processes of perceiving and behaving*, not *descriptions* of how processes in the world or behavior routines appear to an observer [6]. Representations are constantly reinterpreted by new representations in an ongoing sequence of commentary. Their meaning is never fixed or definable [18].
- **Information** is created through individual perceptions that organize and reinterpret behavior and models of the world, as part of a larger process of interaction with the social and physical environment. Thus, the *cognitive* (individual statements and behaviors) is organized by the *social* (the dynamics of group activities and representations). Information is not given as objective pieces of data, but created by interaction of perceptual and external processes [12]. Information is not a sequence of words in a pipeline. To speak is to perceive is to represent is to generate knowledge [6, 9, 18]. This means that problems in the workplace, as well as theories for resolving them, are not supplied or predefined, but rather created anew in every situation by the participants in their interaction [16].
- **Learning** is how we describe and account for change in behavior. Strictly speaking it should not be viewed as acquiring some*thing*, so much as the development of a new routine, a capacity to interact more automatically in some setting. What is learned is not the description of

the behavior, but a process for behaving. We must distinguish between our *"knowledge level" descriptions* as observers and the *internal mechanisms* that couple the individual to his or her environment. Knowledge-level descriptions are the province of the theoretician, accounting for behavior, they are not programs and data stored in the head of the subjects being studied [4]. Therefore, what is learned is not an observer's descriptions of how the performance will appear, but rather a capacity to move, look, and speak in a particular way in some situation.

What are computer tools?

These hypotheses have a dramatic effect on the design and use of computer tools. For example, we are less sanguine about developing an expert system and "transferring the knowledge" to a student. We are reconsidering how *theories* and *behaving* interact, that is, how talking about a situation, conceiving theories about it, and action are related. We do not by any means abandon the modeling methods of knowledge engineering and cognitive science [3, 7]. Rather, *we seek a radical reinterpretation of how knowledge-level models relate to human memory and perception.* We seek appropriate ways of using such models in the workplace in view of the dynamic, social nature of representation construction and reinterpretation. In essence, we move from viewing knowledge as something that can be fully captured and written in a "knowledge base," to accepting and working from the inherent and ongoing change in knowledge in every human perception and social interaction.

How can we design computer tools to enhance the everyday process by which people construct their own understanding of the task and the methods by which they organize their behavior? What assistance might facilitate sharing and reworking of conceptions? What recordkeeping by individuals would help a team develop its understanding of trends in the workplace over time? What mechanical aids would promote questioning, probing, and theory formation in the course of everyday work? How can computer controls be designed to facilitate their responsible use by people? *How can a program make people aware of its potential as a source of information and the weaknesses of its fixed design?* What features enable a team to customize a program to meet their changing goals and specific needs? How can a program help people detect inadequacies in their work and know what information they need to improve their understanding?

In summary, we reframe how intelligent tutoring system (ITS) technology (e.g., knowledge representations, student modeling, explanation) is used by viewing it in the context of computer tools for everyday work. As indicated here, all of the pieces in the puzzle--representations, information, learning--are themselves viewed in a new way. Perhaps the most central theoretical claim driving this new analysis is that knowledge is not stored in networks in the brain, but dynamically generated and reinterpreted in our words, drawings, images, in our ongoing social and physical interactions. Hence, we say that knowledge is situated; it is never in hand or describable apart from the setting in which the individual is interacting. Knowledge is not a thing, but a capacity to behave adaptively.

Reconsidering how ITS technology is used

How can we apply our revised theoretical perspective about the nature of knowledge and representations to the development of a new generation of computer tools? Specifically, we consider here how this conception changes our view of knowledge bases: They are not something to be delivered and memorized by a student, but are representations, created by social organizations and given meaning by the social and physical context. Both organization and context are not definable, and hence cannot be placed inside the knowledge base itself. Our changing perceptions of our social and physical setting change how we interpret a knowledge base, changing what the representations mean and hence their implications for our behavior. Rather than using an ITS program to deliver a predetermined, "finished" model to a student, we ask, *what kind of computer tools would facilitate the student's development of his own models?* We approach this question by considering failure analysis of computer models and how a knowledge base can be made into a "working document."

Glassbox Representations and Failure Analysis. Second generation expert system design (e.g., NEOMYCIN [7]) emphasizes that expert systems contain two kinds of models: a model of some domain system (e.g., a physical device) and a model of some reasoning process (e.g., a diagnostic or design procedure). Thus, kinds of knowledge are stated more abstractly, so they can be better explained, used for multiple purposes, and reused in the design of future programs. For example, "explanation-based learning" techniques enable a program to help users reason about and improve domain models.

Our new hypotheses about human knowledge suggest that explanation and "glass box" design is important because programs lack the flexibility of reinterpretation and reorganization of human perception. Indeed, the computer is "using representations" without giving them meaning. Any "interpretation" by the computer is grammatical. That is, it follows predefined relations between terms and predefined ways of organizing how data is processed. The inherent limitations of today's computer tools, relative to human perception and learning, place a primacy on exposing the computer's line of reasoning so people can be alerted about the limitations of the program's recommendations, and have a basis for improving them. To put this simply, a knowledge base *is just a model.* Models describe or simulate routine, stable interactions in the world. That is, models deal with patterns. Detecting when a model is inadequate and adapting it in subtle ways to the nuances of each new setting cannot be fully automatized by using models alone.

With the perspective that knowledge bases contain models, it is natural to *apply engineering approaches for describing the assumptions behind these models and formalizing the bounds of their applicability.* Recent research, such as in GUIDON-DEBUG, is considering means for integrating multiple representations, as leverage by which the program can detect weaknesses in its models and alert the user [2, 7]. GUIDON-DEBUG is a knowledge acquisition program in which the NEOMYCIN diagnostic expert system is enhanced by including explicit statements of what constitutes a good diagnosis. GUIDON-DEBUG then applies these constraints to evaluate NEOMYCIN's diagnoses and track them back to missing or incorrect facts in the domain model.

We might investigate how representations of strategy and representations of the setting that makes a strategy meaningful enable people to detect failures (in themselves, other people, or machines). Our new conception suggests that a strategy is not like a program that people "apply" by rote. Instead, a strategy is an observer's description of historical patterns in behavior. These patterns emerge from people interacting with each and their environment. They are not predefined plans that serve as instructions, but emergent, constantly adaptive routines and habits [6,17]. One role of strategic representations is to organize how we perceive new information for coping with failures; strategic talk (i.e., descriptions of what we typically do) directs our attention. We hypothesize that this *metacognitive knowledge*--representations of patterns in intelligent behavior as well as representations of the social and physical setting--is essential for expert system design (as glassboxes) that will promote human learning [2].

Notice that improvements to program design are coupled with the awareness of what cannot be incorporated in the program and remains the province of the people who interact with the program. Indeed, the role of glassbox and failure-sensitive designs is not merely to automate work, but to facilitate what only people can do.

The Working Knowledge Base. Knowledge acquisition is the process by which knowledge engineers construct expert systems. There are two basic concerns: interviewing experts and representation of knowledge. *We now realize that knowledge acquisition is a theory formation process, not a matter of simply writing down what has been preconceived and stored in the expert's brain.* In particular, we must make a distinction between an articulated, intellectual understanding of a task and the "activity-based knowledge," by which a person knows what to do when in a particular situation [20]. Such know-how cannot be reduced to theoretical knowing-that. This is similar to the previous observations that a strategy representation (knowing-that) is generated as an observation and is not the mechanism (knowing-how) by which behavior is generated.

For a strategic representation to change behavior, it must be external, out where it can be perceived as words, diagrams, gestures, etc. What can be said about a task and its environment is neither bounded nor definable from any given perspective. Models are inherently subjective and selective [18]. Thus, a knowledge base captures how people *talk* about a task environment from some perspective. Such concepts are useful for orienting and organizing work, but they do not constitute and cannot replace the adaptability of the human expert.

Knowledge acquisition is not delivering knowledge, so much as providing a tool for the community to collect and refine models of their work environment. Can we enhance their collaborative construction of better models and operating procedures by providing communication tools? In this way, we view the knowledge base as a working model, something never finished, always questioned and augmented through experience. Integral to this is a means of recording and indexing a case library-- a record of what has happened in the past, which can be inspected and probed as theories are developed to cope with changing, difficult situations. Can we provide a means for annotating and sharing the detection of problems and the work-arounds people develop? A simple example of this is the use of e-mail bulletin boards. For example, Macintosh special interest groups allow programmers to ask for help and share tips without ever meeting the people they communicate with.

Consistent with our view of knowledge, *we don't view expert systems as a method of delivering what has been preconceived, but of promoting the means by which the team will construct its own models through their workplace interactions.* That is, our methods for knowledge acquisition should be integrated in the tools we provide for the workplace, which are tools for learning. This can be summarized as a variety of concerns:

Team orientation: An individual is not learning the subject matter so much as learning how to participate in a team's ongoing construction of what the subject matter is.

On-the-job learning: Work is inherently an activity of learning, not performing a predefined task, but constantly reinterpreting what is relevant and what constitutes good performance.

Assessment: Evaluation of learning is a natural activity of any community of practice, as people constantly form opinions about individual specialties and foibles, rather than something objective and imposed from outside.

Our theoretical perspective suggests that later reuse of a knowledge base can be enhanced by even informal annotations that record the history of the knowledge base's development and the rationale for design modifications. While they don't often reflect on this fact, knowledge engineers must reinterpret what happened in the past when they develop a new theory.

For example, a rule put in an expert system because of one case might need to be reinterpreted for a later case. Having a record of this first case could greatly change how the second case is interpreted. Thus, a knowledge base or any computer representation might be viewed in terms of accumulating layers of case experience and representation, in which we greatly benefit by being able to see through the entire past history, rather than just having a single layer with no annotations, which is the current case in most expert systems. Crucially, we can make these rationales available to the team on the workplace (or to the expert system itself) so the advantages and limitations of the computer system can be better understood.

What features for annotating and sharing experience in using a computer tool would *reflect back the team's performance,* so they became more aware of the quality of their work, knowing when to question what they are doing and have leads on how to change what they are doing? As an example, consider the disease taxonomy of NEOMYCIN. Our new perspective suggests that we view such medical knowledge not as a *product* to be delivered to a student, but as a model of a practice. In this respect, NEOMYCIN's disease taxonomy is greatly impoverished. Besides learning the various diseases and their relations, we would want the student to learn the following:

- ❑ Why do we taxonomize?
- ❑ What gets glossed? How do we know when this stereotypic view of physiological processes is misleading?
- ❑ Who knows this taxonomy; what is its origin?
- ❑ What is the nature of disagreements; how do they arise; how are they settled?
- ❑ How are taxonomies related to medical research? Are they becoming unnecessary as we develop better mechanistic models?
- ❑ What are good ways to keep a taxonomic perspective up-to-date?

By this view, knowledge is *knowing how to live* in a community, not static facts or procedures that *represent* what people do. We might say that we want a student to learn the *complement* of a conventional expert system knowledge base. Where did it come from? In what setting will it be useful? These matters can be represented, but the world and experience of the physicians (what they perceived and did in the past) cannot itself be put inside the machine. Becoming aware of this tension--an evolving dialectic between our environment and our representations---and skills for managing it are at the heart of *problem framing*, a step above and always outside what we have previously represented. This is much closer to what we want the student to learn and what we want our computer tools to support [16].

Learning Tools: Mediating conversations

The ideas of failure analysis and the working knowledge base emphasize that computers are tools that people use in a social setting. We can vividly illustrate this idea by considering two ways of viewing expert systems. In Figure 1, we show the traditional view of expert systems. The idea is that expertise is rare. Experts cannot always be present where their knowledge is needed. Therefore, the traditional view is to clone the expert by automating his reasoning in an expert system. The user (e.g., a medical intern or nurse) interacts with a computer program in which medical knowledge is formalized.

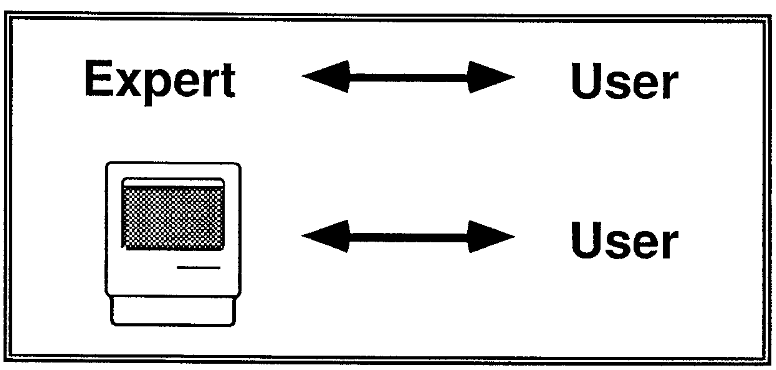

Figure 1. Traditional view of expert systems: Automate the role of the expert, replace the expert.

A more realistic view is that computer programs enable people to perform jobs they couldn't have done as well before. We are not replacing people so much as raising the level of competence of lessexperienced workers (Figure 2). The idea is that the computer is a tool for mediating conversations [14, 19]. In this diagram, the user is a patient (more generally, a client). The service person is an intern or nurse. The expert still exists; he is available for consultation. But he is usually contacted only for difficult situations, those which can't be routinely handled, but require some customization of the standard procedures that have been formalized in the computer.

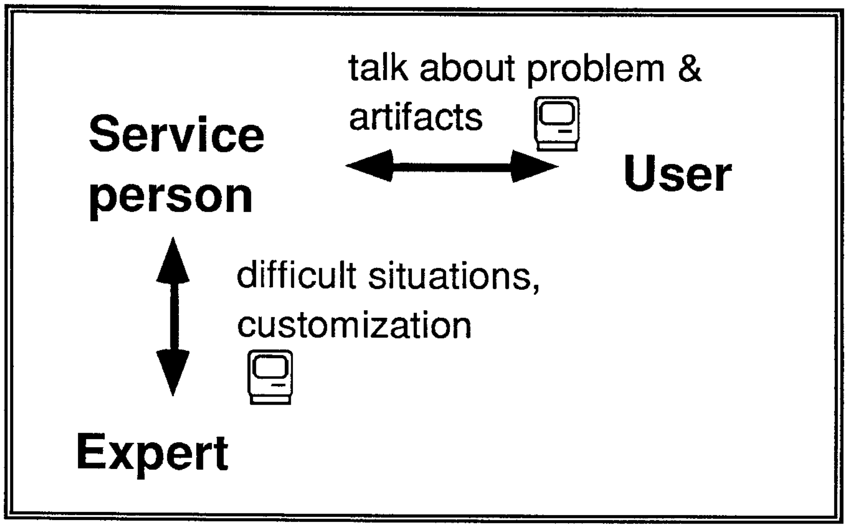

Figure 2. More realistic view of expert systems: Computer representations provide a medium for facilitating conversations.

Crucially, we view the computer tool not as doing anyone's job for them, but as a means for people to represent community practices and the theories that justify them, as well as the details of specific situations, and reflect upon them. The service person uses the computer representations (often displayed as on-line forms) as a means of organizing a problem and focusing her interaction with the patient. Similarly, the same representations form a conversation piece for the service person's interactions with the more senior, expert physicians. When problems arise, the concrete representation of the diagnostic and therapeutic procedures on-line, coupled with the computer's interpretation of the current problem, provides a well-grounded means for the expert to reconsider standard practices, and make whatever modifications are necessary to the knowledge base of the computer program.

In summary, it is more realistic and fruitful to view computer tools as enabling people to do their jobs more efficiently, to carry out new kinds of tasks. Indeed, *rather than removing people from the loop, properly designed job aids help people with different expertise and roles to work together more effectively.*

To be consistent with the theory of situated knowledge, such computer tools must be developed within their context of use. In the workplace environment the creation and reproduction of communities of practice, exemplified by apprenticeship in the hospital setting, are integral, and highly intentional, processes [11]. To paraphrase Lave and Wenger's analysis, the existence of a well-defined social institution, such as a hospital and medical school, facilitates the study of the relation between social and epistemological boundaries. On the one hand, the *institutional boundaries* are delineated by the social formation, including its structural and process organization and its competing constituencies. On the other hand, the *epistemological boundaries* are generated by the practice of the communities and the various forms of membership within them (e.g., hospital roles based on task specialization and seniority shape the topics and curricula of the medical school). Our challenge is to reflect the dynamic character of this social-epistemological dialectic in both our models that justify practice, as well as the tools that will facilitate the ongoing evolution of this interaction.

Conclusion

Our study of knowledge moves from stuffing it into knowledge bases, in which it is viewed as having a static character, to articulating how it arises in a social context and is given new form in every new situation. In parallel, our study of instructional programs moves from stuffing knowledge representations into the heads of students, to providing tools by which people can articulate their own theories. Crucially, these proposed computer tools are designed with the social context in mind by which representations are always being created, shared, recorded, and reinterpreted as people do their work. Programs to facilitate learning are no longer designed as one-dimensionally directed "from us to the students," but are instead designed with the historical, evolutionary, dynamic, and interactional character of knowledge in mind.

Simply put, tools for learning must be more like cameras and video-recorders than merely television screens or monitors. They must become a medium that prompts us to look out into the world in a new way, to make a statement, to adapt and reorganize what others have said, to present it to others, and then to reflect on what we have done. This dynamic view of knowledge provides an exciting, intuitively more balanced perspective for applying computer technology to education, but the work of realizing its potential has just begun.

Acknowledgements

Fitting the theory espoused here, most of the ideas in this paper emerged through my interactions with colleagues at the Institute for Research on Learning. I would particularly like to thank Penny Eckert and Jim Greeno, who co-authored a concept paper from which this paper is drawn. Discussions with Phil Agre, John Seely Brown, Brigitte Jordan, Jean Lave, Charlotte Linde, Jeremy Roschelle, Susan Stucky, and Etienne Wenger were equally influential.

References

1. Bartlett, F. C. *Remembering--A Study in Experimental and Social Psychology*. Cambridge: Cambridge University Press. (First published 1932.) 1977

2. Clancey, W.J. The knowledge engineer as student: Metacognitive bases for asking good questions. In H. Mandl, & A. Lesgold (Ed.), Learning Issues in Intelligent Tutoring Systems Springer-Verlag. 1988

3. Clancey, W.J. Viewing knowledge bases as qualitative models. *IEEE Expert* **4**(2) 9-23, Summer 1989

4. Clancey, W.J. The knowledge level reinterpreted: Modeling how systems interact. *Machine Learning* **4**, 287-293. 1989

5. Clancey, W.J. Why today's computers don't learn the way people do. In P. Flasch and R. Meersman (editors), Future Directions in Artificial Intelligence. Amsterdam: Elsevier. In press.

6. Clancey, W.J. The frame of reference problem in the design of intelligent machines. In K. vanLehn (editor), Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition (Hillsdale: Lawrence Erlbaum Associates). 1991

7. Clancey, W.J. Model Construction Operators. To appear in *Artificial Intelligence*. 1991

8. Dennett, D.C. Precis of "The Intentional Stance." *Behavioral and Brain Sciences* **12**(3), 495-546. 1988

9. Greeno, J.G. *Situations, mental models, and generative knowledge.* In D. Klahr and K. Kotovsky (editors), Complex Information Processing: The impact of H. A. Simon. Hillsdale, NJ: Lawrence Erlbaum Associates. 1988

10. Lave, J. *Cognition in Practice.* Cambridge: Cambridge University Press. 1988

11. Lave J. and Wenger E. Situated Learning: Legitimate Peripheral Participation. New York: Cambridge University Press. In press.

12. Reeke, G.N. and Edelman, G.M. Real brains and artificial intelligence. *Daedalus,* **117** (1) Winter, "Artificial Intelligence," 1988

13. Rommetveit, R. Meaning, context, an control: Convergent trends and controversial issues in current social-scientific research on human cognition and communication. *Inquiry,* **30**:77-79. 1987

14. Roschelle, J. The computer as a medium for conversations. In preparation.

15. Rosenfield, I. *The Invention of Memory: A new view of the brain.* New York: Basic Books, Inc. 1988

16. Schon, D..A. *Educating the Reflective Practitioner.* San Francisco: Jossey-Bass Publishers. 1987

17. Suchman, L.A. *Plans and Situated Actions: The Problem of Human-Machine Communication.* Cambridge: Cambridge Press. 1987

18. Tyler, S. *The Said and the Unsaid: Mind, Meaning, and Culture.* New York: Academic Press. 1978

19. Winograd, T. and Flores, F. *Understanding Computers and Cognition: A new foundation for design.* Norwood: Ablex. 1986

20. Zuboff, S. *In the Age of the Smart Machine: The future of work and power.* New York: Basic Books, Inc. 1988