# SELECTING PLANS FOR CAPITAL ASSET ACQUISITION THROUGH CLASSIFICATION PROBLEM SOLVING

Eugene Wallingford and Jon Sticklen

AI/KBS Group
Department of Computer Science
Michigan State University
East Lansing, MI 48824-1027

## PROBLEM DESCRIPTION

The future existence and profitability of nearly all business entities from small privately-held firms to the largest international corporations depend on the judicious acquisition, maintenance, and disposal of capital assets. No concrete definition exists for the term "capital asset," but generally a capital asset has a useful life greater than one year and is expected to earn income sufficient to cover the operating expenses and amortized acquisition cost associated with it. Among the assets fitting this definition are those usually called "plant and equipment": land, buildings, facilities, equipment, machinery, and vehicles. The importance of capital asset decisions proceeds from the relatively large commitment of resources ordinarily involved and from the fact that funds used in acquisition traditionally remain invested for a long time.

When considering a potential capital asset acquisition , decision-makers must address both the firm's financial goals and the many nonfinancial issues associated with its environment. Financial analysis techniques incorporate ideas such as net present value and discounted cash flow in order to determine the expected monetary value of the asset, given the firm's objectives with respect to risk and return on investment [9,14]. Other issues, though, also impinge on the feasibility of the transaction. Government and industry regulations often come into play, and the company must always consider how its actions will affect its role in the community and its relationships with employees, suppliers, and customers. Once a capital acquisition is demonstrated to be financially acceptable, the firm enters a second major phase of decision-making: selecting the most appropriate way to structure the acquisition as a transaction. A whole set of financial and nonfinancial factors shapes the context of this planning decision, and the elements of uncertainty and risk play even greater roles in complicating future expectations [2,10].

The form of a capital asset acquisition greatly affects a firm's tax status, its financial reporting obligations, and its economic position for future business activity. Amortization of the asset's cost through depreciation and other tax deductions form the most important element of the asset's profitability to the firm, and specific tax treatment depends on many different characteristics of the firm, the asset, and the nature of the transaction [11]. Financial accounting requirements also rely heavily on the nature of the transaction and on the intent of the parties involved [1,13]. Finally, a firm's ability to carry out future decisions can be enhanced or hampered as a result of structuring an asset acquisition in a particular way; decision-makers must exercise great care so as not to constrict the firm's operational freedom unnecessarily. Each of these three elements must be given strong consideration when deciding on an optimal transaction structure precisely because of their effect on the firm's current and future status.

This project implements a classification problem solver in the domain of capital asset acquisition planning. The problem solver selects a plan that best fits the user's context and goals so as to insure the desired tax and financial accounting treatment. Thus, the decision-making process considered is that of the potential acquirer of a capital asset. The so-called "investment decision" whether the asset should be acquired or not involves primarily financial analysis and will be addressed only briefly by the problem solver. The focus on the "financing decision" how to acquire the asset is justified by the richness of the types of knowledge used in selecting an acquisition plan and by the uncertainty inherent in evaluating evidence. Funding and tax planning issues drive the selection process and are stressed early in the decision-making process; financial accounting considerations, viewed by some as playing a less significant role in capital asset acquisition [10], affect the decision only later. Finally, the problem solver does not deal with real estate acquisition, a separate problem in its own right.

## BACKGROUND OF THE PROBLEM SOLVING METHODOLOGY EMPLOYED

Classification problem solving forms the heart of the approach exercised in this research. Within the artificial intelligence community, classification has been demonstrated to be an important modality in diverse problem-solving situations [7,8]. All types of classification problem solving depend critically on the availability of a predefined set of classification categories to be considered. The problem solving then centers on the selection of the category most appropriate to the problem situation under examination.

This paper employs a particular type of classification problem solving known as hierarchical classification [12,6,16]. In addition to the the general requirement for an existing set of categories, hierarchical classification further requires that the categories must be amenable to arrangement into a "specialization hierarchy," with the most general categories at the top of the hierarchy and the most specific categories at the bottom. Hierarchical classification utilizes an establish-refine control strategy. When a specialist in the hierarchy establishes itself (i.e., determines that it provides an accurate description of the problem situation), it attempts to refine the classification by allowing its subordinates in the hierarchy -- which correspond to more specific categories to attempt to establish themselves. Problem solving begins at the top of the hierarchy with the most general category and then alternates between establishment steps and refinement steps until a leaf-level category, one at the bottom-most level of the hierarchy, is established or until no further refinement succeeds.

The refinement step itself is not difficult, but the establishment step poses a more elaborate task. Within each of the classification categories resides a local, structured knowledge base that contains patterns of experience-based information to establish the category. This type of establishment, structured matching, largely follows longstanding results on data abstraction in artificial intelligence [17]. The basic idea of structured matching is that partitions of data can be used to reach a conclusion concerning category viability. For example, in determining whether or not to lease an asset, a manager will generally consider both financial and nonfinancial factors affecting a potential lease transaction. Patterns dealing with financial conditions would be placed in one partition, or knowledge group, and patterns relating to nonfinancial factors would reside in a second knowledge group. Thus the problem solver's knowledge is partitioned into one or more knowledge groups, each of which can be evaluated based on the patterns it contains.

The knowledge group is the basic unit of structured matching. In order for a specialist to establish or reject, it must be able to evaluate its low-level knowledge groups and merge the results of these evaluations into an overall decision about the category's degree of fit. Structured matching accomplishes this by reporting the values of its low-level knowledge groups, those dealing with direct observations, to higher-level knowledge groups that correspond to more general partitions of knowledge. A higher-level knowledge group takes these values, which are abstract in the sense that they represent the degree of fit in terms of a small set of discrete confidence values, and composes an aggregate evaluation of all the knowledge groups reporting to it. As a result, structured matching employs a hierarchy of knowledge groups to perform its task. One must be careful to distinguish the two types of hierarchies that play a role in hierarchical classification. One hierarchy consists of classification categories; within each category exists a hierarchy of structured matching knowledge groups whose duty is to establish or reject the category in which it resides. This two-tiered knowledge decomposition has been discussed in detail by Sticklen et al. [18] and has been shown to possess characteristics favorable for system debugging and system extension.

## CONCEPTUAL DECOMPOSITION OF THE PROBLEM

Decision-making in the domain of capital asset acquisition planning reflects the hierarchical structures of tax and financial accounting regulations. As a result, classification problem solving offers a natural approach for selecting an appropriate acquisition type based on a firm's current status and goals. The Internal Revenue Service (IRS) delineates a hierarchy of transaction types in order to determine the correct tax treatment for each acquisition [11], and the Financial Accounting Standards Board (FASB) has created its own classifications to insure that a firm's financial reports represent its standing as fairly as possible [1,13]. To a lesser extent, the corresponding tax regulations and financial accounting standards impose a natural ordering on the process of selecting one transaction type from among a set of alternative types. These hierarchies can be viewed from two different perspectives. The IRS and the FASB use their own structures to classify completed transactions for the purpose of insuring regulatory compliance by the parties involved. The parties themselves potential buyers and lessees, potential sellers and lessors attempt to achieve the most favorable tax and reporting effects by using a composite hierarchy, assembled using compiled knowledge, to choose a transaction type having the desired characteristics.

Developing a hierarchy that unifies the IRS and FASB classifications with compiled decision-making knowledge reveals a basic feature of the domain: the different objectives of each element prevent a fully cooperative planning effort. The IRS focuses on the issue of income recognition, the FASB concerns itself with this and with many other aspects of fair and representative reporting, and business decision-makers must consider not only these static factors but also the firm's performance and future well-being. Yet even this trichotomy supports the investigation of classification problem solving in this domain. In many ways, the problem of selecting an acquisition plan seems to consist of several separate selection problems related by an overall concern for the fitness of the firm. The final selection is governed by the interaction of the solutions for these smaller problems. The adoption of a classification problem-solving strategy for capital asset acquisition planning emphasizes the hierarchical nature of the compiled knowledge that experts have developed to address this problem and recognizes the contextual element involved in the planning process as the smaller competing problems are solved.

The composite classification hierarchy for this problem appears in Figure 1. Nodes in this hierarchy correspond to transaction types available to the firm at varying degrees of specificity. The top-level node, capital asset acquisition , indicates the feasibility of acquisition and thus addresses the "investment decision" mentioned earlier. A more comprehensive handling of the investment decision by this problem solver would require this node to call upon another decision-support tool to perform a financial analysis of the acquisition under consideration. If acquisition is feasible, the problem solver incrementally refines the selected plans with respect to a set of factors that includes the nature of the asset, the source of consideration given (debt, equity, or internally-generated funds), and external consequences (tax and reporting effects). Tax planning occurs at higher levels in the hierarchy due to its relative economic importance, and financial accounting classification appears in lower-level nodes. One of the interesting issues of this project involves determining whether the replication inherent in this hierarchy is a product of the limitations of classification for solving this problem or of the contextual character of plan selection.

********************** Figure 1 goes about here **********************

The knowledge used to establish nodes in this decomposition has a variety of forms. Since the acquiring firm's main goal is assumed to be profitability, much of the decision-making involves financial analysis, and this characteristic cannot be ignored. But judgmental knowledge plays a crucial role in dealing with abstract goals, in forecasting future business conditions, and in assessing the risk and uncertainty embodied in current projections [2,10]. The ultimate decision of whether a particular acquisition plan should be employed or not must incorporate financial analysis of the transaction's net benefit to the firm; however, performing a financial analysis of every possible acquisition type is not only costly but also wasteful of available information and knowledge. This decomposition permits a high degree of judgment to play an active and accurate role in the problem-solving process of identifying and recommending viable plans best suited to the firm's specific circumstances.

This conceptual decomposition of acquisition plan types provides a framework for a composite knowledge-based system that reflects the expertise of individual specialists and the explicit effects produced by IRS and FASB standards. From these sources of knowledge, one can identify several general principles that support the soundness of the decomposition presented in Figure 1. The fundamental decision made using this hierarchy involves the nature of the acquisition . One can acquire the full rights of ownership to an asset or merely the right to use an asset. This concept recurs throughout the hierarchy, but nowhere is it more prominent than in the purchase versus lease decision. Upon determining the appropriateness of each alternative with respect to the nature of the acquisition, the problem solver can recommend one or more legal forms for the transaction.

The purchase hierarchy decomposes quite naturally around two common-sense notions, the reason for the acquisition and the "funding principle." Acquiring an asset to replace an existing one introduces considerations not present when acquiring for the purpose of expansion. The existing asset provides unique options (recondition, trade-in) and also a unique responsibility (the problem of disposal), and these shape the context of the purchase. Whenever the basic type of purchase has been identified, the problem solver must consult the funding principle, which states simply: "The firm can pay for the asset now, or the firm can pay for the asset later." The decision between outright and financed purchases completes each subtree of the purchase hierarchy that involves a new asset, in each case taking a form suitable for the specific context. This replication, then, does not seem to be a product of classification's limitations in this domain.

In the classification of leases, IRS guidelines and FASB standards provide explicit decomposition of the plan alternatives. These sets of rules essentially embody the "nature of acquisition" principle discussed above, only this time with respect to the intent of the parties involved. Due to the overwhelming effect of income tax on a firm's profitability, the IRS classification receives primary consideration in the hierarchy. The available acquisition types are then refined to address financial reporting issues under FASB standards. Much replication results, including some unlikely combinations; but this repetition derives directly from incorporation of separate classifications, not from context. The final decomposition of this problem as a whole reflects a common intuitive understanding of the acquisition process that is grounded in the economic realities of a complex business environment.

## EXAMPLES OF CLASSIFICATORY RECOMMENDATIONS

The nodes of the classification hierarchy, which correspond to recommended acquisition types, employ only a few different knowledge group organizations, with each node structured in the way most appropriate to the type and amount of knowledge it embodies. In only the simplest nodes, one or two knowledge groups suffice to recommend or reject a plan. Direct lease relies heavily on the context established by its superspecialist and thus requires just one knowledge group to determine its feasibility. When a node includes exactly two knowledge groups, as does recondition , they typically address the feasibility and desirability of the associated plan. "Feasibility" determines if the firm meets the preconditions necessary to select a particular plan (in this case, to recondition an existing asset), and "desirability" then decides whether the firm should adopt the plan. A specialist having a feasibility knowledge group in its organization will generally evaluate other knowledge groups only when "feasibility" returns a favorable value.

Leveraged lease offers a more representative example of specialists in the acquisition hierarchy. The hierarchical organization of this specialist's knowledge groups, as shown in Figure 2, denotes a second level of data abstraction in the evaluation of a potential leveraged lease transaction. The four lower-level nodes deal with specific factors affecting the decision, much as knowledge groups in the simpler nodes do, and these factors are common to many of the more complex nodes. "Feasibility" and "conditions on the lessor" insure the possibility of a leveraged lease in the firm's current environment, with the former determining whether a potential lessor exists and the latter determining whether this potential lessor would be willing to participate in such an arrangement. The acquiring firm's reasons for structuring the transaction as a leveraged lease are evaluated in two knowledge groups. "Motivation" explores environmental factors that make this plan desirable, and "financial factors" tests those conditions likely to have a material impact on the firm's economic status.

********************** Figure 2 goes about here **********************

The job of recommending or rejecting a leveraged lease transaction ultimately belongs to the "summary" evaluation knowledge group, which aggregates the values returned by the four lower-level knowledge groups. While these lower-level groups consider individual data from the firm's case (e.g., trends in interest rates), "summary" deals with the abstract notions of feasibility, motivation, and financial impact. The final decision regarding a leveraged lease is made not on the basis of the numerous individual data, which are often in conflict, but rather on the basis of how they as a whole are expected to affect the future economic position of the firm. Thus "summary" relies on the other knowledge groups to gather and to evaluate specific information relevant to the plan and then considers only the more general concepts in making its decision. "Summary" need consider "motivation" and "financial factors" only in the case that "feasibility" and "conditions on the lessor" indicate that a suitable lessor exists for the transaction; if so, it must strike a balance between the firm's desires and the likely economic impact of the transaction. This abstraction of knowledge levels seems to provide a more accurate model of how a human manager would solve this problem in an accounting domain.

Another example of a common knowledge group structure can be seen in the node conditional sale (Figure 3). Here, "summary" evaluates the two major factors that lead a firm to seek a conditional sale the firm's ability to benefit from ownership and the firm's inability to purchase. While "financing" is able to determine the strength of the second factor by asking only two questions about credit standing, "ownership" must consider more complex issues. This knowledge group gathers some primary data about the asset and also refers to the knowledge group "tax considerations," which evaluates the firm's ability to employ the tax deductions generated by a conditional sale. The knowledge group cognate headed by "tax considerations" occurs throughout the acquisition hierarchy, primarily in the purchase subtree, where tax consequences play such a vital role in transaction planning. Since many of the high-level concepts involved in this type of planning affect the selection of more than one transaction type, several cognates of this variety are used in the hierarchy to provide a consistent and efficient abstraction of detailed information.

********************** Figure 3 goes about here **********************

The most complex knowledge group organizations belong to purchase and lease. The decision to recommend or to reject these acquisition types is fundamental to the problem of capital asset acquisition planning because these nodes provide the basic distinction between acquiring the rights of ownership and acquiring the rights of use. All tax and reporting consequences hinge on this principle, so decision makers must generally consider a wide variety of factors to achieve the results desired by the firm. As a result of this variety, several intermediate concepts are employed by these specialists for the purpose of grouping related information. The knowledge group structure for the purchase node, given in Figure 4, presents an interpretation of how these intermediate concepts relate to one another with respect to the recommendation decision. The four levels of knowledge abstraction in purchase's hierarchy reflect the complexity of the specialist.

********************** Figure 4 goes about here **********************

In reality, the structure consists of four basic cognates related in a standard pattern. The "tax considerations" cognate was discussed earlier in the explanation of conditional sale , but here it appears in its most common setting. The tax consequences of a transaction are only one element, albeit a major one, of the transaction's total financial impact; also of importance are the costs involved and trends in the conomic environment. "Financial factors" refers to these knowledge groups in an effort to evaluate the full economic impact to the firm. Of course, a capital asset acquisition has non-monetary effects as well, and the set of knowledge groups headed by "nonfinancial factors" determines the relative impact of these effects. Finally, the "summary" evaluation knowledge group weighs the financial and nonfinancial implications of the transaction in order to recommend or to reject a purchase plan. These four sets of knowledge groups provide a basic structuring of the acquisition planning process; and, as cognates, they appear throughout the hierarchy in a variety of forms and combinations, depending upon the particular context in which they are evoked. But in purchase they appear together and establish the many contexts of purchase alternatives.

## SIGNIFICANT ISSUES ENCOUNTERED IN PROJECT DEVELOPMENT

This model of capital asset acquisition plan selection has been incorporated into a working knowledge-based system using CSRL, a system-building tool designed specifically to facilitate the construction of hierarchical classification problem solvers. In CSRL, a specialist attempts to establish itself (through evaluation of data by its knowledge groups) and then returns an integer value between +3 and -3, inclusive. These values denote qualitative valuations of a category's plausibility from "highly likely" (+3) to "highly unlikely" (-3). Whenever the specialist establishes itself at a value greater than +1, it attempts to refine the classification by invoking its subordinates. The use of only a small set of discrete qualitative ratings plays an important role in representing the decision-making knowledge of human experts [18].

The application of hierarchical classification to this problem of plan selection introduces difficulties grounded both in the nature of the domain and in the state of classification as a planning tool. As mentioned earlier, little synergy exists between IRS regulations and FASB standards due to the differences between the goals of the two bodies, so the plan hierarchy contains several instances of replicated nodes. This replication results from the contextual nature of plan selection in a domain characterized by a dichotomy of motivating factors. One should expect to see a similar type of replication whenever a domain requires a planner to make two (or more) essentially independent classificatory decisions. So long as the number of individual classifications to be made remains small enough that the composite classification hierarchy is of manageable size, hierarchical classification offers a reasonable strategy for plan selection.

Another obstacle arose in the selection and grouping of domain knowledge. Most of the specialists in this hierarchy cannot be recommended on the basis of one or two facts from a specific firm's case; on the contrary, a specialist must often consider a wealth of "small" factors, carefully weighing their net effect, before arriving at a solution. Even in instances when a single high-level indicator such as "favorable tax consequences" is sufficient to recommend a node, this indicator itself requires inspection of several items of information. As a result of the inherent complexity of many acquisition planning decisions, many specialists have multiple levels of knowledge abstraction.

This feature permits the consideration of high-level concepts which in turn consider more specific factors. A negative effect of such a situation is that one must structure the knowledge groups and their interactions quite carefully in order to provide an acceptable sequence of questions during execution of the system. On the other hand, one might view this difficulty as a positive force, since it requires actually reinforces a greater understanding of the acquisition type and its associated recommendation knowledge. The complexity of the decision, then, is accompanied by an intrinsic depth of understanding into the structure of the decision.

*Interpretation of Classification Results*. The use of hierarchical classification as a planning tool requires particular interpretations of many common classification scenarios. When a specialist establishes in a simple classification hierarchy, one can only conclude that the data presented to the problem solver matches the context of the specialist with some degree of confidence. This degree of fit should then be interpreted in terms of some higher-level task, such as diagnosis or planning, as a part of the problem-solving process. Much of the literature on hierarchical classification [4,7,8] has dealt primarily with diagnostic problem solving; and, as a result, many common interpretations of classification rely on a diagnostic background. One also notices this subtle bias in the command terminology of CSRL. In order for classification problem solving and tools like CSRL to gain acceptance in a broad variety of domains, an effort must be made to cast the language of classification in terms independent of any one type of higher-level problem solving.

********************** Figure 5 goes about here **********************

Consider the two classification hierarchy fragments presented in Figure 5, one taken from a classifier charged with a diagnostic task in the domain of infectious disease and the other taken from the plan selection classifier of Figure 1. In diagnosis , when specialist infection establishes itself but none of its subspecialists can establish (see Figure 5a), the classification essentially means "Infection is one disease category that matches the current case, but more information is needed to determine a more specific category." However, for a plan selection problem, this type of interpretation should probably be reserved for situations in a which the subspecialists can neither recommend nor reject (Figure 5b). When all subspecialists of purchase through replacement reject in the acquisition system presented here, the classification implies "The firm should acquire the asset via purchase through replacement, but the firm may be unable to structure the acquisition as any specific kind of purchase through replacement." The classification given in Figure 6 provides an example of this; it recommends that the firm purchase the asset but also demonstrates the firm's inability to buy the asset outright or to finance the purchase. This distinction between diagnostic and plan selection scenarios arises because one cannot typically assume that all diagnostic hypotheses have been included in a classificatory representation (the closed world assumption ). In developing a plan selection hierarchy, though, the definitional nature of the acquisition plan types make the closed world assumption reasonable.

********************** Figure 6 goes about here **********************

Essentially the same interpretation of this scenario is required in any planning system constructed using classification problem solving: "The user should adopt plan X, but all specific instances of plan X are unsatisfactory." How to resolve this situation seems to be an issue dependent on the domain and on the specific facts of the case; for the case in Figure 6, what the firm should do depends on the original motivation for seeking the acquisition. Other examples of domain-independent and domain-dependent interpretation exist. If a specialist and all of its subspecialists establish in any planning hierarchy, then one can infer that adopting any of the more specific plans would be recommended. Again, resolving such a conflict might be handled differently in different domains; in accounting domain, one might appropriately "mix" two plans (e.g., debt and equity financing) to create a range of possible options. Clearly, though, the planning interpretations of all these situations can stray far from their standard meaning in simple classification.

For reasons such as these, hierarchical classification does not seem self-sufficient as a technique for solving plan selection problems. An additional, higher-level problem solver is needed to shape the problem, and express the solution, in the context of planning. This new problem solver would call on the hierarchical classifier to evaluate case data using the hierarchy of plan types; it would then interpret the result of the classification and report the interpreted recommendation to the user (Figure 7). Relying directly on the confidence values associated with nodes in the classification hierarchy and, when appropriate, on the values of individual knowledge groups within the nodes the interpretation task would take two forms. Such a high-level planning module would use domain-independent knowledge to apply the results of hierarchical classification in cases involving task-specific plan selection conflicts (i.e., Figure 5b). Additionally, domain-dependent knowledge in the planning interface would enable the composite problem solver to resolve domain-specific classificatory conflicts, ones in which the features of the particular domain facilitate specialized interpretation and plan combination (i.e., Figure 6).

********************** Figure 7 goes about here **********************

## CONCLUSION

The primary intent of this project has been to investigate the applicability of hierarchical classification as a problem-solving technique in a suitable accounting domain. As such, this paper demonstrates the feasibility of this approach with respect to selecting capital asset acquisition plans, and the resulting prototype represents a valid framework for a commercially-oriented knowledge-based system. From a theoretical perspective, this investigation has provided a better understanding of how classification can be utilized as a planning tool and has led to an initial proposal for a general planning interface that works with hierarchical classification to solve planning problems. This project indicates several avenues for further research. The idea of a general planning interface requires substantial investigation into the principles that underlie plan selection, especially with respect to the distinction between domain-independent and domain-dependent resolution of specific selection scenarios. Also meriting further study is the notion that accounting knowledge for solving problems of this type may be characterized in a limited number of knowledge group types, which can then be assembled in various ways to establish different contexts.

Finally, the use of hierarchical classification in selecting capital asset acquisition plans seems to offer a means for understanding plan selection problems in this domain. Adopting the classification paradigm enables a consideration of financial problem solving from the level of the task being performed and thus clarifies the nature of the issues involved with the problem solving. In this respect, the research outlined herein provides a new perspective on problem solving in accounting.

## ACKNOWLEDGEMENTS

The authors wish to thank Dr. William McCarthy, Department of Accounting, Michigan State University, for his guidance and comments in the course of preparing this paper for submission and are also grateful to the Artificial Intelligence/Knowledge-Based Systems Group of Battelle Columbus Laboratories for supplying the version of CSRL utilized in the research described here.

Partial support for this work is through the MSU Ameritech Program, which has supported the research of Drs. Sticklen and McCarthy. Research activities in the AI/KBS group at MSU are also currently supported by the McDonnell Douglas Independent Research and Development program, by the Research Excellence Fund of Michigan, and by equipment grants from Apple Computer and Texas Instruments.

## BIBLIOGRAPHY


Accounting Standards: Original Pronouncements , Financial Accounting Standards Board, Stamford, CT, 1985.

Baker, C. Richard and Rick S. Hayes, Lease Financing , John Wiley and Sons, New York, NY, 1981.

Chandrasekaran, B., "Decomposition of Domain Knowledge into Knowledge Sources: The MDX Approach," Proceedings of the 4th National Conference of the Canadian Society for Computational Studies of Intelligence, Saskatchewan, May 1982.

Chandrasekaran, B., "Generic Tasks in Knowledge-Based Reasoning: Building Blocks for Expert System Design," IEEE Expert , Fall 1986, pp. 23-30.

Chandrasekaran, B., "Generic Tasks in Knowledge-Based Reasoning: Characterizing and Designing Expert Systems at the 'Right' Level of Abstraction," Proceedings of the 2nd Annual IEEE Conference on Artificial Intelligence Applications , December 1985.

Chandrasekaran, B., "Towards a Taxonomy of Problem Solving Types," AI Magazine , vol. 4, no. 1, 1983, pp. 9-17.

Chandrasekaran, B., and S. Mittal, "Conceptual Representation of Medical Knowledge for Diagnosis by Computer: MDX and Related Systems," Advances in Computers , vol. 22 (M. Yovits, ed.), Academic Press, 1983, pp. 217-293.

Clancey, William J., "Classification Problem Solving," Proceedings of the AAAI , May 1982, pp. 49-55.

Clark, John J., Thomas J. Hindelang, and Robert E. Pritchard, Capital Budgeting , Prentice-Hall Inc., Englewood Cliffs, NJ, 1984.

Elgers, Pieter T. and John J. Clark, The Lease/Buy Decision , The Free Press, New York, NY, 1980.

Fink, Philip R., Tax Guide for Acquiring, Maintaining and Disposing of Business Equipment, Prentice-Hall Inc., Englewood Cliffs, NJ, 1982.

Gomez, F., and B. Chandrasekaran, "Knowledge Organization and Distribution for Medical Diagnosis," IEEE Transactions on Systems, Man, and Cybernetics , January 1981, pp. 34-42.

Jarnagin, Bill D. and Jon A. Booker, Financial Accounting Standards: Explanation and Analysis , Commerce Clearing House Inc., Chicago, IL, 1984.

Kaufmann, Mike, ed., The Capital Budgeting Handbook , Dow Jones-Irwin, Homewood, IL, 1986.

Mays, Eric, Chidinand Apte, James Griesmer, and John Kastner, "Organizing Knowledge in a Complex Financial Domain," IEEE Expert , Fall 1987, pp. 61-70.

Mittal, S., Design of a Distributed Medical Diagnosis and Data Base System , Ph.D. Thesis, Computer and Information Sciences Department, Ohio State University, 1980.

Samuel, A., "Some Studies in Machine Learning using the Game of Checkers," IBM Journal of Research and Development , 1967, pp. 601-617.

Sticklen, J., J. W. Smith, B. Chandrasekaran, and John R. Josephson, "Modularity of Domain Knowledge," International Journal of Expert Systems: Theory and Applications , 1987, pp. 1-15.