Intellectica, 2000/1, 30, pp. 115-149

Aaron V. CICOUREL*

What counts as data for modeling medical diagnostic reasoning and bureaucratic information processing in the workplace?[1]

A goal of knowledge-based or expert systems in the workplace is the development and implementation of software from laboratory to commercial or public or scientific applications. The above-stated goal requires an examination of the variable and adaptive behavior of personnel in daily work routines, including management support for training personnel and upgrading knowledge-based software. Field research (called "task-oriented ethnography") is needed to understand informal organizational activities that enhance and constrain decision making and bureaucratic information processing. For example, systematic observation and recording in the workplace of routine and special verbal, nonverbal and paralinguistic activities, documents produced by such activities, and the use of routine and special artifacts by personnel. Data fragments from two case studies (a pediatric outpatient clinic and an infectious disease division) within a university hospital are used to illustrate the kinds of training and discourse practices that occur in functional medical education and health care delivery settings. The latter activities are seldom observed on-line by modelers of medical diagnostic reasoning and medical bureaucratic information processing. We identify empirical issues necessary for using or modifying available computational architectures or creating new software.

* Department of Cognitive Science & Human Computer Interaction Lab
9500 Gilman Drive - Dept. 0515 - University of California, San Diego
La Jolla, CA 92093-0515

[1] A revised version of a lecture and seminar presented at a meeting in Paris sponsored by the *Association pour la Recherche Cognitive* and the *Fondation Maison des Sciences de l'Homme*, 11 and 12 March 1998. The general theme of the monthly seminar was: "Cognitive Technology in the Workspace: state of the art & new perspectives." I want to thank Charles Elkan for allowing me to make use of materials from his graduate course on knowledge-based systems (UCSD, Winter, 1998), and Jim Hollan for his helpful suggestions for reorganizing the manuscript.

Keywords: modeling, reasoning, medical diagnosis, bureaucratic information processing, expert systems, work routines, training, task-oriented ethnography; verbal, nonverbal and paralinguistic activities, functional medical education, health care delivery, computational architectures.

Quelles sont les données valables pour la modélisation du raisonnement en diagnostic médical et en traitement bureaucratique de l'information au travail ? L'un des buts des systèmes experts ou systèmes à base de connaissance est de développer et réaliser des logiciels utilisables dans des applications commerciales, publiques ou scientifiques. Ce but requiert que l'on examine le comportement adaptatif du personnel dans des situations de travail quotidiennes ; il requiert aussi le soutien de la direction pour la mise au jour des logiciels de bases de connaissance et pour le processus d'apprentissage. Une recherche de terrain — une "ethnographie orientée vers la tâche" — est nécessaire afin de comprendre les activités organisationnelles informelles qui facilitent et contraignent à la fois les processus de décision et du traitement de l'information bureaucratique. Il s'agit, par exemple, d'une observation systématique, et de l'enregistrement sur le lieu de travail des activités verbales, non-verbales et paralinguistiques particulières, des documents produits par ces activités, et de l'utilisation par le personnel d'artefacts de routine ou spécialement conçus. Des données partielles provenant de deux études de cas (une clinique pédiatrique en externe, et une division de maladies infectieuses) au sein d'un hôpital universitaire sont utilisées pour illustrer les modes d'apprentissage et les pratiques discursives qui se produisent dans la formation médicale fonctionnelle et le livraison des soins. Ce genre d'activité est rarement observé en temps réel par les modélisateurs du raisonnement en diagnostic médical et en traitement bureaucratique de l'information médicale. L'article identifie des questions empriques qui doivent être prises en compte dans l'utilisation ou la modification des architectures computationnelles existantes, ou dans la conception de nouveaux logiciels.

Mots-clés : modélisation, raisonnement, diagnostic médical, traitement bureaucratique de l'information, systèmes experts, routines, apprentissage, ethnographie orientée vers la tâche, activités verbales, non-verbales et paralinguistiques, formation médicale, livraison des soins, architectures computationnelles.

INTRODUCTION

The development and implementation of software from laboratory to commercial or public or scientific applications assumes that knowledge-based systems will function “appropriately” in practical settings. The term “appropriately” is elusive because we seldom can confirm the ecological validity of the software in terms of its survival over time by observing the daily activities of the “expert” and novice in her or his normal work setting. The present work proposes that the validity of knowledge-based systems can be improved by examining the behavior of personnel during their daily work routines, including the training of personnel. For example, the use of a data mining methodology in medical research makes use of historical databases to improve the way subsequent decisions are made about the prescription and successful use of pharmaceutical products, but such databases do not reflect physician-patient interaction, patients' complaints or symptoms about side-effects and routine decisions on how medical records are assembled from different, on-line workplace exchanges and written documentation.

A seldom used methodology in the design of knowledge-based systems (task-oriented ethnography) requires systematic observation and recording of routine and special workplace verbal, nonverbal and paralinguistic activities, examining documents produced as a consequence of such activities, and the common and special artifacts employed by different personnel. The method presupposes knowledge about bureaucratically organized activities.

Individual and collectively derived information is required to construct expert systems and sustain their functional viability and efficacy in a workplace. Ideally, the necessary information to construct a knowledge-based system should also reflect the socially organized settings in which such a system will function. Social organizations invariably are governed by their own formal and informal (often unstated) behavioral patterns, regularities, rules or policies and practices. The idea of software validity and “survival” becomes strained if we index such notions only by interviewing experts and using their metalevel accounts and official clinic records. An alternative is to obtain systematic information about what constitutes the experts' daily round of activities in the workplace.

The goal of this paper is to go beyond identifying procedures or “facts” by consulting experts, for example, or mining databases accumulated for various problem solving tasks, by asking: How do novices and experts function in their organizational, “natural” habitats?

Two empirical settings will be examined; telephone calls to a pediatric outpatient clinic and selected activities within an infectious disease division at a university teaching hospital.

Research on nonhuman primates, uses such terms as “power,” “rank,” and “brute force” to describe activities within a dynamic behavioral ecology. Among human primates, however, the complexity of “authority,” “power,” “rank,” and “obeying orders” has seldom lead to observations based on systematic samples of workplace behavior, and explicit procedures for obtaining, describing and coding information processing and communication in such settings. For example, tension always exists between formal and informal social or communication networks, and knowledge of this tension and how it is displayed routinely in work settings is essential for understanding the way social organizations survive, reproduce and change.

Although knowledge-based and related modeling systems depend on human information processing and communication, complex cognitive, organizational and linguistic mechanisms and processes are often treated as self-evident aspects of designing functionally efficient computational models. Artificial intelligence (AI) research or simulation of natural language processing and comprehension, for example, seldom includes direct observation and recording of task-oriented workplace activities.

KNOWLEDGE MANAGEMENT PROBLEMS IN MEDICAL SETTINGS

The cumulative medical records of particular patients that emerge from the physician's use of elicitation procedures during the medical interview and physical examination create a relentless series of summary statements about a patient's past, present and possible future problems. The records also include medication summaries, laboratory results and radiological information presumed to justify a physician's inferences about a differential diagnosis. The net result is the creation of an archive that can grow steadily as patients mature and remain in the same health care network within a limited geographical area. A serious empirical issue is the way patients engage health care delivery personnel in particular kinds of discourse in order for patients to gain access to a clinical setting and be diagnosed and treated. Instead of asking experts for retrospective accounts about what they do, we examine their activities in clinical settings, and the way the patient obtains an appointment and gains access to the clinic.

A major problem for all health care systems are the judgments required to code clinical practice (medical procedures and services) such that monetary information is the outcome. For example, discretionary assessments are needed to record the frequency and distribution of activities within different clinics, hospital wards, emergency rooms and private practices. In busy emergency rooms, hospital clinics and wards, and private practice settings, the tedious work of transforming health care delivery into monetary outputs can be overwhelming because of the volume of (sometimes ambiguous) information that must be processed.

The accumulation of medical information necessary for functioning health care delivery systems is a relentless process due to a database consisting of different forms of communication (face-to-face, telephone, written histories and laboratory or radiological reports sent by fax), meta level accounts by experts that resemble their discussion of patients with each other formally and informally. The advent of high speed computer hardware and increasingly sophisticated software has not altered significantly the storage of the information contained in a patient's chart (except for the results of particular tests, and the kinds of medications the patient receives currently). Obtaining an adequate database on prior and actual activities of physicians in their work settings remains empirically challenging.

The billing of costs for particular services often means that a printed form (see Figure 3) must be completed that transforms the office procedures (weighing the patient, taking their blood pressure, their temperature) and the physicians activities into (mostly) precoded categories so the billing office can further transform the activities into monetary consequences for the clinic, hospital, department or division and physician. The information in Figure 3, therefore, is supposed to represent the medical services and procedures given the patient. Hospital activities are summarized by detailed computer printout sheets that list the procedures, services, medications and artifacts for which patients are normally charged.

Health care personnel face a continually burgeoning explosion of bureaucratic, clinical, and basic research information, and face a constant struggle to incorporate new, modified, and existing knowledge into their declarative and nondeclarative memories. Physicians' micro knowledge management contributes an additional level of complexity to medical records (written progress notes and medical history). The physician's on-line judgments and abbreviated accounts can be seriously truncated because of constant time constraints.

Healthcare facilities must manage a huge amount of information on the cost of services provided to patients regardless of whether the facility is private or public. Linking the delivery of health care services to “appropriate” costs estimates associated with such services remains a daunting problem.

Calls to a pediatric outpatient clinic

The creators of knowledge-based or expert systems often model a particular aspect of a system that learns from experience (e.g., using historical data from hospital records (“data mining”) to provide advice on which new patients might best respond to particular treatments (Mitchell, 1997). The complex process of how patients initially seek a medical appointment, visit a clinic for a medical interview, diagnostic workup and treatment, submit to laboratory tests and/or radiological and/or ultra-sound procedures has not been modeled.

In the pages that follow, I present fragments of discourse from calls to a pediatric outpatient clinic. The initial process is instructive because medical personnel are often not involved; the “gatekeeping” process in the present case was direct because a person answers the telephone (rather than an automated telephone answering system that can be confusing and irritating to many patients because of a series of options that must be negotiated).

The activities of the pediatric outpatient clinics in which I worked and studied begin with non-medical receptionists who answered all calls. The receptionists have no training in health care delivery yet make decisions that can affect the diagnosis and treatment of patients.

If we were interested in creating an automated system for callers to a pediatric outpatient clinic, and we assumed that receptionists can be replaced by an intelligent KB system, how might we make use of the exchanges in Figures 1 and 2 in order to have a voice activated system that could function as the receptionists do in these Figures?

Figure 1 -
Calls to the receptionist - case 1

1	Re:	Pediatric primary care, can I help you?
2	Ca:	My child is sick.
3	Re:	Wait a minute (calls research over). Your child is sick?

Ca: Yes, he's only 18 days (old?). He woke up this morning sick.
Re: What's his temperature?
Ca: 99.4 or 99.6...He hasn't had a bowel movement since yesterday may[be] last night.
Re: I need your phone number.
Ca: (gives number)
Re: Baby's last name?
Ca: (gives name)
Re: First name?
Ca: (gives name)
Re: Birthdate?
Ca: (gives date)
Re: Brown card number?
Ca: (gives number)
Re: Your name?
Ca: (gives name)
Re: A nurse will call you back.
Ca: When?
Re: In about 45 minutes.
Ca: Okay.
Re: Goodbye.
Ca: Goodbye

The demographic information elicited in Figures 1 and 2 could be obtained by existing softward programs. The caller could be asked if they wish to make an appointment and given a menu of possible dates and times. If the caller indicates an immediate appointment is needed, then an option can be provided in order for someone to answer the call and make a judgment about a date and time, e.g., asking the caretaker to bring an infant or child into the clinic immediately.

Figure 2 -
Calls to the receptionist - Case 2

Re: Pediatric primary care, can I help you?
Ca: Mi nino es muy lloron. (My child cries a lot)
Re: Digame senora (tell me [more] madam)
Ca: Es muy lloron (?) ha tenido viruela (or varicela) (He cries a lot (?) He has had (smallpox [or] chicken pox)
Re: Como se llama el nino? (What is the child's name?)
Ca: (gives name)
Re: Fecha de nacimiento? (Date of birth?)
Ca: (gives date)
Re: Nombre suya? (Your name?)
Ca: (gives name)
Re: Desde cuando esta enfermo su nino? (How long has your child been sick?)
Ca: Desde sabado. Ha tenido diarrea, llorando siempre.

	(Since Saturday. He has had diarrhea, always crying.)
Re:	Fiebre? (Fever?)
Ca:	No creo, pero no tome' la temperature porque no tengo termometro. (I don't think so, but I didn't take his temperature because I don't have a thermometer.)
Re:	(Unable to follow more dialogue.)

In the paragraphs that follow, I describe several conditions that can be addressed by a hypothetical KB receptionist-appointment system we can call *Receptionist*. For example: Assume a caretaker calls a clinic to ask for an appointment.

— The voice-activated system first asks for the name of the patient and a medical record number which will indicate the kind of care they are entitled to receive.

— *Receptionist* would require the activation (with the help of a search engine) of the patient's past clinic appointment records in order to examine its contents and accumulate (or update) demographic information.

— If there has been one or more prior visits to the clinic and no reported problems (e.g., missed appointments or being late frequently), *Receptionist* can ask the caller to specify a desired date (month, day, and time) for the appointment. The last part of the message can ask if this is an emergency; the caller can be asked to enter another number for immediate attention.

— *Receptionist* should be able to retrieve additional information. For example, the child's past innoculations to see if they are up-to-date, to see if a patient has missed one or more appointments or has been consistently late. The system also can remind the caller of missed appointments and late arrivals.

— A new medical record file can be created if this is a first-time caller. The caller is asked if they possess a particular kind of plastic card indicating eligibility for particular services. If the caller is unable to satisfy the request for information about eligibility, *Receptionist* can ask the caller to press another number in order to speak with someone.

An automated system can satisfy routine appointments despite the necessity of resolving the not so simple task of offering a menu of available appointment openings and achieving a good match between the caller's and clinic's activities. Medical appointment times always favor the clinic or physician's work conditions. Patients might request the earliest possible visit and be willing to wait in the reception room to be “squeezed in” between other patients. If the patient is not willing to wait, assuming “squeezing” is an option, they will be told to go to Urgent Care or an Emergency Room.

*Receptionist***,** therefore, should possess some knowledge about how to interpret aspects of the patient's health history in the database if it is to assess a caller's request for immediate care. For example, “mining” a child's past medical appointments for a history of particular conditions (e.g., ear infections, sore throats) that could generate a question about the patient's ears or throat.

Turning to actual receptionists in Figure 1, there may be minimal information overload when the receptionist updates her computer records on the patient, but the next patient may have to wait longer. The possibility of overload begins as the patient is processed through the clinic, first by the nurse or licensed vocational nurse, then, in a training hospital, by the intern or resident, sometimes by an attending physician supervising the interns and residents, and finally by the nurse or Licensed Vocational Nurse (LVN) when treatment plans outlined by the physician are discussed. We can examine actual calls to identify routine and unusual events in calls to the clinic in order to modify of our hypothetical program called *Receptionist*.

The opening lines of Figure 1 are typical; the receptionist identifies the clinic as “pediatric primary care” (line 1), and asks how the caller may be helped, but does not identify the hospital. The caller's reply (line 2, “My child is sick”) is an expected response, and the receptionist calls the researcher over to the telephone before repeating the caller's response as a question (“your child is sick?”) that appears to invite elaboration. The caller (line 4) appears to understand the invitational nature of the interrogative by the receptionist and begins to elaborate (“Yes, he's only 18 days (old?)..He woke up this morning sick.” The receptionist (line 5) now begins to ask about the baby's health (“what's his temperature”). The caretaker provides a response that includes information (and notice the choice of words) about the baby's bowel movements, observing (lines 6-7) that “He hasn't had a movement since yesterday, may[be] last night.” The remainder of the exchange is fairly perfunctory (obtaining a telephone number and demographic and bureaucratic information before telling the patient that a nurse will call back in about 45 minutes.

The second call (Figure 2) begins with the same opening line by the receptionist but the caretaker's response (line 2) is in Spanish and states that “My child cries a lot.” We cannot tell if the caller understood the receptionist's English or simply presumed she could speak in Spanish with anyone at the pediatric outpatient clinic because of past experience and taking for granted that the number is that of a medical facility. The receptionist's response (line 3) in Spanish is also an interrogative that appears to request additional information about the child's physical symptoms.

In the many cases I have observed in the same clinic, the receptionists consistently ask for a few physical symptoms before proceeding to obtain demographic and bureaucratic information. In the present case, however, the receptionist does not ask for a medical card number, but asks (lines 12-13) how long the child has been ill and the caretaker (line14-15) supplies additional physical symptoms (“...diarrhea, always crying”). The symptoms reported by the caretaker were perhaps construed as more than a typical request for an appointment, that is, suggesting that immediate attention was needed, especially since the caretaker reported (lines 17-20) that she could not tell the receptionist if the child had a fever because she did not have a thermometer. The receptionist decided to have the caretaker come to the clinic as soon as possible rather than have a nurse or LVN call back. The receptionist's judgment about the call is not unusual; if a patient's caretaker can appear “convincing,” they may be told to come to the clinic immediately. The receptionist's decision may be influenced by knowledge of the clinic's patient load for the day, but if she perceives a “problem,” she may believe one or more patients might be “squeezed in” on that particular day.

Our hypothetical KB system *Receptionist* would require more sophisticated software in order to assess a patient's reported physical symptoms. This might be done by supplying the caller with a few symptoms considered to be indicators of acute illness. For example, asking if the patient has a fever and if it is above, say, 101 Fahrenheit, asking to report any “restlessness,” “frequent sleeping,” finding that the infant or child has had six or less wet diapers in a 24 hour period, a rash over different parts of the body, “sore throat,” ear “hurting,” “coughing a lot,” and the like. Finally, it is important to note that in some medical practices, the nurse practitioner or physician may give advice over the telephone and refer the patient to a pharmacy for a prescription instead of having them visit the clinic or private office of a medical practice.

The latter possibility often occurs when relatively benign, known symptoms appear to be widespread in a community.

Information based on clinical and bureaucratic judgments begins to multiply when the Pediatric Outpatient Clinic's printed form (Figure 3) is filled out. Note that the form also contains basic demographic information about patients and their history of visits to the clinic, as well as a truncated (annotated) account of the activities that occurred during the clinic visit. The demographic information parallels what is stored in the clinic's network while information pertaining to clinical care that must be interpreted in economic terms falls within the purview of the billing office of the hospital. The patient's larger medical record is seldom in the archive of a central computer server, and consists of typed and handwritten documents, and laboratory reports that are stored in a medical records division.

The medical form serves the important function of summarizing which physician actually administered health care (the “provider”), the supervising physician or faculty member, and the kinds of medical problems observed or suspected (four spaces are allowed). The form does not indicate the amount of money that is to be paid to the physician's department for the services rendered, but each line of the three columns of the form (where relevant) contains code numbers that enables the billing office to transform each line into a monetary figure. Notice that under “visits,” the terms “Brief, Limited, Intermediate, Extended, Comprehensive” refer to the time spent with the patient. In the clinic I observed and worked, the line with “Intermediate” was the most frequently marked.

The number of possible procedures, medications, immunizations, treatments and supplies is large and creates considerable bureaucratic problems vis-a-vis charging the patient or another source such as an insurance company, medicare or medi-caid. The health care provider may send bills to the patient, the health care plan, and medicaid and medicare (if applicable). The problem is especially exacerbated by the volume of patients seen. In the case of patients who are admitted to a hospital, the accounting procedures for charging the patients or others is even more daunting because of the kinds of simple and complex services, procedures and medications can include hundreds of entries of a printout sent to a third party medical plan carrier and/or the patient.

The medical form and others like it appear straightforward. Yet they mask the judgmental decision making uncertainties that remain "invisible" in what otherwise can give the appearance of effortless, routinized practices, when medical personnel fill out such forms. To my knowledge, such practices have not been the object of systematic study. In an operating room or an intensive care unit or ER, the number of procedures, personnel, equipment, medications and other on-line activities would require considerable systematic observation in order to describe them accurately. Filling out medical forms, therefore, occurs after the fact and often includes many discretionary judgments. KB systems can search for patterns in the after-the-fact outcomes of discretionary judgments but the more challenging problems are ignored; the decisions that make up actual, on-line judgments in the clinic, the hospital ward, surgery, intensive care and the Emergency Room.

A troubling issue is the following: Why haven't medical expert systems been implemented successfully despite their presence for over three decades?

1) Whereas data entry and retrieval have been improved, their use varies across settings, and an overload problem remains because no systematic method exists for rapidly scanning and digesting summary statements in medical records. It is difficult to assess the correspondence between

a) how medical personnel initially observe and elicit (frame questions), and

b) how they subsequently summarize and record their findings using technical terms (and everyday language, including metonyms and metaphors).

2) Nonmedical personnel who engage in data entry and/or data retrieval are seldom aware of theories of categorization and how categories shape and constrain what will be seen and interpreted as "information." In many hospitals and clinics, in addition to cryptic notes, there is an increasing reliance on word-processing systems when transcribing dictated medical histories by physicians. The word-processing system, however, remains a difficult to access resource for

the physician (except for information on prior tests and current medications) because it can require considerable scanning that may not be faster than thumbing through a thick sheaf of papers containing familiar documents.

"Progress Notes" add (often small) increments of supplemental information to medical records that are somewhat cryptic and informal. Dictated medical summaries of a medical history and physical examination, however, transform a "selective" memory into more formal medical categories.

An important issue is the correspondence between the medical interview and the physicians' summary remarks. Readers of medical records seldom view such documentation as self-evident, objective, declarative sources of information (Cicourel, 1974; 1982; 1990; 1992). Physicians' nondeclarative (tacit) and declarative (consciously available) expertise enable them to go beyond the information given in the medical records. Physicians, nevertheless, often maintain a kind of dubious perspective about the accuracy or meaning of clinical descriptions in medical records, and are inclined to interview a patient to obtain more details about cryptic clinical descriptions already contained in the medical record in order to confirm their own hypotheses about the relevance of the patient's past medical history.

An important consequence of the physician's concern with the medical record is the trade-off between the record and how good of an "historian" is the patient. The physician's micromanagement of the information that is exchanged, filtered and documented cannot be identified in the official records. The physician's original on-line judgments that subsequently become represented in the medical history (written or dictated) have now transformed the discourse of the initial medical interview into the appearance of compressed time frames consisting of declarative, objective statements.

Physicians invariably are under considerable time pressure to complete a medical interview and physical examination quickly while trying to compile an "appropriate" account that colleagues can use at a later date and that is legally defensible. In addition, when drugs or procedures are used (including complicated surgical ones that require hospital care and a post-hospital nursing home and/or physical rehabilitation center), insurance companies and national health programs exert constraints by demanding justification vis-a-vis the appropriateness of prescribing a drug or treatment or diagnostic procedure and the particular time periods in which patients will be allowed to remain within hospitals, a nursing home and/or rehabilitation center.

Fragments from an exchange in a microbiology laboratory

The fragments of discourse from a microbiology laboratory and infectious disease clinic presented below briefly illustrate aspects of medical school training about infectious diseases, the kinds of explicit knowledge that is used, and how particular factual or conceptual elements can be linked to an exchange between a novice and a patient. The “snapshots” about infectious diseases indicate the kinds of expertise a modeler could use to create a knowledge-based system for novices and nonspecialists.

An important organizational element in the present case is the infectious disease clinic's continual, direct link to the microbiology laboratory (unlike some hospitals and clinics). The infectious disease attendings and the Director of the Microbiology Laboratory belong to the Division of Infectious Diseases, consult regularly with each other and the technical staff about particular patients, and organize and teach two required courses to medical students.

Figure 4 -
Discourse during microbiology lab rounds

PA:	(?) (Low voice level) Is this the same one (we?) (ya?) did
	yesterday?
IDA:	No. This is the eye lady.
PA:	(?)
IDA:	Cellulitis
PA:	Oh.
IDA:	with group A strep..in shock
PA:	In shock. (slight rise in voice level) How about that.
IDA:	I(t) was gonna be more interesting /if she didn't
MR:	/I'm (?)
IDA:	have bacteremia but (laughing and voice level increasing)
	now she's had /bacteremia so
MR:	/there's a little, there's a little (voice level increases)
	problem with that I'll, will go into more as far
IDA:	Yeah.

(exchange continues)

/ = overlapping speech

Figure 4 consists of a fragment of a speech event among a group of physicians about a patient during daily microbiology lab rounds. One of the participants, an intern (first year resident in medicine, MR) had previously examined the patient's medical records and then interviewed the patient. I accompanied him throughout the day. The exchange in Figure 4 is a routine discussion about a patient thought initially to have an unusual clinical condition. The exchange begins with what appears to be casual conversation about a patient but refers to a serious medical condition. I need to say something about the speakers.

The reader would have great difficultly recognizing that the person with the least knowledge (MR) did the most talking, and was reviewing some notes he had taken from the patient's medical chart and his memory of his interview with her earlier the same morning. The exchange in Figure 4, therefore, included two attendings (experts) and one novice. The experts saw each other on a daily basis, speak informally, and their language is sprinkled with a few technical terms that presuppose considerable background knowledge and clinical experience. The resident is trying to simulate this knowledge.

The exchange in Figure 4, however, does not provide us with evidence about the organizational status of the discussants, nor their expertise. Prior observation, participation, and interaction within the microbiology laboratory, and observing residents and medical students in hospital wards and outpatient clinics enabled the author to contextualize the exchange in Figure 4. I have discussed the case in an earlier paper (Cicourel, 1987) and here will only indicate a few highlights.

The exchange between two experts and a novice in Figure 4, lines 1-3, illustrates the kind of metaphorical language use that occurs routinely in clinical settings:

PA: "Is this the same one (we?) (ya?) did yesterday?"
IDA: "No. This is the eye lady".

The first speaker, PA, is the pathology attending and director of the lab; the second speaker, IDA, was the current infectious disease attending on call. Further information by these two experts is provided below. The apparently casual remarks by PA and IDA are not what we might expect experts to provide a computional modeler about the activities of the microbiology lab.

In line 5, Figure 4, the technical term "*cellulitis*" is perhaps our first clue to the medical nature of the discourse when uttered by the IDA. The PA's brief, apparently perfunctory response (line 6, "*Oh*") is followed by IDA's additional technical observation (line 7) that the patient had "*...Group A Strep..in shock*." In line 8, PA seems to express some surprise ("*In shock. How about that*"). The IDA (lines 11-12) continues her remarks with a mixture of informality and the use of a technical term ("*bacteremia*" or bacteria in circulating blood).

The novice (lines 13-14), MR, injects observations he gleaned from the patient's medical records earlier that morning. A kind of paradox occurs; we have a novice using formal talk rather than his mentors. The exchange, however, illustrates the claim made earlier; experts who are part of a daily workplace and social networks employ metaphors during their informal exchanges. They are not likely to use such informal talk when describing their work to outsiders who want to model their activities.

AN INTERVIEW WITH THE CHAIR OF A MEDICAL SCHOOL COURSE ON INFECTIOUS DISEASES

Figure 5 contains part of an exchange between the Director of the Microbiology Laboratory (PA, in Figure 4) and the author. I had asked the PA to describe what medical students should have learned before beginning their internship or first year of residency. The fragment of his response (Figure 5) illustrates a few details about infections that constitute some of the background knowledge the reader would have to know to understand Figure 4. The Director's remarks begin with a long preamble (lines 1-15) and typify the kinds of remarks we would expect to receive from an expert if asked to describe aspects of the microbiology of infectious diseases in order that we might construct a knowledged-based or expert system. The reference (Figure 5) to such technical terms as "streptococci" (line 16), "beta-hemolytic ones" (lines 18-19), "group A strep" (line 19) can be linked to "bacterial pharyngitis" (the folk term is "sore throat") used in line 20. The terms are relevant for understanding the medical condition of the patient alluded to in Figure 4. Computational modelers expect experts to use such terms.

Figure 5 -
Microbiology laboratory director on what medical students should retain from microbiology course

1 Dr: Well the kind of things that, that they, that we think they should
2 carry away from the microbiology course, involve, really what I
3 try to emphasize to the students, is the key information they
4 should get out of the course. And I can summarize that real
5 quickly. And if they all learned only these things, they'd carry
6 away a helluva lot more from the course.
7 One is, what are the important bacteria, and, in terms of human

disease, and what are their unique microbiologic characterisitics. You know, I don't give a damn if any them ever know which bacteria ferments sorbitol. Alright? What I'd like them to know is that pseudomonas is oxidase positive, and is a non-fermenter. Okay? And that uh the enterobaceriaceae are oxidase negative, though grow either aerobically or anaerobically. Okay? And that are fermentative organisms. You know, that kind of things. The streptococci can be differentiated into the hemolytic ones and the non-hemolytic ones and the green ones. Those that are alpha-hemolytic make green colonies. And among the beta-hemolytic ones we have the group A strep, which is really the only accepted cause of bacterial pharyngitis. Okay? Now those are unique microbiologic characteristics, so it's real, simple, little items, you know?

(exchange continues)

The opening lines of Figure 5, therefore, reflect the kinds of information that someone seeking to create an expert system in infectious disease would hope to elicit from an expert. The earlier reference to “cellulitis” (Figure 4) is associated with “group A strep” (line 19 in Figure 5) but the information in Figure 5 was described prior to and independently of the exchange in Figure 4.

A LECTURE BY AN INFECTIOUS DISEASE SPECIALIST TO INTERNS

The lecture and written notes (Figure 6) by the infectious disease attending (IDA of Figure 4) to first year house staff (interns) are consistent with the information presented thus far on the patient of Figure 4. Two key passages in Figure 6 (lines 1-7 and 12-15) should be cited because they are congruent with the discourse in Figures 4 and 5:

1) In Figure 6, lines 1-7, there are references to such key terms as *“Septicemia - microbial agents in the bloodstream,” “Gram-positive bacteremia,” “hemolytic streptococcus or* S. aureus*,” “Gram-negative bacteremia,” and “abscess.”*

2) *“One is led to the diagnosis of bloodstream infection by a sudden change in clinical state. It is difficult to guess the class of organisms without taking the entire clinical setting and recent microbiological history into account”* (lines 12-15)

The patient discussed in Figure 4 had gram-positive cocci (bacteria) and an abscess due to an infection, and experienced “*a sudden change in clinical state*” (very low blood pressure — noted previously by two attending physicians).

The pathologist's remarks reflect prior training and clinical experience that experts acquire and that the builder of a knowledge-based system must identify (and obviously understand in some minimal sense) in order to create an explicit knowledge base.

Notice that the elements of knowledge described in Figures 5 and 6 are implemented in actual clinical practice (Figure 4) in ways that are not self-evident. The remarks by the pathologist in Figure 5 and by the infectious disease specialist in Figure 6 reflect aspects of textbook knowledge and prior clinical experiences. Medical students' try to incorporate such information into their notes from class and laboratory experiences when they study microorganisms noted in Figures 5 and 6. Their acquisition this knowledge base, however, does not automatically translate into clinical practice. More general background knowledge is accumulated by exposure to diverse, novice-based, clinical experiences in the third and fourth year of medical school, and the first year of residency. Additional training may occur if a four or six week rotation in infectious diseases is pursued during the fourth year of medical school or during the internship. Residents pursuing a specialty in internal medicine or pediatric infectious diseases will be exposed to additional information about this sub-speciality.

Figure 6 -
Written notes of infectious disease attending for new interns

Septicemia - microbal agents in the bloodstream

A. Gram-positive bacteremia may occur in "normal" adults with pneumonia (30%), urinary tract infection (enterococcus), skin or soft tissue infection (hemolytic streptococcus or S. aureus). Gram-negative bacteremia usually arises in patients who have some defect in host defense: (1) obstruction (gall bladder, urinary tract, abscess, bowel, lung), 2) denuded surface, 3) absent or defective phagocytes, or 4) defective antibody. Fungemia, especially candidemia, does not occur except in patients severely compromised by disease (uncontrolled diabetes, malingnancy with immunosuppression and multiple antibiotics).

B. One is led to the diagnosis of bloodstream infection by a sudden change in clinical state. It is difficult to guess the class of organisms without taking the entire clinical setting and recent microbiological history into account. All sorts of organisms can cause hypotension, but in the absence of hypoxia (pneumonia) or heart failure, gram-negative bacteremia with endotoxemia is by far the most common cause of sustained hypotension....

.

.

D. Other causes of infectious shock

.
.
33 Strepococcal shock (Group A strep), often in the absence of
34 bacteremia, associated with soft tissue infections (usually severe)
(written notes continue)

The MR's remarks suggest he possesses the appropriate background knowledge (enhanced on this occasion by having read the patient's chart and having spoken to her earlier the same morning), but his background did not include enough clinical experience in infectious diseases to diagnose and treat the patient in question.

The clinical case discussed in Figure 4 closely resembles the remarks in lines 33-34 of Figure 6, namely, *"D. Other causes of infectious shock (group A strep often in the absence of bacteremia, associated with soft tissue infections (usually severe))..."*

The patient in question had gram-positive cocci, group A strep, a soft tissue infection around the eye, bacteremia or bacteria in circulating blood, but apparently had not been in shock, as noted by IDA in Figure 4. No one appears to have confirmed the absence of shock, and the presence of this clinical condition without bacteremia would have made the case particularly interesting for the experts.

In principle, software is available for representing the technical terms employed by the experts and novices. The remarks by PA in Figure 6, therefore, provide elements that can be linked formally to algorithms, and the latter can also reflect the way a medical history and physical examination summarize the medical personnel's exchange with a patient. In the case of the physician's oral and written medical history, however, a large amount of knowledge remains unobservable to the reader of a paper on the subject and equally hidden from the user of the KB or expert system.

As noted earlier, one concern is what we assume the user of the expert system must know in order to benefit from a model? The computational modeler is expected either to make explicit the kinds of unstated knowledge noted by the PA and IDA in the system being modeled, or to assume the user of the system is familiar with basic science and clinical concepts partially revealed in Figures 5 and 6. The model can perform as both a "teacher" as well as a diagnostic tool, permitting the user to pose queries of the system (Clancy, 1983).

Perhaps the most difficult task facing the modeler, however, is the variability of the novice or user's skill in eliciting appropriate information from the patient such that he or she enters relevant data in the software employed.

Aspects of the background knowledge on the patient known to the novice before interviewing the patient

The content of the questions used in MR's interview with the patient can be traced to fragments (shown in Figures 7 and 8) from two medical histories and physical examinations by two physicians (after an inital visit to the Emergency Room) after the patient was admitted to the hospital ward. The information about the patient discussed in Figure 4 and available in the two medical histories, however, do not reveal the kinds of elicitation procedures employed by the two physicians. As noted earlier, the MR and the researcher reviewed the patient's rather extensive medical records before the MR conducted an interview with the patient at bedside. Hence the fragments from the first physician's medical summary (and the subsequent physician after admission to the medical ward) that follow are presumed to have contributed to the MR's clinical knowledge about the patient.

Figure 7 -
Fragment of medical history by first infectious disease attending

5-days prior to admission she was bitten by a spider at night. She did not see the spider, but presumed that it was a spider bite. The following day she called the Poison Control Center, and was told to put hot packs on the lesion. The eye did not improve and she began to have increased swelling over the past 2-days, she had fever to 101 F, the days following the spider bite without chills or sweats. Her vision is unchanged. She has noted pain over the area of the eye extending back into the head. Denies alteration in eye movements. She has a history of insulin dependent diabetes mellitus, and began to note nausea and vomiting following since a spider bite. She has had minimal oral intake, and has held insulin intake since then. She continued to have several beers over this period of time. Patient was seen in the E.R. on [date], and was noted to have systolic blood pressure of 80. In the E.R., she was given IV fluids, with increase to BP to 100, and was subsequently admitted to the floor, had an I & D of this periorbital swelling. On the floor, she was noted to have systolic blood pressures in the 70-80's, thus was transferred to the MICU for further treatment.

(physical examination history follows)

The remarks in the first physician's medical history combine everyday language such as "*increased swelling*" with a few technical terms such as "*I and D*" (like extracting and examining pus oozing from the area around the eye). The ER did not have the patient's medical records available and was alarmed by **"**a systolic blood pressure of 80" which climbed to 100 after she was given fluids. After admission to the medical ward, **"**she was noted to have systolic blood pressure in the 70-80's" which also alarmed the ward personnel because it could signify the onset of shock. Hence the patient was transferred to the intensive care unit. The information in the first physician's medical history presupposes the kinds of information presented earlier from the interview with the Director of the Microbiology Lab and the written lecture notes by the infectious disease attending. A similar knowledge base is presumed by the physician who conducted a second medical history and physical examination on the ward when noting the following:

Figure 8 -
Fragment of medical history by second attending

Mrs. Price was doing fairly well until she felt something crawling on her forehead five days ago and a bite. She did not see a spider, but assumes that this bit her. Apparently she was bitten several times on the left lid. She treated this with local heat, but she did not improve. Her eye swelled shut approximately two days ago. She came to the Emergency Room having had no improvement, as well as a fever to 101 on the day of the bite, and chills and sweats since. She has not noted any problems with vision. On the morning after the spider bite she began to have nausea and vomiting, and did not take her insulin as a result. She says she has had little to eat or drink, but tells one historian that her last drink was one week ago. She had a watery bowel movement yesterday. She says that her glucose by urine dip-stick is usually negative but was positive during the past week. An I & D of the abscess of her lid was performed in the Emergency Room with a substantial amount of pus showing gram-positive cocci. She was treated with Ancef and Gentamicin, but was noted on the Floor to be hypotensive without significant orthostatic changes.

(medical history continues) [physical examination revealed]

There is some whitish exudate in the pharynx.

(physical examination history continues)

The patient's account about being bitten by a spider proved to be a mystery for all of the medical personnel, yet no other explanation was available. There is no reference to calling the Poison Center but the terms **"**local heat" for the patient's initial self-treatment were used. The initial medical history referred to "*The left eyelid is erythematous with* *oozing purulent yellow-green discharge*" and this observation is paraphrased in the second medical history but also notes that "*Her eye swelled shut approximately two days ago*." The second history also contains everyday language but includes more technical terms. Notice that line 18 of Figure 8 refers to "*some whitish exudate in the pharynx*." There was some speculation as to whether the infection was due to bacteria in the throat ("*some whitish exudate in the pharynx*") that may have moved up to the eye, or if a spider bite was the cause of the infection.

The information in the first and second medical histories represent highly truncated accounts that do not reveal the kinds of language used by patient and physician that might permit us to make inferences about the reasoning used by the physician. The interview conducted by the MR, after reading the patient's medical records, appears to reflect elements of the medical records and appears to have become the motivation and knowledge base for MR's attempt to interpret aspects of infectious diseases to the patient.

The general point is that considerable agreement exists between the medical summaries and the clinical evidence cited. The evidence is also in close correspondence with the kinds of expert knowledge found in Figures 5 and 6. The convergence of information presumably should lend itself to modeling.

MR's interview of patient

Having read the patient's medical records a few minutes earlier, MR proceeded to her hospital room where the exchange represented in Figure 9 is (partially) shown. The opening salutation (lines 1-5) announces (line 5) that MR is "*the infectious disease doctor*." The resident indicates that he works for the IDA (of Figures 4 and 6) but does not specify his status as a first year intern learning about clinical aspects of infectious diseases. Knowing that the physician is a novice might make some patients anxious or confused while others might not comprehend its potential significance. Notice that MR's elicitation procedures were facilitated considerably by knowledge of the two prior medical histories and physical examinations, first in the ER, and then after the patient's admission to the medical ward. The lack of focus and its didactic style would not have been as likely in a clinic setting where there might have been little or no prior clinical information.

Figure 9 - MR's (intern's) interview of patient on the ward

MR: Hi! I'm Dr. X
P: (?)
MR: That's okay. This is (researcher).
C: Hello.
MR: I'm the infectious disease doctor.
P: Oh, are you!
MR: I work with Dr. Y, /she's my boss.
P: /I've heard about you. Okay.
MR: It looks like your doctors are doing a real good job treating that infection with penici /llin.
P: /I think so.
MR: I wanna take (?) your blood. (?) blood pressure (?) (5 sec.) (mumbling) A lot of uh people get, bacteria in their blood, and ya get uh shock. [Low monotone intonation] When your blood pressure goes down because it means (?) the bacteria releases certain toxins, depending on (?) area [raised intonation] Has your, blood pressure always run kind of on the low side, mam?
P: Never.
MR: Never. Do you know what your blood pressure has been in the past?
P: Uh well ever since I been operated on they say it's been normal, and sometimes they'll tell you 100 over something and I don't pay any attention /to it.
MR: /All right.

/ = overlapping speech

In lines 9-10, Figure 9, the MR is looking at the patient's swollen, infected eye while commenting on the use of penicillin. The remark about taking the patient's blood pressure (lines 12-17) is curious because it is highly unlikely that the patient could make sense of nor be able to understand the link between low blood pressure and the presence of "bacteria in their blood." Nor the idea that the condition can be a consequence of the notion (Figure 9, lines 14-16) that "*When your blood pressure goes down because it means (?) the bacteria releases certain toxins* [...]" The comments by MR reflect his having read the patient's medical records, attended microbiology rounds similar to those cited in Figure 4, and having read the written lecture notes by the IDA (his "*boss*" in line 6, Figure 9). This reasoning also appears to motivate the leading question in Figure 9, lines 16-17: "*Has your, blood pressure always run kind of on the low side, mam?*" The reply of "*never*" is ambiguous and provides little help in clarifying the lowness of the patient's blood pressure in the ER and medical ward. The ambiguity continues when the patient (lines 21-23) states "*Uh well ever since I been operated on they say it's been normal, and sometimes they'll tell you 100 over something and I don't pay attention to it.*"

The MR did not follow-through to ask what "never" could mean to the patient, nor did he pursue the occasion for her having "been operated on." A KB system would have to assume that information on the patient's "normal" blood pressure would be elicited. Such a system, therefore, presupposes particular interviewing skills by the novice or nonspecialist.

The MR's remarks are a combination of trying to elicit information, simultaneously seeking to "educate" the patient about her condition, and providing the MR with a captive audience for demonstrating his knowledge of the patient's condition to both the patient and researcher. The patient's remarks about her blood pressure was not pursued by the MR in order to see if he could activate the patient's memory about when she might have had her blood pressure taken rather than employing the leading question in line 16-17, Figure 9

The MR's questions were awkward despite his prior examination of the patient's medical records. A knowledge-based system that made use of information provided by the Director of the Microbiology Lab and the infectious disease attending's lecture to first year House Staff or interns could have provided the MR or another novice with the guidelines necessary to interview a patient. The guidelines, therefore, could include examples of the kinds of questions that should be asked of the patient and some indication of how to modify questions if the patient does not provide useful information.

The information in Figures 4-9 consists of fragments from one of a larger corpus of several cases observed and recorded in the Division of Infectious Diseases. The discussion of a particular case after the patient had been seen in the ER and admitted to a medical ward in order to call attention to the interaction of experts and novices and to reveal how the initial discussion of a particular patient with a serious periorbital eye infection is consistent with the formal training received by medical students (as described by the Director of the Microbiology Laboratory) and house staff or interns and residents (the lecture prepared by the infectious disease attending). The information in Figures 4-9 provides the reader with illustrative material with which to understand the kinds of medical knowledge about infectious diseases a computational modeler should be aware of in order to create a knowledge-based or expert system.

Details about an infectious disease patient and aspects of the training of health care personnel were used to illustrate the kinds of clinical judgments or decisions that can contribute to the computational modeler's on-line understanding of local medical ecologies and exchanges. The above material provides a glimpse of the kinds of training and discourse practices that occur within health care education and health care delivery, and the complexity of bureaucratic and clinical activities that should be examined when creating and updating or modifying viable KB or expert systems.

DISCUSSION

The most problematic state of affairs in medical history taking and physical examinations that a KB system must confront is the kind of training and knowledge base needed to enable a novice to pose questions that will activate the patient's memory such that a response addresses the question posed by the physician. Further research is needed on the extent to which we can infer the novice's motivation for asking a question and the claim that the patient's response addresses the intent attributed to the question.

KB systems could also enable a novice to follow the program on a CRT while interviewing the patient on-line. Some physicians would surely note that such a procedure might be too demanding and make it difficult to create and sustain a "bedside manner" or "rapport" with the patient. The demands of the KB system, therefore, could divert the novice's attention away from posing adequate questions and not addressing subtleties or nuances of the patient's responses (e.g., voice intonation, rhythmic aspects of speech, ambiguous polysemic terms of reference, facial expressions).

Activating the patient's memory "appropriately" depends on the language employed by the physician and being able to comprehend and thus influence the patient's reasoning and discourse about past and current symptoms, medications and treatments. Developing KB systems, therefore, presumes that we have some understanding of past training, the knowledge base and routine practices expected within particular organizational circumstances, namely, the conditions in which non-experts are likely to use such systems.

Formal training in clinical medicine in American medical schools is seldom accompanied by explicit instruction about the organization of memory and how to access past experiences yet avoid leading and misleading questions. Most American medical schools provide students with some experience interviewing simulated patients or actual patients, but monitoring the quality of such training (directly reviewing actual interviews with patients early on and during their clinical clerkships) can vary considerably.

One way to use an intelligent system is to view novice-system interaction as both a tutorial and a way to update the system based on the kinds of problems novices encounter as they use the system. Clancy (1983) suggested such an approach several years ago, but I am not aware of general and widespread attempts to implement such a system. The implementation would require detailed knowledge about the organizational circumstances in which a knowledge-based or expert system would be employed. Particular systems that address specialty or sub-specialty areas in medicine could be available in particular clinics and hospitals or geographical areas where there are few specialists or where they do not exist. The systems could also be used as training devices (as suggested by Clancy) in which an attending would subsequently interview the same patient as the novice in order to insure quality medical care and as a means of fine-tuning the system.

Knowledge-based systems, knowledge networking, and their organizational environments

The development of expert or knowledge-based (KB) systems has been in existence for over three decades within AI, but their use has not been as prevalent as might have been expected.

Another related area of AI is called Knowledge Networking (KN). Research in KN seeks to understand the interaction and flow of information and knowledge among people, and across organizations and communities. For example, understanding knowledge generation and use, encouraging collaborative computation and remote interaction (such as distance learning and problem solving) by transforming information into knowledge structures that can be distributed widely.

An "inference engine" is required to program algorithms, logical operators, and heuristics in order to represent knowledge explicitly with a special language. An "inference engine" refers to the use (and manipulation) of explicit knowledge, drawing conclusions and guiding the behavior of an entire system. In order to go beyond the known requirements of compilers, editors, browsers, debuggers, we need to know about the activities of actual work settings, and workplace-oriented software tools for editing, inspecting, and organizing knowledge and tracing inferences.

An expert system, like a servomechanism, e.g., a thermostat, contains implicit knowledge about how the internal workings of such a device would raise or lower the temperature in a room. No special language is required to operate the system. KB systems, however, require explicitly stated knowledge as well as a special language to represent this knowledge (cf. Gonzalez and Dankel, 1993; Stefik, 1995). The "special" language, however, often borrows lexical items employed in natural language, but gives them special meaning for a particular type of application. Confusions between a special language and natural language usage remain a problem.

One source of confusion when we use a special language is that their similarity to natural language lexical items activates nondeclarative background knowledge perceived as having semantic equivalents in natural language. For example, KB and KN systems require a set of "rules" and an object-oriented knowledge base that depend on explicit semantic nets for creating graphs whose nodes are objects or classes, and the use of edges to indicate relationships. Decisions about creating a semantic net and a generator of explanations are not self-evident; potential confusions can emerge when lexical items appear to be cognates of implicit natural language usage within the software.

Textbooks and research papers that address the components of a knowledge-based system tend to take for granted the communication, decisions and research strategies necessary for field research on experts' activities in their organizational settings. Systematic sampling in workplace environments are needed to identify "essential" semantic nets said to constitute an "appropriate" database. The special language needed to represent the expert's knowledge and practices remains difficult because semantic and pragmatic theories in linguistics are not tested in everyday intellectual work environments. Workplace discourse must become a topic of research in AI rather than a taken-or-granted resource if some of its properties are to be made compatible with the logical structure necessary for creating algorithms or heuristics or fuzzy sets used by students of AI.

Within the medical profession, experts are accustomed to dealing with a formal vocabulary that seeks to avoid the polysemic and context-specific usage associated with natural language discourse. Medical diagnostic language, however, has its roots in the use of natural language during physician-patient interviews. Physicians must constantly disentangle the overlapping semantic elements used during medical interviews and physical examinations that are subequently represented (but only partially) in the creation of medical records.

Past and current AI, and experimental research on medical diagnostic modeling and reasoning (Kassirer, 1989; Boshuizen and Schmidt, 1992)) do not address the interpretative skills needed to elicit information from patients and transform the everyday discourse of medical interviews and physical examinations into the formal clinical discourse of experts when responding to the queries of researchers in AI. A paradox exists within professional medicine; clincial practitioners are suspicious of expert systems, yet such systems enjoy a following in the area known as "medical informatics." In order to interest physicians in KB systems, and make knowledge engineers sensitive to how experts solve routine and special problems in workplaces, detailed observations of the local organizational practices between experts, novices and patients are needed.

Students of KB systems recognize the need to create a database by interviewing experts about their work in order to eliminate contradictions and obtain a better grasp of what constitutes expertise in a particular workplace. We can clarify the nondeclarative and declarative knowledge base of experts and novice before and during the use of KB systems if we systematically observe and record on-line their clinical activities. Consider the following:

1) AI modelers recognize the need to look at the physical and social environments in which KB or KN systems are implemented, yet they often minimize the fact that cognitive/linguistic mechanisms are always embedded in functional social networks. The socially organized aspects of networks in which informal discourse occurs, therefore, may obstruct or facilitate collaborative problem solving and lower the attainment of goals within organizational environments, rather than achieving desired, efficacious communication and decision-making. Strained interpersonal relations and rivalries often emerge based on competition for power or scarce social and monetary resources. Organizational goals, therefore, can be rendered problematic because of networks of interpersonal relationships, or network members may overlook flaws or mute them, and may promote colleagues based on friendship. For example:

a) Individuals must "work" at creating and sustaining social networks and acquiring information about the likely knowledge base and interpersonal skills and habits of network members while expecting and seeking an accommodation to perceived individual differences in their interpretation of activities.

b) Within medical settings, for example, health care professionals should be but are not always aware that patients are seeing and receiving medications from several specialists and subspecialists. The medications each physician can prescribe may not be compatible and may create some discomfort and other symptoms. Medical record systems tend to be chronological and details about visits to different specialities can be poorly integrated despite the fact that in some health care systems, a general or primary care practitioner is supposed to be coordinating the patient's visits to specialists. A source of problems is the late arrival of a patient's medical records or the arrival of only some of the records at the time of an appointment.

2) How do workers compare and categorize messages or documents while continually inundated with information and communication? Personnel invariably reorganize some of the information to the exclusion of other materials, but can seldom judge the trade-off between what is captured by reorganization and what is lost or ignored.

a) The social organization of all workplace activities assumes that all personnel can grasp relational similarity in order to engage in abstract thought necessary to comprehend oral and written information or communication. In addition to enabling us to understand patterns of symmetry and mathematical structure and other invariances, relational similarity also allows us to perceive essential analogies and develop links between domains that appear to be unlike superficially (Kotovsky and Gentner, 1996).

b) Commonalities of language comprehension and use (knowledge of relational terms) associated with communication assume that self-evident relational comparisons exist when messages emanate from the executive offices of different parts of an organization. How should we construct messages that will lead to facilitative rather inhibitory links? Messages sent to an entire organization (where n>50) must employ a standardized syntax and semantics, unless it can be assumed that the use of certain lexical items (as in medicine) are peculiar to the organization. When sending messages to particular departments or divisions, management personnel should be sensitive to the kind of language that is used routinely within the division. Division leaders should have a sense of the kinds of social networks that exist and use them when creating a division of labor in which subgroups will engage in problem solving tasks.

3) An example of language use that can challenge assumptions about relational commonalities is the ubiquitous use of analogies and metaphors in messages intended for specific and (and sometimes general) audiences. Analogies and metaphors allow us to use one situation or issue by reference to another. According to Holyoak and Thagard (1995: 4-9), the use of (or thinking with) analogies requires "...a mapping , or systematic set of correspondences, between the elements of the source and the target analog" that are assumed to exist.

a) In addition, "...the analogy is guided to some extent by direct similarity of the elements involved." For example, we assume common properties exist between different parts of an organization as well as objectives or goals and official and unofficial relationships that emerge, exist, or are modified. Analogies and metaphors allow personnel to create a special lexicon that identifies them as members of a particular group or network and facilitates communication, comprehension and inferences across existing and emergent official and unofficial relationships. The identification and study of the use of analogies and metaphors within and between groups represents a major challenge for the development of knowledge-based systems. To what extent should such systems address and reflect the role of social networks and problem solving practices when dealing with information overload and communication oveflow?

b) According to Holyoak and Thagard, "...the analogy is guided by a pressure to identify consistent structural parallels between the roles in the source and the target domain." For present purposes, this can mean the study of structural parallels when employees develop and use such strategies to cope with information overload and communciation overflow. For example, in health care delivery, there is an essential need to create structural parallels when health care personnel scan, retrieve or seek to integrate new or modify older information by reference to their understanding of an existing policy, rule or other constraint about symptoms and laboratory results.

c) For Holyoak and Thagard (1995: 5-6), however, there should be "...a one-to-one constraint: each element of the target domain should correspond to just one element in the source domain (and vice versa)." In the complex, social network, polysemic environments that coexist within large organizations, such a constraint poses challenging empirical problems.

4) In a dynamical, organizational context (in contrast to problem solving in the idealized laboratory), can we readily identify the "goals"

inherent in the use of analogies? Holyoak and Thagard presuppose but do not address organizational environments when they state that "...analogy is guided by the person's goals in using it, which provide the purpose for considering the analogy at all....[and] people generally tend to select especially familiar situations to serve as source analogs" (p. 6). Pursuing a person's goals within an organization can be motivated by their understanding of official constraints as well as the concerns of an informal social network with somewhat different goals.

a) Although organizational messages are often couched in the language of directives or imperatives, and factual or technical statements, analogies enable participants of communication to express their thoughts indirectly and thus go beyond the information given formally.

b) Within social networks, therefore, the use of analogies can incorporate implicit messages. Messages used within informal social networks, therefore, are not only functional, but presuppose unstated experiences or commonalities that constitute challenges to computational modeling. Medical records, for example, often contain declarative statements that are factual or technical, yet also may be expressed informally between health care professionals (as in Figure 4) and thereby take on different semantic connotations. The connotations may be visible occasionally in a coded form in a patient's medical records.

5) A related problem is how should we design field research about work environments to clarify pending changes for workers because of the introduction of an expert or knowledge based system? For example, the research necessary to create an expert system should recognize the constraints that can occur because of formal management goals and practices designed to make such settings "efficient" or "more" cost-effective. We should also understand the way workers pursue informal or unofficial efforts that can make workplaces "friendly" or "pleasant" rather than "tolerable," but at a cost not welcome by management.

The huge literature on workers' reaction to automation, management attempts to increase productivity while often ignoring so-called "human factors" has a long industrial management history. The well-known study of the relay assembly test room and bank wiring room at the Western Electric factory in Chicago (Gillespie, 1991; Homans, 1950; Mayo, 1933; Roethlisberger and Dickson, 1939;) remains an unclarified example of the role of "human needs" that called attention to improved work conditions, workers' social networks, and their apparent resistance to managerial control while exercising economic self-defense in order to explain productivity in the workplace. Workplace conditions create a contextual framework within which KB systems must function.

The creation of a knowledge-based system to address cognitive overload and communicative overflow requires an understanding of

1) how individuals process and reorganize information when engaged in activity within their work settings,

2) the role of communication and information processing in group problem solving environments,

3) and how personnel make use of links within informal systems or social networks in order to comprehend and cope with their perception and processing of information and communication in the workplace.

The formal/informal structures and processes that always permeate routine organizational practices revolve around poorly understood communication and unofficial collaboration necessary to initiate and sustain functional and disfunctional social networks. Formal communication tends to promote and expect explicit work routines that are compatible with the development of knowledge-based systems, but (inadvertantly or by design) can significantly alter or replace aspects of workers' preferred daily informal social network activities. Formal communication (letters, reports, memos) presuppose and build on prior (often unofficial) informal exchanges and agreements or disagreements and negotiation. The informal system (such as the exchange in Figure 4 between the two experts PA and IDA) tends to remain "invisible" yet essential to reproduction and change in the formal system. Informal, often tacit, yet ubiquitous elements of communication and knowledge that exist in all organizations must be addressed when we create and use expert systems.

CONCLUSION

Decisions about the format in which a knowledge-based or expert system or KN system will be implemented presume that structured health care units exist within an organization such that the semantics and pragmatics of certain kinds of messages will be readily understood across departments, groups and subgroups.

Knowledge networks, however, differ from knowledge-based systems because of the necessity of developing representational flexibility; creating distributed representations that can include the entire organization or subsets of it. Individuals, interacting groups and "closed" social networks must be capable of comprehending or processing "reasonable" mappings between predicates with different numbers of arguments, embedded in known or emergent activities that make use of complex metaphorical and analogical structures. A document search program should reflect the kinds of issues, common, everyday, and technical analogies, and metaphors that are comprehensible to everyone in the organization despite the existence of groups and social networks that develop and sustain their own metaphors, metonyms, and analogies.

Individual and collective management-knowledge processing practices required in health care systems are always part of a larger institutional setting which exerts constraints that are often not directly evident to an observer nor to personnel in the system who have worked long enough to take for granted the policies and practices that have become part of their nondeclarative memory. The socialization of workers can include an orientation to elements most likely to be used in an expert system, but also alert employees to the way a system (the use of particular categories) can elide or mask important information. For example, if the socialization of employees is limited to particular work routines, then an expert system may be less relevant for this subset and others like it within an organization. Personnel turnover or routine changes in work schedules can also affect (e.g., weaken) institutional memory about personnel practices and procedures and the way personnel are trained.

The hiring and training of new personnel exacerbates the propagation of information through a bureaucratic system in which hierarchy, informal organizational practices and social networks of communication and assistance create semi-self-contained worlds that can facilitate, filter, mask and distort larger organizational or institutional goals.

Creating uniform socialization practices and periodic updates to review larger and local goals can diminish information overload and communication overflow problems. A form of "principled redundancy" is required to sustain but not block the cross-referencing of tasks performed by different groups and ecologies. The notion of "controlled redundancy" can help avoid indifference or damaging competition between groups (or social networks) that are supposed to share information and technology for improving functional aspects of organizational performance.

Creators of knowledge-based or expert systems rely on heuristics and recognize the relevance of organizational problems such as identifying aspects of background knowledge necessary to address practical issues. Rather than use a step-by-step procedure such as a diagnostic algorithm that is less flexible, one favored strategy is to use pieces of knowledge or heuristics that can be incorporated into rules, and an inference engine that can implement the rules. The idea is to write self-contained rules that can incorporate independent fragments of knowledge about, say, a diagnosis. Elkan (1998) notes that using heuristics is like solving a jigsaw puzzle; you keep adding pieces and hope that a larger picture begins to emerge. Elkan notes that a knowledge base is like a repository of partial cases of problem solving or inference by retrieving and adapting a case that appears to be similar. A "similarity" metric is needed that depends on relevant features in order to create indices

Assigning feature values to cases, however, is not self-evident, nor is the idea of a metric distance that compares pairs of items to give a number or the distance between the two items. We need to know which features are predictive of producing a similar solution and how can we index this solution and place the outcome into a category. Placing an outcome into a category is similar to a notion of rules described by John Rawls (1955). His reference to two concepts of rules specifies a rule and then a rule that allows one to place a case under the specified rule.

The challenging problem, however, is the reasoning that enables someone to decide that a particular case falls under a specifiable rule. A problem associated with Rawls' notion of a rule is how we can identify the cognitive/cultural mental models needed to generalize from a set of perceived conditions or circumstances, namely, the declarative and nondeclarative knowledge needed to frame or index an entry. For example, perception, thought, and action that triggers elements about the world or the local ecology such that the activation of working and long-term memory leads to analogy formation and inferences about a state of affairs. One consequence is emergence of relational reasoning and cognitive/cultural schemata that go beyond any self-evident comprehension of an event or set of speech acts.

The research on LISA by Hummel and Holyoak (1997) and similar models (by Gentner, 1988, for example) rely on sentence-based (context-free) semantic databases in order to satisfy existing architectures for modeling natural language processing. The assumption is that sentences, their grammatical structures and the meaning of the lexical items therein can be decomposed into invariant, smaller units or elements of structural constraints and meaning. The sources of data, however, tend to remain hypothetical and not based on recordings of actual language use and social networks in practical settings.

The developer of knowledge-based or expert systems, however, recognizes the necessity of speaking to experts within an organization and ideally of preparing personnel for the eventuality of introducing a knowledge-based system. The student of analogy formation and metaphorical usage seeks to make complex analogies or metaphors tractable by using a computational architecture that will interface with a compositional semantics while ignoring knowledged-based systems that rely on heuristics.

A curious paradox emerges. The AI student encounters difficulty modeling the structure and use of analogies and metaphors despite their inherent use in organizational social networks. AI students propose that a separate set of language categories is required in order to write software programs to avoid the pitfalls of natural language use in daily life settings. The student of analogy formation and metaphorical usage, however, is dependent on tacit everyday reasoning and language use as a resource for assembling a database necessary for modeling.

Students of AI and researchers of analogies and metaphors often rely on descriptions of experts within the organization but neglect the daily life activities (reasoning and language use) of individual personnel during collective or conflictful and collaborative exchanges and/or problem solving. Modelers in both areas just noted are constrained by the kinds of computational architectures currently available for developing software or simulations of decision making and information processing in actual organizations..

I have underscored the importance of informal organizational structures and processes, particularly the social networks that personnel rely on for learning about new policies and procedures. The study of metaphors and analogies and their use in the creation of future computational knowledge-based systems or simulations can benefit from task-oriented ethnography, namely, a better understanding of routine, on-line, day-to-day problem solving in organizations.

References

Boshuizen, H.P.A. and Schmidt, H.G. (1992) On the role of biomedical knowledge by experts, Intermediates and novices. *Cognitive Science*, 16, pp. 153-184.

Cicourel, A.V. (1974) Interviewing and memory. In C. Cherry (ed.), *Pragmatic aspects of Human Communication*. Dordrecht: Reidel, pp. 51-82.

Cicourel, A.V. (1982) Language and belief in a medical setting. In H.Byrnes (ed.), *Contemporary perceptions of language: Interdisciplinary Dimensions.* Georgetown: Georgetown University Press, pp. 1-41.

Cicourel, A.V. (1990) The integration of distributed knowledge in collaborative medical diagnosis. In J. Galagher, R. Kraut, and C. Egido (eds.), Intellectual teamwork: *The social and technological foundation of cooperative work.* Hillsdale, NJ: Lawrence Erlbaum, pp. 221-242.

Cicourel, A.V. (1995) Medical speech events as resources for inferring differences in expert-novice diagnostic reasoning. In U.M. Quasthoff (ed.), *Aspects of oral communication*. Berlin: Walter De Gruyter, pp. 364-387.

Clancy, W. (1983) The epistemology of a rule-based expert system: A framework for explanation. *Artificial intelligence*. 20, pp. 215-251.

Elkan, C. 1998. Lecture notes on Intelligent systems.

Gentner, D. (1983) Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, pp. 155-170.

Gentner, D. (1988) Metaphor as structure mapping: The relational shift. *Child Development*, 59, pp. 47-59.

Gillespie, R. (1991) *A history of the Hawthorne experiments.* New York: Cambridge University Press.

Gonzalez, A.J. and Dankel, D.D. (1993) *The engineering of knowledge-based system*: Theory and practice. Englewood Cliffs, NJ: Prentice Hall.

Hummel, J.F. and Holyoak, K.J. (1997) Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review*, 104 (n. 3), pp. 427-466.

Holyoak, K.J. and Thagard, P. (1995) *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.

Homans, G.C. (1950) *The human group*. New York: Harcourt, Brace.

Kassirer, J.P. (1989) Diagnostic reasoning. *Annals of internal medicine*, 110 (n. 11).

Kotovsky, L. & Gentner, D. (1996) Comparison and categorization in the development of relational similarity. *Child development*..

Mayo, E. (1933) *The human problems of an industrial civilization*.. New York: Macmillan.

Mitchell, T.M. (1997) *Machine learning*. New York: McGraw-Hill.

Rawls, J. (1955) Two concepts of rules. *Philosophical Review*, 54, pp. 3-32.

Roethlisberger, F. and Dickson, W.J. (1939) *Management and the Worker: an account of a research program conducted by the Western Electric Company, Hawthorne works, Chicago*. Cambridge, MA: Harvard University Press.