INTELLIGENT MAN TALKS TO INTELLIGENT MACHINE: IMPLICATIONS OF DISTRIBUTED-INTELLIGENCE SYSTEMS FOR THE DESIGN OF MAN-MACHINE INTERFACE

W. Zachary*, F. Glenn*, and J. Hopson**

** Analytics, Willow Grove, Pennsylvania*
*** Naval Air Development Center, Warminster, Pennsylvania*

ABSTRACT

Advances in artificial intelligence and hardware miniaturization indicate that the "computer-aided" systems that are the focus of current research will be rapidly displaced by a generation of distributed-intelligence systems-ones where man and computer have complementary, intelligent roles and must work together as a team. The forms of communication which must occur between intelligent, autonomous man and intelligent, autonomous computer in these systems will be categorically different than those forms which occur in the systems of today. Man-machine communication requirements in distributed intelligence systems are described and analyzed in this paper, both in general and in the context of one such system currently under development. Also presented are some approaches to designing man-machine interfaces for these types of systems and a discussion of the research questions which must be answered before these approaches can be brought to fruition.

RÉSUMÉ

Les développements dans l'intelligence artificielle et dans la miniaturisation du matériel indiquent que les systèmes "assistés par ordinateur" qui dominent la recherche seront rapidement devancés par une génération de systèmes à intelligence répartie - systèmes dans lesquels l'homme et l'ordinateur ont des rôles intelligents complémentaires et doivent travailler en équipe. Les formes de communications qui doivent exister entre un homme intelligent et autonome et un ordinateur autonome et intelligent dans ces systèmes seront catégoriquement différentes des formes qui existent dans les systèmes actuels. Les exigences de communication homme-machine dans les systèmes à intelligence répartie sont décrites et analysées dans le présent document tant en général que dans le contexte d'un tel système actuellement en développement. Le document présente également quelques approches à la conception d'interfaces homme-machine pour ces types de systèmes et une étude sur les questions de recherches auxquelles il faut répondre avant que ces approches donnent des résultats concrets.

Introduction

The relationship between cybernetic machines and the humans which use them has evolved along with the capabilities of the machines themselves. As the machines have evolved from primitive regulators and servomechanisms to sophisticated general purpose digital computers with limited cognitive capabilities, their use by humans has increased and broadened into complex interactive partnerships that pervade almost every aspect of everyday life.

The most advanced form of man-computer relationship to achieve widespread application has been the "computer-aided" system, in which specialized computers assist humans in the performance of complex creative and/or intellectual tasks. Computer-aided systems are currently in use for assisting in translation of natural language documents, in the development of designs of digital circuits, in the generaton of workplace designs, in optimization of decision making, and in the detailing of artistic works, to name but a few of the most widely-known applications. In these computer-aided systems the human and computer work as partners, interacting extensively and sharing much information. However, the human alone is given the initiative to act, with the computer merely supplying data, information, and analyses when requested to do so by the man. Without direct requests for action from the human, the computer does nothing, except possibly to alert the human that certain predefined conditions have occured.

Advances in artificial intelligence, although substantial in recent years, have seen few practical applications. A possible reason for this is that most AI systems have been built in research laboratories to work in isolation --without any substantive interactions with other intelligent systems or humans performing related or complementary tasks. The result has been the development of programs which behave "intelligently" but only with regard to certain contrived problems in highly constrained situations. AI researchers have, perhaps, only recently begun to realize that an important aspect to human intelligence is that it does not operate in isolation, but in a complex interactional and communicational milieu that may include intelligent machines and other intelligent humans.

The realization is now growing that the symbiotic man-computer partnership represented by the computer-aided system approach can be profitably applied to the artificial intelligence breakthroughs to create a new type of practical application system. In these systems, unlike computer-aided systems, intelligent functions are performed by both man and computer, and initiative for action is accordingly assigned to both man and computer. Thus, the man and computer interact not as an intelligent entity using an intelligent (albeit a highly complex and sophisticated) tool, but rather as intelligent partners or co-workers, working together in complementary roles to complete complex tasks. Such systems can be termed "distributed-intelligence" systems, because in them the intelligent functions are distributed among all components of the system, regardless of whether they are humans or computers.

Distributed-intelligence systems of this sort will demand an entirely different form of man-computer communication than the present computer-aided systems. The interactive computer-aided systems of today employ the machine as a sophisticated but dedicated servant of the man to whom all decision-making responsibility is allocated. This role for the computer -- as a glorified servomechanism -- allows the communication between man and machine to be basically one of data transfer. The man inputs data, menu-selections, cue responses, and decisions to the computer, while the computer outputs data, calculation results, cues and messages to the man. Complex, abstract communication of the sort which typifies interactions among humans working in teams is not necessary. In the coming generation of distributed-intelligence systems, very different role assignments will prevail.

Allocation of Cognitive Functions

The major distinction between man and computer in distributed-intelligence systems lies in the capabilities and limitations of each to perform specific cognitive functions. For example, humans and computers have very different limitations with respect to memory and computational resources. Whereas computers are particularly effective and reliable in dealing with masses of quantitative data, humans tend to excel at filtering out irrelevant details and in focusing, for memory and manipulation, on associative concepts. Also, the intelligence of computer algorithms is much more restricted by context than is the intelligence of humans. Such differences help to determine the appropriate roles of man and computer in distributed intelligence systems.

While it is recognized that the underlying structures of distributed-intelligence systems are necessarily unique and specific to the situations they address, generalized statements can be made regarding the functions to be performed by human operators and automated processors. Many algorithms exist which permit computers to perform intellectual processing functions but some tasks have not yet yielded (and, indeed, may never yield) to satisfactory algorithmic resolution (hence, the rationale for developing distributed-intelligence systems). For example, algorithms can not effectively perform abstract problem solving tasks, establish heuristic values, and recognize meaningfulness of novel information which is relevant to the decision process. However, computers can acquire, organize, and analyze data, and develop inferences from the data. Based on these processing functions, computers then have the capability to perform other functions such as:

- assisting the operator in determining an optimal course of action,
- advising the operator of problems which are algorithmically insolvable,
- assisting in establishing priorities,
- predicting future consequences based on comparison of historical and present patterns, and
- alerting the operator to new data so it won't be overlooked.

Functions that human operators perform in a distributed-intelligence system are consequently associated with the higher, more abstract decision or problem solving levels. Since computers have limited intelligence, functions designated to operators are to assess the validity of automated processing results, determine the value of information in respect to the total problem, input new values, alter the decision path taken by the automated intelligent system when necessary, and, in general, monitor problem solving analyses and actions taken by the computer. The roles of the man and computer in distributed-intelligence systems thus progress from worker and tool to that of superordinate and subordinate co-workers. Man as the executive oversees the problem solving task while the computer as a subordinate does the "leg work" and advises the executive.

Interface Considerations

With the allocation of intellectual functions to both human and computer components of systems, unique interface mechanisms are needed to facilitate the complex interaction between the two. Man/machine interfaces for conventional systems have only been concerned with presentation of data and choice of input/output channels. Although these concerns are still pertinent, the major issue is how to structure an interface which allows both intelligent entities, human operators and computers, to communicate with each other. At a minimum, mechanisms are needed so that the human operator can monitor the autonomous decisions made and actions taken by the computer (i.e., know what the computer-as-subordinate is doing), interrogate the computer about its actions/decisions (i.e., find out why it is doing what it is doing), and intervene or override computer actions with which the operator disagrees (i.e., correct it when a mistake is made or a more appropriate action is recognized by the human). Similiar types of communications are needed by the computer - it needs to know what it's human supervisor is doing, why he is doing it, and how any inaccuracies and/or inconsistencies can be corrected or brought to his attention. These communicative functions -- monitoring, querying, intervening -- are not new, but distributed-intelligence systems require them to be performed at a higher and more abstract level than anything envisioned in present-day man-computer systems.

It is not the intent of distributed-intelligence systems to eliminate data sources from operators, but rather to alleviate the inundation of data so they can better perform their primary tasks. There are at least three informational levels the operator needs -- dynamic structure of the problem, decisions made by the automated processor and data sources used in the analyses. A holistic or gestalt representation of the situation is needed so that human operators can readily see the scope of the problem without having to integrate individual data sources to determine problem structure. The representation should be dynamic to reflect actions taken, impact of actions (through predictive models), and status of other contributing factors (e.g., enemy actions). For operators to assess algorithm peformance, a mechanism is needed to make the processing functions transparent. One method is to present operators with the results of key decisions as the problem solving analyses are taking place. Should the operator want to override any decision, a mechanism is needed to change the decision output. If the operator needs further detail regarding how the decision was made, a mechanism must be provided to access this level of detail. To provide operators with these necessary capabilities, four different interface subsystems are needed -- a conceptual situation display, an algorithm monitoring display, a query system, and a control input system. The control input system must allow not only for overriding decisions but also for altering the decision path taken by the automated processor or for changing the information on which this decision analysis was based.

Although automated processor functions are predetermined by system design, this design may call for the computer to initiate communciations with the human operator for various intelligent purposes. For instance, automated intelligent systems have the capability to monitor operators' performance and to alert them to critical situations. Or it may be necessary for the computer to query the operator regarding new inputs made to the system.

The structure of the interface subsystems, though not yet designed, is envisioned as rather complex. For example, an underlying component of the distributed-intelligence system requires an executive not only to communicate with operators, but also to communicate with individual processors to "orchestrate" the automated tasks. Results obtained from individual processors must be integrated and organized in a structure which operators can understand in terms of their own methods of problem solving. Conversely, a translator/director is needed to transform operators' inputs and direct them to the appropriate processors.

Consider, for example, a system in which the computer is controlling the positioning of players in a basketball game. If the computer were asked by the human "general manager" what it was doing at the moment, we want it to respond "I recognized that the opponents were going into a zone defense so I responded by overloading the zone to the left side," not "I sent player 36 to coordinates 3.5, 4.7 on the floor, I sent player 39 to coordinates 3.9, 5.2 on the floor, . . ." The first response is a conceptual response, the second is a data response. Similarly, if a computer that was controlling the positioning of air defenses against enemy air attack were asked by a human tactical controller what it was doing, we want it to respond "I recognized the movement of enemy aircraft to the south as a feint and therefore directed my defenses to follow them at a distance sufficient to allow them to respond to the main attack which I anticipate will come from the east." These responses need not be in a natural language, of course; they could be in graphic symbols, mathematical language, or some other form of communication. What is important is that the computer be able to utilize this conceptual level of communication, not the form that this communication takes.

It is also important to plan for flexibility of transition between the levels of conceptual communication employed in the man-computer communication process. While the dialog suggested above focuses on a very abstract level, it is often necessary to "jump" from detailed to abstract conceptual levels during a single dialog. For example, if the attempt to "overload the zone" failed, the human might want to know why, and the computer could appropriately alter the conceptual level to a more detailed one and respond in terms of the failure of player x to perform a specific task properly. Thus it is not merely the need to communicate in conceptual terms that is important, it is also the need to manipulate the conceptual level according to the communicational context.

Utilizing Cognitive Capabilities in System Design

In designing distributed-intelligence systems the processes of function allocation and interface design tend to be confounded with one another. The manner in which functions are allocated determines the information tranfer specifications for the man-machine interface. But operation of the interface requires both man and machine to perform cognitive processes of the same type as is required for accomplishment of the basic decision functions. The concept of distributed-intelligence implies a high degree of interaction between man and machine in order to communicate inputs and outputs of decision-making processes and to enable man and machine to monitor each other's performance. Thus, information processing functions must be allocated with consideration of the additional processing requirements that will be imposed by the interface on both man and machine.

In accomplishing the intertwined tasks of function allocation and interface design, the system designer must deal with three major constraining factors -- decision task requirements, man and machine cognitive capabilities, and characteristics of interface options. In particular, task requirements and interface characteristics must be understood in terms compatible with knowledge of man and machine cognitive capabilities. Understanding of cognitive capabilities assumes a pivotal role in this process because the functions to be allocated and the interfacing operations are primarily cognitive. To facilitate function allocation, it is important for the cognitive capabilities of man and machine to be represented in terms of capabilities to perform generic cognitive processes that can be potentially accomplished by man or machine. By expressing task requirements in terms of such generic cognitive processes it is possible to determine, on the basis of known capabilites, whether man or machine is better suited for each process.

Identification of suitable generic cognitive processes must satisfy several constraints. The processes must be specific enough to permit clear association with particular decision tasks and to foster unambiguous delineation of human and machine capabilities. The processes must be general enough to be accomplishable by man or machine and yet be sufficiently specific to relate to the empirical literature on man and machine performance. The total collection of identified cognitive processes must be comprehensive, dealing with all information processing operations needed for decision making, yet the number of distinct processes must be small enough to permit compilation of a useful inventory of man and machine processing characteristics. Comprehensiveness should apply to the depth as well as the breadth of the set of cognitive processes. The processes should include registration processes such as perception and memory, manipulation processes such as reasoning and modeling, and management processes such as learning and hierarchical planning. Detailed specifications of component cognitive processes should be constructed through analyses of processing requirements for developing distributed-intelligence systems.

The AEW Intercept Planning System

Efforts are currently in progress to develop a distributed intelligence system for the intercept planning mission of a Naval airborne early warning (AEW) aircraft. This mission, which is performed in simultaneous conjunction with several other missions, requires the operators of the AEW system to monitor a large air space for threatening activity and to direct friendly fighters in the investigation and interception of possible threats. Much sophisticated electronic equipment is used in current AEW platforms to detect and classify airborne objects, primarily through the use of radar. However, the enemy can be relied upon to make the job difficult through the use of radar and radio jamming and deceptive tactics. Furthermore, the high speeds of modern military aircraft and the long ranges of many weapons systems imply that there is generally very little time for the AEW operators to develop and implement intercept plans. Thus the mission requires a great deal of intelligence and speed.

Fortunately, virtually all of the component tasks involved in intercept planning can be automated to some degree. Threats can be classified using rule-based inference. Value models can be used to establish intercept priorities. Intercept vectors can be determined algebraically. However, the uncertainties introduced by jamming and deception cannot easily be processed by algorithm. In seeking to create a distributed-intelligence system for intercept planning, design efforts are focusing on development of a system concept in which human and computer intelligences will complement and support one another in the timely accomplishment of a very difficult task.

Only a few high-level decisions concerning the design of a new intercept planning system have yet been made. These include that:

- The computer must be programmed to accomplish all tasks that can be automated and to perform these tasks by default unless the operator intervenes. This feature is necessitated by the possiblity of a saturation raid by the enemy which would present the AEW operators with more threats than they could effectively handle manually.
- The operator must have access to all information appropriate for monitoring automated operations. A graphic display of the tactical situation must be provided incorporating a means for the representation of automated decisions. An easily operated query subsystem must be available to enable the operator to have quick access to a broad range of information at various levels of abstraction.
- The operator must have the option of making minimal interventions in

automated operations rather than having to choose between totally automated and totally manual control.

Continuing research is focusing on the identification of generic cognitive processes required for the accomplishment of intercept planning tasks and on the assessment of the relative capabilities of humans and computers to implement those processes. A collection of appropriate cognitive processes have been tentatively identified through the distillation of generic processes from a comprehensive functional task analysis. Man and machine processing capabilities are currently being evaluated in the context of situational extremes (including jamming and threat saturation scenarios) in order to determine how allocation of functions should be varied across situations in order to achieve optimal performance

Summary and Conclusions

Recent technological developments in artificial intelligence and hardware miniaturization have stimulated the evolution of advanced man-machine systems from "computer-aided" design concepts to "distributed-intelligence" concepts. In distributed-intelligence systems, humans and computers share the responsibilities for intelligent processing of information. This type of system requires new approaches to the problems of function allocation and man-machine interface design. Some general guidelines have been offered for the design of distributed-intelligence systems such as the requirement for two-way (i.e., man to machine and machine to man) interface subsystems for monitoring, querying, and intervening. Continuing research is required and is in progress to determine effective procedures for characterizing man and machine cognitive processing capabilities and for using those characterizations to achieve optimal allocations of intelligent responsiblities between human and computer components of a system.