

The Learning Process in the Epistemology of Medical Information

William J. Clancey

**Institute for Research on Learning
66 Willow Place
Menlo Park, CA 94025**

Methods of Information in Medicine **34**(1/2): 122-30. Special issue, IMIA WG6 Meeting,
Geneve, May 29-June 1, 1994, Jean-Raoul Sherrer (editor), 1995

The Learning Process in the Epistemology of Medical Information

William J. Clancey
Institute for Research on Learning
Palo Alto, California USA

Abstract

Progress in the past few decades in representing medical knowledge, along with the availability of low-cost, powerful workstation computers, has increased interest in encoding all medical records in electronic form. But despite the advantages of legibility, access, and automated performance reviews, computerization may also restrict what can be recorded and rigidify health care interactions with patients. Balancing the beneficial and negative effects requires an understanding of medical practice, especially the difference between human knowledge and today's computer programs. Human perceptual and conceptual capabilities bring an aspect of improvisation and reinterpretation to every human action, which in general computers today cannot replicate. Designers of medical record systems must consider the creative process by which descriptions are generated by people, as well as the collaborative process by which descriptions of past work are reinterpreted for guiding future decisions. This paper provides a framework for understanding these issues, illustrated by design opportunities.

Keywords: Human Learning, Medical Vocabulary, Objectivism, Holistic Analysis, Medical Practice

1 Knowledge and Information

Understanding human learning requires a consideration of the circumstances under which learning occurs and the means used to facilitate learning. In health care, learning circumstances range from the most mundane moments of inquiring about a particular patient's problems to overhauling the health care and insurance programs of a nation. The means to facilitate learning range from literature reviews and informal conversations to partially-automated outcomes analysis provided by an electronic medical record system.

But learning cannot proceed systematically within the medical profession if the epistemology of medical information is misunderstood. By the *epistemology* of medical information, we mean the origin and nature of knowledge responsible for creating and interpreting medical information. By medical *information*, we mean all facts, procedures, and theories concerning health care.

Crucially, by these definitions a statement in a book is not knowledge, but information that a knowledgeable person created and is able to interpret in a practical situation. We can describe what a person knows, for example, by a classification of diseases and symptoms, but such representations are not knowledge itself. Descriptions are themselves the *product* of knowledgeable action, not knowledge itself.

We stand today at the threshold of new computer tools, formalized vocabularies, and health care reform, with an antiquated view of knowledge. Developing and exploiting these requires understanding better the difference between representations (particularly computer representations) and human knowledge. Most people to some extent recognize the crisis in how we talk: What is the difference between a data, information, and knowledge? Today these terms are used more or less equivalently. We say that the statement, "the patient has a fever" is data in the medical record; we say it is information about the patient; and we say it is knowledge about the patient.

Through computerization in everyday life and the rise of computational theories in psychology, the difference between data and human knowledge has become confused [1,21,33]. Very simply, as Dewey put it, we are no longer distinguishing properly between the carpenter and his chisel [8]. By equating human knowledge with *descriptions* such as a medical record (data) or disease models (theory), we lose track of how models are created and used in practice, how computer tools can help people, and how design projects for developing tools should be conceived.

Fortunately, there is a major opportunity today in the design of useful computer tools for medical practice. For the first time, cognitive scientists and social scientists are collaborating in design projects, bringing together holistic and analytic perspectives [13]. In this position paper, I will describe theories of human learning that improve upon computationalist perspectives that equate knowledge with descriptions. I will briefly describe the learning process by which medical knowledge and theories evolve. Finally, I will describe how design practice might be improved by dropping the “either-or” contrast between basic and applied research.

2 The Holistic Approach

The holistic perspective is familiar to biologists, social scientists, and many medical providers [16,29,35]. The holistic perspective adopts a view of mechanism or causality that is process-oriented, viewing the activity of people as:

- interactive — knowledge is created within coordinated activity
- interpretative — knowledge enables giving meaning to descriptions
- collaborative — knowledge is oriented towards shared goals and methods
- developmental — knowledge is always changing in activity
- emergent — knowledge of the group transcends individual control
- participatory — knowledge develops through interactions between communities of practice

For purpose of this discussion about medical information, we will focus on the *development* of different aspects of knowledgeable behavior:

- standards
- theories
- technology
- communities of practice

A holistic approach to cognitive science seeks to understand the developmental dependencies between these aspects of human activity, in order to promote and manage organizational change. These themes are summarized well by Hirschorn [15]:

Systems theory, control engineering, utopian thinking, and Taylorist prescriptions all converge to limit the worker’s skill. In contrast, the principle of flexibility creates a conception of work in which the worker’s capacity to learn, to adapt, and to regulate the evolving controls becomes central to the machine system’s developmental potential. (p. 57-58)

In cybernetic systems machines and workers complement each other with respect to a typology of errors: machines control expected or ‘first-order’ errors, while workers control unanticipated or ‘second-order’ errors.... The worker moves from being the controlled element in the production process to operating the controls to controlling the controls.

The new technologies do not constrain social life and reduce everything to a formula. On the contrary, they demand that we develop *a culture of learning*, an

appreciation of *emergent phenomena*, an understanding of *tacit knowledge*, a feeling for *interpersonal processes*, and an appreciation of our *organizational design* choices. It is paradoxical but true that even as we are developing the most advanced, mathematical, and abstract technologies, we must depend increasingly on informal modes of learning, design, and communication. [emphasis added] (p. 169)

One way to show the relations between standards, theories, technology, etc. is in terms of the *scope* of learning and change involved in different activities (Fig. 1).

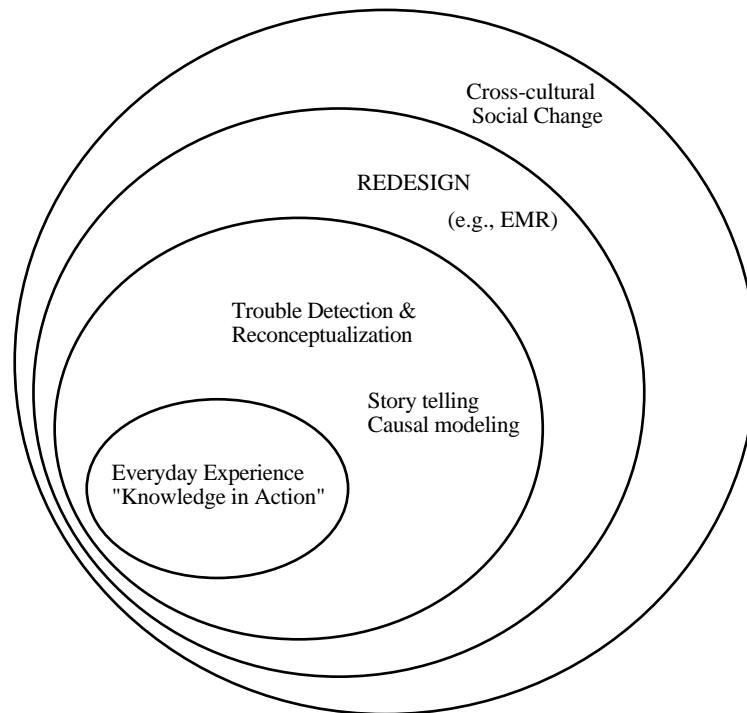


Fig. 1. Scope of learning and change in medicine.

The inner circle represents everyday learning, such as learning about a particular patient and his or her life and environment. Everyday learning involves perceiving symptoms, understanding the patient’s story, conceiving disease relationships—what Schön calls “knowledge in action” and “reflection in action” [27]. This is the realm of inarticulate, implicit, or tacit knowledge—not facts and theories accessed and matched subconsciously, but having a sense of timing, juggling multiple priorities, perceiving objects, and carrying out physically-coordinated procedures in a skilled way. By definition, this aspect of practice is routine.

During this activity, while interacting with one’s own patients or discussing patients of other providers, one may encounter trouble. Here a situation becomes known as problematic, requiring consideration:

- framing, giving a name to a situation.
- recounting the history of what you observed and did
- telling related stories from experience
- ordering events, claiming temporal relationships
- tentatively arranging a causal explanation of underlying processes
- reconceptualizing the meaning of observations, theories, and policies.

Expert systems cannot handle most problematic situations; by definition such situations are ill-defined, nonroutine, and unpredictable. Expert systems require a library of generally static models that fit broad circumstances. In contrast, tools for problematic situations would help

people *create new models*, perhaps by helping them relate past cases to the present situation or helping them communicate with knowledgeable colleagues [18].

More extraordinarily, an individual or group may reflect more broadly on experience, beyond the interactions with particular patients, to consider their practice of interacting with each other, and the tools they use. They now frame and tell stories more generally about their work practice. They may produce a process model, laying out typical events, what tends to go wrong, and how problems are generally handled. In this setting, on the job learning routinely occurs [28].

In this phase of reflection and recoordination, the individual may also craft a new tool, perhaps using a computer system or asking another individual to do something in a different way at a different time. The group may also be deliberately engaged in redesigning their practice, by an initiative of a union or management. A familiar example is the development of an electronic medical record system [20]. Such design may be local to a clinic, regional, or corporate-wide.

Within such a redesign project, multiple perspectives may be aired in considering design and resource tradeoffs, especially in how the project is to proceed. Outcomes and utility measurement advocates will drive the design differently from computer scientists, nurses, and patients. Redesign of work systems often involves hearing and responding to multiple perspectives. Redesign involves reconceptualization of the goals and values of the work, not just delivering a new technology or reorganizing a group according to an analysis of how yesterday's business can be made more efficient [23,26].

More broadly, the redesign effort may require or be driven by broader cross-cultural change. For example, the US Government's health care reform initiative deliberately seeks to recoordinate the perspectives and actions of employers, insurance companies, employees, and health care providers. More typically, these re coordinations occur when departments of a corporation or university begin to work together. For example, the non-regulated environment today forces engineers in telecommunications businesses to develop new products in collaboration with marketing specialists.

Reorganization typically begins by determining who will participate when and where during a redesign process [22]. Participation in a redesign team is the first consideration in recoordination of perspectives, habits, and values. When a new person joins a group from a distant region or competitor, a new cultural perspective is usually introduced that requires changing the group's practice. Cross-cultural change of this type can be unintentional or beneficently-subversive [17].

In summary, Fig. 1 views learning as occurring in different scopes, relative to change in practice, values, tools, models, and policies. Learning about a patient involves little change in any of these. Handling trouble often requires a reconceptualization of the meaning or importance or prioritization in relating models, tools, and policies [35]. For example, what should be done when a patient doesn't respond to a standard therapy? Redesign usually changes all of these: models, tools, organizations, and policies. Cross-cultural change occurs when language and ways of viewing problems changes, usually through participation of new members in the group, forced by stress that forces reorganization.

3 Contrast with Rationalist View of Learning

Figure 1 improves upon the "rationalist" view of learning we followed when developing the first medical expert systems [2]. Here we consider the tenets of the rationalist approach.

First, by the rationalist view *every patient encounter is a problem*. Every process of diagnosis and treatment is called "problem solving." But by definition, expert systems can only handle what is routine to the experts creating the system. In a group with varying degrees of

capability, the program will be of little use, except in automating what they don't need to personally handle. When Mycin was created the number of meningitis cases at Stanford were so few, the cost of developing and installing Mycin was impractical. On the other hand, pneumonia patients *did* present problems for these physicians, and there was no expertise to formalize.

Second, the rationalist view *equates descriptions with concepts*. By identifying “knowledge” exclusively with descriptive facts and theories, the process of understanding and meaning involved in creating descriptions and interpreting them is ignored [31]. All knowledgeable behavior is assumed to proceed from descriptions, suggesting that a theory and facts are necessary in order to act.

The contrary view is that understanding, sense of similarity and difference, and capability to coordinate ways of seeing with ways of acting precede and form the basis for description generation and interpretation [26]. This view is well-known in philosophy, psychology, anthropology, sociology, having been promoted over the past century by Dewey, Ryle, Bartlett, Mead, Bateson, and many others [1,11,19,25,31]. Today this approach is generally called “situated cognition.” The inadequacy of an epistemology which identifies knowledge with descriptions is argued on philosophical grounds, empirically through studies of everyday learning, and by new neurological models [9,24].

In practice, descriptions in the patient record are created and interpreted collaboratively. By equating internal mental processes with description manipulation, the special character of reading and writing is distorted. Especially, conversations for negotiating appropriate descriptions and meanings are reduced to the role of communicating (transmitting) individual ideas and work products. Rather than developing tools for facilitating conversations, the *rationalist approach works to eliminate conversations* and replace human reasoning by automatic deductive programs, such as expert systems [33].

Finally, the rationalist view *downplays and distorts the nature of cultural change*, viewing it as capturing and disseminating better theories and procedures. In practice, changing technology and human interactions requires changing many implicit conceptualizations and ways of coordinating action. Organizational change cannot be accomplished by only promulgating better policies, tools, or techniques. Change needs to be approached systematically and evolutionarily, building on the historical trends of the group [10].

In summary, the learning process in the medicine (as in any complex professional domain) involves a mixture of experience, theorizing, and design activities, such that learning is never just accessing and manipulating facts and theories, but also perceiving, conceiving, participating, and conversing in new ways [27].

As a simple application of this idea, a medical informatics specialist asked to design “supporting technologies for knowledge access” must consider not only search and retrieval of stored facts on-line. *Knowledge* should also be viewed as the capabilities of people to form new conceptions and see things in new ways; *access* should also be viewed as the opportunity to work with someone knowledgeable. Hence, knowledge access might mean having an opportunity to converse with a person.

Although this idea may at first seem mundane, it has tremendous implications for exploiting multimedia technology, including video teleconferencing, bulletin board services, and electronic mail. For example, to address the needs of non-English speaking members of a health care organization, a hospital might be better off investing in cellular phones to get quick access to on-call bilingual caregivers. Without this perspective, one would typically seek to replace people and human conversations by more costly and less effective automated mechanisms.

Another application of situated cognition is in understanding *the inherent conflict between local and global control of situations*. Local control refers to individual and small group

collaboration in patient care, as in a clinic. Global control refers to the policies, standards, and laws established by hospitals, professional associations, insurance companies, and the government. As I have said, one reason problems arise is because previously successful descriptions such as standard disease models and procedures are inappropriate for a given situation. Of special concern is the work required to interpret a policy so it makes sense in a given situation. People always lie between local practice and global theorization, framing situations in terms of past ways of talking, reconceiving the meaning of a theory or standard, and inventing new ways of coordinating resources and actions [7,12,26,29].

4 Inherent Dynamics between Practice and Theory

Figure 2 summarizes the inherent, dynamic relation between practice and theory. The rationalist view is that medical science is involved in a monotonic process of “increasing medical knowledge” by moving all experience and problems to the realm of heuristics, theories, and standards. By this view, all practice will eventually be dictated by scientifically proven methods, and can be automated by deductive machines.

Situated cognition claims that there is an inherent give and take between theory and practice. Good, efficient, appropriate, and effective practice requires theory. Theory does follow from practice, but not just in the sense of supplying data. In effect, knowledge involves knowing how to apply theories and this capability is not formalized in the theory itself. Good practice involves knowing what sources to examine, how to selectively relate guidelines to the resources at hand, and whom to call for further information. Good practice involves knowing when to violate a rule, when a new classification might be appropriate, when an alternate interpretation of a guideline might better fit the values of medicine.

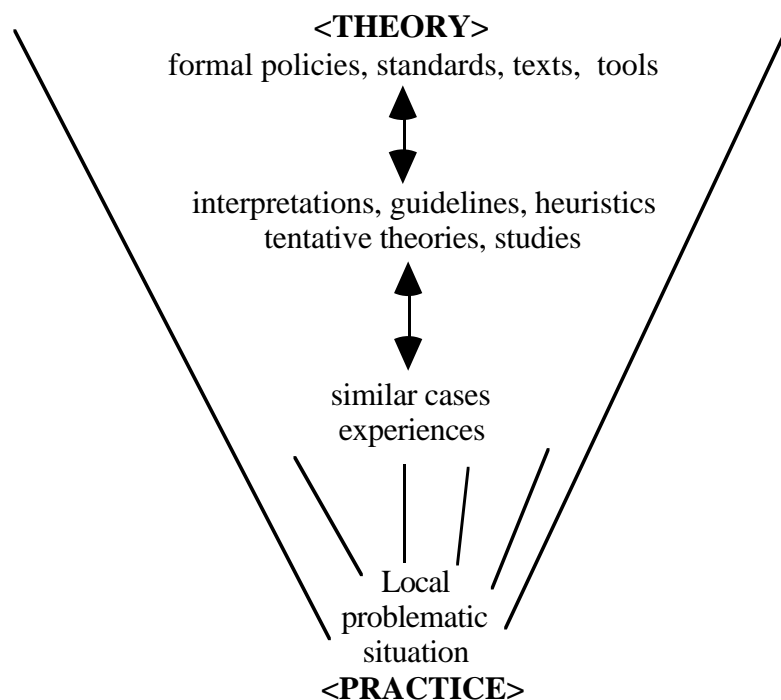


Fig. 2. Relation of descriptions to practice.

Crucially, people are speaking and conceiving every step along the way. People are *remembering similar cases and experiences*, as they encounter a difficult situation. People are referring to guidelines (which they solicit from the literature and by conversing with

colleagues and team members). People are *articulating new, tentative theories* to explain what is happening in a particular patient and what a new trend means. People are *suggesting new policies*, which enables them to balance cost and health care, and new tools for augmenting their capability to control equipment, schedule visits, manage therapy, and the like.

By this view, people are always deciding what is routine and what is non-routine, and how to apply policies and tools so they fit a given situation [34]. A simple example is the process of answering an expert systems' requests for patient data. The program may ask, "Has the patient responded to the erythromycin?" Interpreting the meaning of "responded" cannot be strictly separated from the *actions* alternative interpretations imply. Has there been sufficient time for the patient to respond? Is the dosage appropriate? A practitioner engaged in treating this patient cannot simply respond yes or no to the program. At the very least, the practitioner requires an intimate understanding of the program's capabilities: Will it take these other considerations into account?

The view that users of expert systems are merely data suppliers and drug administrators, which was tacitly assumed in the development of early medical systems, is wrong and also inappropriate. At the very worst, such designs assume that people can be removed from the loop in routine cases and then expected to jump back in and handle difficult cases that the program cannot manage [15,30,35].

5 Relation of Standard Vocabularies to Learning

The disagreement between the rationalist view and situated cognition is no more clear than in the attempt to formalize medical knowledge in a vocabulary, such as SNOMED. The rationalist view suggests that eventually we can develop a single language that describes all medical situations any human being will ever encounter on this planet. Presumably, the exceptions will be so rare as to be unimportant. By this view, the problems are more in the realm of storing and retrieving descriptions, rather than in developing the knowledge to create and interpret descriptions.

At a recent medical informatics workshop, a physician said, "The real question is whether a coding system like SNOMED is sufficient to represent medical concepts." I would suggest that the underlying question is actually, "Can human concepts ever be fully described, that is, replaced by words?" Situated cognition argues that human understanding will always exceed what is written down.

Consider, for example, the SNOMED term F92248 corresponding to "apprehensive." The circumstances under which the person is apprehensive or says he or she is apprehensive are not recorded. Indeed, we cannot fully record why the observer says that this patient is apprehensive. Neither the patient's experience of apprehension, nor the observer's experience of the patient, rests on a bedrock of words.

Attempts to develop standard medical vocabularies, as all efforts at theorization and standardization, are important and necessary. However, such efforts are misconceived and probably poorly managed if they assume that their task is to convert all knowledge into atomic terms. In most cases, the inability to record "all the details" may not matter. But we must be philosophically sophisticated enough to know that in practice it is impossible to record "all the details," and there will be practical implications that ensue.

As in the example of "knowledge access," the push to formalize medical descriptions into a standard vocabulary is blind to the nature of conceptualization and interpretation. In particular, besides formalizing medical processes in a vocabulary, we must augment such descriptions with indications of doubt, uncertainty, source of information, etc. Any cursory examination of physician's progress notes indicates copious use of question marks, the word "doubt," and many graphic symbols such as arrows and sketches of graphs. Yet typical

efforts to formalize medical records proceed as if “doubt” is not part of a physician’s vocabulary.

As another example, how would one code a patient report, “Hears dogs barking in the night”? One might select the SNOMED term corresponding to “auditory hallucinations.” But suppose we later find out that there are actually dogs in the neighborhood? Again, we might augment the vocabulary, but there is no end to this process. In the 1960s one might have needed to add, “Hears rock bands playing in the night.” And so on. Since social practices are open to change, there is *in practice* no vocabulary that can once and for all describe human activities, intentions, and experiences. Thus, we must study and understand the implication that medical records are only descriptions, which are always abstracted and open to interpretation.

Instead the development of formal vocabularies moves about in its own world of formalisms. Now that many groups have developed competing vocabularies, the problem has been transformed to *formally relating formal descriptions*. The original problems faced by a medical practitioner of *producing* good descriptions and *interpreting* past work are lost. To the vocabulary enthusiast, we need only produce finer-grained vocabularies and all will be complete and rational.

But there is no way to *completely* record subjective experiences and events. All coding requires abstracting from experience and interpreting terminology. The primary information and its context are necessarily left out. We must focus on the practice of how people use descriptions, perhaps augmenting them by informal representations (e.g., free text, photos, video, sound recordings) and ensuring that meaning can be reconstructed collaboratively.

But the problem is not just in the content of the recorded facts, per se, but the context in which facts are recorded. If the health care provider is to learn by reflecting on past work, the chronology and associations between observations and decisions must also be recorded. Such a model of *medical practice* is disjoint from a vocabulary of symptoms, tests, diseases, and procedures. In particular, an electronic medical record must provide ways for the physicians and nurses to record the linkages in their thinking and actions. How to do this is not obvious:

Despite the fact that the Weed [PROMIS] system was designed precisely to promote a synthesis of scientific and practical thinking, its automated form was not open and it did not leave enough room for the frequent tentative rearrangement of facts and hypotheses that are part of cognition in real world problem solving. [20; p. 164]

What MUMPS could not do, and the industry failed to do, was to provide an adequate platform that would address the interactive needs of top professionals properly: to be able to enter their own data and to navigate through a context rich data base with the express purpose of solving clinical problems.... These records do no more than document a series of specific transactions. They do not support spreadsheet-like interactions that deal with tentative inference and with volatile, revisable situations. (p. 178)

That is, what is required is much more than a vocabulary. Medical practice requires a *notational structure*, not just a set of words. Lincoln [20] says, “There are deeper issues about how individual observations should be labeled to meet different objectives” (p. 171). Indeed, inventing labels is a creative problem. Furthermore, there is no reason to restrict notation to one-dimensional labels. We may need to invent kinds of representations to organize observations and actions. Conceptual *graphs* are more likely to be of use [3].

In effect, the learning process requires better records to reflect on, during and after a medical encounter. But today’s tools are either too unstructured, like word processors or too structured, like forms. The difficulty in designing good tools partially stems from an inadequate appreciation of the cognitive processes involved, that is, the epistemology of medical information. Viewing learning as storage and retrieval of facts and theories misses

the point. Providers need *tools for constructing models of patients*, and for reinterpreting old models. Lincoln [12] continues:

Judgment will only return to its own if the logic behind it is captured and becomes subject to review, personal improvement, and teaching by better example. For clinical experience to be self-correcting, data must be approached in a manner similar to research: hypotheses must be recorded, together with actions taken in response to each, and the expected outcome predicted—all in sufficient detail (as a part of the patient chart)—to be the equivalent to a laboratory notebook. Only then will it be clear why a particular care policy succeeded or failed in a given circumstance, and whether the logic that was used was appropriate. (p. 173)

To this end, Lincoln has suggested the use of SGML (standard generalized markup language) to retain prose details of the computerized patient record, to allow open-ended classification and post-processing on a formal set of standard and locally-defined tags.

In summary, the medical record needs to document the etiology and progression of the patient’s disease story, as well as the circumstances of the diagnostic process itself. The idea of a medical vocabulary is useful, but limited, suggesting the pathologist’s view that medical experience and knowledge can be captured by labels attached to cases. The alternative view is that the electronic record is a *narrative*, not just a series of names, but a story, linked by dimensions of observation, location, treatment, etc.

Indeed, one need only look at the term “Computerized patient record” (CPR) to see how the emphasis has been placed on *data about the patient*, rather than data about the diagnosticians and the circumstances of the patient encounter. For example, one frequently finds that family doctors and specialists have limited or no access to each other’s activities; a physician in an outpatient clinic often wastes time getting routine information such as mammograms from the hospital. Providers in the hospital may say, “It’s negative, why do you need to see it?” But learning from past experience via records, requires including who, how, what, where, and why descriptions were created and interpreted. Members of a team are not merely exchanging medical descriptions (“it’s negative”) but cross-checking each other’s work and examining primary data from different perspectives. In order to interpret the reasoning processes of other providers, the links of time and place by which one experience leads to another needs to be retained.

6 Understanding Dependent Hierarchies

Learning at the broadest level (Fig. 1) is helped immensely by not framing alternatives in terms of binary oppositions: local vs. global, cost vs. quality, holistic vs. analytic, individual vs. social. Instead, one must view these as dependent conceptualizations, within a hierarchy [32].

Consider for example the arguments in medical informatics about how to relate learning processes to departmental boundaries. Practitioners usually begin by framing their experience in terms of dichotomies (Table 1).

Table 1

Common dichotomization of activities into formal and informal components.

Formal	Informal
departments	research projects
degrees	professions
training curricula	practice
research	service

Planning and problem solving are often structured by such binary oppositions: Should *research projects* be required to provide service to other *departments*? (Research projects and service are viewed as mutually exclusive.) Should medical informatics be a degree? (Roles in professional collaborations are contrasted with theoretical specializations, which are marked by degree areas.) Should training involve practice? (Learning and doing are viewed as separate activities.) How should indirect costs be allocated for multidisciplinary projects? (Organizational activities sustaining research are viewed as independent of research.)

Throughout our experience, we find conflicts between organizations and how we conceive of our activities and interests. Just as I indicated in relating practical knowledge to policies, we should not start by assuming that the solution is to “find the right organization and degree programs.” Practice will always require *ad hoc*, improvised collaborations and reconceptualizations across departmental boundaries.

This is not to say that we don’t make progress by formalizing our activities into new policies or organizations. On the contrary, it is important to recognize when a discipline (such as medical informatics) is so complex and involves such a recurrent suite of tools and collaborations, that it should be formalized as a department, a degree program, and a research arena in its own right.

One way to understand the dynamics of cross-cultural learning is to examine what is wrong about the “basic” versus “applied” opposition. Western culture traditionally tends to define disciplines along this spectrum: Classic science such as chemistry and physics are basic, while engineering, medicine, and social sciences are inherently applied. This distinction is useful, but like all binary oppositions it fails to acknowledge that one person’s ceiling is another person’s floor: Basic and applied research are not only relative along a spectrum of activity, they are *dependent* upon each other.

If we examine what falls under the rubric of “basic research,” we find development of modeling notations (as in mathematics), modeling tools (such as knowledge engineering [2]), and models of physical and biological phenomenon, apart from their use (such as organic chemistry and neurobiology). Basic research is the realm of tidy categories of “knowledge,” the realm of disciplines, formal methods, and languages. From this perspective, “service” or practical work is viewed as mere application, a distraction from the business of developing fundamental understanding. Collaboration in practical projects might be viewed as *validation* of the tools or a compromise to secure funding or help the overall organization in the face of competitive pressures.

On the contrary, the focus of “applied research” is not on theories in isolation but relating theories to some practical context. Multiple views are brought to bear to define problems and design new tools and organizations. Developing an electronic medical record clearly falls under this rubric. Here the emphasis is not on designing tools in isolation but on *designing processes* for solving problems. Indeed, from the perspective of engineering, business, and the social sciences, this work is *basic research on the nature of managing and promoting organizational change*. Collaborative processes are the methods, and they must be invented and designed on the spot, in context. In some respect, applied work is always a process of “redesign” or “reengineering” as existing tools, methods of participation, languages, and organizations must be reorganized into a new practice.

To this point, the opposition between basic and applied might be restated as being a difference between *designing methods that are formal tools* versus *designing methods that are practical processes*. To carry this analysis further, we must consider how even theoretical learning exists within a cross-cultural environment of dependent activities, goals, and values.

Wilden suggests that we use a hierarchical, “both-and” view rather than an “either-or” spectrum for understanding how an activity or process is related to its environment (Fig. 3).




Fig. 3. Either-or view, a binary opposition of activities.

Showing disciplines as *dependent hierarchies* produces the following pictures instead (Fig. 4). By the view of a dependent hierarchy, we can view the disciplines on levels, but they are *defined with respect to each other*. Broadly speaking, molecular biology provides theories to physiology, whose purpose is partially to inform the practice of medicine. Each level can be viewed this way.

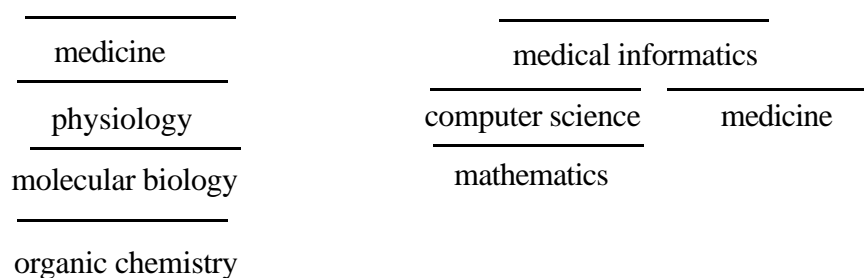


Fig. 4. Dependent hierarchies of disciplines.

An objective view of science, suggesting that the disciplines exist only to arrive at the truth of reality, assumes that the disciplines are not defined with respect to each other. An objective view suggests that the disciplines exist in isolation, defined by and validated by the subject matter [14,31]. Scientific descriptions are viewed as independent of human perceptions, purposes and values [19]. In the extreme, this view of science claims that the role of science is to produce the most basic, purified descriptions, in which an observer’s perspective and interests (“bias”) are fully extinguished.

In practice, sciences may appear more “basic” because they would not cease to exist if the contextual enterprises that interpret them didn’t exist. For example, organic chemistry would not be meaningless and undirected if molecular biology didn’t exist. But the identity of the organic chemistry discipline today—its methods, funding, training, and role—are all partially understood and practiced with respect to the disciplines that lie “above” it.

Not understanding the *practical* dependence of disciplines, scientists may view that something is wrong when they are obligated to seek practical projects. Rather than a nefarious attempt to extinguish “basic” work, this shift in broader cross-cultural demands may be viewed as recognizing the dependence of the disciplines on each other. The idea that computer science, for example, might proceed fully independently of services to other fields fails to recognize the source of new ideas, as well as the social fabric that makes funding for any discipline possible. Trying to wall off a discipline from the rest of the world is at best a temporary illusion, driven by the rationalist idea that knowledge itself can be packaged and preserved as isolated facts and theories.

The patterns mentioned above might be generalized further (Fig. 5), showing how the learning processes in “basic” and “applied” areas are dependent.




Fig. 5. Dependent hierarchies of theory-packaging vs. design activities.

By this view, projects, practice, and designing provide the environments that make disciplines, theories, and models (tools) meaningful. For example, the development of a standardized medical vocabulary is justified and given direction by knowledge and experiences of the nurses, physician assistants, and MDs who will use electronic medical records. A more formal activity always seeks to define itself as more fundamental and hence “real science,” for concrete descriptions always appear to be more rigorous and precise than the tacit, implicit coordinations of human knowledge and judgment. But ultimately, without social activities—either research projects or efforts to redesign practice—these disciplines, theories, and languages would not exist.

7 Conclusion

When we consider the learning processes of medicine, it is useful to adopt a holistic epistemology, that is an understanding of the social character of knowledge and activity by which information is created. A sound philosophical, psychological, and social background is required to disentangle the ways in which organizations and disciplines today define the learning process. Without this understanding, the management of the electronic medical record design process, as well as the design of the tools themselves, will fall to those who identify knowledge and information with descriptions. If learning is reduced to storage of descriptions (theories, models, data), technology will be inadequately exploited and glaring problems will remain.

Medical informatics research doesn’t often explicitly discuss issues of the nature of knowledge, though important assumptions always hover in the background. A participant at a recent workshop recently expounded the rationalist view by saying, “The controlled trial is the least biased, you’re trying to get to the truth,” not recognizing how handicapped such a perspective leaves him as a participant in the design of practical tools. Many researchers do not appear to have an adequate philosophical background for understanding the issues of truth and objectivity that they associate with medical practice.

Nevertheless, there is reason to believe that the majority of successful professionals, especially in the medical arena, appear to have an intuitive understanding that medical practice is inherently unformalizable and truth-constructing, and that it doesn’t fit the objectivist view of classic science. This tension between the way medical professionals talk and their understanding animates their conversations, as they attempt to reconcile the names in their theories and the formal policies of their organizations with the practical issues of training and health care reform.

To understand human learning, we must begin by acknowledging that practice and theory form a dependent hierarchy of languages, tools, models, and designs. Attempting to define the work of inventing new tools and new ways of collaborating in terms of university departments will place the interests of disciplines and projects in binary opposition, and fail to acknowledge that what is most problematic for us today is learning how to work together.

Acknowledgments

Funding for this work is provided in part by grants from Nynex Science & Technology, Inc., and Kaiser-Permanente. I have benefited from extensive discussions with my colleagues at the Institute for Research on Learning.

References

1. Bateson G. *Steps to an Ecology of Mind*. New York: Ballentine Books, 1972.
2. Buchanan BG, Shortliffe EH. *Rule-Based Expert System: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley, 1984.
3. Campbell KE, Das AK, Musen MA. A logical foundation for the representation of clinical data. Tech report *KSL-94-02*, Medical Computer Science, Stanford University, 1994.
4. Clancey WJ. Why today's computers don't learn the way people do. In: Flach PA, Meersman RA, eds. *Future Directions in Artificial Intelligence*. Amsterdam: Elsevier, 1991;53-62.
5. Clancey WJ. Guidon-Manage revisited: A socio-technical systems approach. *J of AI and Educ* 1993;4(1):5-34.
6. Clancey WJ. Situated action: A neuropsychological interpretation (Response to Vera and Simon). *Cog Sci* 1993;17(1):87-116.
7. Clancey WJ. Practice Cannot be Reduced to Theory: Knowledge, Representations, and Change in the Workplace. In: Bagnara S, Zuccermaglio C, Stucky S, eds. *Organizational Learning and Technological Change*. In press.
8. Dewey J. The criteria of experience (1938). In: McDermott JJ, ed. *The Philosophy of John Dewey*, Chicago: University of Chicago Press, 1981;511-23.
9. Edelman GM. *Bright Air, Brilliant Fire: On the Matter of the Mind*. New York: Basic Books, 1992.
10. Engeström Y. *Learning by Expanding*. Helsinki: Orienta-Konsultit Oy, 1987.
11. Gardner H. *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic Books, 1985.
12. Gasser L. Social conceptions of knowledge and action, *Artif Int* 1991;47(1-3):107-38.
13. Greenbaum J, Kyng M. *Design at Work: Cooperative Design of Computer Systems*. Hillsdale, NJ: Lawrence Erlbaum, 1991.
14. Gregory B. *Inventing Reality: Physics as Language*. New York: John Wiley & Sons, Inc., 1988.
15. Hirschorn L. *Beyond Mechanization: Work and Technology in the Postindustrial Age*. Cambridge, MA: The MIT Press, 1984.
16. Jordan B. *Birth in Four Cultures*. London: Eden Press, 1983.
17. Kling R. Cooperation, coordination and control in computer-supported work. *Comm of the ACM* 1991;34(12):83-88.
18. Kukla CD, Clemens EA, Morse RS, Cash D. Designing effective systems: A tool approach. In: Adler PS, Winograd TA, eds. *Usability: Turning Technologies into Tools*. New York: Oxford University Press, 1992;41-65.
19. Lakoff G. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago: University of Chicago Press, 1987.

20. Lincoln T, Essin, Ware. The Electronic Medical Record: A challenge for computer science to develop clinically and socially relevant computer systems to coordinate information for patient care and analysis. *The Information Society* 1993;9:157-88.
21. Lave J. *Cognition in Practice*. Cambridge: Cambridge University Press, 1988.
22. Lave J, Wenger E. *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press, 1991.
23. Nonaka I. The knowledge-creating company. *Harvard Business Review* 1991;November-December:96-104.
24. Reeke GN, Edelman GM. Real brains and artificial intelligence. *Daedalus* 1988;117(1):143-73.
25. Ryle G. *The Concept of Mind*. New York: Barnes & Noble, Inc., 1949.
26. Schön DA. Generative metaphor: A perspective on problem-setting in social policy. In: Ortony A ed. *Metaphor and Thought*, Cambridge: Cambridge University Press, 1979;254-83.
27. Schön DA. *Educating the Reflective Practitioner*. San Francisco: Jossey-Bass Publishers, 1987.
28. Scribner S, Sachs P. Knowledge acquisition at work. *Inst on Educ and the Econ IEE Brief* 1991;Number 2:1-4.
29. Suchman LA. *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge: Cambridge Press, 1987.
30. Theureau J, Filippi G, Gaillard I. Traffic control activities and design: Case studies. In: *Workshop on Work Activity Analysis in the Perspective of Organization and Design*, Paris, October, 1992.
31. Tyler S. *The Said and the Unsaid: Mind, Meaning, and Culture*. New York: Academic Press, 1978.
32. Wilden A. *The Rules are No Game: The Strategy of Communication*. New York: Routledge and Kegan Paul, 1987.
33. Winograd T, Flores F. *Understanding Computers and Cognition: A New Foundation for Design*. Norwood: Ablex, 1986.
34. Wynn E. Taking Practice Seriously. In J Greenbaum, M Kyng (eds), *Design at Work: Cooperative Design of Computer Systems*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1991;45-64.
35. Zuboff S. *In the Age of the Smart Machine: The future of work and power*. New York: Basic Books, 1988.

Address of the author:

William J. Clancey, Ph.D.,
 Institute for Research on Learning,
 2550 Hanover Street,
 Palo Alto, CA,
 USA