

Actionable Intelligence for the Warfighter

Nicholas Morizio, Sergio Gigli, and Angela Pawlowski
Lockheed Martin Advanced Technology Laboratories
Cherry Hill, NJ 08002
[nmorizio, sgigli, apawlows]@atl.lmco.com

1. ABSTRACT

Lockheed Martin Advanced Technology Laboratories (LM ATL) has researched and developed Situation Understanding technologies to provide tailored, Actionable Intelligence to the individual warfighter. Situation Understanding (SU) is a core requirement of the Future Combat Systems and programs such as the Distributed Common Ground Station – Army (DCGS-A) and the Aerial Common Sensor (ACS).

LM ATL has developed an SU Engine to automatically fuse multiple intelligence reports, e.g., IMINT, SIGINT, MASINT, HUMINT and OSINT, with track data, into a Common Relevant Operating Picture (CROP) of the battlespace. The SU Engine augments the CROP with hypotheses as to the relationships that may exist between entities, environment, and events within the battlespace. These relationships are then used as the basis for inferring the most likely and most dangerous courses of actions that the enemy may be pursuing. The Future Force is actively trading weight for intelligence, while at the same time supporting a broader range of missions, with fewer operators and greater volumes of information. Automated SU systems that aid the warfighter are, therefore, essential to enabling the Army to “see first, understand first, act first and finish decisively.”

The SU Engine maintains the context of the various warfighters that the system is supporting. A warfighter’s context includes location of the warfighter, the warfighter’s mission, and the state of the battlespace surrounding the warfighter. The SU Engine, based on any explicit information requests provided by the warfighter combined with needs inferred by the SU Engine based on the warfighter’s context, dynamically composes multi-level fusion services to convert raw sensor and report data into higher level relationships and ultimately into predictions of enemy courses of action. The SU Engine can access sensor and report data from a range of sources including service-enabled net-centric systems similar to those planned for the DCGS Integration Backbone (DIB). Services within the SU Engine are described using industry open standards augmented with semantic definitions to support just-in-time service composition. The SU Engine is built around the SU Virtual Battlespace, a data store that supports the persistence of a range of SU products from the underlying tracks through to actual courses of action.

The SU Engine, through both the needs driven invocation of fusion services and the filtering and tailoring of the resulting products, is able to match the information needs of the warfighter, thereby providing the warfighter with Actionable Intelligence. Experiments with the SU Engine have created tailored information for warfighters at various echelons including an S2 and a Platoon Leader.

2. OPERATIONAL DRIVERS

When faced with the agile and creative adversaries that are becoming more common in contemporary military environments, decision-makers at all echelons are burdened with difficult challenges. Among these challenges are the needs to quickly observe and understand large amounts of data provided by a

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 MAY 2005		2. REPORT TYPE		3. DATES COVERED 00-05-2005 to 00-05-2005	
4. TITLE AND SUBTITLE Actionable Intelligence for the Warfighter				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) Lockheed-Martin Advanced Technology Laboratories,3 Executive Campus, 6th Floor,Cherry Hill,NJ,08002				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	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

wide range of sensor platforms, and to produce from that understanding, Actionable Intelligence that will provide aid to the warfighter.

These challenges are further complicated by the often very large gap between unprocessed data and the types of high-level conclusions that can be used to enable tactical counteraction. While sensor and track fusion can be used to provide a CROP to decision-makers and warfighters, additional intelligence is needed to add context to observed track movements and to infer courses of action (COA). Without tools to provide substantiated estimates of likely and dangerous enemy COA, the decision making cycle of the enemy may be too rapid for an effective and decisive response.

Often, COA can only be substantiated using information available from multiple perspectives. These perspectives can range from those that can be directly obtained via sensors, such as the Mobility of enemy units, to those that must come primarily from human analysts, such as the effects that the enemy's Culture may have on a given situation. The combination of these perspectives into one cohesive picture of the battlespace is typically a manual and tedious process with little automated fusion available to aid the decision-maker.

3. TECHNICAL CHALLENGES

A number of technical obstacles exist on the path from data to a cohesive and complete picture of the battlespace. Data is typically obtained from multiple, possibly overlapping sources and must be combined into a single representational structure in order to be shared by various reasoning processes. These sources will typically have different formats, protocols, and update rates that require mediation strategies that are customized to each specific source.

The quality of the data provided can also vary among the sources. Different sensor platforms can provide the same measurements with different accuracy, and some systems may be inherently unreliable while others are consistently effective. Thus all data has some pedigree that must be taken into account in order to avoid errors in propagation of inferences and subsequent compromising of the inferential accuracy. This pedigree must also be used to disambiguate contradictions in reports from multiple sources.

Further, no one fusion algorithm has been found to accomplish all the inferences and evidence combination needed to form the complete battlespace picture. Thus multiple fusion processes are required to combine the data available to the fusion system. One important technical challenge is how best to compose these processes. This is especially difficult given the large quantity of data that can be gathered in a small amount of time by modern sensors. Accurate evidence combination relies on correct composition.

Just as humans evaluate information and make conclusions from multiple perspectives, evidence accumulation in inferential reasoning may be approached from multiple perspectives. The need to assess the situation from multiple perspectives and combine these into one cohesive picture suggests that eventually all data must enter into a central data store for processing. The data store must then support the combination of these perspectives.

In addition, this data store must be intuitive enough to be maintainable by analysts as well as robust enough to accurately represent data that is very different in nature. Sensor reports must be captured along with the high-level inferences to which they contribute, and intangible relationships between events and abstract concepts must be linked to tracks that are constantly on the move. Ultimately, the data store must

support efficient querying by analysts and reasoning algorithms while at the same time support the representation of arbitrarily complex situations.

Since the needs of both the warfighters and the decision-makers change frequently with mission and status, it is also important to focus the fusion on the production of those pieces of information that are most imperative at a given point in time. When a rapid response is required from the fusion system, often there will not be time to apply every fusion algorithm to the entire set of data available. Meta-control techniques based on the context of the warfighter must be applied to focus the fusion performed by the system.

4. APPROACH

These formidable technical challenges led LM ATL to design and implement the Virtual Battlespace (VBS) as the central component of the SU Engine (Figure 1). The VBS is an ontology representing all that is known about the particular domain and the current situation.




Figure 1. The Virtual Battlespace

Since the VBS must represent multiple perspectives of the environment in addition to historical data on past situations, it is important that the representation be both robust and flexible. In order to perform reasoning across the complete set of data that is entered into the VBS the semantics of that data must be captured. Conceptual Graphs (CG) [Sowa92], when reinforced by an underlying ontology of types, are a natural and powerful way to fulfill these requirements. To capture uncertainty within the CG itself, LM ATL added confidence values to all of the nodes. This confidence is interpreted existentially so that the confidence of any concept or relationship is independent of the confidence on any other node to which it might be connected.

To enable reasoning over events, we chose to represent an event as a special type of concept within the CG. Situation nodes are concepts that contain subgraphs describing an event at a certain point in time. Situations all have a time window in which they occur. Concepts that exist in multiple situations are referenced out to the first subgraph in which they first appear. This representation allows the VBS to easily store sequences of events, and allows for fusion processes that perform temporal analysis.

The concept and relation hierarchies used by the CG were implemented in the Web Ontology Language (OWL) as a class hierarchy. These definitions contain meta-information that allows incoming data to be easily mapped to a concept type. In addition, type definitions can refer to other more primitive types to ease the creation of these hierarchies. This approach allows the semantics of the concepts and relationships to be defined alongside the labels, and these semantics can then be used during the reasoning process.

Since the VBS needs to maintain and provide persistence for the sensor reports that are used to create the concepts and relationships, the VBS has the ability to represent tracks and many types of sensed data as first order objects in the knowledge base. The concepts in the graph can then be linked to these data items, forming a separate layer within the VBS. This provides a way of merging the tactical picture of the battlespace with the intelligence picture.

In the example in Figure 2, the aggregate “Platoon Z” is linked to the tracks that are members of the aggregate, which are in turn linked to the sensor reports that generated the tracks. The “City Delta” has an associated spatial region that is represented in an IMINT report.




Figure 2. Concepts in the CG linked to the data items from which they are composed

Although the VBS gives us a common representation for all of the information available to the reasoning process, different inference procedures are required for the SU Engine to ultimately derive COA [Pawlowski02]. Each of these inference procedures is implemented as a separate algorithm that can be executed over some subset of concepts in the VBS to further refine and enhance the set of concepts and relationships. As an example, a Mobility algorithm can be executed over the set of PhysicalThings in the VBS to discover relationships such as MovingToward, MovingAway, TravelingAlong, etc. These new relationships can then be used to refine the COA predictions maintained by the SU Engine.

Rather than have these algorithms applied to the VBS at arbitrary intervals or whenever new data (which may be irrelevant to the current mission) enters the VBS, we chose an approach to meta-control that is a significant extension of the Joint Directors of Laboratories Level 4 Fusion Process Refinement functional concept. Our approach is to use Information Needs generated both implicitly (from the roles and missions of the decision-makers and warfighters known to the SU Engine) and explicitly (from the users working with the system) to constrain which algorithms are run over the VBS, and over which subset of concepts they are run.

The selected algorithms are composed into an Agent that is responsible for a sequence of inferences that will satisfy one or more Information Needs. This composition is dynamic in that it does not rely on a pre-enumerated set of processing workflows. We label this collection of functionality Process Refinement.

In the simple example in Figure 3, a new Information Need is obtained, and Process Refinement infers the satisfaction of that need to be a computation of Mobility information for the new enemy unit, followed by a Capabilities analysis of “Platoon Z” and “Enemy Unit A.” An Agent is composed to perform both of these calculations on the referenced nodes in the CG. The results of the execution of this Agent will be stored in the VBS for use by other inference algorithms.




Figure 3. Agent workflow generation as the result of a new Information Need

Once evidence is combined and inferences have been made, the SU Engine needs to disseminate those results to the decision-makers and warfighters that need the information. This process of information sharing is called Shared Understanding. The sheer size of the VBS can preclude sharing every SU Engine output with every user due to limitations in bandwidth and node processing. In addition, results only need to be sent to those warfighters and decision-makers who have an Information Need for that particular result. Our approach to this dissemination, Shared Understanding, leverages LM ATL’s patent-pending Grapevine technology [Farrell03], first developed for the DARPA Small Unit Operations program and matured via six Army programs and internal funding.

Our overall approach to the SU Engine combines data mediation, Information Need generation, Process Refinement, evidence combination and inference algorithms, the VBS, and Shared Understanding into one integrated system.

5. SU ENGINE

The SU Engine is built on industry open standards and designed so that it can be easily integrated into net-centric environments (Figure 4). All of the components that form the SU Engine are designed as services that can be accessed and invoked by other services using the Simple Object Access Protocol (SOAP). These components have been used in web services environments by describing their services using the Web Service Definition Language (WSDL). Semantic definitions for these services are provided in OWL-S to allow other reasoning systems to understand and use the SU Engine services. The result is a system that can be deployed and interoperate with other open standards based systems. The SU Engine’s services also support network security initiatives such as DISA’s Net-Centric Enterprise Services Security Services.




Figure 4. SU Engine (Summary View)

All data entering the SU Engine services is first mediated into a common internal schema using Semantic Mediators available in the SU Engine. The data is also marked with initial pedigree and confidence information at this time. Pedigree contains source, validity, and history information, along with justification for inferred data to aid in substantiating a conclusion to the user.

Shared Understanding maintains the context of users who have logged into the system. This context contains the user's role, security level, mission, and plan, in addition to the user's location and current state. Process Refinement reasons over this context, examining spatial and temporal bounds in the plan, mission type, and the user's echelon and role to form a set of Information Needs. These needs are then prioritized and broken down into a fusion workflow consisting of evidence combination and inferential algorithms. These workflows are instantiated into an Agent and scheduled alongside other agents. Scheduling is performed using algorithms similar to those in [Boddy94], in combination with performance profiles for the components of the workflow.

Inferential algorithms add or delete properties in the data and/or nodes in the CG. These algorithms compute a wide range of values including Kinematics, Affiliation, Capabilities, and Vulnerabilities. They also attempt to discover, through inference, relationships among the concepts in the CG. The absence of relationships on concepts that can be related is an indicator that evidence has either not yet been obtained or not yet been combined. For example, one would expect a Vehicle to have a Driver, and the algorithm could detect this by noting that the type ontology defines the Vehicle class as having a potential relationship called "driver," which can be connected to one or more instances of class Person.

Temporal analysis algorithms attempt to examine sequences of events in the CG and how they impact entities within the VBS. Spatial and temporal aggregation is executed on tracks to look for coordinated behavior and to create Concepts in the CG that represent units [Farrell04]. Unit Comparison algorithms look for the strengths and weaknesses of various pairings of the units in the CG. Mobility algorithms look for likely end points of observed unit movements and estimate time of arrival and expected paths.

The goal of these evidence combination and relationship discovery algorithms is to provide situation understanding via the VBS and to enable the generation of potential enemy COA (ECO). Since we are

interested in producing ECOA for only certain types of entities performing some action with respect to certain types of targets, the attention of the ECOA generation algorithms is restricted to subsets of the CG containing those types of entities, such as enemy aggregates and friendly assets. ECOA generation continues by looking for key chains of relationships that might indicate an ECOA. The potential ECOA are then filtered through a multi-perspective ranking process. The SU Engine takes a decision-theoretic approach to ECOA generation by ranking ECOA based on several factors including confidence and their impact to the Blue force plans. These ECOA results are maintained in the VBS as the situation unfolds.

An example of this ranking is shown in Figure 5. The CG layer of the VBS is shown here with two Situations representing the current situation and an event in the past. Two potential ECOA are being ranked by each of the Tactical, Historical, and Cultural perspectives. Each of these perspectives is looking for key relationship chains that exist in the graph, fall into the domain of that perspective, and support any of the potential ECOA. Although both ECOA are possible from a tactical perspective, historical and cultural analysis support only the “occupy City Delta” ECOA. In this example showing only the likelihood ranking, evidence accumulation favors one ECOA over the other by a large margin. Additional perspectives not shown in this diagram may further influence the ranking.




Figure 5. Multiple Perspective ECOA ranking based on the Conceptual Graph layer of the VBS

When new results are obtained, the Grapevine’s proxies for decision-makers and warfighters determines whether the results overlap with the Information Needs of those users and prioritizes sharing the results. Results are bundled with a confidence and an explanation, allowing recipients to understand the impact of those results without having seen all the underlying data and without having knowledge of the exact fusion algorithms applied.

To provide users a view into the VBS and to allow the entry of additional information, the SU Engine provides graphical interfaces for each layer of the VBS. Examples of these graphical interfaces are shown in Figure 6.




Figure 6. Virtual Battlespace Graphical Interfaces

Through data mediation, the construction of the VBS, fusion of a wide range of data, inference over those fusion products, and the dissemination of inferred results, the SU Engine provides Actionable Intelligence to decision-makers and warfighters.

6. CREDITS

Much of the research, development, and demonstration of the SU Engine has been performed in collaboration with the Army Research Laboratory (ARL). LM ATL would like to thank Mr. Tim Hanratty, Mr. John Dumer, Dr. Windell Ingram, Mr. Robert Winkler, Mr. Larry Tokarcik and Mr. Randy Woodson, all of ARL, for their support.

7. REFERENCES

- [Bobby94] Boddy, M. and Dean, T., "Decision-Theoretic Deliberation Scheduling for Problem Solving in Time-Constrained Environments," *Artificial Intelligence*, 67(2):245-286, 1994.
- [Farrell03] Farrell, W.J., Jameson, S., and Stoneking, C., "Shared Situation Awareness for Army Applications," *Proceedings of 2003 National Symposium on Sensor and Data Fusion*, San Diego, CA, June 2003.
- [Farrell04] Farrell, W.J., "Hierarchical Entity Aggregation for Situation Assessment: Identifying Asymmetric Aggregates," *MSS NSSDF Conference*, Laurel, MD, June 7, 2004.
- [Pawlowski02] Pawlowski, A., Gigli, S., and Vetesi F., "Situation and Threat Refinement Approach for Combating the Asymmetric Threat," *MSS NSSDF Conference*, San Diego, CA, August 2002.
- [Sowa92] Sowa, J.F., "Conceptual Graphs Summary," *Conceptual Structures: Current Research and Practice*, P. Eklund, T. Nagle, J. Nagle, and L. Gerholz, eds., Ellis Horwood, 1992, pp. 3-52.