An Integration Framework for Multiple-Expert Systems

Fatemeh (Mariam) Zahedi
Research Professor
University of Wisconsin Milwaukee

Jamshid Hosseini
Associate Professor
Marquette University

March 1994

Abstract

As expert systems become a commonplace component of information systems, they pose issues related to the development of large and real-world expert systems that have multiple experts as well as those related to combining multiple-expert systems and distributed-expert systems. This paper is a preliminary report on a framework for integrating multiple-experts views and multiple-expert systems using a decision hierarchy and the AHP.

1. Introduction

In large expert systems, one seldom encounters a case where the knowledge of only one expert would be sufficient for creating and processing the knowledge base. In expert systems with a mission-critical nature, the participation of multiple experts is an important factor for ensuring the the knowledge base is comprehensive and complete. Furthermore, as expert systems become more commonplace and are created and used across geographically dispersed networks, integrating multiple-expert systems and distributed-expert systems will be one of the main directions of growth in this field.

Multiple experts and multiple-expert systems pose some unique issues in systems development and implementation, such as:

- Variability in experts' expertise in different areas of the knowledge base
- Conflicting views of experts regarding the rules that apply to various situations
- Disagreement among experts on the reliability of a given rule in achieving a given objective of the expert system

One way to avoid the problems of multiple experts is to create different modules of the knowledge base, each based on one expert or a group of experts with similar opinions. In this approach, as long as the results of different modules match, the system's recommendation can be considered strong and reliable. However, when the results do not match, it is not easy to reconcile the conflicting recommendations. One may say that the conflict shows the uncertainty of the recommendations. But if the experts are indeed categorized into homogeneous groups, one may get conflicting recommendations every time. This reduces the usefulness of the system for its customers. After all, expert systems are expected to be a decision-making aid.

In this paper, we address the issue of expert systems with multiple experts and distributed-expert systems, and develop a methodology that can be helpful in combining the knowledge of multiple experts or multiple-expert systems into one cohesive system. The outcome of this methodology is to have rules that are rated based on their roles in achieving the system's objectives. In this sense, these ratings could also be used in context-sensitive search heuristics in optimal processing of the knowledge base.

This paper is organized as follows. Section 2 provides a background review of the literature. Section 3 identifies the problem components. Section 4 develops the hierarchy that combines the components of multiple experts and systems. Section 5 discusses the rating of rules in expert systems. Section 6 discusses how the rated rules could be used in determining the reliability, certainty, or importance of rules in expert systems with multiple experts, aiding the rule selection heuristics in an expert system, and integrating multiple-expert systems. Section 7 provides the concluding remarks.

2. Background

Real-world applications of expert systems normally require the involvement of multiple experts or even multiple-expert systems. It is an established fact that there is no single truth in any field, and experts differ in their opinions and approaches. The issue is how to combine experts' views into a consistent and workable knowledge base. Furthermore, if there are expert systems developed in the same domain and for the same type of decision, then the integration of multiple-expert systems adds a great deal to the value of the recommendations obtained from these systems.

The expert-systems and artificial-intelligence literature appears relatively limited with respect to acquiring knowledge from multiple experts. Liou and Nunamaker [1993] and La Salle and Medsker [1991] recommended computer conferencing for knowledge acquisition from multiple experts. Dayou et al. [1992] developed three methods of consistency treatment and two methods of synthesizing multiple views of experts.

Research in multiple-expert systems has been somewhat more frequent than that of multiple experts. Hodge [1989] discussed conceptually the issue of integrating expert systems. Fraser et al. [1989] developed a conflict resolver for cases when competing expert systems make conflicting recommendations. Their method which uses a Delphi-like approach, works iteratively to pass from a system of optimal recommendations by various expert systems, which are usually inconsistent, to eventually a set of close-to-optimal, consistent recommendations. Zinser et al. [1989] developed a communication architecture for the exchange of knowledge between multiple-expert systems. Their system, called the GRADIENT (GRaphic DIalogue EnvironmENT, is essentially the architecture for a distributed-expert system. Other environment building for multiple experts has also been suggested (e.g., Al-Schamma and Chen, 1990).

Several applications of expert systems with multiple experts or multiple systems have been reported. These applications include: 1) intelligent control systems [Miaoliang, et al. 1989], and flexible manufacturing systems [Long, et al. 1992]; pavement diagnosis and rehabilitation [Corby et al. 1990], planning and coordination for future fighter aircraft [Smith and Broadwel 1988], and military strategic planning [Zeng, et al 1989].

In developing systems with multiple experts or multiple-expert systems, a number of issues have to be dealt with. These issues include uncertainty, consistency assessment, conflict resolution, and search methodologies. A number of studies have attempted to address these topics.

Uncertainty. Uncertainty is a multifaceted concept. Assessment and analysis of uncertainty with respect to multiple experts (or multiple systems) is of vital importance to the success of coordination and integration of multiple-expert systems. There is virtually no literature addressing this aspect of uncertainty effectively. As a side issue the uncertainty in multiple experts (or multiple-expert systems) has been touched on by researchers in the consensus-building and consistency-evaluation areas (e.g., Ionnidis and Sellis 1989; Liebowitz, et al. 1991). Hagen [1992] suggested methods for improving the probabilistic assessment in knowledge acquisition. Grim [1986], Kriz [1986], and Deng et al. [1990] have addressed uncertainty in facts, rules, and meta rules.

Conflict Resolution. Conflict resolution may be done through either human interface, using consensus building and group methodologies, or using analytic tools to seek an optimal or, at least, heuristically sound compromise solution or position for multiple parties. Conflict resolution is extremely important in expert systems in general. In a typical ES hundreds of rules and facts are imbedded. The consistency of decisions reached from different rules must be verified. In cases where rules give conflicting results, this conflict must be resolved. In a multi-experts or multi-system expert system, another level of potential conflict is added: The experts or expert systems may give contradictory solutions given the same set of facts. There has been very little research in this area. Liebowitz et al. [1991] discuss strategies for conflict resolution in an expert scheduling system. They use the analytic hierarchy process (AHP) for conflict resolution in a case study involving NASA satellites' scheduling problems. Most other research on conflict resolution is more general in nature, imbedding the issue of conflict resolution in ES (see, e.g., Dutta and Bonissone, 1986; Ioannidis and Sellis, 1989; and McLean, et al., 1992)

Search Methodologies. ES search methodologies are heuristics and context-insensitive. Webster [1991] provides simple descriptions and examples for various heuristic search methodologies (e.g., depth-first, breadth-first, best-first, etc.) used in expert systems. Sharma et al. introduce a depth-first algorithm in a distributed environment in the sense that multiple parallel processors are at work. Their basic idea may be implemented in a multi-expert (system) environment. Other recent articles mostly have concentrated on the algorithmic aspects of search methods in multi-tasking environments (see, e.g., Colbrook and Smythe. 1990; Li and Marlin, 1989; Song and Donovan 1989). We are not aware of any work on the development of search heuristics that utilizes the specific knowledge in the knowledge base in order to optimize the search process.

Validation and Verification. Probably one of the most important issues in multi-level, multiple-system, or multiple-expert expert systems are validation and verification. Meseguer and Verdaguer [1993] present a comprehensive theoretical and practical discussion of several aspects of validation and verification including uncertainty and control of information, inconsistency, redundancy, and circularity.

As is evident in this brief review of the literature, issues related to knowledge bases with multiple experts and multiple-expert systems have not yet been thoroughly addressed within a unified framework. On the other hand, many expert-system professionals and researchers agree that the future of AI and expert systems is in the integrated and distributed systems involving multiple experts. We provide a framework for dealing with the multiple-expert systems issues.

3. Problem Components

In the analysis phase of expert systems, the definition of the system goal is one of the first and most important steps. The system goal is used to identify the concrete objectives of the system. For example, an expert system may have as a goal "helping the bank loan officer in granting mortgage loans." At the design phase of the expert system, this goal becomes more concrete in the form of an objective: granting the loan, denying the loan, postponing the decision pending on more information. The identification of these specific system objectives is an essential step in guiding the identification of rules that would lead the system to recommend a particular system objective. In other words, the general goal of the system is broken down into specific system objectives.

The objectives themselves may have a hierarchy of sub-objectives. For example, the objective of "grant a loan" may have sub-objectives of "the amount of the loan," "duration of the loan" and "monthly mortgage." Depending on the complexity and size of the system, the hierarchy of the objectives may have more levels.

After the identification of the system objectives, one has to decide on the rules that make each one of them fire. In other words, one has to identify the rules that lead to the decision of "grant a loan," or "the amount of the loan." In the design of expert systems, we may use experts to identify general goals or concrete objectives of the system. One would expect little disagreement or variability among experts in identifying the system objectives. However, when it comes to defining the rules that would cause a particular objective become true, experts may have different opinions regarding the validity of particular rules, and their expertise may vary from one objective to another.

For example, expert bank officers may vary in their opinions on the rules leading to the decision to "grant a loan." Furthermore, the expert who has expertise on deciding to grant a loan may not have enough expertise on deciding about "the amount of the loan." In other words, experts should be rated for their expertise related to particular objectives. We should be able to capture the degree of experts' expertise in various areas of knowledge domain.

Another complicating issue in multiple experts is the treatment of experts' disagreement about the uncertainty of a rule. Rarely can a rule produce an objective with 100 percent certainty, because there are a number of sources that lead to uncertainty [Zahedi, 1993, Chapter 11].

One of the sources of uncertainty is the ambiguity and uncertainty within the knowlege domain. That is, a particular rule itself may be uncertain in the eye of an expert. The question is how to incorporate the uncertainty of various experts into the knowledge-base rules. Here, we allow the reliability or uncertainty of a rule to be evaluated by each expert separately. The reliability of the rule should reflect the combined ratings of all experts who have identified that rule in their reasoning process.

In the next section, we will see that the structure of the problem lends itself naturally to a decision hierarchy that can be evaluated using the AHP approach.

4. The Expert-System Hierarchy

The discussion in the previous section identifies a hierarchy that starts with the general goal of the system and leads to the system objectives. Domain experts identify rules for achieving each one of system objectives. However, since experts vary in their expertise regarding each system objective, their views should be weighted by the extent of their expertise in each objective. Finally, once each expert identifies rules, they should be evaluated by experts for their reliability. This leads to a four-level decision hierarchy as shown in Figure 1. The following subsections elaborate on this hierarchy.

Level 1. Expert-System Goal

The underlying model for rating the rules in the knowledge base has an inherently hierarchical structure (Figure 2). At the uppermost level of the hierarchy, we have the general goals of the system. In the example of the bank mortgage loan, it would be "mortgage loan decision."

Level 2. Expert-System Objectives

The second level of the hierarchy has the concrete system objectives, and the subsequent levels would accommodate the sub-objectives. These objectives are the final recommendations of the system, such as granting or rejecting the loan. In some expert systems, however, there could be intermediary objectives that may be as important as the final conclusions. For example, in a medical diagnosis, the intermediary conclusion that a CAT scan is needed could be an important intermediate step. The objectives identify the significant output of the system that the user could use, one way or another, in his or her decision process. Similarly, in a loan-application expert system, although the final objectives are granting or rejecting a loan application, the intermediate objective of the credit-worthiness of the loan applicant may also be of interest to the users of the system. In Figure 2, each box representing an objective could be considered a mini-hierarchy itself, containing a hierarchy of system sub-objectives.

Level 3. Multiple Experts

Experts' expertise could be categorized in many ways, such as for their sub-specialty, their experience, their level of generality or specialization, and their breadth and depth of knowledge. The knowledge domain also determines the types and extent of specialization and capability of experts.

The knowledge domain is divided into sub-areas that are reflected as system objectives. These objectives are indicative of the areas or categories of knowledge that lead to various types of decisions. Therefore, if we were to decide on the capability, expertise, and experience of an expert, one basis of their evaluation or specialization is their ability to identify rules for each system objective.

For example, an expert loan officer may be most capable of identifying the rules for granting or rejecting a loan. He or she may not be the most suitable or knowledgeable person to decide on the amount of loan insurance if the loan is granted or about the intermediate objective of determining the credit-worthiness of a loan applicant. Similarly, a physician who can identify symptoms of flu may not have the same ability to identify the symptoms of AIDS.

Therefore, the third level of the decision hierarchy contains the experts whose expertise would be used in the expert system. This level allows the incorporation of multiple views within one system. An alternative case is one in which each expert's opinion is incorporated in a separate knowledge base, and the ultimate goal of the analysis is to combine the rules fired by various knowledge bases in order to arrive at a conclusion that integrates the views of many experts in the field.

Level 4. Knowledge-Base Rules

The fourth level of the hierarchy contains the decision tree of rules expressed by multiple experts for arriving at the objectives of the system. The depth of the hierarchy of rules depends on the complexity of the system and the number of ways in which each objective of the system could be accomplished. Multiple experts evaluate each rule for its degree of reliability and uncertainty.

5. Rating of Rules in Expert Systems

We use the hierarchy to attain the rating of rules in the knowledge base for achieving the general goal and specific objectives of the expert system. The rating starts at the second level of the system. We rate system objectives for their role in achieving the general system goal. In the mortgage loan example, the general goal of the system was to arrive at a loan decision for loan applications. We can compare and rate the objectives of the expert systems in achieving this goal. The decision to grant or reject the loan may have equal importance in the loan decision. However, the objective "postpone the decision," or "credit-worthiness of the loan applicant," although important, may not have the same rating as the first two objectives.

The difference in the ratings of objectives is more dramatic in risk-oriented expert systems, such as medical diagnosis. The objective of a cancer diagnosis is far more important than diagnosing the common cold or flu. Rating the objectives at the second level of the hierarchy allows the system developers to distinguish among objectives and give a higher weight to those objectives that are critical to the main purpose for which the system is designed.

At the second level of the hierarchy, the objectives are rated in a relative fashion and their weights sum to one. Therefore, the method of rating is eigenvalue method at this stage.

At the third level of the hierarchy, we rate the experts on their expertise, knowledge, or capability to identify rules that lead to the recommendation of an objective. This level allows developers to incorporate multiple views of experts within the system, each rated or weighted by the degree of experts' expertise in a particular area of the domain knowledge that is identified by the system objective.

For example, expert A could be very knowledgeable in the rules for accepting or rejecting a loan, but his or her capability is less reliable than expert B in deciding the amount of loan insurance needed, if the loan is granted to an applicant. In that case, expert A gets a high rating for the first two objectives, and a comparatively lower rating than expert B in the third objective. Another example is the case in which a medical expert system contains the objectives of diagnosing liver cancer and a heart disease. Obviously, the degree of expertise varies for each one of the objectives. A cardiologist would be highly qualified for the objective of heart disease and less for the diagnosis of liver cancer, whereas an oncologist will have a higher rating for the objective of liver cancer and lower rating for the objective of heart disease.

At the third level of the hierarchy, experts are compared in their relative expertise in various areas of the domain knowledge identified by the system objectives. Experts' expertise could be evaluated more easily in comparison with one another. Therefore, the rating method for this level also is the eigenvalue method.

While the above discussion concerns the case of multiple experts and the incorporation of their diverse knowledge into a single expert system, the method applies equally to the case where the knowledge bases of multiple-experts systems have been used in order to create a unifying framework for combining the multiple systems' recommendation. In other words, knowledge bases could be humans living outside the machine in the form of experts,

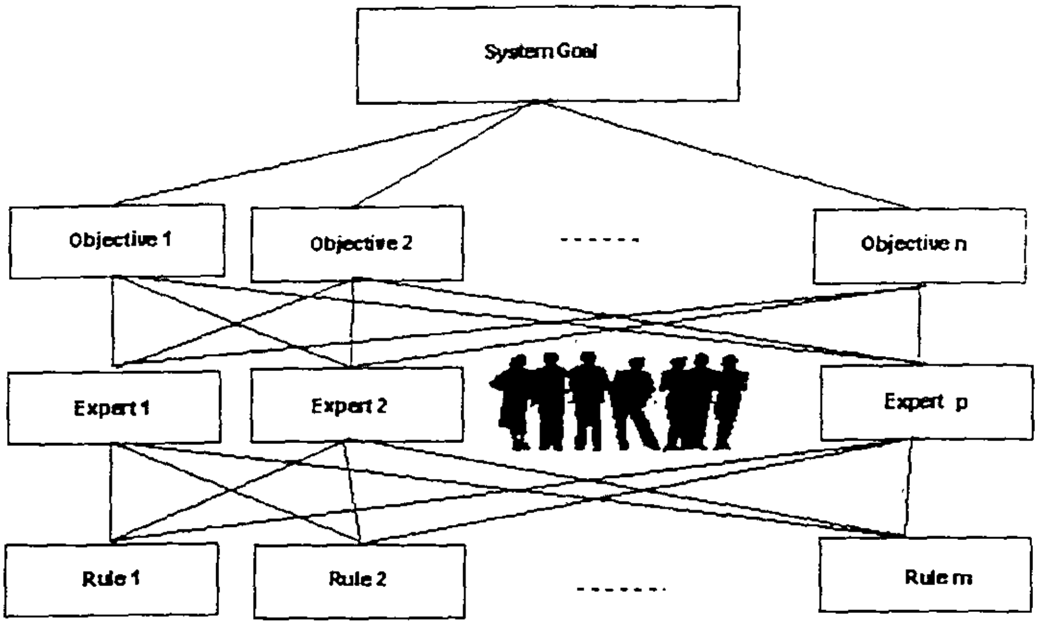

Figure 2:: Expert-System Hierarchy
in Expanded Form

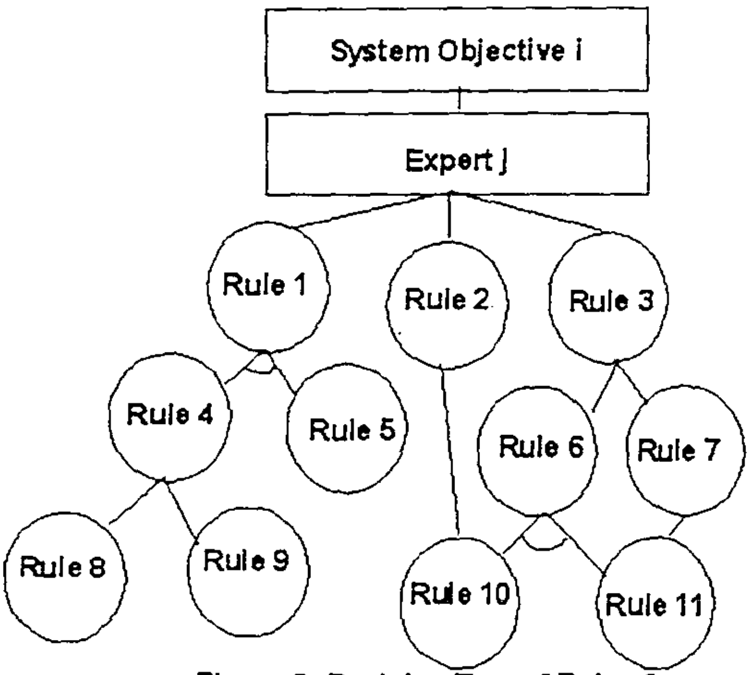

Figure 3: Decision Tree of Rules for
Objective i and Expert j

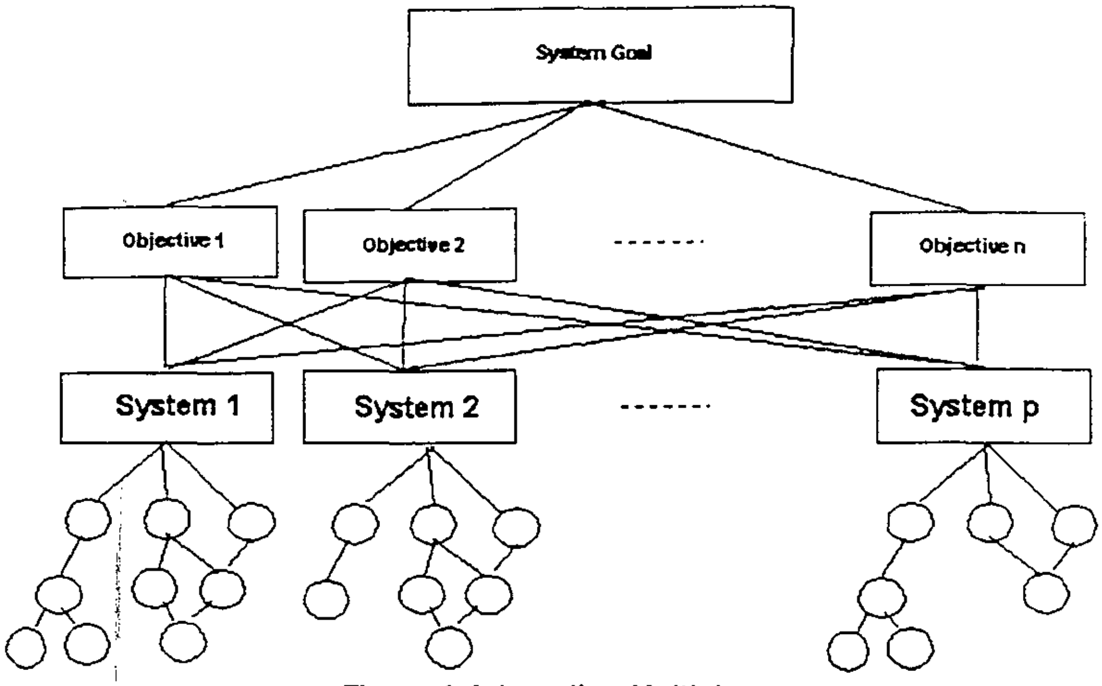

Figure 4: Integrating Multiple- or Distributed- Expert Systems

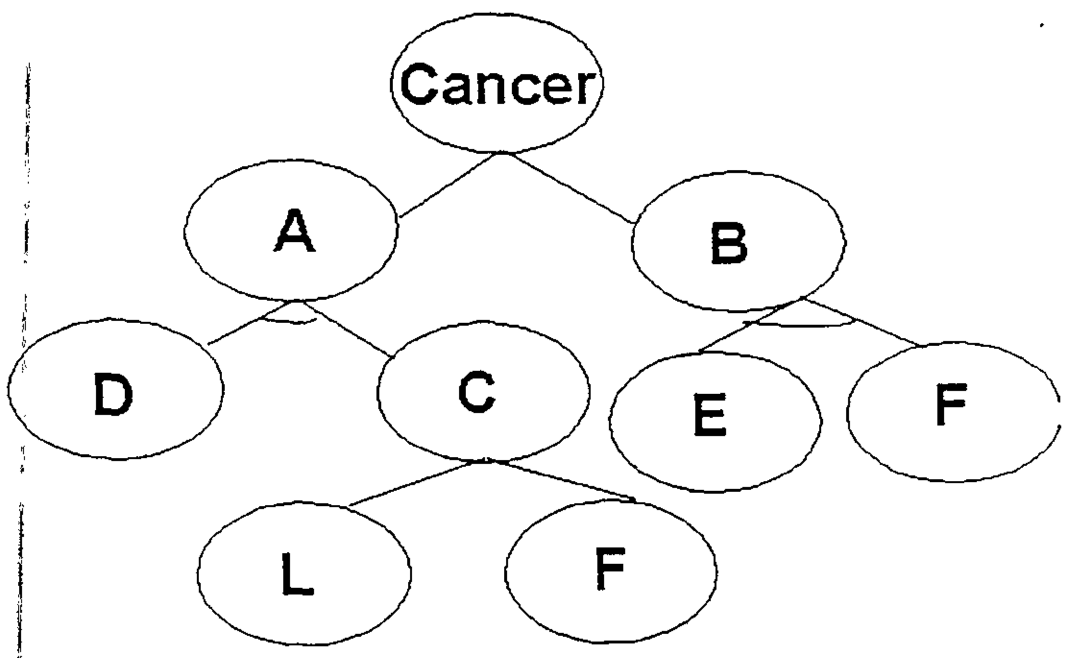

Figure 5: Rules for Cancer Diagnosis