

Fusion-Based Knowledge for the Objective Force

Gerald M. Powell, Ph.D.
U.S. Army CECOM/I2WD
Fort Monmouth, NJ

Barbara Broome
U.S. Army Research Laboratory
Aberdeen, MD

1.0 Introduction.

Army Vision 2010 identifies information superiority as the key enabler for such force characteristics as dominant maneuver and precision engagement. These concepts are also central to the design and implementation of the Army's Future Combat System (FCS) and Objective Force. To establish and maintain information superiority, analysts and decision-makers need to identify, analyze and interpret pertinent information relative to achieving their task requirements. Currently, the sheer volume of information presented to Army intelligence analysts significantly exceeds their capabilities to fully analyze and interpret it in a timely manner. Consequently, the answers to commanders' critical information requirements (CCIRs) and priority intelligence requirements (PIRs) (FM 34-130) are typically based on a hasty, partial analysis of the information available. This condition of information overload experienced by analysts has the potential to significantly worsen for various reasons. First, our capabilities to collect, communicate and store data/information are steadily rising. Second, faster, more precise, and more lethal battlespace systems of the adversary cause an increase in operational tempo, as well as an increased risk to one's own forces, thereby resulting in more severe time constraints on analysis and decision-making. The nature of the analytical and interpretive tasks required to answer PIRs, and our ability to explain and justify their derivation, have largely been outside the realm of current machine capabilities. In recent years, a number of technologies and approaches have been developed (or matured) that show promise for addressing some of the key sources of difficulty characterizing this set of complex tasks either by emulating human methods or by providing automated support for aspects of these tasks that strain or exceed human cognitive capacities.

To address this set of complex military intelligence problems, the U.S. Army Communications-Electronics Command and the U.S. Army Research Laboratory have submitted a collaborative proposal that would be carried out under the Army's Science and Technology Objective Program starting in FY03. One perspective for viewing this project is the Joint Directors of Laboratories (JDL) Data Fusion Model (Steinberg et al., 1998). With respect to this model, the present project will focus on problems associated primarily with data fusion Levels 2 and 3. However, it is our belief that data fusion problems are more likely to be understood and solved if they are approached more holistically by utilizing data fusion at any or all levels, if appropriate, to help solve a problem on a given level. The present paper provides a description of the technical challenges facing this project, and our current views on addressing them. The remainder of this paper begins by sketching the intelligence cycle and the military decision making process. Next, we discuss operational problems this project will address. This is followed by a description of some of the approaches and technologies we consider to have merit in tackling these problems in the context of a candidate approach representing how they might be employed in this project. Next, we describe issues, and candidate approaches, regarding metrics and operational evaluations. The final section briefly discusses work we have identified as

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 14 AUG 2002		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Fusion Based Knowledge for the Objective Force				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) HQ, U.S. Army CECOM, Intelligence and Information Warfare Directorate, AMSEL-RD-IW-BM, Fort Monmouth, NJ 07703				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

closely related to this project that is being carried out in the Army, in other services, and at the level of the U.S. Department of Defense.

2.0 The Intelligence Cycle, and the Military Decision Making Process

The activities used to gather, analyze and interpret battlespace information are collectively referred to as the intelligence cycle. The intelligence cycle occurs concurrent with, and is logically related to, the military decision making process (MDMP). For this reason, we believe investigations into automated support for fusion should be carried out by considering the problem-solving contexts (planning, decision-making, controlling, etc.) in which fusion occurs. In summary form, the MDMP can be grouped into several major phases: receive and analyze the mission; develop courses of action (COAs); wargame friendly COAs against enemy COAs, and select the most preferred COA; generate and disseminate the operations order; and assess and manage execution of the operation.

The intelligence cycle also consists of several major phases: direct, collect, process, produce, and disseminate. In the Direct Phase, the intelligence staff analyzes the battlefield environment to determine its effects on operations and develops the COAs available to the enemy using a procedure called intelligence preparation of the battlefield (IPB). Wargaming determines which intelligence requirements become priority intelligence requirements (PIRs) as the mission is carried out. The intelligence staff helps identify "trigger criteria;" it is these that become PIRs. Each PIR is stated as a question that must be answered before its associated decision point can be carried out during battle. PIRs are those intelligence requirements critical to the accomplishment of the mission. Only the commander has the authority to select or approve nominated PIRs. PIRs are situation dependent. For a PIR to be considered good, it must ask a question that is rather narrowly scoped such as "Will the opposing force use chemical agents on our reserve in avenue of approach Charlie?" Asking a specific question about what the threat will do, to what part of the force, and where, allows the collection manager to assess the feasibility of whether this PIR can be planned and collected against. The PIRs must be translated into specific information requirements (SIRs). The SIRs provide observable, or inferable, evidence in direct support of the PIRs. The level of description of the SIRs is too low to be useful to commanders. During the Collection Phase, the SIRs are converted into a format more appropriate for collection. A collection plan is developed by comparing the SIRs to available collection resources. The plan specifies collection against the set of SIRs in the form of specific orders or requests (SORs). In the Processing Phase, the raw information generated by the collection resources is transformed into a form suitable for the production of intelligence. During the Producing Phase, processed intelligence is analyzed to generate intelligence conclusions in light of the particular battlefield context.

These conclusions represent answers to each PIR. (In a later section of this paper, the process of moving from a statement of a given PIR through hypothesizing answers to it, and gathering support for/against each PIR will be elaborated). In the Dissemination Phase, the conclusions are distributed to battlefield entities having a need to know the answers to the PIRs. It should be noted that the foregoing descriptions of the MDMP and the intelligence cycle are based on doctrinal sources (FM 101-5; FM 34-8). An empirical analysis could reveal that, in practice, there are deviations from doctrine.

3.0 Operational Problems

The scope of this project presently includes tasks carried out by intelligence analysts (principally the G2/S2 and Collection Manager) in a U.S. Army Division All-Source Analysis and Control Element (ACE), as well as personnel who will conduct intelligence analysis supported by the use of Distributed Common Ground Station – Army (DCGS-A) (Objective), and personnel to be responsible for intelligence analysis in the Army's Unit of Employment (UE) and Unit of Action (UA). Because the Division ACE exists and is well documented, we have much more to say about it than the others at this time. As the other contexts become more defined, we will focus more of our attention on them.

Figure 1 is a slightly modified version of an illustration developed by Walsh (Walsh 2002). We utilize this figure to try to show the focus of our project within the much larger context of Army fusion. The figure provides a perspective that attempts to characterize the Army's fusion problem space. Our interpretation of it is as follows. The x-axis depicts a loosely ordered set of problems associated with fusion that are characterized by the nature of differing tasks involved in collection management, sensor collection and processing, fusion levels, visualization and dissemination. The y-axis is partitioned into categories that represent a hierarchy of intelligence processing activities ranging from single source to all sources of intelligence. The Tactical Unmanned Aerial Vehicle (TUAV), Aerial Common Sensor (ACS), and Prophet represent a sample of sensor platforms associated with single source intelligence, multiple source intelligence, and all source intelligence functions, respectively. The z-axis reflects the fact that fusion problems appear at all levels of the command hierarchy whether it be an individual soldier or elements at EAC. This 3-dimensional space of problems is enormous. Every point in this space has some problem characteristics that differentiate it from all other points. This presents a significant challenge in terms of our ability to provide a generalized solution to any given point in the space.

If we consider the three intelligence systems shown on the y-axis, which represent the systems and intelligence contexts of particular interest in this project, we envision different types of requirements for fusion. We anticipate the FCS UA requiring an ability to carry out Level 1 and perhaps some Level 2 fusion in order to develop an interpretation of the composition and disposition of the local threat forces and their current activities. Due to its computation-intensive nature as well as a more global focus on the battlefield, we expect fusion Levels 1-3 to be carried out at the FCS UE, and for the actionable results of that processing to be communicated to the UA. We would expect that, along with Level 4 processing (probably located at the UE), Levels 1 and 2 (those globally as well as locally oriented) and Level 3 would all be working together in a cooperative manner to answer PIRs. The nature of task allocation and cooperation in this regard should be influenced by, and influence, concepts of operations, staff organization, etc. for the FCS. We expect some subset of fusion tasks will be carried out only at the UA and another subset only at the UE due to such factors as limitations in organic computing power, communication bandwidth constraints, and because information may be locally available to the UA but not the UE. DCGS-A will need to perform Level 1 fusion within a given intelligence discipline (such as IMINT) and across multiple, or all, disciplines based on its requirements. DCGS-A may also need to do some lower-level Level 2 fusion in the form of object aggregation. The All Source Analysis System (ASAS) will need to be expanded to also carry out the full range of Levels 2 and 3 fusion (and, ideally Level 4) in order to identify and adequately interpret threat activities, events and intent in the METT-T context. The remainder of this section characterizes the severity of the information overload problems faced by analysts.




Figure 1. The Army's Fusion Problem Space (after Walsh 2002)

Today, the ACE of a U.S. Army Division has access to approximately 10,000 messages (reports and database documents) per hour in a major theater of war (MTW) scenario. Every intelligence collection system generates USMTF reports. USMTF messages consist of fields that are specified and known; these message types are parsed into the All Source Correlated Database. Other reports, such as Spot Reports, have free text and remarks sections in addition to several fixed fields; the free text and remarks sections are not machine-parsed. The Intelligence Information Report, which is based on HUMINT, is the most problematic type to handle because it is all free text.

Of the approximately 10,000 messages (reports) per hour referred to above, it is estimated that approximately 1,000 of the messages are analyzed superficially and approximately a few hundred are fully analyzed. To say that a message is fully analyzed in the context of the commander's PIRs means that the analysts, in conjunction with the planners, have considered all reasonable implications of each message (both individually, and in the context of all previously received information) in relation to which possible answers (hypotheses) to the PIRs are best supported or refuted by evidence or lack thereof. To say that a message is superficially analyzed indicates that the analysis fails to properly consider the context provided by METT-T in attempting to answer PIRs. Note also that the incoming messages may suggest hypotheses that were not considered in wargaming.

In contrast to a MTW scenario, in a stability and support operation (SASO), the situation is exacerbated because approximately 70% of the reports received are based on HUMINT. It is estimated that only approximately 5% of the incoming messages are fully analyzed in a SASO scenario. In the FCS, an Armor Company Commander (Unit of Action - UA) is anticipated to do approximately 50% of reporting via voice; this will need to be digitized.

In the FCS UA, the MI officers and non-commissioned officers are replaced by multi-functioned staff officers. Although the relevant documents are still in draft form, this concept suggests there may be a need, even beyond that of today, for automated analysis and interpretation to replace some of the human expertise required in going from a specialist in a single functional area (intelligence) to someone who will need knowledge and skills in multiple functional areas.

4.0 A Process to Answer PIRs

During mission analysis, IPB produces a set of threat models including an initial event template and supporting matrix. The collection manager uses these products to focus collection on identifying the COA the threat will carry out. This process continues during wargaming, but includes more focus on particular aspects of the battle.

During wargaming, the intelligence staff role-plays the threat by "fighting" multiple enemy COAs against each friendly COA already generated. During this process, the intelligence staff helps determine the "trigger criteria" (enemy actions) for each decision point (DP) within the Battlefield Operating Systems (BOS) Synchronization Matrix, and Decision Support Template (DST). These trigger criteria become PIRs. For example, "enemy units (BN-strength or greater) extend beyond phase line (PL) Bravo" would signify to the commander during battle to commit the reserve forces. The reserve may have multiple triggers, each one having a corresponding action for the reserve to take. Each logical trigger-action pairing is called a DP. The DPs are recorded by the staff on the BOS Synchronization Matrix (a table specifying a temporal organization of friendly actions associated with trigger criteria). This same type of information is depicted as a map overlay referred to as the DST. Doctrine specifies a one-to-one mapping between DPs and PIRs. An example PIR is "Will the enemy's main defense be along PL Delta or PL Echo?" These PLs would correspond to lines of defensible terrain (LDT) on the modified combined obstacle overlay. For each PIR, the intelligence staff must develop answers (plausible hypotheses). These hypotheses represent enemy activities or enemy COA fragments (solutions) such as "the main effort will be to the north of mountain range Delta along avenue of approach Foxtrot" or "Three enemy tank companies will defend abreast along LDT Echo at avenue of approach X-ray and avenue of approach Yankee." The hypothesis set for a given PIR should be rank-ordered by the intelligence staff and reflects their estimate of likelihood of occurrence for each.

4.1 Differentiating Hypotheses

Analysts need to be able to differentiate hypotheses such that they can recognize which hypothesis appears to most closely reflect what the enemy is actually carrying out based on the available evidence and how it is used in inferencing. (Note that analysts must remain open-minded to the possibility that the enemy could pursue some hypothesis outside those in the set.) The intelligence staff needs to develop a set of indicators for each hypothesis that provides evidence in support of it and uniquely identifies it as different than the other hypotheses. These indicators are typically defined in terms of specific events, activities and entities that should be present if the hypothesis is true. (It should be noted that war gamers often analyze each decision in terms of what specific intelligence will support making it. This analysis may be complex especially if there is even a moderate degree of uncertainty regarding what the threat models are.) These indicators are usually specified in general terms such as "forward deployment of ADA." The requirements manager for collection should coordinate closely with the mission manager to understand the types of SIRs and degree of specificity required to support mission planning and

execution. Each indicator must be further specified by determining where, specifically, to collect on the battlefield. For example, a specific named area of interest (NAI) would replace the general location indicated by “forward.” A similar degree of specificity must be determined for “when to collect” and “what to collect.” If the mission manager requires it, the “what” should be specified in further detail such as types of equipment (e.g., M-109 self-propelled artillery system), numbers of each equipment type, and behavior (e.g., use of a specific radio signal) of entities in the NAI.

4.2 Reasoning with Evidence and Assumptions

As the battlefield situation evolves, information is reported that potentially provides evidence for, or against, the hypotheses under consideration. Typically, the information analysts receive (reports and databases) has been analyzed such that it: (a) has been correlated to resolve ambiguities about which entity is being referenced (entity in this case refers to a platform such as an APC or a missile launcher), (b) indicates which observed parts belong to a given entity (such as a particular radio is linked to a particular weapon system), and (c) identifies entities in terms of type and class (such as a T72 tank). Reports and databases would also contain information communicated about events and activities observed, i.e., not just about entities. In addition, analysts have maps and map overlays available.

The information available to analysts may be inaccurate, incomplete, and otherwise uncertain due to factors such as imprecision in collection assets. Analysts must consider the information in light of these characteristics and use an approach that allows them to estimate likelihoods in terms of the existence and location of key events, activities and entities. This task may be quite complex in that the analyst must apply knowledge of mission, enemy, terrain, troops and time available (METT-T) to properly analyze and interpret each element of information; first, to determine if it is pertinent. Second, if an element of information is deemed to be pertinent, it must be analyzed to determine how it relates to existing information. The analyst needs to construct an interpretation of the battlefield and relate elements of this interpretation to the indicators and SIRs associated with the set of hypotheses. The likelihoods the analyst needs to estimate should be incorporated in the inferencing process and result in an overall likelihood associated with each hypothesis. These overall likelihoods would provide a basis for ranking-ordering the hypotheses in the set. This rank-ordered set is provided to the commander. In addition, analysts need to be able to explain and justify to the commander how each hypothesis was derived. It should be noted there are times when needed information is not obtainable for various reasons. Consequently, assumptions made by analysts become a part of the inferencing process; their truth values need to be monitored for their impacts on the process.

4.3 Examples of Inferencing Types

Analysts need to infer relationships between observed entities in terms of the enemy command hierarchy and in terms of coordinated behavior between units (such as units x, y and z are conducting a reconnaissance operation, or are expected to initiate such an operation during a certain time-interval relative to H-hour). Analysts also need to be able to accurately infer the presence and likely locations of parent entities from observations of child entities. Conversely, knowledge of threat models would apparently be used to guide collection assets to detect, track and identify unobserved child entities. Analysts also need to be able to hypothesize plausible enemy objectives and plausible COAs by which they could be achieved.

The foregoing description of tasks intelligence staff must perform would indicate that the process of developing hypotheses, and sets of indicators and SIRs to represent their validity, have at least the following problem-solving characteristics: a) an ability to infer the set of most

plausible COAs (or COA fragments) the enemy will adopt (these comprise the alternative answers to a given PIR), b) an ability to depict each COA (or COA fragment) in terms of what objective the enemy will attempt to achieve, what events and activities will be required to achieve a particular objective, which types of entities will be involved, where the activities and events will occur and when their presence should be observable and c) analyzing and interpreting information from a large volume of reports and databases in an attempt to find support for (or against) the set of hypotheses.

The tasks are also characterized by: a) knowledge and information that is often incomplete, uncertain and inaccurate; b) the need to reason about time and space; c) the need to deal with large volumes of information that is represented heterogeneously and at different levels of granularity; and d) the stress of making life and death decisions in time-critical situations.

The ability to carry out these tasks requires contextual reasoning drawing on historical and recent knowledge of the enemy, the terrain, the weather, one's own forces, the current mission and situation, and time available. However, the overall nature of the contextual reasoning required will need elucidation. Cognitive task analyses conducted in an operational context should shed light on this.

4.4 An Overall Architecture

Section 4.4 begins by discussing the major architectural elements for this project as we currently envision them. In Section 4.4.1, we describe the Fusion component at the architecture level. This is followed by a discussion of technical issues and requirements associated with this class of interpretation problems, and candidate approaches to address each of them. In particular, we discuss issues and potential approaches related to uncertainty, knowledge representation, time, space, assumption-based reasoning, explanation of hypotheses, and hypothesis management.

In Section 4.4.2, we discuss the Knowledge Management component in the same manner, i.e., architecture, technical issues and requirements, and candidate approaches. Section 4.4.4 discusses how we currently envision accomplishing the integration of the Fusion Component and the Knowledge Management Component. An overall architecture is shown in Figure 3. Section 4.4 ends by briefly discussing the human-computer interface.

4.4.1 Fusion Elements

The JDL Data Fusion Model approaches fusion problems by decomposing the functionality required into multiple levels. Another perspective is that the model decomposes the overall problem into sets of subproblems with a different set assigned to each level. Most of the progress to date has been on Level 1 fusion. As mentioned in the Introduction, we believe the most fruitful approach to solving problems on a given level is to make use of some or all of the other levels as well, i.e., taking a holistic approach. A problem-solving model that uses this approach is called the Blackboard Model. This model has been used successfully in solving other military interpretation problems (e.g., Nii and Feigenbaum 1982). A basic property of the model is to use contextual information on one level (e.g., the plausible behavior of enemy entities) to help resolve ambiguities and/or fill in missing pieces of the solution being addressed on another level (e.g., identifying what unobserved enemy units may be present in a particular area of interest). The model includes a data structure called a blackboard. This is a global database that keeps all of the problem-solving state data (input data, partial solutions, alternative and final solutions). The knowledge required to solve the overall problem is partitioned into knowledge sources that are separate and independent. The knowledge sources cause changes to the blackboard that

incrementally result in a solution to the overall problem. Changes on the blackboard result in opportunistic activation of the knowledge sources; this is the nature of control. An extension of the Blackboard Model, called the Blackboard Framework, resulted from similarities emerging from applications that used the Blackboard Model to build Blackboard applications. At the present time, the Blackboard Framework is the leading candidate for developing a fusion system architecture in the present project. In the Blackboard Framework there exists a set of control modules that monitor the blackboard and have knowledge to decide what actions should happen next, i.e., where attention should be focused. Any type of reasoning approach (data driven, goal driven, model driven, etc.) can be employed at each step of solution formation; problem solving is opportunistic.

Figure 2 depicts a highly notional example of how the levels of blackboard and knowledge sources might appear. Note that the solution space, the blackboard, is hierarchically organized into different levels of analysis and abstraction (from platform level to higher echelon COAs). This is a characteristic of Blackboard problem solving. Corresponding to each level on the blackboard is a knowledge source that solves problems at its own level, but also can contribute solution fragments to other levels. Control knowledge can range from simple to sophisticated. For example, control could incorporate goal-driven strategic problem-solving knowledge with respect to what type of reasoning step should be used next.




Figure 2. A Notional Blackboard Architecture for Fusion

We anticipate a significant amount of knowledge acquisition will be required to develop the knowledge sources. The experiential knowledge used by intelligence analysts is expected to present the most significant challenge for knowledge acquisition. Knowledge acquisition is on the critical path of the project and is recognized as a risk. We hope to mitigate that risk with the environments DARPA is developing in their Rapid Knowledge Formation Program (RKF 2001).

At this point, no decision has been made regarding how to proceed with handling uncertainty. We will be better informed about this decision when we conduct problem analyses in the operational contexts. A number of uncertainty calculi (e.g., Dempster-Shafer, Bayesian, Certainty Factors) have been shown to be effective in interpretation problems that require knowledge-based solutions. In fact, some recent work directly related to some of the problems being addressed by the present project has used Bayesian Belief Nets (Wright et al. 2002; Gonsalves and Cunningham 2000) to represent and propagate uncertainty. Some intelligence analysts believe Bayesian Belief Nets provide an accurate representation of certain aspects of the way that Army analysts reason about answering PIRs (Schlabach 2002).

It is likely that multiple knowledge representation formalisms will be needed. We have already mentioned that Bayesian Belief Nets are a candidate. They provide a probabilistic approach. They structure knowledge as networks wherein nodes represent variables denoting solution fragments while links represent probabilistic relations between nodes. To support reasoning about objects (such as weapon platforms or battalions) in terms of their attributes, we anticipate a need for some form of structured representation such as frames. To represent heuristic knowledge, such as that which may represent inferencing between levels on the blackboard, we anticipate the use of production rules. Rules may be used in various ways; for example, to represent problem decomposition knowledge (non-terminal rules), and to represent knowledge that produces a solution state (terminal rules).

We anticipate a need for an inferencing mechanism that implements the use of assumptions in reasoning. This would likely be part of a truth maintenance system employed to preserve the logical integrity of the conclusions inferred. As beliefs expressed by clauses in the knowledge base are revised, it is necessary to recompute the values of the inferencing structure dependent on those beliefs. It will likely be desirable to maintain multiple possible states of belief at once; an assumption-based truth maintenance system provides this advantage (deKleer 1984).

The battlefield is dynamic. Observations are collected over time. Time is an element used to plan and execute single agent behaviors, and to coordinate plans among multiple cooperating agents; etc. We anticipate the requirement to model events and activities in terms of absolute times, relative times, and durations. One approach to temporal reasoning is with probabilities as in stochastically modeling the progression of a system through a sequence of states. A number of different temporal logics have been developed to support temporal reasoning. McDermott developed a temporal logic for reasoning about plans and actions (McDermott 1982; also see Allen, 1981 and 1984). A better understanding of requirements will guide us in making decisions about how to deal with temporal issues.

It is anticipated that spatial reasoning will play a key role in the fusion tasks to be addressed by this project. Various approaches have been used successfully (e.g., quadrees and fuzzy spatial templates). Fuzzy spatial templates may be used for recognition and possible identification of complex aggregate objects. No commitment to a particular approach has been made at this time. As for the effects of terrain on entity attributes such as location and mobility, we are hoping to utilize battlefield terrain analysis software developed for the Army.

Since uncertainty characterizes information and knowledge in this domain, we will want to maintain multiple, simultaneous hypotheses (e.g., Jones et al. 2002). In answering PIRs, analysts attempt to generate a set of plausible solutions the enemy can adopt. Each solution corresponds to a hypothesis. Moreover, each level of the blackboard contains its own set of hypotheses corresponding to the solution types developed on that level.

To give the user insight into the interpretive process, we plan to provide an explanation facility. This type of facility has been successfully incorporated into various knowledge-based systems by allowing the user to see the inferencing chain used to reach conclusions. If the formalisms used for inferencing are not understandable by the users, then a translation into a more natural language format (and perhaps visual format) will be required.

4.4.2 Knowledge Management Elements

In recent years, methods for harnessing an organization's knowledge have converged in a practice referred to as *Knowledge Management (KM)*. Simply put, KM is the process of capturing an organization's collective information and expertise and providing access to them in a manner that produces a payoff. The information may be explicit, residing in databases, or on paper; or it may be tacit, residing in people's heads. The goal of KM is to help people work better together, using and managing combinations of their explicit and tacit knowledge to have a more effective impact. (Hibbard 1997; Excalibur 1999; Liebowitz 1999) We use the term *knowledge* here to refer to information on which one may act in order to perform a given task. It is distinguished from *data*, the raw facts associated with a task, and *information*, summaries of that data, simply by the degree to which it supports the user's decision process. As such, the software environment must be more tuned to the user than traditional database management and information management systems of the past.

The proposed effort addresses not just Level 2 and Level 3 fusion, but a Knowledge Environment for the Intelligence Analyst (KE-IA) that supports Level 2 and Level 3 fusion requirements. The Army has become increasingly aware of the knowledge-oriented nature of its mission and operations. Recent plans include knowledge among its five Research Focus Areas (lethality, survivability, agility, sustainment, and knowledge) to achieve the leap-ahead capabilities anticipated for the Objective Force Warrior. (Andrews, Beatrice et al. 2002) Information Technology, as evidenced in the Internet, has enabled entirely new ways of managing business knowledge for the commercial world, broadening our notion of the types of information that can be readily accessed and the techniques by which that access is achieved. These same dynamics, applied to the DoD, can transform our military from a platform-centric force to a network-centric force in information that can be readily shared among geographically distributed forces including sensors, decision-makers, and shooters. (OSD2 2001) In fact, in applying these techniques to domain-specific tasks, the knowledge sharing function can be even more efficient and effective than in the more general-purpose Internet environment. In this context, then, a *KE-IA* should put together a coherent situation description, alert the analyst when certain events of interest can be hypothesized from available observations, suggest answers to PIRs, and evaluate them against the evolving scenario, providing:

- Rapid access to widely distributed heterogeneous data and information systems, that may change or increase over time.
- Information push techniques, identified from an analysis of the user's task, that offload much of the information pull typically associated with today's systems.

- Tools that support both automated and user-directed Knowledge Discovery, based on the integration of information across sources

The Intelligence Community's current All Source Analysis System (ASAS) can be thought of as a collection of analysis systems that support tasks like terrain analysis, weather forecasting, sensor correlation, and search engines are a few examples. They represent data sources ranging from structured to semi-structured to unstructured data formats. The CECOM fusion techniques represent another software package that sits between the user and a series of widely varying data sources. For this project, as depicted in Figure 3, we will pick at most 5 data sources to work with. The sources will be refined as the project unfolds, but will most certainly involve the ASAS database (of sensor data), the Internet, the IMETS weather forecaster (Hooock and Giever 2000), and a terrain database, again ranging from structured to unstructured. Our primary task will be to accomplish the interface between the software and the data for CECOM.




Figure 3. Components of the Fusion Based Knowledge for the Objective Force STO

The integration task will be accomplished via a community of agents, as outlined in a later section (4.4.4). Figure 3 illustrates that this interface will provide flexible access between all software tools and all data sources. In addition, it will provide access between all software tools. Thus, the advanced fusion algorithms and representations may access directly from the OLAP system, or from the data sources that feed the OLAP tool.

4.4.3 Ontology-Based Source Integration

A great deal of research is currently focused on the Semantic Web concept. While the techniques are not completely defined at this point, it is clear that the use of XML with Ontologies to integrate heterogeneous sources has been used successfully and will improve with time (Berners-Lee, Hendler et al. 2001). Advances in techniques for extracting data from the web or for retrieving information from databases look promising (Abiteboul, Buneman et al. 2000; Goldman and Widom 2000). Information agents are already being used to support a push paradigm over the standard user-intensive pull paradigm of the past (Burke, Hammond et al. 1996; Delgado 2000)]. It is not clear at this point whether the best approach for ontology development is to try to define a single ontology for the entire intelligence task, an Interlingua, or individual Ontologies for specific intelligence categories, or a hybrid of the two (Wache, Vogele et al. 2001). It is also not clear how deep the analysis must be to build an effective Ontology. The approaches described for Internet projects vary greatly from those described by database developers, and those vary greatly from those described by linguists. ARL's work in MT and Ontology Algebras can directly impact this STO and, potentially, vice-versa.

4.4.4 Information Agents

This is a system in which the information required for decision support will not belong entirely to the user. The data within many of the sources will be extremely volatile, and we expect many users will access the system at once. Database technology has introduced a number of techniques for accessing such, including migration, mediation, migration-mediation-hybrid and agent-based architectures. (Subrahmanian, Bonatti et al. 2000) Each of these approaches has strengths and weaknesses. For simplicity, the migration approach is hard to beat, but the continual migration efforts required for volatile data sources can prove quite expensive in terms of overhead. A single mediator can eliminate that overhead. In a mediator system the original data sources are tapped when data is required. That leaves the burden of maintaining the data on its originator, thereby eliminating the overhead associated with the migration approach. But it introduces a bottleneck, the mediator itself, that can render the system ineffective (Subrahmanian, Adali et al. 1995). A third approach, a hybrid of the two, might be to migrate the most often used data to ameliorate the bottleneck problem, and then access the most volatile data via mediators. But when there are lots of users and lots of sources, the relief is minimal. Recent work (Eiter, Subrahmanian et al. 1999) has focused on an agent-based approach, with multiple mediators organized by a supervisor or responding to broadcast queries to accomplish the data access tasks. While this agent architecture is more complex than the others, when there are many volatile sources and many users the relief from bottlenecks and constant updates more than makes up for the increased architectural complexity. Based on the current and projected number of sources and users for the Intelligence task, the agent-based approach is probably most appropriate for this problem.

4.4.5 Knowledge Discovery Tools

During the 1970's E.F. Codd (Codd 1970) introduced the relational data model, earning him the Turing Award a decade later and serving the foundation for today's standard database industry (Pedersen and Jensen 2001). Twenty years later, Inmon (Inmon 1992) and Codd (Codd 1993) observed that standard relational databases and their associated operational-level online transaction processing (OLTP), could not efficiently co-exist with decision support applications, due largely to their very different transaction characteristics. While standard relational databases and OLTP is effective in supporting an organization's current asset summary requirements, when users attempt to identify trends and predict future requirements, wider ranging, often historical, sources and more sophisticated data structures and access techniques are required. (Jarke, Lenzerini et al. 2000). Over the past ten years On-Line Analytical Processing (OLAP) has emerged as a powerful Knowledge Discovery tool to address forecasting and trend analysis issues within the Decision Support environment.

The Multi-Dimensional Database MDDDB was developed to facilitate exploration of data from a variety of perspectives. Those perspectives are built directly into the data structure as dimensions. The dimensions of the data structure are used for selecting and aggregating data. To avoid unnecessary duplication of data, the developer can define dimensional hierarchies. (Pedersen and Jensen 2001) This MDDDB structure used within the OLAP system provides a more natural, more flexible storage and retrieval mechanism than the more traditional 2-dimensional table or spreadsheet structure of the OLTP. This representation better reflects the way decision makers think about their data, and it creates a natural environment for applications that involve time-series analysis, cross-sectional analysis, and forecasting (Pottle 2000). Typically the MDDDB is maintained independently of an organization's operational databases. It contains data consolidated from a variety of such databases, so it is often orders of magnitude larger than a standard database. It is developed principally to support decision support applications, providing historical records that summarize the contents of the operations databases that feed it. (Chaudhuri, Dayal et al. 2001).

The use of MDDDBs for trend analysis and forecasting addresses a user-directed approach to many of the same problems addressed by the project's proposed fusion algorithms. In order to accomplish these tasks a system must maintain records over both time and space. Since OLAP was developed in the early-to-mid 90's it is not surprising that it is not prevalent in today's ASAS, a system designed in the early 80's. One part of this program will be the tailoring of this relatively mature technology to the intelligence task. One problem we foresee with this approach is the incorporation of volatile data sources. While OLAP is optimized for roll-up, drill-down, trend analysis, and forecasting, I/O is not its strong point. One benefit we foresee is the potential to incorporate data mining techniques into this research. The MDDDB structure is commonly used for data mining as well as OLAP. While data mining is not a deliverable of this project, it is a potential "extra" that can be tied to either internal ARL research or to the data mining activities associated with the DARPA Anti-terrorism program.

4.4.6 Human-Computer Interface

So far we have described this KE-IA as one that minimizes the user burden to acquire, access, and mentally combine all the data and information required to accomplish the Intelligence task. But a Knowledge Environment, that is a system of tools that support the sharing of data and

information in order to provide more timely and efficient decisions, must be user centered. That is, in a system that supports as many users and functions as the one proposed, the system must be capable of adapting its response to the task at hand and to the competencies of the current users. Visual and organizational structures should match the nature of the information requested and should accommodate both multi-modal and collaborative interaction. While these issues must be addressed in the long-term, these capabilities are not funded in the current effort. We have, however, identified several very promising efforts, both in ARL and CECOM, that are targeted to the human-computer interface issues of the Intelligence Analyst. One of the challenges of this project will be to integrate those programs to provide a user interface that facilitates the analyst's reasoning process.

4.5 Cognitive Engineering

Aligned with our holistic view toward solving data fusion problems that appear in the JDL Model, is our belief that it is essential that the entire system requirements, design and development process be focused on understanding human problem-solving in context. We believe a detailed understanding of the sources of difficulty facing analysts will reduce the risk and cost of repeatedly developing systems that fail to support operational personnel in meeting the most critical challenges of military intelligence such as analysis under conditions of information overload.

In the present project, Cognitive Task Analyses will be employed in an attempt to reveal the overall flow of the problem-solving process, the classes of problems addressed, the information requirements, the key problem-solvers and the nature their collaborative problem-solving in the process, the flow of information, the use of visual information such as maps, map overlays and imagery (perceptual processing), the knowledge required and how it is used, as well as the types of intermediate and final solutions generated. This type of analysis, coupled with experimentation, should also reveal some of the difficulties encountered by staff due to sources of complexity presented by the problems as well as constraints on the human information processor. These areas of difficulty would become potential candidates for machine solutions or support.

The techniques and technologies we have outlined in Section 4.4 represent our current best hypotheses about how to address what we believe are the major sources of difficulty, but we have made these determinations without the benefits that will be derived from performing cognitive task analyses. The goal is to develop a human-computer cooperative problem-solving system that identifies appropriate roles for the users, and leads to a human-machine system design that increases overall performance.

5.0 Evaluation and Metrics

In a sense, the information access agents provide an information retrieval (IR) system, and as such, we will rely on IR techniques to assess the agent development over time. Precision and recall are the two most commonly used metrics for evaluating IR systems, where *precision* is the *fraction of relevant documents retrieved*, and *recall* is the *fraction of relevant documents in the answer set*. They are popular in large part because they support quantitative assessments of both the quality of the answer set and the breadth of the retrieval algorithm, and are widely used in the literature. They have come under criticism recently, in large part because they are not easily obtained in a large interactive environment like that on the Web or in our Intelligence Analyst environment.(Baeza-Yates and Ribeiro-Neto 1999) Nevertheless, they provide a mechanism for early laboratory assessments of progress on agent component development. As later integration

efforts are in place a more task based assessment of effectiveness will be incorporated with system-level evaluations.

The evaluation of the OLAP user-directed knowledge discovery tool will require both subjective and objective assessments. Since the multi-dimensional technique is new to the user community, one question we must ask is whether the user can, within a reasonable amount of time, become comfortable using this tool in place of the more familiar 2-dimensional relational database tools. In addition, more objective tests will compare the test scores of users with and without access to the OLAP system when addressing questions that require the complex analysis of information involving multiple parameters.

The main cognitive task in intelligence analysis involves developing the most plausible explanation for uncertain and incomplete information. Due to this uncertainty, the product of interpretation is characteristically arguable. However, some hypotheses and their derivations could be argued to be better (more plausible) than others. This can sometimes be recognized by human experts who perform intelligence analysis.

To evaluate changes in interpretation effectiveness and performance, we would like to compare the ability of analysts to answer PIRs using their current methods versus the methods they will be able to use given machine capabilities developed in the present project. Some comparisons can be made on the basis of ground truth whereas others will require an independent group of human experts in intelligence analysis.

Example measures of interpretation (fusion) performance and effectiveness that we are considering include:

- accuracy of hypotheses regarding aggregate entities including their echelon, classification, functional grouping, and location (especially if they are expected to be present in a NAI). These may be measured with fidelity scores against ground truth.
- accuracy of hypotheses regarding plans and sub-plans (e.g., defending abreast along LDT Bravo; conducting a reconnaissance operation; conducting a supporting attack along Avenue of Approach Charlie)
- latency to detection of critical events (e.g., seizing key terrain such as a particular bridge)
- missing the presence of critical events (or indicators); measured by per cent and type
- false alarms (incorrect hypotheses of all types)
- strength of evidence supporting hypotheses generated
- latency for Commander to respond to identified critical enemy activities

We deem scenario development to be a key element of our approach to evaluation. We need to have scenarios that reveal how well our project's software supports analysts in meeting challenges from the domain. Scenarios need to be usable. That is, analysts under evaluation must be able to understand them (this, in itself, will need to be assessed) and they must be designed to allow us, as researchers, to interpret analysts' performance without ambiguity.

6.0 Related Work

This research sits within the context of a number of recent initiatives focused on various aspects of the Level 2/3 Fusion problem. Much of that work is DARPA funded, since many of the techniques required to address this complex problem are still in the basic research phase. However, this effort places the Army in an excellent position to transition those research components into the Intelligence-focused component of the Unit of Employment while addressing many of the application-specific issues of today's analyst. DARPA's recently formed Information Awareness Office is focused on providing instant access to surveillance and information analysis systems to support Homeland Security, while DARPA's Information Exploitation Office has established research to shorten the time between when an enemy target is located and when it is attacked on the battlefield (Markoff 2002). Their INSCOM led Information Dominance Center provides a testbed that addresses transitions of this research to the Theater-Level Intelligence problem. DARPA's Future Combat System program recently awarded a short-term contract focused on Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR) for Knowledge Management and Fusion (DWDU 2002) that will explore the effectiveness of a number of select technologies, used in combination, on solving problems at Levels 2 and 3 of the JDL Data Fusion Model. In addition, a number of DARPA efforts can impact the various components of the FBKOF program. The work associated with multi-dimensional databases and user directed knowledge discovery provides an excellent transition vehicle for portions of DARPA's Knowledge Discovery, Data Mining and Machine Learning (KDD-ML) effort (Goldszmidt and Jenson (Ed.) 1998). The approach to the overall architecture is impacted both by DIA's Virtual Knowledge Base Concept of Operations (DIA 2002) and by the Joint Intelligence Virtual Architecture (JIVA) project (FAS 2000) as well as by the DoD's Network Centric Warfare and Horizontal Fusion concepts (OSD2 2001). DARPA also recently initiated a short-duration seedling project aimed at investigating the utility of a combination of a particular set of technologies aimed at problems associated with fusion Levels 2 and 3 (Kessler 2002). We are coordinating closely with DARPA with respect to this seedling; there is an excellent opportunity for the results to be used in further shaping the FBKOF program. The Air Force Office of Scientific Research is sponsoring new basic research efforts in upper levels (JDL Model) of Information Fusion targeted at image analysis, command and control, and support for natural and man-made disasters (Hinman 2002a). The Air Force Research Laboratory, via the Small Business Innovation Research Program, is pursuing computational approaches for situation and impact assessment (Hinman 2002b). There are excellent opportunities for the Air Force and the Army to mutually take advantage of the results of these Air Force Programs and the FBKOF Program.

7.0 Summary

In sum, the volume and nature of information reported to analysts and decision-makers exceeds their capabilities to process it in a manner that satisfies the time-constraints and level of situational understanding desired for planning and acting within the adversaries decision cycle. The overall objective of this science and technology project is to develop an advanced knowledge generation and explanation capability (automated decision-support) for answering war fighting commanders' critical intelligence requirements in a timely manner. The scope of the project will address particular requirements and issues associated with intelligence analysis and decision-making conducted at the U.S. Army Division level today, as well as the Army's Future Combat System's Unit of Employment and Unit of Action. Clearly, the problems addressed by this project intersect with many of the critical problems characterizing the war against terrorism as well. This paper characterized the nature of the problems and challenges currently faced by Army analysts and the decision-makers they support. It identified issues and requirements associated with these problems and described our planned technical approach including: the technologies we plan to explore and how they may be utilized (e.g., ontologies, Bayesian belief

networks, rule-based systems, and knowledge discovery); an initial candidate system architecture; metrics and methods for system evaluation; and the central role of cognitive engineering in our approach to human-computer system design. We also identified a number of key projects directly related to this one both within DARPA and the U.S. armed services that we believe provide an excellent opportunity for cross-fertilization and synergy.

8.0 References

- Abiteboul, S., P. Buneman, et al. (2000). Data on the Web: From Relations to Semistructured Data and XML. San Francisco, Morgan Kaufmann Publishers.
- Allen, J.F. (1981). An Internal-based Representation of Temporal Knowledge. Proceedings 7th IJCAI, 221-226.
- Allen, J.F. (1984). Towards a General Theory of Action and Time. Artificial Intelligence, 23(2), 123-154.
- Andrews, A. M., P. Beatrice, et al. (2002). A Vision for the Objective Force Warrior. Army AL&T.
- Baeza-Yates, R. and B. Ribeiro-Neto (1999). Modern Information Retrieval. New York, Addison-Wesley ACM Press.
- Berners-Lee, T., J. Hendler, et al. (2001). "The Semantic Web." Scientific American(284): 35--43.
- Burke, R., K. Hammond, et al. (1996). Knowledge-based navigation of complex information spaces. 13th National Conference on Artificial Intelligence, Menlo Park, CA, AAAI Press.
- Chaudhuri, S., U. Dayal, et al. (2001). "Database Technology for Decision Support Systems." IEEE Computer 34(12): 48-55.
- Codd, E. F. (1970). "A relational model for large shared data banks." Communications of the ACM 13(6): 370-387.
- Codd, E. F. (1993). Providing OLAP (On-line Analytical Processing) to User-Analysts: An IT Mandate (White Paper), E.F. Codd Associates, Commissioned by Arbor Software (now Hyperion Solutions).
- Delgado, J. (2000). Agent-Based Information Filtering and Recommender Systems on the Internet. Engineering. Nagoya, Nagoya Institute of Technology.
- DIA (2002). Virtual Knowledge Base: A DoD Intelligence Interoperability Framework Supporting Joint Vision 2020 (Concept of Operations). Washington, DC, Defense Intelligence Agency.
- DWDU (2002). Subcontractors Win Work on Future Combat System. Defense Week Daily Update.
- Eiter, T., V. S. Subrahmanian, et al. (1999). Regular Agent Programs in IMPAC. Proceedings of the International Workshop on Distributed and Internet Programming with Logic and Constraint Languages, Las Cruces, New Mexico.
- Excalibur (1999). Knowledge Retrieval: the Critical Enabler of Knowledge Management, Excalibur Technologies.
- FAS (2000). Joint Intelligence Virtual Architecture (JIVA). 2000.
- FM 101-5 (1997) Staff Organization and Operations. HQ, Department of the Army.
- FM 34-130 (1994) Intelligence Preparation of the Battlefield. HQ, Department of the Army
- FM 34-8 Combat Commanders Handbook on Intelligence.
- Goldman, R. and J. Widom (2000). WSQ/DSQ: a practical approach for combined querying of databases and the web. ACM SIGMOD International Conference on Management of Data, Dallas, Texas.

- Goldszmidt, M. and D. Jenson (Ed.) (1998). General Findings of the KDD-ML Workshop. DARPA Workshop on Knowledge Discovery, Data Mining, and Machine Learning (KDD-MI), Pittsburgh, PA, Carnegie Mellon University.
- Gonsalves, P. G. and R. Cunningham (2000). Automated ISR Collection Management Systems (AICMS): Final Report for US Army Communications-Electronics Command. Fort Monmouth, NJ, Charles River Analytics.
- Hibbard, J. (1997). "Knowing What We Know." Information Week Online.
- Hinman, M. (2002a). AFOSR Basic Research in Information Fusion. G. Powell
- Hinman, M. (2002b). Some Computational Approaches for Situation Assessment and Impact Assessment. Fifth International Conference on Information Fusion, Annapolis, MD.
- Hoock, D. W. and J. C. Giever (2000). Weather Data Visualization - Decision Aid Tools for Army C4I. Battlespace Atmospheric and Cloud Impacts on Military Operations, Fort Collins, Colorado.
- Inmon, W. H. (1992). Building the Data Warehouse. Wellsley, MA, QED Information Sciences.
- Jarke, M., M. Lenzerini, et al. (2000). Fundamentals of Data Warehouses. New York, Springer.
- Jones, E.K., Denis, N. and D. Hunter (2002). Hypothesis Management for Information Fusion. Fifth International Conference on Information Fusion, Annapolis, MD.
- Kessler, O. (2002). PBA Seedling Project Kickoff Meeting. G. Powell
- Liebowitz, J., Ed. (1999). Knowledge Management Handbook. Washington, D.C., CRC Press.
- Markoff, J. (2002). Chief Takes Over New Agency to Thwart Attacks on U.S. New York Times. New York, NY.
- McDermott, D. (1982). A Temporal Logic for Reasoning About Plans and Actions, Cognitive Science, 6, 101-155.
- Nii, P. H. and E. A. Feigenbaum (1982). Signal-to-Symbol Transformation: HASP/SIAP Case Study. AI Magazine: 23-35.
- OSD2 (2001). Network Centric Warfare. 2001. <http://www.c3i.osd.mil/NCW>.
- Pedersen, T. B. and C. S. Jensen (2001). "Multidimensional Database Technology." IEEE Computer 34(12): 40-46.
- Pottle, B. (2000). Oracle Express Foundation, Kelly Lee Publisher.
- RKF (2001). Rapid Knowledge Formation Project for the Stanford Knowledge Systems Laboratory. 2001. <http://www.ksl.stanford.edu/projects/RKF/>.
- Schlabach, J. L. (2002). Intelligence Analysis. G. Powell.
- Steinberg, A.N., Bowman, C.L. and F.E. White, Jr. (1998). Revisions to the JDL Data Fusion Model. Proceedings 3rd NATO/IRIS Conference, Quebec City, Canada.
- Subrahmanian, V. S., S. Adali, et al. (1995). HERMES: A heterogeneous reasoning and mediator system. College Park, MD, University of Maryland.
- Subrahmanian, V. S., P. Bonatti, et al. (2000). Heterogeneous Agent Systems. Cambridge, Mass., MIT Press.
- Wache, H., T. Voegelé, et al. (2001). Ontology-Based Integration of Information -- A Survey of Existing Approaches. Workshop on Ontologies and Information Sharing, IJCAI 2001, Seattle, Washington.
- Walsh, D. (2002). A Notional Level Two Fusion Processing Architecture. Military Sensing Symposium, San Diego, CA.
- Wright, E., Mahoney, S., Laskey, K., Takikawa, M. and T. Levitt (2002). Multi-entity Bayesian Networks for Situation Assessment. Fifth International Conference on Information Fusion, Annapolis, MD.