

Decision Support Systems

Marek J. Druzdzel and Roger R. Flynn

Decision Systems Laboratory
School of Information Sciences
and Intelligent Systems Program

University of Pittsburgh
Pittsburgh, PA 15260

`{marek,flynn}@sis.pitt.edu`

`http://www.sis.pitt.edu/~dsl`

To appear in

Encyclopedia of Library and Information Science,
Second Edition,

Allen Kent (ed.), New York: Marcel Dekker, Inc., 2002

Contents

Introduction	3
Decisions and Decision Modeling	4
Types of Decisions	4
Human Judgment and Decision Making	4
Modeling Decisions	5
Components of Decision Models	5
Decision Support Systems	6
Normative Systems	7
Normative and Descriptive Approaches	7
Decision-Analytic Decision Support Systems	8
Equation-Based and Mixed Systems	10
User Interfaces to Decision Support Systems	11
Support for Model Construction and Model Analysis	11
Support for Reasoning about the Problem Structure in Addition to Numerical Calculations	11
Support for Both Choice and Optimization of Decision Variables	12
Graphical Interface	12
Summary	12

Introduction

Making decisions concerning complex systems (e.g., the management of organizational operations, industrial processes, or investment portfolios; the command and control of military units; or the control of nuclear power plants) often strains our cognitive capabilities. Even though individual interactions among a system's variables may be well understood, predicting how the system will react to an external manipulation such as a policy decision is often difficult. What will be, for example, the effect of introducing the third shift on a factory floor? One might expect that this will increase the plant's output by roughly 50 percent. Factors such as additional wages, machine wear, maintenance breaks, raw material usage, supply logistics, and future demand need also be considered, however, as they all will impact the total financial outcome of this decision. Many variables are involved in complex and often subtle interdependencies and predicting the total outcome may be daunting.

There is a substantial amount of empirical evidence that human intuitive judgment and decision making can be far from optimal, and it deteriorates even further with complexity and stress. Because in many situations the quality of decisions is important, aiding the deficiencies of human judgment and decision making has been a major focus of science throughout history. Disciplines such as statistics, economics, and operations research developed various methods for making rational choices. More recently, these methods, often enhanced by a variety of techniques originating from information science, cognitive psychology, and artificial intelligence, have been implemented in the form of computer programs, either as stand-alone tools or as integrated computing environments for complex decision making. Such environments are often given the common name of *decision support systems* (DSSs). The concept of DSS is extremely broad, and its definitions vary, depending on the author's point of view. To avoid exclusion of any of the existing types of DSSs, we will define them roughly as interactive computer-based systems that aid users in judgment and choice activities. Another name sometimes used as a synonym for DSS is *knowledge-based systems*, which refers to their attempt to formalize domain knowledge so that it is amenable to mechanized reasoning.

Decision support systems are gaining an increased popularity in various domains, including business, engineering, the military, and medicine. They are especially valuable in situations in which the amount of available information is prohibitive for the intuition of an unaided human decision maker and in which precision and optimality are of importance. Decision support systems can aid human cognitive deficiencies by integrating various sources of information, providing intelligent access to relevant knowledge, and aiding the process of structuring decisions. They can also support choice among well-defined alternatives and build on formal approaches, such as the methods of engineering economics, operations research, statistics, and decision theory. They can also employ artificial intelligence methods to address heuristically problems that are intractable by formal techniques. Proper application of decision-making tools increases productivity, efficiency, and effectiveness and gives many businesses a comparative advantage over their competitors, allowing them to make optimal choices for technological processes and their parameters, planning business operations, logistics, or investments.

While it is difficult to overestimate the importance of various computer-based tools that are relevant to decision making (e.g., databases, planning software, and spreadsheets), this article focuses primarily on the core of a DSS, the part that directly supports modeling decision problems and identifies best alternatives. We will briefly discuss the characteristics of decision problems and how decision making can be supported by computer programs. We then cover various components of DSSs and the role that they play in decision support. We will also introduce an emergent class of *normative systems* (i.e., DSSs based on sound theoretical principles), and in particular, decision-analytic DSSs. Finally, we will review issues related to user interfaces to DSSs and stress the importance of user interfaces to the ultimate quality of decisions aided by computer programs.

Decisions and Decision Modeling

Types of Decisions

A simple view of decision making is that it is a problem of choice among several alternatives. A somewhat more sophisticated view includes the process of constructing the alternatives (i.e., given a problem statement, developing a list of choice options). A complete picture includes a search for opportunities for decisions (i.e., discovering that there is a decision to be made). A manager of a company may face a choice in which the options are clear (e.g., the choice of a supplier from among all existing suppliers). She may also face a well-defined problem for which she designs creative decision options (e.g., how to market a new product so that the profits are maximized). Finally, she may work in a less reactive fashion and view decision problems as opportunities that have to be discovered by studying the operations of her company and its surrounding environment (e.g., how can she make the production process more efficient). There is much anecdotal and some empirical evidence that structuring decision problems and identifying creative decision alternatives determine the ultimate quality of decisions. Decision support systems aim mainly at this broadest type of decision making, and in addition to supporting choice, they aid in modeling and analyzing systems (such as complex organizations), identifying decision opportunities, and structuring decision problems.

Human Judgment and Decision Making

Theoretical studies on rational decision making, notably that in the context of probability theory and decision theory, have been accompanied by empirical research on whether human behavior complies with the theory. It has been rather convincingly demonstrated in numerous empirical studies that human judgment and decision making is based on intuitive strategies as opposed to theoretically sound reasoning rules. These intuitive strategies, referred to as *judgmental heuristics* in the context of decision making, help us in reducing the cognitive load, but alas at the expense of optimal decision making. Effectively, our unaided judgment and choice exhibit systematic violations of probability axioms (referred to as *biases*). Formal discussion of the most important research results along with experimental data can be found in an anthology edited by Kahneman, Slovic, and Tversky [16]. Dawes [2] provides an accessible introduction to what is known about people's decision-making performance.

One might hope that people who have achieved expertise in a domain will not be subject to judgmental biases and will approach optimality in decision making. While empirical evidence shows that experts indeed are more accurate than novices within their area of expertise, it also shows that they also are liable to the same judgmental biases as novices and demonstrate apparent errors and inconsistencies in their judgment. Professionals such as practicing physicians use essentially the same judgmental heuristics and are prone to the same biases, although the degree of departure from the normatively prescribed judgment seems to decrease with experience. In addition to laboratory evidence, there are several studies of expert performance in realistic settings, showing that it is inferior even to simple linear models (an informal review of the available evidence and pointers to literature can be found in the book by Dawes [2]). For example, predictions of future violent behavior of psychiatric patients made by a panel of psychiatrists who had access to patient records and interviewed the patients were found to be inferior to a simple model that included only the past incidence of violent behavior. Predictions of marriage counselors concerning marital happiness were shown to be inferior to a simple model that just subtracted the rate of fighting from the rate of sexual intercourse (again, the marriage counselors had access to all data, including interviews with the couples). Studies yielding similar results have been conducted with bank loan officers, physicians, university admission committees, and so on.

Modeling Decisions

The superiority of even simple linear models over human intuitive judgment suggests that one way to improve the quality of decisions is to decompose a decision problem into simpler components that are well defined and well understood. Studying a complex system built out of such components can be subsequently aided by a formal, theoretically sound technique. The process of decomposing and formalizing a problem is often called modeling. Modeling amounts to finding an abstract representation of a real-world system that simplifies and assumes as much as possible about the system, and while retaining the system's essential relationships, omits unnecessary detail. Building a model of a decision problem, as opposed to reasoning about a problem in a holistic way, allows for applying scientific knowledge that can be transferred across problems and often across domains. It allows for analyzing, explaining, and arguing about a decision problem.

The desire to improve human decision making provided motivation for the development of a variety of modeling tools in disciplines of economics, operations research, decision theory, decision analysis, and statistics. In each of these modeling tools, knowledge about a system is represented by means of algebraic, logical, or statistical variables. Interactions among these variables are expressed by equations or logical rules, possibly enhanced with an explicit representation of uncertainty. When the functional form of an interaction is unknown, it is sometimes described in purely probabilistic terms; for example, by a conditional probability distribution. Once a model has been formulated, a variety of mathematical methods can be used to analyze it. Decision making under certainty has been addressed by economic and operations research methods, such as cash flow analysis, break-even analysis, scenario analysis, mathematical programming, inventory techniques, and a variety of optimization algorithms for scheduling and logistics. Decision making under uncertainty enhances the above methods with statistical approaches, such as reliability analysis, simulation, and statistical decision making. Most of these methods have made it into college curricula and can be found in management textbooks. Due to space constraints, we will not discuss their details further.

Components of Decision Models

While mathematically a model consists of variables and a specification of interactions among them, from the point of view of decision making a model and its variables represent the following three components: a measure of preferences over decision objectives, available decision options, and a measure of uncertainty over variables influencing the decision and the outcomes.

Preference is widely viewed as the most important concept in decision making. Outcomes of a decision process are not all equally attractive and it is crucial for a decision maker to examine these outcomes in terms of their desirability. Preferences can be ordinal (e.g., more income is preferred to less income), but it is convenient and often necessary to represent them as numerical quantities, especially if the outcome of the decision process consists of multiple attributes that need to be compared on a common scale. Even when they consist of just a single attribute but the choice is made under uncertainty, expressing preferences numerically allows for trade-offs between desirability and risk.

The second component of decision problems is available decision options. Often these options can be enumerated (e.g., a list of possible suppliers), but sometimes they are continuous values of specified policy variables (e.g., the amount of raw material to be kept in stock). Listing the available decision options is an important element of model structuring.

The third element of decision models is uncertainty. Uncertainty is one of the most inherent and most prevalent properties of knowledge, originating from incompleteness of information, imprecision,

and model approximations made for the sake of simplicity. It would not be an exaggeration to state that real-world decisions not involving uncertainty either do not exist or belong to a truly limited class.¹

Decision making under uncertainty can be viewed as a deliberation: determining what action should be taken that will maximize the expected gain. Due to uncertainty there is no guarantee that the result of the action will be the one intended, and the best one can hope for is to maximize the chance of a desirable outcome. The process rests on the assumption that a good decision is one that results from a good decision-making process that considers all important factors and is explicit about decision alternatives, preferences, and uncertainty.

It is important to distinguish between good decisions and good outcomes. By a stroke of good luck a poor decision can lead to a very good outcome. Similarly, a very good decision can be followed by a bad outcome. Supporting decisions means supporting the decision-making process so that better decisions are made. Better decisions can be expected to lead to better outcomes.

Decision Support Systems

Decision support systems are interactive, computer-based systems that aid users in judgment and choice activities. They provide data storage and retrieval but enhance the traditional information access and retrieval functions with support for model building and model-based reasoning. They support framing, modeling, and problem solving.

Typical application areas of DSSs are management and planning in business, health care, the military, and any area in which management will encounter complex decision situations. Decision support systems are typically used for strategic and tactical decisions faced by upper-level management—decisions with a reasonably low frequency and high potential consequences—in which the time taken for thinking through and modeling the problem pays off generously in the long run.

There are three fundamental components of DSSs [22].

- *Database management system (DBMS)*. A DBMS serves as a data bank for the DSS. It stores large quantities of data that are relevant to the class of problems for which the DSS has been designed and provides logical data structures (as opposed to the physical data structures) with which the users interact. A DBMS separates the users from the physical aspects of the database structure and processing. It should also be capable of informing the user of the types of data that are available and how to gain access to them.
- *Model-base management system (MBMS)*. The role of MBMS is analogous to that of a DBMS. Its primary function is providing independence between specific models that are used in a DSS from the applications that use them. The purpose of an MBMS is to transform data from the DBMS into information that is useful in decision making. Since many problems that the user of a DSS will cope with may be unstructured, the MBMS should also be capable of assisting the user in model building.
- *Dialog generation and management system (DGMS)*. The main product of an interaction with a DSS is insight. As their users are often managers who are not computer-trained, DSSs need to be equipped with intuitive and easy-to-use interfaces. These interfaces aid in model

¹As Benjamin Franklin expressed it in 1789 in a letter to his friend M. Le Roy, “in this world nothing can be said to be certain, except death and taxes” (*The Complete Works of Benjamin Franklin*, John Bigelow (ed), New York and London: G.P. Putnam’s Sons, 1887, Vol. 10, page 170).

building, but also in interaction with the model, such as gaining insight and recommendations from it. The primary responsibility of a DGMS is to enhance the ability of the system user to utilize and benefit from the DSS. In the remainder of this article, we will use the broader term *user interface* rather than DGMS.

While a variety of DSSs exists, the above three components can be found in many DSS architectures and play a prominent role in their structure. Interaction among them is illustrated in Fig. 1. Essentially, the user interacts with the DSS through the DGMS. This communicates with the DBMS

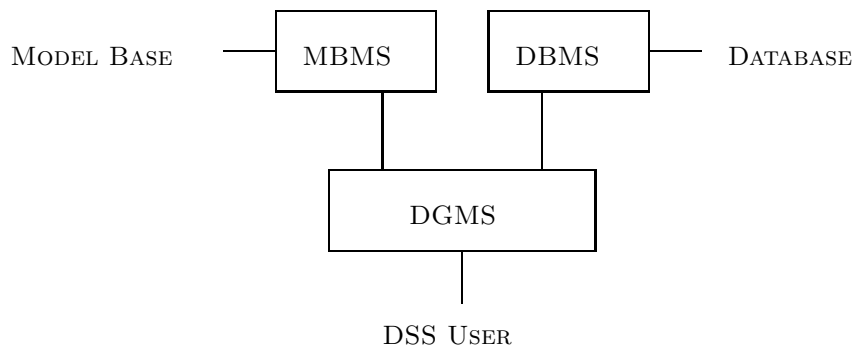


Figure 1: The architecture of a DSSs (after Sage, Ref. [22]).

and MBMS, which screen the user and the user interface from the physical details of the model base and database implementation.

Normative Systems

Normative and Descriptive Approaches

Whether or not one trusts the quality of human intuitive reasoning strategies has a profound impact on one's view of the philosophical and technical foundations of DSSs. There are two distinct approaches to supporting decision making. The first aims at building support procedures or systems that imitate human experts. The most prominent member of this class of DSSs are *expert systems*, computer programs based on rules elicited from human domain experts that imitate reasoning of a human expert in a given domain. Expert systems are often capable of supporting decision making in that domain at a level comparable to human experts. While they are flexible and often able to address complex decision problems, they are based on intuitive human reasoning and lack soundness and formal guarantees with respect to the theoretical reliability of their results. The danger of the expert system approach, increasingly appreciated by DSS builders, is that along with imitating human thinking and its efficient heuristic principles, we may also imitate its undesirable flaws [13].

The second approach is based on the assumption that the most reliable method of dealing with complex decisions is through a small set of normatively sound principles of how decisions should be made. While heuristic methods and ad hoc reasoning schemes that imitate human cognition may in many domains perform well, most decision makers will be reluctant to rely on them whenever the cost of making an error is high. To give an extreme example, few people would choose to fly airplanes built using heuristic principles over airplanes built using the laws of aerodynamics enhanced with probabilistic reliability analysis. Application of formal methods in DSSs makes these systems

philosophically distinct from those based on ad hoc heuristic artificial intelligence methods, such as rule-based systems. The goal of a DSS, according to this view, is to support unaided human intuition, just as the goal of using a calculator is to aid human's limited capacity for mental arithmetic.

Decision-Analytic Decision Support Systems

An emergent class of DSSs known as *decision-analytic DSSs* applies the principles of decision theory, probability theory, and decision analysis to their decision models. Decision theory is an axiomatic theory of decision making that is built on a small set of axioms of rational decision making. It expresses uncertainty in terms of probabilities and preferences in terms of utilities. These are combined using the operation of mathematical expectation. The attractiveness of probability theory, as a formalism for handling uncertainty in DSSs, lies in its soundness and its guarantees concerning long-term performance. Probability theory is often viewed as the gold standard for rationality in reasoning under uncertainty. Following its axioms offers protection from some elementary inconsistencies. Their violation, on the other hand, can be demonstrated to lead to sure losses [23]. Decision analysis is the art and science of applying decision theory to real-world problems. It includes a wealth of techniques for model construction, such as methods for elicitation of model structure and probability distributions that allow minimization of human bias, methods for checking the sensitivity of a model to imprecision in the data, computing the value of obtaining additional information, and presentation of results. (See, for example, Ref. [27] for a basic review of the available techniques.) These methods have been under continuous scrutiny by psychologists working in the domain of behavioral decision theory and have proven to cope reasonably well with the dangers related to human judgmental biases.

Normative systems are usually based on graphical probabilistic models, which are representations of the joint probability distribution over a model's variables in terms of directed graphs. Directed graphs, such as the one in Fig. 2, are known as Bayesian networks (BNs) or causal networks [19]. Bayesian networks offer a compact representation of joint probability distributions and are capable of practical representation of large models, consisting of tens or hundreds of variables. Bayesian networks can be easily extended with decision and value variables for modeling decision problems. The former denote variables that are under the decision maker's control and can be directly manipulated, and the latter encode users' preferences over various outcomes of the decision process. Such amended graphs are known as *influence diagrams* [15]. Both the structure and the numerical probability distributions in a BN can be elicited from a human expert and are a reflection of the expert's subjective view of a real-world system. If available, scientific knowledge about the system, both in terms of the structure and frequency data, can be easily incorporated in the model. Once a model has been created, it is optimized using formal decision-theoretic algorithms. Decision analysis is based on the empirically tested paradigm that people are able to reliably store and retrieve their personal beliefs about uncertainty and preferences for different outcomes, but are much less reliable in aggregating these fragments into a global inference. While human experts are excellent in structuring a problem, determining the components that are relevant to it and providing local estimates of probabilities and preferences, they are not reliable in combining many simple factors into an optimal decision. The role of a decision-analytic DSS is to support them in their weaknesses using the formal and theoretically sound principles of statistics.

The approach taken by decision analysis is compatible with that of DSSs. The goal of decision analysis is to provide insight into a decision. This insight, consisting of the analysis of all relevant factors, their uncertainty, and the critical nature of some assumptions, is even more important than the actual recommendation.

Decision-analytic DSSs have been successfully applied to practical systems in medicine, business,

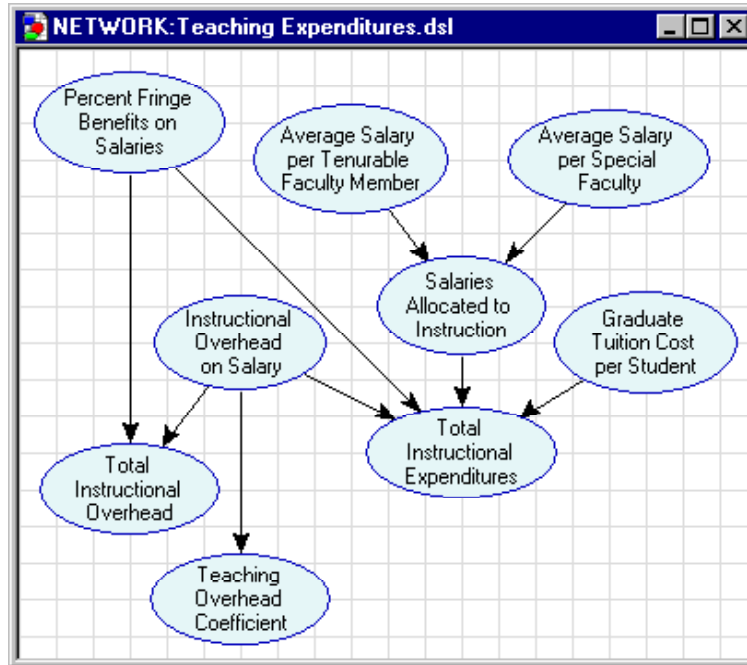


Figure 2: Example of a Bayesian network modeling teaching expenditures in university operations.

and engineering.² As these systems tend to naturally evolve into three not necessarily distinct classes, it may be interesting to compare their structure and architectural organization.

- *Systems with static domain models.* In this class of systems, a probabilistic domain is represented by a large network encoding the domain's structure and its numerical parameters. The network comprising the domain model is normally built by decision analysts and domain experts. An example might be a medical diagnostic system covering a certain class of disorders. Queries in such a system are answered by assigning values to those nodes of the network that constitute the observations for a particular case and propagating the impact of the observation through the network in order to find the probability distribution of some selected nodes of interest (e.g., nodes that represent diseases). Such a network can, on a case-by-case basis, be extended with decision nodes and value nodes to support decisions. Systems with static domain models are conceptually similar to rule-based expert systems covering an area of expertise.
- *Systems with customized decision models.* The main idea behind this approach is automatic generation of a graphical decision model on a per-case basis in an interactive effort between the DSS and the decision maker. The DSS has domain expertise in a certain area and plays the role of a decision analyst. During this interaction, the program creates a customized influence diagram, which is later used for generating advice. The main motivation for this approach is the premise that every decision is unique and needs to be looked at individually; an influence diagram needs to be tailored to individual needs [14].

²Some examples of applications are described in a special issue of *Communications of the ACM* on practical applications of decision-theoretic methods (vol. 38, no. 3, March 1995). The readers can experiment with GeNIe [7], a development system for decision-analytic DSSs developed at the Decision Systems Laboratory, University of Pittsburgh, available at <http://www2.sis.pitt.edu/~genie>.

- *Systems capable of learning a model from data.* The third class of systems employs computer-intensive statistical methods for learning models from data [1, 11, 12, 21, 26]. Whenever there are sufficient data available, the systems can literally learn a graphical model from these data. This model can be subsequently used to support decisions within the same domain.

The first two approaches are suited for slightly different applications. The customized model generation approach is an attempt to automate the most laborious part of decision making, structuring a problem, so far done with significant assistance from trained decision analysts. A session with the program that assists the decision maker in building an influence diagram is laborious. This makes the customized model generation approach particularly suitable for decision problems that are infrequent and serious enough to be treated individually. Because in the static domain model approach an existing domain model needs to be customized by the case data only, the decision-making cycle is rather short. This makes it particularly suitable for those decisions that are highly repetitive and need to be made under time constraints.

A practical system can combine the three approaches. A static domain model can be slightly customized for a case that needs individual treatment. Once completed, a customized model can be blended into the large static model. Learning systems can support both the static and the customized model approach. On the other hand, the learning process can be greatly enhanced by prior knowledge from domain experts or by a prior model.

Equation-Based and Mixed Systems

In many business and engineering problems, interactions among model variables can be described by equations which, when solved simultaneously, can be used to predict the effect of decisions on the system, and hence support decision making. One special type of simultaneous equation model is known as the structural equation model (SEM), which has been a popular method of representing systems in econometrics. An equation is structural if it describes a unique, independent causal mechanism acting in the system. Structural equations are based on expert knowledge of the system combined with theoretical considerations. Structural equations allow for a natural, modular description of a system—each equation represents its individual component, a separable and independent mechanism acting in the system—yet, the main advantage of having a structural model is, as explicated by Simon [24], that it includes causal information and aids predictions of the effects of external interventions. In addition, the causal structure of a structural equation model can be represented graphically [24], which allows for combining them with decision-analytic graphical models in practical systems [9, 20].

Structural equation models offer significant advantages for policy making. Often a decision maker confronted with a complex system needs to decide not only the values of policy variables but also which variables should be manipulated. A change in the set of policy variables has a profound impact on the structure of the problem and on how their values will propagate through the system. The user determines which variables are policy variables and which are determined within the model. A change in the SEMs or the set of policy variables can be reflected by a rapid restructuring of the model and predictions involving this new structure [25].

Our long-term project, the Environment for Strategic Planning (ESP) [6], is based on a hybrid graphical modeling tool that combines SEMs with decision-analytic principles. ESP is capable of representing both discrete and continuous variables involved in deterministic and probabilistic relationships. The powerful features of SEMs allow ESP to act as a graphical spreadsheet integrating numerical and symbolic methods and allowing the independent variables to be selected at will without having to reformulate the model each time. This provides an immense flexibility that is not afforded

by ordinary spreadsheets in evaluating alternate policy options.

User Interfaces to Decision Support Systems

While the quality and reliability of modeling tools and the internal architectures of DSSs are important, the most crucial aspect of DSSs is, by far, their user interface. Systems with user interfaces that are cumbersome or unclear or that require unusual skills are rarely useful and accepted in practice. The most important result of a session with a DSS is insight into the decision problem. In addition, when the system is based on normative principles, it can play a tutoring role; one might hope that users will learn the domain model and how to reason with it over time, and improve their own thinking.

A good user interface to DSSs should support model construction and model analysis, reasoning about the problem structure in addition to numerical calculations and both choice and optimization of decision variables. We will discuss these in the following sections.

Support for Model Construction and Model Analysis

User interface is the vehicle for both model construction (or model choice) and for investigating the results. Even if a system is based on a theoretically sound reasoning scheme, its recommendations will be as good as the model they are based on. Furthermore, even if the model is a very good approximation of reality and its recommendations are correct, they will not be followed if they are not understood. Without understanding, the users may accept or reject a system's advice for the wrong reasons and the combined decision-making performance may deteriorate even below unaided performance [17]. A good user interface should make the model on which the system's reasoning is based transparent to the user.

Modeling is rarely a one-shot process, and good models are usually refined and enhanced as their users gather practical experiences with the system recommendations. It is important to strike a careful balance between precision and modeling efforts; some parts of a model need to be very precise while others do not. A good user interface should include tools for examining the model and identifying its most sensitive parts, which can be subsequently elaborated on. Systems employed in practice will need their models refined, and a good user interface should make it easy to access, examine, and refine its models. Some pointers to work on support for building decision-analytic systems can be found in [8, 10, 18, 28].

Support for Reasoning about the Problem Structure in Addition to Numerical Calculations

While numerical calculations are important in decision support, reasoning about the problem structure is even more important. Often when the system and its model are complex it is insightful for the decision maker to realize how the system variables are interrelated. This is helpful in designing creative decision options but also in understanding how a policy decision will impact the objective.

Graphical models, such as those used in decision analysis or in equation-based and hybrid systems, are particularly suitable for reasoning about structure. Under certain assumptions, a directed graphical model can be given a causal interpretation. This is especially convenient in situations where the DSS autonomically suggests decision options; given a causal interpretation of its model,

it is capable of predicting effects of interventions. A causal graph facilitates building an effective user interface. The system can refer to causal interactions during its dialogue with the user, which is known to enhance user insight [3].

Support for Both Choice and Optimization of Decision Variables

Many DSSs have an inflexible structure in the sense that the variables that will be manipulated are determined at the model-building stage. This is not very suitable for planning of the strategic type when the object of the decision-making process is identifying both the objectives and the methods of achieving them. For example, changing policy variables in a spreadsheet-based model often requires that the entire spreadsheet be rebuilt. If there is no support for that, few users will consider it as an option. This closes the world of possibilities for flexible reframing of a decision problem in the exploratory process of searching for opportunities. Support for both choice and optimization of decision variables should be an inherent part of DSSs.

Graphical Interface

Insight into a model can be increased greatly at the user interface level by a diagram representing the interactions among its components; for example, a drawing of a graph on which a model is based, such as in Fig. 2. This graph is a qualitative, structural explanation of how information flows from the independent variables to the dependent variables of interest. As models may become very large, it is convenient to structure them into submodels, groups of variables that form a subsystem of the modeled system. Such submodels can be again shown graphically with interactions among them, increasing simplicity and clarity of the interface. Fig. 3 shows a submodel-level view of a model developed in our ESP project. Note that the graph in Fig. 2 is an expanded version of the *Teaching Expenditures* submodel in Fig. 3. The user can navigate through the hierarchy of the entire model in her quest for insight, opening and closing submodels on demand. Some pointers to work on user interfaces of decision-analytic systems can be found in [4, 5, 28].

Summary

Decision support systems are powerful tools integrating scientific methods for supporting complex decisions with techniques developed in information science, and are gaining an increased popularity in many domains. They are especially valuable in situations in which the amount of available information is prohibitive for the intuition of an unaided human decision maker and in which precision and optimality are of importance. Decision support systems aid human cognitive deficiencies by integrating various sources of information, providing intelligent access to relevant knowledge, aiding the process of structuring, and optimizing decisions.

Normative DSSs offer a theoretically correct and appealing way of handling uncertainty and preferences in decision problems. They are based on carefully studied empirical principles underlying the discipline of decision analysis and they have been successfully applied in many practical systems. We believe that they offer several attractive features that are likely to prevail in the long run as far as the technical developments are concerned.

Because DSSs do not replace humans but rather augment their limited capacity to deal with complex problems, their user interfaces are critical. The user interface determines whether a DSS

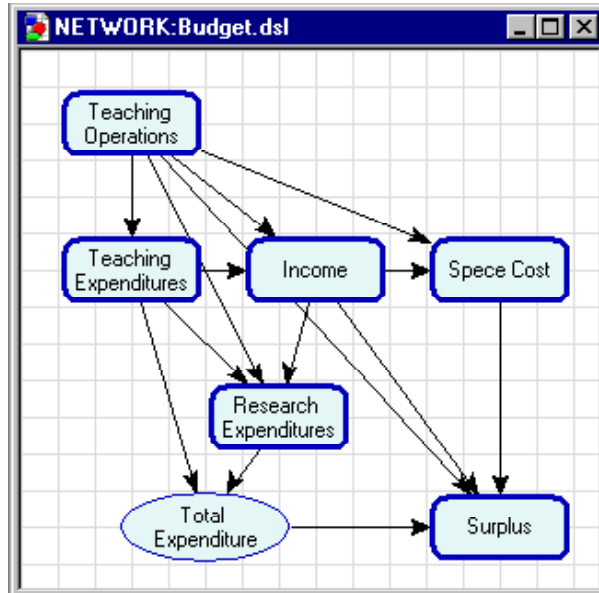


Figure 3: A submodel-level view of a decision model.

will be used at all and if so, whether the ultimate quality of decisions will be higher than that of an unaided decision maker.

Acknowledgments

Work on this article was supported by the National Science Foundation under Faculty Early Career Development (CAREER) Program, grant IRI-9624629, by the Air Force Office of Scientific Research under grants F49620-97-1-0225 and F49620-00-1-0112, and by the University of Pittsburgh Central Research Development Fund. Figures 2 and 3 are snapshots of GeNIe, a general purpose development environment for graphical decision support systems developed by the Decision Systems Laboratory, University of Pittsburgh and available at <http://www.sis.pitt.edu/~genie>. We would like to thank Ms. Nanette Yurcik for her assistance with technical editing.

References

- [1] Gregory F. Cooper and Edward Herskovits. A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9(4):309–347, 1992.
- [2] Robyn M. Dawes. *Rational Choice in an Uncertain World*. Hartcourt Brace Jovanovich, Publishers, 1988.
- [3] Marek J. Druzdzel. *Probabilistic Reasoning in Decision Support Systems: From Computation to Common Sense*. PhD thesis, Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, December 1992.

- [4] Marek J. Druzdzel. Explanation in probabilistic systems: Is it feasible? will it work? In *Proceedings of the Fifth International Workshop on Intelligent Information Systems (WIS-96)*, pages 12–24, Dęblin, Poland, 2–5 June 1996.
- [5] Marek J. Druzdzel. Five useful properties of probabilistic knowledge representations from the point of view of intelligent systems. *Fundamenta Informaticæ, Special issue on Knowledge Representation and Machine Learning*, 30(3/4):241–254, June 1997.
- [6] Marek J. Druzdzel. ESP: A mixed initiative decision-theoretic decision modeling system. In *Working Notes of the AAAI-99 Workshop on Mixed-initiative Intelligence*, pages 99–106, Orlando, FL, 18 July 1999.
- [7] Marek J. Druzdzel. SMILE: Structural Modeling, Inference, and Learning Engine and GeNie: A development environment for graphical decision-theoretic models. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, pages 902–903, Orlando, FL, July 18–22 1999.
- [8] Marek J. Druzdzel and F. Javier Díez. Criteria for combining knowledge from different sources in probabilistic models. In *Working Notes of the workshop on Fusion of Domain Knowledge with Data for Decision Support, Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 23–29, Stanford, CA, 30 June 2000.
- [9] Marek J. Druzdzel and Herbert A. Simon. Causality in Bayesian belief networks. In *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI-93)*, pages 3–11, San Francisco, CA, 1993. Morgan Kaufmann Publishers, Inc.
- [10] Marek J. Druzdzel and Linda C. van der Gaag. Building probabilistic networks: “Where do the numbers come from?” guest editors’ introduction. *IEEE Transactions on Knowledge and Data Engineering*, 12(4):481–486, July–August 2000.
- [11] Clark Glymour and Gregory F. Cooper, editors. *Computation, Causation, and Discovery*. AAAI Press, Menlo Park, CA, 1999.
- [12] David E. Heckerman, Dan Geiger, and David M. Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20(3):197–243, 1995.
- [13] Max Henrion, John S. Breese, and Eric J. Horvitz. Decision Analysis and Expert Systems. *AI Magazine*, 12(4):64–91, Winter 1991.
- [14] Samuel Holtzman. *Intelligent Decision Systems*. Addison-Wesley, Reading, MA, 1989.
- [15] Ronald A. Howard and James E. Matheson. Influence diagrams. In Ronald A. Howard and James E. Matheson, editors, *The Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA, 1984.
- [16] Daniel Kahneman, Paul Slovic, and Amos Tversky, editors. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, 1982.
- [17] Paul E. Lehner, Theresa M. Mullin, and Marvin S. Cohen. A probability analysis of the usefulness of decision aids. In M. Henrion, R.D. Shachter, L.N. Kanal, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence 5*, pages 427–436. Elsevier Science Publishers B.V. (North Holland), 1990.
- [18] Tsai-Ching Lu, Marek J. Druzdzel, and Tze-Yun Leong. Causal mechanism-based model construction. In *Proceedings of the Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 353–362, San Francisco, CA, 2000. Morgan Kaufmann Publishers, Inc.

- [19] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1988.
- [20] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2000.
- [21] Judea Pearl and Thomas S. Verma. A theory of inferred causation. In J.A. Allen, R. Fikes, and E. Sandewall, editors, *KR-91, Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, pages 441–452, Cambridge, MA, 1991. Morgan Kaufmann Publishers, Inc., San Mateo, CA.
- [22] Andrew P. Sage. *Decision Support Systems Engineering*. John Wiley & Sons, Inc., New York, 1991.
- [23] Leonard J. Savage. *The Foundations of Statistics (Second Revised Edition)*. Dover Publications, New York, NY, 1972.
- [24] Herbert A. Simon. Causal ordering and identifiability. In William C. Hood and Tjalling C. Koopmans, editors, *Studies in Econometric Method. Cowles Commission for Research in Economics. Monograph No. 14*, chapter III, pages 49–74. John Wiley & Sons, Inc., New York, NY, 1953.
- [25] Herbert A. Simon, Jayant R. Kalagnanam, and Marek J. Druzdel. Performance budget planning: The case of a research university. In preparation, 2000.
- [26] Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. Springer Verlag, New York, 1993.
- [27] Detlof von Winterfeldt and Ward Edwards. *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge, 1988.
- [28] Haiqin Wang and Marek J. Druzdel. User interface tools for navigation in conditional probability tables and elicitation of probabilities in Bayesian networks. In *Proceedings of the Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 617–625, San Francisco, CA, 2000. Morgan Kaufmann Publishers, Inc.