PERSONAL CONSTRUCT THEORY AND THE TRANSFER OF HUMAN EXPERTISE

John H. Boose

Boeing Computer Services,
7A-03, PO Box 24346, Seattle, Washington, 98124

Abstract. The bottleneck in the process of building expert systems is retrieving the appropriate problem-solving knowledge from the human expert. Methods of knowledge elicitation and analysis from psychotherapy based on enhancements to George Kelly's Personal Construct Theory are applied to this process. The Expertise Transfer System is described which interviews a human expert and then constructs and analyzes the knowledge that the expert uses to solve his particular problem. The first version of the system elicits the initial knowledge needed to solve analysis problems without the intervention of a knowledge engineering team. Fast (two hour) initial prototyping of expert systems which run on KS-300_{tm}* (an extended version of EMYCIN) and OPS5 is also performed. Conflicts in the problem-solving methods of the expert may also be enumerated and explored.

Index Key Words: *artificial intelligence, expert systems, factor elicitation, knowledge acquisition, knowledge engineering, Personal Construct Theory.*

Introduction. An *expert system* is a computer system that uses the experience of one or more experts in some problem domain and applies their problem-solving expertise to make useful inferences for the user of the system (Waterman and Hayes-Roth, 1982). This knowledge consists largely of rules of thumb, or *heuristics*. Heuristics enable a human expert to make educated guesses when necessary, to recognize promising approaches to problems, and to deal effectively with incomplete or inconsistent data.

Eliciting problem-solving knowledge from an expert is one of the critical problems in building expert systems. A long series of interview, build, and test cycles are necessary before a system achieves expert performance. The time required to build an expert-level prototype is typically six to twenty-four months. *Knowledge engineering* is the process of acquiring knowledge and building an expert system (Feigenbaum, 1977).

The goal in building the Expertise Transfer System (ETS) is to provide a tool set to significantly shorten the knowledge acquisition process, and improve the quality of the elicited problem-solving knowledge. To do this, ETS automatically interviews the expert, and helps construct and analyze an initial set of heuristics and parameters for the problem. Experts need no special training to use ETS; an initial fifteen minute explanation of the basic idea and how to use the workstation is usually all that is necessary. No initial knowledge base is necessary.

ETS is capable of automatically producing KS-300_{tm} and OPS5 (Forgy, 1981) knowledge bases. KS-300_{tm} is an extended version of EMYCIN developed by Teknowledge, Incorporated, of Palo Alto, California. EMYCIN, an expert system building tool, was extracted from MYCIN, developed at Stanford University (Shortliffe, 1976; van Melle et al., 1981).

In-depth analyses of ETS's knowledge base listings and reports by both the expert and knowledge engineers provide problem-solving information prior to any human interviewing. Most of the initial interviewing process is eliminated, which is typically the most painful and time-consuming part of the knowledge acquisition process.

First, aspects of Personal Construct methodology will be discussed which are relevant to the Expertise Transfer System. Then, the system itself is described, along with its relation to the knowledge engineering process. Finally, results and limitations of the methodology are discussed.

Personal Construct Theory. ETS employs clinical psychotherapeutic interviewing methods originally developed by George Kelly who was interested in helping people categorize experiences and classify their environment. A person can not only use this organization to predict events more accurately and act more effectively, but can also change the organization to fit specific perceived needs (Shaw, 1981). George Kelly's (Kelly, 1955) theory of a *personal scientist* was that each individual seeks to predict and control events by forming theories, testing hypotheses, and weighing experimental evidence.

Certain techniques for use in psychotherapy were developed by Kelly based on this philosophy. In a Repertory Grid Test for eliciting role models, Kelly asked his clients to list, compare, and rate role models

*KS-300_{tm} is a trademark of Teknowledge, Inc., of Palo Alto, California.

to derive and analyze character traits. Aspects of these role models were used to build a *rating grid.* A non-parametric factor analysis method was then used to analyze the grid (Kelly, 1963). The results helped Kelly and his client understand the degree of similarity between the traits. He named a trait and its opposite a *construct,* and hypothesized that each construct represented some internal concept for the client.

Following construction and analysis of the grid, the clinician entered an interviewing phase. Typically, in this phase, the interviewer would attempt to help the subject expand on and verify the relationships between concepts pointed out by the grid analysis.

One interviewing technique was known as *laddering*. This was a method which helped connect the elicited concepts in their *superordinate* and *subordinate* relationships by asking the client "how" and "why" questions.

Hinkle (Hinkle, 1965) developed a taxonomy of implication types. He suggests *ambiguity* may arise when a subject has an incomplete abstraction of the differences between the contexts in which the concept was used, or a subject uses one concept label for two independent traits. He also felt that the processes of psychological movement, conflict resolution, and insight depend on locating and resolving such points of ambiguous implication into *parallel* or *orthogonal* forms using techniques similar to laddering.

More recently, elicitation and analysis of repertory grids has been made available through interactive computer programs (Shaw, 1979). A variety of grid analysis techniques using distance-based measures between vectors - either rows or columns of the grid - have been used, in which both elements and constructs may be graphically compared by the subject to find similarities and differences. Some of these techniques include principal components analysis (INGRID, Slater, 1977), a Q-Analysis of the grid in a cluster-analyzed hierarchical format (QARMS, Atkin, 1974), and a linear cluster analysis (FOCUS, Shaw, 1980).

A more formal description of implication relationships is presented by Gaines which is based on logic (Gaines and Shaw, 1981). Instead of looking at grid rows and columns as vectors in space, Gaines views them as assignments of truth-values to logical predicates. In binary rating systems such as those used in Kelly's original grid methodology, an "X" would simply mean true, and a blank would mean false. A grid, then, can be seen as a matrix of truth values. Gaines goes on to show a method of deriving implications from grids which use rating scales rather than binary scales. The method is based on multi-valued logics (Rescher, 1969) using fuzzy set theory (Zadeh, 1965). Using this method, the implication relation can be extended to include *implication strength.* The program ENTAIL achieves this and produces graphs which show entailment relations among constructs and elements (Gaines, 1981).

<u>Human and Computer Interviewing.</u> It is almost always difficult for the expert to articulate problem-solving knowledge in terms which can be utilized by an expert system. Human interviewing processes elicit knowledge which is incomplete, inconsistent, and imprecise. The knowledge is often subconscious, and the expert may not be reliable when introspecting about problem-solving. The expert must come to trust the interviewer enough to overcome any fears or insecurities felt about the expert system building process. He may feel insecure about losing his job, or feel threatened by the encroachment of computers into his private domain, or he may not want to subject his problem-solving methods to the scrutiny of other human experts. Gaines points out that using a computer to interview subjects alleviates many of these difficulties (Gaines and Shaw, 1981).

<u>Expertise Transfer System.</u> In an effort to apply grid methodology techniques to knowledge acquisition, the Expertise Transfer System (ETS) has been developed. ETS runs in Interlisp-D on a Xerox 1100 Dolphin Lisp Machine, using the high-resolution bitmap windows and mouse interaction capabilities provided.

In the following example session, an expert will attempt to build a knowledge base for a Database Management System Advisor. The completed expert system would be able to advise a software engineer as to which database management system to use for an application problem.

First, the Expertise Transfer System elicits conclusion items from the expert. Kelly referred to these items as *elements.* An expert system would be expected to recommend some subset of these items based on a given set of problem characteristics. In this case, the elements consist of all the databases which the expert believes the expert system should be knowledgeable about (see Figure 1).

```
Do you already have a list of things which you wish to classify?
ETS** YES
Please enter a list of database types, one to a line. When you're
done, enter STOP, or a RETURN.
ETS** IMS
ETS** SYSTEM-2000
ETS** INQUIRE
ETS** IDMS
ETS** TOTAL
ETS** SQL/DS
ETS** MODEL-204
ETS** RAMIS-II
ETS** DMS-170
ETS** RIM
ETS** ORACLE
ETS** INGRES
ETS** ADABAS
ETS** SIR
ETS** EASYTRIEVE
ETS** CREATABASE
```

Figure 1. ETS Eliciting Elements from the Expert - List Mode.

If the expert can not verbalize the set of conclusion elements initially, ETS enters an incremental interview mode of operation based on a program called DYAD (Keen et al., 1981), where elements are elicited one at a time, based on differences between and similarities to other elements (see Figure 2).

```
What is the name of a database you'd like to consider?
ETS** CREATABASE
What is the name of another database to consider?
ETS** EASYTRIEVE
What is the name of a third database to consider?
ETS** SIR
Think of an important characteristic that two of CREATABASE,
EASYTRIEVE, and SIR share, but that the other one does not. What
is that characteristic?
ETS** TEXT RETRIEVAL
What is that characteristic's opposite as it applies in this case?
ETS** NOT TEXT RETRIEVAL
Please rate these things on a scale of 5 to 1, where 5 means more like
TEXT RETRIEVAL and 1 means more like NOT TEXT
RETRIEVAL.
    (CREATABASE) ** 5
    (EASYTRIEVE) ** 5
    (SIR) ** 1
What is the name of another database which you feel is different
from SIR in some important attribute?
ETS** ADABAS
What is that attribute?
ETS** INVERTED    . . .
```

Figure 2. ETS Element and Construct Elicitation - Incremental Mode.

After the expert has listed the database management systems to be considered, ETS asks him to compare successive groups of three databases, and name an important trait and its opposite which distinguishes two members of this triad from the third one (see Figure 3). The result of this first phase of the interview process is a list of elements to be classified, and a list of classification parameters, all of which were derived from the expert.

```
Think of an important characteristic that two of CREATABASE,
EASYTRIEVE, and SIR share, but that the other one does not. What
is that characteristic?
ETS** TEXT RETRIEVAL

What is that characteristic's opposite as it applies in this case?
ETS** NOT TEXT RETRIEVAL

Think of an important attribute that two of EASYTRIEVE, SIR, and
ADABAS share, but that the other one does not. What is that
attribute?
ETS** INVERTED

What is that attribute's opposite as it applies in this case?
ETS** NOT INVERTED

Think of an important trait that two of SIR, ADABAS, and INGRES
share, but that the other one does not. What is that trait?
ETS** RUN ON VAX

What is that trait's opposite as it applies in this case?
ETS** DO NOT RUN ON VAX
```

Figure 3. Problem Trait (Construct) Derivation.

As Kelly points out, an initial set of constructs will probably not be a sufficient window into an individual's construct system. Later in the interviewing process, Kelly's technique of laddering is used, as well as construct volunteering and further triad formation to expand the construct network. So far, these techniques have been sufficient for building rapid prototype systems with reasonable behavior. Laddering may also be continued later on in the manual interviewing process as the expert and knowledge engineers work together to refine the knowledge base.

Next, the system asks the expert to rate each element against each construct (see Figure 4), thereby forming a rating grid (see Figure 5). In addition to allowing numerically scaled ratings, ETS accepts the ratings "N" - neither trait applies - and "?" - both traits apply - as described by Landfield (Landfield, 1976).

```
Please rate these things on a scale of 5 to 1, where 5 means more like
THEY ARE HIERARCHICAL and 1 means more like NON
HIERARCHICAL. If neither one seems to apply, just enter an "N." If
both seem to apply equally, enter a "?" (a-XOtl)
 THEY ARE HIERARCHICAL(5) NON HIERARCHICAL(1)
   (CREATABASE) ** 1
   (EASYTRIEVE) ** 1
   (SIR) ** 5
   (ADABAS) ** 1
   (INGRES) ** 1
   (ORACLE) ** 1
   (RIM) ** 1
   (DMS-170) ** 1
   (RAMIS-II) ** 3
   (MODEL-204) ** 3
   (SQL/DS) ** 1
   (TOTAL) ** 3
   (IDMS) ** 3
   (INQUIRE) ** 1
   (SYSTEM-2000) ** 4
   (IMS) ** 5
```

Figure 4. Rating Constructs with Elements.

Once this grid is established, several analytic methods are invoked to structure the knowledge. A non-parametric factor analysis is performed to cluster similar constructs into constellations (Kelly, 1963). Then, an entailment graph of implication relationships is built using the methodology of ENTAIL (see Figure 5). The constellation analysis shows constructs which are nearly functionally equivalent, given the elements chosen by the expert.

ETS and Conflict Analysis. The entailment graph shows implication relationships between various halves of the constructs. For instance, "RUN ON VAX," implies both "NOT TEXT RETRIEVAL" and "NOT INVERTED." These graph arcs should correspond with paths in the expert's construct system. However, as Kelly points out, logical representations do not necessarily correspond with a person's internal construct hierarchy. Typically, the expert is surprised at many of the relationships which are revealed by the graph.

In this instance, the expert initially disagreed with several of the relationships. ETS helps the expert explore them. When the expert points at an arc and presses a mouse button, the corresponding ratings which produced that entailment are highlighted on the grid. The expert may edit these ratings until he is satisfied with their values. If this does not correct the perceived problem shown by the entailment graph, ETS asks the expert if he can think of any elements which would be exceptions to the entailment. If the expert can think of an exception, then that element is added to the knowledge base, rated against all the constructs, and the entailment graph is regenerated. If this still does not correct the problem, ETS asks the expert if he would like to refine the constructs involved in the implication relation. This involves breaking up one of the constructs into two or more new constructs. To do this, ETS invokes a simple laddering method which asks "why?" and "how?" questions concerning the constructs. New constructs are added to the knowledge base, and the rating grid and entailment graph are regenerated.

If the entailment still exists on the graph, the expert may indeed agree that the entailment relationship is sound, and that he just never thought about the problem characteristics in that manner before. In this way, ETS helps the expert structure the problem-solving knowledge based on his own operative internal construct network. On the other hand, if the entailment relationship still exists, and the expert still disagrees with it, then this represents an inconsistency in the way the expert thinks about the problem. In effect, what has happened is that ETS has captured an important internal conflict in the expert's construct hierarchy.

Ambiguous construct relations also point out internal conflicts. When the expert is finished correcting entailment arcs, ETS searches for ambiguous relations, and again invokes the laddering method to try and refine these points of conflict into parallel or orthogonal forms.

Both ambiguous relations and relations with which the expert disagrees are important points of conflict in the expert's problem-solving methods. These may be resolved with ETS or in later discussions between the knowledge engineers and the expert. The process of resolving them involves "psychological movement, conflict resolution, and insight" (Hinkle, 1965). These are points of interest in which further exploration is necessary both in producing the expert system, and in refining the expert's own problem-solving processes. A set of conflict points is generated as part of the knowledge base report listings.

Rule Generation. After the entailment graph has been constructed, ETS generates two types of heuristic production rules: *conclusion rules* and *intermediate rules*. Each production rule is generated with a belief strength or *certainty factor*. Certainty factors are used to represent a relative *strength of belief* which the expert would associate with the conclusion of the rule. Once generated, all rules may be reviewed and modified by the expert.

Conclusion rules are created from individual ratings in the grid. Each rating has the potential for generating a rule. The expert is first asked to rate the relative importance of each construct in terms of its potential importance in solving the problem. Then, ETS employs an empirical algorithm to generate certainty factors for each rule. The algorithm takes into account grid ratings, relative construct importance, and the certainty factor combination

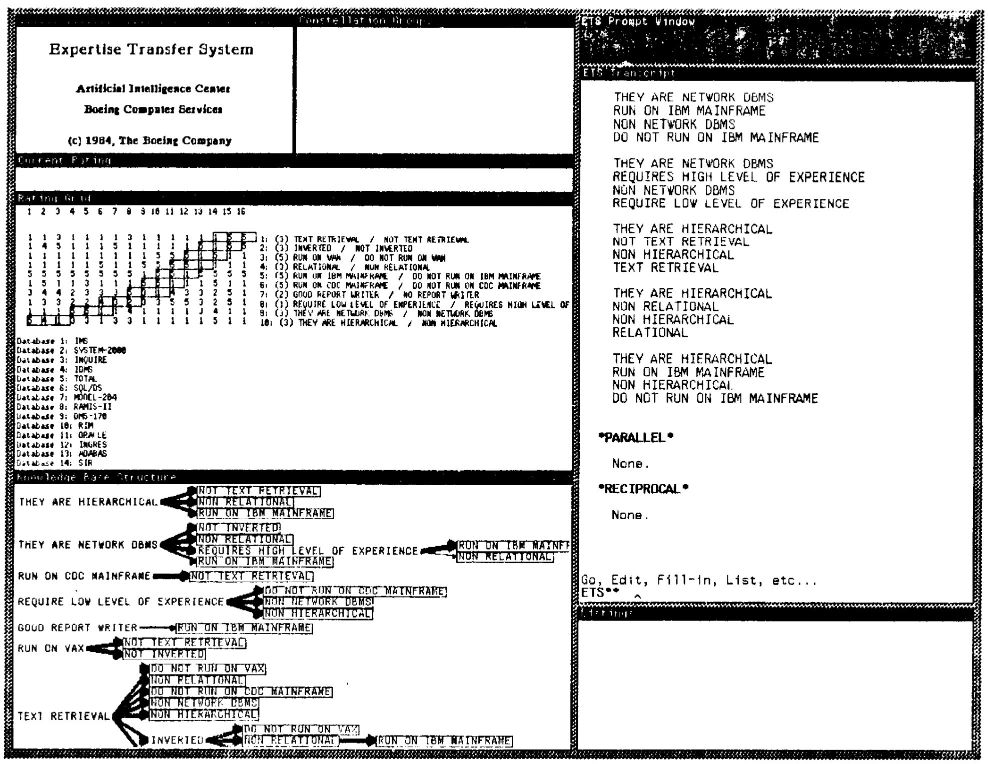

Figure 5. Screen Snapshot of ETS Showing Rating Grid and Entailment Graph.

algorithm in the target expert system building tool.

Intermediate rules are based on relations in the entailment graph. For each entailment, one rule is generated. The strength of the rule's certainty factor is based on the relative strength of the entailment. These rules generate intermediate pieces of evidence at a higher conceptual level than those of conclusion rules.

Multiple Knowledge Representations. Allowing the expert to work with multiple forms of his stated problem-solving knowledge is an important aspect of ETS. Lenat (Lenat, 1982) argues that knowledge representations should shift as different problem-solving needs arise. Each different representation method in ETS potentially helps the expert think about the problem in a new way, and tends to point out conflicts and inconsistencies over time. Rather than trying to force the expert to eradicate inconsistencies, this methodology takes advantage of the important psychological and problem-solving aspects of inconsistencies by helping the expert explore them.

Knowledge Expansion. In addition to exploring conflicts to add new elements or constructs, the knowledge base may be expanded in a number of ways. Information may be modified and volunteered. New element triads may be created for comparison. Incremental interviewing may be continued, and laddering may be invoked to expand construct hierarchies.

Listings and reports generated by the system are useful in later manual interviewing phases of the knowledge engineering process. The knowledge engineering team does not need to begin from scratch when beginning discussions with the expert. They have basic vocabulary, important problem traits, an implication hierarchy of these traits, and conflict areas where discussions may begin. This has been an important aid in streamlining the knowledge acquisition process in building expert systems. Knowledge engineers may also use associated interviewing techniques from personal construct methodology such as laddering and the resolution of ambiguous construct relationships.

Testing - Rapid Prototyping of an Expert System. Once the rules have been generated, ETS has enough information to automatically generate a knowledge base for an expert system building tool based on production rules. Currently, ETS can generate knowledge bases for KS-300$_{tm}$ and OPS5. Consultations are then run from these prototypes to test the knowledge base for necessity and sufficiency. An example consultation using the knowledge base generated for the Database Management Advisor is shown in Figure 6.

Manual interviewing and incremental knowledge refinement are still necessary to produce a system that performs at an expert level. The initially

```
--------DATABASE-1--------
 1) What is the name of DATABASE-1
    for this specific application?
 ** TEST 7
 2) What is the REPORT-WRITER attribute
    of DATABASE-1 (GOOD-REPORT-WRITER NO-
    REPORT-WRITER)?
 ** GOOD-REPORT-WRITER
 3) What is the NETWORK attribute of
    DATABASE-1 (THEY-ARE-NETWORK-DBMS
    NON-NETWORK-DBMS)?
 ** NON-NETWORK-DBMS
 4) What is the EXPERIENCE trait of
    DATABASE-1 (REQUIRE-LOW-LEVEL-OF-
    EXPERIENCE REQUIRES-HIGH-LEVEL-OF-
    EXPERIENCE)?
 ** REQUIRES-HIGH-LEVEL-OF-EXPERIENCE
          .
          .
          .
The values of DATABASE-1 are as follows:
    IMS (.89).
    IDMS (.8).
    MODEL-204 (.79).
    SQL/DS (.78).
    RAMIS-II (.77).
    ADABAS (.73).
    INQUIRE (.73).
    CREATABASE (.67).
    TOTAL (.66).
    EASYTRIEVE (.61).
    SYSTEM-2000 (.5).
    DMS-170 (.42).
    ORACLE (.37).
    INGRES (.37).
Debug/review phase, Rules, Parms, or other option (? for help)
```

Figure 6. Rapid Prototyping: A KS-300$_{tm}$ Knowledge Base Developed in Two Hours

generated knowledge base may not be similar in structure to later ones. However, fast prototyping can be used to help analyze the sufficiency of the initial knowledge base.

Other Efforts in Developing Knowledge Acquisition Systems. ETS could be used in combination with any of the knowledge acquisition tool families described below as a front-end processor to elicit initial traits and heuristics.

TEIRESIAS (Davis and Lenat, 1982) is a subsystem of EMYCIN, which aids the expert and knowledge engineers when they attempt to refine an existing knowledge base. ETS can be used to supply the initial knowledge base since TEIRESIAS is not capable of eliciting such information on its own.

META-DENDRAL (Buchanan and Feigenbaum, 1978) and AQ11 (Michalski, 1980) both perform classification analyses of training examples from their respective knowledge bases in order to produce generalized rules using inductive inference strategies. META-DENDRAL learns rules that predict how classes of compounds fragment in a mass spectrometer, and AQ11 formulates rules from traits and test cases. Both of these systems need an initial set of problem traits before classification can begin; it is up to the expert and knowledge engineer to produce the initial list of applicable traits and relevant training examples. Again, ETS methodology could be used as a front-end for these systems to elicit the problem characteristics.

NANOKLAUS (Hass and Hendrix, 1981) attempts to elicit a classification hierarchy from an expert through a natural language dialog. The information is then used for certain classes of deductive data retrieval. The expert uses NANOKLAUS to enter *IS-A* hierarchy relations and object descriptions. The expert needs to have such a hierarchy in mind before using NANOKLAUS. ETS's methods could be combined with this system to help elicit initial relevant concepts in terms of constructs derived from objects, and to produce heuristic rules.

Discussion. Over sixty prototype systems have been built with ETS, ranging over a variety of problems. Although no formal studies have yet been performed, knowledge engineers associated with these projects feel that anywhere from two to five calendar months of knowledge acquisition time are saved using ETS. Experts claim that these systems generally demonstrate reasonable behavior, although it is expected that "expert" expert systems can only be built by using ETS in combination with more traditional methods.

Frequently, an expert begins using ETS without a clear problem structure in mind. A "false start" might occur, for example, when the expert gives conclusion items at varying levels of abstraction. ETS forces the expert to deal with this problem during item triad comparison. After spending fifteen or twenty minutes trying to generate traits for these items, the expert realizes the nature of the expected responses, and starts over again. In these and similar processes, the expert is trained to think in terms that are useful for problem-solving using production systems.

As a rule, experts are enthusiastic about using ETS. Typically, an expert will want to use the system again the following day after having had a chance to think about the problem in the system's terms. This enthusiasm is important in starting projects quickly.

ETS is best suited for analytic problems whose solutions may be based on production systems. The system can not readily handle synthesis class problems or problems which require a combination of analysis and synthesis. However, ETS can handle the analytic portions of these problems, and it should be noted that many planning and design problems involve synthesizing the results of several analytic components (eg., R1, in McDermott, 1980a and 1980b). These components may be investigated with ETS.

It is difficult to apply grid methodology to elicit deep causal knowledge, procedural knowledge, or strategic knowledge, although some alternate forms of interviewing techniques are being explored in this area for use with ETS. For instance, the expert may be asked for problem-solving strategies rather than conclusion items, and traits of these strategies could then be elicited.

One assumption of grid methodology (Kelly, 1955) is that the elicited set of elements will be a sufficient representation of the problem conclusion set. It must be assumed that the expert knows what these conclusions are, or that the relevant set will be built with ETS knowledge expansion methods or subsequent manual interviewing.

It is difficult to verify that a sufficient set of constructs have been elicited. Inappropriate constructs are relatively easy to weed out of the system, but errors of omission are harder to detect. Some important constructs which are missing may be elicited using ETS's knowledge base expansion techniques, but there is no guarantee that a sufficient set will be found. This is a problem with knowledge acquisition in general. Expert-level performance of the final expert system is critically dependent on obtaining and effectively using a sufficient set of problem-solving knowledge.

Many enhancements are being considered to improve ETS's utility. These include expansion of the interview methods, inclusion of more analytic tools to identify the relative importance and validity of elicited constructs and elements (such as in Hinkle, 1965, and principle components analysis), and development of feedback paths between ETS and the target expert system building tool. Other psychological techniques such as multi-dimensional scaling are being analyzed. A knowledge engineering guide, illustrating the use of ETS methodology and its associated manual interviewing techniques, is also being prepared.

In conclusion, ETS and its related methods have been useful aids during knowledge engineering serving to streamline the knowledge acquisition process.

Acknowledgements. Thanks to Roger Beeman, Keith Butler, Alistair Holden, Earl Hunt, Art Nagai, Steve Tanimoto, Lisle Tinglof, Rand Waltzman, and Bruce Wilson for their contributions and support. This work was performed at the Artificial Intelligence Center of Boeing Computer Services in Seattle, Washington.

References

Atkin, R. H., Mathematical Structure in Human Affairs, London: Heinemann, 1974.

Bannister, D., and Mair, J. M. M., The Evaluation of Personal Constructs, Academic Press, 1968.

Barstow, D. R., Aiello, N., Duda, R., Erman, L., Forgy, C., Greiner, R.,Lenat, D. B., London, P., McDermott, J., Nii, H. P., and Weiss, S., "Languages and Tools for Building Expert Systems," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (eds.), Building Expert Systems, Addison-Wesley, 1983.

Buchanan, B. G., and Feigenbaum, E. A., "DENDRAL and META-DENDRAL: Their Applications Dimension," *Artificial Intelligence*, 11, 1978.

Buchanan, B. G., Barstow, D., Bechtal, R., Bennet, J., Clancey, W., Kulikowski, C., Mitchell, T. M., and Waterman, D. A., "Constructing an Expert System," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (eds.), Building Expert Systems, Addison-Wesley, 1983.

Davis, R., and Lenat, D. B., Knowledge-Based Systems in Artificial Intelligence, New York: McGraw-Hill, 1982.

Feigenbaum, E. A., "The Art of Artificial Intelligence: I. Themes and Case studies of Knowledge Engineering," Proceedings, *Fifth*

International Joint Conference on Artificial Intelligence, Massachusetts Institute of Technology, 1977.

Forgy, C. L., OPS5 User's Manual, Department of Computer Science, Carnegie-Mellon University, 1981.

Gaines, B. R., and Shaw, M. L. G., "New Directions in the Analysis and Interactive Elicitation of Personal Construct Systems," in M. Shaw (ed.), Recent Advances in Personal Construct Technology , New York: Academic Press, 1981.

Hass, N., and Hendrix, G., "Learning by Being Told: Acquiring Knowledge for Information Management," in R. S. Michalski, J. Carbonell, and T. M.Mitchell (eds.), Machine Learning: An Artificial Intelligence Approach, Palo Alto, Calif: Tioga Press, 1983.

Hinkle, D. N., The Change of Personal Constructs from the Viewpoint of a Theory of Implications, Ph.D. Dissertation, Ohio State University, 1965.

Keen, T. R., and Bell, R. C., "One Thing Leads to Another: A New Approach to Elicitation in the Repertory Grid Technique," in M. Shaw (ed.), Recent Advances in Personal Construct Technology , New York: Academic Press, 1981.

Kelly, G. A., The Psychology of Personal Constructs, New York: Norton, 1955.

Kelly, G. A., "Non-parametric Factor Analysis of Personality Theories," *Journal. of Individual Psychology*: 19, 1963.

Landfield, A., "A Personal Construct Approach to Suicidal Behavior," in Slater, P. (ed.), Dimensions of Intrapersonal Space, Vol. 1, London: Wiley, 1976.

Lenat, D. B., "The Nature of Heuristics," *Artificial Intelligence* 19, 1982.

McDermott, J., "R1: An Expert in the Computer Systems Domain," in *AAAI* 1, 1980a.

McDermott, J., "R1: An Expert Configurer," Rep. no. CMU-CS-80-119, Computer Science Department, Carnegie-Mellon University, Pittsburgh, Pa., 1980b.

Michalski, R. S., "Pattern Recognition as Rule-guided Inductive Inference," *IEEE Transactions of Pattern Analysis and Machine Intelligence* 2, no. 4, 1980.

Rescher, N., Many-Valued Logic, New York: McGraw-Hill, 1969.

Shaw, M. L. G., and Gaines, B. R., "Externalizing the Personal World: Computer Aids to Epistemology," in Improving the Human Condition: Quality and Stability in Social Systems, Kentucky: Society for General Systems Research, 1979.

Shaw, M. L. G., and Gaines, B. R., "Fuzzy Semantics for Personal Construing," in Systems Science and Science, Kentucky: Society for General Systems Research, 1980.

Shaw, M. L. G., and McKnight, C., "ARGUS: A Program to Explore Intra-Personal Personalities," in M. Shaw (ed.), Recent Advances in Personal Construct Technology, New York: Academic Press, 1981.

Shortliffe, E. H., Computer-Based Medical Consultants: MYCIN, New York: Elsevier, 1976.

Slater, P. (ed.), Dimensions of Intrapersonal Space, Vol. 2, London: Wiley, 1977.

van Melle, W., Shortliffe, E. H., and Buchanan, B. G., "EMYCIN: A Domain-Independent System that Aids in Constructing Knowledge-Based Consultation Programs," *Machine Intelligence*, Infotech State of the Art Report, Series 9, No. 3, 1981.

Waterman, D. A., and Hayes-Roth, F., An Investigation of Tools for Building Expert Systems, R2818-NSF, Rand Corporation, 1982.

Zadeh, L. A., "Fuzzy Sets," Information and Control, 8, 1965.