IFAC/IFORS/IIASA/TIMS

The International Federation of Automatic Control
The International Federation of Operational Research Societies
The International Institute for Applied Systems Analysis
The Institute of Management Sciences

SUPPORT SYSTEMS FOR DECISION AND NEGOTIATION PROCESSES

Preprints of the IFAC/IFORS/IIASA/TIMS Workshop

Warsaw, Poland

June 24-26, 1992

Editors:

Roman Kulikowski
Zbigniew Nahorski
Jan W.Owsiński
Andrzej Straszak

Systems Research Institute
Polish Academy of Sciences
Warsaw, Poland

VOLUME 1:

Names of first authors: **A-K**

SYSTEMS RESEARCH INSTITUTE, POLISH ACADEMY OF SCIENCES

SUPPORT SYSTEMS FOR DECISION AND NEGOTIATION PROCESSES
Preprints, IFAC/IFORS/IIASA/TIMS Workshop, June 24-26, 1992, Warsaw, Poland

RULES AND ENVIRONMENTS: A HYBRID REPRESENTATION FOR DECISION PROBLEMS

GREGORY E. KERSTEN[+] AND STAN SZPAKOWICZ[++]
[+] School of Business, Carleton University, Ottawa, Canada
[++] Department of Computer Science, University of Ottawa, Ottawa, Canada

Abstract: We consider a hybrid representation paradigm for decision problems, in which the structural aspects of a problem are expressed in rule form, and the quantitative aspects are captured in a mixed symbolic/numeric language. We present informally the essential concepts of agents, goals, decision processes, decision environments, and flexibility. The discussion is couched in terms of NEGOPLAN, a decision support system embedded in the logic programming language Prolog. We show examples of problems successfully modelled in NEGOPLAN.

Keywords: Negotiation, Sequential decisions, Hybrid modelling, Rule-based systems, Decision support, Restructurable modelling

1. Structure of Decision Problems

People make and implement decisions because they want to achieve goals and to prepare for future decisions. Decision making is, therefore, an evolutionary, causally motivated reasoning process. It comprises the evaluation of the decision maker's (the agent's) goals and expectation, the analysis of the decision environment, the choice among alternative decisions. Alternatives are evaluated from a dual perspective, with a possibility of conflicts: the agent's current goals may be at odds with his evolving (and not always fully determined) future goals and with future (not always fully predictable) changes in the environment. For example, a trade union may want to improve work conditions and increase wages within the company, while trying to gain a strong political position in a region with significant unemployment.

Modelling and support of decisions has been traditionally concerned with the development of *closed* problem representations. Those are representations that do not evolve, do not permit changes from within. A closed representation of a decision problem may be preferred because it offers a conceptually simple way of solving the problem: it freezes the world, so that the agent, the environment and the problem itself are static. For instance, a manager may develop an understanding of the market and the company's position, and use that to consider how both would react to the initiation of a new project. The rapidly growing complexity of analysis would make it unlikely for him to take into account reactions to reactions, reactions to reactions to reactions, and so forth.

Closed formal representations may not be appropriate for complex strategic decisions because many aspects are not, and cannot be, accounted for. This is one of the reasons why decision support systems (DSS) only help model and analyze more detailed elements of a decision problem, but the decision maker is responsible for building and maintaining the "big picture". (Similarly, a typical expert system contains highly specialized and detailed knowledge.)

DSS work in tandem with the agent. This is equally true of intelligent DSS that do inference on a knowledge base. The system maintains quantitative representations, solves them, analyzes and compares solutions, and assists the agent in this "number-crunching" side of decision making. The agent supplies information about the problem, chooses modes, devel ps and enters structural aspects of the problem, selects alternatives, and generally speaking guides the system in cooperative work.

The strength of the DSS paradigm is in having the computer-based system as a participant in problem solving. A weakness is that numerical models are difficult to develop and not readily intelligible. This may be somewhat facilitated by giving the agent tools for building numerical models, perhaps based on model development methods proposed by Geoffrion (1987). It is also a weakness that the agent is required to supply the initial structures that go into the problem representation, and to describe modifications of structures and connections for future changes. We propose restructurable modelling (Kersten *et al.*, 1991) as a remedy for some of that weakness.

Model development methods work at the detailed level of problem solving. They support the building of an appropriate quantitative representation of the agent's decision problem. This is sufficient for problems that can be solved with one quantitative model. Suppose, however, that a problem can only be analyzed by constructing a series of more and more adequate models, and each solution contributes to the next model. Such models can be mutually compared, related and transformed in a structural framework, in a representation that treats them as relatively small building blocks. Model development methods are also sufficient when the agent supplies information required for "what-if" analysis, comparison of efficient (e.g. Pareto optimal) solutions, and sensitivity analysis. The agent mentally converts his goals into control information that the system needs to model the achievement of the goals. This conversion—and the preceding goal evaluation—can be seen as high-level restructuring of the problem, which it is desirable to represent explicitly. The agent's evolving image of the decision problem, and the fluid state of the environment, ought to be part of the problem itself.

2. Representing Qualitative Elements

In the NEGOPLAN project (Kersten and Szpakowicz, 1990; Matwin *at al.*, 1989) we have put forward a method of building *open*, changeable representations of the agent's perception of the decision problem and the decision environment. We model both the agent's goals and the decision process, during which the goals may change. We make two fundamental assumptions in NEGOPLAN: (1) a decision problem can be hierarchically decomposed into subproblems; (2) the uniqueness of a decision lies in its unique composition (from partial decisions or other elements, usually related to the environment), but the components themselves need not be—and normally are not—unique. This means that problem representations can be constructed from a predefined, sufficiently rich set of elements.

NEGOPLAN is both a methodology and a prototype computer system. The agent supported by the system has a principal goal that can be broken down into lower-level goals. Decomposition stops at elementary goals, referred to as *facts*. Facts, which correspond to decision variables in decision analysis, are directly verifiable in the problem domain, or can be treated as common knowledge. For example, when one models union-management negotiations, overtime, paid vacation time, and the company's contribution to the retirement plan could be seen as elementary, and the employees' income as non-elementary.

Decomposition results in a *generic goal representation*. It is a directed acyclic graph whose nodes are goals expressed by *predicates*. For example, a company's plan of action may be decomposed into individual projects, each project into elements, each element into descriptions of required resources. Another example: in union-management negotiations the union may require that paid vacation time and the company's contribution to the retirement plan increase, but that the overtime be paid or exchanged for additional vacation.

Parameters of predicates may be underspecified: a range rather than a single value is specified. Such parameters must be instantiated when a generic problem representation is adapted to a concrete problem. An instantiated description, referred to as a *goal representation*, can be interpreted as a logical formula in which nodes correspond to logical connectives AND and OR.

If each fact is associated with a truth value (*true* may be treated as "achieved", *false* as "not achieved"), the value of the principal goal will be determined. An assignment of truth values that satisfies the principal goal is referred to as a *goal solution*. A goal solution may assign a truth value only to selected facts; the remaining facts represent decision issues whose values are irrelevant, because the principal goal will be satisfied regardless of those values. For example, the economic goal of the union may be achieved by gaining a significant salary increase, and this makes the union flexible with respect to the company's contribution to health care.

A *metafact* is a fact annotated with its truth value and an indication of the agent who accepts this truth value. For example, "significant salary increase" could be true according to the union (and false according to the management), "modest salary increase"—according to the management. A decision-maker's perception of the problem (or, in the case of negotiations, a negotiator's position) is represented as a collection of metafacts stored in a knowledge base. Metafacts are also used to represent the state of the environment. In negotiation, the environment may include opponents with their own positions. The environment is complex: some of its elements (e.g. co-agents, opponents) are responsive to the agent's decisions, other elements change irrespective of the agent's actions. A market, an economy, an organization, a natural environment are examples of decision environment. To simplify the discussion, we say that the agent, the opponents (if any), and the environment make decisions represented as positions composed of metafacts.

A position affects the situation by eliciting a reaction from other participants of the decision process and perhaps from the agent himself. Such reactions are modelled by *metarules*. A metarule lists metafacts whose presence in NEGOPLAN's knowledge base triggers the reaction; it may also specify tests and actions that must be done on the parameters of facts before the metarule can be applied. An application causes the addition of new metafacts to the knowledge base. For example, when the management accepts a significant salary increase and a small contribution to health care, the union will respond with a concession regarding the pension plan contribution. The union's response will be represented as a new metafact.

3. Modelling Quantitative Aspects

Decision problems that can be modelled with NEGOPLAN are complex by virtue of their multi-level hierarchy, and their dependence on the environment whose parts may have to be modelled separately and linked with the problem's representation. Temporal links are presupposed between certain parts of problem representation. Structure is represented with rules. Links and dependencies are normally represented with ancillary quantitative elements.

The quantitative aspects of the decision problem are those that involve numerical calculations. Three kinds of quantitative information are described by parameters of rules and meta-rules. We discuss them in terms of the changing representations of a decision problem.

(1) A parameter appears in a rule in one representation, and its value must be determined (typically, by a procedure embedded in the rule). Examples: the management must calculate the budget for the union's negotiation proposal; the union computes a salary increase for its offer from the current salaries, inflation, and settlements recently negotiated in other companies.

(2) Several representations share a parameter *and* its value. This can be illustrated with the union's initial problem representation stating the salary "reservation price" (Raiffa, 1982). This value is then used to determine concessions in subsequent problem representations.

(3) Several representations share a parameter whose value changes in time. A simple example: a parameter used to describe the cycles, phases, or rounds in negotiation.

The quantitative aspects are expressed in a mixed symbolic-numeric language. A predicate's name suggests the nature of the event or situation represented by this predicate—e.g. `strike_readiness`—and summarizes the relationship between arguments—e.g.

`strike_readiness(union, high)`. An argument may be a symbol that stands for an entity in the domain (e.g. `union`). It may also be a number, a measure of something—e.g. `wage_hike_pct(4.5)`.

The current prototype of NEGOPLAN is implemented in Prolog. Operations on parameters (both symbolic and numeric) are carried out in this language. They take the form of Prolog calls embedded in metarules. This is a significant extension of the usual expressive power of rule-driven systems.

4. Decision Processes

A decision problem is usually part of a larger process. A decision leads to other decision-making. Even a single decision problem may require "what-if" analysis or sensitivity analysis, or it may be necessary for the agent to repeatedly adjust and modify problem representations in search for the most appropriate one (Mayer, 1989). The process need not happen in real time—it may evolve subjectively, in the agent's perception and understanding. If it does change in time, we refer to it as a *sequential* decision problem. Obviously, such problems also encompass subjective changes of representation at any specific moment.

In a hybrid representation of decision problems, the structural and quantitative aspects and the representation of the decision environment are integrated into a coherent formal model. Integration with the quantitative aspects of the decision environment is usually present in models typically built in MS/OR. The hybrid approach makes it possible to consider different (not necessarily uniform) environments, to model interaction between the agent's decisions and the environment, and to build and modify representations.

In sequential decision-making the subsequent problem representations must reflect the changing perspectives and conditions. We assume that a computer system based on restructurable modelling would maintain a repository of past cases, for example, previously constructed representations of union/management negotiation. The structural elements of such cases (rules, metafacts, meta-rules) could be retrieved and used as a departure point in an interactive construction of a representation for a similar or related case. Sometimes large fragments of a previous case might be reused, sometimes only the barest structural shape would be kept. Interaction would produce a generic goal representation (GGR).

GGR usually contains competitive elements that describe alternative ways of achieving goals. Assume that we can select from the GGR the principal goal (with one of its alternative decompositions), and all facts that contribute to achieving this goal (again, always choosing one decomposition). If exactly one assignment of truth values to facts leads to the *true* value of the principal goal, this selection of goals and facts is called a *goal solution* (GS). A GS is, in a sense, isomorphic with one of potentially many possible outcomes of the decision problem.

Facts and goals from the GGR that have not been selected for a GS are referred to as *flexible*. This is because, in this solution, their values do not influence the value of the principal goal. Decision flexibility will be the higher the more redundancy has been built into the problem representation. Redundancy can be measured by the number of alternative solutions. We talk about flexibility when certain aspects or characteristics of the problem need not be taken into account in a given solution. When the solution is implemented, these aspects may be present or absent without influencing the principal goal. In negotiation, they may be considered as bargaining chips.

Implementing a solution may cause changes in the environment, and trigger reactions of other agents. These, in turn, cause modification of the current problem representation. A particular modification is chosen from a set of candidates on the basis of metafacts which describe the present state of the agent and the environment.

Restructuring may be simple—a small change achieved by adjusting metafacts. If a flexible fact is assigned a truth value by a metarule, it becomes temporarily bound. A more profound restructuring is described by metarules that take previously unused structural elements and add them to the GGR, or remove elements from the current GGR. This gives a modified GGR, and consequently a new GS.

5. Case Studies

5. 1. Disaster management

We have considered the disaster management and control responsibilities of a manager of a chemical plant (Michalowski *et al.*, 1991). The plant, producing a toxic, flammable gas, operates around the clock. The manager has at his disposal written procedures that describe actions to be taken should an accident occur. In response to an accident, the manager draws upon limited resources located in the plant and in the city. He can request emergency equipment from neighbouring counties, but he has no control over its availability. Weather conditions may prevent access from these counties.

The possible disasters are fire, chemical leak, and industrial accident. They may be big or small, and happen during the day or night. Depending on weather conditions, the city's population may be affected by an accident in the plant. The number of casualties and the amount of damage depend on the time necessary to contain the accident, and on the resources put into action. The longer it takes to combat a disaster, the higher the casualties and the greater the damage. The time spent on dealing with an accident depends on the amount of emergency equipment that reached the scene of the accident. The accident may happen at any time, its magnitude may change, and it is possible for an accident to trigger another of a different type. The manager may have to deal with many accidents simultaneously.

The disaster manager must establish effective responses to an accident. The NEGOPLAN model was developed as a first step toward the assessment of managerial skills in such a situation. The use of restructurable modelling allows modification of the initial decision context in order to introduce unforeseen events. The quantitative mechanisms make it possible to evaluate the effectiveness of the manager's control.

5.2. Simulation of a robot's mission

This case was developed in order to simulate the decision making processes of an industrial robot performing a mission in an unknown environment (Kersten *et al.*, 1990). The environment was described by four parameters: radiation, temperature, the surface, and the level of light. The mission consisted in collecting a number of samples and photographs, and returning to the base. If the robot's limited energy were depleted, the mission would fail: samples and photographs could not be delivered. Energy use and individual actions (take a picture or pick up a sample) depended on the state of the environment, modelled as unpredictable.

Simulation of the robot's movements in a dynamic, random environment required numerical control procedures. A purely symbolic representation was enhanced by quantitative choice mechanisms, in the form of embedded computations used to calculate parameter values. Similar embedded mechanisms were used to access information about the environment and to maintain the logical consistency between parameters describing the same aspect at different time or parameters that describe interconnected aspects. Another issue was to determine a sequence of positions on the basis of historical data—the robot's previous positions and states of the environment.

6. Conclusion

Realistic, evolving decision problems imply shifts of perspective, and require the feeding of intermediate decisions back into the later stages of the decision process. We have introduced the concept of open representations, in which incarnations of the decision problem form a sequence: each is rooted in the previous decision and arises from its modification. We argue for logic-based modelling of the structural aspects of the problem and for inference-driven modelling of the decision process. The agents—decision-makers—and the decision environment influence each other. Such influences can be represented as reasoning in the NEGOPLAN formalism that embodies the essential elements of our methodology.

The quantitative aspects of the decision problem are represented by number-valued parameters of logical expressions, and by procedures that calculate new values. Such procedures add a new dimension to the more traditional, rule-oriented problem representations typical of expert systems.

Competitive elements of the representation—alternative ways of structuring and solving the problem—allow us to introduce an important notion of flexibility.

Test cases, described here informally but implemented in NEGOPLAN, demonstrate the expressive power of our method.

References

Geoffrion, A. (1987),"An Introduction to Structured Modeling", *Management Science*, Vol. 33, No. 5, (pp. 547-588).

Kersten, G. E., W. Michalowski, S. Szpakowicz and Z. Koperczak (1991), "Restructurable Representations of Negotiation", *Management Science*, Vol. 37, No. 10, (pp. 1269-1290).

Kersten, G. E., S. Szpakowicz and Z. Koperczak (1990) "Modelling of Decision Making for Decision Processes in Dynamic Environments", *Computers and Mathematics with Applications*, Vol. 20, No. 9/10, (pp. 29-43).

Kersten G. E., and S. Szpakowicz (1990), "Rule-Based Formalism and Preference Representation: An Extension of NEGOPLAN". *European Journal of Operational Research* Vol. 45, (pp. 309-323).

Matwin, S., S. Szpakowicz, Z. Koperczak, G. E. Kersten and W. Michalowski (1989) "NEGOPLAN: An Expert System Shell for Negotiation Support", *IEEE Expert*, Vol. 4, No. 4, (pp. 50-62).

Mayer, R. E. (1989) "Human Nonadversary Problem Solving", *Human and Machine Problem Solving*. K.J. Gilhooly, Ed., New York: Plenum.

Michalowski, W., G. E. Kersten, Z. Koperczak and S. Szpakowicz (1991), "Disaster Management with NEGOPLAN", *Expert Systems with Applications*, Vol. 2, (pp. 107-120).

Raiffa, H. (1982) *The Art and Science of Negotiation*. Cambridge, MA: Harvard Univ. Press.