FINANCIAL EXPERT SYSTEMS

Peter E. Hart
Syntelligence

Panelists: Norton Greenfeld, Applied Expert Systems;
Walter Reitman, Palladian Software; Chuck Williams, Inference Corp.

Financial decision-making problems represent an important new set of applications for expert systems. While these problems share certain similarities with the scientific and engineering domains that have historically been the focus of expert system research and development, they have their own special characteristics. What distinguishes financial applications? What demands do they place on expert system technology? What research opportunities do they suggest?

Most applications of expert systems in finance center on computing an assessment of the risks and rewards associated with alternative courses of action. These assessments, though they may involve the analysis of large amounts of data, nonetheless typically involve a great deal of uncertainty. A banker considering making a loan, or a corporate officer considering building a new factory, cannot have complete knowledge about the myriad of events that over time govern the outcome of a financial decision. Private individuals planning their financial futures face analogous problems.

A designer of expert systems for financial applications is confronted by several challenges. First, as with any expert system problem, the knowledge representation and inference procedures selected must be well-matched to the structure of the domain. Second, the design of the user interface must reflect the fact that the system will be used daily by a large population of business-oriented people. Third, for many applications the system must be embedded in a larger hardware and software environment that may not include the machines, operating systems and languages that are typical of AI research organizations. Let us examine these challenges in greater detail, noting how financial applications differ from more-familiar scientific and engineering domains.

An obvious characteristic of financial problems is the mixture of numerical and symbolic reasoning that they demand. In contrast, domains like medicine and geology place make few demands on reasoning quantitatively.

A second characteristic of financial problems is the absence of strong underlying causality models of the domain. While there may be simple algebraic models (such as net present value calculations) for certain components of an analysis, there are few problems where a global model is available. This stands in sharp contrast to certain engineering problems, such as trouble-shooting digital systems, where excellent multi-level causal models exist. The availability of deep models in engineering problems is useful even when they are not represented explicitly, since they inform the knowledge engineering process that produces surface models. The absence of deep models increases the difficulty of knowledge engineering in financial domains.

A third characteristic of many financial problems is the apparent emphasis by experts on non-procedural decision-making knowledge. This may reflect the difficulty of constructing accurate deep models, or it may reflect the fact that complete case data is virtually never available in practice.

The business orientation of users of financial system poses special problems for system designers. For certain applications, the user population may have gained familiarity with the use of computers through work with spreadsheets or other packages. For other applications, business professionals may be skeptical or even hostile regarding the use of computers, which heightens the need for user interfaces that closely mirror familiar paper-based work habits. This contrasts with a variety of engineering and technological applications where the user population may be comprised of computationally sophisticated people.

Many financial applications require the expert system to be tightly integrated with a variety of existing hardware and software resources. Mainframes and operating systems, teleprocessing systems, transaction and accounting systems, and database systems may all be "givens" with respect to a proposed expert system application. Accommodating to these needs may mean that the "expert system component" is only a small part of an overall application.

What are the implications for future research? Each of the issues noted above suggest a number of obvious research problems. One issue *not* mentioned above yet deserving attention is the weakness of explanation subsystems in current expert systems. The ability of expert systems to explain their reasoning has long been an extremely attractive feature. However, as knowledge representations have become richer the difficulty of designing an explanation subsystem has grown. Financial applications place particular stress on this weakness since the user, not the system, bears ultimate responsibility.