

Kochevar, Johnson

Problem Solving is What You Do When You Don't Know What to Do

Laura K. Kochevar
Department of Psychology

Paul E. Johnson
Carlson School of Management
University of Minnesota

Experts make it look easy. They perform their tasks efficiently, and usually, correctly. Tasks which novices perform only at the expense of a great deal of time and effort are accomplished almost immediately, seemingly intuitively by the expert. In developing these skills individuals typically lose the ability to tell us what they do (Johnson, 1983). The fluency of expert performance and the tacit nature of expert knowledge can be understood if viewed as the result of adaptations which facilitate performance in a specific task environment.

Adaptations can be described in terms of environmental specificity, goal relevance, and function. The specific adaptations experts develop are collectively referred to as "knowledge". An agent's knowledge can be described in terms of how the agent knows what can be done to solve a problem. In this case the goal is problem solution, the environmental information is "what can be done" and the function is "knowing". In most theoretical descriptions one or more of these aspects of performance remain implicit. For example, saying "the giraffe uses its long neck to reach food" focuses on the relation of the adaptation to goal directed function -- leaving environmental specificity implicit. Theories of expert performance usually address the issue of how the expert knows how to solve the problem -- also leaving environmental specificity implicit.

Of course, in order for the giraffe's long neck to be facilitate feeding, the giraffe's environment must be structured so that food exists in high places. If knowledge is adaptive then the problem environment must also be structured so that knowing facilitates goal attainment. If the structure of the environment does not specify goal relevant actions then there is nothing to adapt to, nothing to know.

Not all animals eat "leaves in high places". Anteaters have long tongues rather than long necks. The anteaters tongue is a different adaptation to a different feeding environment. Both the anteater's tongue and the giraffe's neck function as "reachers" to accomplish the "feeding" goal. The different adaptations correspond to different environmental

structures. If we were to develop a model of "food reachers" without concern for the structure of "food" we would not readily propose two such radically different models.

The critical question for those constructing models of cognitive processes is "what constitutes food for thought?" We propose that the expert and the novice function in very different task environments. Expert level performance is the result of adaptations which allow extraction of information from the environment which differs from information available to naive performers. It is our assertion that current cognitive theories fail to capture the qualitative difference between expert and novice level performance because they assume individuals function within the same nominal environment.

An environment contains information which is structured in a way that indicates what you can do with it (Gibson, 1979/1986). Most environments are informationally "rich". That is, there is quite a lot you can do with them, hence, they have many structural descriptions. Even the very restricted "environment" which we label "tree" presents various information structures which specify "climb up-able", "obstacle", "hide behind-able", etc. These information structures indicate the real physical characteristics of objects and events that make them useful in task performance. A physician is able to diagnose a patient's disorder because the disease produces physiological deficits which in turn produce symptoms. In this sense, the environment is objective -- the relationship between symptoms and disease reliably exists apart from an observer. The task environment however, is subjective in that it consists of the information meaningful to a particular individual seeking to attain a particular goal.

The environment can be described in terms of its "surface structure" or nominal features. The relationships among features can also function as sources of information. Certain sets of features and feature relationships specify goal relevant characteristics of the environment. Given a fixed goal, and a fixed behavioral potential, there is some set of features and feature relationships which specifies the optimal goal directed behavior. This is the "deep structure" of the task environment.

Experts perform by selecting from the environment information specifying the "deep structure" of the domain. Novices are only capable of accessing the "surface structure". Since novices do not access task relevant feature relationships they must mentally combine surface level information or "deduce" task solution. To the extent that the novice's reasoning consists of invoking processes functionally similar to

environmental structures, the novice will be able to perform the task. Experts need not reason. They can behave intuitively because they can access relational information from the environment.

The distinction between expert and novice level performance is paralleled by the nature of the knowledge underlying each. The expert has task specific adaptations which correspond to the environmental structure related to task performance. "Expertise" is this body of operative knowledge. The novice's adaptations do not allow direct knowledge of the information which adequately specifies performance. Therefore, the novice's operative knowledge must be supplemented by procedural knowledge. The procedural knowledge is the set of cognitive processes which manipulate operative knowledge to derive predictions about the environment.

The fluent, implicit nature of much of an individual's operative knowledge presents a serious problem in the study of expertise (Berry, 1987). Most investigations of expert performance attempt to discover the structure of operative knowledge by placing the expert in a task environment where his or her adaptations fail (Simon, 1969). This forces the expert to invoke procedural knowledge. Although an explicit behavior trace can be generated by this method we cannot be sure that the knowledge observed is similar to that normally used.

Most cognitive theories resolve the problem of implicit knowledge by assuming intrinsic, combinatorial semantics (Fodor and Pylyshyn, 1988). The meaning of a particular representation is assumed to be reducible to the meaning of its constituent features. Since meaning at any level of description can be derived from meaning at another level, it can be assumed that the explicit knowledge trace produced by an expert individual performing in a novel environment is functionally equivalent to the underlying expertise.

In the model we have presented the semantics are environment referenced. This allows varying degrees of relationships between knowledge levels. Given that knowledge is viewed as an adaptation to the environment, and what one knows about the environment is what one can do with it, we can assess the structure of implicit knowledge directly by studying the relationship between environmental information structures and behavior.

We have applied this model to understanding the knowledge used by an expert in the diagnosis of congenital heart diseases. We expected that by analyzing the history of a specific individual's interactions with the environment we could describe this individual's functional environment. Eventual outcome measures were also available, allowing us to describe a

hypothetical optimal knowledge structure. Since multiple individuals treated each patient we could construct a knowledge description of a composite "other"; in this case, a composite novice.

METHOD

For this analysis, 15 patient charts were chosen from the records of the University of Minnesota Heart Hospital. The same expert physician had diagnosed all 15 cases. Each case had also been diagnosed by one or more less experience physicians (novices). After physician diagnosis each patient's actual disease state was determined by cardiac catheterization. Each chart contains the record of symptoms noted, the diagnosis each physician assigned to the case, and the actual underlying disease state. There were three major diseases represented: transposition of the great vessels (TGV), ventricular septal defect (VSD) and atrial septal defect (ASD). It is possible for any patient to have either disease in isolation, both diseases simultaneously, or neither disease. Each of these 15 cases had at least one disease.

Of the three general types of analytical methods available for describing patterns of symptom variation (clustering, multidimensional scaling, and factor analysis); factor analysis is conceptually most similar to the theory we have presented. Unlike clustering methods, factor analysis provides ways to use one symptom in combination with others to indicate more than one underlying pathology. For example, while the combination of blood pressure and age may indicate a cardiac abnormality; the combination of weight and age may signal that the patient was born prematurely or is suffering from malnutrition. Multidimensional scaling techniques are based on dissimilarities among cases while factor analysis methods identify patterns of symptom covariance. The use of factor analysis allows us to identify factors which indicate a particular disease without implying the absence of other disorders.

Factor analytic methods are normally applied to a limited number of variables measured on a large numbers of cases. The goal is to explain relationships among variables in terms of fewer, more general, hypothetical constructs. In this example we are using the technique to mathematically describe this particular sample of 15 cases. We are not concerned at this point about stability of factor loadings or generalizability to other samples. Explanation and data reduction are accomplished by determining which symptom patterns are related to disease state and diagnosis.

What the physician knows about the environment is how symptom structures are related to physical defects. To describe this knowledge

three separate analyses are performed. Two of these relate the expert or novice physician's diagnosis of each case to available symptom patterns. The third uses disease state as determined by catheterization to show which of the available symptom patterns are related to actual pathology. In each analysis the diagnosis or disease of each case is treated as an additional "symptom". All other symptom values are constant across analyses for each case. Any resultant differences in factor structure are due to diagnostic differences and reflect the relationship of symptom patterns to diagnosis and disease.

The factor analyses only indicate which symptom interrelationships covary with diagnostic behavior or disease. They do not tell us what information was actually used. Nor does the analysis indicate what the information means to the physician.

To address these issues, sample cases were presented to a second expert physician. Case pairs in which the cases differed with respect to a single expert knowledge factor were selected from the original data. In addition, other cases from the original data set were presented paired with a fictitious case constructed as a distortion of the original with respect to a single factor. In all, 6 case pairs were presented for the expert's interpretation. We expected that the expert's behavior would conform to factor based predictions. We also expected that the physician would be able to tell us how cases differed from each other, giving conceptual labels to the empirically derived factors.

RESULTS AND DISCUSSION

Figure 1 shows a simplified view of the optimal, expert, and novice knowledge structures. The ovals represent diseases, the circles represent orthogonal symptom patterns. A circle within an oval is a symptom pattern which specifies the disease. When two diseases are related to the same factor, the factor is bipolar and discriminates between them. A set of important nominal symptoms appears to the right. The symptoms which are associated with factors across descriptions are indicated at the far right. The numbers represent the symptom patterns to which they relate. The lines between the knowledge structures and symptoms indicate additional symptoms contributing to each factor. Each analysis accounted for approximately 90% of the diagnostic or disease variance.




Figure 1 Knowledge structures for optimal diagnosis, expert diagnosis, and composite novice diagnosis.

The optimal knowledge structure shows four factors which specify TGV, two factors which specify ASD and two factors which specify VSD. ASD and VSD share a discriminating factor. TGV is independent of the other diseases. This relationship between two diseases and the independence of the third is consistent with the known physiology of the diseases. ASD and VSD are both variants of "left to right shunts" while TGV is classified as an admixture lesion.

The expert's diagnostic behavior is consistent with two of the optimal factors. In this structure VSD and TGV are specified by independent, single factors. ASD diagnosis is related to both factors. The composite novice description shows a single TGV factor, two ASD factors, and three VSD factors. In this case TGV is independent of ASD and related to VSD. If an individual exhibited a knowledge structure similar to this composite novice we would expect he or she would have to combine information about these four factors to derive a diagnosis. We would expect the expert's behavior to appear more fluent as the expert would only have to combine two factors.

The results suggest that although much of the relevant symptom information is constant across analyses there are some important differences. For example, the factor which discriminates between ASD and VSD for the expert is nearly identical to that in the optimal description. The nature of murmurs which occur early in the cardiac cycle

discriminates between these diseases. The composite novice description shows a similar factor which includes information about diastolic murmurs. This suggests that novices use all murmurs to discriminate between VSD and ASD. We may label the expert factor "systolic turbulence" and the novice factor "turbulent blood flow".

The knowledge descriptions also suggest how information may be used. The information which independently specifies TGV in the optimal description discriminates TGV from ASD for the expert and VSD for the composite novice. The factor which indicates VSD in the optimal description is a combination of negative distress findings and positive EKG data. The corresponding novice factor is composed of positive distress findings and positive EKG data. The novice factor predicts ASD diagnosis. The correlation between the optimal factor and the novice factor is .95. This is an example of a minor feature difference producing a major knowledge difference.

We expected that an expert would perceive variance along a symptom factor in terms of a domain concept (such as "systolic turbulence"). When paired cases were presented to the second test expert we were surprised to find that he was unable to articulate any conceptual difference between cases other than diagnosis. When presented with a pair representing low vs high scores on the ASD-VSD discriminator the expert identified each case appropriately. When presented with cases representing moderately high vs very high scores on the ASD-VSD factor he immediately diagnosed both cases as VSD's, and, pointing to the case with the higher score, said "that one's a large one".

The expert also appropriately diagnosed cases which varied along the other factor. Two of the three pairs presented were diagnosed as ASD's and TGV's. For the third pair, the expected TGV was diagnosed as Truncus, a different admixture lesion. This difference may indicate that the "TGV-ASD" factor is actually an "Admixture-ASD" factor. An alternative explanation is that this expert's knowledge structures differ slightly from those on which the case constructions were based.

These results suggest that the factors obtained from the data provide an accurate description of the knowledge the expert uses to intuitively diagnose a patient's disease. This knowledge provides him with a way of treating large amounts of diagnostically relevant information as one or two chunks. Since none of the individual elements of the novice's knowledge sufficiently species an appropriate diagnosis, the novice must reason out a solution by combining information.

References

- Berry, D. C. (1987). The problem of implicit knowledge. Expert systems, 4(3), 144-151.
- Fodor, J. A., and Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture: A critical analysis. Cognition, 28(1988), 3-71.
- Gibson, J. J. (1986). The ecological approach to visual perception. Hillsdale, New Jersey: Lawrence Erlbaum. (Originally published 1979).
- Johnson, P. E. (1983). What kind of expert should a system be?" Journal of medicine and philosophy, 8, 77-97.
- Simon, H. A. (1969). The sciences of the artificial. Cambridge, Massachusetts: MIT Press.