

PSYCHOLOGICAL TOOLS FOR KNOWLEDGE ACQUISITION

Henry H. Rueter

Vector Research, Incorporated
P. O. Box 1506
Ann Arbor, Michigan 48106

Judith Reitman Olson

School of Business Administration
The University of Michigan
Ann Arbor, Michigan 48109-1234

Knowledge acquisition is said to be the biggest bottleneck in the development of expert systems. The problem is getting the knowledge out of the expert's head and into a computer. In cognitive psychology characterizing mental structures and why experts are good at what they do is an important research area. Is there some way that the tools that psychologists have developed to uncover mental structure can be used to benefit knowledge engineers? We think that the way to find out is to browse through the psychologist's toolbox to see what there is in it that might be of use to knowledge engineers.

Expert system developers have relied on two standard methods for extracting knowledge from the expert: (1) the knowledge engineer engages in an intense bout of interviews with the expert or experts, or (2) the knowledge engineer becomes an expert himself, relying on introspection to uncover the basis of his own expertise. Unfortunately, these techniques have the difficulty that often the expert himself isn't consciously aware of the basis of his expertise. If the expert himself isn't conscious of how he solves problems, introspection is useless.

Cognitive psychology has faced similar problems for many years and has developed exploratory methods that can be used to discover cognitive structure from simple data.

We will skip over what we call "direct" methods for knowledge acquisition. Direct methods include interviews, questionnaires, protocol analysis, interruption analysis, and inferential flow analysis. Our goal here is to expose the reader to "indirect" methods, methods which are likely to be a good deal less familiar to the practicing knowledge engineer: multidimensional scaling, hierarchical clustering, general weighted networks, ordered trees, and repertory grid analysis. But first, a few points need to be made about the variety of ways that experts can organize information.

Simple OBJECT-ATTRIBUTE-VALUE triples and forward and backward search strategies are only a few of the knowledge structures and search strategies that human experts seem to have. Expertise is primarily a skill of recognition, of "seeing" old patterns in a new problem. Chess experts, for instance, have the same limited abilities as novices to hold information for analysis; their non-chess memory abilities are not exceptional. They excel because they have hundreds of thousands of chess configurations in their heads, and because they can quickly encode the current situation into constellations of previously-seen chess patterns. The choice of candidate good moves for the expert is thus restricted to a small set of know-good moves that fit the patterns, whereas the novice has no such expert pattern knowledge to filter out bad candidates. [1, 2]

There is also evidence to suggest that experts see more richly encoded patterns than novices do. They have organized the concepts in their knowledge bases with much more depth and with many more central associations than novices. For example, in the laboratory we found that expert ALGOL programmers had much more structure in their concept relationships than did novice programmers. Furthermore, the experts' mental organizations were highly similar, whereas the novices had scattered, idiosyncratic organizations for ALGOL-specific concepts. [3]

Not only do experts have information organized in a highly structured way, they also use a variety of different kinds of knowledge structures. For instance, some things are stored in simple lists like the months of the day and the days of the week. Other information fits a table better, information such as calendar appointments and the periodic table. Some information is stored as a flow diagram, such as decision trees, for example, representing the routing of telephone messages to people who can handle them. There is information stored

in hierarchies of relationships, nested categories or clusters, such as animal taxonomies or familial relationships. Networks store richly connected language associations. Information concerning room arrangements or maps may be stored as a physical model or physical space. And, some information may be stored about a device's internal components and how they are causally related as a physical model, commonly referred to as a mental model. Thus, experts may hold what they know about objects and their relationships in many different representations, each suitable for a particular kind of reasoning or retrieval.

METHODS FOR KNOWLEDGE ACQUISITION

There are two general classes of methods for revealing what experts know. "Direct Methods" ask the expert to report on knowledge that can be directly expressed. This set of methods includes interviews, questionnaires, simple observation, thinking-out-loud protocols, interruption analysis, and inferential flow analysis. In contrast, "Indirect Methods" do not rely on experts' abilities to articulate the information that is used, or how it is used. Instead, indirect methods use other kinds of behavior, such as recall from memory or rating scales, as the basis for inferences about what the expert must have known (and, perhaps, the form in which it must have been represented) in order to produce the responses that were observed. Indirect methods include multidimensional scaling, hierarchical clustering, general weighted networks, ordered trees, and repertory grid analysis.

INDIRECT METHODS

All of the direct methods mentioned above ask the expert directly what he knows. They rely on the availability of the information to both introspection and articulation. Of course, it is not always the case that the expert has access to the details of his mental processing. In fact, it is not uncommon for experts to perceive complex relationships or come to sound conclusions without knowing exactly how they did it. In these cases, indirect knowledge elicitation methods are required.

In the following methods, experts are asked not to express their knowledge directly. Instead they are given a variety of other tasks, e.g., to rate how similar two given objects are, or to recall a collection of objects several times from prescribed starting points. From the results, the analyst then infers the underlying structure among the objects rated or recalled. All the indirect methods discussed here have been validated in experimental studies that have convincingly demonstrated their psychological validity.

To make progress, these different techniques must make assumptions about the way the data were produced. Assumptions must be made about

the nature of the mental representation: Is it physical space, lists, networks of association, tables, etc.? Furthermore, the stronger the assumptions that the analyst is able to make, the stronger the conclusions that can be made. Thus, it is important for the analyst to make a good guess about what form the expert's underlying representation is likely to take. An informed guess can be made after initial interviews with the expert, as well as from careful questioning and noting of object names and notations that the expert uses.

Of the methods to be discussed, multidimensional scaling, hierarchical clustering, and general weighted networks are the most general, in the sense that they make the weakest assumptions about the data being analyzed. These three methods can be reasonably applied to any similarity judgments, while repertory grid analysis and ordered tree analysis make strong psychological assumptions about the kind of mental structure and processes under investigation.

Multidimensional Scaling

Multidimensional scaling (MDS) is a technique that should be used only on similarity data that can be assumed to have come from stored representations of physical n-dimensional space [4]. The subject provides similarity judgments on all pairs of objects in the domain of interest. These judgments are assumed to be both symmetric and graded. This means that the similarity of A to B must be the same as the similarity of B to A (symmetry) and that there must be a continuum of possible similarity values relating A and B, not merely a simple judgment of similar or dissimilar (gradedness).

A computer program is required to perform the multidimensional scaling analysis. The result is a configuration of the objects in space. The dimensionality of the space and the metric that obtains in it are selected by the analyst, usually on the basis of trial and error.

Of course, it may not be possible to find a configuration that exactly represents the generating similarity matrix. In fact, each MDS solution has a "stress" value associated with it that provides a measure of the degree to which the computer-produced configuration and the input matrix differ. In practice, the analyst looks for the lowest stress solutions with the fewest dimensions.

The MDS technique is good for producing a diagram that the expert can later inspect and describe in more detail. It can reveal interesting clusters of objects, neighbor relations, and outlier, or "fringe" objects. One difficulty with this technique, as well as the others that we describe that require a similarity matrix, is the tedious and time-consuming process of collecting the pair-wise judgments. For n objects $n(n-1)/2$ judgments are

required, a number that soon grows quite large, even for motivated subjects.

A second difficulty with the technique is discovering the right space: in particular, the right dimensionality, the right distance metric, the right starting configuration, and the right interpretation of clusters and dimensions. Once the data are in hand performing the analysis is fairly straightforward, but interpreting the results requires some expertise.

Cluster Analysis

Like MDS, cluster analysis starts with a matrix of symmetric similarity judgments. There are many clustering algorithms, developed for many purposes, but for psychological investigations Johnson hierarchical clustering is the method of choice, because the result of this clustering technique is sensitive only to the ordinal properties of the similarity judgments and not to magnitude [5]. This insensitivity to judgment magnitude reflects the prudence required in interpreting psychological judgments.

Johnson hierarchical clustering produces hierarchical representation of the items of interest; the hierarchy take the form of a rooted tree in which the items are the "leaves." Each subtree forms a cluster and the path that connects two items in the tree is a measure of the diameter of the smallest cluster that contains them both.

Hierarchical clustering is ordinarily done using either the "minimum" method, in which the similarity between two clusters is that of the most similar items in either, or the "maximum" method, in which the similarity between two clusters is that of the least similar items in either. The minimum method tends to give long, stringy clusters, the maximum method tight, spherical ones.

General Weighted Networks

This is a third method using a symmetric similarity matrix obtained from experts' pair-wise similarity judgments. In this case there is a somewhat more theoretical basis for the analysis: We assume that in producing the judgments the expert is traversing some mental network of associations, a network in which there is a single primary path between every two items, and, for some of them, a differently encoded, secondary path as well.

The object of this method, which was developed by Schvaneveldt, et al. [6, 7], is to reconstruct the associative network through the similarity judgments. In attempting this, Schvaneveldt, et al., recently investigated the nature of expertise in airplane pilot performance using networks.

The method requires a computer and works as follows: First, a Minimal Connected Network (MCN) is formed by connecting the most similar items, then the next most similar items, etc., with arcs until there is a unique path between any two items (a minimal spanning tree). In the second stage of the analysis more links are added to the MCN to form the Minimal Elaborated Network (MEN). To form the MEN we add a link between two items to the MCN if and only if it is shorter than the path between them in the MCN.

The interpretation of the MCN and MEN involves looking for:

- (1) dominating concepts—those that have a large number of connections to many other nodes; and
- (2) members of cycles—collections of items that are fully linked in circular paths.

In their exploration of the MCN and MEN for both expert and novice pilots Schvaneveldt, et al., collected similarity judgments on a set of flying terms having to do with "split-plane concepts." The analysis of the judgments revealed:

- (1) Expert's structures are simpler than students'.
- (2) Elaborated links connected larger integrated conceptual structures.
- (3) Experts could easily identify link relations in the networks, relations such as "affects," "is-a," "desirable," and "acceptable."

The fact that the experts were so clearly different from novice fliers suggests that this GWN technique can reveal significant aspects of expertise, aspects that clearly should be encoded into an expert system.

The object of the ordered tree technique is to induce a subject's mental structure for the set of to-be-recalled items from his recall orders. The structure will be an ordered tree, that is, a tree which reflects the subject's clustering and prioritization of the items of interest.

Unlike hierarchical clustering, the ordered tree technique is based on a detailed psychological model of how the recall orders are produced by the subject: It assumes that people recall all items from a stored cluster before recalling items from another cluster. (This is the hypothesis implicit in the concept of "chunks" in memory.)

This assumption builds on data from people recalling from known (learned) organizations.

Regularities found throughout a set of orders are taken as evidence of responsible mental structure and processing. Sets of orders need not come from recall; they can be obtained simply by asking subjects to order items so that items that are related are placed together.

The computer program that conducts ordered tree analysis examines all orders for sets of item that form connected suborders. The set of all such connected item sets forms a lattice of chunks, where the elements of the lattice are ordered by set inclusion. The lattice is converted into an ordered tree structure in which a node may be marked as unidirectional (only one order of its constituents was seen), bidirectional (only one order and its inverse), or nondirectional (more than two orders observed). The program can also perform certain advanced analyses in addition, such as calculating an index of organization and looking for anomalous, or "outlier," orders, whose exclusion from the analysis yields a new tree structure with significantly more structure.

This technique has been used in a variety of studies of expert-novice differences. In [3], for example, novice, intermediate, and expert ALGOL-W programmers were asked to recall ALGOL keywords many times from many different starting points while their performance orders were recorded. Experts differed remarkably from the novices. They showed much more organization, and the similarity among the expert structures (ordered trees) was far greater than that among the novices. In [21], furthermore, the pauses between recalls of successive items was accounted for by the number of chunk boundaries crossed in the inferred memory organization. There have been a variety of studies that have used this technique to reveal organization in different domains of expertise; all have shown a convergence among experts in their mental organization of the concepts.

Repertory Grid Analysis

This technique as used in [9] is the most integrated cognitive tool for knowledge acquisition of those presented here. It includes an initial dialog with the expert, a rating session, and analyses that both cluster the objects and the dimensions on which they were rated. Essentially, it is a free-form recall and rating session in which the analyst makes inferences about the relationships among objects and the relatedness of the dimensions that the expert finds important.

Since the use of repertory grid analysis as an expert system-building tool is beautifully covered in [9] we will give only a brief outline here, referring the reader to [9] for details.

Repertory grid analysis is a technique that comes from personal construct theory, a clinical tool intended to reveal the structure of a patient's emotional system. As used in knowledge engineering the first step in the analysis is an open interview with the expert, in which some important objects in the domain of expertise are elicited. Once a set of items is available, the analyst picks sets of expertise are elicited. Once a set of items is available, the analyst picks sets of three elements and asks "What trait distinguishes any two of these objects from the third?" The expert-supplied trait identifies a "dimension" in the domain. Then the expert is asked to rate all three objects along the named dimension. This process of asking for salient dimensions for further triples continues until the analyst is satisfied that the major dimensions of the system have been uncovered.

The analyst now constructs a matrix, or grid, with objects labeling columns and dimensions labeling rows. Then the expert is asked to fill in all the missing values, so that all objects are rated on all dimensions.

It is now possible to perform a cluster analysis on both objects and dimensions, using an appropriate similarity measure between the vectors of interest. Such analyses are used to identify prototypical dimensions and items.

CONCLUSION

Just as a statistician makes judgments about the suitability of a data set to the assumptions of a proposed analysis, the knowledge engineer must make judgments of the suitability of a method for knowledge elicitation to the kinds of knowledge the expert is assumed to possess. There are a number of ways these techniques can be misapplied for scientific discovery of mental organizations. However, if used as exploratory tools, these techniques can bring a great deal of information to the knowledge engineer [10]. With them, knowledge engineers can hope to uncover more of what experts know than can be learned through interviews or introspection.

REFERENCES

- [1] Chase, W. G., and Simon, H. A. "The Mind's Eye in Chess." In W. G. Chase (Ed.), VISUAL INFORMATION PROCESSING. New York: Academic Press, 1973.
- [2] DeGroot, A. D. THOUGHT AND CHOICE IN CHESS. Paris. Mouton Co., 1965.
- [3] McKeithen, K. B., Reitman, J. S., Rueter, H. H., and Hirtle, S. C. KNOWLEDGE ORGANIZATION AND SKILL DIFFERENCES IN COMPUTER PROGRAMMERS. Cognitive Psychology, 1981, 13, 306-325.

- [4] Shepard, R. N., Romney, A. K., and Nerlove, S. B., MULTIDIMENSIONAL SCALING: THEORY AND APPLICATIONS IN THE BEHAVIORAL SCIENCES. Vol. I New York: Seminar Press, 1972.
- [5] Johnson, S. C. HIERARCHICAL CLUSTERING SCHEMES. Psychometrika, 1967, 32, 241-254.
- [6] Schvaneveldt, R. W., Durso, F. T., and Dearholt, D. W. Pathfinder: "Scaling with network structures." Memoranda in COMPUTER AND COGNITIVE SCIENCE, Report No. MCCS-85-9, Computer Research Laboratory, New Mexico State University, 1985.
- [7] Schvaneveldt, R. W. Anderson, M., Breen, T. J., Cooke, N. M., Goldsmith, T. E., Durso, F. T., Tucker, R. G., and DeMaio, J. C. STRUCTURES OF MEMORY FOR CRITICAL FLIGHT INFORMATION. Air Force Human Research Laboratory, Technical Report 81-46, 1982.
- [8] Reitman, J. S., and Rueter, H. H. "Organization Revealed by Recall Orders and Confirmed by Pauses." COGNITIVE PSYCHOLOGY, 1980, 12(4), 554-581.
- [9] Boose, J.H. EXPERTISE TRANSFER FOR EXPERT SYSTEM DESIGN. New York: McGraw Hill, 1986.
- [10] Hart, A. KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS. New York: McGraw Hill, 1986.