THE PRAGMATICS OF EXPERTISE IN MEDICINE

David A. Evans
Departments of Philosophy and Computer Science
Carnegie Mellon University

INTRODUCTION

A central goal of cognitive science research in problem solving has been to illuminate the nature of expertise (Chi & Glaser, in press; Johnson, et al., 1981). Some of this research has centered on the acquisition of specific skills that are prerequisites of expert performance (Feltovich, et al., 1984); some, on the differences in knowledge—in scope and organization—found among experts, intermediates, and novices (Chase & Simon, 1973; Chi, Feltovich & Glaser, 1981; Larkin, et al., 1980; Patel, 1984; Patel & Frederiksen, 1984).

The picture of expertise that emerges from these studies is of individuals who *know more* than non-experts; but whose knowledge differs not only in *quantity* but also *quality*. In particular, the knowledge of such experts is *dense* and *context-sensitive*: experts not only seem to have *more* domain-specific concepts, but more relations among concepts; and experts have a greater facility in identifying appropriate *contexts* in which to interpret the data in a problem (Evans, Gadd & Pople, in press).

In a domain like medicine, where problems may involve complex details and solutions may be grounded in incomplete models of pathophysiological processes, knowing more can lead to a tension between the need for both effective problem-solving behavior and thorough investigation of rival hypotheses. How such tensions are resolved reveals aspects of the pragmatics of medical expertise. This paper identifies two associated phenomena resulting from domain pragmatics—*selective ignorance* and *disciplinary ecological equilibrium*—that account for apparently sub-optimal behavior in medical expert problem solving.

ON PRAGMATICS IN MEDICINE

Pragmatics generally refers to the collection of effects in the processing of information that do not derive alone from either the *surface data* or the *semantics of the code* in the apprehended message. This characterization applies equally to linguistic and non-linguistic communications. Thus, pragmatics subsumes the inferences, interpretations, grounding of references, and other effects of context—whether the 'visible' co-occuring phenomena that accompany a communication or the 'invisible' mental models—that are mobilized in the pursuit of interpretations.

Pragmatics in medicine deserves special attention because successful problem solving depends on identifying the appropriate contexts in which to interpret the typically overwhelming number of possible features in any particular medical case. Such contexts are discovered by organizing surface observations (e.g., symptoms, signs, laboratory data) into patterns suggestive of diagnostic hypotheses; but the patterns themselves reflect components of disease models—idealizations of pathophysiological processes, often expressed as pseudo-causal relations among individual observations in a case. Since the available data in most medical cases are too numerous, and the available correlations too complex, to support direct inductive reasoning, physicians must rely on *a priori* models that may be underdetermined by the immediately accessible information (Evans, Gadd & Pople, in press). One natural result is that physicians *hypothesize* candidate diagnoses; and one effect is that candidate diagnoses define contexts in which data may be *over-interpreted.*

ON EXPERTISE IN MEDICINE

Experts in medicine are individuals who have achieved a high level of competence, perhaps as measured by peer assessment—individuals who have completed all *normal* phases of training and have been practicing medical problem solving for some time. Typically, such experts are not expert across-the-board, in all phases of medical problem solving, but are 'specialists'—experts in solving domain-specific problem types, e.g., problems in internal medicine or pathology. But though specialized, medical expertise is not special-case or extraordinary expertise—is not cognitively anomalous—rather, is attainable by normally intelligent individuals. In particular, medical non-experts (novices, intermediates) are not cognitively immature individuals or laypeople. They are medical students, interns, and residents: individuals who have attained a high level of general competence and some degree of specialized training. The focus of research in medical problem solving is on how such non-experts acquire *domain-specific* knowledge and competence.

One measure of expertise is the ability to manage an amount of information that would overwhelm the novice. Indeed, that may be the single most striking and consistent characteristic of experts across subdomains of medicine. But there are aspects of medical expert behavior (and expert/novice differences) that indicate that such expert ability to manage data may not lead to commensurably improved reasoning abilities in problem-related tasks. For example, experts make accurate diagnoses even when they *explain* their diagnoses incorrectly (Patel & Groen, 1986). This suggests that different relations among the features in a case may be involved in actual problem solving than are revealed in the normative models of pathophysiological processes that form the community standard for medical explanations. As a related example, experts in one medical domain have difficulty solving problems in another medical domain, even when presented with the *basic science* information required to account for the pathophysiological processes involved in the problem (Patel, Arocha & Groen, 1987; Patel, Evans & Groen, in press). This suggests that the inferences that experts make are not based on 'first-principle' reasoning, but are guided by other, domain-specific, reasoning processes. As a final example, experts are typically more intolerant of uncertainty than novices, though they have the knowledge required to see a greater number of possibilites in interpreting data (Patel, Evans & Kaufman, in press). This suggests that the processing of information by medical experts is at least partly *categorial*, and that the categories effectively filter variation in data.

In short, expert performance is not perfect; and aspects of expert performance fall short of standards that experts, themselves, may acknowledge as desirable. (This is especially the case in the emphasis on basic science as a basis for medical reasoning.) More problematic than simple failures in performance, though, is the fact that the observed behavior of experts—in their approach to problems, in their explanations—is consistent, coherent, and superficially rational. What can account for this?

EFFECTS OF PRAGMATICS IN MEDICAL PROBLEM SOLVING

If we take the above observations at face value, we must conclude that medical experts are *not* capable of flawless judgment, even in their areas of specialization; that, in a sense, they practice *false* science, or at least a different science than they imagine. In either case we should be concerned. If there is no method at all, there is no basis for the decisions physicians make. And if there is a method, but it is different from the one used to inform the discipline, there is no basis for predicting the effects of change—modifications of the normative models in the *façade* method—on actual practice.

On the surface, this has the form of a dilemma. I suggest, however, that it is the natural consequence of problem-solving behavior 'in the limit'; and that this situation may actually reflect an optimal adaptation to the needs of communication in complex domains. In particular, I would claim that the observed sub-optimal performance on some problem-related tasks can be attributed to the pragmatics of information management in medicine. One pragmatic effect is *local* (and 'low-level', in actual problem solving): the need to discriminate phenomena so that they can be analyzed as elements of complex categories. This leads to *selective ignorance*—a failure to see variation and to distinguish 'source' data from interpretations. Another pragmatic effect is *global* (and can have consequences for perfomance at all levels): the need to organize information in problem-solving situations is greater than the need to establish accurate relations among elements. This leads to an emphasis on principles of organization and on compliance with established standards of coherence (at the expense of principles of science) giving rise to *disciplinary ecological equilibrium*—where degredations in community standards pose a greater threat to performance than degredations in models. The following sections offer further reflections on these points.

Selective Ignorance

One source of error in expert performance derives directly from a source of efficiency: experts may fail to see certain relations in data because invoked, candidate interpretations effectively account for many of the relations in the data, obscuring possible alternative interpretations and inviting an under-interpretation of anomalous data. This is seen, for example, in studies of novices, intermediates, and experts in doctor/patient interviews (Patel, Evans & Kaufman, in press). The principal problem for novices is *seeing too much*—they cannot discriminate the relevant from irrelevant details; and they cannot organize all the information they gather in interviews (thought they may retain it). The principal problem for experts, on the other hand, is that they may *mis-manage* the flow of information by over-interpreting possible cues and, consequently, ignoring possibly important details.

In actual practice, the salvation in the first case is knowledge; in the second, discipline.

Through knowledge, the novice acquires prototypes that serve both as targets for the interpretation of details and also as predictors of additional relations or details that may be relevant. Prototypes become signposts on the way to the successful organization of the information in a case; and prototypes are *learnable*. Through discipline, the expert checks the tendency to premature closure—in particular, by following procedures that will typically require that details be examined from several points of view before a conclusion can be drawn. This forces the expert to re-examine the preferred hypothesis, though it may not force him or her to abandon it in the face of counterevidence.

It is interesting to note that experts we have studied do not recover well from defeated hypotheses. They frequently discard methodology at such points and generate inferences that are based on possible associations of data—not necessarily probable or plausible ones. Indeed, attempts to appeal to 'basic scientific' reasoning at such times is actually counterproductive (Patel, Evans & Groen, in press).

This may be partially explained by a related phenomenon, viz., that expert understanding involves what appears to be simultaneous derivation of a *situation model* along with the processing of *surface data*—and that the two are integrated in recall (Patel, 1988). This is distinct from the results with novices, where there appears to be a separation in memory of the facts and the derived situation.

Both phenomena—the over-interpretation of cues (and consequent under-interpretation of variation) and the integration of source data and interpretations—bear a striking similarity to what we see in *categorial perception*—where discrete phenomena are <u>not</u> noticed if they fall below a *threshold of perception*; but are homogenized and integrated into a *perceptual* or *processing category*, if they exceed the threshold. Typically, in such cases of perception, the source stimulus is lost and only the category and its prototypical predictions are recoverable.

In the face of noisy or complex data, categorial perception gives rise to efficient performance—but also selective ignorance: one has the ability to make clear judgments but one loses the ability to make (or see) fine distinctions. One important effect of selective ignorance is that all people who share the perceptual categories will agree on their classifications of phenomena. (This phenomenon assures agreements in phonological interpretations, for example, among native speakers of the same language.) But another effect is that when categorization proves incorrect, it may be difficult to see data objectively, without partial interpretations.

Disciplinary Ecological Equilibrium

A series of studies on the basis of medical-concepts have recently suggested that actual medical expertise may in some instances be based on *false* information (Coulson, Feltovich & Spiro, 1986; Feltovich, Spiro & Coulson, in press; Spiro, Feltovich & Coulson, 1987). Such observations clearly present a challenge to the science of medicine, but also to our understanding of effective problem solving. If the acknowledged basis for problem solving is flawed, what accounts for the persistence in use of inaccurate models in medcial textbooks, in medical curricula, and in expert discourse? Is this one source of the sub-optimal behavior we observe in expert problem solving?

One analysis is that false models arise, in part, through simplifications that facilitate learning. The strong claim in this case would be that such practice leads to a *pernicious degradation* in the validity of a science. In the worst case, the cycle becomes: Learning *efficiency* leads to performance *deficiency*, which later leads to *new* learning *inefficiency*. More concretely, the simplifications that make the learning of complex concepts easier on first exposure actually become the bases of models that are flawed. These, in turn, give rise to performance anomalies; and inhibit the acquisition of new—and possibly 'correct'—models of biomedical phenomena.

An alternative interpretation is possible. Simplifications (even falsifications) are both adaptive and necessary in complex domains such as medicine. The claim would be that the need for coherence outweighs the need for accuracy; and coherence is assured only if the normative models are widely accepted, remembered, and used as the basis for the description, analysis, and discussion of problems. To the extent that detail is complex and renders models difficult to learn and use, actual accuracy can lead to degradations in agreement on coherence standards. Rather than promote performance deficiency, changes in practice (e.g., emphasizing scientific accuracy when we train physicians) may present a greater threat of reducing global expertise than hope of improving disciplinary performance.

The tension here is between the need for global coherence and problem solving efficiency (purchased via selective ignorance), on the one hand, and accuracy of knowledge (purchased via attention to greater detail—and perhaps intolerance of variation), on the other. The claim would be that in domains where perfomance grows asymptotic in the limit—and, more importantly, where expertise is *collective* in practice, not individual—there is a disciplinary ecological equilibrium—in which the need for great sensitivity to coherence and great conformity with paradigms (to promote agreement in problem solving behavior) is more important than accuracy of the accepted models. In such a case, the superfical biases we observe would actually be the symptoms of constructive adaptive behavior.

CONCLUSION

The source of the problems we observe in medical experts derives from limitations in human intelligence (and epistemological models) in capturing the ontology of biomedicine. The uniformity of performance across experts reflects the 'sociology' of the discipline—the adaptive behavior associated with shared and reinforced principles of communication and explanatory adequacy.

One lesson is that what experts do well may necessarily determine what experts do poorly. The medical expert purchases efficiency with selective ignorance; and achieves effective communication and coherence through a conspiracy of simplifications. These are two of the pragmatic features of expertise in medicine.

The reflections offered here require elaboration beyond the scope of this paper; but they raise questions that have implications for our understanding of learning and skill mastery in other complex problem-solving domains—not just in *medicine*.

REFERENCES

Chase, W.G. & Simon, H.A. (1973). Perception in chess. *Cognitive Psychology, 4*, 55–81.

Chi, M.T.H., Feltovich, P.J. & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*, 121–152.

Chi, M.T.H. & Glaser, R. (Eds.) (In Press). *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Coulson, R.L., Feltovich, P.J. & Spiro, R.J. (1986). Foundations of a misunderstanding of the ultrastructural basis of myocardial failure: A reciprocating network of oversimplifications (Report No. 1). Springfield, IL: Conceptual Knowledge Research Project, Southern Illinois University School of Medicine.

Evans, D.A., Gadd, C.S. & Pople, H.E., Jr. (In Press). Managing coherence and context in medical problem-solving discourse. In D.A. Evans & V.L. Patel (Eds.), *Cognitive science in medicine*. Cambridge, MA: The MIT Press.

Feltovich, P.J., Johnson, P.E., Moller, J.H. & Swanson, D.B. (1984). LCS: The role and development of medical knowledge in diagnostic expertise. In W.J. Clancey & E.H. Shortliffe (Eds.), *Readings in medical artificial intelligence: The first decade*. Reading, MA: Addison Wesley.

Feltovich, P.J., Spiro, R.J. & Coulson, R.L. (In Press). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D.A. Evans & V.L. Patel (Eds.), *Cognitive science in medicine*. Cambridge, MA: The MIT Press.

Johnson, P.E., Duran, A.S., Hassebrock, F., Moller, J., Prietula, M., Feltovich, P.J. & Swanson, D.B. (1981). Expertise and error in diagnostic reasoning. *Cognitive Science, 5*, 235–283.

Larkin, J.H., McDermott, J., Simon, D.P. & Simon, H.A. (1980). Expert and novice performance in solving physics problems. *Science, 208*, 1335–1342.

Patel, V.L. (1984). *Prior knowledge and comprehension of medical texts: Expert-novice differences*. Centre for Medical Education, McGill University, Montreal, Canada. CME Report # CME84-CS3.

Patel, V.L. (1988). *The nature of understanding and problem solving in medical expertise*. Centre for Medical Education, McGill University, Montreal, Canada. CME Report # CME88-CS4.

Patel, V.L., Arocha, J.F. & Groen, G.J. (1986) Strategy selection and degree of expertise in medical reasoning. *Proceedings of the Eight Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.

Patel, V.L. & Frederiksen, C.H. (1984). Cognitive processes in comprehension and knowledge acquisition by medical students and physicians. In H.G. Schmidt & M.C. de Volder (Eds), *Tutorials in problem-based learning*. Assen, Holland: Van Gorcum.

Patel, V.L. & Groen, G.J. (1986). Knowledge-based solution strategies in medical reasoning. *Cognitive Science, 10*, 91–115.

Patel, V.L., Evans, D.A., & Groen, G.J. (In Press). Biomedical knowledge and clinical reasoning. In D.A. Evans & V.L. Patel (Eds.), *Cognitive science in medicine.* Cambridge, MA: The MIT Press.

Patel, V.L., Evans, D.A. & Kaufman, D. (In Press). A cognitive science framework for doctor/patient interaction. In D.A. Evans & V.L. Patel (Eds.), *Cognitive science in medicine.* Cambridge, MA: The MIT Press.

Spiro, R.J., Feltovich, P.J. & Coulson, R.L. (1987). *Seductive reductions: The hazards of oversimplification of complex concepts* (Report No. 4). Springfield, IL: Conceptual Knowledge Research Project, Southern Illinois University School of Medicine.