AN AI APPROACH TO THE INTEGRATION OF INFORMATION

Technical Note No. 318

April 1984

By: Thomas D. Garvey, Program Director, AI Technology
John D. Lowrance, Computer Scientist

Artificial Intelligence Center
Computer Science and Technology Division
SRI International

This paper is a revised version of a talk entitled "Artificial Intelligence Applications to Electronic Warfare," presented to the 1983 DoD/AOC EW Technical Symposium, November 3, 1983, in Virginia Beach, Virginia.

333 Ravenswood Ave. • Menlo Park, CA 94025
(415) 326-6200 • TWX: 910-373-2046 • Telex: 334-486

ABSTRACT

This paper discusses the use of selected artificial intelligence (AI) techniques for integrating multisource information in the process of assessing and understanding an ongoing situation. The approach takes an active, "top-down" view of the task, projecting a situation description forward in time, finding gaps in the current model, and assigning sensors to acquire data to fill these gaps. Information derived from such sensors and other sources is combined by means of new, non-Bayesian inference techniques.

This active approach seems essential for solving the problems posed by the low-emission signatures anticipated for near-future threats. Simulation experiments lead to the conclusion that the utility of ESM system operation in future conflicts will depend on how effectively onboard sensing resources are managed by the system.

The view of AI that will underlie the discussion is that of a technology with the potential of extending automation capabilities from the current "replace human hands" approach to that of replacing or augmenting human cognitive and perceptual capabilities. Technology transfer issues discussed in the presentation are the primary motivation for emphasizing this view. The paper will conclude with a discussion of unresolved problems stemming from the introduction of AI technology into military systems.

I INTRODUCTION

Artificial Intelligence (AI) techniques are beginning to be exploited in support of a variety of intellectual endeavors. It is appropriate to consider using them to solve difficult problems in electronic warfare. Here we discuss an application of AI techniques to the problem of integrating information from diverse sources for the purpose of understanding a developing threat situation. The work discussed here was drawn from a number of projects that have been conducted during the past several years in SRI's Artificial Intelligence Center with support from the Office of Naval Research and other government agencies. The techniques used to confront this electronic-support-measure (ESM) problem are relevant to any task in which information about a situation is provided by *evidential* sources or sensors that have an incomplete view of a situation, may be uncertain in their determinations, and, in some cases, may be incorrect.

The advantages of being able to combine evidential information intelligently include lowered false alarm rates, reduction in ambiguities, robust identifications, and improved overall system operation. Applications of the work include threat warning, ESM collection, C^3I support, indications and warning (I&W), and sensor management. The specific methods used belong to an area of AI called *evidential reasoning*. The work itself derives from the field of *knowledge-based systems*.

II EXPLOITATION OF INFORMATION

The goal of this work is to compile an air-defense order of battle (OB), determining the identity of threats with an acceptable degree of confidence, locating threats within a minimal region of uncertainty, and determining the status of a threat during the time of interest (when the sensor platform is in its vicinity). The difficulties involved with detecting, locating, and countering current and projected threat systems will require the exploitation of all available sources of relevant information. The degree to which we can efficiently organize the use of our information sources to acquire a comprehensive model of a developing situation will determine how successfully we can anticipate threats and organize our response. The situation model we develop is used as an aggregating medium for collecting relevant information, while the current model in turn enables us to organize and focus our efforts for gathering new data.

There are numerous difficulties inherent in this task: (1) threat systems have diminishing signatures -- certain threats may operate entirely passively, acquiring their target data from remote sites; (2) weapon systems will emit for the minimum possible time, often only after a weapon is already launched; (3) many threats will be agile in several parameters, therefore difficult to correlate and track; (4) finally, many threat systems will use regions of the electromagnetic spectrum outside the range of traditional RF intercept systems, thereby requiring alternative means for their detection.

It is, therefore, imperative that threats be anticipated prior to exposure. Furthermore, because of the potentially high data rates, sensor resources must be managed effectively so as to optimize the collection of vital information, while, at the same time, irrelevant data are filtered from the stream. Finally, to reduce false alarms and ambiguities, we must have methods for combining threat information that exploit the associations and correlations of many types of parameters.

There are several capabilities that appear to be crucial to the task of situation assessment. The first of these is the ability to relate our understanding of the situation to the satisfaction of system goals. Mission requirements (for example, the need to acquire information about the upcoming threat situation) establish these *goals* for the system, which, in turn, drive the system functions. Next we must be able to correlate current information with models of situations and the players in those situations. For example, a hot spot in a FLIR image might be related to a developing model of an antiaircraft artillery site where RF emissions have already been intercepted. Since available information is likely to be ambiguous, it is important to be able to maintain competing explanations or *hypotheses.* Comparison of competing hypotheses provides a means for determining just which information could help resolve the ambiguity. Gaps in our understanding will be detected by noticing which "slots" in a situation model are not filled in. It is critical to understand what we *do not* know about a situation as well as what we do know. Finally, we must be able to combine and draw effective conclusions from information known to be inexact, incomplete, and possibly incorrect.

We must capitalize on our understanding of a given situation by using the pertinent information to manage scarce resources, particularly our sensors. Since a threat may be emitting for only a brief period, it is important to make sure that an appropriate sensor is monitoring likely regions of the spectrum during the time of expected emission. Similarly, passive threats may need to be actively sought out, with either active or passive means used for sensing. The effective control of resources requires the ability to reason over time and space so as to plan for resource needs, recognize potentially harmful interactions, and make estimates regarding the utility of resources in the current situation.

True understanding of a situation presupposes the ability to project the scenario into the future, thereby anticipating likely events. A view of upcoming events is much more important than full understanding of what's happening now. To make these projections, we need to be able to represent and understand *dynamic* processes as well as static situations. In addition, we must recognize the effects of current uncertainties upon future projections. Typically, present uncertainties translate into greater future uncertainties until, after a certain point, the number of possible events becomes so large that any meaningful estimates are rendered impossible without further information.

The simple architecture shown in Figure 1 represents our basic assessment loop. This architecture emphasizes ***active*** acquisition of information. Instead of waiting for information to come pouring in and then trying to sort through it, the system actively seeks out high-value information in a top-down fashion. Such a focused approach is essential for coping with data overload and keeping it under effective control.

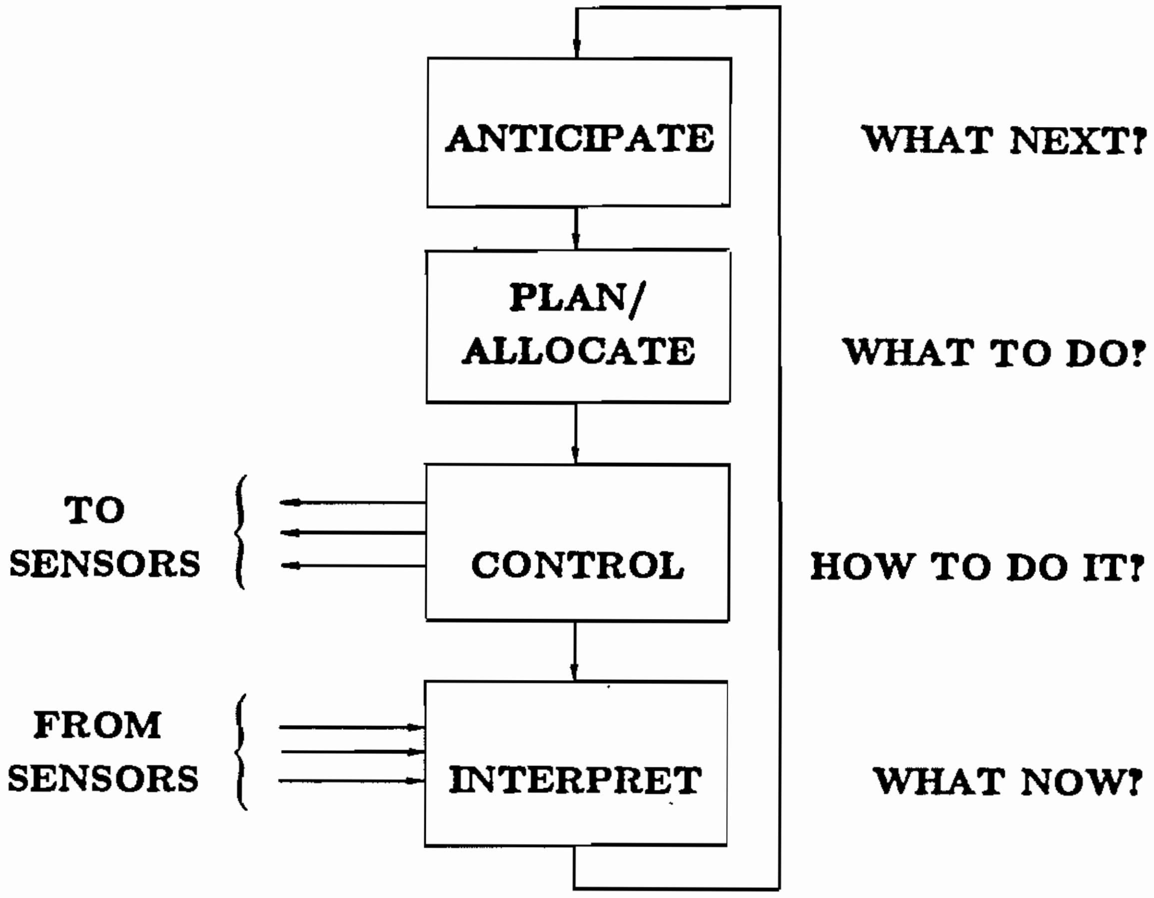

Figure 1: THE BASIC LOOP

The basic loop shown is applicable to a number of information integration applications. The process starts by using current information to ANTICIPATE prospective significant events. The system is in effect asking "What will happen next?" Information used in this operation includes process models for actions such as threat operation sequences, and known associations of threats derived from knowledge of typical deployments. Essentially this module attempts to hypothesize possible but as yet unseen threats for which there are no current sensor data. This list of hypothesized threats is passed to the next module.

The PLAN/ALLOCATE module determines "What to do next" by examining the list of possibilities, ordering the list according to the system's current requirements (for threat warning, the system uses an estimate of the lethality of the threat in the current environment), examining models of its sensor resources to determine which sensors can provide the necessary data, and then selecting an optimal allocation of sensors to possible threats. These sensors are assigned to collect specific information about not only the selected threats, but also any others that might fall within their purviews.

The CONTROL module determines what data the sensors need to carry out their tasks. This could include pointing data for optical and electro-optical sensors, tuning data for intercept receivers, or operating programs that specify parameter ranges and dwell times or statistics for computer-driven receivers.

The INTERPRET module takes sensor reports, compares them with the model of the current situation, and updates the model on the basis of the new information. This could lead to a report of a new threat, elimination of an earlier false alarm, disambiguation of an earlier report, or a stronger belief in the presence of a threat already in the model. The updated situation model provides the basis for the next ANTICIPATION step, and so the process iterates.

Carrying out the operations just described requires access to a variety of information. Indeed, we would like to be able to represent and access *all* information that may be relevant. Requirements for the situation model include the set of *players* (typically the threat systems), their interactions with one another, and their capabilities, tactics, doctrine, and operating procedures. We use strongly structured representations for this information, an example of which is shown in Figure 2.

One critical piece of information is a description of the processes and operations that represent typical threat behavior. This provides the means of anticipating future behavior on the basis of a current understanding of the threat's activity. These processes and operations are represented by process models depicting allowable states (for example, a state might represent a

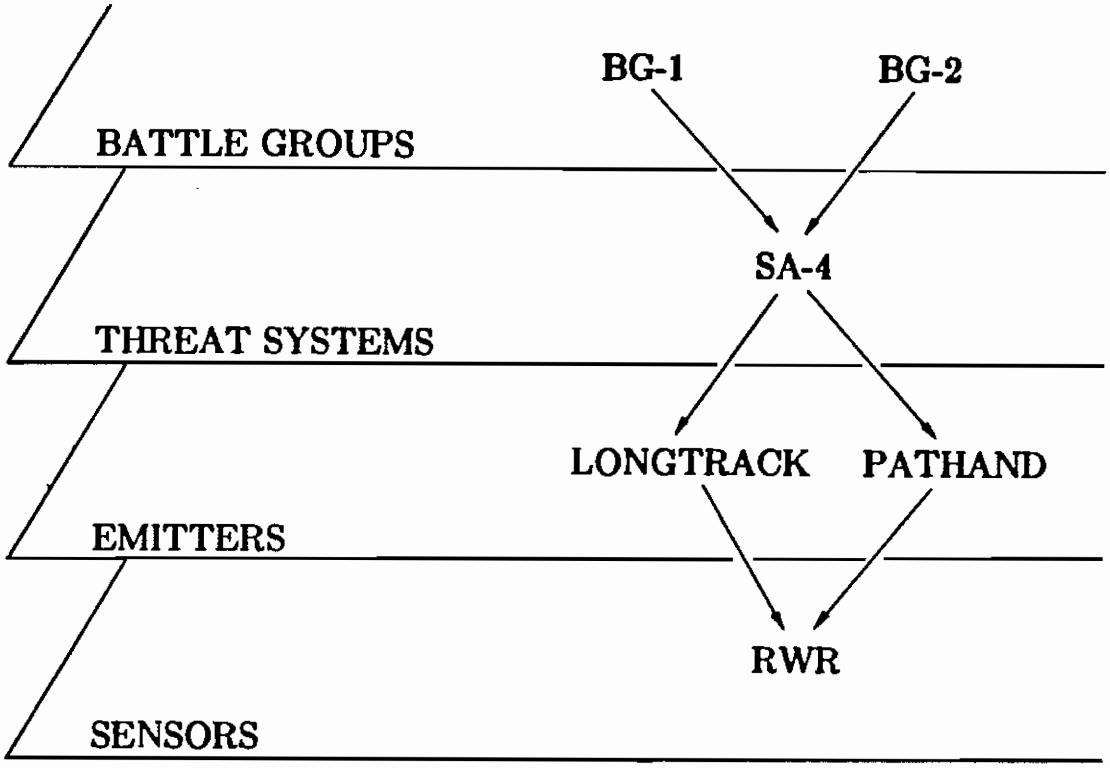

Figure 2: TYPICAL SITUATION STRUCTURE

particular SAM in target-tracking mode), the new states that each state could pass to, the conditions for making such a transition, and any observable phenomena associated with the state, to which our sensors might respond.

We must be able to represent the capabilities of our information sources, particularly our sensors. Extended probability tables are currently used to store this information; entries in the table indicate the possible reports a sensor could produce in response to a specific observable. These tables provide the means for estimating the effectiveness of selected sensors for acquiring desired information.

The final class of information that is crucial to our understanding of situations is environmental effects. These include terrain constraints on mobility and siting, atmospheric effects on both threat system sensors and for our own onboard resources, and effects of factors such as pulse densities and both hostile and friendly ECM. Two important ways in which environmental information is used are to lower our expectations of sensor performance because of degrading environmental influences (for example, an RWR will be generally less effective, in high-pulse-density environments) and to modify our estimates regarding the current state of a threat system (for example, we would not anticipate the use of a weapon's optical target-tracking mode under IFR conditions).

Figure 2 shows a fragment of generic information about an SA-4 system stored in the WORLD KNOWLEDGE BASE. This model represents the SA-4's range of capabilities in general, but says nothing about a specific example or *instance* of the system -- these are stored in the situational model. This model is a *layered* representation that relates SENSOR REPORTS to EMITTERS to THREAT SYSTEMS -- and could be extended to include BATTLE GROUPS. It shows that a particular report from an RWR could be evidence for the target-tracking mode of a PATHAND radar. This operating mode is in turn linked to the SA-4 system (which has an associated *process model* in the THREAT SYSTEM plane). The model also describes the use of a LONGTRACK radar for providing the initial acquisition information. The advantages of this type of model are numerous, but one especially important benefit is the ability to propagate information around the structure; it thus uses information about our current belief in the LONGTRACK, for example, to update our belief in the presence of the SA-4. The representation also makes it possible to extract information from the data structure that is commensurate with a user's needs. For example, important information for selecting ECM techniques may reside at the EMITTER level, from which an ECM system could then draw information. The crew member may be interested in the specific threat systems present and may therefore extract information from that level. The intelligence analyst may be interested in groupings, mobility vectors, and probable objectives; accordingly, he may draw information from the BATTLE GROUP level.

The situation model (SM) is the structure that stores our current situation information. An entry in the SM is created for each potential threat by copying relevant information from the world knowledge base whenever newly obtained data are not explained by anything already in the SM. The specific data recorded in the SM include mode information, time stamps, and location data. SM also tracks competing possibilities; additional data may support the existence of one or more of these.

The SM is the first place the system checks to begin assessing the import of newly acquired information. If the latter is consistent with a threat in the SM, it is combined (at the appropriate level) with the information already present. This integration is performed by a method known as Dempster's rule, and yields interval measures of belief associated with each node in the structure. The results are then propagated throughout the model. Once again the layered representation provides an important capability: information may be entered at a level consistent with the source. For example, a pilot report might provide information concerning an existing threat system and be entered at that level, while new receiver data can be merged with this pilot report by entering it at the SENSOR REPORT level and propagating the combined information to the THREAT SYSTEM level.

With these models as background, the integration procedure as shown in Figure 3 is relatively straightforward. Observables are converted into reports by the sensors; these reports are checked against the SM to see whether they are consistent with current information. If the new reports match an expected threat, the data are combined and the results propagated throughout the structure; if the report represents information about a previously unknown threat, the proper information is extracted from the world knowledge base and entered into the SM, where it will be combined with future data.

Let us reiterate several advantages of this approach. The combination of information occurs at a level consistent with the source and type of information available. Information may be extracted at a level in keeping with the user's requirements. Furthermore, since the combination rule and the representation of belief is non-Bayesian, the system can represent world knowledge more faithfully

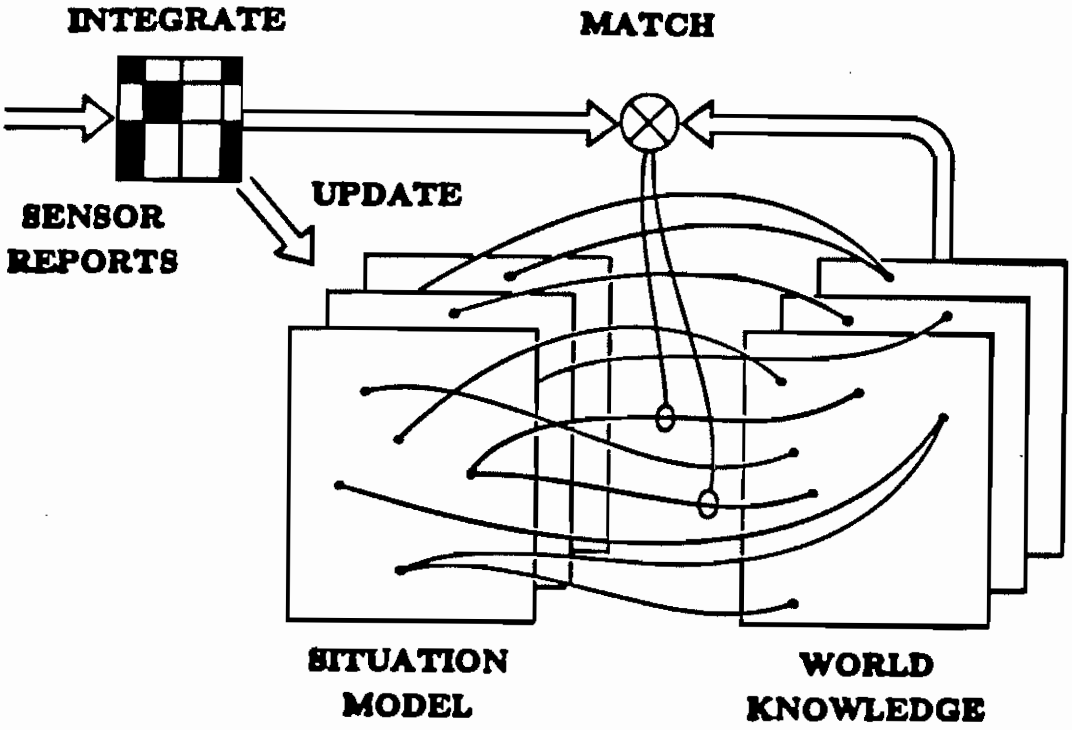

Figure 3: THE INTERPRETATION PROCESS

than would otherwise be the case. In particular, the representation encodes ignorance explicitly, thereby enabling us to distinguish between ignorance and disbelief. This is crucial in determining, for example, whether a small likelihood associated with a threat system is due to our having looked for it and not seen it (disbelief) or to our not having looked, yet having no particular reason to believe it is there (ignorance). A standard Bayesian system cannot distinguish between these two important cases.

III SIMULATION RESULTS

To verify the principle discussed in this paper, we have constructed a demonstration simulation system. The simulator flies the aircraft along a preselected path over defended hostile territory, operates threat systems encountered by the aircraft en route, and determines the reactions of sensors to threat system observables. Threats are selected from a dozen different types, with a typical scenario involving possibly 20 to 50 different systems. The sensor suite is selected from approximately ten different types of sensors. Environmental conditions may be varied to stress the sensors or modify the scenario. Results of operations, along with ground truth, are shown on a color graphics display (an example may be seen in Figure 4). A scoring technique enables us to relate the effectiveness of the system at acquiring necessary threat data in a timely manner from one scenario or operational mode to the next. In particular, we can turn off AI-based information integration and planning so as to be able to compare the operation of the AI system with a more conventional approach.

The results of many simulations indicated that we could generally improve scores by adding more situation and threat knowledge to the system. This knowledge was used by the system to optimize its control of available resources. Adding sensors in either AI or conventional mode tended to produce increases in scores, with one interesting exception -- if the sensor was likely to produce a high rate of false alarms, the overall system score in conventional mode tended to decrease because of an inability to remove these errors. The AI-based system, however, could access other sources of data to verify a threat and thereby eliminate most false alarms. Similarly, the system could resolve ambiguities in AI mode by assigning sensors to acquire the specific data needed. Finally, the AI system tended to degrade gracefully as environmental conditions worsened because it could supplement poor information from one source with information derived from other modalities, thereby reducing reliance upon any single sensor.

The keys to success of the AI-based approach are the active approach to acquisition of specific information and the effective management of scarce sensor resources. These are made possible by the ability to anticipate events sufficiently in advance of their actual occurrence, thus maximizing the certainty of a timely response.

IV UNRESOLVED ISSUES

We have described one approach to information integration that is based on AI technology. The ultimate promise of this technology will probably be determined as much by the pragmatics of applications as by the capabilities of the technology itself. Some of the issues that must be resolved are the following: how to evaluate the true effectiveness of a knowledge-based system for EW operations; how large the knowledge base must be for effective operation; how to achieve the processing speeds necessary to cope with threat time lines; what languages to choose for development and implementation; how to deal with the difficulties of maintaining, debugging, modifying, and extending the program.

Because we regard these as significant issues, and because the information-processing technologies are very new, we believe that the introduction of these technologies into fielded EW equipment will be a gradual process. This view is in keeping with the concept of AI technology as an extension of current approaches to automation. While this means that AI cannot be considered as either short-term magic or as a panacea, it is nevertheless our firm conviction that, over the next several years, AI technology will completely revise our approaches--not only to EW problems, specifically, but to information processing in general.

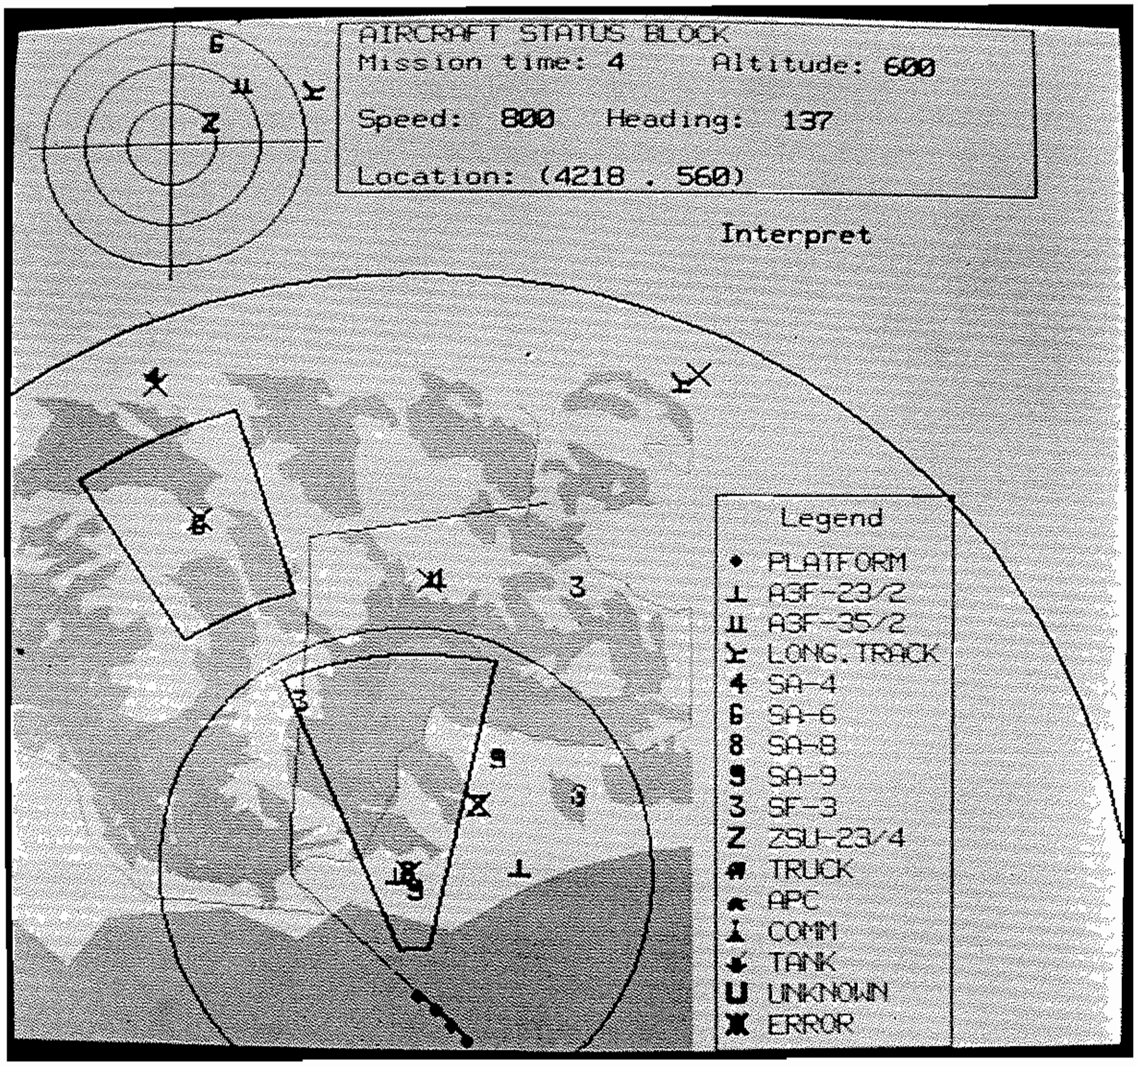

AIRCRAFT STATUS BLOCK

Mission time: 4 Altitude: 600

Speed: 800 Heading: 137

Location: (4218 . 560)

Figure 4: SCREEN OUTPUT FROM SIMULATOR RUN

Addendum--WHAT IS ARTIFICIAL INTELLIGENCE?

There are numerous definitions of artificial intelligence (AI), none completely satisfactory. For this discussion, it is convenient to consider two perspectives: the "scientific view" and the "engineering view." In the scientific view, AI work is seen as a study of intelligence generally, but especially with an eye to uncovering those computational structures and algorithms that are necessary for modeling or demonstrating intelligent behavior. Intelligence usually implies the ability to acquire information directly from the environment, to understand that information in the light of goals and desires, and then to act on the basis of that information so as to influence or change the state of the environment, thereby completing the perceive-interpret-decide-act loop. According to this view, there are several functional capabilities that are normally considered essential to intelligent behavior. These include the ability to perceive events in the world by means of the senses, the ability to abstract events so that problem-solving and planning activities may take place, the ability to convert plans into actions, the ability to communicate with other entities, and the ability to learn and adapt. AI researchers are addressing every one of these issues and, in some instances, making rapid progress towards emulating the desired capabilities.

In the engineering view, AI represents the logical end to current trends in automation. We have been extraordinarily successful in automating a variety of capabilities, especially for sensing raw information and in manipulating the local environment. The extension of automation to intellectual and cognitive tasks is seen as a prospective role for AI. The engineering view is appropriate for applications of AI to electronic warfare problems, as it emphasizes a more gradual infusion of technology--i.e., *incremental* improvements in capabilities. Among the important objectives to consider from this perspective are better ways of sensing the environment, techniques for integrating sensed information to arrive at a comprehensive understanding of the current situation, and the planning and execution of constrained responses to the environment.