In *Working Notes of the AAAI-99 Workshop on 'Mixed-initiative Intelligence,'* pages 99-106, Orlando, Florida, 18 July 1999

ESP: A Mixed-Initiative Decision-Theoretic Decision Modeling System

Marek J. Druzdzel

Decision Systems Laboratory
School of Information Sciences, Intelligent Systems Program,
and Center for Biomedical Informatics, University of Pittsburgh
740 SIS Building, Pittsburgh, PA 15260
marek@sis.pitt.edu
http://www.pitt.edu/~druzdzel

Abstract

Systems built on decision-theoretic principles typically rely on models that contain a fixed in advance set of objectives and decision variables. This paper describes a decision-theoretic decision modeling system under development in the Decision Systems Laboratory that departs from this tradition and allows for flexible construction and later reconstruction of models according to the current objectives of the user. Given the knowledge of causal mechanisms acting in the system, a set of objectives, costs, and preferences encoded in the model, the system can take the initiative in discovering decision opportunities, framing decision problems, and solving them.

Introduction

Decision-theoretic systems are increasingly applied in various domains because of their sound foundations, ability to combine existing data with expert knowledge, and intuitive framework of directed graphical models, such as Bayesian networks (Pearl 1988) and influence diagrams (Howard & Matheson 1984). Some of the applications of decision-theoretic systems are: medical diagnosis and therapy planning, machine diagnosis, natural language processing, vision, robotics, planning, fraud detection, processing of military intelligence data in the context of battle damage assessment, and many others (March 1995 issue of *Communications of the ACM* lists several practical applications of Bayesian networks; others can be found in the electronic proceedings of the Annual Conference on Uncertainty in Artificial Intelligence, available on-line at http://www.sis.pitt.edu/~dsl/uai.html).

Typically, decision-theoretic systems use domain models that contain a fixed in advance set of objectives and decision variables. For example, a medical therapy choice model may encode of a number of possible therapies from which, given a patient's data, the system suggests one that is optimal for the patient. Sometimes, in case of complex real world systems, it is not obvious for a decision maker what decision options are available. In fact, framing decision problem and finding creative decision options is much harder than providing a numerical solution for it. Most decision-theoretic systems rely on human decision analysts and knowledge engineers to build decision models. Since models are built by humans and solved by computers, the dialog between the two is rather predetermined and it is hard to talk about initiative on the part of the computer system.

Many applications, such as strategic planning in organizations, require a close collaboration of humans and machines. Computers provide the ability to combine quantities of information and to implement practically error-free formal methods of reasoning. Humans are rich in intuition and common sense, both extremely valuable in structuring decision problems and evaluating computer recommendations.

This paper describes an ongoing effort in the Decision Systems Laboratory to develop a system that allows for a flexible model specification and later model reconstruction according to the current user objectives. This system, called ESP (Environment for Strategic Planning) is mixed initiative in the sense that both the system and its user can engage in a dialog aiming at model development, identifying decision alternatives, and finding the optimal alternative. There are four aspects of ESP that make it particularly suitable for mixed initiative interaction with the user. These are: (i) model building that can combine various types of information, (ii) identification of opportunities and generation of decision options, (iii) choice among identified options, and (iv) intuitive interaction with its user. We will outline the foundations of each of these in the remainder of the paper.

We first briefly review the principles underlying decision theoretic modeling, directed graphical models, causality, changes in structure, and choice of decision variables. We demonstrate that system's awareness of causal mechanisms encoded in the model allows for restructuring models and tailoring them to the decision-maker's needs. We then summarize the foundations of the ESP project and the type of activities that underlie its interaction with the user. Finally, we discuss our insight into the basic properties that a decision support system must have in order to be able to exhibit autonomy or initiative in its interaction with the user.

Decision-Theoretic Modeling, Causality, and Changes in Structure

Modeling

An essential step in applying decision-theoretic principles in intelligent systems is modeling. Models are abstractions of real world systems, consisting of variables and specification of interactions among them. Variables in a model can be divided into two basic classes: exogenous and endogenous. Exogenous variables are those, whose values are determined outside the model. Endogenous are those whose values are determined inside the model from the values of exogenous variables and the specification of interactions among the model's variables. In decision-theoretic systems, exogenous variables can be further divided into two subclasses: *decision variables*, which are variables that are under decision maker's control, and *environment variables* that are outside the influence of the decision maker. The latter are specified usually by probability distributions over their values. Decision-theoretic models include also the concept of preferences, a measure of which, *utility*, is a function of the possible values of some objective variables. Given a model, decision theory prescribes how to solve it, i.e., how to identify the actions (i.e., values of the decision variables) that will maximize the expected utility.

Decision analysis, which is an applied branch of decision theory, has developed a graphical modeling tool, known as influence diagram (Howard & Matheson 1984). Influence diagrams, closely related to Bayesian networks, are directed acyclic graphs in which nodes represent variables and arcs represent direct probabilistic dependencies among them. Rectangle-shaped nodes represent decision variables, diamond-shaped nodes represent utility functions. All other nodes, usually oval-shaped, represent other variables. An example of an influence diagram is presented in Figure 1.

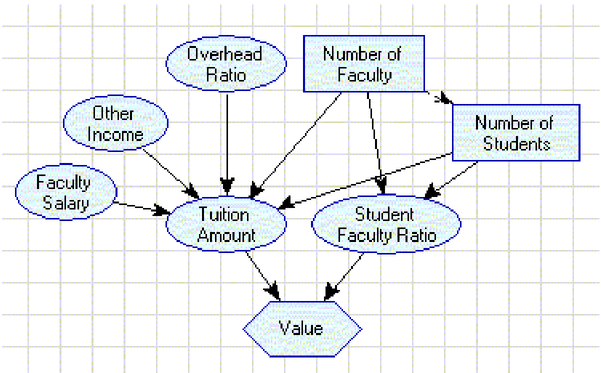

Figure 1: An influence diagram.

Nodes *Number of Faculty* and *Number of Students* are decision variables, nodes *Faculty Salary*, *Other Income*, and *Overhead Ratio* are exogenous environment variables, and nodes *Tuition Amount* and *Student Faculty Ratio* are objective variables. The model expresses also a utility function (represented by the node *Value*) over the objective variables.

Causality

Probability theory, that underlies decision theory, does not put any restrictions on the direction of the arcs between nodes in an influence diagram. Arcs, roughly speaking, specify only the direction of conditional probability and they can be reversed by means of Bayes theorem. As long as the only goal of using a model is prediction of a probability distribution given some evidence (this is the case in typical diagnostic tasks), the notion of causality is technically not useful. Consider, for example, a model consisting of two variables: *weather* and *barometer*. Given the outcome of one of the variables, we can do extremely well in predicting the probability distribution of the other.

Causality becomes crucial when we want to predict the effect of a "change in structure" of our system, i.e., the change in some mechanism in the system through an external intervention. Without knowing the direction of the causal relation between the *weather* and the *barometer*, we cannot tell whether a manual manipulation of the *barometer* will affect the *weather*. The effect of a structural change in a system cannot be induced from a model that does not contain causal information. Having the causality right is crucial for any decision making.

In decision analysis, typically one never deals with changes in structure – all decision options and instruments that are expected to affect a system are explicitly included in the decision model. Whatever causal knowledge is necessary for building this model is assumed to be possessed by the decision maker, and is captured in the conditional probability distributions in the model. The decision maker is assumed to know that, for example, manipulating the barometer will not affect the weather. As long as the probability distribution over the influence diagram nodes conditional on the value of decision nodes is computed correctly, decision theorists and analysts see no theoretical reason for assigning a causal interpretation to directed arcs. Obviously, they are willing to make an exception for the decision nodes – their outgoing arcs roughly speaking model the causal impact that the decision, an external manipulation of the system, will have on the model's variables.

Causality is extremely useful for the sake of user interfaces and can be referred to at the knowledge engineering stage and in explanation of the results of the system. In situations where the human-model loop is loose, for example in autonomous or mixed initiative systems, causal knowledge becomes crucial. To be able to autonomously generate and evaluate various decision options, intelligent systems need a way to compute the effect of imposing values or probability distributions on some of the variables in a model. A causal model gives a sound basis for this and frees the system from the need to store a combinatorially large set of pairs *action* → *effect*. The result of an external manipulation on model variables can be predicted directly from the model.

Causal mechanisms, causal ordering

One way of structuring domain knowledge is in terms of causal mechanisms, elementary, conceptually distinct, single mechanisms active in a system. The notion of a mechanism can be operationalized by providing a procedure for determining whether the mechanism is present and active or not. Sometimes a mechanism is visible and tangible. One can, for example, expose the clutch of a car and even touch the plates by which the car's engine is coupled with the wheels. One can even provide a graphic demonstration of the role of this mechanism by starting the engine and depressing the clutch pedal. Often, especially in systems studied in social sciences, a mechanism is not as transparent. Instead, one often has other clues or well-developed and empirically tested theories of interactions in the system that are based on elementary laws like "no action at a distance" or "no action without communication" (Simon 1977, page 52). Causal mechanisms may be identified entirely on the basis of a theory or consist of principles derived from observations, knowledge of legal and institutional rules restricting the system (such as tax schedules, prices, or pollution controls), technological knowledge, physical, chemical, or social laws. They may, alternatively, be formed on a dual basis: a theory supported by systematically collected data for the relevant variables.

In econometrics, causal mechanisms have been traditionally captured by structural equations (e.g., Koopmans 1950, Hood & Koopmans1953). A model composed of structural equations exhibits asymmetry among its variables that can be captured in a directed acyclic graph and given a causal interpretation. This was observed and explicated by Simon (1953), who developed an algorithm for extracting this graph and argued that, if each equation in the model is structural and each variable in the model that is assigned a constant value is an exogenous variable, then this ordering has a causal interpretation. This structure will correspond to the interactions in the real world as far as the model corresponds to the real world.

Consider a simple structural equation model of an aspect of university financial operations based on (Simon 1967):

$$\mathit{facsal}\ \mathit{nfac}\ (1 + \mathit{overh}) = \mathit{tuition}\ \mathit{nstud} + \mathit{oinc}$$

$$\mathit{stratio} = \mathit{nstud} / \mathit{nfac} \qquad (1)$$

The first equation states that the faculty salary (*facsalar*) is bound with the number of faculty (*nfac*), number of students (*nstud*), tuition amount (*tuition*), other income (*oinc*) such as grants and endowments, and overhead ration (*overh*). Student faculty ratio (*strat*) is a simple function of the number of students (*nstud*) and the number of faculty (*nfac*).

If we add five more equations describing the values of exogenous variables: *nstud*, *nfac*, *oinc*, *facsal*, and *overh*, we will obtain a self-contained system of simultaneous equations that we will be able to solve for the remaining, endogenous variables *tuition* and *stratio* (Figure 2).

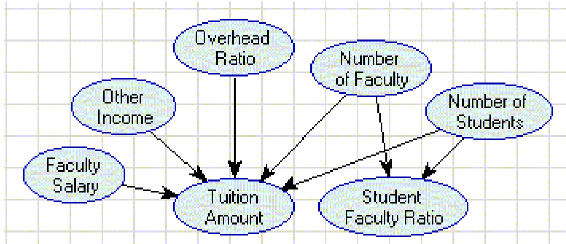

Figure 2: One mode of operation of the university model: tuition amount is endogenous.

If, however, the tuition amount is set by a policy decision, it may determine the faculty salaries. Note that Equation 1 is a hard constraint on where the money comes from and where it goes. This is modeled easily by providing equations for *nstud*, *nfac*, *oinc*, *tuition*, and *overh*, and solving for *facsal* and *stratio* (Figure 3).

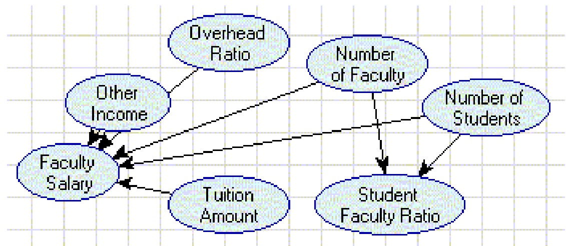

Figure 3: Another mode of operation of the university model: tuition amount is exogenous.

Once a model is causal, the process of prediction of effects of manipulation is straightforward: we remove the mechanisms that are no longer active in the system, add new mechanisms, and derive the new causal structure. Models that are based on structural equations supports changes in structure directly in the sense that the solution of the system of equations after the replacement will be the predicted state of the system after the change in structure (Hood & Koopmans 1953, Koopmans 1950). Ease of prediction of effects of manipulation almost matches that of human beings. Foundations for graphical support for prediction of effects of manipulation have been laid by Simon's work.

Causal manipulation and causal reversibility

Druzdzel and Simon (1993) have shown that Bayesian networks and structural equation models are related – any Bayesian network model can be represented by a corresponding structural equation model and we can view a

Bayesian network as a collection of causal mechanisms described by conditional probability distributions (these are described by functions in structural equation models). A collection of such mechanisms forms a causal model, which can be used to predict the effects of actions. This model can be automatically reconfigured for different possible action scenarios to account for a different ways of achieving a strategic objective and effectively deriving creative decision and action options.

In the context of directed graphical probabilistic models, such as Bayesian networks, two foremost formalizations of manipulation are due to Spirtes *et al.* and Pearl. Spirtes *et al.* (1993) proposed a theorem, known as *causal manipulation theorem*, specifying the effect of imposing externally a value on any node in a graphical model. Pearl (1995), who built his recent work on structural equation models, has an equivalent formalization of manipulation. While these formalizations are very useful and bring much clarity into statistics by making a clear distinction between observation and manipulation, they suffer from a modeling weakness. Both approaches imply that if the effect of an external manipulation of a node *n* imposes a deterministic value on it (for the sake of simplicity, we will consider only such interventions in this paper), the effect of this manipulation on other nodes in the network can be predicted by removing the arcs from the parents of *n* to *n*. This operation, popularly known as *arc cutting semantics*, is based on the premise that each of the mechanisms in the model represents an asymmetric relationship among their variables. This poses a limitation on modeling reversible causal mechanisms, i.e., mechanisms that are capable of working in several directions, depending on which of their variables are manipulated exogenously, such as the university example in Figures 1 through 3.

The reader can observe that a variable can act as a cause in one system but it can reverse its role and be an effect in another. This was the case for the variable *Tuition Amount* in our example model. More generally, it is often up to the context whether a variable is exogenous or endogenous. Many model variables can be manipulated directly and, if the decision maker chooses so, be designated as decision variables. An example of a reversible causal mechanism is the power train of a typical car. When a car goes up the hill, the engine moves the transmission, which, in turn, moves the wheels. When the car goes down the hill, the driver may want to use the power train to slow the car, i.e., let the wheels move the transmission, which then moves the engine. The interactions between the engine and the transmission and between the transmission and the wheels are independent of how the power train is actually used. Knowledge of these interactions can be reused whenever these mechanisms are parts of a larger system. Similarly, one might write a balance equation for the total number of aircrafts in an airbase. The total number of aircrafts is a strategic decision that impacts the base's ability to perform military missions. On the other hand, when planning a mission, the number of aircrafts sent is a strategic decision and it impacts the total number of aircrafts in the base. A good planning system should be able to represent just one balance equation and use it appropriately in different modes of operation. The importance of dealing with this algorithmically becomes especially clear in models consisting of tens or hundreds of variables.

Bayesian networks and the approach based on structural equation models taken by Pearl do not currently address the issue of reversible causal mechanisms. We have shown in (van Leijen & Druzdzel 1998) that causal reversibility can be modeled in Bayesian networks. In (Dash & Druzdzel 1999), we show that arc cutting semantics fails to model manipulation in equilibrium models when the underlying system has a dynamic nature.

The Environment for Strategic Planning (ESP) Project

We have employed causal manipulation procedures in a general purpose decision-theoretic modeling system that we call ESP (Environment for Strategic Planning). The goal of the ESP project is to develop a modeling system that will allow its user for flexible model building, model reuse, and identification of decision variables in addition to finding their optimal values. ESP employs both structural equation models and Bayesian networks. One of the project foundations is the theory of causal ordering (Simon 1953) enhanced with a causal interpretation of Bayesian networks (Druzdzel & Simon 1993). This approach allows for a normatively correct treatment of various types of information, uncertainty, and utility. ESP is especially powerful in complex situations where the available information is heterogeneous and consists of a mixture of deterministic and uncertain relationships among discrete and continuous variables.

Support for interactive model building

Model building is probably the hardest and most time consuming part of applying decision-theoretic principles. Model building can be aided or partly automated, but it still remains an activity where human experts or knowledge engineers take the leading role.

ESP allows for an interactive construction of decision models that is aided directly by a database of causal mechanisms and variables of interest. Given that equations describing system's interactions and conditional probability distributions are treated as causal mechanisms, i.e., elementary building blocks of larger systems, model construction amounts to selecting appropriate building blocks and adding or adjusting their numerical parameters. These building blocks are included in a database of mechanisms that is accessed by ESP during a modeling session. Some causal mechanisms in the database are reversible (such as the balance equation for the number of aircrafts in an airbase), others work always in the same direction (e.g., aircrafts consume and never produce fuel). Information about the nature of each of the mechanisms has to be included in the database. In addition, variables have to be marked with utility (if never of concern, then the utility is zero), cost of manipulation (if it is not possible to manipulate the variable, then the cost is infinite), and cost of observation (if it is not possible to observe a variable, the cost of observing it is infinite). Utility, cost of manipulation and observation are often dependent on the state of other system variables and form complex multi-attribute cost and utility structures.

Given initial information about user objectives, this process can be to a certain degree autonomous, although typically selection from a general database of mechanisms will be greatly aided by user's hints. The user has the opportunity to compose the model out of these building blocks, browse the database of mechanisms, and override system's suggestions. The system makes use of human intuition and experience. System's suggestions for building blocks can be overridden by the user. More details of the model building module of ESP are presented in (Druzdzel *et al.* 1998).

Learning

In case of data-rich domains, model building.can be aided by learning the model structure or the model parameters from previous cases stored in a database (e.g., Spirtes *et al.*, 1993, Cooper & Herskovitz, 1992). Also in this case it is important that the knowledge engineer be able to override the results of the learning module. This can be accomplished by entering prior knowledge (both structural and numerical) that the algorithms of the learning module have to take into account. Another way of overriding the learning module is modifying its results, such as the learned graphical connections and numerical parameters.

ESP's learning module is highly interactive and allows for both entering prior knowledge and overriding the learned parameters. The user can use one of several methods for discretizing continuous variables, discretize them manually, indicate presence or absence of arcs between nodes, and choose one of several available learning algorithms. A learned model can be manually refined by the user to accommodate possible differences between the available data and the system under study.

Search for opportunities

A decision maker confronted with a complex system may not know what actions should be taken to achieve a strategic goal — the set of decision variables and their optimal values may become clear only in the course of the analysis. Typical systems require the user to decide beforehand which variables are decision variables and which are intermediate or goal variables. A change in the set of decision variables may have a profound impact on the structure of the problem and on how their values will propagate through the system. A change in the set of decision variables has often to be followed by rebuilding the model, a process too cumbersome to be of exploratory nature. We believe that supporting changes in structure is essential for strategic planning systems and the user interface to such systems should make changes in structure easy and efficient.

One of the strengths of building a planning system on structural equation models is that it allows simple switching between dependent and independent variables, provided that the overall system remains consistent. The system then rebuilds the model by resolving functional references in the desired direction and draws the new model structure.

Even more importantly, this process of selecting a set of decision variables (or actions) can be done by the system autonomously. Once the system has a model, knows the user's objectives, their desirability in terms of a utility function over their possible values, which variables can be manipulated in the model and which cannot, how much the manipulation costs, it can start searching for a set of variables that when changed (by external manipulation) will achieve the objective at a minimum cost. Search for the strategy consists of a cycle: select decision variables, restructure the model using the method of causal ordering, and find their optimal values, expected costs, and benefits. This search can be interrupted by the user and restarted after a partial selection of decision variables. The insight obtained from the interaction with the system can lead to revision of the utilities and quite possibly the structure of the model.

Our approach opens ways for integrating heuristic AI techniques of searching for the optimal subset of decision variables to manipulate. Once the objectives and their utility are known, the system can engage in a creative process of choosing variables to be manipulated, restructuring the model, and computing the optimal settings of manipulated variables along with their costs. With the flexibility to reconfigure the system and to be able at each point to predict the expected cost, effect, and utility of actions, the system engages in a creative process of plan or policy generation.

The search process may be in general very costly, as there is a combinatorial number of groups of variables that can be manipulated. To reduce the computational complexity, the system can be guided by sensitivity analysis (i.e., what impact manipulation of individual variables has on the objectives) following the most promising gradient. In addition, the computation of optimal values can be performed using stochastic simulation. As the search aims at finding large effects in the first place, precision is of less importance. Precise values can be found for those plans and strategies that are most promising. The complexity of the search can be reduced by allowing for a system interaction with its user — often an expert with considerable experience and intuition regarding the most promising courses of action. Alternatively, it can be performed off-line, which may be especially attractive in case of search for opportunities, when there is not much pressure for an immediate course of action. Obviously, this procedure can apply AI heuristic search methods.

Consider the simple model of university operations illustrated in Figure 2. Suppose we want to decrease the student/faculty ratio in order to improve the quality of teaching and to make the university more competitive. The system can figure out that the best way of doing this is by admitting fewer students and at the same time compensating for the lost income by increasing other income. This will impact the tuition amount. Another solution is to reduce the number of students and to increase the tuition. Yet another solution in to reduce faculty salary while keeping the tuition constant. Each of these solutions originates from a different causal ordering among the variables (two such orderings are presented in Figures 2 and 3), because the causal mechanisms in this simple system are reversible.

Sensitivity analysis and value of information computation

Application of decision-theoretic techniques allows for implementation of three powerful decision-analytic tools: sensitivity analysis, importance analysis, and value of information. *Sensitivity analysis* (von Winterfeld & Edwards 1988) is used to identify those parts of the model that are critical to the decision. The main idea here is that one does not need to worry about those elements of the model that will not change the decision anyway, but invest effort in collecting information about those elements that can potentially impact the decision. *Value of information* analysis (e.g., von Winterfeld & Edwards 1988) looks at the model at a meta-level and computes the maximal costs that a decision maker should be willing to invest to obtain more information about some part of the model. *Importance analysis* (Henrion 1982) identifies those sources of uncertainty that contribute most to the uncertainty in the objective variable.

Sensitivity analysis, value of information computation, and importance analysis offer systematic ways of dealing with reduction of uncertainty and costs related with it. Each of them can be performed at the initiative of the system.

User interface

Application of formal techniques without attention to the interaction with the humans runs the risk of failure. Even though they guarantee correctness of computation (up to possible errors in the reasoning algorithms!) optimality, this optimality is illusory, as the search is conducted within a model, which is always at best a good approximation of reality.

We believe that it is crucial for normative decision support systems to be equipped with effective human-computer interfaces. Such systems can then be used not only as valuable decision aids in critical situations, but also as powerful training devices by allowing the learner to experiment with the inputs to the system and observe the consequences. Further, we believe that humans trained with such systems will in the long run be able to use them effectively to make rapid decisions in situations that are both complex and stressful.

We are focusing on three aspects of user interfaces: support for automatic and interactive model building, discussed earlier, intuitive presentations of the model structure. and explanation of results.

Understanding the structure of a problem, i.e., what variables are involved and how exactly they interact is usually very enlightening to the user. There is also anecdotal and some empirical evidence that in case of decision-theoretic models the structure of a model matters much more than the numbers for the quality of decisions (Henrion et al. 1996). Fostering insight into the structure of the problem can help the user to find possible model problems, usually coming down to wrong premises (note that the reasoning in normative systems is always correct, unless of course there is a bug in the program and it is the model than is critical for the quality of decisions). Graphical presentation of the model structure can also form a basis for group decision making sessions in electronic meeting rooms. We can derive the causal structure of the system in its every mode or operation, based on Simon's (1953) causal ordering. In addition to displaying the structure of the model and allowing the user to browse it, we plan to support queries about the chain of influence between any chosen pairs or groups of variables.

Once the system has proposed a decision, the user may question it. One of the first steps is to ask the system for a justification. Very often a graphical presentation of results, such as a graph of functional dependence between two variables, a tornado diagram, or a break-even point plot can offer enormous helps in understanding the recommendation. We plan to support this by verbal explanations, such as those proposed in (Henrion & Druzdzel 1990, Druzdzel 1996).

Discussion

Mixed initiative interaction between an intelligent decision support system and its users can enhance the performance of both. Systems are equipped with high computational power that does not diminish rapidly with complexity, while humans bring in domain knowledge, intuition, and common sense. Decision-theoretic systems are suitable for normative decision support and are capable of supporting mixed initiative interactions.

The ESP system that we have been developing in the Decision Systems Laboratory integrates automatic and interactive model construction from causal mechanisms stored in a knowledge base of mechanisms. It supports changes in structure as desired by the user. In addition, it supports autonomous and mixed initiative search for creative solutions through an automatic exploration of the possible ways that a given objective can be achieved. It supports sensitivity analysis of the objective to various parts of the model and value of information computation. Finally, it is equipped with a user interface that involves human decision-makers into the planning process, increase their insight into the problem, and support human intervention.

Four important aspects of ESP make it particularly suitable for mixed initiative interaction with the user. These are: (i) model building that can combine various types of information, (ii) identification of opportunities and generation of decision options, (iii) choice among identified options, and (iv) intuitive interaction with its user. We believe that the approach taken in the ESP project is unique in that it supports all these.

The prototype of the ESP system, described in this paper is currently under development. We hope to demonstrate some ESP modules during the workshop.

Our work has confronted us with the interesting question: what do machines need to exhibit autonomous behavior? We believe that flexible models that encode causality and a proper treatment of manipulation are important in mixed initiative systems and, we believe, crucial in decision support, as it frees the user from having to rebuild the model each time a different set of decision variables is considered. Within manipulations, it is important in many applications that the formalism used supports causal reversibility. Knowledge of causal asymmetries in a system is necessary in predicting the effects of changes in the structure of the system and, because of the role of causality in human reasoning, is essential in human-computer interfaces to decision support systems.

Acknowledgments

This research was supported by the Air Force Office of Scientific Research, grant F49620-97-1-0225 and by the National Science Foundation under Faculty Early Career Development (CAREER) Program, grant IRI-9624629. All illustrations included in this paper were created using **GeNIe**, a graphical modeling environment developed at the Decision Systems Laboratory, University of Pittsburgh and available at http://www2.sis.pitt.edu/~genie.

References

Cooper, G.F. and Herskovits, E. (1992). A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9(4):309-347.

Dash, D.H. and Druzdzel, M.J. (1999). An inconsistency between formalisms for causal discovery and causal reasoning. Under review.

Druzdzel, M.J., Lu, T.C. and Leong, T.Y. (1998). Interactive construction of decision models based on causal mechanisms. In *Working notes of the AAAI 1998 Spring Symposium on Interactive and Mixed-Initiative Decision-Theoretic Systems*, pages 38-44, Stanford, CA, March 23-25.

Druzdzel, M.J. (1996). Qualitative verbal explanations in Bayesian belief networks. *Artificial Intelligence and Simulation of Behaviour Quarterly*, special issue on Bayesian networks, 94:43-54, 1996.

Druzdzel, M. J. & Simon, H.A. (1993). Causality in Bayesian belief networks. In *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI-93)*, pp. 3-11, Washington, DC.

Henrion, M. (1982). The value of knowing how little you know: The advantages of a probabilistic treatment of uncertainty in policy analysis. Ph.D. dissertation, Carnegie Mellon University.

Henrion, M. & Druzdzel, M.J. (1990). Qualitative propagation and scenario-based approaches to explanation of probabilistic reasoning. In *Uncertainty in Artificial Intelligence 6*, P.P. Bonissone, M. Henrion, L.N. Kanal, and J.F. Lemmer (eds), Machine Intelligence and Pattern Recognition 12, pp. 17-32, Elsevier, North Holland: Amsterdam.

Henrion, M., Pradhan, M., Del Favero, B., Huang, K., Provan, G. & O'Rorke, P. (1996). Why is diagnosis using belief networks insensitive to imprecision in probabilities? In *Proceedings of the Twelfth Annual Conference on Uncertainty in Artificial Intelligence (UAI-96)*, pp. 307-314, Portland, OR.

Hood, W.C. & Koopmans, T.C. (Eds.) (1953). Studies in econometric method. *Cowles Commission for Research in Economics. Monograph no. 14*. New York, NY: John Wiley & Sons, Inc.

Howard, R.A. & Matheson, J.E. (1984). Influence diagrams. In R.A. Howard & J.E. Matheson (Eds.), *The Principles and Applications of Decision Analysis* (pp. 719-762). Menlo Park, CA: Strategic Decisions Group.

Koopmans, T.C. (Ed.) (1950). Statistical inference in dynamic economic models. *Cowles Commission for Research in Economics. Monograph no. 10,* New York, NY: John Wiley & Sons, Inc.

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

Pearl, J., (1995). Causal diagrams for empirical research. *Biometrika*, 82)(4):669-710.

Simon, H.A. (1953). Causal ordering and identifiability. In (Hood&Koopmans, 1953), Chap. III, pp. 49-74.

Simon, H.A. (1967). The Job of a College President. *The Educational Record*, pages 68-78, American Council on Education, Washington, D.C.

Spirtes, P., Glymour, C., & Scheines, R. (1993). Causation, Prediction, and Search. New York: Springer Verlag.

van Leijen, H. and Druzdzel, M.J. (1998). Reversible causal mechanisms in Bayesian networks. In *Working notes of the AAAI 1998 Spring Symposium on Prospects for a Commonsense Theory of Causation*, pages 24-30, Stanford, CA, March 23-25.

von Winterfeldt, D. and Edwards, W. (1988). *Decision Analysis and Behavioral Research*. Cambridge University Press: Cambridge.