

Introduction

The increasing complexity of technological systems has shifted the demands on human performance from a mechanical/physical level to a cognitive level. The role of the human in complex systems is to act as a supervisor of automated or semiautomated resources with goal setting, problem solving, and decision making as the primary tasks. While the negative consequences of possible malfunctioning have grown, the tolerances for performance variability have been reduced and the demands to accurate and safe performance increased, thereby leaving the human operator in a very difficult position. At the same time, advances in computer science (e.g. automated decision makers, object oriented programming, and expert systems) have provided new possibilities to support human performance that challenge established man-machine system design principles.

Decision making has emerged as the focal point for the often conflicting demands to human action. To ease the operator's task one must reduce these demands by incorporating basic intelligence functions in the man-machine interface. But it is essential that an intelligent interface supports the operator's decision making rather than replaces parts of it, i.e. that it is a tool rather than a prosthesis. The former may improve the task, while the latter will surely aggravate it. The logic behind decision making has been considerably extended in recent years, for instance by fuzzy-set theory and possibility theory. A clear distinction has also emerged between domain specific knowledge and decision making strategies. Decision making is no longer regarded as simply the strict following of a single set of rules or a strategy, but must also include the selection of the appropriate rules and the possible switch between them as the decision evolves, i.e. metalevel decisions.

The aim of this ASI was to identify the knowns and unknowns of intelligent systems that can support human decision making in process environments. The focus was the functional rather than the analytical aspects of such systems. The emphasis on and the development of systems that perform cognitive tasks require a corresponding shift in the multiple disciplines that support effective man-machine systems. It requires contributions from decision theory, theoretical and applied; from philosophy and logic; from process control theory and information science; from cognitive psychology and the study of human performance and human error; and from artificial intelligence and computer science. This was achieved by describing (1) the foundation provided by decision theory, (2) the problems of decision making in process environments, (3) the cognitive aspects of decision making, and (4) the possibilities of artificial intelligence and advanced computer applications. The ASI aimed to synthesise and integrate developments in all four areas.

As mentioned in the preface, this ASI joined a series of other meetings that all had dealt with problems of man-machine systems. The five major previous meetings were:

1976	Monitoring Behaviour and Supervisory Performance	Berchtesgaden, West Germany
1977	Theory and Measurement of Mental Load	Mati, Greece
1979	Changes in Nature and Quality of Working Life	Thessaloniki, Greece
1981	Human Detection and Diagnosis of System Failures	Roskilde, Denmark
1983	Human Error	Bellagio, Italy

The present ASI used the interdisciplinary cooperation and knowledge foundation established in the previous conferences to address the problems of cognitive man-machine systems. During the two weeks specialists from the above-mentioned fields worked together to advance the synthesis of knowledge needed for the development and evaluation of such joint cognitive systems, present the current unsolved problems, and point to possible solutions. The outcome is documented in this book.

Whereas most readers of this book probably will have a good understanding of what decision making and decision support systems are, it may be useful to describe briefly what is meant by process environments and how decision making in process environments differs from other types of decision making.

Process environments are first of all characterised by having short time constants. The decision maker must constantly pay attention to the process because the state of it changes dynamically; if he fails to do so he will lose control. If time was not that critical it would be possible to leave the process and look for information in other places, consult available experts and knowledge bases, etc. To use an example from medicine, if time is not critical one is dealing with a population phenomenon, for example the effects of smoking; if time is critical one is dealing with an individual phenomenon, for example toxification. If there is no time pressure one can always refer to data from the population and neglect the process without endangering it.

Formal decision theories do generally not contain any element of time (cf. the paper by Giuseppe Volta), hence fail to recognise an essential attribute of process environments. They are therefore not directly applicable but require further development and modification.

Another characteristic of process environments is the uncertainty about data, i.e. one is uncertain about what the next datum may be. There is less than complete predictability of what evidence will appear, both on the level of individual process parameters and on the level of the process as a whole. The uncertainty does not only arise from the imperfection or breakdown of single components as in normal reliability analysis or Probabilistic Risk Analysis. Far more important is the uncertainty of the evidence that comes from the dynamics and complexity of the process as such. This lack of certainty is in conflict with some of the main assumptions of both normative and descriptive decision theory and naturally makes decisions more difficult (cf. the comments by George Apostolakis).

Process environments are further characterised by being dynamic, by having multiple and possibly conflicting goals, and by having incomplete information. The latter may lead to tradeoffs between the speed and quality of the decision. Decision making in process environments normally means multiple decisions based on partial information, with little chance to consider all the alternatives or to make revisions. It is thus very different from the orderly world that is assumed by conventional decision theories – and provided by most of the experiments carried out in the behavioral sciences.

The presentations and discussions at the ASI, as well as the papers in this book, are organised in four sections corresponding to the four main aspects of Intelligent Decision Support systems. These are *decision theory*, *cognitive engineering*, *systems engineering*, and *artificial intelligence*. The development and application of Intelligent Decision Support systems must be based on a combination of these four aspects, as none of them is sufficient on its own. Each of them refer again to other scientific disciplines some of which, like cybernetics, are related to several aspects of Intelligent Decision Support. The presentations and discussions at this ASI clearly showed the insufficiency of any one of these aspects to account completely for the functioning of Intelligent Decision Support systems, but also demonstrated how interdisciplinary cooperation can bring about new ideas and solutions.

Decision theory, whether normative or descriptive, is a natural starting point for this book. It provides a basic understanding of the theoretical and practical issues and of the currently unresolved problems. Cognitive engineering gives the view of a recent combination of a number of disciplines that all focus on the cognitive aspects of system functioning, and in particular addresses the relation between knowledge and decision making. This view is closely related to the current developments in systems engineering which, together with decision theory, represent the more established views. The section on systems engineering summarises and identifies some of the shortcomings of a traditional engineering approach, and points to ways in which they can be solved. Finally, the section on artificial intelligence provides the experience from trying to incorporate

intelligent capabilities in artificial systems. Artificial intelligence has produced a fascinating repertoire of ideas and tools which can be applied in developing Intelligent Decision Support systems, and the influence on cognitive engineering and systems engineering is obvious.

We believe that these four aspects are both necessary and sufficient to give a comprehensive description of Intelligent Decision Support systems, to point out the major problems, and indicate how the solutions can be achieved. It is also our strong belief that solutions can only be achieved by a combination of two or more of these disciplines. The insufficiency of an approach to Intelligent Decision Support systems based purely on either of them is demonstrated both in the papers and in the summarised panel discussion. Decision theory and systems engineering both focus on a description of the problem domain, decision making in process environments. Cognitive engineering and artificial intelligence provide an outline as well as examples of the methodological and functional solutions that can be applied on the problem domain. This ASI was an expression of the clearly recognised need for an interdisciplinary cooperation, and we trust that the book demonstrates the benefits of mixing minds and metaphors as a foundation for further advances into a new and challenging field.