Knowledge Acquisition Opportunities Resulting from Verification, Validation and Maintenance

Daniel E. O'Leary
Graduate School of Business
University of Southern California
Los Angeles, CA 90089-1421
oleary@mizar.usc.edu

ABSTRACT

This paper argues that verification, validation and maintenance (VV&M), and knowledge acquisition, are not independent activities in the expert system life cycle. Instead VV&M offer many knowledge acquisition opportunities, resulting from the identification of different errors, anomalies and changes.

A model of human learning is proposed as the basis of the analysis for the relationship between VV&M and knowledge acquisition for knowledge-based systems. In that model feedback is found to be a necessary condition for learning i.e., knowledge acquisition. VV&M provides the feedback function in that model.

1. Introduction

The purpose of this paper is to investigate the relationship between knowledge acquisition and other "phases" of the expert system life cycle, including verification, validation and maintenance (VV&M). This is done by first noting that the important life cycle model for knowledge-based system is the human model of learning. It is only after the equivalent activity of learning is accomplished that an appropriate system will have been developed. In that model of human learning, feedback is a necessary condition for learning, i.e., knowledge acquisition. Thus, acquisition and feedback are tightly integrated processes. In systems development that feedback is provided using VV&T. Second, the nature of that feedback process is investigated for its ability to elicit knowledge. Different aspects of VV&M are found to provide different types of knowledge acquisition opportunities.

1.1 Dependence of Verification, Validation and Maintenance, and Knowledge Acquisition

Virtually all treatments of the issues of knowledge acquisition, verification, validation and maintenance of expert systems have viewed each of those activities as distinct portions in the expert system life cycle. This paper argues that those activities are not independent. Instead the basis of much knowledge acquisition is in the errors, anomalies and changes determined by or required for VV&M. This suggests that systems which have received only limited VV&M are also systems for which the knowledge acquisition process is incomplete. In addition, the paper suggests that different errors, anomalies and changes offer the opportunity for the acquisition of different types of knowledge. Thus, identification and characterization of errors, anomalies and changes can give insight into the knowledge that can be acquired. Such insights can guide the knowledge acquisition process.

The processes of VV&M are integrated with KA by functioning as feedback devices. As such they provide knowledge to the knowledge acquisition process about the quality of the acquired knowledge and its representation. In particular, VV&M feedback information about "problems" encountered.

1.2 Problems in Knowledge Bases: Errors, Anomalies and Changes

In the VV&M literature, there is reference to "... checking the knowledge base for existing and potential problems" (Nguyen et al. [1987, p. 69]. There are a number of "problems" that can occur in the development and use of knowledge bases. For exposition purposes those problems are characterized as errors, anomalies and changes.

Errors occur under a number of circumstances. For example, knowledge may be represented incorrectly or incorrect knowledge may be placed into the system.

Anomalies include unusual circumstances that fall outside the scope of the system's knowledge that result in the system indicating an incorrect answer or the system's inability to provide an answer. They are anomalies since the system's model of the world either does not account for them correctly or cannot account for them. The inability of the system to account for the anomalies indicates the need for additional knowledge or changing the existing knowledge.

Changes include updates to the knowledge base to reflect changes in the state of the world. The dynamic nature of the world requires that systems have updates to their knowledge. That is reflected in changes to the knowledge base.

The distinction between errors, anomalies and changes is somewhat superficial. The need for a "change" in the knowledge base may be seen as an "anomaly." An "anomaly" may be seen as an "error" in the knowledge base. An "error" may result in "anomalous performance." However, the distinction between these terms is helpful in eliciting the nature of problems associated with the use of a knowledge base.

1.3 Contributions

The primary contribution of this paper is to contend that those "problems" can indicate a need for additional knowledge acquisition: VV&M feedback to the KA process. In addition, the nature of the "problem" found using VV&M can indicate, to a certain extent, the type of knowledge that needs to be acquired.

This paper also contributes to the literature that is aimed at the factors associated with knowledge acquisition. Previous researchers have found that domain and the type of knowledge play an important part in the knowledge acquisition process. For example, Boose [1989] found a number of domain specific tools for knowledge acquisition (KA). Further, Guber [1990] discussed the acquisition of a specific type of knowledge: strategic knowledge.

1.4 This Paper

This paper proceeds as follows. Section 2 discusses some of the relationships between KA and VV&M. Sections 3, 4, and 5 analyze the types of knowledge acquisition facilitated by verification, validation and maintenance, respectively. Section 6 provides a brief summary of the paper.

2. On the Relationship between KA and VV&M

The close relationship between knowledge acquisition and verification, validation and maintenance results from at least two factors. First, VV&M provide feedback to the KA process. Feedback is an integral part of models of human behavior and necessary for learning. Second, at some level each of KA and VV&M employ and exploit knowledge about the form of knowledge representation, e.g., rules.

2.1 Feedback

The generally accepted model of human judgement-action learning situations (e.g., Hogarth [1989]) includes a pivotal role for the process of feedback (figure 1). In that model, "outcomes" of events and "evaluation of outcomes" are fed back to both "judgement, predictions and hypotheses," and to "actions, choices and experiments." As noted by Hogarth [1989, p. 128], "In short, an understanding of the learning of relations used in predictive judgement necessitates an understanding of the task environment, actions, outcomes, coding of outcomes in memory, interpretation of outcomes in memory and the different feedback loops." Feedback is a necessary condition for learning, i.e., acquiring knowledge.

In the development of an expert system, the life cycle mirrors human learning since the process of building an expert system requires learning the necessary knowledge in order to place it in the system. Thus, feedback must play a critical part in that development process. That feedback is accomplished by using VV&M.

[Picture]

Figure 1 - Model of Human Judgment (from Hogarth [1989, p.127])

2.2 Knowledge Representation

A portion of the learning model in figure 1 is "coding, storage and retrieval." In the development of an expert system this activity is, in part, captured with the process of knowledge representation. For example, acquired knowledge is represented as rules or frames, etc.

The structure of the knowledge ("coding, storage and retrieval") impacts the evaluation in figure 1 and in the processes of VV&M. Choice of a given structure defines a set of verification processes. For example, in rule-based systems typically the rules-base should not contain cycles, i.e., there should not be circular reasoning. In addition, any maintenance to the system should be done so as to not introduce any cycles into the system.

3. Verification

Verification has been defined as the determination of the consistency, correctness and completeness of the computer program (Adrion et al. [1985]). Nguyen et al. [1987], among others, have developed a number of tests of rule_based systems, under the issues of consistency, correctness and completeness. Throughout the discussion will focus on the verification findings of rules, although a similar set of findings could be established for other forms of knowledge representation, such as frames or objects.

3.1 Consistency

Conflicting Rules. One of the primary issues of consistency is ensuring that there is no conflicting knowledge. In the case of two conflicting rules, additional knowledge acquisition is required to determine if any of the conflicting rules are correct or incorrect. The conflict provides an hypothesis that knowledge acquisition will resolve.

Redundant Rules. Another issue is the potential existence of redundant rules. There are a number of explanations for redundant rules. First, the rules are just redundant, one does not belong in the group and it was put there by accident. Second, one of the redundant rules should be a different rule, and an error was made. Third, the expert from whom the rules are being solicited is not aware of the fact that there are redundant rules in the model that is being used and it is only through the process of mapping the rules into a rule-base that such an awareness occurs. Finding the redundancy helps determine which of these alternatives is appropriate, thus, guiding knowledge acquisition.

Subsumed Rules. (If A and B then C; If A then C). Subsumed rules can indicate uncertainty as to which model is appropriate or a need for the use of uncertainty in the rule-base. For example, subsumed rules are appropriate if there are certainty factors associated with the two different rules. Subsumed rules may indicate that one of the rules is in error because of too much generality or too much specificity. In either situation, additional knowledge acquisition, focusing on generality or specificity, is required to determine which is appropriate.

3.2 Correctness

Circularity. As part of the correctness of the knowledge in a rule-based system, typically, it is necessary to ensure that there are no cycles in the rules (A _ B _ C _ A). If one is found then one of the arcs must be removed or else the underlying network is incongruent with the requirement of an acyclic rule network. However, what does it mean if a cycle is found? It can indicate that the system designer or expert from whom the knowledge was solicited just made an error. However, it may also mean that the expert's model has a cycle in it, in which case the expert's model still has some unresolved ambiguity in it. In any case, it is unclear which arc must be removed. That ambiguity indicates the area where knowledge acquisition is required.

Dead End Goals and Dead End If Conditions. Dead End Goals are associated with rules that do not lead to other rules or do not lead to goals identified as ultimate system decisions. Dead end If Conditions are the conditions that cannot be arrived at by any other rule. These problems point to possible "gaps" of knowledge that are missing from the knowledge base. Such knowledge gaps may be missing from the expert's model of the world, may represent unsolicited knowledge or may be errors in the development of the knowledge base. In either situation additional knowledge acquisition is necessary to find those chunks of knowledge needed to fill the gaps.

3.3 Completeness

Illegal Value Referenced. Given the list structure discussed above, if there is an attempt to use a concept that is not on the list that may indicate that there is a need to update the system to include the additional knowledge implicit in the omitted item. Alternatively the incomplete list may mean an incomplete understanding of the knowledge. In either case the source of omission guides the knowledge acquisition.

Unreferenced Value. Given the list structure, there could be a situation where there is an attribute on the list that no rule references. In that situation, the question becomes should there be a rule that employs that attribute or should the attribute be on the list. In either case, this situation will guide the knowledge acquisition process.

4. Validation

Validation is more concerned with ensuring that we build the "right system" (e.g., O'Keefe et al. [1987]). In addition, O'Leary [1987, p. 468] indicates that validation requires ascertaining "... what the system knows, does not know, or knows incorrectly." From these definitions we can anticipate the close relationship of validation with knowledge acquisition. Ensuring it is the right system means ensuring that we have acquired the appropriate knowledge. Ascertaining what the system knows, does not know or knows incorrectly, means being able to acquire the appropriate knowledge and know when you have not.

Although there are a number relationships between KA and validation, this section will explore only a few of those interfaces where portions of the validation process lead to additional knowledge acquisition. In particular, content validity tests, different actors in the validity process and reliability are analyzed in some detail.

4.1 Content Validity: Tests of the System as a Basis for Knowledge Acquisition

Knowledge acquisition arising from content validation is consistent with the notion of learning from feedback, either positive or negative. If the system is validated by comparing the system to an expert or a set of experts, then the comparative performance of the system provides a basis for knowledge acquisition. If the system and the expert provide similar responses then that would provide positive evidence of the quality of the knowledge. However, if there is disagreement then that could indicate that there is additional knowledge required, incorrect knowledge or inconsistent knowledge. In either of these situations, there is a need for additional knowledge acquisition. The validation process points to the need and the source.

Some research in human learning has found that feedback that emphasizes the process and structure of the judgemental task is more effective than outcome feedback (Hammond et al. [1973]). The first, provides insight into the content of the knowledge used in the judgement process, while the second only elicits the judgement. This finding indicates the importance of VV&M based on determining why a system makes specific judgements, not just the specific judgements or outcomes.

For example, consider the following description of the validation process of a system developed by the Internal Revenue Service (Archer and Peer [1990]).

> In the interim a person (the "user") was hired at an advanced clerical level to run the expert system. ... She was asked to obtain a sample of worked cases (worked by the experts) and run through the system. ... As she started putting finished cases ... into the computer, cases continually came up where the computer differed from the results that the original (expert) analyst had determined. At this point the user ... asked the analyst who was right and why. (p. 4)

4.2 Different Actors in the Validation Process

There are a number of different actors involved in various aspects of the validation process. Depending on the group doing the validation, we can anticipate establishing different knowledge acquisition needs. Different validators are sources of different knowledge.

If the validation source is the same expert, then either: the expert's knowledge was implemented correctly (no new knowledge to correct it); the expert's knowledge was implemented incorrectly (in which case revised knowledge must be obtained); the knowledge has changed (new knowledge is required and that expert will provide it); or knowledge is missing (as a result of recognition of this omission, new knowledge will be solicited).

In many situations, there will be multiple experts available as part of the validation process. In the case of complete agreement of the experts there are no keys to further knowledge acquisition. However, the situations where there is disagreement offer additional knowledge acquisition opportunities. If there is consensus, then that indicates a level of agreement and disagreement. If a core of the experts continues to disagree with a larger group, then this can indicate additional knowledge acquisition effort in the form of alternative solutions. On the other hand, if there is not any consensus, then this can indicate a situation where knowledge acquisition is impossible.

In other situations, the validator may be a less than an expert user, possibly in parallel use of the system in the specific decision making environment. In these situations, their use of the system can bring out the existence of unstated assumptions taken for granted by domain experts (O'Leary [1991]). This points the way for additional knowledge acquisition to establishing necessary knowledge to resolve those assumptions.

4.3 Reliability

Reliability tests, such as sensitivity analysis also are a source of knowledge acquisition. In the case of sensitivity analysis, if the system does not exhibit any unusual behavior on manipulation of a variable or patterns of variables then that provides positive feedback on the quality of the system. However, if the system begins to behave in an unusual manner on variation of a parameter or set of parameters, then that would indicate additional knowledge acquisition related to that parameter or may ultimately question the overall validity of the model. It is analysis of this type that indicates a close relationship between catastrophe theory (resulting from small variations in model parameters) associated with validation and knowledge acquisition.

5. Maintenance Portion of Life Cycle

An issue for systems that must change over time is knowledge acquisition in the maintenance process. In particular, what triggers the acquisition of new knowledge and what are sources of that new knowledge. The identification of the changing knowledge is necessary to ensure systems adapt to changes in the environment. However, little in the literature is aimed at KA during maintenance.

Maintenance can lead to a number of different sources of knowledge acquisition opportunities, deriving from the anomaly; the development of new products, the issuance of new laws, and the issuance of new standards; and new information deriving from new technology.

5.1 Anomaly

The anomaly plays an important role in knowledge acquisition from the maintenance process. O'Leary and Watkins [1991] discuss one system's movement through the maintenance portion of the life cycle. The system was developed by a large bank to analyze foreign currency trades for unusual trades. An initial version of the system had about 65 to 70 rules. During use of the system a number of trades were identified as those deserving of investigation, i.e., "unusual." This set of transactions was found to originate with a specific set of banks and occurred at a specific time of the day. The identification of these transactions as unusual led to the finding that some of the banks did foreign currency trades at different hours than did others because of their location in different time zones. In order to account for this effect, an additional 30 to 35 rules were added to the system.

5.2 New Products, New Laws and New Standards

The nature of the system will indicate other sources of knowledge acquisition. In many cases, new product information or new standard information or new laws will indicate the need for knowledge acquisition.

The development of new products must be accounted for by every expert system that includes information about the product line. Not only are marketing systems affected, but so are configuration systems. For example, one of the best known expert systems, XCON was built by DEC (Digital Equipment Corporation) to configure computer systems. An interview with a DEC scientist associated with that project led to the finding that the primary source of new knowledge for the system is product changes.

New laws can have a similar impact. In financial systems if the tax law changes, then those tax laws must be reflected in the systems that make decisions affected by those laws. One of the primary sources of knowledge acquisition in EXPERTax (Sphilberg and Graham [1986]) occurs with the tax law. When there were substantial changes in the 1986 tax law, there were substantial changes in the system knowledge base.

Another source of change is new standards. For example, the system EDAAS (Feinstein and Siems [1985]), which was designed for the Environmental Protection Agency (EPA) is dependent on governmental disclosure requirements. As those requirements change so will the knowledge base.

Additionally, new data may spur new problem solving. For example, in financial applications, new accounting standards are constantly being issued. Those standards lead to different disclosure requirements for external financial reports (such as income statements). Those different standards can lead to different financial data being disclosed. Adapting the system to those requirements is part of the KA process.

5.3 New Technology

Another source of new knowledge occurs when new technology changes the production process. O'Leary and Watkins [1991] discuss the development of an expert system, that led to the ability to audit multiple days transactions, where in the past only a single day's transactions could be examined. This resulted in a need to assess how that new capability could be exploited. In the past, there had been no ability to examine the potential of multiple interacting days, so new knowledge had to be created to take advantage of that availability.

6. Summary and Extensions

In the past the KA process has focused on documentation, interviews with experts, observation, and a number of other approaches. However, this paper suggests that one of the primary sources of knowledge acquisition are the activities of VV&M. It was found that

VV&M each provided different potential knowledge acquisition opportunities, arising from a different set of events. In addition, the errors, anomalies, and changes in VV&M point to different types of knowledge that can be acquired.

Understanding and taking advantage of those opportunities is critical since they allow us to anticipate not only the types of knowledge that can be acquired, but also the sources of that knowledge and the potential difficulties associated with that knowledge.

The results of this paper can be extended by examining other phases of the life cycle to determine how they facilitate knowledge acquisition. In addition, the results could be extended empirically beyond the theoretical and case evidence presented. In particular, empirical analysis of problems can be made to determine if they are error, anomalies or changes, and what kinds of knowledge acquisition they determined.

REFERENCES

Adrion, W., Branstad, M. and Cherniavsky, J. (1982). Validation, Verification and Testing of Computer Software. Computing Surveys, Volume 14, Number 2 (pp. 159-192).

Archer, C. and Peer, R. (1990). Rewrites of an Expert System Lessons Learned. Unpublished paper, Internal Revenue Service, Artificial Intelligence Laboratory.

Feinstein, J. and Siems, J. (1986). EDAAS: An Experimental System at the US Environmental Protection Agency for Avoiding Disclosure of Confidential Business Information. Expert Systems, Volume 2, Number 2 (pp. 72-85).

Guber, T. (1990). Acquiring Strategic Knowledge from Experts. In The Foundations of Knowledge Acquisition. J Boose and B. Gaines, Academic Press, San Diego, CA.

Hayes-Roth, F., Waterman, D., Lenat, D. (1983). Building Expert Systems. Addison-Wesley Publishing, Reading, MA.

Hamilton, D., Kelly, F and Culbert, C. (1991). State-of-the-Practice in Knowledge-Based System Verification and Validation. Unpublished paper NASA.

Hammond, K., Summers, D., Deane, D. Negative Effects of Outcome-Feedback on Multiple-Cue Probability Learning. Organizational Behavior and Human Performance, Volume 9 (pp. 107-118).

Nguyen, T., Perkins, W., Laffery, T. and Pecora, D. (1987). Knowledge Base Verification. AI Magazine, Volume 8, Number 2 (pp. 67-79).

O'Keefe, R., Balci, O. and Smith, E. (1987). Validating Expert System Performance. IEEE Expert, Volume 2, Number 4 (pp. 105-113).

O'Leary, D. (Summer 1987). Validation of Expert Systems. Decision Sciences, Volume 18, Number 3 (pp. 468-486).

O'Leary, D. (1991). Design, Development and Validation of Expert Systems: A Survey of Developers. In Verification, Validation and Testing of Expert Systems. John Wiley.

O'Leary, D. and Watkins, P. (1991). Expert Systems in Internal Auditing. Institute of Internal Auditors.

Sphilberg, D. and Graham, L. (1986). EXPERTax: An Expert Systems for Corporate Tax Planning. Expert Systems, Volume 3, Number 3 (pp. 136-151).