



Faculty Publications

1995

Case-Based Reasoning: Application Techniques for Decision Support

James V. Hansen
Brigham Young University

Rayman D. Meservy
Brigham Young University

Larry E. Wood

Follow this and additional works at: <https://scholarsarchive.byu.edu/facpub>



Part of the [Management Information Systems Commons](#)

BYU ScholarsArchive Citation

Hansen, James V.; Meservy, Rayman D.; and Wood, Larry E., "Case-Based Reasoning: Application Techniques for Decision Support" (1995). *Faculty Publications*. 3244.
<https://scholarsarchive.byu.edu/facpub/3244>

This Peer-Reviewed Article is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in Faculty Publications by an authorized administrator of BYU ScholarsArchive. For more information, please contact ellen_amatangelo@byu.edu.

Case-Based Reasoning: Application Techniques for Decision Support

James V. Hansen, Rayman D. Meservy and Larry E. Wood

Marriott School of Management and Department of Psychology, Brigham Young University, USA

ABSTRACT Decision-support systems can be improved by enabling them to use past decisions to assist in making present ones. Reasoning from relevant past cases is appealing because it corresponds to some of the processes an expert uses to solve problems quickly and accurately. All this depends on an effective method of organizing cases for retrieval. This paper investigates the use of inductive networks as a means for case organization and outlines an approach to determining the desired number of cases—or assessing the reliability of a given number. Our method is demonstrated by application to decision making on corporate tax audits.

INTRODUCTION

Case-Based Reasoning (CBR) is a form of approximate reasoning that relies on past cases to aid in deriving solutions or decisions for current problems. CBR is particularly relevant to human decision making, since decision makers nearly always rely on **prior experience** in solving problems (Rich and Knight, 1991).

Generally, CBR systems consist of a case memory, an indexing scheme, matching and retrieval mechanisms, and an interpretation component. The matching and retrieval mechanisms, driven by the current decision context, return the most similar cases (to a case on which a decision must be made) from the case memory. Similarity among cases is based on an evaluation of salient and relevant features.

As this suggests, CBR systems are concerned mainly with finding solutions to present problems through the examination of similar problems that have been solved in the past. A problem, its solution and the results of the solution are stored together in a case library where they can be accessed when a similar

problem is encountered. In the context of decision support, CBR can be used to find previous decisions stored as solutions to problem descriptions that are similar to a current situation (Duncan *et al.*, 1991).

Our paper is focused on the indexing problem. Effective methods of indexing are necessary for CBR to support decision making. Our particular interest is the method of inductive networks, its potential and limitations, as well as a methodology for estimating the reliability of the corresponding case base.

INDEXING: THE KEY TO EFFECTIVE CBR

Perhaps the most important issue in CBR and the design of CBR decision-making aids is the retrieval of appropriate cases. This issue has been termed the **indexing problem** (Slade, 1991).

The indexing problem is composed of two parts. The first concerns assigning appropriate labels to cases so that they can be retrieved at appropriate times (Kolodner, 1993). For

example, if we wished to retrieve cases pertaining to tax law, one useful method of labeling might be to distinguish a case as partnership or corporate. In this way, the set of retrieved cases could easily be limited to the type of case—partnership or corporate—relevant to the case at hand.

The second is the focus of this paper: organizing cases so that a search through the case library can be done effectively and efficiently (Kolodner, 1993). The principal contributions of our work are: (1) to elaborate on inductive networks as a method for organizing cases and (2) to demonstrate a technique for estimating the reliability of the related case base.

Case Features and Indexing

Research on the indexing problem has included studies of the type of features most useful for indexing. The results point to at least two levels of features found useful in indexing: surface features and structural features (Rich and Knight, 1991). Surface features are those that are represented by individual values. Structural features are more abstract and express relationships among surface features. Examples of surface features for a decision on a loan to a bank customer might be 'income \$75000', 'credit-history good', with no assessment of saliency or relationship among these or other features. All cases characterized by these respective values (or values close to these) might be judged as similar. Examples of structural features for the same situation might be 'income \geq \$75000' and 'owns-home yes' or 'credit-history good' and 'income > \$50000'. In the latter case, not all customers with incomes of \$75000 and good credit history might be classified as similar. Surface features are less complex because they are fact-based. Structural features offer a richer and more powerful representation scheme and are the main concern of this paper.

To use memory effectively, a powerful organizing mechanism for the selected features is required so that when decision makers are confronted with problems, the retrieved cases will be relevant to the decision makers' needs. This necessitates a strategy for distinguishing important indices from less important ones.

A number of indexing methods have been developed or proposed to address this need (the algorithm(s) associated with each method of case organization is shown in parentheses): flat memory (serial or parallel search), shared-feature networks (breadth-first graph search), prioritized discrimination networks (depth-first graph search), redundant discrimination networks (breadth-first graph search), and inductive networks (parallel search) (Kolodner, 1993).

Inductive Networks

The work focuses on inductive networks, which are hierarchically organized, shared-feature networks. The principal reasons for interest in inductive networks are twofold:

- (1) According to Kolodner (1993), the biggest technological issue of CBR is scale-up. How can one extend retrieval algorithms that work for hundreds of cases to work efficiently enough for tens of thousands? The machine-learning community has developed several inductive clustering algorithms that enable cases to be organized hierarchically so that only a small subset needs to be considered during retrieval. These may be useful in addressing scale-up needs.
- (2) The use of hierarchically organized, shared-feature networks lends itself to formal methods of estimating the reliability of the case base. Knowing the reliability of a set of cases has important implications for both novice and expert decision makers.

Inductive networks allow retrieval of stored cases by propagating the attribute values of the current case to a decision tree that has been induced from the stored case base. Each branch of the decision tree leads to a cluster of example cases at the leaf nodes. Since not all cases in a cluster necessarily have identical solutions, and the fact that the objective is to provide analogs to the decision maker, it is not appropriate to simply store the majority solution at each leaf. The set of attribute values from the current case will determine a path through the decision tree to a leaf node where cases having similar attribute values are stored. Each of these cases is then compared to the

current case using nearest-neighbor matching to establish a similarity score within the retrieved set of cases.

The inductive indexing mechanism parallels inductive learning algorithms that generate decision trees, except that the latter store a single classification decision at each leaf node. In this manner, rules are represented by paths through the decision tree. This is a different objective from storing a set of cases that are similar in order to facilitate efficient retrieval for inspection by a decision maker.

Rissland and Skalak (1990) have elaborated on the value of inductive machine-learning techniques to generate and refine indices and other ingredients of a CBR system. They tested a machine-learning algorithm on a set of legal cases involving tax law litigation, with the following results:

- The inductive machine-learning algorithm unearthed the surprising result that the taxpayer's occupation was the attribute that provided the most 'information gain'. That is, the taxpayer's occupation is a very salient feature in tax cases, yet the legal statutes omit this factor entirely.
- The inductive machine-learning algorithm derived features as to the number of hours per week the home office was used and whether the home office was in a separate structure as other salient features in determining the success of a plaintiff's claim. These features were indirectly contained in the statutes, but their relative importance was not established.
- In general, the features identified by the inductive machine-learning algorithm were among the most important in determining the success of plaintiff's claims.

Similar evidence is provided in the work of Selfridge (1990) and Laffey *et al.* (1991). Generally, it was found that in most problem domains not every case attribute will be relevant to indexing. For example, in retrieving cases to help in determining whether or not to grant a loan, the customer's name or social security number will not likely affect the applicant's ability to repay a loan. Attributes such as income, assets and occupation have a more direct bearing on ability to repay a loan. Consequently, these attributes might form a

basis for indexing cases, and their relative importance could be determined from inductively indexing a set of past cases.

Inductive networks can also be important in CBR applications where the surface characteristics of two situations may be very similar, yet the underlying problem is quite different; or the problems can be similar, while the situations share few surface similarities. Finding appropriate military strategies from chess problem solutions is an example of the latter (Duncan *et al.*, 1991).

Finally, inductive networks are based on theoretical concepts that facilitate analysis of the number of cases required to attain specified levels of indexing accuracy. This is discussed in the following section.

INDUCTIVE NETWORKS AND DECISION SUPPORT

In a decision-support system, CBR can aid in the following important decision areas identified by Libby (1981):

- By retrieving a set of similar cases to the one being considered by the decision maker, a decision can be made simply by observing the solutions of the retrieved cases. That is, if all (or the majority of) the similar cases had the same outcome, the decision maker may simply decide based upon expectation of that outcome. For example, if all the retrieved cases showed an outcome of 'loan repaid', the decision maker might, in the absence of other information, make a decision to grant a loan. This is particularly important for novice decision makers. It helps them to be consistent with past decisions, as well as to learn those combinations of attribute values that tend to be associated with loan repayment.
- For the expert decision maker, similar benefits of consistency apply. The expert decision maker is, however, more likely to combine information from past cases with other knowledge, and perhaps new information, in making a decision. Consider the same example as above. The expert decision maker may recognize that the loan applicant has a high income, but is employed by an auto-

maker that is noted for cyclical plant shut-downs. Consequently, the decision maker may want to limit the cases he or she compares to those loans made to applicants working for the same automaker. Such flexible indexing is available in current CBR software.

In both the above situations, knowing what the likely error is in the predicted outcome—as suggested by the set of retrieved cases can be important to determining the level of reliability to be placed upon that information. The following section presents a method of formal analysis of reliability from both *ex post* and *ex ante* points of view.

DETERMINING THE NUMBER OF CASES TO ACHIEVE DESIRED RELIABILITY LEVELS

Fundamentals

Inductive indexing of cases allows one to adapt recent theoretical work in assessing the likelihood of error in the resulting classifications. This information can be valuable to the decision maker in determining the reliability of retrieved cases in representing the desired information of interest, that is, the indexed classification of a case. Following Hausler (1988) and Tsai and Koehler (1993), this indexed classification is called a **concept**.

Let D represent a problem domain of interest to a decision maker. A **target concept**, f , is any subset of D . In the loan example, one concept would be the set of loan cases that were repaid. To identify the target concept, a set of cases is drawn randomly with replacement from D according to a fixed but arbitrary probability distribution, P . If a loan case resulted in repayment, it is called a positive example of the target concept. If it was not repaid, it is called a negative example. Thus, a case example for f is a pair (x, y) , where $x \in D$, $y \in \{0,1\}$, and $y = f(x)$.

A **learned concept**, g , is a subset of D consisting of cases that will be classified as positive according to the inductive rule:

$$g = \{x \in D \mid f(x) = 1\}$$

The error, e , of a learned concept is the

probability of the symmetric difference between the learned concept and the target concept, that is,

$$P(f\Delta g) = P\{(f - g) \cup (g - f)\} \text{ (Blumer et al., 1989)}$$

A confidence parameter, b , is an upper bound on the likelihood of an error (Angluin and Laird, 1988). That is,

$$\Pr\{P(f\Delta g) \leq e\} \geq 1 - b \quad (1)$$

where \Pr denotes probability.

Informally, the requirement for desired case indexing is that the probability of the difference between the inductive classification of a case and the true classification be small (e) with high probability ($1 - b$) (Hausler, 1988). We shall term an index to be (e, b) -probably approximately correct (PAC) if equation (1) is satisfied (Tsai and Koehler, 1993).

It is easy to see that the class of all possible concepts that could be specified based only on the attributes of the cases in D is $C = 2^{|D|}$, where C contains the target concept f . A hypothesis space, $H \subseteq C$, is composed of those concepts that are consistent with the language used by a learning algorithm that uses a finite set of cases from D . That is, a concept is consistent if all the positive and negative examples in an actual sample are the same as those predicted by the concept.

Now a labeling function for H can be defined in the following way:

Definition 1. A labeling function for H is the maximum number of ways that the concepts in H can label a set of m cases. This function is denoted as $L_H(m)$.

Definition 2. The Vapnik–Chervonenkis dimension of H , $D_{VC}(H)$ is the largest integer m such that $L_H(m) = 2^m$.

We now have the tools to establish the reliability of case classification.

Theorem 1 (Hausler, 1988). Denote the size of the hypothesis space as $|H|$. For a given e and b ($0 \leq e, b \leq 1$), the required number of cases is

$$\min \left\{ (1/e)[1n(1/b) + 1n|H|] \right. \\ \left. (1/e)[4\log_2(2/b) + 8D_{VC}(H)\log_2(13/e)] \right\}$$

Determining the Required Number of Cases

Suppose that for the loan situation, one wished to classify applications according to the following attributes and possible values:

<i>Attribute</i>	<i>Values</i>
DEBT	Yes/No
INCOME	Low/Medium/High
HOMEOWNER	Yes/No
CREDIT-RATING	Poor/Fair/Good
EMPLOYMENT-HISTORY	Poor/Fair/Good

We are typically interested in conjunctive concepts, where each attribute is either a term of a concept or not, the number of conjunctive concepts is $|H| = 3^2 4^3 = 586$. That is, there are two attributes whose values can be 'Yes', 'No' or 'Makes no difference', and three attributes that can take on four values (e.g. 'Poor', 'Fair', 'Good' or 'Makes no difference'). Haussler (1988) has shown that

$$n \leq D_{VC}(H) \leq 2n$$

where n is the number of attributes. We then have

$$5 \leq D_{VC}(H) \leq 10$$

Thus, for $e = 0.1$, $b = 0.05$ and $D_{VC}(H) = 10$ (to be conservative), the number of cases required by Theorem 1 is

$$m = \text{Min} \{94, 755\} = 94$$

With 94 cases, we are guaranteed to derive a concept that is within 0.1 of the true concept, with 0.95 probability. Formally, (0.1, 0.05)-PAC indexing is guaranteed.

From the above, there is a sound theoretical basis to determine the number of cases required, given *ex ante* reliability requirements. This is very useful when the required number of cases are available.

Assessing Reliability When the Desired Number of Cases are Not Available

It is often the situation that the desired number of cases are not available. This situation has recently been reported concerning applications at Apple Computer (Laffey *et al.*, 1991). In such instances, the decision maker can benefit by knowing how closely the concepts developed from the available cases approximate the true

concept. For the expert decision maker, this information can suggest the degree to which other information or judgement should be applied. For the novice decision maker, it can suggest when additional help should be sought.

Suppose that a decision maker finds that q cases are misclassified in a test sample of size m . Tsai and Koehler (1993) have shown that

$$\Pr\{b \geq e\} \leq \exp[-2(e - q/m)^2 m] \text{ for } q/m \leq e \quad (2)$$

and

$$\Pr\{b \leq e\} \leq \exp[-2(e - q/m)^2 m] \text{ for } q/m \geq e \quad (3)$$

Suppose that a set of cases is inductively indexed, and that this set is used to suggest decisions on 14 new cases, resulting in two incorrect decisions. Choose e to be 0.1. Since $q/m = 1/16 = 0.0625$, use equation (2) above, yielding

$$\Pr\{b \geq 0.1 | 2, 16\} \leq 0.95599$$

This means that there is as much as a 96% chance that the error of the concept is greater than 0.1. This might suggest to the decision maker that there may be factors about the cases that require careful application of the decision maker's judgement. On the other hand, suppose that the set of tested cases was 150, with five errors in the decisions suggested by the indexed cases. Then one would compute

$$\Pr\{b \geq 0.1 | 5, 150\} \leq 0.265399$$

The increased size of the test sample greatly improves the results. This might lead the decision maker to be more reliant on the decisions suggested by the retrieved cases.

FLEXIBLE SUPPORT FOR DECISION MAKING—AN EXAMPLE APPLICATION

A set of audit cases recorded by the state of Pennsylvania was utilized for the application. This was an ideal set of cases in that it represented a real-world unsolved problem with noisy data. Noisy data—data that is not always consistent with itself—is characteristic of many real-world problems. Due to noise and ill-defined relationships among the independent variables, the data in the audit cases had not satisfactorily yielded to other methods

of analysis (Nagin, 1988). The reader may recall that Rissland and Skalak (1990) used an inductive network for tax cases. Their work, however, involved tax litigation, and we know of no evidence to suggest that tax settings are any more suited to inductive networks than other problem domains.

Briefly stated, the state of Pennsylvania wishes to make efficient decisions on which corporate tax returns to audit in order to maximize the return per auditor hour. For the application, a set of 200 cases was drawn from a population of 5913 audits that were performed over a 30-month period, January 1984 to June 1986. Seven attributes were recorded that were hypothesized to affect the expected return per auditor hour. These are listed in Table 1.

Elsewhere (Denna *et al.*, 1992), the advantages of CBR for both **novice** and **expert** decision makers were outlined. These ideas were briefly exemplified above and can be summarized in the following way:

- Novice decision makers are presumed to be more reliant on the outcomes of the retrieved cases to guide their decisions. That is, the novice decision maker brings less outside information to the decision process and may benefit more from regularities that are found

in the case clusters of an inductive network. This is particularly true if the predictive accuracy based on the outcomes of the retrieved cases is very high (or very low, for that matter).

- Expert decision makers are expected to utilize the information in the retrieved cases in concert with other information drawn from their experience. In fact, they may utilize an informal cause-and-effect model that has been validated from experience.

In the analysis, equation (1) was addressed by simply utilizing the predictions suggested by the retrieved cases on a set of holdout cases. Support for expert decision makers was addressed by utilizing a study conducted by a group at Carnegie-Mellon University (Nagin, 1988), which generated the type of informal model that might be used by an expert.

The method of testing was resampling using a tenfold cross-validation method, which produces reliable estimates of the true error rate (Weiss and Kulikowski, 1991). In this application, 20 different sets of indexed cases were created, one at a time. Each set omitted 10 randomly selected cases, with each set of 10 being disjoint from the others. Each of the holdout cases was supplied to the CBR system

Table 1 Salient attributes for audit decision

	Variables	Description
Outcome	Deficiency per hour	The amount of additional tax assessed (after appeals, but exclusive of penalties and interest) divided by the total number of hours spent by the auditors
Attributes	Tax amount	Tax, in addition to the original tax liability reported by the business, assessed as a result of the audit not including penalty and interest
	Tax type	Sales tax deficiency assessed, use tax deficiency assessed, or both
	Deficiency type	Code identifying the type of reassessment payment made on the original tax deficiency assessed
	Reassessment amount	Reassessed amount of tax as determined in the appeal process
	Wages to gross ratio	Total wages divided by the total gross sales
	Gross	Reported gross sales
	Prior audit tax deficiency	Additional tax assessed in prior audit, not including penalty and interest

created from the remaining 190 cases for appropriate retrieval.

The actual return per auditor hour on the 200 cases was divided into two categories: **high** and **low**. Based on the prior work at Carnegie-Mellon (Nagin, 1988), the following assignment was made:

Expected return per auditor hour =
{ low, if value \leq \$36
{ high, otherwise

CBR Support for Novice Decision Makers

A profile of an indexing tree that was induced for stored cases is illustrated in Figure 1. The intent is to convey, without detail, a sense of the complexity of the induced network. An enlarged fragment of the indexing tree is




Figure 1 Indexing tree induced from case data

reproduced in Figure 2. The results of using the outcome of the retrieved stored cases for predicting the outcome of the 10 holdout cases over 20 trials is shown in Table 2(a). The numeric values represent relative weighting units. One hundred and fifty-five predictions were correct, for 76.5% prediction accuracy.

Interestingly, for all cases where the expected return per auditor hour was greater than \$65, the predictions were 100% accurate (Table 2(a)).

Of the 45 errors in prediction, 23 resulted from a **low** prediction when the result was **high**, 20 of the errors resulted from a **high** prediction when the result was **low**. These results suggest neutrality in the type of error that results. An erroneous decision of the first type above could result in the loss of additional tax revenue due to a 'no-audit' decision. An erroneous prediction of high might lead to expending auditor time on a low return audit. Since the inductive network was very accurate in retrieving cases when the expected return per auditor hour was over \$65, the second type of error seems the least worrisome for the novice decision maker.

CBR Support for Expert Decision Makers

As previously noted, an expert decision maker often brings experience to the decision analysis that may suggest a more flexible use of the indexed cases. In a number of problem areas such as financial analysis, there is well-established theory and empirical evidence of causal relationships among case attributes. For example, when a customer applies for a home-equity loan, the decision made by the bank may depend on various financial calculations and ratios to determine the applicant's ability to repay. But there may be qualitative information that is of equal or greater relevance: job stability, work skills, general creditworthiness, and so on.

In the case of the audit choice decision, prior work by the Carnegie Group (Nagin, 1988) suggested a qualitative model of the form shown in Figure 3. This model was used as a surrogate for an expert decision maker. The case retrieval results again are good (Table 2(b)), yielding 73.5% accuracy

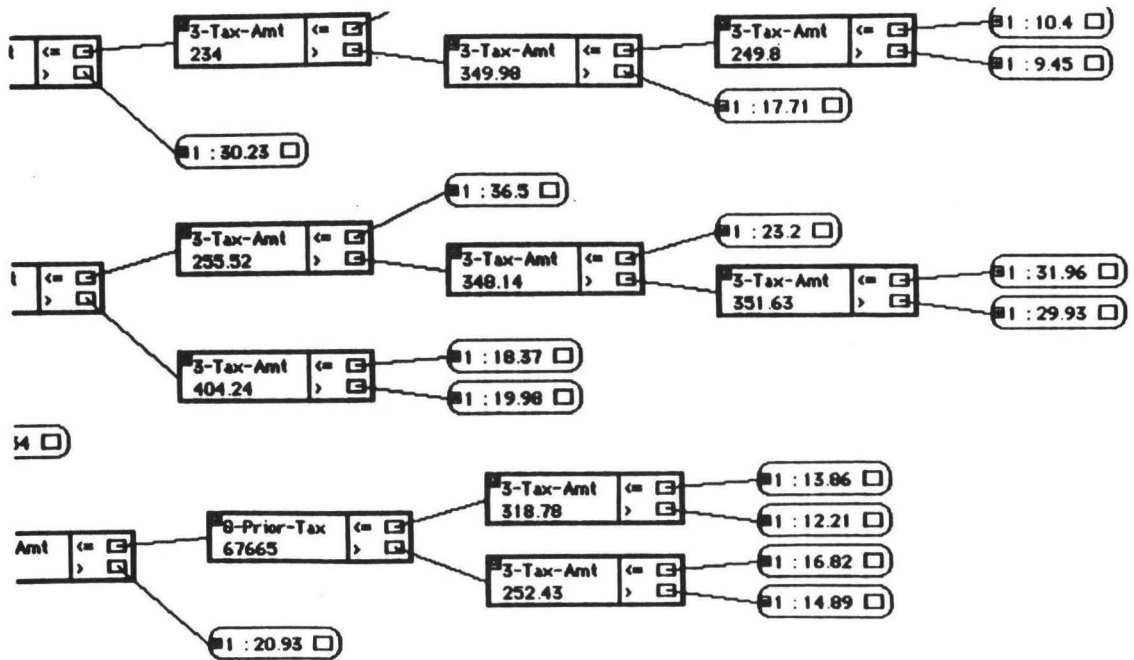


Figure 2 Fragment of indexing tree illustrated in Figure 1

Table 2 Comparative performance results
(a) Induced indexing tree

Predicted audit class	Actual audit class	
	High	Low
High	70	25
Low	20	85

(b) Qualitative indexing tree

Predicted audit class	Actual audit class	
	High	Low
High	73	23
Low	30	74

over the tenfold tests; however, they are not quite as good as those generated from the model induced from all the data in each case (76.5%). Weiss and Kulikowski (1991) suggest the following test for comparing models such as these:

- (1) Compute the standard error for the best-performing model, as follows:

$$SE = [E(1 - E)/n]^{1/2}$$

Thus, for our results, we have

$$SE = [0.235(0.765)]^{1/2} = 0.42$$

- (2) Select the simplest model (if any) that is within one standard error of the best-performing model.

Since the qualitative model is the simplest and falls within one standard error (0.03) of the inductive model ($0.765 - 0.735 = 0.03$), by the above rule it would be selected for implementation.

How Reliable are the Results?

In this application we were not able to select an optimum number of cases, so the evaluation of reliability must follow the *post-hoc* method described earlier.

Assume that the qualitative model is being evaluated. The evaluation parameters are as follows:

$$\text{Number of errors} = q = 53$$

$$\text{Number of cases tested} = m = 200$$

$$e = 0.3$$

We then have

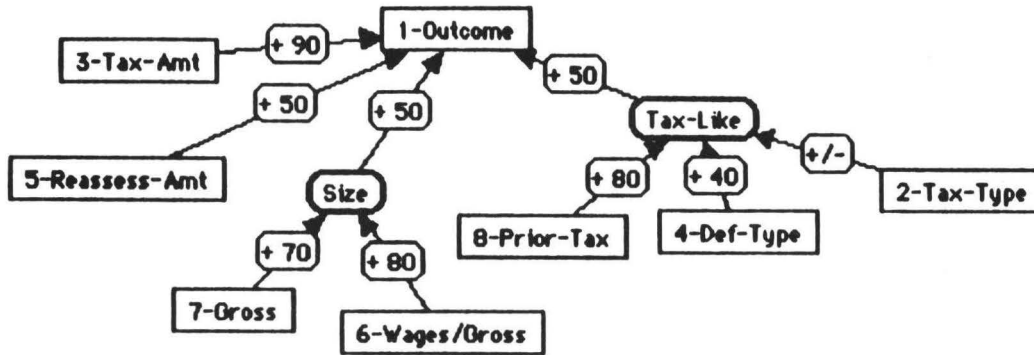


Figure 3 Qualitative indexing model

$$\Pr\{b \geq e\} \leq \exp[-2(e - q/m)^2m]$$

$$= \exp[-2(0.3 - 0.265)^2 200] = 0.613$$

That is, there is as much as a 61.3% probability that the error in the indexing concept is greater than 30%. This suggests that a larger set of cases may be necessary or, at the very least, that the model's performance be carefully monitored over time.

SUMMARY

The underlying processes of CBR have emanated from research in artificial intelligence (AI)—particularly machine learning. While expert systems are probably better known as AI's contribution to decision making, CBR is perhaps better suited to providing **support** to the decision maker through retrieval of relevant analogs, as well as adapting new information as it becomes available. These capabilities are particularly valuable for complex and ill-structured problem domains.

The **indexing problem** is the most pressing challenge in the implementation of CBR systems. Our exposition focused on the case organization aspect of this problem as represented by inductive networks. Inductive networks are of particular interest because of their capability to represent case similarities at a number of levels, thereby reducing the search time for large case bases. Moreover, inductive networks can support both expert and novice decision makers. Inducing indexing trees, based upon only the case library, provide support for

novice decision makers via the similarities discovered during indexing. For the expert decision maker indexing trees can be created that are based upon the prior knowledge of the expert about relationships among case attributes.

Inductive networks further provide a sound theoretical basis for assessing the likelihood an error resulting from the classification of a case based on similarities to previous ones. Ideally, when large numbers of prior cases are available, one can decide *a priori* the number of cases needed to provide a desired level of reliability. Alternatively, when a limited set of cases is available, it is informative to assess the level of reliability achieved with that set. We have outlined applicable methods for dealing with both possibilities.

References

- Angluin, D. and Laird, P. Identifying k-CNF formulas from noisy examples. *Machine Learning*, 2, No. 4, 1988, 319-342.
- Blumer, A., Ehrenfeucht, A., Haussler, D. and Warmuth, M. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36, No. 4, 1989, 929-965.
- Denna, E., Hansen, J., Meservy, R. and Wood, L. Case-based reasoning and modeling auditor decisions. *International Journal of Intelligent Systems in Accounting, Finance, & Management*, 1, No. 3, 1992, 1-13.
- Duncan, N., Blevis, E. and Glasgow, J. A case-based approach to decision support. Technical Report No. 91-311, 1991, Department of Computing and Information Science, Queen's University, Kingston, Canada.

- Haussler, D. Quantifying inductive bias: AI learning algorithms and valient's learning framework. *Artificial Intelligence*, **36**, 1988, 139-154.
- Kolodner, J., Improving human decision making through case-based decision aiding. *AI Magazine*, **12**, No. 2, 1991, 52-68.
- Kolodner, J., *Case-Based Reasoning*, Morgan-Kaufmann, San Mateo, CA, 1993.
- Laffey, J., Machiraju, N. and Chandhok, R., Organizational memory as a support for learning and performance: prototypes and issues. *Proceedings of the International Conference on the Learning Sciences*, 1991, Association for the Advancement of Computing in Education, Evanston, IL.
- Libby, R., *Accounting and Human Information Processing: Theory and Applications*, Prentice Hall, Englewood Cliffs, NJ, 1981.
- Nagin, E., Graduate consulting group. Unpublished report, SUPA, Carnegie-Mellon University, 1988.
- Rich, E. and Knight, K., *Artificial Intelligence*, 2nd edn, McGraw-Hill, New York, 1991.
- Rissland, E. and Skalak, D. Opportunities for learning with and within CBR. 1990 Spring Symposium Series on Case-Based Reasoning, Stanford University.
- Selfridge, M., Why do adults tell stories? Why do children play? For the same reason: re-indexing old cases under new generalizations. 1990 Spring Symposium Series on Case-Based Reasoning, Stanford University.
- Slade, S., Case-based reasoning: a research paradigm. *AI Magazine*, **12**, No. 1, 1991, 42-55.
- Weiss, S. and Kulikowski, C. *Computer Systems That Learn*, Morgan Kaufmann, San Mateo, CA, 1991.
- Tsai, L. and Koehler, G. The accuracy of induction for expert systems construction. *Decision-Support Systems* (in press, 1993).