PERCEPTUAL REASONING IN A HOSTILE ENVIRONMENT

Thomas D. Garvey and Martin A. Fischler

Artificial Intelligence Center
SRI International, Menlo Park, CA 94025

ABSTRACT

The thesis of this paper is that perception requires reasoning mechanisms beyond those typically employed in deductive systems. We briefly present some arguments to support this contention, and then offer a framework for a system capable of perceptual reasoning, using sensor-derived information, to survive in a hostile environment. Some of these ideas have been incorporated in a computer program and tested in a simulated environment; a summary of this work and current results are included.

I INTRODUCTION

Living organisms routinely satisfy critical needs such as recognizing threats, potential mates, food sources, and navigable areas, by extracting relevant information from huge quantities of data assimilated by their senses. How are such "relevant" data detected?

We suggest that a reasoning approach that capitalizes on the goal-oriented nature of perception is necessary to define and recognize relevant data. Perception can be characterized as imposing an interpretation on sensory data, within a context defined by a set of loosely specified models. The ability to select appropriate models and match them to physical situations appears to require capabilities beyond those provided by such "standard" paradigms as logical deduction or probabilistic reasoning.

The need for extended reasoning techniques for perception is due to certain critical aspects of the problem, several of which we summarize here:

* The validity of a perceptual inference (interpretation) is determined solely by the adequacy of the interpretation for successfully carrying out some desired interaction with the environment (as opposed to verification within a "closed" formal axiomatic system).
* Since it is impossible to abstractly model the complete physical environment, the degree to which purely abstract reasoning will be satisfactory is limited. Instead, perception requires tight interaction between modeling/hypothesizing, experimenting (accessing information from the environment), and reasoning/verifying.
* Reasoning processes that embody concepts from physics, geometry, topology, causation, and temporal and spatial ordering are critical components of any attempt to "understand" an ongoing physical situation. Explicit representations appropriate to these concepts are necessary for a perceptual system that must provide this understanding. These representations are incommensurate and it is not reasonable to attempt to force them into a single monolithic model.
* There is typically no single, absolutely correct interpretation for sensory data. What is necessary is a "maximally consistent" interpretation, leading to the concept of perception as an optimization problem [1, 2] rather than a deductive problem.

Research in perception and image processing at SRI and elsewhere has addressed many of these issues. An early effort focused upon the goal-directed aspect of perception to develop a program capable of planning and executing special-purpose strategies for locating objects in office scenes [3]. Research addressing interpretation as an optimization problem includes [1, 2, 4]. Current research on an expert system for image interpretation [5] has also considered the strategy-related aspects of determining location in situations involving uncertainty.

The most recent work (at SRI) on perceptual reasoning has addressed the problem of assessing the status of a hostile air-defense environment on the basis of information received from a variety of controllable sensors [6]. This work led us to attempt to formulate a theory of perceptual reasoning that highlighted explicit reasoning processes and that dealt with those aspects of perception just described. In the following section, we will use this work as a vehicle to illustrate a paradigm for perceptual reasoning.

----

* This work was supported by the Defense Advanced Research Projects Agency under Contracts MDA903-79-C-0588, F33615-77-C-1250, and F33615-80-C-1110.

II PERCEPTUAL REASONING IN A SURVIVAL SITUATION

The specific problem addressed was to design a system able to interpret the disposition and operation (i.e., the order of battle or OB) of hostile air-defense units, based on information supplied by sensors carried aboard a penetrating aircraft [6]. The situation may be summarized as follows. A friendly aircraft is faced with the task of penetrating hostile airspace en route to a target behind enemy lines. Along the way, the aircraft will be threatened by a dense network of surface-to-air missiles (SAMs) and antiaircraft artillery (AAAs). The likelihood of safe penetration and return is directly related to the quality of acquired or deduced information about the defense systems.

Partial information is furnished by an initial OB, listing known threats at, say, one hour before the flight. Additional knowledge is available in the form of descriptions of enemy equipment, typical deployments, and standard operating procedures. Since the prior OB will not be completely accurate, the information must be augmented with real-time sensory data. The OB forms the starting point for this augmentation.

The explicit goal of the overall system is to produce and maintain an accurate OB, detecting and identifying each threat prior to entering its lethal envelope. The density of threats means that this goal will result in conflicting subgoals, from which selection must then be made to ensure that critical data will be received. This must be accomplished by integrating data from imperfect sensors with prior knowledge. The paradigm that was developed for this task is summarized below:

(1) Available knowledge is used to create an hypothesized OB that anticipates the developing situation.

(2) A plan that attempts to allocate sensors to detect or verify the presence of threats, in an optimal way, is constructed. Sensors are then allocated and operated.

(3) Information returned from the sensors is interpreted in the context established by the anticipated situation. Interpretations modify the current OB, and the process is iterated.

We will briefly discuss each of these steps.

A. Anticipation (Hypothesis Formation; Model Selection)

In the anticipation phase, the system calls upon a variety of types of internally available knowledge to hypothesize the current OB. The specific goal of this step is to determine where critical knowledge is lacking and to produce a set of information requests; these will become goals for the resource allocation procedure in the experimental planning phase. The primary knowledge source is the initial OB, which lists the geographic locations of all threats known prior to flight time. Since these threat systems are mobile and the situation is changing rapidly, the system determines probable new locations for each threat by comparing data about the mobility of weapon vehicles with a terrain map of the area and computing "uncertainty regions" likely to contain the threats.

Partial information about systems of threats (i.e., collections of weapons deployed in mutually supporting ways) is used by a program that performs "geometric reasoning" to predict the location of missing components. The anticipation phase attempts to reduce the difficult problem of detecting all threats to the simpler problem of verifying hypothesized threats.

B. Experimentation (Accessing Information from the Environment)

The goal of this step is to access information needed to detect or verify the presence of threats inferred in the anticipation step, but not available in the "internal" knowledge base of the system. In general, it might be necessary to define and execute one or more experiments to extract this needed information from the environment. In the more limited context of model instantiation by "passive" sensing, the problem reduces to that of allocating sensor resources to maximize the overall utility of the system; sensing is a specific instance of the more general process of experimentation.

First the list of data-acquisition goals is ordered, based on the current state of information about each threat and its lethality. The allocator attempts to assign (a time sliced segment of) a sensor to satisfy each request based on the expected performance of the sensor for that task.

Sensor detection capabilities are modeled by a matrix of conditional probabilities. These represent the likelihood that the sensor will correctly identify each threat type, given that at least one instance thereof is in the sensor's field of view. This matrix represents performance under optimal environmental conditions (for the sensor) and is modified for suboptimal conditions by means of a specialized procedure. This representation is compact and circumvents the need to store complete, explicit models describing sensor operation in all possible situations. Similar models describe each sensor's identification and location capabilities.

The sensor models are used to compute the utility of allocating each sensor to each of the highest priority threats. These utilities form the basis for the final allocation, which is carried out by a straightforward optimization routine. At the same time, the program determines how the sensor should be directed (for example, by pointing or tuning). Appropriate control commands are then sent to the simulated sensors.

C. Interpretation (Hypothesis Validation; Model Instantiation)

In this phase, the program attempts to interpret sensor data in the context of threats that were anticipated earlier. It first tries to determine whether sensor data are consistent with specifically anticipated threats, then with general weapon types expected in the area. Since sensor data are inherently ambiguous (particularly if environmental conditions are suboptimal), this step attempts to determine the most likely interpretation.

Inference techniques used for interpretation include production rule procedures, probabilistic computations, and geometric reasoning. Production rules are used to infer probable weapon operation (e.g., target tracking, missile guidance), on the basis of such information as past status, environmental conditions, and distance from the aircraft. Probabilistic updating of identification likelihoods is based on the consistency of actual sensor data with expected data, and on agreement (or disagreement) among sensors with overlapping coverage. Geometric reasoning introduces a concept of global consistency to improve identification by comparing inferred identifications and locations of threat system components with geometric models of typical, known system deployments.

The interpretation phase brings a great deal of a priori knowledge to bear on the problem of determining the most likely threats the sensors are responding to. This results in much better identifications than those produced by the sensors alone. Confident identifications are entered into the OB and the entire process is continued.

D. Performance

An experimental test of the system, using a simulated threat environment, allowed a comparison between two modes of operation--an "undirected" mode and one based on perceptual reasoning. A scoring technique that measured the effectiveness with which the system detected, identified, and located hostile systems in a timely fashion was used to grade performance. The ability of the perceptual reasoning system to use external knowledge sources effectively, and to integrate information from multiple sensors, produced superior capabilities under this measure. These capabilities showed themselves even more prominently in situations where environmental conditions tended to degrade sensor performance, rendering it critical that attention be focused sharply.

III DISCUSSION

Our approach to perceptual reasoning suggests that the problem of perception actually involves the solution of a variety of distinct types of subproblems, rather than repeated instances of the same general problem. The system we described utilizes a nonmonolithic collection of representations and reasoning techniques, tailored to specific subproblems. These techniques include both logical deduction and probabilistic reasoning approaches, as well as procedures capable of geometric reasoning and subjective inference.

We have discussed several key aspects of the general problem of perceptual reasoning, including the assertion that perception is goal oriented, and inductive and interpretative rather than deductive and descriptive; that because complete modeling of the physical world is not practical, "experimentation" is a critical aspect of perception; and finally, that multiple representations and corresponding reasoning techniques, rather than a single monolithic approach, are required.

The specific system discussed above constitutes an attempt to address the reasoning requirements of perception in a systematic way and, to our knowledge, represents one of the few attempts to do so. While systems that truly interact with the physical world in an intelligent manner will certainly assume a variety of forms, we believe they will all ultimately have to resolve those aspects of the problem that have been described here.

REFERENCES

1. M. A. Fischler and R. A. Elschlager, "The Representation and Matching of Pictorial Structures," IEEE Transactions on Computers, Vol. C-22, No. 1, pp. 67-92 (January 1973)

2. H. G. Barrow and J. M. Tenenbaum, "MSYS: A System for Reasoning About Scenes," Artificial Intelligence Center Technical Note No. 121, SRI International, Menlo Park, California (1976)

3. T. D. Garvey, "Perceptual Strategies for Purposive Vision," Artificial Intelligence Center Technical Note No. 117, SRI International, Menlo Park, California (1976)

4. A. Rosenfeld, "Iterative Methods in Image Analysis," Pattern Recognition, Vol. 10, No. 4, pp. 181-187 (1978)

5. R. C. Bolles et al., "The SRI Road Expert: Image-to-Database Correspondence," Proc. DARPA Image Understanding Workshop, Science Applications, Inc., Report No. SAI-79-814-WA, pp. 163-174 (November 1978)

6. T. D. Garvey and M. A. Fischler, "Machine-Intelligence-Based Multisensor ESM System," Technical Report AFAL-TR-79-1162, Air Force Avionics Laboratory, Wright-Patterson AFB, OH 45433 (Final Report DARPA Project 3383) (October 1979)