A Knowledge-Based Decision Aid for Enhanced Situational Awareness

Dan R. Ballard
Reticular Systems, Inc.
Lisa Rippy
U.S. Army AATD

Abstract

This paper describes a multi-year research and development effort to develop a system for performing situation assessment in next-generation Army helicopters. We first provide a formal definition of situation assessment and describe the motivation for the architecture based on studies in human cognition and attention. The paper describes the overall architecture and the processing paradigm used in performing situation assessment. In particular, we show how extensive knowledge about the battlefield, the threat, terrain, enemy and friendly doctrine can be used to aid in performing situation assessment. We also show how the overall inferencing process can be controlled in such a way as to bound the requirements for scarce computational resources. We describe a system composed of three independent reasoning subsystems performing recognition, evaluation and prediction. We also describe the knowledge bases and important data structures used in developing the system.

Introduction

Significant improvements have been made in modern rotorcraft avionics and communications systems. Advanced on-board sensors provide the pilot with a much improved picture of the aircraft's environment. Modern battlefield information and communications systems can deliver real-time and near-real-time information about many aspects of the battlefield. In addition, the aircraft and its weapons systems have become considerably more capable and more complex. Finally, the aircraft missions are more complex since they often occur at night and in adverse weather and require extensive crew coordination during target acquisition, hand-off, etc.

All of these changes have resulted in increases in the pilot's cognitive workload. In battlefield situations, pilot workload reaches a peak at the same time that large amounts of sensor and battlefield information are being presented to him. In such a condition a pilot can suffer from cognitive overload; he becomes too busy to recognize and attend to all important information and events that have the potential to affect his aircraft and his mission.

The Army has initiated a number of R&D efforts to address these issues. Some of these efforts have included the Advanced Rotorcraft Technology Integration (ARTI) program, an improved Obstacle Avoidance System (OASYS), the Air-to-Air Mission Equipment Package/Weapons Demonstration (AAMWD), the Advanced Pilotage Sensor Program (APSP), and the Combined Arms Tactical Command and Control Demonstration (CATC2D). The focus of the Aviation Applied Technology Directorate (AATD) efforts has been on the development of associate systems and cognitive decision-aiding systems. AATD programs include the Day/Night Adverse Weather Pilotage System (D/NAPS), the Weapons and Targeting Expert System (WATES) and the Rotorcraft Pilot's Associate (RPA) program.

This paper describes research, development and implementation of a Situation Assessment system for use aboard Army helicopters. This work was performed by Reticular Systems, Inc. under the sponsorship of AATD [1-5]. The focus of this research was to develop Situation Assessment (SA) technology that will be useful for a wide variety of applications including command and control centers, autonomous vehicles, avionics systems and other associate and cognitive decision-aiding systems. Indeed, the technology is extremely useful where any agent (human or software) must make complex decisions in a dynamic, time-constrained environment.

System Architecture

An SA system must monitor the external environment for entities of interest, recognize those entities and then infer high-level attributes about those entities. In the case of the rotorcraft application, the SA subsystem must develop enhanced situational awareness in the aircrew by providing an accurate and coherent description of all external entities of interest (EEOI). An EEOI has the potential for affecting the planned rotorcraft mission. It is important to recognize that an EEOI is not necessarily a threat entity. Indeed, EEOIs must include friendly forces as well as threat forces and the SA system must be able to distinguish between the two as well as disambiguate "gray" threats (i.e., potential threats using the same kinds of weapons platforms as friendly forces).

The SA system uses data gathered by external environmental sensors such as radar, ESM, and the EO systems, and from intelligence updates provided over a number of communications links. This data is used to describe the external environment. The SA system can provide valuable feedback for the sensor control mechanisms and can be used to "focus" sensors on areas and objects of interest. In addition, the SA system uses extensive information about terrain, weather and the battlefield situation (e.g., locations of forces, FEBA description, FARP locations, political situation). This information is quasi-static (i.e, it does not change much over the course of a mission). Finally, in order for the system to perform the inferences required to develop an assessment of the current situation, it must utilize extensive knowledge about the EEOIs including knowledge about their capabilities, probable mission objectives, intentions, plans and goals. Knowledge from the pre-mission brief, intelligence updates, weather and terrain are all used for forming inferences about the battlefield situation. These low-level sources of information are used to infer EEOI status, probable mission objectives, intentions and activity. All of these elements combine to form a complete situational state description.

<u>System Functions</u>

The SA system consists of three major functions. These are a *recognition* function, an *assessment* function and a *prediction* function. These three functions work in concert to develop a situation state description. The recognition function develops a description of external entities of interest. This function provides knowledge sources for recognizing external entities by:

- developing environmental state descriptions
- computing geometric and algorithmic attributes
- performing formation and maneuver recognition

The assessment function performs five major activities. These include:

- generating a local situation description
- generating hypotheses about the external entities and their mission objectives
- developing hypotheses about the external entities and their intentions and plans
- characterizing the external entities of interest
- developing an Environmental State Description

The prediction function utilizes the current assessment of the situation and projects forward in time using a short-term horizon to estimate the external entities future position, course, intent and determine potentially dangerous or interesting events before they occur.

The prediction function provides the following major subfunctions:

- monitoring and updating the Environmental State Description
- monitoring posture of External Entity
- developing anticipated threat responses
- generating Predicted Situation Description

System Implementation

Figure 1 illustrates the basic architecture of the system. This architecture is based on the use of a blackboard model for problem solving. A blackboard model is made up of three separate components - a global database (the blackboard), multiple knowledge sources and a control mechanism [6, 7].

The blackboard is a global database that contains the problem-solving state data (objects from the solution space). Knowledge sources produce changes to the blackboard that lead incrementally to a solution to the problem. Communication and interaction among the knowledge sources takes place solely through the blackboard. The objects on the blackboard may be input data, partial solutions, alternatives, and final solutions (and, possible control data).

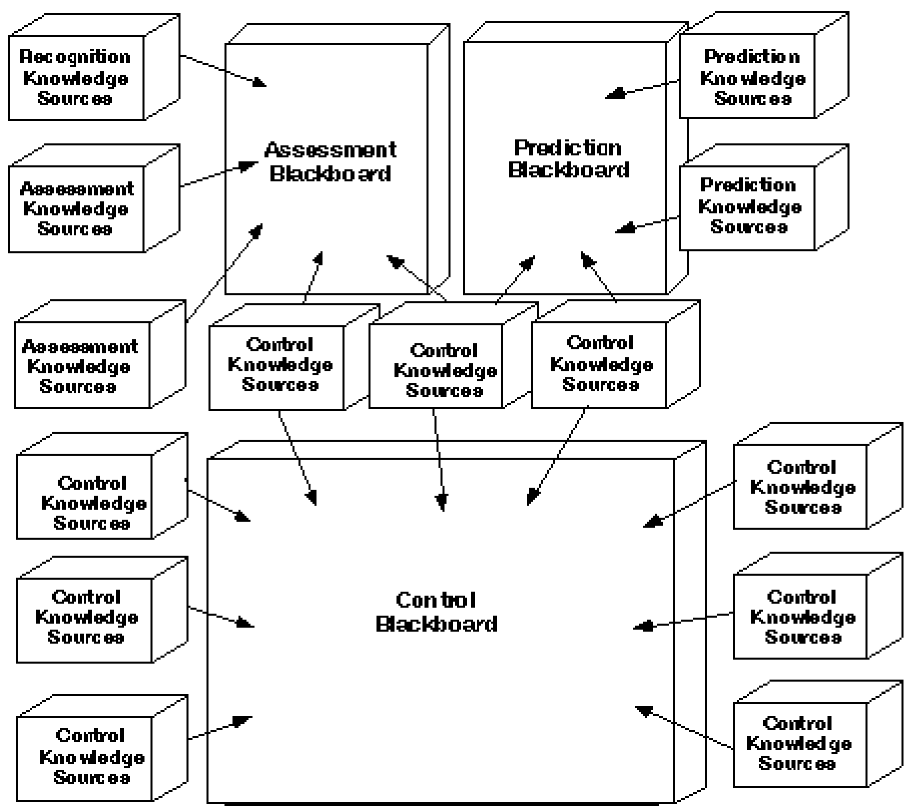

Figure 1. *Basic SA System Architecture*

The objects on the blackboard are hierarchically organized into levels of analysis. Information (properties) of objects on one level serve as inputs to a set of knowledge sources, which, in turn, place new information on the same or other levels. Blackboards typically have multiple blackboard panels. That is, the solution space can be partitioned into multiple hierarchies.

Knowledge sources are independent entities that have access to the contents of the blackboard. The knowledge required to solve a problem (the domain knowledge) is partitioned into these knowledge sources which are separate and independent.

Each knowledge source contributes information that will lead to a solution to the problem at hand. The knowledge sources modify only the blackboard or control data structures (which may also be on the blackboard) and only the knowledge sources are allowed to modify the blackboard. Each knowledge source is responsible for knowing the conditions under which it can contribute to a solution.

Control of the blackboard problem-solving process can be accomplished in a number of ways. The essence of the control problem is determination of when to apply a certain knowledge source and to what part of the blackboard. The control problem is one of determining which of the possible actions the system should perform at each point in the problem-solving process. In developing solutions to the control problem, a system must decide what problems it will attempt to solve, what knowledge will be used in attempting the solution and what problem-solving methods, techniques and strategies will be applied [8]. The SA implementation uses a separate control blackboard for controlling the problem-solving process.

Three blackboards are used in the SA implementation. An assessment blackboard is operated on by recognition and assessment knowledge sources. The prediction blackboard uses information from the assessment blackboard and also has independent knowledge sources posting and reading information on it. A control blackboard provides for overall control of SA problem-solving behavior.

The system was implemented using a reasoning under uncertainty mechanism based on Dempster-Shafer reasoning [9, 10]. Rules in the knowledge sources form a belief network with a measure of belief and disbelief. The reasoning under uncertainty mechanism enables the system to resolve data and inferencing contradictions. Also, the system is aware of the level of its own uncertainty and is less likely to produce errorful output.

The Assessment Blackboard

The assessment blackboard is divided into a six-level hierarchy. These levels are concerned with developing an environmental state description (i.e., a survey), characterizing the EEOI, interpreting EEOI plans, roles and intent and developing a summary description of the overall situation. Figure 2 illustrates the assessment blackboard hierarchy.

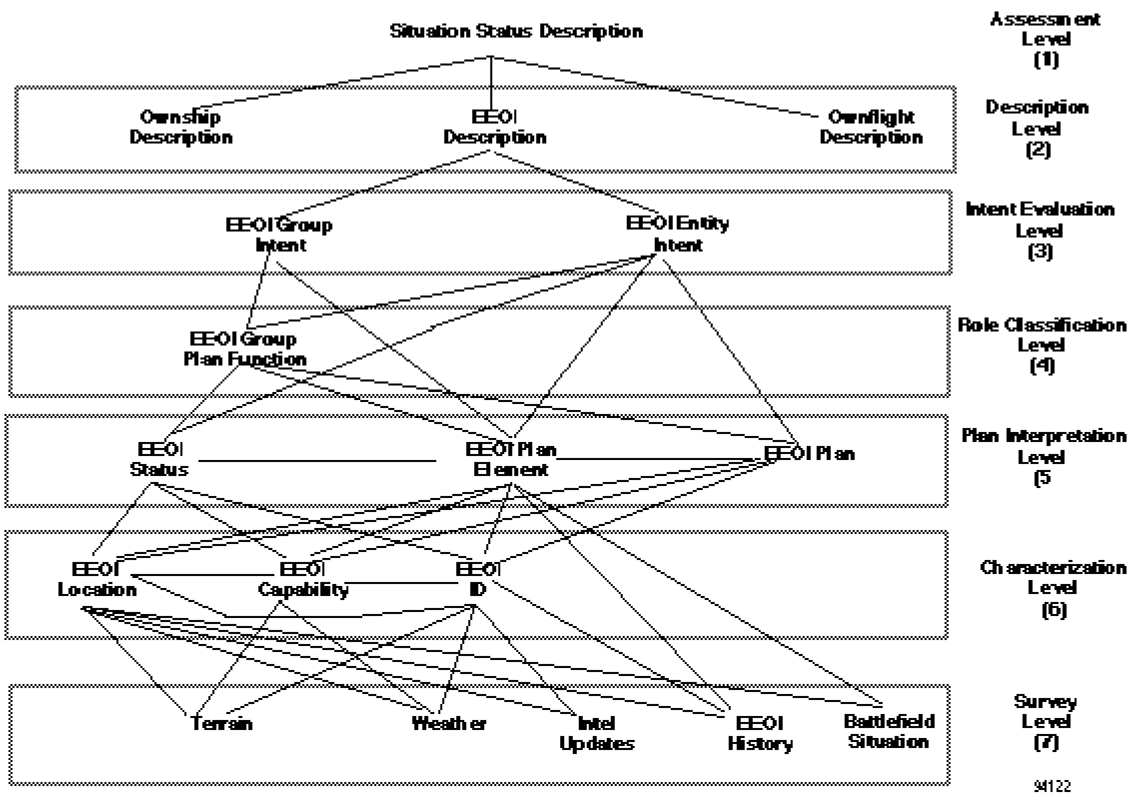

Figure 2. *The Assessment Blackboard Hierarchy*

<u>Survey Level</u>

The environmental state description is developed on the lowest level of the assessment blackboard. The environment is characterized by a number of different kinds of data including intelligence reports, terrain, weather, the battlefield description as well as updates from the sensors about entities of interest. Knowledge sources use this information to develop a description which characterizes the environment external to the aircraft. Note that this description utilizes information provided during the premission brief and information obtained during the mission.

<u>Characterization Level</u>

At the next level in the assessment blackboard hierarchy, the system develops an assessment of the location, identification and capability of each EEOI. *EEOI location* specifies the physical location of an EEOI at a specified instant in time. While sensor information or intelligence reports can provide a location in some instances, often the precise location can only be inferred using information about terrain, weather, EEOI history and the battlefield situation. *EEOI capability* reflects the assessed capability of an entity to affect ownship operations. Capability is defined in terms of three primary criteria - the ability of the EEOI to see ownship, the ability of the EEOI to move and the ability of the EEOI to react to ownship by either reporting (and thereby eliminating covert aspects of a mission) or engaging with hostile fire. *EEOI identification* describes the assessed identification of an external entity. Identification is based on comparison of recognized attributes with the attributes associated with a particular entity. Note, that only on rare occasions will it be possible for identification to be established by the sensor system. Normally, identification will require inferencing which considers terrain, intel reports, weather, battlefield situation, the EEOI history log and sensor data as well as higher-level inferences that recognize the plans of a particular entity.

Plan Interpretation Level

The levels above the characterization level address inferencing to determine the high-level attributes of an entity. Interpretation of a tactical situation can be defined in terms of problems in plan recognition. The SA system must interpret the activities of external entities by hypothesizing about their goals and inferring the plans that are being carried out in order to achieve these goals [11]. SA attempts to match the OPFOR plan by both observing single entities and plan elements as well as recognizing composite plans of multiple entities. The system attempts to match those things that are observed or inferred with plans that describe the behavior of an entity of interest.

It is in the Plan Interpretation level that the problem of situation assessment begins to be characterized as a problem in the AI domain of plan recognition. Within the Plan Interpretation Level, there are three main modules: *EEOI Status*, *EEOI Plan Element*, and *EEOI Plan*. These three modules work together to construct hypotheses on what the EEOI is currently doing. Figure 3 illustrates the Goal/Plan structure for the SA system.

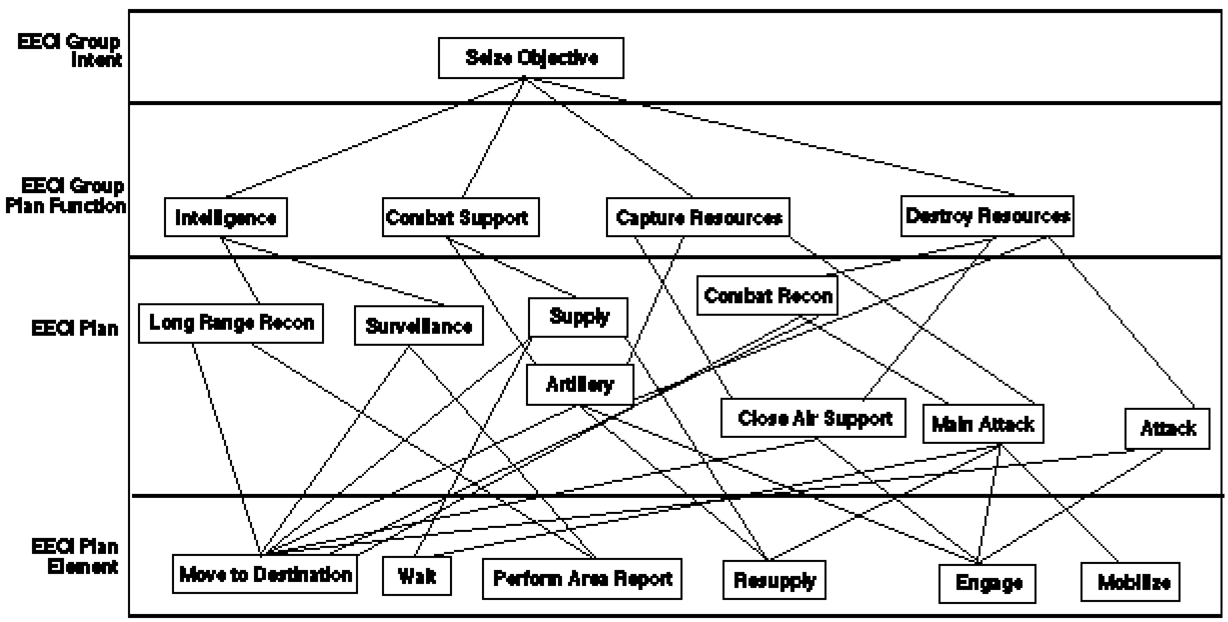

Figure 3. *Goal/Plan Structure*

Role Classification Level

The focus of the role classification level is on constructing an hypothesis about the plan function of the entity. A plan function can be considered as a description of what role the entity, in conjunction with other entities, plays in achieving some objective.

Structurally, a plan function consists of a set of plans which must be completed in order to satisfy the plan function. This construct is a multi-entity construct in that it specifies a set of single-entity plans.

This level is particularly important because here top-down hypotheses are combined with bottom-up data-driven inferences. At this level the system develops hypotheses about plan functions based on intelligence preparation of the battlefield, doctrine, and terrain. Plan functions are developed based on all factors except entity observations. If entity plans do not match the plans in the plan function hypothesis, the system will develop an hypothesis about a new plan function (and thus possibly new goals at the intent level) that account for the data. There are several situations where multi-agent hypothesis requests are made. These include situations when:

- ownship becomes aware of an EEOI that does not fit into any of the current multi-agent hypothesis objects
- ownship is flying into a geographic area where doctrine dictates the OPFOR might be present
- ownship is flying near key terrain
- a history log entry will cause a hypothesis object to be created if it does not fit into a current one

Thus the system establishes a set of expectations and then uses the entity observations to confirm or deny the hypotheses.

Intent Evaluation Level

The Intent evaluation level is the level of the assessment blackboard at which SA attempts to construct a hypothesis about the entity's goals. The entity's goals are used to hypothesize plan functions. *EEOI Group Intent* describes the goal of a group of entities which, by virtue of being a member of the group, the EEOI shares. This type of intent can be considered "global" in the sense that it extends beyond the external entity and is a goal which requires more than one entity and possibly more than one *EEOI Group plan* function. The *EEOI Entity intent* can be considered a "local" type of intent in that it deals with a single entity; it is the goal of a single entity.

Description Level

The Description level is where SA attempts to build an overall summary of ownship, the single EEOI, and ownflight based on information from the lower levels. This summary on a single entity is used in conjunction with other summaries on other entities to build the situation status description defined in the next level. We use Dempster-Shafer evidence combination rules to combine *EEOI posture*, *EEOI capability*, *EEOI predicted plan element*, and *EEOI predicted capability* to calculate an *Entity Assessment Rating* (EAR). The EAR is given by:

$$\text{EAR} = [\mathbf{bel}(\text{threat}),\ 1-\ \mathbf{bel}\ (\sim\text{threat})]$$

In addition to considering entities that SA has detected, SA hypothesizes about potential threats that might affect ownship's mission. The EAR represents the degree of threat that an EEOI either currently, or shortly in the future, represents to ownship. For instance, until ownship is almost within maximum weapons/detection range of the set of IDs hypothesized for the EEOI, the EAR belief function is [0,0]. The first number is the degree of confirmed danger to ownship, and the second number is the plausibility of danger. While still out of maximum weapons/detection range, the plausibility of danger is zero because the EEOI does not yet represent a threat. But as ownship comes into range, the plausibility of danger rises. The belief in danger remains at zero unless ownship actually has a confirmed entity sighting. This type of inferencing, based not on observations but on hypotheses, aids in the formation of probable enemy goals and *Global Plan Functions*. Finally, as sensors gather information about the area, the belief in danger and the plausibility of danger will converge about a single point that represents the danger of the EEOI to ownship.

The EAR uses a belief function to show not only the threat that an EEOI represents but uncertainty about the threat. This enables the system to efficiently filter information in several ways. The system must favor information about EEOIs that threaten ownship. In addition, it should favor information about EEOIs whose status is uncertain but have a high potential threat. The system also favors information about new EEOIs (i.e., EEOIs about which there is no previous knowledge). The system favors information about nearby EEOIs and rejects duplicate information. These heuristics are entity-specific and are posted at the policy level of the control blackboard.

Assessment Level
This level builds an assessment of the situation based on the combination of single entity descriptions developed in the Description level. The result of the Assessment level is an assessment of the situation at a particular time slice. This assessment is defined as a particular score attached to areas labeled as Dangerous, Cautious, Fair, or Safe.

Prediction Blackboard
The Prediction blackboard is used to reason about future entity location, capability, plan element, and intent. Additionally, future survey level data such as terrain descriptions, weather forecasts, and battlefield situation are also present on this blackboard. The levels that reason about location, capability, and plan element reason about a single entity. The Prediction blackboard does not make predictions about EEOI Group plan functions or EEOI Group intents. This blackboard is partitioned into five separate levels. Figure 4 illustrates the various levels on the prediction blackboard.

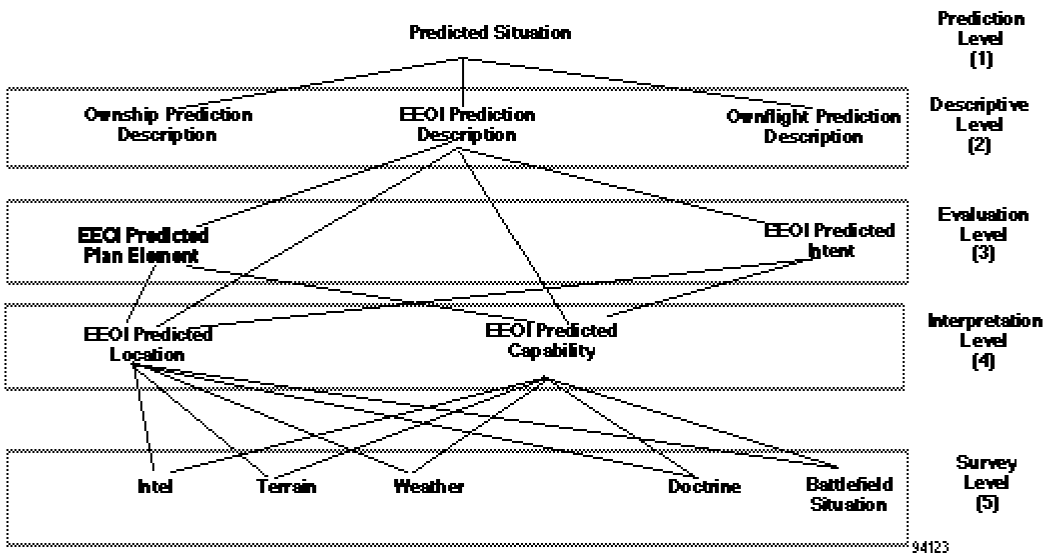

Figure 4. *Prediction Blackboard Hierarchy*

Survey Level
Predicted terrain description is a function of the predicted location of the EEOI. Other data items such as weather and much of the battlefield situation are assumed to be fairly static and unchanging from the current situation assessment.

Characterization Level
The prediction function is concerned with hypothesizing about probable alternatives not necessarily choosing a single one. Predicted location is perhaps the most important item being inferred on this blackboard. This is because most of the other objects depend on this inference. For example, predicted terrain description and EEOI predicted capability depend heavily on the EEOI's predicted location.

Once the EEOI's predicted location is inferred, predicted terrain description is simply recalled from the terrain database and EEOI capability is recalculated based on this general location and other factors. The Predicted location module uses information from the Assessment blackboard such as *EEOI Group plan* function and *EEOI Plan element* to infer *EEOI Predicted location*. The location predicted is not a simple set of x-y coordinates. Rather, it is an avenue of approach which SA infers that the EEOI will take.

<u>Evaluation Level</u>
This level attempts to build a hypothesis about what the EEOI is likely to be doing at a given future time. Two types of actions or plan elements are inferred. The first is *EEOI Predicted Plan Element* which is the inferred step in a plan that the entity will be performing. This predicted plan element is closely related to the predicted location in that the two modules interact closely to develop a solution.

The second type of predicted plan element is called *EEOI Predicted intent*. It is called predicted intent because it refers to a mental state of the EEOI yet to be realized. That is, SA is predicting that this entity will have some intent in the future to perform one of several activities-engage, evade or some other activity. EEOI Predicted Intent can therefore be thought of as an "Actions on Contact" type of activity in that it describes what an entity would do in the future given a certain set of circumstances. This module uses ownship's future location as a crucial input when calculating how the EEOI will react.

<u>Description Level</u>
This level is a summarization of a single EEOI's predicted attributes. Namely, it is used to summarize the *EEOI Predicted location*, *EEOI Predicted capability*, *EEOI Predicted plan element*, and *EEOI Predicted intent*. EEOI Predicted Intent is developed for an action-on-contact type of activity given knowledge of ownship and whether ownship is in range of an EEOI. Thus, predicted intent will normally be the same as the intent description developed by the assessment blackboard except when ownship will be moving into an EEOI's area of influence.

<u>Prediction Level</u>
This level builds a predicted situation description based on more than one entity's description level summary. It is therefore a multi-entity summarization of the predicted attributes.

<u>The Control Blackboard</u>
The SA system solves the control problem using an opportunistic control model to control its cognitive and perceptive processes. SA opportunistically makes plans to assess both the present situation and future situations. This model requires development of a balance of commitment to the execution of these plans with variable sensitivity to run-time conditions [12]. This balance is a function of the uncertainty of ownship's environment.

SA opportunistically generates plans to solve a problem and places these plans on the control blackboard as they are generated. These plans are then executed in order to solve a particular domain problem. The control blackboard is divided into levels of varying abstraction These levels are illustrated in Figure 5 and further described below.

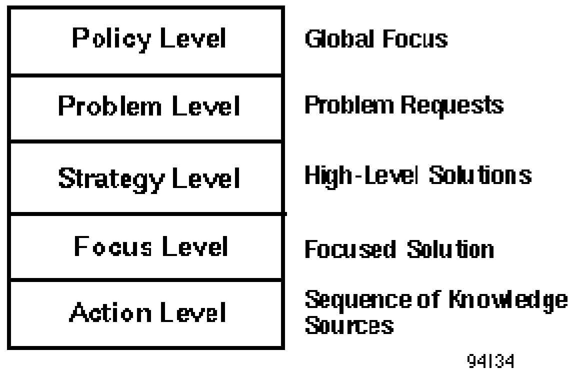

Figure 5. *The Control Blackboard*

The *policy level* is where global focusing decisions are stored. These are heuristically determined at run-time depending on the state of the situation. The *problem level* is where problems presented to the system are placed. Problems can be presented to SA by an external agent (e.g., a system planner or the pilot) or by SA itself. The *strategy level* is where strategies (high-level general descriptions of sequences of actions) are stored. Generating a strategy is the first step in making a plan. Strategy decisions are posted at the strategy level for all accepted problems. A strategy decision is nothing more than a constraint on future actions. The *focus level* is equivalent to the strategy level but is populated by more specific strategies called foci. The *action level* represents the actual sequence of actions chosen by the system to solve the problem [12]. SA has two control shells-one domain-independent and one domain-dependent. These rule sets are distinct and easily separated. It is a relatively simple matter to port the domain-independent portion over to another problem domain.

In order to solve the control problem, SA must opportunistically accept problems and make efficient plans to solve them. Often there will be several plans competing for completion. The system may decide to perform the elements of a partially completed plan and then return to finish the plan later. It may start solving a new problem before finishing the one that it is currently solving. However, before a plan is formulated to solve a problem, the problem must be accepted by the system. Following are the six problems that can be posed to SA:

1. Integrate new information into the Domain Blackboard
2. Focus sensors
3. Generate a current situation assessment
4. Generate a predicted assessment for some opportune time in the future
5. Generate a predicted assessment for time *T* from the present

A real-time SA system is continuously supplied with sensor updates, pilot requests for information, anticipation of possible future events, and a plethora of cognitive actions that must be executed in order to assess the current situation. Thus, most of the work of the control part of the system is in deciding what problem to solve. Following are details on each of the six problems that the system can choose to solve and the criteria for problem acceptance.

<u>1. Integrate New Information into the Domain Blackboard</u>

This problem appears whenever new information becomes available to the system. The system develops an Entity Assessment Rating (EAR) for every previously encountered entity using the Dempster-Shafer methods of evidence combination described previously. The system uses the EARs to strategically place incoming information into the problem acceptance hierarchy.

2. Focus Sensors

This problem is almost equivalent to determining the SA system's focus of attention. This problem request is always on the agenda with a low priority. Generating a focus sensor plan is almost equivalent to determining where the pilot should focus his attention. We added this problem because it is an example of an action (non-cognitive act) that SA requests. Following are situations that favor the acceptance of the focus sensor problem:

- Focus on local entities with high EARs
- Focus on local entities about which we have incomplete information. These entities will have large gaps in their EAR belief function rating.
- Focus on local areas where high belief multi-agent hypothesis object *formation* slot says that as yet unencountered entities might be
- Focus on local key terrain

The system literally uses the formation information from high belief multi-agent hypothesis objects to anticipate general OPFOR locations. This focus recommendation can be combined with local terrain information, individual EEOI doctrine objects, detailed sensor models, and detailed sensor histories to opportunistically direct the sensors. The system aids in directing the sensors to focus on an area that is most likely to provide information that best enhances understanding of both the present and the future situation.

3. Generate a Current Situation Assessment

This problem is a request to evaluate the current situation. SA combines all of the relevant high-level entity specific information with local hypotheses to formulate a situation assessment rating. Situation assessments are frequent and such a problem request is always on the agenda (with varying priority, of course). Favorable situations for the situation assessment problem request being granted are:

- Little unprocessed local EEOI data
- Amount of time since the last situation assessment. The more time since the last assessment, the more favorable it is to generate a new one.

4. Generate a Predicted Assessment for Time *T* Seconds from the Present

This is the problem that an external reactive planner would normally pose to SA. In the existing implementation, SA only asks this of itself. Priority is a function of who requested the problem and the availability of computational resources.

5. Generate a Predicted Assessment for Some Opportune Time in the Future

This is a request to do a predicted assessment and to pick the best time offset from the present for the prediction. Once this problem is accepted and the optimal time from the present to make a prediction is determined, this problem requests problem 5 above. This problem is always under consideration with at least a low priority. It is generally what the system is doing whenever there are slack resources available. Control is always necessary for predictions because time is a continuous function and so there are an infinite number of possible time offsets from the present for which to request predictions. There is a temporal focus on the policy level of the blackboard that has three values-present, short-term future, and long. The future foci are posted when enough time has passed since the last prediction, when the current situation is not dangerous, when there are no predictions available, or when important new data has entered the system.

Once a problem has been accepted, it may be rejected if other more important problems appear. Alternatively, SA may temporarily suspend execution of the less important plan in order to create a plan to solve the more important problem.

Use of Terrain and Terrain Reasoning

For a pilot, knowledge and utilization of terrain is key to survival on the battlefield. The SA implementation uses terrain data and knowledge about the uses of terrain to develop a number of important inferences, such as the following:

- Given the reported/inferred identification of the EEOI, and the reported range and bearing information, validate that the location specified is plausible.
- If the location is found to be unlikely, then use the terrain knowledge base to update characterization level of the assessment blackboard to indicate nearest plausible position.
- Determine whether or not the reported/inferred EE identification is congruent with allowable threat classifications for the surrounding terrain.
- The terrain world model data bases and the Area of Interest (AOI) knowledge base aid in determining the most economical sensor focus pattern.
- Using knowledge from the current assessment characterization level of the blackboard and data from the world model infer the current capabilities of an EEOI.
- Using the terrain knowledge infer a current Volume of Movement (VOM) for each EEOI including Ownship. (This information is useful in predicting location).

Any terrain that might represent a threat to ownship is treated like any other entity except that in this case the entity's status is unconfirmed. SA draws inferences about what might be on the terrain based on doctrine and other nearby EEOIs and develops an Entity Assessment Rating (EAR) for the terrain. The EAR represents the degree of threat that the terrain currently or shortly in the future represents to ownship. The terrain EAR is placed at the Situation Assessment level along with the EARs for the observed entities where it will compete for sensor resources.

Summary and Conclusions

We have implemented a situation assessment system which utilizes extensive domain knowledge for developing inferences about the battlefield environment. While other SA implementations have focused primarily on the sensor data fusion, track generation and track monitoring aspects of SA, our approach has been to utilize extensive background knowledge about the terrain, the battlefield and the enemy to draw extensive high-level inferences about the current situation. The approach closely models the way pilots reason about the battlefield and provides a better situational description since all reasoning is accomplished within the total context of the battlefield rather than considering only sensor information. This approach provides a further significant advantage in that it allows the pilot to "see" and reason beyond the ranges of his sensors.

Acknowledgements

Mr. Joe Marcelino was instrumental in the implementation of the prediction and assessment blackboards. Mr. Travis Bryson implemented the control blackboard systems. Mr. Richard Warren implemented the terrain reasoning system. Mr. Jerry Clark served as the rotorcraft domain expert on this project. Mr. Clark and Mr. Bryson were instrumental in developing the reasoning under uncertainty mechanism used in SA.

Bibliography

[1] D. R. Ballard, "Situation Assessment by Plan Assumption Monitoring," Reticular Systems, Inc. - Final Report (3 Volumes): DAAJ02-90-C-0030 CLIN 0001AF, March 24, 1991.

[2] D. Ballard, "Artificial intelligence in the helicopter cockpit of the future," in *Proceedings of 10th Digital Avionics Systems Conference*, Los Angeles, CA, 1991, pp. 125 - 130.

[3] D. R. Ballard and L. Owsley, "Situation assessment in rotorcraft associate systems," in *Proceedings of Associate Technology Conference: Opportunities and Challenges*, Fairfax, VA, 1991, pp. 111 - 121.

[4] D. Ballard, "Situation Assessment by Plan Assumption Monitoring," Reticular Systems, Inc. - Final Report: DAAJ02-92-C-0034, August 30, 1994.

[5] T. Bryson and D. Ballard, "PalymSys (TM) - An extended version of CLIPS for constructing and reasoning with blackboards," in Proceedings of Third Annual CLIPS Conference, Houston, TX, 1994.
[6] H. P. Nii, "Blackboard Systems - Part I," *AI Magazine*, vol. 7, no. pp. 38 - 53, 1986.
[7] H. P. Nii, "Blackboard Systems - Part II," *AI Magazine*, vol. 7, no. 4, pp. 82 - 107, 1986.
[8] B. Hayes-Roth, "A blackboard architecture for control," *Artificial Intelligence*, vol. 26, no. pp. 251 - 321, 1985.
[9] G. Shafer, <u>A Mathematical Theory of Evidence</u>. Princeton, NJ: Princeton University Press, 1976.
[10] G. Shafer and J. Pearl, "Uncertain Reasoning," San Mateo, CA: Morgan Kaufmann, 1990.
[11] J. Azarewicz, G. Fala and C. Heithecker, "Template-based multi-agent plan recognition for tactical situation assessment," in *Proceedings of Fifth Conference on Artificial Intelligence Applications*, Miami, FL, 1989.
[12] B. Hayes-Roth, "Opportunistic control of action in intelligent agents," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 23, no. 6, pp. 1575 - 1587, 1993.
282.