

Cognitive Elements of Human Decision Making

Jens Pohl

Collaborative Agent Design Research Center, California Polytechnic State University, San Luis Obispo, California, USA

Summary. This chapter presents some understandings of the human problem-solving activity that a group of researchers in the Collaborative Agent Design Research Center at the California Polytechnic State University, San Luis Obispo, California has gained over the past two decades. Based on the premise that the human decision-maker should be an integral component of any computer-based decision-support system, it follows that the elements that appear to be important to the user should be incorporated in the design of these systems. The complexity of the human cognitive system is evidenced by the large body of literature that describes problem-solving behavior and the relatively fewer writings that attempt to provide comprehensive explanations of this behavior. The contributions of this chapter are confined to the identification of important elements of the problem-solving activity and exploration of how these elements might influence the design of a decision-support system.

2.1 Introduction

One could argue that among all attributes it is the intellectual capabilities that have allowed human beings to gain superiority over all animal species on planet Earth. These intellectual capabilities have allowed us humans to adapt to our environment, protect ourselves from predators, and ensure the availability of an adequate supply of food and water under all but the most extreme circumstances. Furthermore, these intellectual capabilities have evolved from very primitive beginnings into much more sophisticated and specialized skills (e.g., observation, planning, coordination, and problem solving), over the relatively short period of a few thousand years.

In this chapter¹ the author will explore the strengths and limitations of human beings within the context of an evolving information society in which the ability to analyze problem situations, formulate and evaluate solution alternatives, and make accurate and timely decisions, are highly valued attributes. It would appear appropriate that the designers and developers of intelligent software systems should not only seek to advance the state-of-the-art of artificial intelligence (AI), but also consider how AI-based decision-support systems can best compliment human decision-making capabilities. In this respect it is also necessary to explore the problem solving techniques that we devised before the advent of computer technology, over many centuries, to suit our human cognitive capabilities. The implication is that computers may be a necessary component of human evolution by overcoming some of the limitations of our intellectual assets.

2.2 Will Technology and Biology Merge?

The principal enabling characteristics of the Information Society are revolutionary advances in computer, bio-electronic, and communication technologies. By utilizing these technological advances a single person is able to achieve today what entire organizations struggled to accomplish only three decades ago. However, at the same time, these new opportunities are placing unprecedented pressure on the individual to perform at a significantly higher level of expectation. How will the human intellectual capabilities that have served us so well in the past measure up in this new era of unprecedented opportunities and corresponding expectations? To what degree can AI-based software systems extend our intellectual capabilities and where should this assistance be best focused?

Kurzweil (1999) argues rather convincingly that technology and biology will merge over the next millennium to significantly accelerate human evolution. Recent developments in subcutaneous sensors and prosthetics (Finn, 1997), bio-engineered materials (Kelly, 1994), brain scanning (Kiernan, 1998; Hübener, 1997; Powledge, 1997), and unraveling of the human genome (DOE, 2000), appear to be only the beginning of bio-electronic advances that promise profound extensions to the quality, productivity and longevity of human life (Brooks, 2002). In Kurzweil's words (Brockman, 2002) "... We are entering a new era. I call it the Singularity. It's a merger between human intelligence and machine intelligence ..."

¹ Portions of the material included in this chapter have appeared previously in keynote addresses presented at two InterSymp Conferences (Pohl, 2002, 2006) and two Technical Reports of the Collaborative Agent Design Research Center at the California Polytechnic State University, San Luis Obispo, California (Pohl et al., 1994, 1997)

2.3 Some Human Problem Solving Characteristics

Human beings are inquisitive creatures by nature who seek explanations for all that they observe and experience in their living environment. While this quest for understanding is central to our success in adapting to a changing and at times unforgiving environment, it is also a major cause for our willingness to accept partial understandings and superficial explanations when the degree of complexity of the problem situation confounds our mental capabilities. In other words, a superficial or partial explanation is considered better than no explanation at all. As flawed as this approach may be, it has helped us to solve difficult problems in stages. By first oversimplifying a problem we are able to develop an initial solution that is later refined as a better understanding of the nature of the problem evolves. Unfortunately, this requires us to contend with another innately human characteristic, our inherent resistance to change and aversion to risk taking. Once we have found an apparently reasonable and workable explanation or solution we tend to lose interest in pursuing its intrinsic shortcomings and increasingly believe in its validity. Whether driven by complacency or lack of confidence, this state of affairs leads to many surprises. We are continuously discovering that what we believed to be true is only partly true or not true at all, because the problem is more complicated than we had previously assumed it to be.

The complexity of the problems faced by human society in areas such as management, economics, marketing, engineering design, and environmental preservation, is increasing for several reasons. First, computer-driven information systems have expanded these areas from a local to an increasingly global focus. Even small manufacturers are no longer confined to a regionally localized market for selling their products. The marketing decisions that they have to make must take into account a wide range of factors and a great deal of knowledge that is far removed from the local environment. Second, as the net-centricity of the problem system increases so do the relationships among the various factors. These relationships are difficult to deal with, because they require the decision-maker to consider many factors concurrently. Third, as the scope of problems increases decision-makers suffer simultaneously from two diametrically opposed but related conditions. They tend to be overwhelmed by the sheer volume of data that they have to consider, and yet they lack information in many specific areas. To make matters worse, the information tends to change dynamically in largely unpredictable ways.

It is therefore not surprising that governments, corporations, businesses, down to the individual person, are increasingly looking to computer-based decision-support systems for assistance. This has placed a great deal of pressure on software developers to rapidly produce applications that will overcome the apparent failings of the human decision-maker. While the expectations have been very high, the delivery has been much more modest. The expectations were simply unrealistic. It was assumed that advances in technology would be simultaneously accompanied by an understanding of how these

advances should be applied optimally to assist human endeavors. History suggests that such an a priori assumption is not justified. There are countless examples that would suggest the contrary. For example, the invention of new materials (e.g., plastics) has inevitably been followed by a period of misuse. Whether based on a misunderstanding or lack of knowledge of its intrinsic properties, the new material was typically initially applied in a manner that emulated the material(s) it replaced. In other words, it took some time for the users of the new material to break away from the existing paradigm. A similar situation currently exists in the area of computer-based decision-support systems.

2.4 Human Limitations and Weaknesses

Deeply embedded in the evolution of the human intellect is the rationalistic approach to problem solving. At face value this approach appears to be entirely sound. It suggests that problem solving should proceed in a logical sequence of clearly defined steps. One begins by defining the problem and then decomposes the defined problem into sub-problems. Decomposition appears to make a great deal of sense because the parts of a problem are intrinsically easier to solve than the whole problem. The reason for this is that the complexity of a problem is normally due to the nature and number of relationships among the elements of the problem and not due to the elements themselves. Decomposition allows us to temporarily neglect consideration of many of these relationships. However, this over-simplification of the problem is valid only as long as the problem remains in a decomposed state. As soon as we try to integrate the separate solutions of the parts into a solution of the whole the relationships that we so conveniently disregarded reappear and invalidate many if not most of our neatly packaged sub-solutions. We find to our consternation that the characteristics of a part of a problem situation considered in isolation are not necessarily similar (let alone the same) as the behavior of that part within the context of the whole problem.

Why have we human beings come to rely so heavily on this flawed approach to problem solving? The reasons are related primarily to the biological nature of our cognitive system. While the biological basis of human cognition is massively parallel (i.e., millions of neurons and billions of connections) our conscious reasoning capabilities are largely sequential. The fact is that our short term memory has a severely limited capacity of only a few chunks of data at any one time. Therefore, we can differentiate among only a small number of objects at any one point in time, even though we continuously move new data chunks from long term memory into short term memory. As a consequence we have great difficulty dealing with more than three or four relationships concurrently.

Secondary limitations and tensions that contribute to our human problem solving difficulties include our tendency to seek a degree of accuracy that is

often unrealistic and usually unnecessary. Our aversion to risk and instinctive need to survive drives us to try to predict the future with great accuracy. In this respect, as mentioned previously, we place a great deal of reliance on mathematics even though mathematical models often fail due to oversimplification of the problem situation and incorrect boundary assumptions (Pohl, 1999).

We often seek to produce an optimum solution even though the problem conditions are continuously changing and, therefore, we have no benchmark that would allow us to judge whether a particular solution is in fact optimal. In other words, under dynamic conditions there is no static benchmark available. This creates related difficulties, because our ability to interpret and judge any situation is necessarily based on comparative analysis. Subject to the experiential basis of the human cognitive system we normally have no alternative but to measure new situations with existing metrics based on past experience. However, the further the new situation deviates from past experience the more misleading the available metrics are likely to be. As a result, since we have no effective metrics for assessing new situations, we typically require a considerable period of time to correctly evaluate such situations. Accordingly, it is not unreasonable to conclude that human judgments are more influenced by the past than the present.

More comprehensively stated, the essentially experience-based nature of human cognition forces us almost always (i.e., at least initially) to apply existing methods, notions, and concepts to new situations. Therefore, our most effective problem solving capabilities utilize prototype solutions based on past experience. While we have become quite skilled in adapting, modifying and combining such prototype solutions, we find it very difficult to create new prototypes. As a consequence we invariably apply existing solution methods to new problem situations and develop new methods only through painful trial and error. This also leads us to generally underestimate the complexity and impact of new situations.

2.5 Human Strengths

So far the discussion has centered on the apparently numerous limitations and weaknesses of human beings, particularly in respect to intellectual and emotional capabilities. Surely we human beings also have intellectual strengths. The answer is yes, of course, but with some qualifications. Certainly human learning capabilities, supported by a very large associative long-term memory, are vast. However, our rate of learning is rather slow and appears to lack efficiency. While some of this inefficiency is undoubtedly due to human communication inadequacies, the very process of progressively collecting experience by building onto existing associative knowledge structures would appear to be cumbersome and rather time consuming. It is not simply a matter of

adding new knowledge elements or associating existing elements by inserting linkages, but instead patterns of neural activations (i.e., firings) have to be repeated many times before they are literally grooved into long-term memory. It is therefore not surprising that formal education takes up one quarter to one third of a human life span and involves a great deal of concentration, as well as assistance from other human beings who have acquired special teaching skills.

An important part of the human learning capability is the ability to conceptualize experiences into knowledge that we consider to be true in most cases. In this way we place emphasis on being able to deal with general conditions and consider the exceptions to the general rules to be much less important. This again exemplifies the human tendency to oversimplify a situation for the sake of being able to reach a quick solution to a problem or an explanation of an observed phenomenon. In fact, as we discover to our surprise time and again, the exceptions are often more important than the generalizations (Minsky, 1990).

It must also be noted that much of human learning is involuntary and therefore virtually effortless. This applies in particular to the acquisition of low-level, largely involuntary skills such as sensory pattern matching that allows us to automatically convert data to information. For example, when we enter a restaurant we immediately recognize the furniture in the room. In fact, our eyes see only image patterns. However, these are automatically interpreted as tables and chairs by our cognitive system which has by experience related these image patterns to the appropriate symbolic entities.

At a higher level, symbolic reasoning allows us to infer knowledge from information. When our reasoning capabilities are unable to cope in complex situations that include many relationships, conceptual pattern matching (i.e., intuition) allows us to assess situations without resorting to logical reasoning. However, again there is evidence that this process is greatly facilitated by experience. Klein (1998) found that highly experienced fire captains will resort to the situation analysis methods employed by novices when they are confronted with situations outside their sphere of experience.

While the creation of new knowledge is normally the province of individuals, once such an intellectual leap has been accomplished we collectively excel in the technological exploitation of this contribution. Typically, this exploitation proceeds incrementally and involves a large number of persons, coordinated in a self-organizing fashion but willing to collaborate to leverage the capabilities of individual contributors.

However, finally, perhaps one of our greatest human strengths is the discovery early on in our evolution of the usefulness of tools. Since then we have been successful in the development and application of more and more powerful tools. Today, we appear to be on the verge of merging computer-based tools with the biological fabric of our very existence.

2.6 The Rationalistic Tradition

To understand current trends in the evolution of progressively more sophisticated decision-support systems it is important to briefly review the foundations of problem solving methodology from an historical perspective. Epistemology is the study or theory of the origin, nature, methods and limits of knowledge. The dominant epistemology of Western Society has been technical rationalism (i.e., the systematic application of scientific principles to the definition and solution of problems).

The rationalistic approach to a problem situation is to proceed in well defined and largely sequential steps as shown in Fig. 2.1: define the problem; establish general rules that describe the relationships that exist in the problem system; apply the rules to develop a solution; test the validity of the solution; and, repeat all steps until an acceptable solution has been found. This simple view of problem solving suggested a model of sequential decision-making that has retained a dominant position to the present day. With the advent of computers it was readily embraced by first Wave software (Fig. 2.2) because of the ease with which it could be translated into decision-support systems utilizing the procedural computer languages that were available at the time.

The close correlation between the rationalistic approach and what is commonly referred to as the scientific method is readily apparent in the series of basic steps that are employed in scientific investigations: observe the phenomenon that requires explanation; formulate a possible explanation; develop a method capable of predicting or generating the observed phenomenon; interpret the results produced by the method; and, repeat all steps until an acceptable explanation of the observed phenomenon has been found. Scientific research typically attempts to establish situations in which observable actions (or reactions) are governed by a small number of variables that can be systematically manipulated. Every effort is made to keep the contrived situation

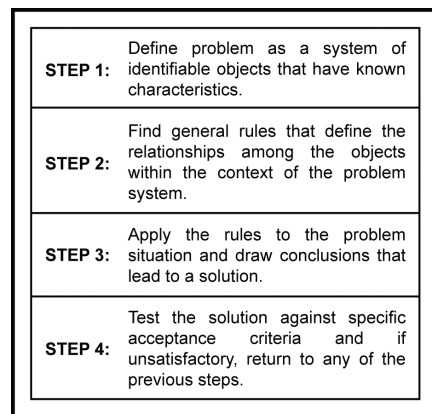


Fig. 2.1. Solution of simple problems

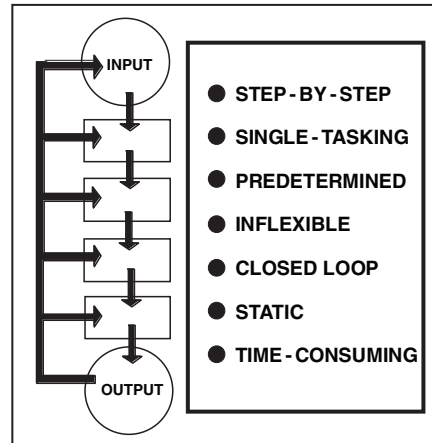


Fig. 2.2. Sequential decision-support

simple, clear and deterministic, so that the results of the simulation can be verified.

However, natural phenomena and real world problems are often very complex involving many related variables. Neither the relationships among the variables nor the variables themselves are normally sufficiently well understood to provide the basis for clear and comprehensive definitions. In other words, problem situations are often too complex to be amenable to an entirely logical and predefined solution approach. Under these circumstances the analytical strategy has been to decompose the whole into component parts, as follows:

- Decompose the problem system into sub-problems.
- Study each sub-problem in relative isolation, using the rationalistic approach (Fig. 2.1). If the relationships within the sub-problem domain cannot be clearly defined then decompose the sub-problem further.
- Combine the solutions of the sub-problems into a solution of the whole.

Underlying this problem-solving strategy is the implicit assumption that an understanding of parts leads to an understanding of the whole. Under certain conditions this assumption may be valid. However, in many complex problem situations the parts are tightly coupled so that the behavior of the whole depends on the interactions among the parts rather than the internal characteristics of the parts themselves (Bohm, 1983; Senge, 1993). An analogy can be drawn with the behavior of ants. Each ant has only primitive skills, such as the ability to interpret the scent of another ant and the instinctive drive to search for food, but little if any notion of the purpose or objectives of the ant colony as a whole. In other words, an understanding of the behavior of an individual ant does not necessarily lead to an understanding of the community behavior of the ant colony of which the ant is a part.

Decomposition is a natural extension of the scientific approach to problem solving and has become an integral and essential component of rationalistic methodologies. Nevertheless, it has serious limitations. First, the behavior of the whole usually depends more on the interactions of its parts and less on the intrinsic behavior of each part. Second, the whole is typically a part of a greater whole and to understand the former we have to also understand how it interacts with the greater whole. Third, the definition of what constitutes a part is subject to viewpoint and purpose, and not intrinsic in the nature of the whole. For example, from one perspective a coffee maker may be considered to comprise a bowl, a hotplate, and a percolator. From another perspective it consists of electrical and constructional components, and so on.

Rationalism and decomposition are certainly useful decision-making tools in complex problem situations. However, care must be taken in their application. At the outset it must be recognized that the reflective sense (Schön, 1983) and intuition of the decision-maker are at least equally important tools. Second, decomposition must be practiced with restraint so that the complexity of the interactions among parts is not overshadowed by the much simpler behavior of each of the individual parts. Third, it must be understood that the definition of the parts is largely dependent on the objectives and knowledge about the problem that is currently available to the decision-maker. Even relatively minor discoveries about the greater whole, of which the given problem situation forms a part, are likely to have significant impact on the purpose and the objectives of the problem situation itself.

2.7 Decision-Making in Complex Problem Situations

As shown in Fig. 2.3, there are several characteristics that distinguish a complex problem from a simple problem. First, the problem is likely to involve many related issues or variables. As discussed earlier the relationships among the variables often have more bearing on the problem situation than the variables themselves. Under such tightly coupled conditions it is often not particularly helpful, and may even be misleading, to consider issues in isolation. Second, to confound matters some of the variables may be only partially defined and some may yet to be discovered. In any case, not all of the information that is required for formulating and evaluating alternatives is available. Decisions have to be made on the basis of incomplete information.

Third, complex problem situations are pervaded with dynamic information changes. These changes are related not only to the nature of an individual issue, but also to the context of the problem situation. For example, a change in wind direction during a major brushfire may have a profound impact on the entire nature of the relief operation. Apart from precipitating an immediate re-evaluation of the firefighting strategy, it may require the relocation of firefighters and their equipment, the re-planning of evacuation routes, and possibly even the relocation of distribution centers. Certainly, a change in the

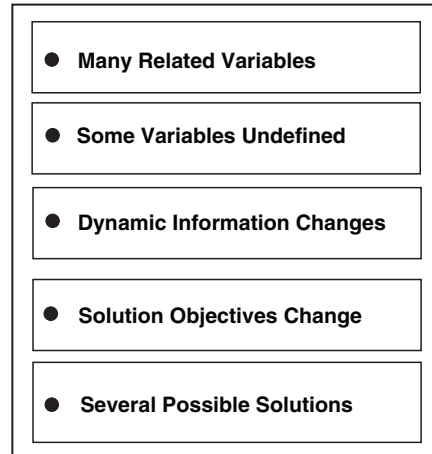


Fig. 2.3. Character of complex problems

single factor of wind direction could, due to its many relationships, call into question the very feasibility of the existing course of action (i.e., the firefighting plan). Even under less critical conditions it is not uncommon for the solution objectives to change several times during the decision-making process. This fourth characteristic of complex problem situations is of particular interest. It exemplifies the tight coupling that can exist among certain problem issues, and the degree to which decision-makers must be willing to accommodate fundamental changes in the information that drives the problem situation.

Fifth, complex problems typically have more than one solution (Archea, 1987). A solution is found to be acceptable if it satisfies certain performance requirements and if it has been determined that the search for alternatives is no longer warranted. Such a determination is often the result of resource constraints (e.g., availability of time, penalty of non-action, or financial resources) rather than a high level of satisfaction with the quality of the proposed solution.

While human decision-making in complex problem situations has so far defied rigorous scientific explanation, we do have knowledge of at least some of the characteristics of the decision-making activity.

- In search of a starting point for assessing the nature of the problem situation, decision makers will typically look for prominent aspects that are likely to have significant impact on the solution space. These aspects or problem features are then used to at least temporarily establish boundaries within which the initial problem solving efforts can be focused. However, the qualifying terms temporarily and initial are appropriately chosen since both the selected features and the early boundaries are likely to change many times due to the highly dynamic nature of the decision-making process.

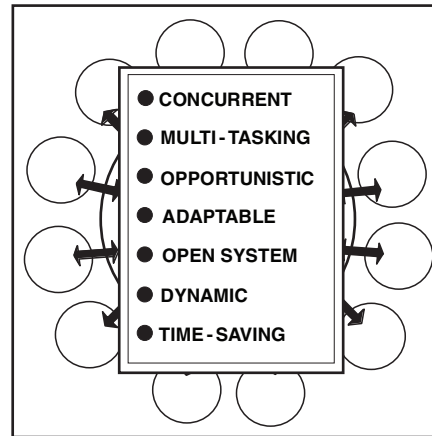


Fig. 2.4. Parallel decision-support

- The complexity of the decision-making activity does not appear to be due to a high level of difficulty in any one area but the multiple relationships that exist among the many issues that impact the desired outcome. Since a decision in one area will tend to influence several other areas there is a need to consider many factors at the same time. This places a severe burden on the human cognitive system. As mentioned previously, although the neurological mechanisms that support conscious thought processes are massively parallel, the operation of these reasoning capabilities is largely sequential. Accordingly, decision-makers tend to apply simplification strategies for reducing the complexity of the problem-solving activity. In this regard it becomes readily apparent why second Wave software with multi-tasking capabilities provides a much more useful architecture for decision-support systems (Fig. 2.4).
- Observation of decision-makers in action has drawn attention to the important role played by experience gained in past similar situations, knowledge acquired in the general course of decision-making practice, and expertise contributed by persons who have detailed specialist knowledge in particular problem areas. The dominant emphasis on experience is confirmation of another fundamental aspect of the decision-making activity. Problem-solvers seldom start from first principles. In most cases, the decision-maker builds on existing solutions from previous situations that are in some way related to the problem under consideration. From this viewpoint, the decision-making activity involves the modification, refinement, enhancement and combination of existing solutions into a new hybrid solution that satisfies the requirements of the given problem system. In other words, problem-solving can be described as a process in which relevant elements of past prototype solution models are progressively and collectively molded

into a new solution model. Very seldom are new prototype solutions created that do not lean heavily on past prototypes.

Finally, there is a distinctly irrational aspect to decision-making in complex problem situations. Schön refers to a “...reflective conversation with the situation...” (Schön, 1983). He argues that decision-makers frequently make value judgments for which they cannot rationally account. Yet, these intuitive judgments often result in conclusions that lead to superior solutions. It would appear that such intuitive capabilities are based on a conceptual understanding of the situation, which allows the problem solver to make knowledge associations at a highly abstract level.

Based on these characteristics the solution of complex problems can be categorized as an information intensive activity that depends for its success largely on the availability of information resources and, in particular, the experience and reasoning skills of the decision-makers. It follows that the quality of the solutions will vary significantly as a function of the problem-solving skills, knowledge, and information resources that can be brought to bear on the solution process. This clearly presents an opportunity for the useful employment of computer-based decision-support systems in which the capabilities of the human decision-maker are complemented with knowledge bases, expert agents, and self-activating conflict identification and monitoring capabilities.

2.8 Principal Elements of Decision-Making

Over the past two decades that our Collaborative Agent Design Research Center has been developing distributed, collaborative decision-support systems some insights have been gained into the nature of the decision-making activity. In particular, we have found it useful to characterize decision-making in terms of six functional elements (Fig. 2.5): *information*; *representation*; *visualization*; *communication*; *reasoning*; and, *intuition*.

2.8.1 The Information Element

Decision-making in complex problem situations is a collaborative activity involving many sources of information that are often widely dispersed. Seldom is all of the information required for the solution, or even only a component of the problem, physically located in the immediate vicinity of the decision-maker. In fact, much of the information is likely to reside in remote repositories that can be accessed only through electronic means, the telephone, e-mail, or the temporary relocation of a member of the problem-solving team (Fig. 2.6). If the desired information requires expert advice the services of a consultant may be required in addition to, or instead of, access to an information resource.

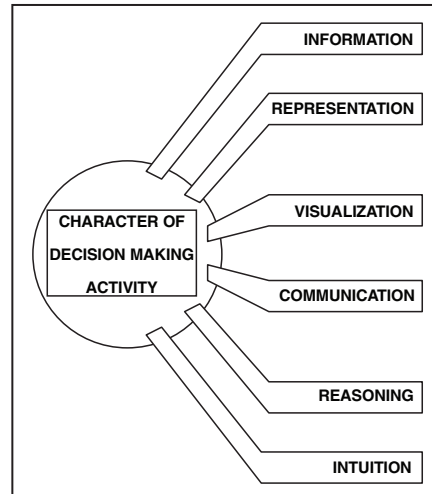
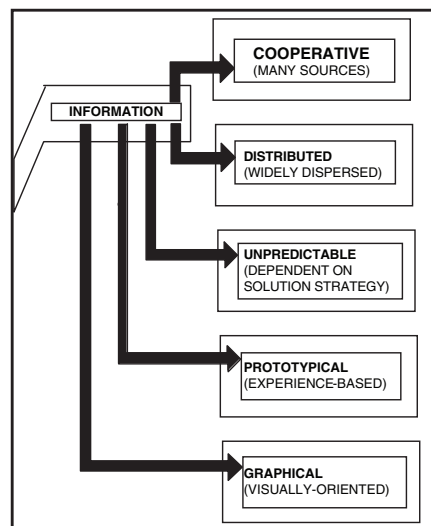


Fig. 2.5. Decision-making elements

Fig. 2.6. The *information* element

The term information is used here in the broadest sense to include not only factual data and the progressively more comprehensive and detailed description of the problem system, but also the many knowledge bases that are part of the local and global environment within which the problem situation is constituted. In this regard, we are concerned with the knowledge of the individual members of the problem-solving team, the knowledge of peripheral players (e.g., colleagues, associates and consultants), the collective knowledge

of the profession (such as the various engineering professions, the military establishment, or the management profession) and industry, and beyond that those aspects of what might be referred to as global knowledge that impact the problem context.

Typically, the problem specifications (i.e., constraints, criteria, and objectives) evolve with the problem solution as the decision-makers interact with the problem situation. Accordingly, the information requirements of the problem solver are not predictable since the information needed to solve the problem depends largely on the solution strategy adopted (Fischer and Nakakoji, 1991). In this respect problem solving is a learning process in which the decision-maker progressively develops a clearer understanding of the problem that is required to be solved. Much of the information that decision-makers use in the development of a problem solution is gleaned from experience with past projects. In fact, it can be argued that solutions commonly evolve out of the adaptation, refinement and combination of prototypes (Gero, 1988). This argument suggests that the more expert human decision-makers are the more they tend to rely on prototypical information in the solution of complex problems. It would appear that the accumulation, categorization and ability to apply prototype knowledge are the fundamental requirements for a human decision-maker to reach the level of expert in a particular domain. Based largely on the work of Gero (1988) and Rosenman and Gero (1993) the following techniques used by engineering designers to develop solutions through the manipulation of prototypes can be identified as being universally applicable to other problem domains:

- *Refinement.* The prototype can be applied after changes have been made in the values of parameter variables only (i.e., the instance of the prototype is reinterpreted within the acceptable range of the parameter variables).
- *Adaptation.* Application of the prototype requires changes in the parameters that constitute the description of the prototype instance, based on factors that are internal to the prototype (i.e., a new prototype instance is produced).
- *Combination.* Application of the prototype requires the importation of parameter variables of other prototypes, producing a new instance of a reinterpreted version of the original prototype.
- *Mutation.* Application of the prototype requires structural changes to the parameter variables, either through internal manipulations or the importation of parameter variables from external sources (i.e., either a reinterpreted version of the original prototype or a new prototype is produced).
- *Analogy.* Creation of a new prototype based on a prototype that exists in another context, but displays behavioral properties that appear to be analogous to the application context.

For application purposes in knowledge-based decision-support systems prototypes may be categorized into five main groups based on knowledge content (Schön, 1988; Pohl and Myers, 1994):

1. Vertical prototype knowledge bases that contain typical object descriptions and relationships for a complete problem situation or component thereof. Such a knowledge base may include all of the types that exist in a particular problem setting, for example: an operational template for a particular kind of humanitarian relief mission; a certain type of propulsion unit; or, a building type such as a library, sports stadium, or supermarket.
2. Horizontal prototype knowledge bases that contain typical solutions for sub-problems such as commercial procurement practices, construction of a temporary shelter, or techniques for repairing equipment. This kind of knowledge often applies to more than one discipline. For example, the techniques for repairing a truck apply equally to the military as they do to auto-repair shops, engineering concerns, and transportation related organizations.
3. Domain prototype knowledge bases that contain guidelines for developing solutions within contributing narrow domains. For example, the range of structural solutions appropriate for the construction of a 30-story office building in Los Angeles is greatly influenced by the seismic character of that region. Posed with this design problem structural engineers will immediately draw upon a set of rules that guide the design activity. Similarly, an acoustic consultant engaged to determine the cause of noise transmission between two adjacent office spaces will diagnose the problem based on experience with previous situations. The possibility of the existence of indirect sound transmission paths, such as a false ceiling, is likely to receive early consideration.
4. Exemplar prototype knowledge bases that describe a specific instance of an object type or solution to a sub-problem. Exemplary prototypes can be instances of vertical or horizontal prototypes, such as a particular building type or a method of welding a certain kind of steel joint that is applied across several disciplines and industries (e.g., building industry and automobile industry). Decision-makers often refer to exemplary prototypes in exploring solution alternatives to sub-problems.
5. Experiential knowledge bases that represent the factual prescriptions, strategies and solution conventions employed by the decision-maker in solving similar kinds of problem situations. Such knowledge bases are typically rich in methods and procedures. For example, a particularly memorable experience such as the deciding event in a past business negotiation or the experience of seeing for the first time the magnificent sail-like concrete shell walls of the Sydney Opera House, may provide the basis for a solution method that is applied later to create a similar experience in a new problem situation that may be quite different in most other respects. In other words, experiential prototypes are not bound to a specific type of

problem situation. Instead, they represent techniques and methods that can be reproduced in various contexts with similar results. Experiential knowledge is often applied in very subtle ways to guide the solution of sub-problems (e.g., a subterfuge in business merger or take-over negotiations that is designed to mislead a competing party).

The amount of prototypical information is potentially overwhelming. However, the more astute and experienced decision-maker will insist on taking time to assimilate as much information as possible into the problem setting before committing to a solution theme. There is a fear that early committal to a particular solution concept might overlook characteristics of the problem situation that could gain in importance in later stages, when the solution has become too rigid to adapt to desirable changes. This reluctance to come to closure places a major information management burden on the problem solver. Much of the information cannot be specifically structured and prepared for ready access, because the needs of the problem solver cannot be fully anticipated. Every step toward a solution generates new problems and information needs (Simon, 1981).

One of the early targets of ontology-based multi-agent systems is the integration of existing information sources (i.e., databases and legacy applications) into comprehensive decision-support systems. The initial focus of such efforts is the linkage of existing databases. This is not a trivial task since these existing information resources typically were implemented in many different ways. Consequently, any integrating system will be required to support the conceptual integration of a variety of data resource models. This can be accomplished through the provision of several internal-level database representations, requiring a number of additional mapping functions to link the internal and conceptual representation levels. In this way, any externally linked database can be removed or replaced by another database, simply by recoding the internal to conceptual level mapping. Since this will not affect the external data representation, the user-interfaces built at the external level will also remain unchanged.

The scope of database query facilities desirable for the kind of multi-agent, decision-support environment discussed here far exceeds traditional database management system (DBMS) functions. They presuppose a level of embedded intelligence that has not been available in the past. Some of these desirable features include: conceptual searches instead of factual searches; automatically generated search strategies instead of predetermined search commands; multiple database access instead of single database access; analyzed search results instead of direct (i.e., raw) search results; and, automatic query generation instead of requested searches only.

A traditional DBMS typically supports only factual searches. In other words, users and applications must be able to define precisely and without ambiguity what data they require. In complex problem situations users rarely know exactly what information they require. Often they can define in only

conceptual terms the kind of information that they are seeking. Also, they would like to be able to rely on the DBMS to automatically broaden the search with a view to discovering information.

This suggests, in the first instance, that an intelligent DBMS should be able to formulate search strategies based on incomplete definitions. It should be able to infer, from rather vague information requests and its own knowledge of the requester and the problem context, a set of executable query procedures. To facilitate this process the DBMS should maintain a history of past information requests, the directed search protocols that it generated in response to these requests, and at least some measure of the relative success of the previous search operations.

A traditional DBMS normally provides access to only a single database. A knowledge-based decision-support environment is likely to involve many information sources, housed in a heterogeneous mixture of distributed databases. Therefore, through the internal-level database representations discussed earlier, the DBMS must be able to access multiple databases. Using the mapping functions that link these internal representations an intelligent DBMS should be capable of formulating the mechanisms required to retrieve the desired data from each source, even though the internal data structures of the sources may differ widely. Particularly when search results are derived from multiple sources and the query requests themselves are vague and conceptual in nature, there is a need for the retrieved information to be reviewed and evaluated before it is presented to the requester. This type of search response formulation facility has not been necessary in a traditional DBMS, where users are required to adhere to predetermined query protocols that are restricted to a single database.

Finally, all of these capabilities (i.e., conceptual searches, dynamic query generation, multiple database access, and search response formulation) must be able to be initiated not only by the user but also by any of the computer-based agents that are currently participating in the decision-making environment. These agents may be involved in any number of tasks that require the import of additional information from external databases into their individual knowledge domains.

A conceptual model of an intelligent DBMS interface (Fig. 2.7) with the capabilities described above should be able to support the following typical information search scenario that might occur in an integrated and distributed, collaborative, multi-agent, decision-support environment. Queries that are formulated either by the user or generated automatically by a computer-based agent are channeled to a Search Strategy Generator. The latter will query a Search Scenario Database to determine whether an appropriate search strategy already exists from a previous search. If not, a new search strategy is generated, and also stored in the Search Scenarios Database for future use. The search strategy is sent to the Database Structure Interpreter, which automatically formulates access protocols to all databases that will be involved in the proposed search. The required access and protocol information, together

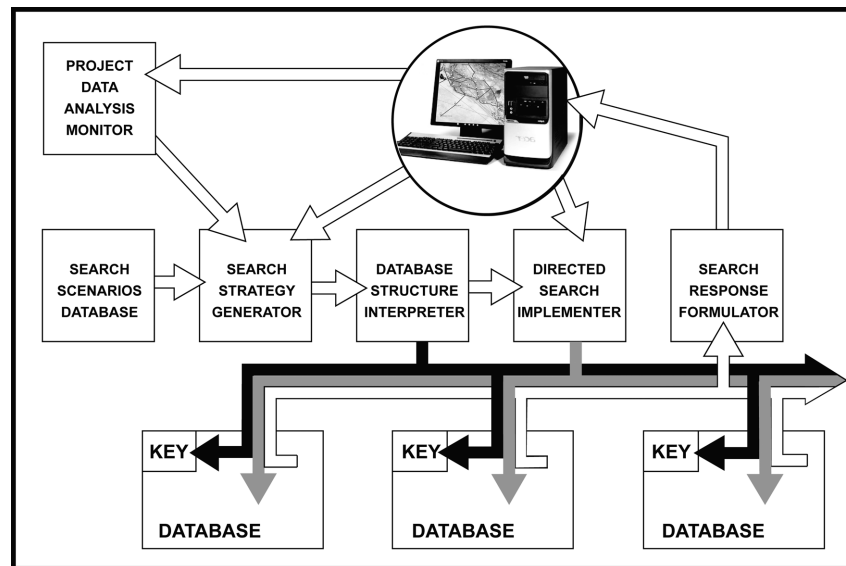


Fig. 2.7. Conceptual model of an intelligent DBMS interface

with the search strategy, are sent to the Directed Search Implementer, which conducts the required database searches. The results of the search are sent to a Research Response Formulator, where the raw search results are analyzed, evaluated and combined into an intelligent response to be returned to the originator of the query.

The proposition that the DBMS interface should be able to deal with incomplete search requests warrants further discussion. When searching for information, partial matching is often better than no response. In traditional query systems, a database record either matches a query or it does not. A flexible query system, such as the human brain, can handle inexact queries and provide best guesses and a degree of confidence for how well the available information matches the query (Pohl et al., 1994). For example, let us assume that a military commander is searching for a means of trapping a given enemy force in a particular sector of the battlefield and formulates a something like a choke point query. In a flexible query system a something like operator would provide the opportunity to match in a partial sense, such as: terrain conditions that slow down the movement of troops; unexpected physical obstacles that require the enemy to abruptly change direction; subterfuge that causes enemy confusion; and so on. These conditions can all, to varying extent, represent something like a choke point that would be validated by a degree of match qualification.

Flexible query processing systems are fairly common. For example, most automated library systems have some level of subject searching by partial keyword or words allowing users to browse through a variety of related topics.

Even word-processing programs include spelling checkers, which by their very nature search for similar or related spellings. However, even a flexible query system cannot automatically form hypotheses, since the system does not know what to ask for.

The ability to search for something like is only a starting point. How can the system be prompted to search for vaguely or conceptually related information? For example, how can the system discover the intuitive connection between a physical choke point, such as a narrow cross-corridor in a mountainous battlefield, and a precision fire maneuver aimed at concentrating enemy forces in an exposed area? In other words, how can the system show the commander that the precision fire maneuver option can satisfy the same intent as the cross-corridor option? In addition, the system must not overwhelm the commander with an unmanageable number of such intuitive speculations. To discover knowledge it is necessary to: form a hypothesis; generate some queries; view and analyze the results; perhaps modify the hypothesis and generate new queries; and, repeat this cycle until a pattern emerges. This pattern may then provide insight and advice for intuitive searches. The goal is to automate this process with a discovery facility that repeatedly queries the prototype knowledge bases and monitors the reactions and information utilized by the decision-maker, until the required knowledge is discovered.

2.8.2 The Representation Element

The methods and procedures that decision-makers utilize to solve complex problems rely heavily on their ability to identify, understand and manipulate objects (Fig. 2.8). In this respect, objects are complex symbols that convey meaning by virtue of the explicit and implicit information that they encapsulate within their domain. For example, architects develop design solutions by

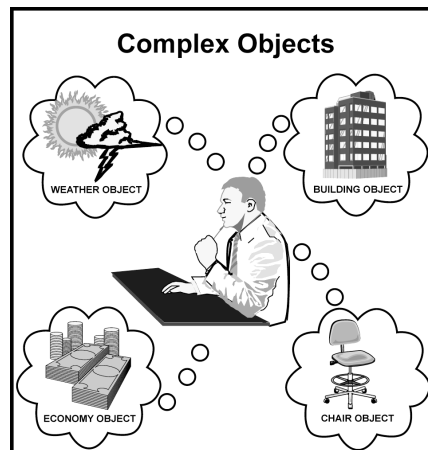


Fig. 2.8. Symbolic reasoning with objects

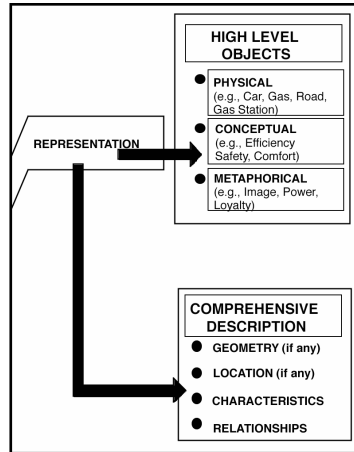


Fig. 2.9. The *representation* element

reasoning about neighborhoods, site characteristics, buildings, floors, spaces, walls, windows, doors, and so on. Each of these objects encapsulates knowledge about its own nature, its relationships with other objects, its behavior within a given environment, what it requires to meet its own performance objectives, and how it might be manipulated by the decision-maker within a given problem scenario (Fig. 2.9). This knowledge is contained in the various representational forms of the object as factual data, relationships, algorithms, rules, exemplar solutions, and prototypes.

The reliance on object representations in reasoning endeavors is deeply rooted in the innately associative nature of the human cognitive system. Information is stored in long-term memory through an indexing system that relies heavily on the forging of association paths. These paths relate not only information that collectively describes the meaning of symbols such as building, car, chair, and tree, but also connect one symbol to another. The symbols themselves are not restricted to the representation of physical objects, but also serve as concept builders. They provide a means for grouping and associating large bodies of information under a single conceptual metaphor. In fact, Lakoff and Johnson (Lakoff and Johnson, 1980) argue that "...our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature." They refer to the influence of various types of metaphorical concepts, such as "...desirable is up" (i.e., spatial metaphors) and "...fight inflation" (i.e., ontological or human experience metaphors), as the way human beings select and communicate strategies for dealing with everyday events.

Problem-solvers typically intertwine the factually based aspects of objects with the less precise, but implicitly richer language of metaphorical concepts. This leads to the spontaneous linkage of essentially different objects through

the process of analogy. In other words, the decision-maker recognizes similarities between two or more sub-components of apparently unrelated objects and embarks upon an exploration of the discovered object seeking analogies where they may or may not exist. At times these seemingly frivolous pursuits lead to surprising and useful solutions of the problem at hand.

The need for a high level representation is fundamental to all computer-based decision-support systems. It is an essential prerequisite for embedding AI in such systems, and forms the basis of any meaningful communication between user and computer. Without a high level representation facility the abilities of the computer to assist the human decision maker are confined to the performance of menial tasks, such as the automatic retrieval and storage of data or the computation of mathematically defined quantities. While even those tasks may be highly productive they cannot support a partnership in which human users and computer-based systems collaborate in a meaningful and intelligent manner in the solution of complex problems.

The term high level representation refers to the ability of computer software to process and interpret changes in data within an appropriate context. It is fundamental to the distinction between data-centric and information-centric software. Strictly speaking data are numbers and words without relationships.² Software that incorporates an internal representation of data only is often referred to as data-centric software. Although the data may be represented as objects the absence of relationships to define the functional purpose of the data inhibits the inclusion of meaningful and reliable automatic reasoning capabilities. Data-centric software, therefore, must largely rely on predefined solutions to predetermined problems, and has little (if any) scope for adapting to real world problems in near real-time.

Information, on the other hand, refers to the combination of data with relationships to provide adequate context for the interpretation of the data. The richer the relationships the more context and the greater the opportunity for automatic reasoning by software agents. Software that incorporates an internal information model (i.e., ontology) consisting of objects, their characteristics, and the relationships among those objects is often referred to as information-centric software (Pohl et al., 2005). The information model provides a virtual representation of the real world domain under consideration.

For example, let us assume that we wish to represent a component of a building such as a conference room in the computer. Until recently, in a data-centric software environment, we would have treated the conference room as a three-dimensional geometric entity that can be described in terms of points (i.e., x-y-z coordinates), lines, or surfaces. While this may be satisfactory for

² Even though data are often stored in a relational database management system, the relationships that are stored with the data in such a database are structural in nature and do not provide any information on how the data will be used (i.e., the *context* of the data)

displaying different internal views of the building space and even generating animated walk-through sequences, it does not provide a basis for the computer to reason about any aspect of the space, such as that a conference room must have a door for it to be usable. To provide the computer with such a reasoning capability the particular entity, in this case the conference room must be represented in the computer as an information structure that constitutes the context of a building. This can be achieved quite easily by storing in the computer the word building and associating this word with some characteristics such as: physical object; made of material; has height, width and length; consists of one or more floors; has spaces on floors; and so on. Then further defining spaces with characteristics such as: enclosed by walls, floor and ceiling; with wall having at least one opening referred to as a door; and so on.

In such an information-centric software environment the same conference room would be presented to and stored in the computer as part of the building information structure (i.e., ontology) to support the following kind of interaction with the user:

<i>Computer user:</i>	I would like to represent a component of a building.
<i>Computer software:</i>	Loads its stored building ontology into memory. Asks user: "What kind of a building component?"
<i>Computer user:</i>	A space of type conference.
<i>Computer software:</i>	For how many persons?
<i>Computer user:</i>	Up to 16 persons.
<i>Computer software:</i>	Suggested space size is: 16 ft (length), 14 ft (width), 8 ft (height). Suggested furniture: 6 ft by 3 ft table, 16 chairs, screen, white board. Other features: There must be at least one door.

As can be seen from the interaction between the user and the computer software, by virtue of the information structure the computer has some understanding of the meaning of a building within the context of its characteristics and the relationships of its components (i.e., floors, spaces, walls, openings, and furniture). This endows the computer software with the ability to collaborate and assist the user by reasoning about the relationships between the data entered by the user and the context contained in the simple information representation provided by the building ontology. Accordingly, driven by the context provided in the ontology, software agents are able to spontaneously reason about the characteristics of a conference room for 16 persons. Beyond establishing the need for at least one exit and the kind of furniture normally required, this could easily extend to the evaluation of the impact on equipment and functionality of an external window.

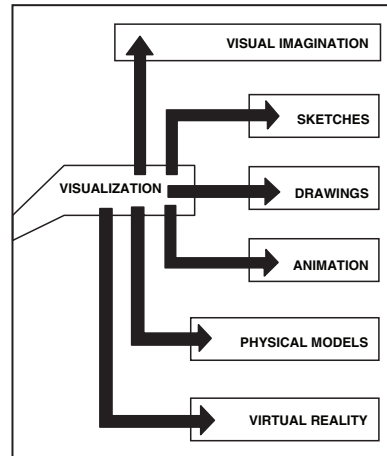


Fig. 2.10. The visualization element

2.8.3 The Visualization Element

Problem solvers use various visualization media, such as visual imagination, drawings and physical models, to communicate the current state of the evolving solution to themselves and to others (Fig. 2.10). Drawings, in particular, have become intrinsically associated with problem solving. Although the decision-maker can reason about complex problems solely through mental processes, drawings and related physical images are useful and convenient for extending those processes. The failings of the drawing as a vehicle for communicating the full intent of the decision-maker do not apply to the creator of the drawing. To the latter the drawing serves not only as an extension of short-term memory, but also as a visual bridge to the associative indexing structure of long-term memory. In this way, every meaningful part of the drawing is linked to related data and deliberation sequences that together provide an effectively integrated and comprehensive representation of the artifact.

From a technical point of view a great deal of headway has been made over the past two decades in the area of computer-based visualization. However, without high-level representation capabilities even the most sophisticated computer generated images are nothing but hollow shells. If the computer system does not have even the simplest understanding of the nature of the objects that are contained in the image then it cannot contribute in any way to the analysis of those objects. On the other hand, visualization in combination with high-level representation becomes the most powerful element of the user-interface of a decision-support system. Under these circumstances, visualization promotes the required level of understanding between the user and the computer as they collaborate in the solution of a problem.

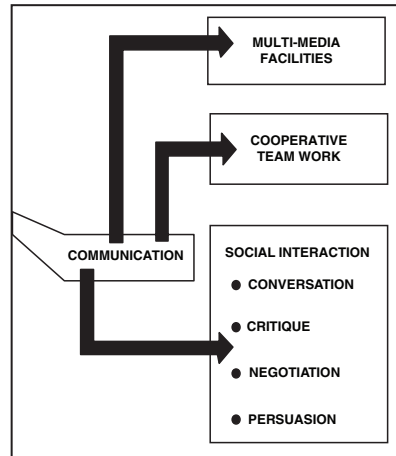


Fig. 2.11. The communication element

2.8.4 The Communication Element

The solution of complex problems is typically undertaken by a team of decision-makers. Each team member contributes within a collaborative decision-making environment that relies heavily on the normal modes of social interaction, such as conversation, critique, negotiation, and persuasion (Fig. 2.11). Two aspects of such an interactive environment are particularly well catered for in computer-based systems. The first aspect relates to the ability of computer-driven communication networks to link together electronically based resources located anywhere on Earth or in space. Technical advances in the communication industry have greatly enhanced the ability of individuals to gain access to remotely distributed information sources, and to interact with each other over vast distances. In fact, connectivity rather than geographical distance has become the principal determinant of communication.

In this respect, the notion of presence is being decisively redefined in an information society. In recent years we have seen the gradual acceptance of a new concept of presence that does not require the physical relocation of persons. Major sporting events and entertainment shows are more conveniently viewed on television from the home. Typically, in the case of sporting events, the quality of the televised presentation of the competition is greatly improved by technical enhancements such as instant replays and close-up shots of particularly skillful maneuvers, explanations and analyses by informed commentators, and short profile films of the best competitors.

Electronic mail, Internet chat groups, telephone and video conferencing facilities, and facsimile (FAX) transmissions, have reduced the need for face-to-face meetings. Commercial companies are gradually being forced to reassess the need for a centralized workplace. Why pay the considerable overhead

costs associated with maintaining office space for employees, if the employees could equally well perform their work at home? Computer-based messaging services and global connectivity have already reached a level of reliability and convenience that is more than adequate for business communications.

The second aspect is interwoven with the first by a major advance in the telecommunication industry that occurred some 20 years ago and allowed all types of data to be converted into digital form. Through the use of digital switching facilities modern communication networks are able to transmit telephone conversations and graphical images in the same way as data streams have been sent from one computer to another over the past 40 years.

As a direct result of these advances in communication systems the convenient and timely interaction of all of the members of a widely dispersed problem-solving team is technically assured. It is now incumbent on software developers to produce computer-based decision-support systems that can fully support collaborative teamwork, which is neither geographically nor operationally limited. Such systems will integrate not only computer-based information resources and software agents, but also multiple human agents (i.e., the users) who will collaborate with the computer-based resources in a near real-time, interactive, distributed environment.

2.8.5 The Reasoning Element

Reasoning is central to any decision-making activity. It is the ability to draw deductions and inferences from information within a problem-solving context. The ability of the problem solver to reason effectively depends as much on the availability of information, as it does on an appropriately high level form of object representation (Fig. 2.12). Decision-makers typically define complex problems in terms of issues that are known to impact the desired outcome. The relative importance of these issues and their relationships to each other change dynamically during the decision-making process. So also do the boundaries of the problem space and the goals and objectives of the desired outcome. In other words, the solution of complex problems is an altogether dynamic process in which both the rules that govern the process and the required properties of the end-result are subject to continuous review, refinement and amendment (Reitman, 1964; Reitman, 1965; Rittel and Webber, 1984).

As discussed previously, the complexity of a problem is normally not due to a high degree of difficulty in any one area but the multiple relationships that exist among the many issues that impact the desired outcome. Since a decision in one area will tend to influence several other areas there is a critical need for concurrency. However, the reasoning capabilities of the human problem solver are sequential in nature. Accordingly, decision-makers find it exceedingly difficult to consider more than three or four issues at any one time. In an attempt to deal with the concurrency requirement several strategies are

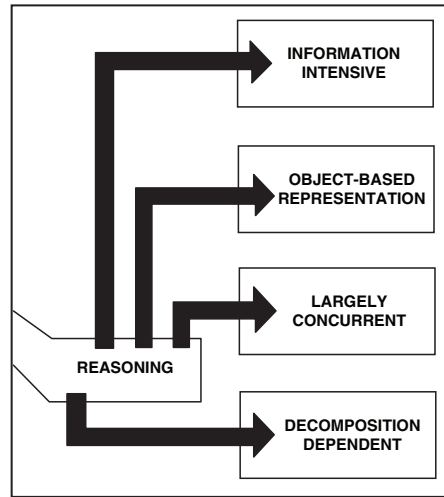


Fig. 2.12. The reasoning element

commonly employed to reduce the complexity of the reasoning process to a manageable level.³

- *Constraint identification.* By sifting through the available information the problem-solver hopes to find overriding restrictions and limitations that will eliminate knowledge areas from immediate consideration.
- *Decision factor weighting.* By comparing and evaluating important problem issues in logical groupings, relative to a set of predetermined solution objectives, the decision-maker hopes to identify a smaller number of issues or factors that appear to have greater impact on the final solution. Again, the strategy is to reduce the size of the information base by early elimination of apparently less important considerations.
- *Solution conceptualization.* By adopting early in the decision-making process a conceptual solution, the problem-solver is able to pursue a selective evaluation of the available information. Typically, the problem-solver proceeds to subdivide the decision factors into two groups, those that are compatible with the conceptual solution and those that are in conflict. By a process of trial and error, often at a superficial level, the problem-solver develops, adapts, modifies, re-conceives, rejects and, often, forces the preconceived concept into a final solution.

In complex problem situations reasoning proceeds in an iterative fashion through a cycle of analysis, synthesis and evaluation (Fig. 2.13). During

³ Reasoning is a logical process that proceeds in a step-by-step manner. In this respect reasoning is quite different from intuition, which allows humans to spontaneously come to conclusions that are neither consciously formulated nor explainable at the time of their first appearance

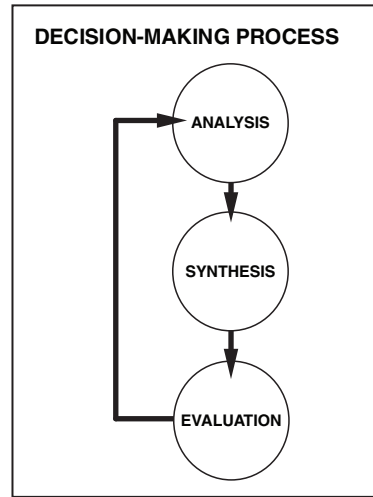


Fig. 2.13. Reasoning methodology

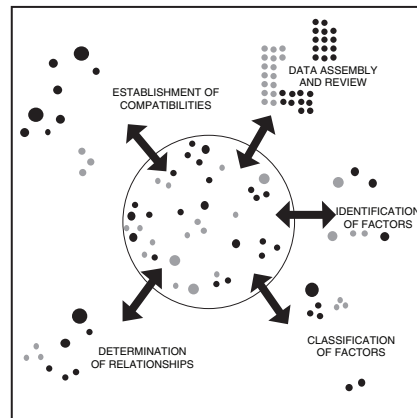


Fig. 2.14. Analysis stage of reasoning

the analysis stage (Fig. 2.14) the problem-solver interprets and categorizes information to establish the relative importance of issues and to identify compatibilities and incompatibilities among the factors that drive these issues.

During synthesis (Fig. 2.15) solution boundaries and objectives are continuously reexamined as the decision-maker develops narrow solutions to sub-problems and combines these narrow solutions into broader solutions. Initially, these solution attempts are nothing more than trial balloons. Or, stated in more technical terms, explorations based on the development of the relationships among the principal issues and compatible factors identified during the analysis stage. Later, as the problem-solving activity progresses,

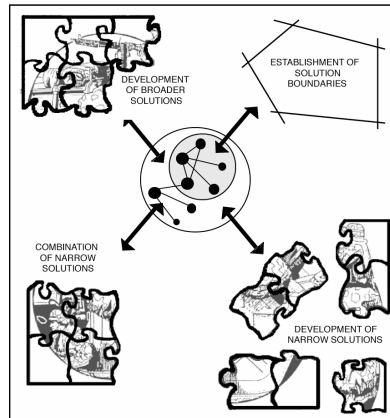


Fig. 2.15. Synthesis stage of reasoning

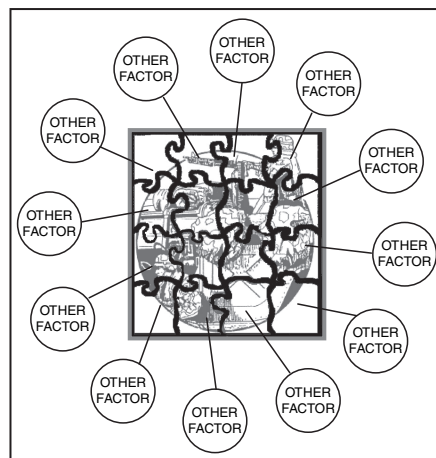


Fig. 2.16. Evaluation stage of reasoning

firmer conceptual solution strategies with broader implications emerge. However, even during later cycles the solution strategies tend to be based on a limited number of issues or factors.

During the evaluation stage (Fig. 2.16) the decision-makers are forced to test the current solution strategy with all of the known problem issues, some of which may have been considered only superficially or not at all during the formulation of the current solution proposal. This may require the current solution concepts to be modified, extended or altogether replaced. Typically, several solution strategies are possible and none are completely satisfactory. Archea (1987), in his description of the architectural design activity refers to this activity as "...puzzle-making", suggesting by implication that the

decision-maker utilizes the reasoning cycle more as a method for exploring the problem space than as a decision-making tool for forcing an early solution.

2.8.6 The Intuition Element

Schön (1983, 1988), has written extensively about the intuitive aspects of decision-making. Although he focused primarily on engineering design as an application area, his views provide valuable insight into the solution of complex problems in general. Design has all of the common characteristics of complex problem situations, and some additional ones such as the desire for solution uniqueness in architectural design, that make it a prime candidate for computer-based assistance (Pohl et al., 1994).

In Schön's (1988) view designers enter into "... design worlds" in which they find the objects, rules and prototype knowledge that they apply to the design problem under consideration. The implication is that the designer continuously moves in and out of design worlds that are triggered by internal and external stimuli. While the reasoning process employed by the designer in any particular design world is typically sequential and explicitly logical, the transitions from state to state are governed by deeper physiological and psychological causes. Some of these causes can be explained in terms of associations that the designer perceives between an aspect or element of the current state of the design solution and prototype knowledge that the designer has accumulated through experience. Others may be related to emotional states or environmental stimuli, or interactions of both (Fig. 2.17).

For example, applying Schön's view to the broader area of complex problem solving, a particular aspect of a problem situation may lead to associations in the decision-maker's mind that are logically unrelated to the problem under

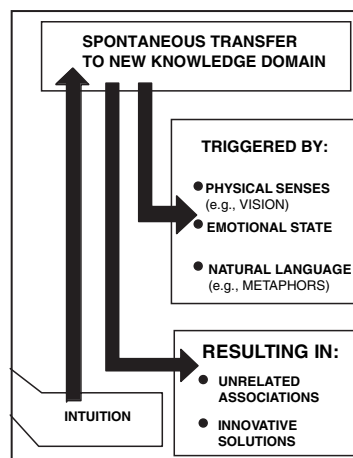


Fig. 2.17. The *intuition* element

consideration. However, when the decision-maker pursues and further develops these associations they sometimes lead to unexpected solutions. Typically, the validity of these solutions becomes apparent only after the fact and not while they are being developed. In popular terms we often refer to these solutions as creative leaps and label the author as a brilliant strategist. What we easily forget is that many of these intuitions remain unrelated associations and do not lead to any worthwhile result. Nevertheless, the intuitive aspect of decision-making is most important. Even if only a very small percentage of these intuitive associations lead to a useful solution, they still constitute one of the most highly valued decision-making resources.

The reasons for this are twofold. First, the time at which the decision-maker is most willing to entertain intuitive associations normally coincides with the most difficult stage in the problem solving process. Typically, it occurs when an impasse has been reached and no acceptable solution strategy can be found. Under these circumstances intuition may be the only remaining course of action open to the decision-maker. The second reason is particularly relevant if there is a strong competitive element present in the problem situation, for example, during a chess game or during the execution of military operations. Under these circumstances, strategies and solutions triggered by intuitive associations will inevitably introduce an element of surprise that is likely to disadvantage the adversary.

The importance of the intuition element itself in decision-making would be sufficient reason to insist on the inclusion of the human decision-maker as an active participant in any computer-based decision system. In designing and developing such systems in our Center over the past two decades we have come to appreciate the importance of the human-computer partnership concept, as opposed to automation. Whereas in some of our early systems (e.g., ICADS (Pohl et al., 1998) and AEDOT (Pohl et al., 1992)) we included agents that automatically resolve conflicts, today we are increasingly moving away from automatic conflict resolution to conflict detection and explanation. We believe that even apparently mundane conflict situations should be brought to the attention of the human agent. Although the latter may do nothing more than agree with the solution proposed by the computer-based agents, he or she has the opportunity to bring other knowledge to bear on the situation and thereby influence the final determination.

2.9 The Human-Computer Partnership

To look upon decision-support systems as partnerships between users and computers, in preference to automation, appears to be a sound approach for at least two reasons. First, the ability of the computer-based components to interact with the user overcomes many of the difficulties, such as representation and the validation of knowledge, that continue to plague the field of machine learning (Forsyth, 1989; Thornton, 1992; Johnson-Laird, 1993).

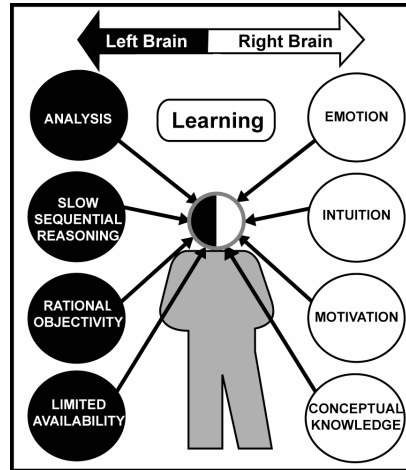


Fig. 2.18. Human abilities and limitations

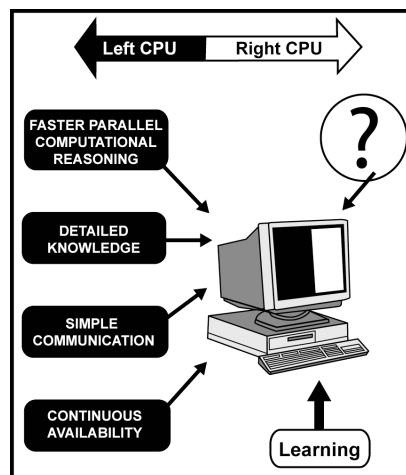


Fig. 2.19. Computer abilities and limitations

Second, human and computer capabilities are in many respects complementary (Figs. 2.18 and 2.19). Human capabilities are particularly strong in areas such as communication, symbolic reasoning, conceptualization, learning, and intuition (Fig. 2.18). We are able to store and adapt experience and quickly grasp the overall picture of even fairly chaotic situations. Our ability to match patterns is applicable not only to visual stimuli but also to abstract concepts and intuitive notions. However, as powerful as these capabilities may appear to be they are nevertheless flawed by those innately human tendencies that

were discussed at the beginning of this chapter under the rubric of human limitations and weaknesses.

Decision-making based on analysis requires not only a great deal of rather tedious and time consuming work, but also the unbiased and unemotional evaluation of past experience and possible future outcomes. This is indeed a tall order since emotions are a driving force in virtually all human activities. Pomerol and Adam (2008), in Sect. 2.5 of Chap. 1, discuss in some detail the critical role that emotions play in decision-making. Due to the influence of emotions, coupled with our aversion to hard work, our experience-based cognitive facilities, and our desire for instant gratification, we easily fall prey to over-reliance on intuition. In contrast to the painstaking sequential logic that must be applied in an analytical process, intuition is almost entirely subconscious and produces almost immediate results rather effortlessly. However, intuition can easily lead to false conclusions (Bonabeau, 2003). Unfortunately, we often see patterns that do not exist in reality. The greater the complexity of a situation the more likely that our intuitive assessments may be incorrect, since we are so easily misled by our almost uncontrollable desire to maintain the status quo. As a result we tend to judge new circumstances in the context of past conditions, particularly when these judgments are made intuitively. Examples of such human failings have been provided by Hammond et al. (2002) and include the following types of deceptions:

The anchoring deception. We tend to use the first information received as a reference point for comparing all subsequent information. For example, posing a question in the form of “Is the distance between Chicago and New York greater than 2,000 kilometers?” is likely to produce an answer that is biased toward a distance of 2,000 km. Clearly, intelligent software that operates in partnership with its human user would not only serve as a convenient source of factual data, but also assist the user in viewing a problem from several perspectives by providing access to information sources and reference points. The intelligence of the software is then related to its ability to automatically detect the need for such assistance, rather than the assistance functions themselves.

The status quo deception. We feel most comfortable with current conditions and practices unless compelled to change by threatening events. The tendency to delay decisions is fundamental to common expressions such as “... let’s wait until things settle down” or “... I’ll rethink this later”. In this case, an intelligent decision-support system should be able to not only alert the user to threatening events, but also assist in tracing the historical path that has led to the status quo conditions and undertake a detailed analysis of alternative courses of action.

The confirming evidence deception. We often favor a particular course of action without adequate justification. In such situations we tend to rely heavily on ad hoc confirming evidence, instead of undertaking the necessary analysis. The factual analysis and evaluation capabilities of an intelligent decision-support system are particularly useful as a counterpoint to the bias and relative laziness of the human decision-maker.

The sunken costs deception. We have difficulty admitting to past errors in judgment and may stubbornly insist on the perpetuation of a decision path that is fundamentally flawed. Under these circumstances the evidence generated through the analysis, evaluation and consequence determination capabilities of an intelligent software partner may be the only effective method of exposing the deception.

The forecasting deception. We can easily become overconfident without corroborating experience or too prudent by relying on a worst case scenario. In this case intelligent decision-support software can assist through the very nature of its disciplined approach to problem solving. It is in this area that the human decision-maker is particularly vulnerable because the absence of experience with future events will force reliance on past experience that may only partially, or not at all, apply.

The framing trap. A poorly framed problem can easily bias our decisions since we tend to be unduly influenced by risks associated with potential losses, even if these risks are remote. The absence of emotions in intelligent decision-support systems for the foreseeable future should be helpful in this regard. They allow the decision-maker to consider a problem from several different reference points. However, care must be taken by the designers of the software to ensure that the results of the computer-based analysis are presented to the human user in a neutral manner, so that potential gains and losses are more likely to be considered on an equal basis.

As these examples indicate, intelligent software systems can be particularly helpful in complementing human capabilities by providing a tireless, fast and emotionless problem analysis and solution evaluation capability. Large volumes of information and multi-faceted decision contexts tend to easily overwhelm human decision-makers. When such an overload occurs we tend to switch from an analysis mode to an intuitive mode in which we have to rely almost entirely on our ability to develop situation awareness through abstraction and conceptualization. While this is perhaps our greatest strength it is also potentially our greatest weakness, because at this intuitive meta-level we become increasingly vulnerable to emotional influences.

The capabilities of the computer are strongest in the areas of parallelism, speed and accuracy (Fig. 2.19). Whereas the human being tends to limit the amount of detailed knowledge by continuously abstracting information to a higher level of understanding, the computer excels in its almost unlimited capacity for storing data. While the human being is prone to impatience, loss of concentration and panic under overwhelming or threatening circumstances, the computer is totally oblivious to such emotional influences. The most effective implementation of these complementing human and machine capabilities is in a tightly coupled partnership environment that encourages and supports seamless interaction.

In conclusion, it is certainly appropriate to revisit Kurzweil's hypothesis mentioned early on in this chapter that computing devices are a natural ingredient and necessary requirement for accelerating the intellectual evolution of

human beings. For this hypothesis to hold true, intelligent software systems would need to be able to compensate for at least three recognized limitations of the human cognitive system; namely: poor performance in the absence of experience; emotional interference with logical processes; and, a distinct lack of motivation for proactive endeavors. Existing computer capabilities that show promise in this regard include: information storage in context building ontological structures; symbolic reasoning; pattern matching; computation speed; virtually unlimited parallelism; low level learning; analogy detection; and, tireless unemotional task performance. However, several additional capabilities would appear to be required. These include at least the following three capabilities that are likely to challenge the developers of AI-based decision-support systems for the next several decades.

First, there is a need for automatic context generation to form the basis of higher level learning capabilities. While much headway has been made during the past two decades in the representation of context using rich information structures such as ontologies, these are still largely static, predefined virtual models of real world knowledge domains. What are needed are methods for extending and merging ontologies dynamically during software execution (i.e., extensible information representation models). Current industry research efforts in this area such as the WebFountainTM project (IBM, 2002; Chase, 2002), are interesting but have not yet led to breakthrough advances in AI.

Second, there is a need for an interface to seamlessly link intelligent software with the human nervous system. Currently available interface devices and virtual reality capabilities are still very primitive. While some very promising advances have been made in the bio-engineering field in recent years with implanted sensors for artificial limbs, artificial hearing devices, and the invasive monitoring of bodily functions, much more progress needs to be made before we can contemplate the feasibility of a practical implanted user-interface.

Third, there is a need for new methodologies that will allow the development of software that can support the creation of knowledge through analogous reasoning or other as yet unknown processes. The notion of a conceptual database search, discussed previously in the context of the *information element* of the decision-making process (Fig. 2.6), is an example of such a capability. The realization of this kind of AI capability is likely to be the furthest goal to reach.

References

- Archea J., Puzzle-Making: what architects do when no one is looking, in Kalay (ed.) Principles of Computer-Aided Design: Computability of Design, Wiley, New York, NY (1987)
- Bohm D., Wholeness and the Implicate Order, Ark Paperbacks, London, England (1983)

- Bonabeau E., Dont Trust Your Gut, *Harvard Business Review*, May (2003)
- Brockman J., The singularity: a talk with ray kurzweil, *Edge* 99 – March 25, 2002, Edge Foundation (2002) (<http://www.edge.org>)
- Brooks R., *Flesh and Machines: How Robots Will Change Us*, Pantheon Books, Random House, New York, NY (2002)
- Chase V., Made to Order: IBM Makes Sense of Unstructured Data, Think Research Web Site, http://www.research.ibm.com/thinkresearch/pages/2002/2002/0308_unstructured.shtml (2002)
- DOE, Research Abstracts from the DOE Human Genome Program: Contractor-Grantee Workshop VIII, US Dept. of Energy, February 27 to March 2, Santa Fe, New Mexico (2000) (<http://www.ornl.gov/hgmis/publicat/00santa>)
- Finn R., Neural prosthetics come of age as research continues, *The Scientist*, 11/19 (1997) 13–16
- Fischer G. and K. Nakakoji, Making design objects relevant to the task at hand, *Proc. AAAI-91 Ninth National Conference on Artificial Intelligence*, MIT Press, Cambridge, MA (1991) 67–73
- Forsyth R., *Machine Learning: Principles and Techniques*, Chapman & Hall, New York, NY (1989)
- Gero J., M. Maher and W. Zhang, Chunking Structural Design Knowledge as Prototypes, Working Paper, The Architectural Computing Unit, Department of Architectural and Design Science, University of Sydney, Sydney, Australia (1988)
- Hammond J., R. Keeney and H. Raiffa, The Hidden Traps in Decision Making, *Harvard Business Review*, September–October (1998)
- Hübener M., D. Shoham, A. Grinvald and T. Bonhoeffer, Spatial relationships among three columnar systems in Cat Area 17, *Journal of Neuroscience*, 17 (1997) 9270–9284
- IBM, WebFountainTM, [http://www.almaden.ibm.com/webfountain/\(2002\)](http://www.almaden.ibm.com/webfountain/(2002))
- Johnson-Laird P., *Human and Machine Thinking*, Erlbaum, Hillsdale, NJ (1993)
- Kelly K., *Out of control: the new biology of machines*, Social Systems and The Economic World, Perseus Books, Reading, MA (1994) 298–302
- Kiernan V., Brains at work: researchers use new techniques to study how humans think, *Chronicle of Higher Education*, 44/20 (1998) A16–17
- Klein, G. *Sources of Power: How People Make Decisions*, MIT Press, Cambridge, MA (1998)
- Kurzweil R., *The Age of the Spiritual Machines*, Viking, Penguin Group, New York, NY (1999)
- Lakoff G. and M. Johnson, *Metaphors We Live By*, University of Chicago Press, Chicago, IL (1980)
- Minsky M., Logical vs. analogical or symbolic vs. connectionist or neat vs. scruffy, in *Artificial Intelligence at MIT*, Expanding Frontiers, Patrick H. Winston (ed.), Vol. 1, MIT Press, Cambridge, MA (1990) 218–243
- Pohl J., Some notions of complex adaptive systems and their relationship to our world, Intersymp-99, in *Advances in Collaborative Decision-Support Systems for Design, Planning and Execution*, Baden-Baden, Germany, August 2–7 (1999)
- Pohl J., Some Thoughts on Human Nature: A Discussion of Human Strengths and Weaknesses, InterSymp-2002, Focus Symposium on Collaborative Decision-Support Systems, Baden-Baden, Germany, July 29 to August 30, Preconference Proceedings (June 2002)

- Pohl J., Intelligent software systems in historical context, in Jain L. and G. Wren (eds.), *Decision Support Systems in Agent-Based Intelligent Environments*, Knowledge-Based Intelligent Engineering Systems Series, Advanced Knowledge International (AKI), Sydney, Australia (2005)
- Pohl J., Elements of Human Decision-Making, *Proceedings InterSymp-2006, Advances in Computer Cybernetics*, Vol. XIII, The International Institute for Advanced Studies in Systems Research and Cybernetics, Baden-Baden, Germany, August 7–12 (2006)
- Pohl J. and L. Myers, A distributed cooperative model for architectural design, in G. Carrara and Y. Kalay (eds.), *Knowledge-Based Computer-Aided Architectural Design*, Elsevier, Amsterdam (1994) 205–242
- Pohl J., A. Chapman, L. Chirica, R. Howell and L. Myers, Implementation Strategies for a Prototype ICADS Working Model, Technical Report, CADRU-02–88, CAD Research Center, Design and Construction Institute, College of Architecture and Environmental Design, Cal Poly, San Luis Obispo, CA (1988)
- Pohl J., J. La Porta, K. Pohl and J. Snyder, AEDOT Prototype (1.1): An Implementation of the ICADS Model, Technical Report, CADRU-07–92, CAD Research Center, Design and Construction Institute, College of Architecture and Environmental Design, Cal Poly, San Luis Obispo, CA (1992)
- Pohl J., L. Myers and A. Chapman, Thoughts on the Evolution of Computer-Assisted Design, Technical Report, CADRU-09–94, CAD Research Center, Design and Construction Institute, College of Architecture and Environmental Design, Cal Poly, San Luis Obispo, CA, September (1994)
- Pohl J., A. Chapman, K. Pohl, J. Primrose and A. Wozniak, Decision-Support Systems: Notions, Prototypes, and In-Use Applications, Collaborative Agent Design Research Center, Technical Report CADRU-11–97, Cal Poly, San Luis Obispo, CA 93407 (1997) 10–11
- Pomeroy J.C. and F. Adam, Understanding human decision making – a fundamental step towards effective intelligent design support, in G. Phillips – Wren, L.C. Jain (eds.), *Intelligent Decision Making: An AI-Based Approach*, Chapter I (2008)
- Powledge T., Unlocking the secrets of the brain, *Bioscience*, 47 (two-part article) July (1997) 330–334 and August (1997) 403–409
- Reitman W., Heuristic decision procedures, open constraints, and the ill-defined problems, in Shelley and Bryan (eds.) *Human Judgments and Optimality*, Wiley, New York, NY (1964) 282–315
- Reitman W., *Cognition and Thought*, Wiley, New York, NY (1965)
- Rittel H. and M. Webber, Planning problems are wicked problems, in Cross (ed.) *Developments in Design Methodology*, Wiley, New York, NY (1984) 135–144
- Rosenman M. and J. Gero, Creativity in design using design prototypes, in Gero and Maher (eds.) *Modeling Creativity and Knowledge-Based Creative Design*, Erlbaum, Hillsdale, NJ (1993) 111–138
- Schön D., *The Reflective Practitioner: How Professionals Think in Action*, Basic Books, London (1983)
- Schön D., Designing: rules, types and worlds, *Design Studies*, 9(3), (1988) 181–190
- Senge P., *Transforming the Practice of Management*, Human Resource Development Quarterly, Jossey-Bass, San Francisco, CA (1993)
- Simon H., *The Sciences of the Artificial*, MIT Press, Cambridge, MA (1981)
- Thornton C., *Techniques in Computational Learning*, Chapman & Hall, Computing Series, London, England (1992)

Intelligent Decision Making: An AI-Based Approach

Phillips-Wren, G.; Ichalkaranje, N. (Eds.)

2008, XIV, 410 p., Hardcover

ISBN: 978-3-540-76828-9