

Knowledge Acquisition Techniques for Decision Analysis Using *Axotl* and *Aquinas*

Jeffrey M. Bradshaw, Stanley P. Covington, Peter J. Russo, and John H. Boose
Advanced Technology Center, Boeing Computer Services
P.O. Box 24346, M/S 7L-64, Seattle, Washington 98124 USA (206) 865-3422;
jbrad@atc.boeing.com

ABSTRACT

The effective application of current decision tree and influence diagram software requires a relatively high level of sophistication in the theory and practice of decision analysis. Research on intelligent decision systems aims to lower the cost and amount of training required to use these methods through the use of knowledge-based systems; however, application prototypes implemented to date have required time-consuming and tedious hand-crafting of knowledge bases. This paper describes the development of *DDUCKS*, an “open architecture” problem-modeling environment that integrates components from *Axotl*, a knowledge-based decision analysis workbench, with those of *Aquinas*, a knowledge acquisition workbench based on personal construct theory. The knowledge base tools in *Axotl* can be configured with knowledge to provide guidance and help in formulating, evaluating, and refining decision models represented in influence diagrams. Knowledge acquisition tools in *DDUCKS* will allow the knowledge to be efficiently modeled, more easily maintained, and thoroughly tested.

1. INTRODUCTION

1.1. Progress in Automated Decision Analysis

The Boeing Company has an urgent need for advanced automated decision support applications. Rapid growth in the complexity of strategic and tactical decisions has

outstripped the capacity of conventional decision aids. Informal, checklist, and rating methods are helpful for simple decisions, but are inadequate for the analysis of tradeoffs involving allocations of critical resources. Spreadsheet, data base, and linear programming models for decision making likewise break down with large amounts of uncertain, incomplete, or conflicting data. Such approaches cannot effectively embody the intuition and flexibility of human decision makers.

Knowledge-based systems have been widely hailed as a solution to the problem of modeling expert problem solving knowledge. Unfortunately, typical knowledge-based approaches also have their limitations. Rule-based methods employing heuristic certainty factors have been shown to perform poorly in problems involving large amounts of uncertainty or risk, and the kinds of complex tradeoffs that inevitably emerge in strategic decision making (Horvitz, Breese, and Henrion, 1988). Furthermore, knowledge-based approaches are not sufficiently flexible for many decisions, since tradeoffs may vary greatly across individual cases (Langlotz, Shortliffe, and Fagan, 1986). Finally, knowledge-based system development environments do not generally provide facilities for integrating historical data with expert judgment (Spiegelhalter, Franklin, and Bull, 1990).

One of the most promising approaches to dealing with decision complexity in a consistent, general-purpose manner is decision analysis (Howard, 1966; Howard & Matheson, 1984; Keeney & Raiffa, 1976; Raiffa, 1968). In the past few years, researchers and developers have made important theoretical advances and have implemented several successful systems for automated support of the decision analysis process (reviews by Horvitz, Breese, & Henrion, 1988; Pearl, 1988; Neapolitan, 1990). Although a thorough discussion is beyond the scope of this paper, we wish to review three of the important developments that have led to the current state of the art. Following this, we will explain why we think that the development of automated knowledge acquisition tools is crucial to the future of efforts to deliver decision analysis methodology to a wider spectrum of decision makers and domains.

Decision tree software. The development of decision tree software (Figure 1) represented an important milestone in automated decision analysis support (Olmsted, 1982; McNamee & Celona, 1987). Through the use of general-purpose commercial tools, decision analysis techniques have become more widely known and used than ever before. At the same time, there are several drawbacks to the use of decision trees as a representation device. For one thing, they grow exponentially with problem size, making them impractical for problems of significant size. Additionally, the tree metaphor for decision problems often leads participants to confuse chronology with the order of probabilistic expansions.




Figure 1. Decision analysis techniques have become more widely used through the availability of decision tree software.

Influence diagram software. Conceived by Howard and Matheson (1981), influence diagrams represent an important advance in the representation of decision problems (Figure 2). While retaining the essentials of the decision theoretic mathematical foundation developed for decision tree manipulation, they provide several advantages over decision trees. Technically, they are superior in that they can explicitly represent and exploit conditional independence relationships between variables. From a practical point of view, they provide a clear and intuitive way of communicating the structure of a decision model. With the advent of commercial software (e.g., Shachter, 1986b) and the assumption of continued success by researchers developing tractable methods for evaluating large influence diagrams (review by Henrion, 1990), we expect influence diagrams to eventually replace decision tree representations for most applications.




Figure 2. Influence diagrams provide both technical and practical advantages over decision trees.

Intelligent decision systems. Unfortunately, the effective application of both decision tree and influence diagram software requires a relatively high level of sophistication in the theory and practice of decision analysis. These tools contain some of the *algorithms* of decision analysis practice, but cannot embody the *experience* and *intuition* of decision analysis professionals in formulating and appraising decision models. Also, because these tools cannot conveniently store and reuse domain expertise, they cannot exploit the similarity between recurring decisions in the same domain. New decisions are typically modeled from scratch.

The importance of these issues led Holtzman (1989) to define an approach for a third generation of automated decision analysis software called *intelligent decision systems (IDS)*¹. An IDS combines a set of automated decision analysis tools with a knowledge-based system that helps decision makers without extensive training in decision analysis build, evaluate, and refine decision models in some well specified domain (Figure 3). To build an IDS application for a class of decision problems, decision analysts and domain experts work with a knowledge engineer to configure the system with “application-independent”² knowledge (i.e., knowledge of decision analysis tools and methodology) and application-specific knowledge (i.e., knowledge about a particular domain). Once built, these knowledge bases can be used again and again to provide guidance and help during consultations with decision makers. Task-level consultation interfaces pose questions and interpret results in language

¹ Related approaches to combining decision analysis and knowledge-based systems have been developed by Breese (1987), Keeney (1986), Moore & Agogino (1987), and Wellman (1986).

² Of course, no knowledge-base is truly application-independent; perhaps “multiple-use” would be a better term.

and graphical presentations tailored to the decision maker, rather than in terms of standard decision analysis concepts.

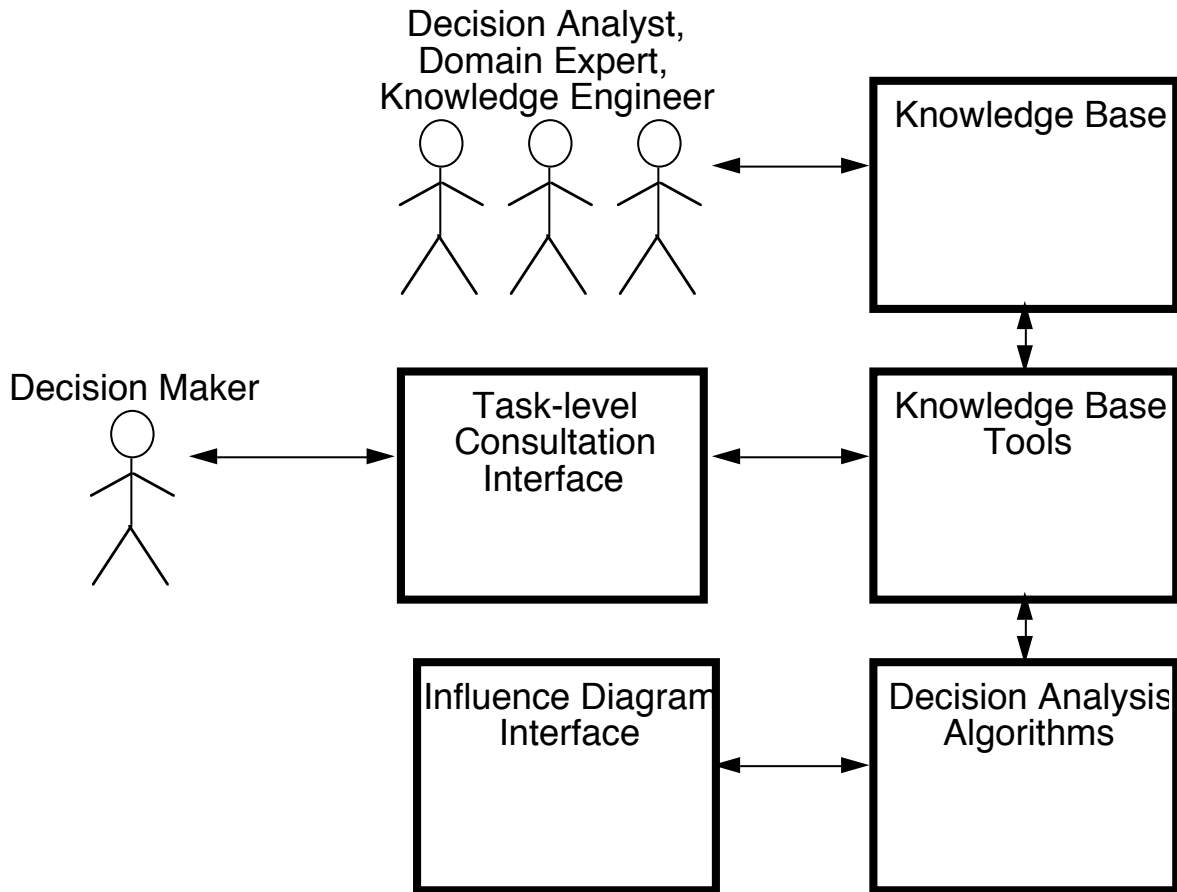


Figure 3. Intelligent decision systems (IDS) combine decision analysis tools with a knowledge-based system

Through implementing a general-purpose IDS architecture (*Axotl*) and applying it to the domain of R&D project selection within The Boeing Company (*PIE*) we have become convinced that decision analysis can be delivered effectively and economically through the use of knowledge-based systems (Bradshaw & Holtzman, 1987; Bradshaw, Covington, & Russo, 1989). We found that the development of knowledge bases for such a system was relatively straightforward compared to other knowledge-based systems we had created because we were dealing with a methodology (decision analysis) that was mature and rigorously defined. However, we were dissatisfied that so much hand-crafting of knowledge bases had to be performed by knowledge engineers. Through the development of knowledge

acquisition tools, we hoped that much more of the knowledge base could be constructed by domain or decision analysis experts themselves.

1.2. Knowledge Acquisition for Fourth-Generation Decision Analysis Support Tools

Figure 4 shows how automated knowledge acquisition tools would fit within an architecture for advanced decision analysis support tools. Knowledge acquisition tools could help cut down the lengthy and error-prone revise-and-review cycle in the development of such systems, making delivery of IDS applications practical on a large scale. While it is unrealistic to expect that the role of knowledge engineers would entirely disappear, their participation in many aspects of knowledge base development and maintenance could be minimized. Figure 4 explicitly includes interfaces to conventional software and external data. The success of future systems in practical applications will depend on whether they can be effectively integrated with other software such as databases, spreadsheets, and hypermedia environments.

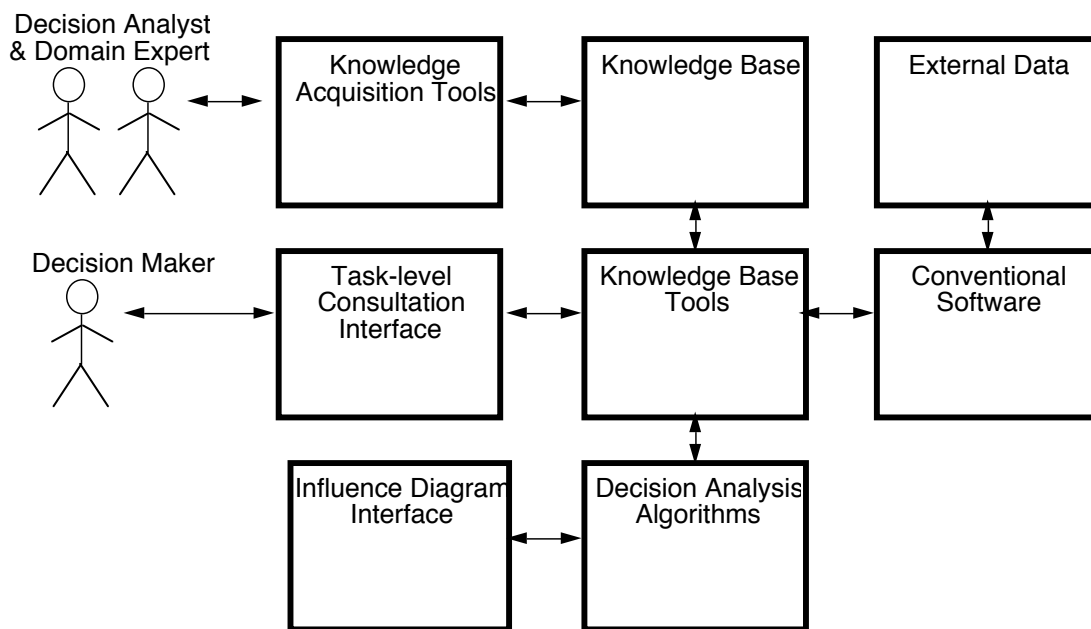


Figure 4. Knowledge acquisition tools and interfaces to conventional software will be an important part of future decision analysis support systems.

2. APPROACH

2.1. A Framework for Knowledge Acquisition

2.1.1. Automating Knowledge Acquisition at the Boeing Advanced Technology Center

The ideal environment for the growth of new tools is one that provides a balance between technology push and application pull. Developers have ideas they feel might be useful for solving a future application problem —*technology push*— and applications present more immediate problems that need to be solved —*application pull*. Technology push tends to be farsighted and application pull tends to be shortsighted. When technology push is the only force, there is a danger that any tools developed will not be useful since they are not anchored in real problems, only a developer's vision of what the problems may be. When application pull is the only force, tools tend to be developed for special purposes and are hard to generalize to other problems. In such a demand-driven environment, revolutionary (or even evolutionary) breakthroughs are rare. Ideally, push and pull forces should both act on a technology development project at the same time, leading to tools that are farsighted, general, and yet useful for a variety of problems.

Boeing's Advanced Technology Center provides a unique environment that fosters the interplay of technology push and application pull. The Center sponsored an Associates' Program that helped transfer artificial intelligence technology to the rest of Boeing. Associates spent one year at the Center, where they attended courses and developed prototypes to solve a specific problem in their home division. When the year's training was complete, Associates returned to their groups where they continued to develop and eventually field the applications. Over 95 Associates graduated from the program; the majority of them built knowledge-based systems. At the same time, the Center was working with universities to import advanced technology projects and was developing its

own projects in the areas of vision, robotics, speech understanding, natural language, machine learning, and knowledge-based systems.

The large projected number of knowledge-based system projects in Boeing provided motivation for a strong program in knowledge acquisition. Knowledge acquisition, the most labor-intensive phase of knowledge-based system development, was seen as a way of minimizing the major cost drivers in the development of such systems as well as a means of ensuring reliability and maintainability of the resultant knowledge bases. The Advanced Technology Center has supported the organization of annual AAAI-sponsored knowledge acquisition workshops in Banff, Canada; annual European knowledge acquisition workshops at various locations; and, beginning in 1990, a Pacific Rim knowledge acquisition workshop. Scientists at Boeing have participated in the development, application, and evaluation of several automated knowledge acquisition tools, including ETS (Expertise Transfer System; Boose, 1985, 1986), *Aquinas* (Boose & Bradshaw, 1987; Boose, Shema, & Bradshaw, 1989), SALT (Marcus, 1987; Marcus, 1988), and KNACK (Kitto, 1988; Klinker, Boyd, Genetet, & McDermott, 1988).

As automated knowledge acquisition tools have been applied across a number of domains within The Boeing Company, impressive increases in efficiency have been demonstrated. Computer-based verification and validation methods as part of the knowledge acquisition process have improved the performance of knowledge-based systems by focusing attention on gaps and inconsistencies in the knowledge base. Knowledge engineers can maintain large knowledge bases much more easily and efficiently within such automated knowledge acquisition tools than through conventional methods.

Because the implementation of automated knowledge acquisition tools is itself labor intensive, their development can often be justified only if it can be demonstrated that they can be easily applied to more than a single application. Knowledge acquisition tool

developers interested in deriving the most benefit from their tools may look for areas consisting of several problems that can each be characterized by a general task model (Boose, 1989). For example, ETS and *Aquinas* have been applied to a number of diverse problems that fit a heuristic classification approach (Boose, Bradshaw, Kitto, & Shema, 1990); SALT has been generalized to work for scheduling and design problems (Marcus, 1989; Stout, Green, & Marcus, 1989); and KNACK has been used to produce a series of reporting systems (Klinker, Boyd, Dong, Maiman, & McDermott, 1989).

2.1.2. What is a Knowledge Acquisition Tool?

Over the years, many of our views on knowledge acquisition have changed. Current work in knowledge acquisition emphasizes that creation of knowledge bases is a constructive modeling process and not simply a matter of “expertise transfer” (Bradshaw & Boose, 1990; Akkermans *et al.*, 1990). Additionally, we used to think of knowledge acquisition as something that occurred mainly in the early stages of knowledge-based system development. Now we have come to realize that work in automated knowledge acquisition addresses the problem of designing appropriate representations and procedures to facilitate creation, validation, verification, and maintenance of knowledge bases over the *lifetime* of a knowledge-based system (cf. Ford *et al.*, 1990c). Thus, it might be said that researchers are attempting to do for knowledge engineering what CASE is doing for traditional software engineering (Bradshaw & Boose, 1989; Gaines, 1988)

Commercial expert system shell vendors have begun to realize the need for knowledge acquisition capabilities and many have been claiming to provide such support within their tools. How can a potential purchaser of such a product tell if an expert system shell truly supports automated knowledge acquisition?

A knowledge acquisition tool is more than a good graphical interface. Automated knowledge acquisition tools augment the facilities typically available in knowledge-base development environments by providing interfaces that focus the interaction between the tool and the expert or knowledge engineer. Some earmarks of a good knowledge acquisition tool are the following:

- role-limiting methods
- mediating knowledge representations
- interviewing techniques
- analysis tools.

Role-limiting methods. Many knowledge acquisition tools derive their power from relying on a well defined problem-solving model that establishes and controls the sequences of actions required to perform some task (Gruber, 1989; Karbach, Linster & Voss, 1990; Klinker, 1989). For example, SALT (Marcus, 1988) is based on a method for design called “propose-and-revise”, while MOLE (Eshelman, 1988) uses a method of heuristic classification called “cover-and-differentiate”. The problem-solving method defines the kind of knowledge applicable within each step, thereby making explicit the different roles knowledge plays. Once these roles are defined, knowledge acquisition tools appropriate to each kind of knowledge are designed.

Mediating knowledge representations. Knowledge acquisition tools attempt to minimize the problem of representation mismatch, the disparity between a person’s task-level description of the problem and its realization in some computable form (Gruber,

1989). This problem is addressed through the design of mediating representations that provide a bridge between verbal data and typical knowledge representation schemes such as production rules:

“[Mediating representations] are able to provide a formal representation that can be explicitly mapped to the internal architecture and knowledge of an expert system and which is easily readable by those who were not involved in the original development programme...” (Diaper, 1989, p. 34)¹

Knowledge acquisition tools employing mediating representations take one of two approaches: either they contain interfaces that bear a strong resemblance in appearance and procedure to the original task (e.g., cancer-therapy protocol forms in OPAL (Musen, 1988); engineering notebooks in vmacs (Sivard, Zweben, Cannon, Lakin & Leifer, 1989)) or they rely on some familiar, high-level, generic knowledge representation metaphor (e.g., repertory grids in ETS, (Boose, 1985, 1986); concept maps in ICONKAT (Ford *et al.*, 1990b)). Structured high-level representations are used to maintain the knowledge base, but if necessary they can be transformed into other forms such as rules (Gaines & Shaw, 1986; Ford *et al.*, 1990b). Transformation of knowledge into multiple forms or perspectives for visualization purposes is useful as a means of facilitating insight.

Interviewing techniques. A number of interviewing techniques have been developed or borrowed from fields such as psychology in order to guide experts through the

¹ Some differences in the literature exist over the terms “mediating” and “intermediate” representations. Johnson (1989, p. 184) defines an intermediate representation as one “which only exists between flanking representations and is bound to them by clearly defined projection rules which map one representation to the next” while the term mediating representation is preferred to “convey the sense of synthesis and coming to understand through the representation.” When we employ the term “mediating representation” to refer to devices used within our knowledge acquisition tools we mean something of both senses. Of course, a most important function of the representation in the early stages of knowledge engineering is better communication and insight among the participants. However, the iterative nature of knowledge refinement requires frequent, coupled cycles between elicitation, analysis, and performance elements of the system, which presupposes some sorts of mappings between various human and machine-oriented representations. These mappings, should not be regarded as having solved the philosophically troublesome issues surrounding the nature of knowledge and expertise, but simply as an application of well-known principles of human interface design to the problem of making knowledge structures more intelligible to domain experts.

knowledge acquisition process (Hoffman, 1987; Meyer, Booker & Bradshaw, 1990). For example, some of the techniques originally developed by Kelly (1955) to discover the conceptual structure of clients in psychotherapy have been applied to knowledge acquisition. Interviewing strategies and representational devices developed by Novak in an educational setting (1977; Novak & Gowin, 1984) have also proven useful in this regard.

Analysis tools. Both formal and heuristic modes of analysis are available within many knowledge acquisition tools as a means of verifying and refining the knowledge base. Shaw and Gaines (1988), for instance, have developed a methodology for analyzing areas of consensus, conflict, correspondence, and contrast between the conceptual systems of two or more experts. Many knowledge acquisition tools have also been integrated with consultation systems, making extensive performance testing of the system possible (e.g., Shema & Boose, 1988; Ford *et al.*, 1990c; Gaines & Rappaport, 1989).

2.2. DDUCKS: An Integrated Environment for Knowledge Acquisition and Decision Analysis

Early knowledge-based tools generally made strong assumptions about their operating environment. At first, dedicated, stand-alone applications were the rule. Over time, as the value of adding “hooks” for access to external applications and data was realized, most knowledge-based tools still operated under the assumption that they were in ultimate control of the system as the highest level executive. Currently, the greatest potential for use of knowledge-based systems is in areas requiring close interaction with traditional software applications and data. An application that assumes it is in ultimate control will be ineffective in such environments.

Brodie (1989) has discussed the need for “intelligent interoperability” in information systems. He defines the term to mean intelligent cooperation among systems to optimally achieve specified goals. While there is little disagreement that future computing environments will consist of multiple heterogeneous software systems running on multiple heterogeneous machines, most current computer systems are disjoint: they do not communicate (Figure 5). Until fairly recently, computer systems that could communicate nearly always used *ad hoc* interfaces for their particular connection. The recent growth in popularity of object-oriented approaches and the development of a few important agreed-upon protocols (e.g., SQL, the standard database interface language for relational database management systems) has nurtured the hope that encapsulated connectivity might someday become a reality.

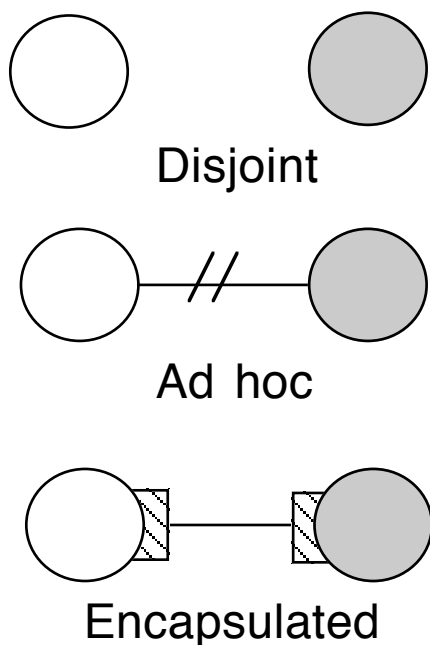


Figure 5. Evolution of system connectivity (Adapted from Brodie, 1989)

A high level of interoperability requires knowledge of the capabilities of each system, so that task planning, resource allocation, execution, monitoring, and, possibly, intervention between the systems can take place. Ideally, a mechanism functioning as a global planner or resource agent would manage cooperation activities (Figure 6).




Figure 6. Cooperating systems with global planner (Adapted from Brodie, 1989).

While a global planner would be workable for small networks of systems, such a scheme would quickly become impractical as the size of the network grew. The activity of the global planner would become a bottleneck for the (otherwise distributed) system. A final step toward intelligent interoperability would be to embed a resource agent within each cooperating system (Figure 7). Agents (i.e., the cooperating systems in the network) would ask their resource agent for the needed resources, thus providing a level of encapsulation at the planning level analogous to the encapsulation provided at the level of protocol.




Figure 7. Cooperating systems with distributed planner (Adapted from Brodie, 1989)

To make progress toward the goal of intelligent interoperability between *Axotl* and other potential system components, we defined an “open architecture” integrating environment that allows for a high degree of connectivity along with access to a global planner to facilitate intelligent cooperation. We call this environment DDUCKS (Decision and Design Utilities for Comprehensive Knowledge Support)¹. DDUCKS will both integrate knowledge acquisition tools and also facilitate coordination of local or networked applications such as spreadsheets, databases, or hypermedia software (Bradshaw, Covington, Russo & Boose, 1990; Figure 8). One of DDUCKS’ components, MANIAC (MANager for InterApplication Communication) supports asynchronous and synchronous communication between any number of multitasking applications. DDUCKS runs on Apple Macintosh II hardware; subsets of DDUCKS run in other environments.

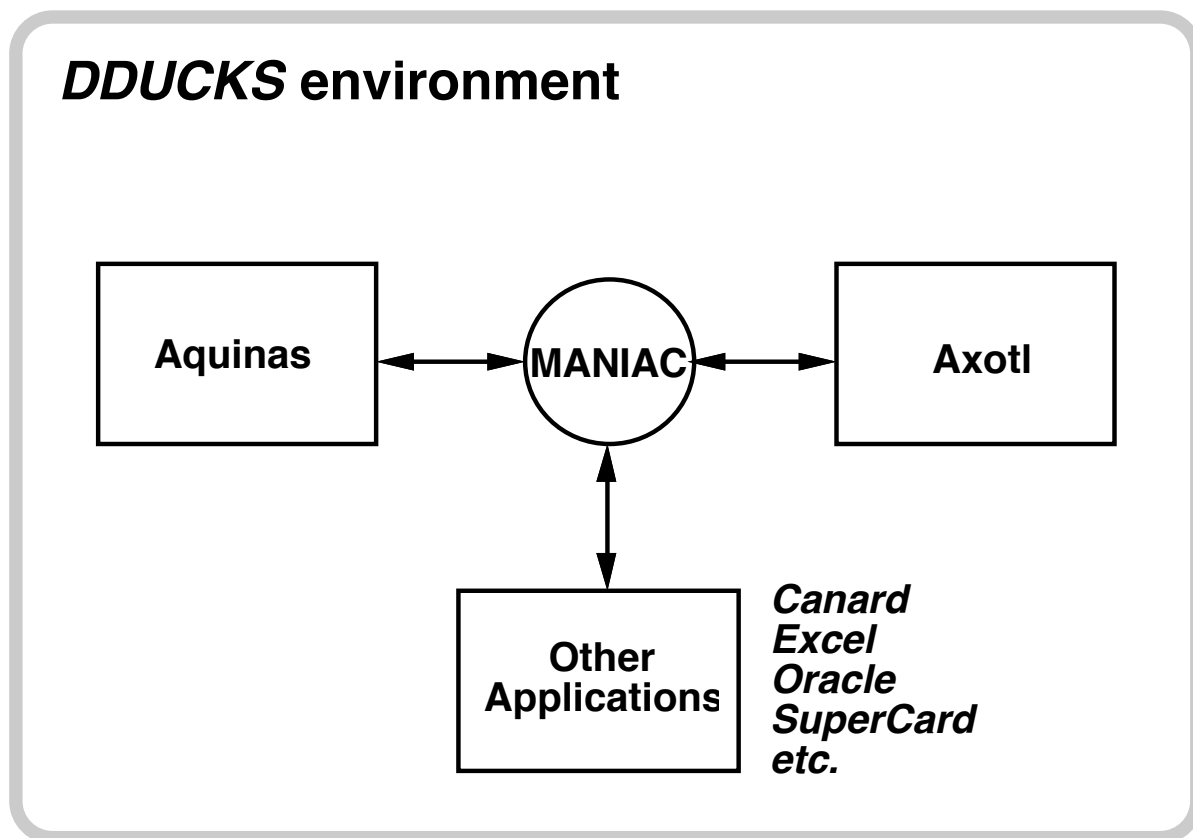


Figure 8. DDUCKS integrates components from *Axotl*, *Aquinas*, and other applications.

¹ Either the first or second *D* in *DDUCKS* is silent, depending on whether one using the tool in a decision or design context.

While integration of tools and approaches is a desirable and necessary objective, Ford *et al.* (1990b) wisely caution against the dangers involved in indiscriminate combining of different tools and techniques:

“Integration is the battle cry as tool-makers rush to produce hybrid knowledge acquisition tools. However, it is becoming widely realized that ad hoc combinations of techniques and tools—a sort of Occam’s hell—may not contribute much to ameliorating the knowledge acquisition bottleneck.”

To counteract the problems inherent in conceptual ad hoc’ism, they argue that “tool-makers should exploit theory as a means of building their tools on a sound footing and as a framework in which to make explicit their epistemological assumptions. (Ford & Adams-Webber, 1990a)”. The relationship between the theoretical foundations of *Aquinas*, *Axotl*, and *Canard* in personal construct theory, decision theory, and design methods theory are discussed in Bradshaw and Boose (1990), Bradshaw, Boose, Covington, and Russo (1989), and Shema, Bradshaw, Covington, and Boose (1990).

In the remainder of this paper, we will describe the decision analysis workbench and knowledge base tools of *Axotl*. Then we will outline the knowledge acquisition tools (e.g., *Aquinas* and *Canard*) being developed to support the creation of content and process knowledge bases for applications of *Axotl*.

2.3. *Axotl*: A Knowledge-Based Decision Analysis Workbench

Axotl combines a general-purpose decision analysis workbench with knowledge base tools to assist individuals consulting with the system about a specific problem. The first version of *Axotl* was developed as part of a joint effort by Boeing Computer Services and Strategic Decisions Group (SDG) to build a flexible, general-purpose intelligent decision system that

could be applied to internal R&D management decisions. The application is called PIE, for Project Investment Evaluator. Boeing Computer Services, with its experience in building knowledge-based systems, contributed expertise in software engineering, knowledge engineering, and Boeing-specific knowledge on this project. SDG has a major practice in applying decision analysis to R&D management decisions. In addition, SDG staff have pioneered the theory and practice of decision analysis (Howard & Matheson, 1984) and had previously implemented intelligent decision system application prototypes (Holtzman, 1989).

Following the successful demonstration of knowledge-based guidance of a decision analysis consultation in mid-1988, it was decided that Boeing and SDG would pursue further development work on the system independently. Since that time, we have made a number of extensions to *Axotl*, and have developed *DDUCKS* as the integrating framework for linking *Axotl* to *Aquinas* and other applications. We have loaned *DDUCKS* to a medical non-profit organization for evaluation of its applicability to bone marrow transplant follow-up care (Sullivan & Shulman, 1989). We are also applying portions of the environment to facilitate demonstrations of design knowledge capture for a NASA-sponsored Corporate Memory Facility project (Bradshaw, Boose, Covington & Russo, 1989; Boose, Bradshaw, Shema, & Covington, 1989; Boose, Bradshaw & Shema, 1990; Shema, Bradshaw, Boose & Covington, 1990), and in business process management in a Boeing quality improvement context (Bradshaw, Kipersztok, Nguyen, & Holm, 1990). SDG has also enhanced their version of the software, and is developing a commercial application to R&D management, called *R&D Analyst*^{TM1}. Holtzman and Seiver (1988) have developed a successful application of the system that assists critical care clinicians in ventilator management decision-making.

Axotl was developed on the Apple Macintosh II (*MacXotl*) and runs on all platforms that support ParcPlace Smalltalk-80 (e.g., Sun, Apollo, Hewlett-Packard, IBM '386 compatibles,

¹ *R&D Analyst* is a trademark of Strategic Decisions Group. All rights reserved.

IBM RS/6000). We will first describe the decision analysis workbench, then the knowledge base tools.

2.3.1. Decision Analysis Workbench

Many knowledge-based systems in use at Boeing are prescriptive in nature. They aim not only to describe situations but also to recommend specific actions. Recommendations made by such systems depend on the three kinds of knowledge shown in Figure 9: the *alternatives* available, *information* about consequences associated with the alternatives, and *preferences* among these consequences.

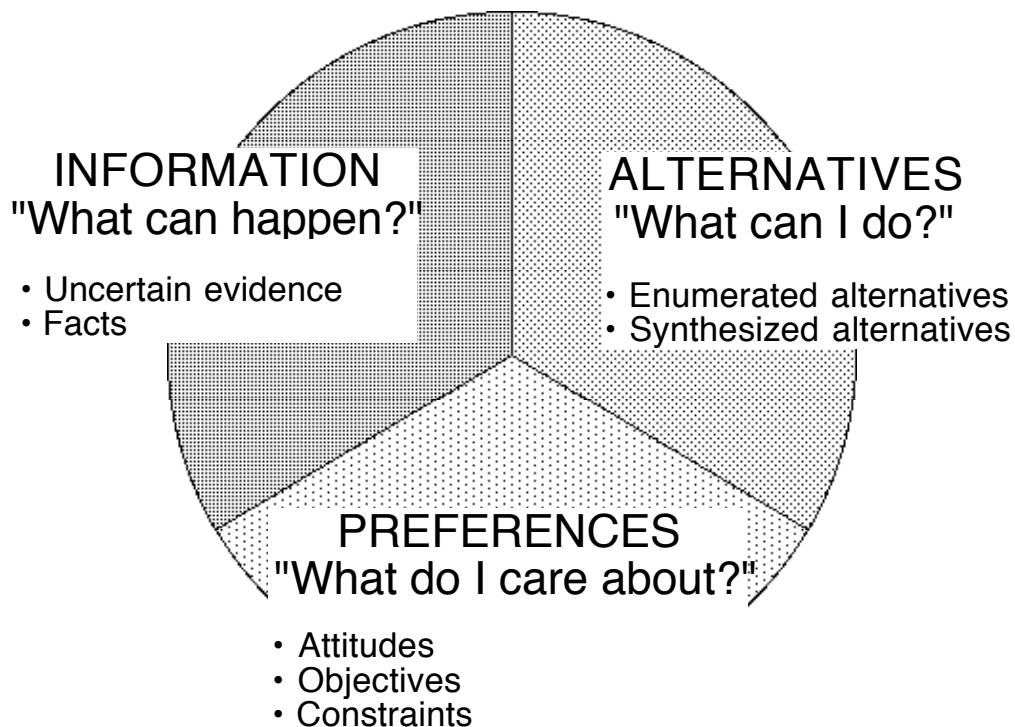


Figure 9. Decision analysis depends on three kinds of knowledge.

The decision analysis workbench contains a graphical editor that is used for creating and refining models of alternatives, preferences, and uncertainties relevant to a specific decision. These models, called *influence diagrams*, are solved to obtain recommended actions in a way that is consistent with probability and utility theory (Howard & Matheson, 1981).

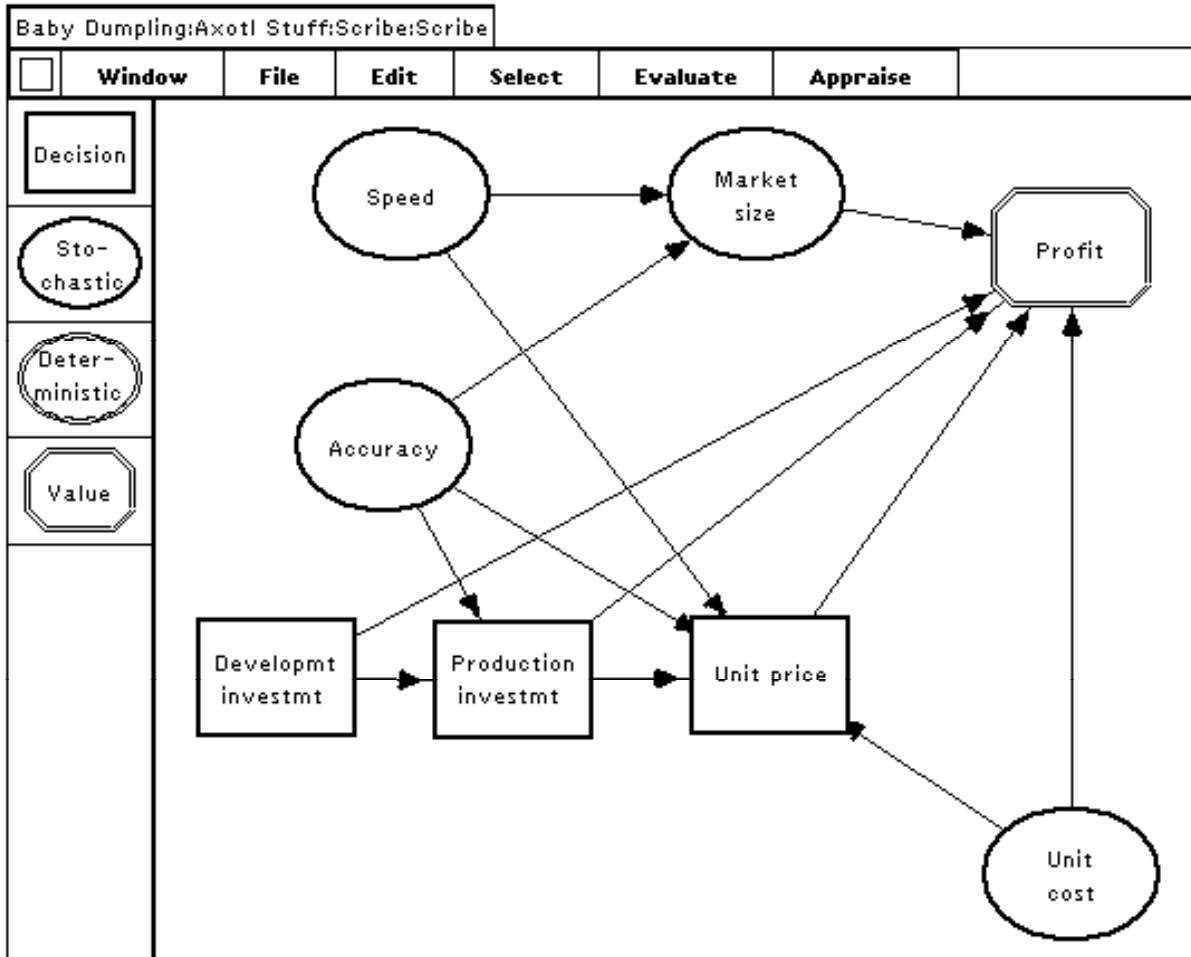


Figure 10. An influence diagram for an R&D investment decision.

Figure 10 shows a screen snapshot of an *Axotl* influence diagram (Howard & Matheson, 1981) representing a prototype R&D project decision problem. The problem is to determine an investment strategy for “Scribe” an automated speech-to-text transcriber, taking technical risks and market uncertainties into account. The investment strategy is composed of three decisions (development investment level, production investment level, and unit price), represented by rectangular nodes on the diagram. Oval nodes represent technical risk variables (accuracy, speed), production uncertainties (unit cost), and market uncertainties (market size). The eight-sided node labeled “Profit” has been designated as the criterion to maximize in evaluating the decision model to determine an optimal policy. Arrows between nodes represent relationships where influence or information is imparted from one variable to

another. For example, arrows from “Speed” and “Accuracy” to “Market size” represent the assertion that the degree of technical achievement in these areas will affect the size of the market for “Scribe”. The arrows from “Speed” and “Accuracy” to “Unit price” indicate that these uncertainties will be known at the time a pricing decision is made. An additional type of node, not shown in this diagram, can represent a deterministic function. This facilitates transparent links from the influence diagram to external procedures or to programs such as spreadsheets and databases.

The method of solving influence diagrams implemented in *Axotl* incorporates a new approach developed by Smith (1988) that allows a wide range of questions to be answered directly from the diagram. It preserves the entire underlying joint distribution during solution and inference procedures, rather than just the value lottery and decision policy as is usually done (e.g., Shachter, 1986a). The distribution editor, used to structure conditional probability distributions, is another feature unique to *Axotl*. It introduces the concepts of *coalescence* (i.e., sharing of atomic distributions) and *clipping* (i.e., explicit pruning of impossible or unnecessary conditioning paths and their atomic distributions).

One of the most significant results of *Axotl* development was the formulation of generic procedures for the use of influence diagrams in valuing information and control (J. Matheson (1988) discusses some of these issues). An approach was developed to allow the automated conversion of any influence diagram to “Howard Canonical Form” so that value-of-perfect-information or -imperfect-information questions could be answered for any variable. Based on an understanding that value-of-control calculations can readily be interpreted only for causal influence diagrams, we developed heuristic techniques that specified when it was appropriate to ask value-of-perfect-control or -imperfect-control questions and to formulate generic procedures for answering them. Eventually, these techniques could be fully automated.

2.3.2. Knowledge Base Tools

When configured with the appropriate knowledge, the knowledge base tools in *Axotl* guide decision makers through the process of formulating, evaluating, and refining a decision model in a specific domain. The model is kept to tractable size by deliberately limiting the problem domain for a given application and using heuristic methods represented in knowledge bases to select and prune variables during influence diagram construction. These heuristics can be regarded as instances of the types of categorical methods described by Szolovitz and Pauker (1978):

”When the complex problems need to be addressed—which treatment should be selected, how much of the drug should be given, etc.—then causal or probabilistic models are necessary. The essential key to their correct use is that they must be applied in a limited problem domain where their assumptions can be accepted with confidence. Thus, it is the role of categorical methods to discover what the central problem is and to limit it as strongly as possible; only then are probabilistic techniques appropriate for its solution.”

We distinguish between two major types of knowledge in the knowledge base: content and process (Figure 11). Each kind of knowledge poses a different set of problems for decision makers. The problems are addressed through the use of the knowledge base tools, in conjunction with the knowledge acquisition tools.

CONTENT KNOWLEDGE

PROCESS KNOWLEDGE

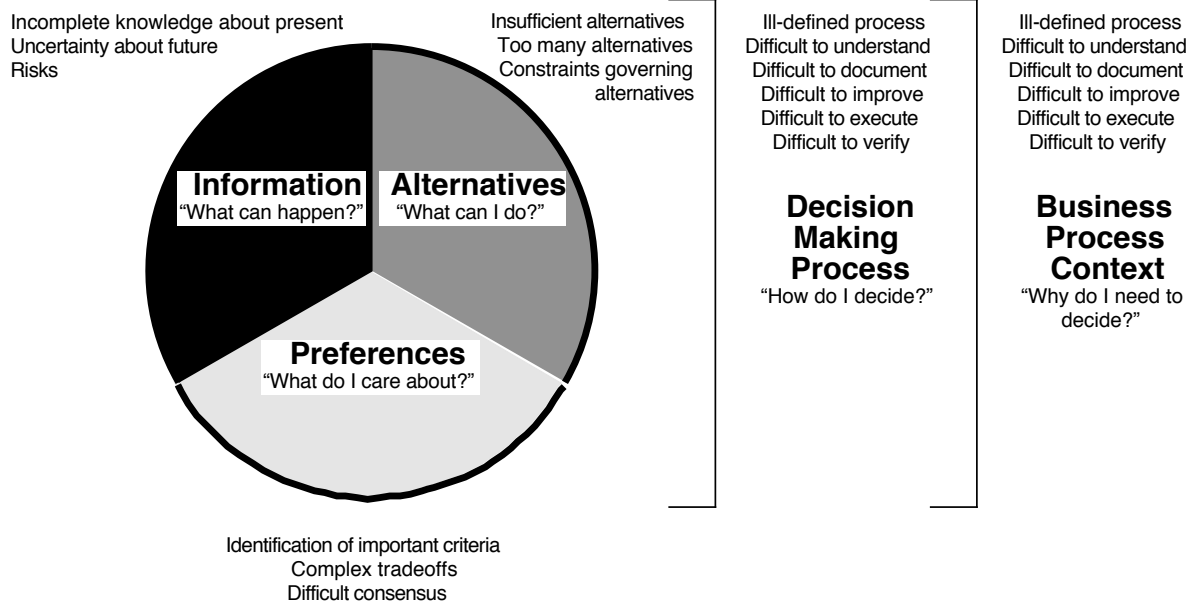


Figure 11. Different kinds of problem solving knowledge, and problems that people might face in dealing with each of them.

Content knowledge. Content knowledge is substantive; it is the *what* of decision model building during a consultation with a knowledge-based system. Internally, content knowledge is represented as a set of partially defined influence diagram variables for a class of decisions, their interrelationships, and the conditions under which they may be added to or removed from the influence diagram being constructed.

During the knowledge acquisition phase of the PIE project, we constructed a "big influence diagram" showing all of the structure that would be considered during the consultation process. Once agreement was reached about this structure, it was further broken down into smaller, overlapping groups of related variables along with some possible value models. We found that the "big influence diagram" for R&D project decisions had three distinct stages (research, development, application) and three possible value models (contract, market, simple value).

These value models and groups of influence diagram variables serve as the building blocks for the first-cut decision model that is constructed during subsequent consultations with a decision maker. Through a process called *template development*, the system ask the decision maker questions about the project that will help it determine the stage of development and the appropriate value model (D. Matheson, 1988). The system uses this project information to construct a generic, skeletal template. The dialogue continues to help decision makers expand and refine the initial model according to their specific circumstances.

Process knowledge. Process knowledge is strategic; it is the *how* of decision model building during a consultation. It consists of generic and domain-specific expertise about what to do at each stage of the decision-making process. Additionally, it may embody knowledge about the business process context. Internally, process knowledge is represented by activity graphs and rules. Figure 12 represents a view of the process of decision making for R&D projects that we used in developing the PIE application.

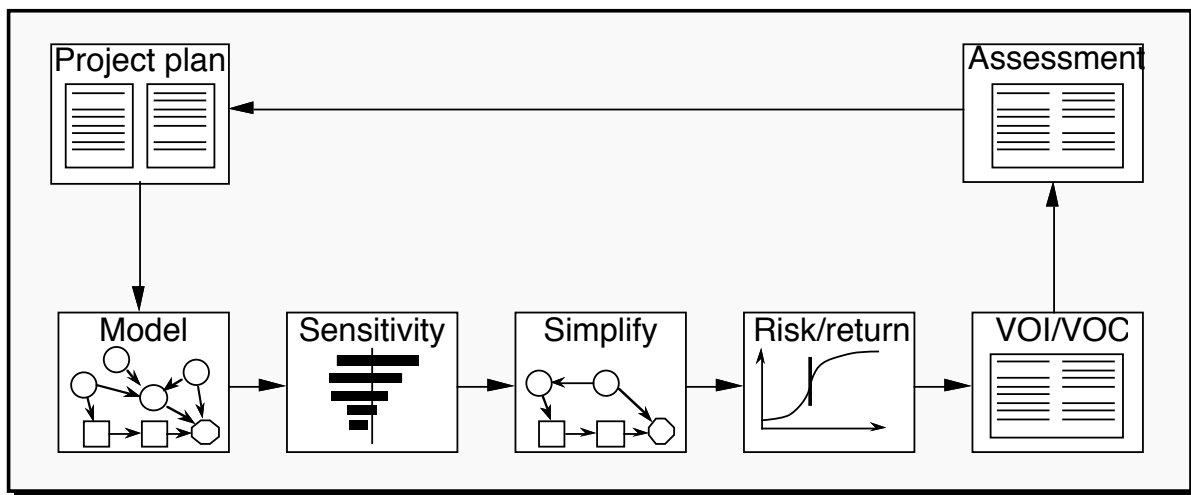


Figure 12. An overview of project investment decision making.

The cycle shown in Figure 12 represents general process knowledge about what to do at each stage of the R&D decision making process. We used the detailed process knowledge we acquired from decision analysis and Boeing experts to build a prototype R&D project selection decision activity graph (Figure 13) and rule base. Consultations using this

knowledge base proved the ability of PIE to pose a series of questions to a Boeing manager about a project under consideration and to formulate, evaluate, and appraise an appropriate influence diagram model based on answers to these questions. The decision analysis tools ran under control of the knowledge base, and the appropriate decision analysis task modules were invoked by the knowledge base at different times to evaluate the model and provide answers to value-of-information and value-of-control questions. The influence diagram was linked to a Microsoft Excel spreadsheet model containing financial formulas. Facilities in *Axotl* allowed Excel to be launched and the spreadsheet computations made automatically during influence diagram solution.

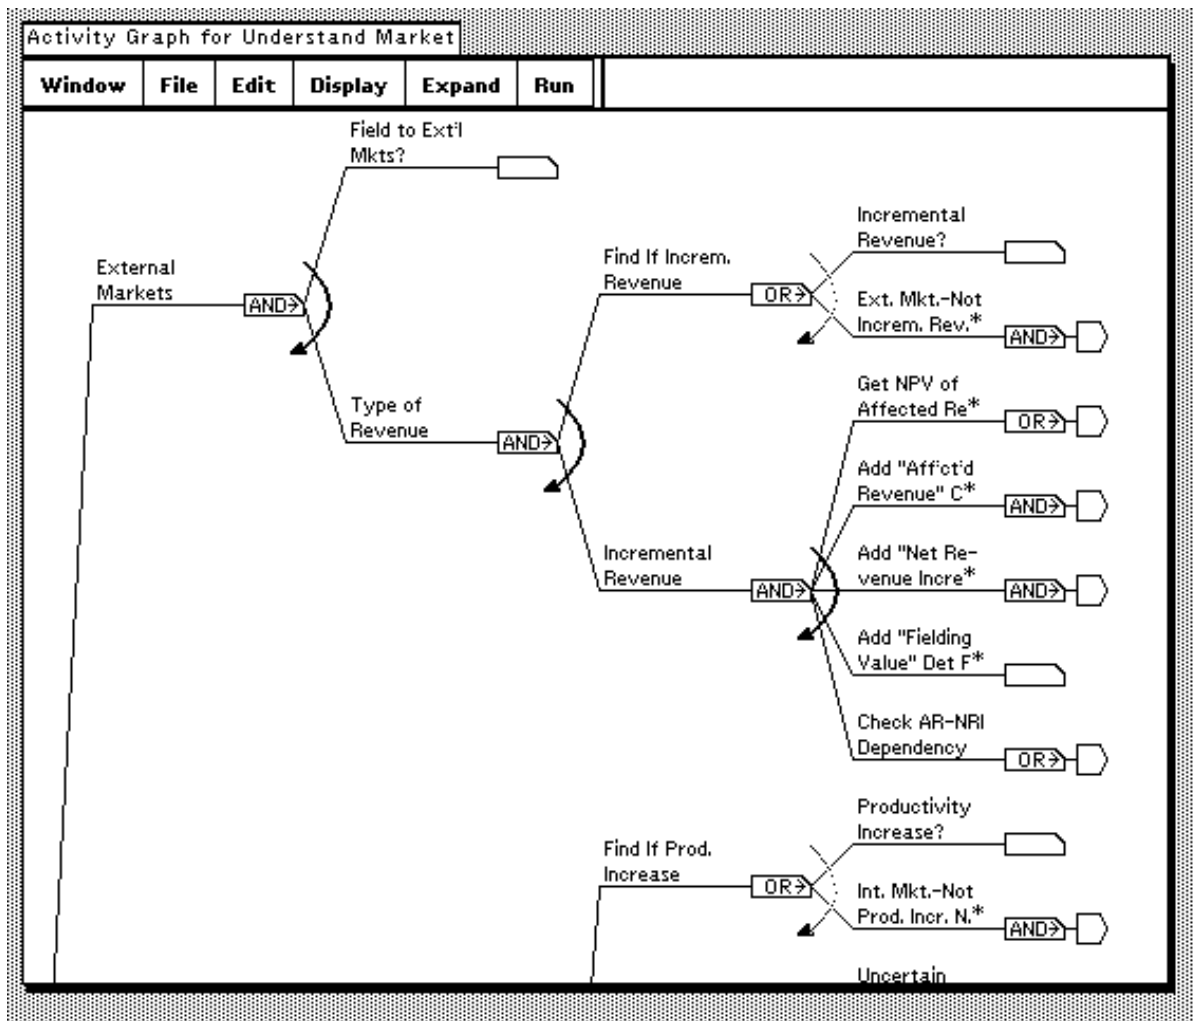


Figure 13. A portion of the R&D project selection activity graph.

An activity graph is a representation of the consultation process as a hierarchy of goals and activities (Holtzman & Russo, 1988). There are two kinds of nodes: goals and activities. The topmost goal in the hierarchy represents the successful completion of a consultation; any number of subgoals may be added. At the leaves of the graph, activities represent procedures that may be executed in support of goals during the course of a consultation (e.g., “assess probability of technical success,” “calculate value-of-information on market size”). Successful completion of a consultation requires that a sufficient set of goals be satisfied, either through the execution of the supporting activities or through being explicitly declared satisfied by the individual.

The activity graph facility includes components for viewing and editing goals and activities graphically. Activity graphs can call other activity graphs. During a consultation, an agenda is constructed from a “cut set” through the activity graph. The activities are executed one by one until the failure of an activity or some other change in conditions necessitates a modification of the agenda. The activity graph and agenda are invisible to the decision maker, who sees only the consultation interface.

A knowledge base tool called the heuristic advisor selects and modifies activity graphs during consultations. The knowledge base for the heuristic advisor is currently represented as rules and facts within an MRS-like inference engine (Russell, 1985), implemented in Smalltalk-80.

2.4. Knowledge Acquisition Tools for *Axotl*

An understanding of the roles that different kinds of knowledge play in problem solving has begun to drive requirements for knowledge acquisition tool development. Figure 14

is a view of how various knowledge acquisition tools under development at Boeing support different aspects of the problem¹.

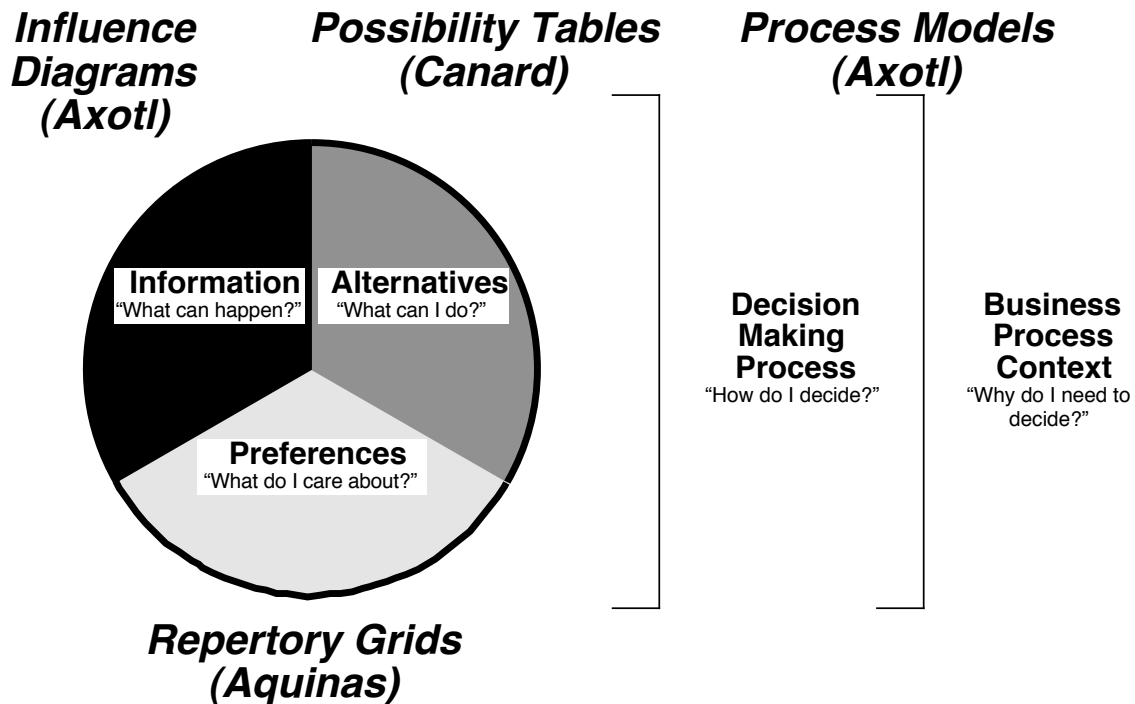


Figure 14. An idealized view of how different tools and representations support different knowledge requirements.

These tools are organized around three different roles for content knowledge and two for process knowledge. Each role has one or more forms of representation associated with it; some have interviewing techniques or analysis tools as well.

2.4.1. Content Knowledge

Elements of the decision basis can be separated into three major components: alternatives, information, and preferences (Figure 15). Possibility tables, grids, and graphs are the major mediating representations used to facilitate knowledge acquisition.

¹ This is, of course, an idealized view of how the tools work together. Although integration of functionality is an important goal of our work, some of the tools currently overlap more than others, and some are better integrated than others. The tools are discussed in more detail below.

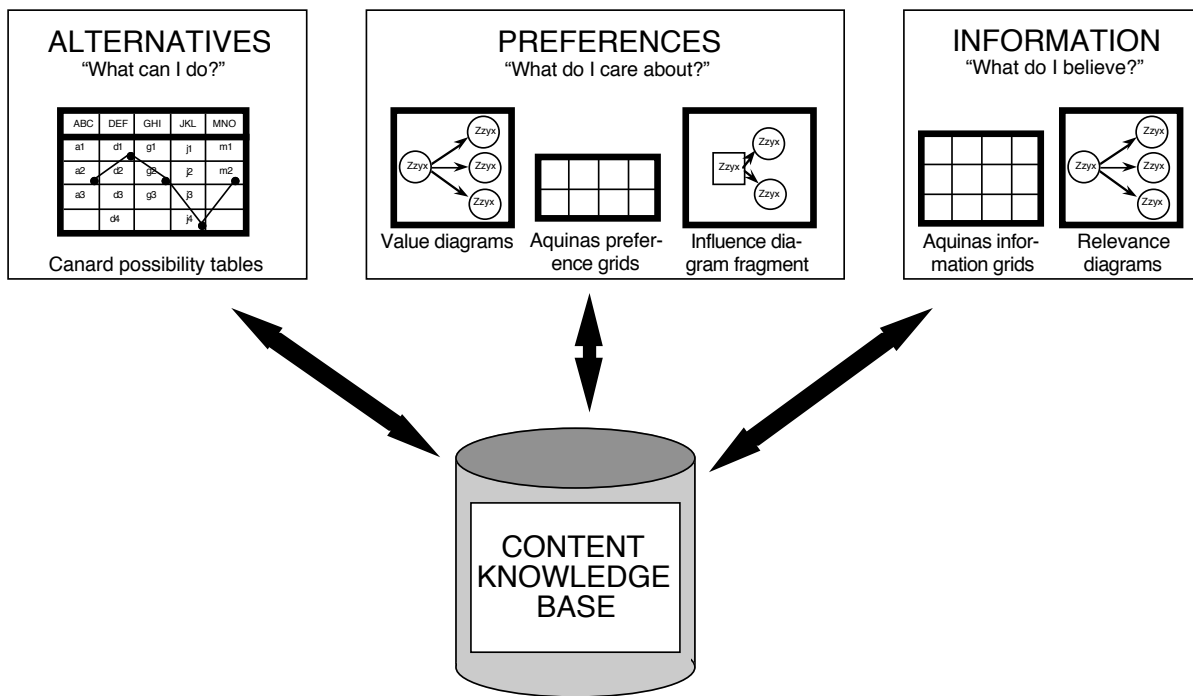


Figure 15. Three roles for knowledge in the content knowledge base with associated forms of representation.

Alternatives. *DDUCKS* includes *Canard*¹, a knowledge acquisition tool based on possibility tables that can be used to generate and structure complex alternatives (Shema, Bradshaw, Covington, and Boose, 1990). Links are maintained between the tables and decision nodes that are part of the content knowledge base. The possibility table representation is based on the manually developed strategy tables (McNamee & Celona, 1987) and morphological charts (Zwicky, 1969) that have been used by decision analysts and designers for many years. *Canard* automates this representation and extends its logic and structure to allow knowledge-based inference and the representation of more complex problems (e.g., hierarchical tables, explicit representation of constraints).

At one point, we used a manually developed possibility table similar to the one in Figure 16 to help us define alternatives for *Axotl* system development. Alternatives are shown in the leftmost column. Other columns represent important components of the system and various options for development within each one. The path of outlined boxes traced through the

¹ A canard is an airplane that sports a tail in front rather than in back—a symbol of our “backward” analytical approach to synthesis.

columns defines the “base case” alternative. Time and budget requirements and constraints could also be associated with variables in the table. Within *Canard* these constraints are used to “gray out” infeasible options during the definition of a path.

POSSIBILITIES	COMPONENTS									
	Spreadsheet	Home-Brew Prolog	Document'n of Code	Probability Encoder	DA Knowledge Base	Journal and Transcript	Dialog Boxes	UNDO Facility	Tutorial and User Manual	
Tight Resources	None 0	Use XQP only 0	No add'l document'n 0	Leave as is 0	None 0	None 0	Sequential menus only 0	None 0	None 0	0
Base Case	Modify existing spreadsheet 2	Straight port of IE w/o XQP emulation 2	Cursory update of PD document 3	Complete wheels and bar charts 3	Single, skeletal KB 5	Journal only 1	Ad hoc dialog boxes (no generality) 2	Support only a few commands 1	Terse, thorough primer 3	3
Flashy Demo	Enhance existing spreadsheet 9	IE with XQP syntax emulation 3	Thorough update of PD document 5	Add cumulative display 5	Single, saleable KB 12	Basic transcript only 7	Ad hoc but flexible dialog boxes 3	Support most commands 3	Verbose primer 5	5
All-out	Include some fancy features 15	IE with XQP & optimization 5		Add discretization 4	Two saleable KB's 15	Journal and fancy transcript 35	Generic dialog box facility 7	Multiple UNDO's up to a definable threshold 5	Full-fledged manual with tutorial 10	10
	Comm'l qual spreadsheet 40									

Figure 16. A possibility table defining *Axotl* development alternatives.

A major concern is helping people better explore the space of alternatives. Cognitive scientists have long known that humans typically retrieve only a small fraction of available alternatives when generating hypotheses. People tend to anchor on initial guesses, giving insufficient regard to subsequent data. For a variety of reasons, people may not be able to visualize whole classes of possibilities. Although it would be impossible in practice to guarantee that all relevant alternatives were indeed identified, *Canard* can help people consider a richer set of alternatives. An iterative search procedure that proposes new alternatives based on permutations of the constraint space assists in generation of alternatives. Through an analogous procedure, the system can hypothesize new constraints based on examples of previously defined alternatives.

Preferences. Knowledge acquisition tools can be used to help people determine important dimensions of value associated with the alternatives. Within *Aquinas*, a number of useful representation, interviewing, and analysis techniques have been developed. Many of these techniques are based on the research of George Kelly (1955), a clinical psychologist who

emphasized the foundational role of distinctions (*personal constructs*) underlying the processes of perception and reasoning. For example, using a *triadic elicitation* interviewing technique, *Aquinas* would ask people to define preferences by considering groups of alternatives presented three at a time: “Think of an important consequence that two of A1, A2, and A3 share, but that the other does not. What is that characteristic?” After giving B as that characteristic, the person would be asked about A2, A3, and A4 and come up with a second characteristic C, and so forth.

One way of representing this information is through a *repertory grid* (Figure 17). A grid is a matrix with elements (i.e., alternatives or outcomes) ranged along the bottom and constructs (i.e., dimensions of similarity and difference between elements), defined by extension, as a horizontal row of point values (or probability distributions on those values) within the matrix. Although the grid and network representations in Figure 17 are logically equivalent, we find both of them useful and complementary as problem clarifiers from a human factors point of view (Jones, 1981). The grid presentation allows the person to see patterns of similarity and difference that would otherwise be difficult to grasp. Analysis techniques in *Aquinas* (e.g., similarity analysis, cluster analysis) exploit these patterns to help users discriminate more carefully among similar concepts as part of model refinement. Implication analysis helps users discover important dependencies between constructs. A more complete discussion of the relationship of personal construct and decision analytic methods is given in Bradshaw and Boose (1990).




Figure 17 A preference grid and its corresponding influence diagram representation.

Information. Many of the same representation, interviewing, and analysis techniques mentioned above can be used to develop information models for the knowledge base.¹ Interviewing techniques (e.g., laddering) develop the structure of the networks by a recursive expansion procedure that terminates when all leaf nodes are directly observable. Heckerman has developed a related method called *similarity networks*, which uses techniques similar to those used in personal construct methodology to identify and display relationships indicating constraints on conditional independence relationships (Horvitz, Breese, & Henrion, 1988). Figures 18 and 19 show equivalent information grid and relevance diagram representations. When the structure of the dependencies between variables is simple, the transformation between grids and influence diagrams is straightforward. For complex dependencies, some of the methods described by Howard (1988) in his paper on *knowledge maps* may be useful. These methods allow complex probabilistic assessments to be broken down into a series of simpler ones by manipulating the conditioning relationships between chance nodes. We are sponsoring a project with Seattle University to develop procedures to make the use of disjoint and redundant knowledge maps practical in *Axotl*.

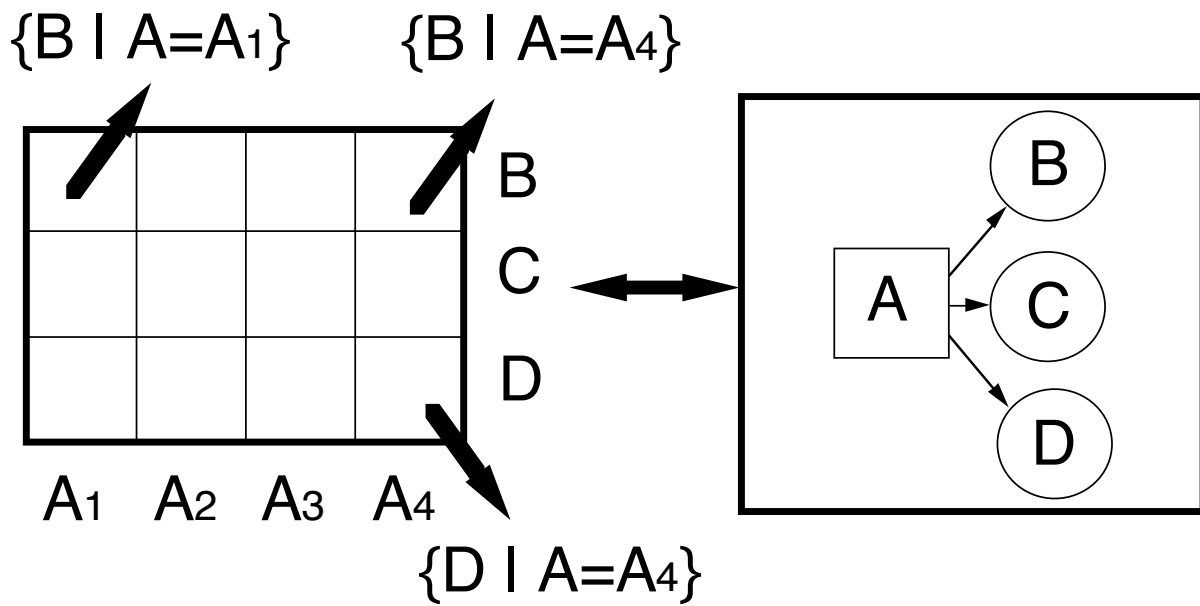


Figure 18. A simple information grid and its corresponding influence diagram representation.




Figure 19. A more complex information grid and its corresponding influence diagram representation.

Figure 20 illustrates some of the relationships possible between elements in the content knowledge base and the structure of a particular influence diagram. Note that the direction of conditioning between chance nodes during the knowledge acquisition process is often the reverse of the direction that is required when the influence diagram is constructed.

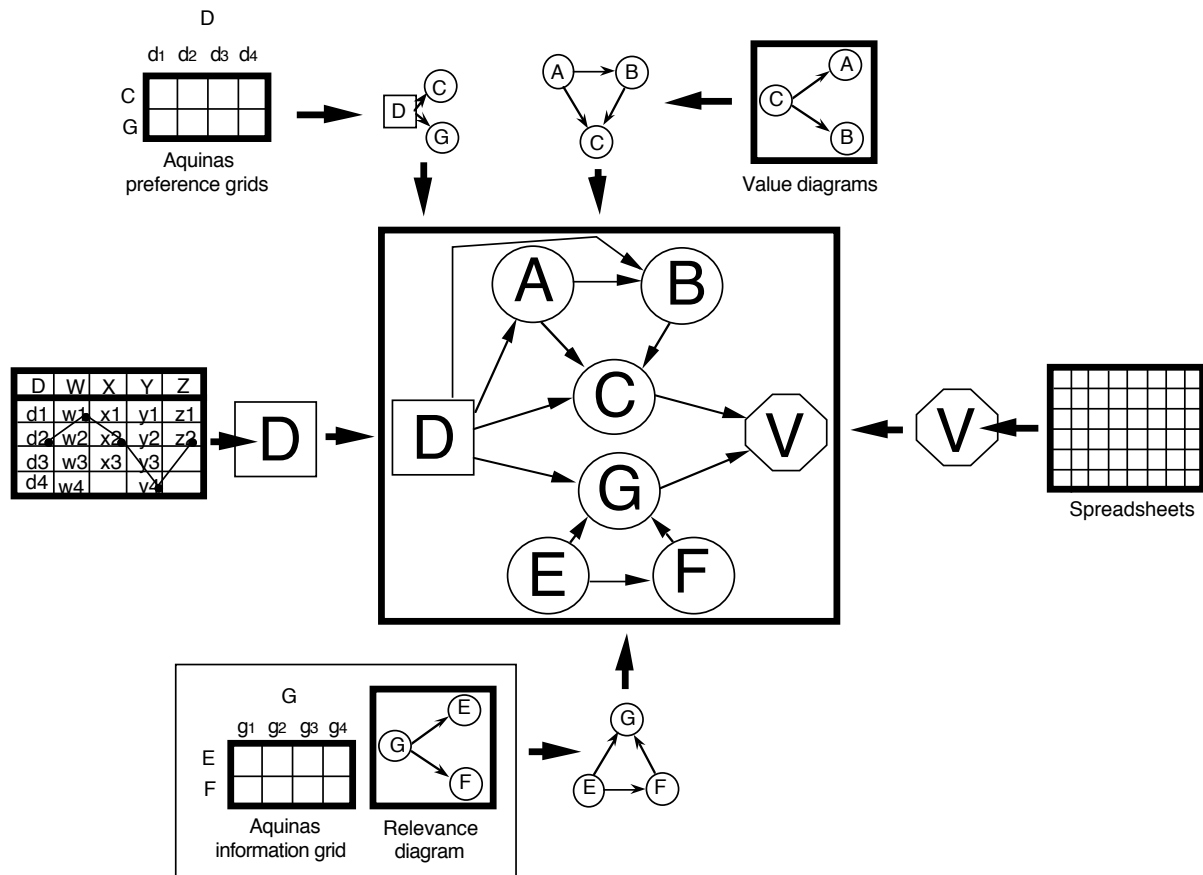


Figure 20. Relationships between elements in the content knowledge base and the structure of an influence diagram.

2.4.2. Process Knowledge

We distinguish between two major types of process knowledge: plans and plan selection knowledge (Figure 21).

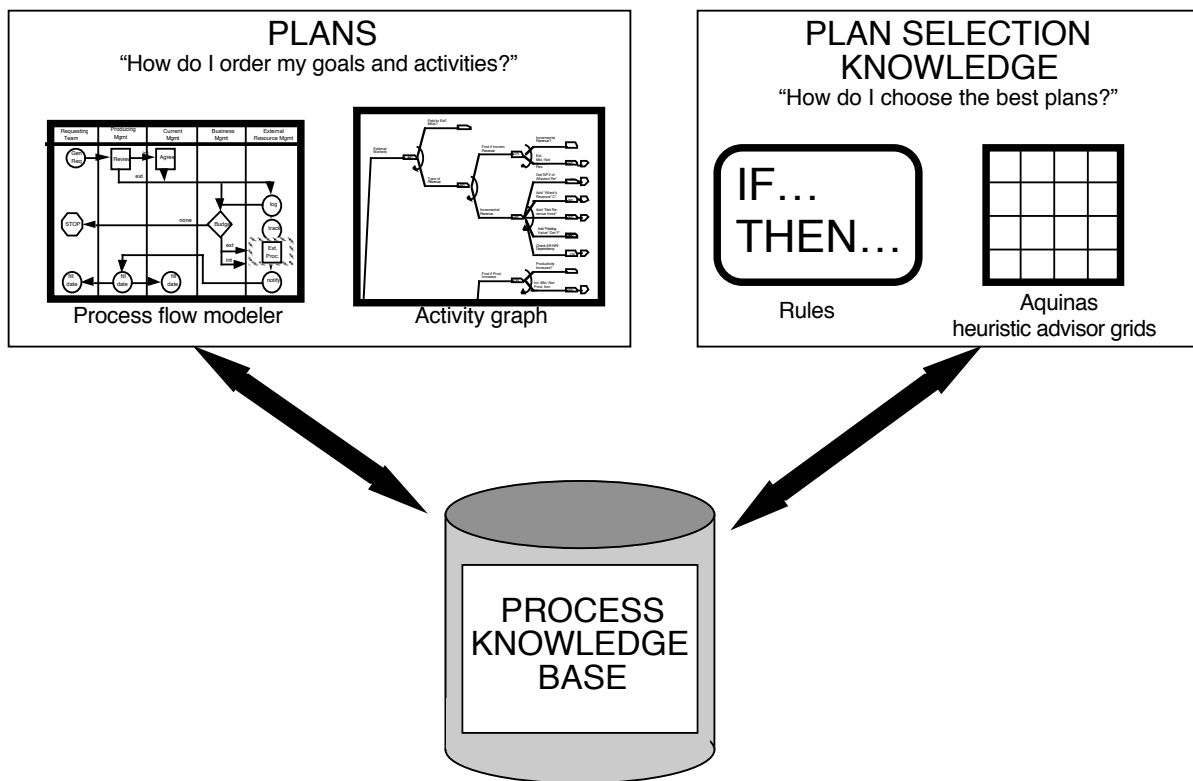


Figure 21. Two roles and sets of representations for knowledge in the process knowledge base.

Plans. Despite many advantages as a representational device, activity graphs have proved to be unworkable as a knowledge acquisition interface. For this reason, we are building an alternative means of activity graph entry based on hierarchical process flow charts (Bradshaw, Kipersztok, Nguyen, & Holm, 1990). Tasks described within process flows are automatically transformed into executable activity graphs (Figure 22). During a consultation, this activity graph is further compiled into an agenda, as described in section *** above¹. Analysis and simulation tools provide verification and validation of the process model.

¹ In addition to their use in controlling a consultation, these tools are being applied to documentation and improvement of general business processes.

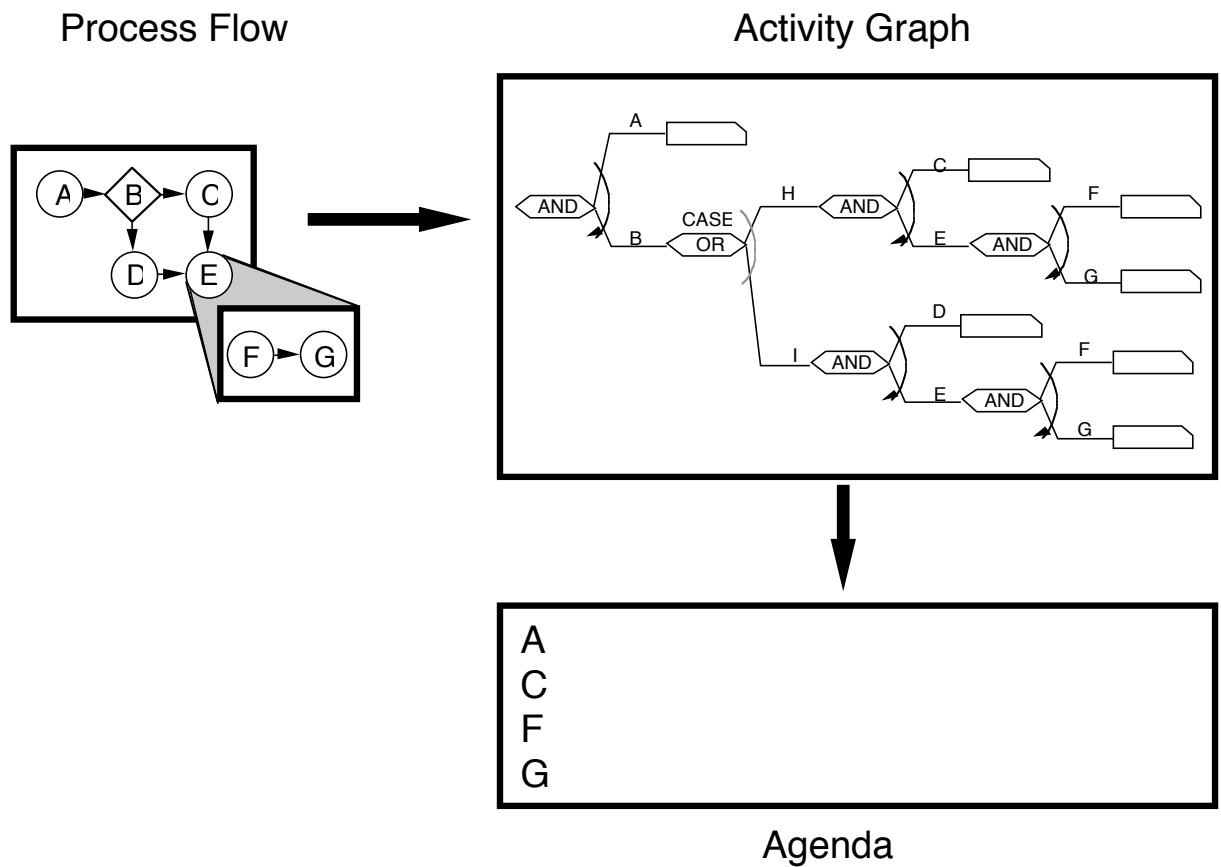


Figure 22 Process flows can be compiled into activity graphs¹.

Plan selection knowledge. Recently, we have begun to design knowledge acquisition tools for this type of knowledge. We have created graphical editors for world trees, rules, queries, and attributes that have simplified the construction of rules in knowledge bases. In the future, links to *Aquinas* will facilitate the construction of knowledge bases for the heuristic advisor. The heuristic advisor is a knowledge base tool that evaluates activities with respect to criteria such as completeness, balance, precision, and cost. Based on this evaluation, the heuristic advisor modifies the activity graph to re-order activities dynamically during a consultation. A related effort has dealt specifically with reasoning about computational resource tradeoffs (Horvitz, 1989). In contrast to our heuristic approach, Horvitz has developed these ideas within a formal framework.

¹ This figure is somewhat simplified for illustrative purposes. In reality, a separate activity graph would be compiled for the main process flow and the subprocess (F->G). This subgraph would be called by the main graph dynamically during execution.

3. DISCUSSION

The heart of our approach to making knowledge acquisition simpler for intelligent decision systems is “divide and conquer”. Having experienced the development of many knowledge-based system prototypes using automated knowledge acquisition tools, we are convinced that such a strategy will permit experts to do much of their own knowledge engineering, without requiring a great deal of specialized training (Boose, Bradshaw, Kitto, & Shema, 1990). To the experienced decision analyst or knowledge engineer, the development of these tools may seem completely unnecessary; there is little question that such a person would find it more efficient to work directly with influence diagrams and activity graphs than with the knowledge acquisition tools we are defining. In addition, one could argue that by making it too easy on the naive users of a system, putting everything in their terms without requiring them to come to grips directly with the underlying methodology, we are promoting a “black box” mentality that makes it impossible for them to step in when the system breaks down or encounters a problem it can’t solve.

There is no easy answer to these objections; there seems to be a real and inevitable tradeoff between the “acquirability” and the expressive power of knowledge representations (Gruber, 1989). Figure 23 shows this tradeoff as a dark curved line. The most powerful means of getting an idea into a computer is programming; unfortunately, even with advances in software engineering, programming remains the most difficult form for nonspecialists to express their knowledge in. On the other hand, form-filling interfaces that correspond to the way a user normally enters information on paper are very easy to learn but are very rigid and limit the applicability of the tools to the very specific problems that the system designer has foreseen.

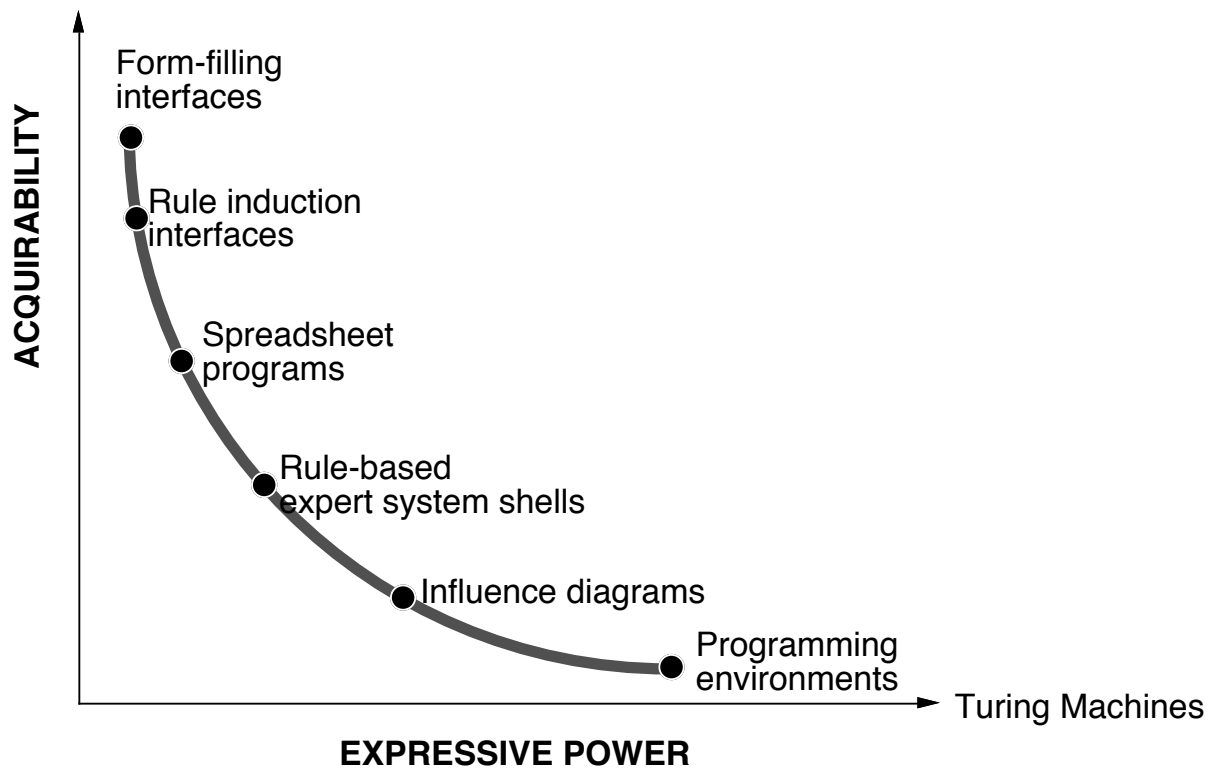


Figure 24. The tradeoff between acquirability and expressive power. (Figure adapted from Gruber, 1989)

Knowledge acquisition tools do not eliminate the competition between acquirability and expressive power, but they can act as a kind of magnet to help pull the curve out (Figure 24; light gray line). Through the application of automated techniques, acquirable interfaces can become more powerful, and powerful interfaces can be more easily learned and used. *Aquinas* and *Axotl* make the power of influence diagrams accessible to a wider range of people.

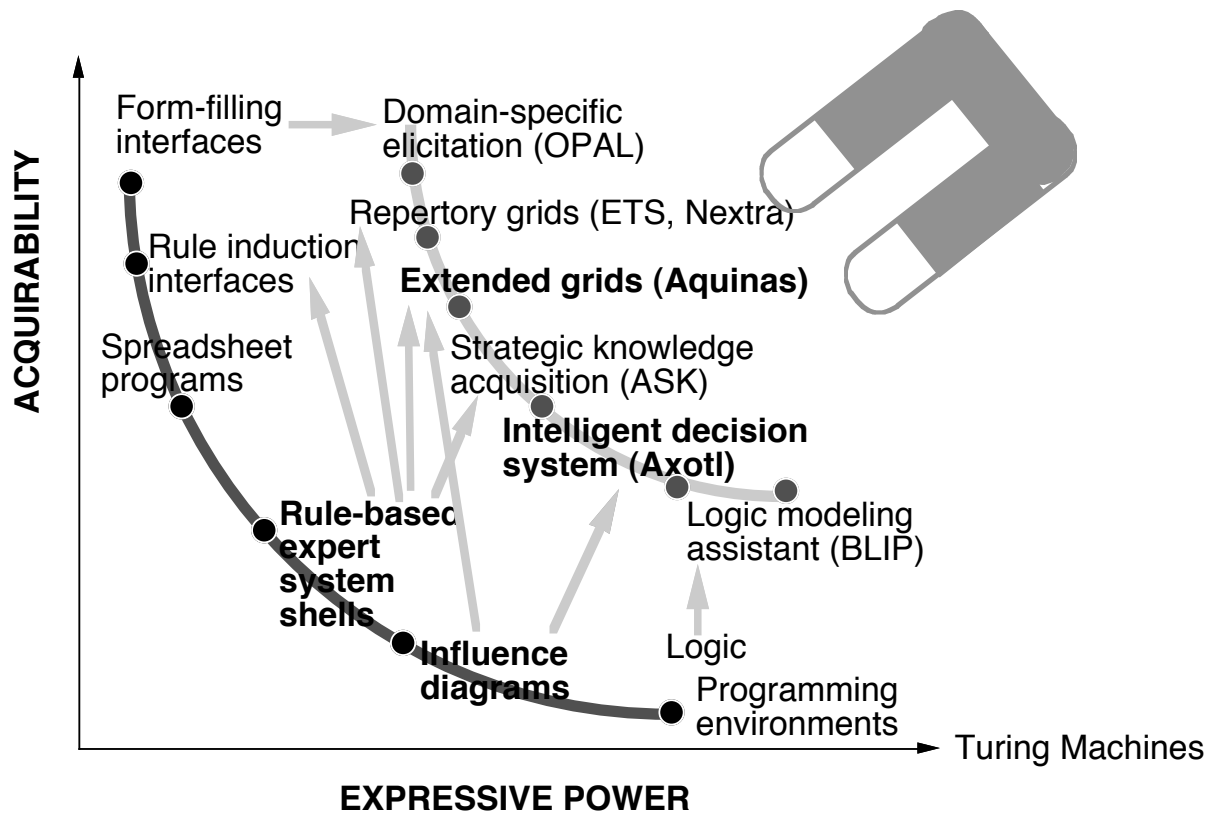


Figure 24. Automated knowledge acquisition tools can help make acquirable interfaces more powerful, and powerful interfaces more easily learned and used (Figure adapted from Gruber, 1989).

Because new users will always prefer acquirable interfaces and experienced users will prefer powerful ones, the best strategy to adopt when adding new layers of mediating representation between the new user and the underlying problem-solving model is to leave the hooks to the lower level representations intact. Designers of such systems should also find ways to encourage the less experienced users to keep learning and move on to more expert modes of interaction as they become more familiar and comfortable with the system. Experienced users, on the other hand, should be able to go directly to the heart of the system without interference from a “friendly” interface. This is the approach taken in the nested “Russian doll” interface of HyperCard, where users can graduate from browsing to painting to authoring to scripting at their own rate while performing useful work with the system at each level. This is also consistent with the “glass box” idea advocated by Wenger and Brown at IRL: to design software that promotes understanding of and access to the inner workings of the system (Feinstein, 1989).

Our conjecture is that high-level representations and interviewing and analysis tools will be helpful to experienced as well as naive users. This is a question that we hope to address as the tools are evaluated in real-world domains.

We are cautiously optimistic about the future of *DDUCKS*. Although current and future applications will no doubt continue to bring new challenges and difficult problems to solve, the tools embody many years of experience in decision analysis automation and knowledge-based system development. Furthermore, the system derives power from its integration of elements that we feel will be the building blocks of future systems for automated support of complex decisions: a decision analysis workbench for sound reasoning under uncertainty and resolution of preference issues; a knowledge-based system for provision of help in building formal decision models and to promote the re-use of domain knowledge; and a collection of knowledge acquisition tools tailored to the domain that allow such systems to be efficiently modeled, more easily maintained, and thoroughly tested.

ACKNOWLEDGEMENTS

We express our appreciation to Kit Bakke, Miroslav Benda, Kathleen Bradshaw, Beverly Clark, Karolyn Covington, Mike Freeman, Peter Friedland, Pete and Cindy Holm, Earl Hunt, Oscar and Sharon Kipersztok, Cathy Kitto, Joe Koszarek, Sandra Marcus, Allen Matsumoto, Art Nagai, Thom Nguyen, Guy Olney, Steve Poltrock, George Roberts, Doug Schuler, Kish Sharma, Dave Shema, Greg Swietek, Lisle Tinglof-Boose, and Bruce Wilson for their contributions and support. Thanks also to reviewers who provided many helpful suggestions. Samuel Holtzman, Ronald Howard, Keh Shiou Leu, James Matheson, Léonard Bertrand, and James Smith of Strategic Decisions Group were co-designers and co-developers of the original *Axotl* prototype. This work has also benefitted from ongoing discussions with Brian Gaines, Ken Ford, Max Henrion, and Mildred Shaw. *Aquinas*, *Axotl*, *Canard*, and *DDUCKS*

are under development at the Boeing Advanced Technology Center. The original development of *Canard* was supported by NASA contract NAS2-12108.

REFERENCES

AKKERMANS, H., VAN HARMELEN, F., SCHREIBER, G. & WIELINGA, B. (1990). A formalization of knowledge-level models for knowledge acquisition. In K. Ford & J. Bradshaw (Eds.), special knowledge acquisition issue of the *Journal of Intelligent Systems*, in press.

BOOSE, J.H. (1985). A knowledge acquisition program for expert systems based on Personal Construct Psychology. *International Journal of Man-Machine Studies*, 23, 495-525.

BOOSE, J.H. (1986). *Expertise Transfer for Expert System Design*. New York: Elsevier.

BOOSE, J.H. (1990). A survey of knowledge acquisition techniques and tools. *Knowledge Acquisition Journal*, in press.

BOOSE, J.H. & BRADSHAW, J.M. (1987). Expertise transfer and complex problems: Using *Aquinas* as a knowledge-acquisition workbench for knowledge-based systems. *International Journal of Man-Machine Studies*, 26, 3-28. Also in J. Boose and B. Gaines (Eds.), *Knowledge Acquisition Tools for Expert Systems*. London: Academic Press, pp. 39-64.

BOOSE, J. H., BRADSHAW, J.M., KITTO, C.M., & SHEMA, D.B. (1990). From ETS to *Aquinas*: Six years of knowledge acquisition tool development. *Knowledge Acquisition Journal*, in press.

BOOSE, J. H., BRADSHAW, J.M., SHEMA, D.B., COVINGTON, S.P. (1989). Design Knowledge Capture for a Corporate Memory Facility, *IJCAI-89 Workshop on Knowledge Acquisition: Practical Tools and Techniques*, Detroit, August, 1989; *IJCAI-89 Workshop on Knowledge Acquisition*, Detroit, August, pp. 5-6.

BOOSE, J.H., SHEMA, D.S., & BRADSHAW, J.M. (1989). Recent progress in *Aquinas*: A knowledge acquisition workbench. *Knowledge Acquisition Journal*, 1(2): 185-214.

BOOSE, J.H., SHEMA, D.S., & BRADSHAW, J.M. (1990). Capturing design knowledge for engineering trade studies. In B. Wielinga, J. Boose, B. Gaines, G. Schreiber, & M. Van Someren (Eds.), *Current Trends in Knowledge Acquisition*. Amsterdam: IOS Press.

BRADSHAW, J. M. & BOOSE, J.H. (1989). Knowledge acquisition as CASE for knowledge-based systems. *Presentation at the Third International Workshop on Computer-Aided Software Engineering (CASE-89)*, London, England, July 17-21, 1989.

BRADSHAW, J. M. & BOOSE, J.H. (1990). Decision analysis techniques for knowledge acquisition: Combining information and preferences using *Aquinas* and *Axotl*. *International Journal of Man-Machine Studies*, 32(2): 121-186. Also in J.H. Boose and B.R. Gaines (Eds.), *Progress in Knowledge Acquisition for Knowledge-Based Systems..* London: Academic Press.

BRADSHAW, J.M., BOOSE, J.H., COVINGTON, S.P., & RUSSO, P.J. (1989). How to do with grids what people say you can't: The application of decision analysis methods in *Axotl* and personal construct methods in *Aquinas* to design problems. *Third AAAI Knowledge Acquisition for Knowledge-based Systems Workshop*, Banff, Canada, November, 1988.

BRADSHAW, J.M., COVINGTON, S.P. & RUSSO, P.J. (1989). PIE: An application of *Axotl* to project selection decisions. *Proceedings of the Boeing AI Conference*, April.

BRADSHAW, J.M., COVINGTON, S., RUSSO, P., & BOOSE, J.H. (1990). Knowledge acquisition techniques for Intelligent Decision Systems: Integrating *Axotl* and *Aquinas* in

DDUCKS.. In M. Henrion, R.D. Shachter, L.N. Kanal, and J.F. Lemmer (Eds.), *Uncertainty in Artificial Intelligence 5*. North-Holland: Elsevier.

BRADSHAW, J. M. & HOLTZMAN, S. (1987). Artificial intelligence and decision analysis. Presentation made to the Nineteenth Annual Meeting of the Society for Information Management, Seattle: October 17.

BRADSHAW, J.M., KIPERSZTOK, O., NGUYEN, T. & HOLM, P. (1990). Acquiring knowledge for process management using *eQuality*. Presentation to NORTHCON, October, 1990.

BREESE, J.S. (1987). Knowledge representation and inference in intelligent decision systems. Doctoral dissertation, Department of Engineering-Economic Systems, Stanford University. Reprinted as Research Report 2, Palo Alto, California: Rockwell International Science Center.

BRODIE, M. L. (1989). Future intelligent information systems: AI and database technologies working together. In J. Mylopoulos and M. L. Brodie (eds.) *Readings in Artificial Intelligence and Databases*. San Mateo, California: Morgan Kaufmann.

DIAPER, D. (1989). Designing expert systems—from Dan to Beersheba. In D. Diaper (Ed.) *Knowledge Elicitation: Principles, Techniques and Applications*. New York: John Wiley.

ESHELMAN, L. (1988). MOLE: A knowledge acquisition tool for cover-and-differentiate systems. In S. Marcus (ed.), *Automating Knowledge Acquisition for Expert Systems*. Boston, Massachusetts: Kluwer Academic Publishers.

FEINSTEIN, D. (1989). Glass boxes: Lifting the veil on information. *Benchmark*, Spring, 12-15.

FORD, K.M. & ADAMS-WEBBER, J.R. (1990a). Knowledge acquisition and constructivist epistemology. In this volume.

FORD, K.M., PETRY, F.E., ADAMS-WEBBER, J.R., & CHANG, P.J. (1990b). An approach to knowledge acquisition based on the structure of personal construct systems. *IEEE Transactions on Knowledge and Data Engineering*, in press.

FORD, K.M., STAHL, H., ADAMS-WEBBER, J.R., NOVAK, J. & JONES, J.C. (1990c). ICONKAT: An integrated constructivist knowledge acquisition tool. Manuscript in preparation.

GAINES, B.R. (1988). Software engineering for knowledge-based systems. *Proceedings of the Second International Workshop on Computer-Aided Software Engineering (CASE-88)*.

GAINES, B.R. & RAPPAPORT, A.T. (1989). The automatic generation of classes, objects and rules at the interface between knowledge acquisition tools and expert system shells. *IJCAI-89 Workshop on Knowledge Acquisition: Practical Tools and Techniques*, Detroit, Michigan, August 1989.

GAINES, B. R. & SHAW, M. L. G. 1986b). Induction of inference rules for expert systems. *Fuzzy Sets and Systems*, 8(3), 315-328 (April).

GRUBER, T.R. (1989). *The Acquisition of Strategic Knowledge*. New York: Academic Press.

HENRION, M. (1990). An introduction to algorithms for inference in belief nets. In M. Henrion, R.D. Shachter, L.N. Kanal, and J.F. Lemmer (Eds.), *Uncertainty in Artificial Intelligence 5*. North-Holland: Elsevier.

HOFFMAN, R. (1987). The problem of extracting the knowledge of experts from the perspective of experimental psychology. *AI Magazine* (Summer).

HOLTZMAN, S.H. & RUSSO, P.J. (1988). Activity graphs. Unpublished PIE project internal technical memo.

HOLTZMAN, S. H. (1989). *Intelligent Decision Systems*. Reading, Massachusetts: Addison-Wesley.

HOLTZMAN, S. & SEIVER, A. (1988). Decision analysis: A framework for critical care decision assistance. *International Journal of Clinical Monitoring and Computing*, in press.

HORVITZ, E.J. (1989). Reasoning about beliefs and actions under computational resource constraints. In L. Kanal, T. S. Levitt, & J. F. Lemmer (eds.), *Uncertainty in Artificial Intelligence*, volume 3. Amsterdam: North-Holland.

HORVITZ, E. J., BREESE, J. S., & HENRION, M. (1988). Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, 2, 247-302.

HOWARD, R. A. (1966). Decision analysis: Applied decision theory. In D. B. Hertz and J. Melese (eds.), *Proceedings of the Fourth International Conference on Operational Research*. Reprinted in Howard & Matheson, 1984.

HOWARD, R. A. (1988). Knowledge maps. *Management Science*, in press.

HOWARD, R. A. & MATHESON, J. E. (1981). Influence Diagrams. Reprinted in Howard & Matheson, 1984.

HOWARD, R. A. & MATHESON, J. E. (Eds.) (1984). *Readings on the Principles and Applications of Decision Analysis*. Menlo Park, California: Strategic Decisions Group.

JOHNSON, N. E. (1989). Mediating representations in knowledge elicitation. In D. Diaper (Ed.) *Knowledge Elicitation: Principles, Techniques and Applications*. New York: John Wiley.

JONES, J.C. (1981). *Design Methods*. New York: John Wiley & Sons.

KARBACH, W., LINSTER, M. & VOSS, A. (1990). A confrontation of models of problem solving. *Knowledge Acquisition Journal*, in press.

KEENEY, R. L. (1986). *Value-driven Expert Systems For Decision Support*. Los Angeles: Systems Science Department, University of Southern California.

KEENEY, R. L. & RAIFFA, H. (1976). *Decisions With Multiple Objectives: Preferences and Value Tradeoffs*. New York: Wiley.

KELLY, G. A. (1955). *The Psychology of Personal Constructs*. 2 volumes. New York: Norton.

KITTO, C. M. (1988). Progress in automated knowledge acquisition tools: How close are we to replacing the knowledge engineer? *Proceedings of the Third AAAI-sponsored Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, Canada, November 1988.

KLINKER, G. (1989). A framework for knowledge acquisition. *Proceedings of the Third Annual European Knowledge Acquisition Workshop*, Paris, France, July.

LANGLOTZ, C. P., SHORTLIFFE, E. H. & FAGAN, L. M. (1986). Using decision theory to justify heuristics. *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia.

KLINKER, G., BOYD, C., DONG, D., MAIMAN, J., MCDERMOTT, J. & SCHNELBACH, R. (1989). Building expert systems with KNACK. *Knowledge Acquisition Journal*, 1(3), 299-320.

KLINKER, G., BOYD, C., GENETET, S. & MCDERMOTT, J. (1988). A KNACK for knowledge acquisition. *Proceedings of the Sixth National Conference on Artificial Intelligence*, Seattle, WA, July.

MARCUS, S. (1987). Taking backtracking with a grain of SALT. *International Journal of Man-Machine Studies*, 26, 383-398.

MARCUS, S. (1988). SALT: A knowledge acquisition language for propose-and-revise systems. In S. Marcus (ed.), *Automating Knowledge Acquisition for Expert Systems*. Boston, Massachusetts: Kluwer Academic Publishers.

MARCUS, S. (1989). Understanding decision ordering from a piecemeal collection of knowledge. *Knowledge Acquisition Journal*, 1(3), 279-298.

MATHESON, D. (1988). Knowledge acquisition results on template development. Unpublished PIE project internal technical document.

MATHESON, J.E. (1988). Using influence diagrams to value information and control. *Proceedings of the Conference on Influence Diagrams for Decision Analysis, Inference, and Prediction*, Berkeley, California, University of California at Berkeley, May 9-11.

MCNAMEE, P. & CELONA, J. (1987). *Decision Analysis for the Professional—With Supertree*. Redwood City, California: Scientific Press.

MEYER, M.A., BOOKER, J.M. & BRADSHAW, J.M. (1990). A flexible six-step program for defining and handling bias in knowledge acquisition. In B. Wielinga, J. Boose, B. Gaines, G. Schreiber, & M. Van Someren (Eds.), *Current Trends in Knowledge Acquisition*. Amsterdam: IOS Press.

MOORE, E. A. & AGOGINO, A. M. (1987). INFORM: An architecture for expert-directed knowledge acquisition. *International Journal of Man-Machine Studies*, 26, 213-230.

MUSEN, M.A. (1988). *Generation of model-based knowledge acquisition tools for clinical-trial advice systems*. Doctoral dissertation, Computer Science Department, Stanford University. Available as KSL-88-06 and STAN-CS-88-1194.

NEAPOLITAN, R.E. (1990). *Probabilistic Reasoning in Expert Systems: Theory and Algorithms*. New York: John Wiley.

NOVAK, J.D. (1977). *A Theory of Education*. Ithaca, N.Y.: Cornell University Press.

NOVAK, J.D. & GOWIN, D.B. (1984). *Learning How to Learn*. New York: Cambridge University Press.

OLMSTED, S. M. (1982). SUPERTREE - decision tree processing program. Menlo Park, California: Strategic Decisions Group.

PEARL, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, San Mateo, CA: Morgan Kaufmann.

RAIFFA, H. (1968). *Decision Analysis: Introductory Lectures On Choices Under Uncertainty*. Reading, Massachusetts: Addison-Wesley.

RUSSELL, S. (1985). The complete guide to MRS. Stanford Knowledge Systems Laboratory, Report No. KSL-85-12. Stanford, California: Stanford University Computer Science Department.

SHACHTER, R. D. (1986a). Evaluating influence diagrams. *Operations Research*, 34: 871-882.

SHACHTER, R. D. (1986b). DAVID: Influence diagram processing environment for the Macintosh. In *Proceedings of the RCA/AAAI Workshop on Uncertainty and Probability in Artificial Intelligence*.

SHAW, M.L.G. & GAINES, B.R. (1988). A methodology for recognizing conflict, correspondence, consensus, and contrast in a knowledge acquisition system. *Proceedings of the Third AAI Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Banff, Canada, November.

SHEMA, D.B. & BOOSE, J.H. (1988). Refining problem-solving knowledge in repertory grids using a consultation mechanism. *International Journal of Man-Machine Studies*, 29, 447-460.

SHEMA, D. B., BRADSHAW, J.M., COVINGTON, S.P. & BOOSE, J.H. (1990). Design knowledge capture and alternative generation using possibility tables in *Canard. Knowledge Acquisition Journal*, in press.

SIVARD, C., ZWEBEN, M., CANNON, D., LAKIN, F. & LEIFER, L. (1989). Conservation of design knowledge. *Proceedings of the 27th Aerospace Sciences Meeting of the American Institute of Aeronautics and Astronautics*, Reno, Nevada, January 9-12.

SMITH, J. E. (1988). On the representation and evaluation of influence diagrams. Unpublished PIE project internal technical document.

SPIEGELHALTER, D.J., FRANKLIN, R.C.G. & BULL, K. (1990) Assessment, criticism and improvement of imprecise subjective probabilities for a medical expert system. In M. Henrion, R.D. Shachter, L.N. Kanal, and J.F. Lemmer (Eds.), *Uncertainty in Artificial Intelligence 5*. North-Holland: Elsevier.

STOUT, J., GREEN, L., & MARCUS, S. (1989). Using a KA tool to test strategy variations. *Proceedings of the Fourth AAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, Canada, October, 1989.

SULLIVAN, K. M. & SHULMAN, H.M. (1989). Cyclosporine and long-term follow-up care of the marrow transplant recipient: Clinical histories of chronic graft-versus-host disease, obliterative bronchiolitis and graft-versus-leukemia effect. *Transplant Proc.*, in press.

SZOLOVITZ, P. & PAUKER, S.G. (1978). Categorical and probabilistic reasoning in medical diagnosis. *Artificial Intelligence*, 11:115-144.

WELLMAN, M. (1986). Reasoning about assumptions underlying mathematical models. In J. S. Kowalik (ed.), *Coupling Symbolic and Numeric Computing*. Amsterdam: Elsevier.

ZWICKY, F. (1969). *Discovery, Invention, Research through the Morphological Approach*. New York: Macmillan.