The Warfighter Associate: Decision-support software agent for the management of intelligence, surveillance, and reconnaissance (ISR) assets

Norbou Buchler*[a], Laura R. Marusich[a], Stacey Sokoloff[b]
[a]Human Research and Engineering Directorate, US Army Research Laboratory, Aberdeen Proving Ground, USA; [b]VELOXITI, Inc., 3650 Brookside Parkway, Alpharetta, GA 30022

ABSTRACT

A unique and promising intelligent agent plug-in technology for Mission Command Systems— the Warfighter Associate (WA)— is described that enables individuals and teams to respond more effectively to the cognitive challenges of Mission Command, such as managing limited intelligence, surveillance, and reconnaissance (ISR) assets and information sharing in a networked environment. The WA uses a doctrinally-based knowledge representation to model role-specific workflows and continuously monitors the state of the operational environment to enable decision-support, delivering the right information to the right person at the right time. Capabilities include: (1) analyzing combat events reported in chat rooms and other sources for relevance based on role, order-of-battle, time, and geographic location, (2) combining seemingly disparate pieces of data into meaningful information, (3) driving displays to provide users with map based and textual descriptions of the current tactical situation, and (4) recommending courses of action with respect to necessary staff collaborations, execution of battle-drills, re-tasking of ISR assets, and required reporting. The results of a scenario-based human-in-the-loop experiment are reported. The underlying WA knowledge-graph representation serves as state traces, measuring aspects of Soldier decision-making performance (e.g. improved efficiency in allocating limited ISR assets) across runtime as dynamic events unfold on a simulated battlefield.

Keywords: data to decisions, mission command, decision support, ISR, resource management, OODA loop, cognitive workload, metrics of performance, associate system, intelligent agent, force synchronization, system interoperability

1. INTRODUCTION

A major tenet of the U.S. Office of Secretary of Defense's "data to decisions" (D2D) initiative and a primary challenge for military commanders and their staff is to shorten the cycle time from data gathering to decisions. The sequence of steps in any "data-to-decisions" pipeline includes data, information, knowledge, decision-making, and action. Performance and effectiveness are curtailed by failures or bottlenecks at any step in this D2D sequence. Effectively managing the entire process requires new and broad systems-level approaches to conceptualize, design, and engineer collaborative systems that flexibly support workflow cycles (i.e. full data-to-decisions pipeline) in a manner that is tailored to multiple end-users to support and achieve organizational goals (i.e. mission) at any given point in their specific and collaborative workflows. In this paper we describe a systems-level prototype capable of providing decision-support—the Warfighter Associate (WA)— developed by ARL, CERDEC, and Veloxiti Inc.. The results of a scenario-based laboratory experiment to capture and validate Mission Command staff performance (intelligence, fires, and maneuver) are described with and without the prototype WA technology.

2. WARFIGHTER ASSOCIATE SYSTEM

The Warfighter Associate (WA) is an intelligent decision support tool under development for Mission Command that considers user's intent, state of the world, and domain specific knowledge to recommend a course of action.[1] The WA uses a doctrinally-based knowledge representation to model role-specific workflows and continuously monitors the state of the operational environment to enable decision-support, providing the right information to the right person at the right time. As a prototype technology, the WA was developed to address the challenges of providing dynamic, mission-tailored decision-support to the Mission Command staff. The Warfighter Associate seeks to improve Mission Command performance and also serves as a novel metric framework to capture and assess staff performance.[1,2] The intent of the WA is to measurably shorten the cycle time from data gathering to decisions. The scientific basis of our "staff-in-the-loop" laboratory validation study is that the WA enables agility by providing dynamic decision-support to the Mission Command staff allowing them to cycle quickly through military decision-making cycles, such as the Observe, Orient,

*norbou.buchler.civ@mail.mil; phone 410-278-9403

Decide, and Act (OODA) loop, with the goal of achieving a high operational tempo and reacting faster than the enemy.[3] The following sub-sections describe the fundamentals of the WA prototype in terms of: (1) knowledge engineering, (2) interactive decision-support, and (3) system dynamics.

2.1 Knowledge Engineering of Military Workflows

Mastery of domain-knowledge engineering is a major enabler of an associate system approach, designed to work in conjunction with a human operator. This approach takes advantage of the fact that cognitive work flows are relatively constrained by the physics of the battlefield. That is, the decision-space is fairly circumscribed by the overarching goals, mission objectives, and the sets of elemental actions and plans to achieve them. As a brief example, given a reported improvised explosive device (IED) event, the associate system provides to the user— here, an intelligence officer (S2) —a recommended list of nearby surveillance assets that could be repositioned to quickly provide eyes-on the event location. The associate system's knowledge representation 'knows' that in response to an IED event this given user is likely to want to reposition known nearby surveillance assets with particular known capabilities. The term 'known' indicates that the associate system is aware of the current state of the world such as the position of and capabilities of the available surveillance assets. In this manner, the associate system is proactive in aligning the likely goals and decision-making authority of the warfighter with the means to achieve them (plans). As currently developed, the Warfighter Associate is able to respond to many normal doctrinal operations, including intelligence surveillance and reconnaissance (ISR) asset management.[1]

The WA is composed of a set of functions that create a closed loop cognitive engine that uses explicitly declared knowledge bases to perform situation assessment, dynamic planning, action execution and coordination across multiple actors, and the development of situation awareness.[7] A properly designed and implemented multi-agent intelligent system must be capable of representing knowledge relevant to all three steps in this process. There are three related knowledge structures in the Warfighter Associate: the Observe-Orient (O-O) graph, the Decide-Act (D-A) graph, and scripted plans that encapsulate procedural knowledge (see Figure 1).

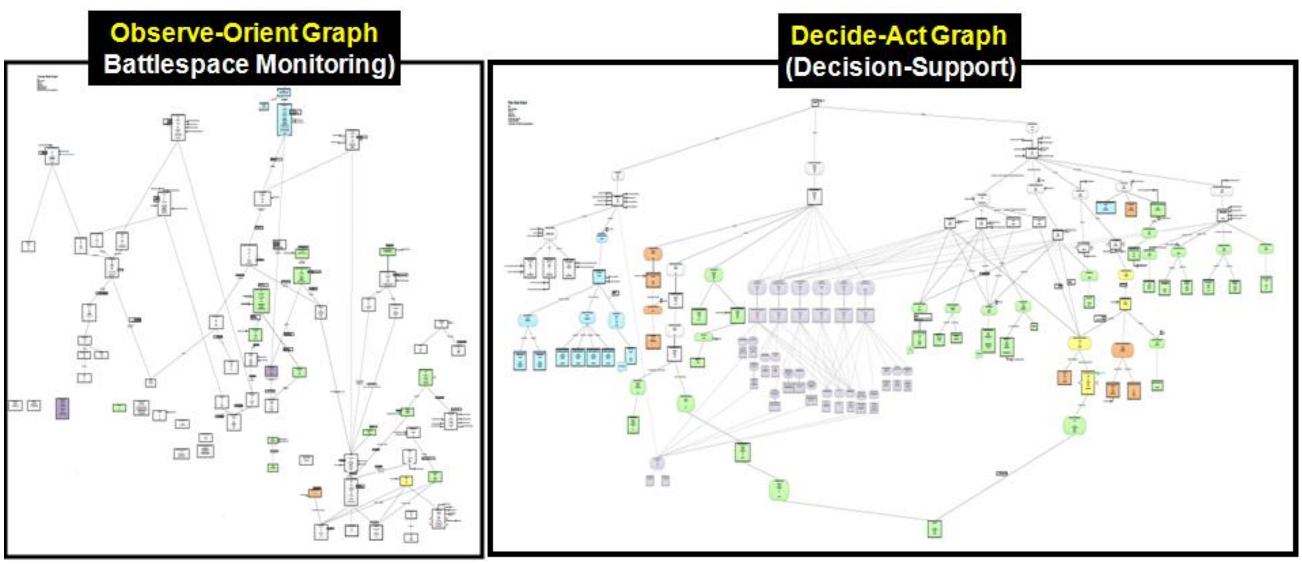

Figure 1. The underlying knowledge-based architecture of the Warfighter Associate is composed of two inter-related and computable hierarchical graphical model structures. The Observe-Orient graph monitors the state of the battlefield in relation to sensor and information network inputs and represents a dynamic description of the factual or perceived state of the world. The Decide-Act graph represents the social-cognitive dimension and is composed of hierarchical plan-goal structures of Mission Command staff workflows.

Plans require information about the current state of the world, and the O-O graph contains general, hierarchical knowledge about the world and is dynamically updated as the operational environment changes. The O-O graph represents general and situational knowledge about the current world state ("who", "what", "when", "where", and "how much”), and linkages among beliefs represented. It is a hierarchical and dynamic description of the factual or perceived state of the world and can include many types of information. This graph provides a means for distinguishing between beliefs about the state of the environment and its true state. It also supports the representation of uncertain relationships between dynamic concepts using Bayesian logic. Beliefs are dynamically updated from observations of incoming data (sensors, information systems, and human inputs).

The D-A graph is a model of means (plans) to achieve end states (goals) and is represented as a collection of hierarchical plan-goal task structures. The D-A graph allows a principled separation between the desired future state (goal) from the means of attaining it (plan). A collection of goals and possible plan to achieving each goal makes up a course of action. Each child plan of a goal is a possible means to achieve the parent goal. Plan nodes are then decomposed into sub-goals, with decomposition in this manner continuing until the level of basic human-machine interactions (actions) is reached.

Scripted Plans represent procedural knowledge—the specific set of steps to pursue a particular course of action. These three knowledge structures work together in a loop, the current state of the world is used to reason about plans, steps for plans are provided, and the execution of plans (recommended or not) can change the state of the world.

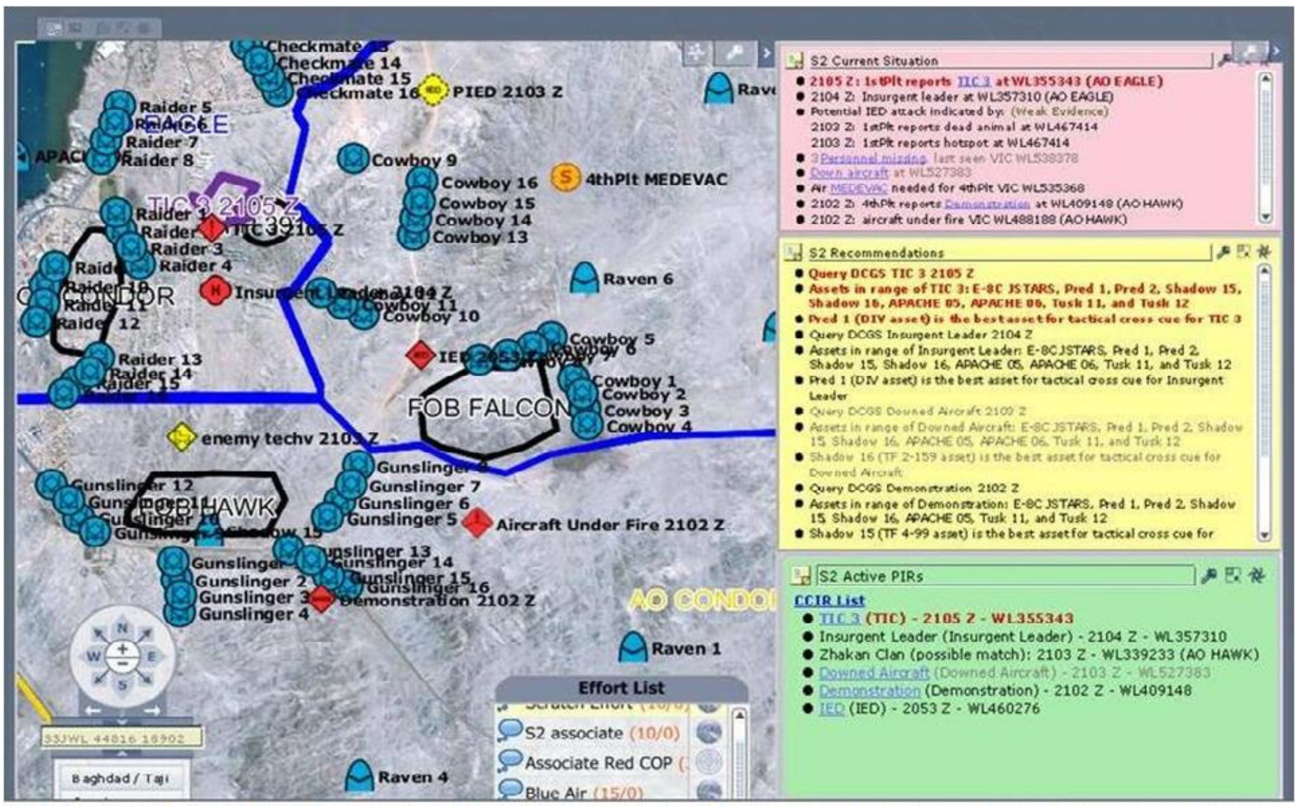

Figure 2. Warfighter Associate screenshot for the Intelligence Officer (S2) as implemented in the Command Post of the Future (CPOF) Army Mission Command System for our experimental scenario. In response to significant events on the battlefield reported in military chatrooms, maneuver graphics are automatically generated and highlighted (purple arrow). Three “stickies” are automatically populated with information corresponding to: (a) the current situation (pink box), (b) decision recommendations (yellow box), and (c) active priority information requirements (green box).

2.2 Interactive Decision Support

The demands to human information processing, which involve observing and monitoring communication streams, as well as organizing, combining, and evaluating data and intelligence, can quickly overwhelm cognitive capabilities, given limited working memory.[4,5,6] As the basis for interactive decision support, the OODA loop defines an essential cognitive cycle and collaborative dimensions of Mission Command performance. It defines how the sensor and information network inputs (Observe-Orient) feed into the social-cognitive (Decide-Act). A properly designed and implemented decision-support system must be capable of representing knowledge in all four stages of this process.

The WA is run within Intelligent Presentation Services (Veloxiti Inc.), which is a Data Dissemination Service (DDS) based Service Oriented Architecture (SOA). In our configuration, the WA obtains asset positions from a Publish and Subscribe (PASS) Server and learns about combat events, such as IEDs (Improvised Explosive Devices) and TICs (Troops in Contact) by monitoring chatroom inputs with natural language processing.

A screenshot of the WA is shown in Figure 2, as implemented in our experimental scenario for the intelligence officer (S2). The Warfighter Associate— via natural language processing— interprets tactical reports or chat messages to establish real-time situational awareness. As an event occurs, the WA monitors and assesses information for relevancy based on location and status. This processing occurs in the Observe-Orient graph and serves as a knowledge filter that can also fuse disparate information. In fact, by reading chatroom inputs the WA can be aware of the event several minutes before it is physically entered into an Army Mission Command System. If considered relevant, the WA triggers the Decide-Act Graph and automatically provides updates to the common operational picture, implemented here in the Command Post of the Future (CPOF), and highlights the significant event for the user with a purple arrow (see Figure 2). A series of information panels (i.e. "stickies") are populated with relevant information to support situational awareness, provide decision support, and track critical pieces of information to relay to the Commander; these are shown in the right colored panels of Figure 2 and correspond to the Current Situation, Recommendations, and Active Priority Information Requirements (PIRs), respectfully. In our screenshot example, the WA provides recommendations for ISR tasking for a Troops in Contact (TIC) event to the intelligence officer, since no ISR assets are currently providing coverage of the event. How the system provides this recommendation is addressed in the next section.

2.2.1 Intelligence, Surveillance and Reconnaissance (ISR) Recommendations

One of the most important types of decision-aiding provided by the WA prototype is intelligence, surveillance, and reconnaissance (ISR) recommendations. When the WA learns about a combat event, it calculates how well each ISR asset can cover the event and recommends the top three assets. The ISR asset selection algorithm scores each asset based on several factors, including time on target and sensor capabilities. The scores for the individual factors are weighted based upon importance and combined into a simple percentage to provide an easily understandable rating to the user.

The decision-support algorithm for ISR recommendations rules out certain assets as inappropriate for a specific event, for instance if the wind speed exceeds performance limits or if the time needed for the ISR asset to travel to target exceeds a predefined time-limit. The remaining candidate set of ISR assets are scored on each of the following factors:

- Arrive Time Score (time on target): Assets are scored based on how quickly they can be within sensor range of the target. The assets' average speed, sensor coverage area, and distance to the target are considered. When evaluating sensor coverage area, the WA accounts for assets requiring a certain offset from the target.
- Asset Manning: Manned assets are considered preferable to unmanned assets
- Sensor Fidelity: Each asset type is assigned a maximum and minimum sensor fidelity rating on a scale of 1-10. This is used to degrade an asset's sensor fidelity score for targets at the edge of its coverage radius. For asset types where this distinction is inappropriate, the same value is assigned to the minimum and maximum sensor fidelity ratings.
- Stealth Rating: Asset types are rated on a scale of 1-10. Assets with a higher stealth rating are given higher scores.
- Survivability Rating: Asset types are rated on a scale of 1-10. Assets with higher survivability rating are given higher scores.

After scoring each factor, the individual scores are combined based on relative importance. For example, an asset's sensor fidelity score has a larger impact on its combined score than the asset's survivability score. The combined scores are adjusted onto a scale of 1-100 for ease of interpretability.

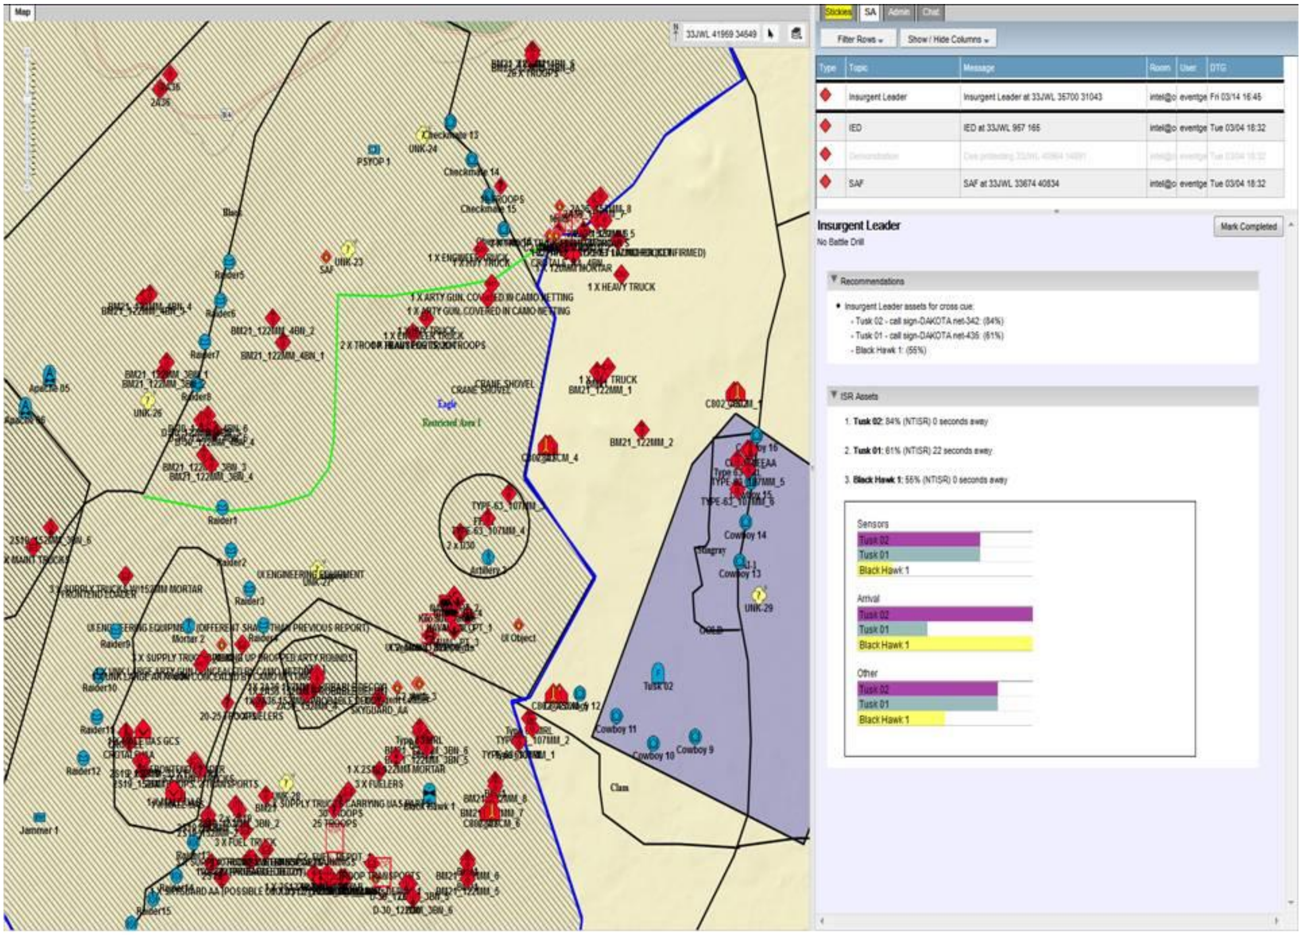

Figure 3. ISR Asset Recommendations provided by a configuration of the Warfighter Associate called the Commanders Critical Awareness Tool (C-CAT).

2.2.2 User Interactions

The WA recommends the three highest scoring assets for each collection target. Each asset recommendation includes the asset's score, a simple explanation of how the score was calculated, and a communication link to the asset owner, if known (see Figure 3). As the situation changes, the WA re-evaluates asset scores and outputs updated recommendations, when necessary. While much of what the WA considers when tasking assets is inherently fixed for each asset type, certain characteristics may change over time or even the state of the battlefield. New ISR collection assets may become available and existing collection assets may be grounded or moved outside of the area of operations. The WA provides users with the ability to add or remove asset types and adjust ratings for existing asset types (see Figure 4). In sum, the WA ISR asset recommendation algorithm provides a mechanism to capture ISR asset expertise from advanced users and make it available to subscribers needing to quickly task assets in response to battlefield events, helping inexperienced, tired, or stressed Soldiers function similarly to ISR experts.

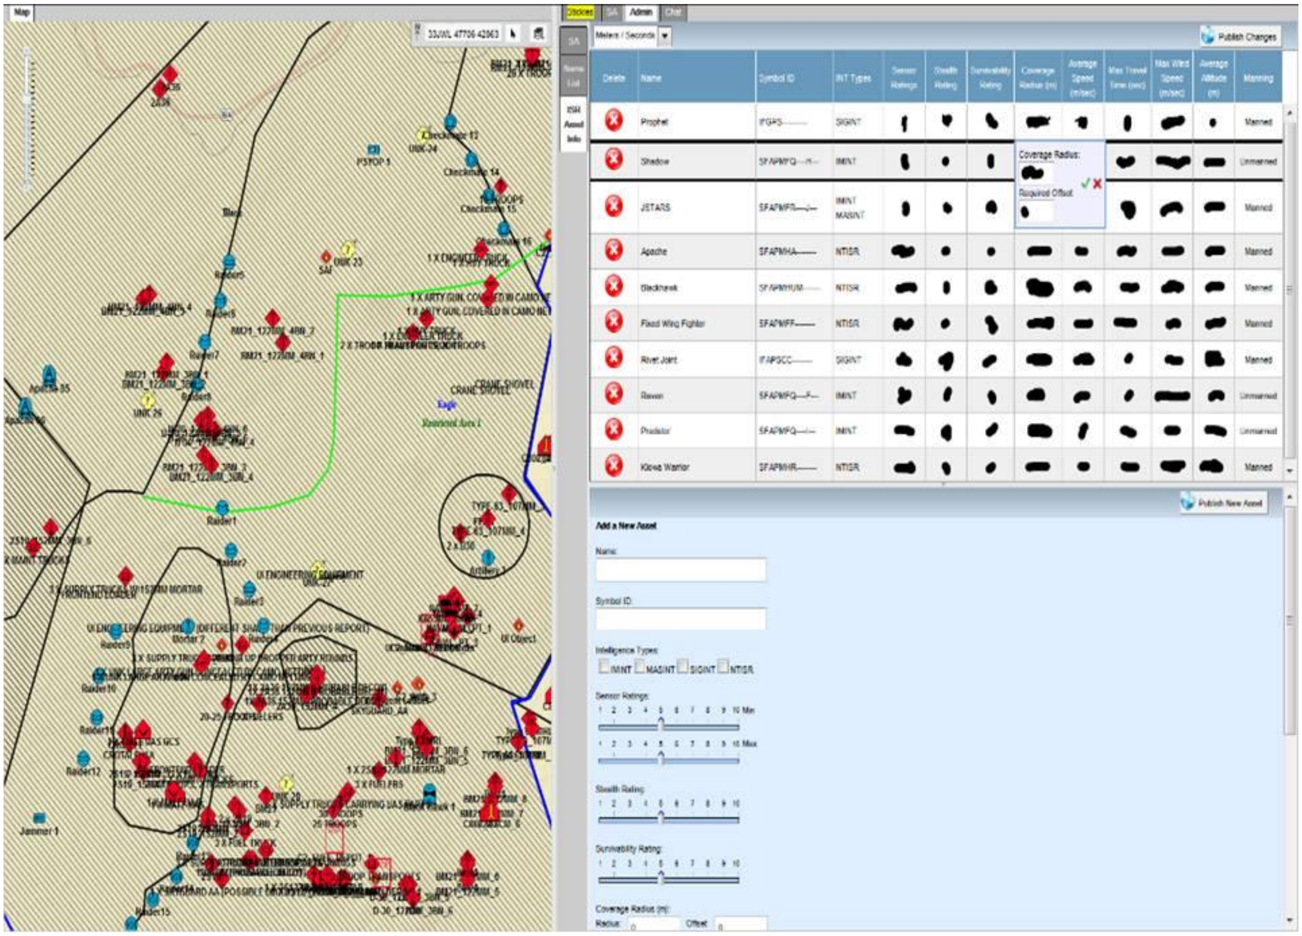

Figure 4. ISR Asset Configurations panel in the Warfighter Associate C-CAT toolset with values hidden.

2.3 Reciprocal System Dynamics

Feedback loops are essential to achieve dynamical system behavior. Nodes (Beliefs, Plans, and Goals) have dynamic life cycle states in the WA representational architecture. In the O-O graph, life cycle states represent the prominence of a belief, allowing beliefs that are no longer supported by evidence to become forgotten. The feasibility and desirability of any particular course of action (path through the D-A graph) will depend on the activated set of beliefs based on the ongoing processing of system inputs. A belief monitoring mechanism provides the dynamic connection between the state of the information environment and possible courses of action. For example, a high-level plan “to survive” may subscribe to messages about threats. If an IED is reported, the subscription monitor will fire, triggering planning. The Soldiers may consider multiple ways to accomplish the goal to diffuse the IED. The specific way selected will activate monitors for the relevant information in the O-O graph. For example, a plan to send a platoon will monitor the O-O graph for obstacles on the route between the platoon’s location and the location of the IED. If a conditional statement in a monitor is found to be true, the detected event may be used to transition a plan or goal to a different life cycle state. In this way, plans and goals that are no longer feasible can be replaced as the state of the beliefs change over time. Actions, by definition, change the world state, which causes the O-O graph to be updated. This may result in re-planning, which may then cause additional actions to be performed. Examples of actions include invoking a route planner, database query, unit movement, and re-tasking a sensor asset.

3. METHODS

The experimental design consisted of two groups, one run with and one without the prototype WA technology. Each humans-in-the-loop experiment was conducted in our laboratory facility at Aberdeen Proving Ground and featured a detailed four-hour military scenario. In separate weeks, two groups of three non-commissioned officer (NCO) participants completed the experiment, assuming the role of an operations officer (S3), an intelligence officer (S2), and a fire support officer (FSO). The level of analysis and human agency was the brigade. All of our Soldier participants were experienced veterans with at least one combat tour (average of two tours) in counter insurgency (COIN) operations in Iraq and/or Afghanistan and were accomplished in their role-position, with an average time in service of 6.2 years (range of 3.2 to 9.8 years). These soldiers were from the same organic unit, understood their role positions, and performed well and cohesively within the combat vignettes of the scenario. The chatroom inputs and unit actions of parent units (higher echelons) and sub-ordinate units were scripted, for instance four sub-ordinate battalions as well as host-nation military and police units acting within an area of interest to the larger, parent military unit, a brigade. The simulation environment included unit and asset movements (i.e. simulating blue force tracker), chatroom inputs, and a repository full of intelligence products, ISR schedules, preplanned graphics, etc.

The experiment was conducted as a two-day event. The first day was devoted to reading the soldier participants into the scenario, which included a "road to war" briefing by the scenario co-author and officer-in-charge (Lieutenant Colonel, O-5). Soldiers were given a full dossier of information including intelligence summaries, preplanned fires, ISR synchronization matrices, and key coordinating maneuver graphics supporting their upcoming responsibilities. All three soldiers were knowledgeable of the mission command system used to maintain a common operational picture during the experiment, the Command Post of the Future (CPOF). Soldiers received three hours of training using CPOF to be assured of a base level of soldier performance using the system. The second day consisted of the four hour experimental run (interspersed with three 20 minute breaks or pauses) followed by an after-action review conducted by the officer-in-charge. The experimental scenario consisted of COIN operations, and key events/threads included: (a) a host nation convoy transporting raw "uranite" fissile material through their area of operations, (b) finding, fixing and conducting a cordon and search operation versus a high value target (HVT), (c) a planned enemy attack on host-nation parade dignitaries using an improvised explosive device (IED), (d) responding to enemy mortar attacks, (e) supervising the repositioning of a critical radar asset, (f) missing ammunition, and (g) recovery of a downed civilian aircraft in a hostile border area.

Since the action took place within a brigade area of interest, the soldiers assumed responsibilities for the operation of a brigade tactical operation center (TOC). They performed several functions. First, they monitored the battlefield environment and Mission Command inputs via six chat rooms. Four of the chat rooms were organized by Mission Command function (manuever, intelligence, fires, and sustainment), one represented Joint Special Operations Command and higher echelons of command (JSOC/HICON), and another simulated liaison lash-ups with host nation units (Gorgasnet). The experimental scenario consisted of three hours of approximately 800 scripted chat room inputs from over 40 scripted actors representing both routine chat room messages and those in response to scenario events. Second, the soldiers processed all of the incoming chat room information and used CPOF to maintain a common operational picture (COP) of significant events occurring within their area of operations. Third, the soldiers interacted with the synthetic (i.e. scripted) actors in the chat rooms, for instance by requesting additional information, tasking subordinate battalion units, retasking ISR assets, coordinating with the host nation units, or issuing regular reports to higher echelons of command. Responses to requests for information were handled by an experimental team, the "white cell", who inserted inputs into the various chat rooms assuming the names of the various synthetic actors. Friendly unit locations were displayed on the COP to simulate the Mission Command functionality of Blue Force Tracker.

The Soldier participants each sat in separate sound attenuated experimental chamber in front of two screens, one populated with the inputs to the six chat rooms and the other with the CPOF Mission Command System that functioned as the brigade COP. The isolation from one another forced the participants to communicate via voice (headset) or the chat room interface. Voice communications were in a common channel. The data recorded during the scenario experiment consisted of a log of all voice and chat room communications, screen video recordings, and the log output of Warfighter Associate. In the condition without WA decision aiding, the WA was run only in the background, allowing for data collection.

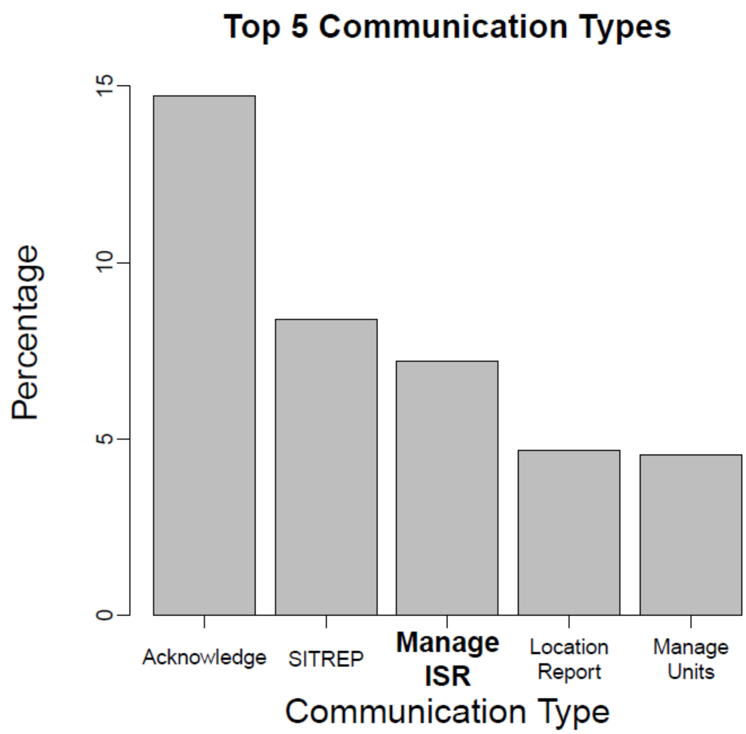

Figure 5. Five most frequent types of communication (including messages from participants, experimenters, and script) throughout the experimental scenario.

4. RESULTS

4.1 Staff Communications

The experimental scenario proved to be quite challenging to the Soldier participants, requiring a high-degree of staff collaboration over chat and voice communications. Soldier chatroom inputs and voice communications were transcribed and categorically coded. Figure 5 shows the five most common categories of communication in one of the experimental runs. It is evident that in this experimental scenario the Soldier participants were required to devote a large part of their communication and collaboration to managing ISR assets.

Figure 6 provides a raw timeline of staff voice and chat room communications, sorted by experimental scenario thread and corresponding ticks of significant chat room inputs. The significant chat room inputs were often reports relayed into chat rooms by battalion operators in contact with subordinate units at the tactical level or collection managers of ISR assets. It is evident that the Soldier participants were challenged with dealing with multiple overlapping events. However, note that the communication pattern of the maneuver (S3) NCO suggests that he was able to successfully manage multiple communications threads about multiple threats as they occurred. In addition, the data also suggests that the maneuver (S3) NCO sustained attention to particular scenario threads such as the uranite convoy moving through the area of operations, perhaps indicating command relevance and prioritization of effort.

4.2 Mission Command Metrics of Performance: State Trace of Knowledge Activation

The WA system captures Mission Command performance by providing detailed state traces of activated domain knowledge (i.e. specific node activations) over the course of the scenario that correspond to particular staff workflows. The significance of our approach is that it provides novel continuous and real-time metric of staff cognitive workload (Figure 7) and ISR resource management (Figure 8).

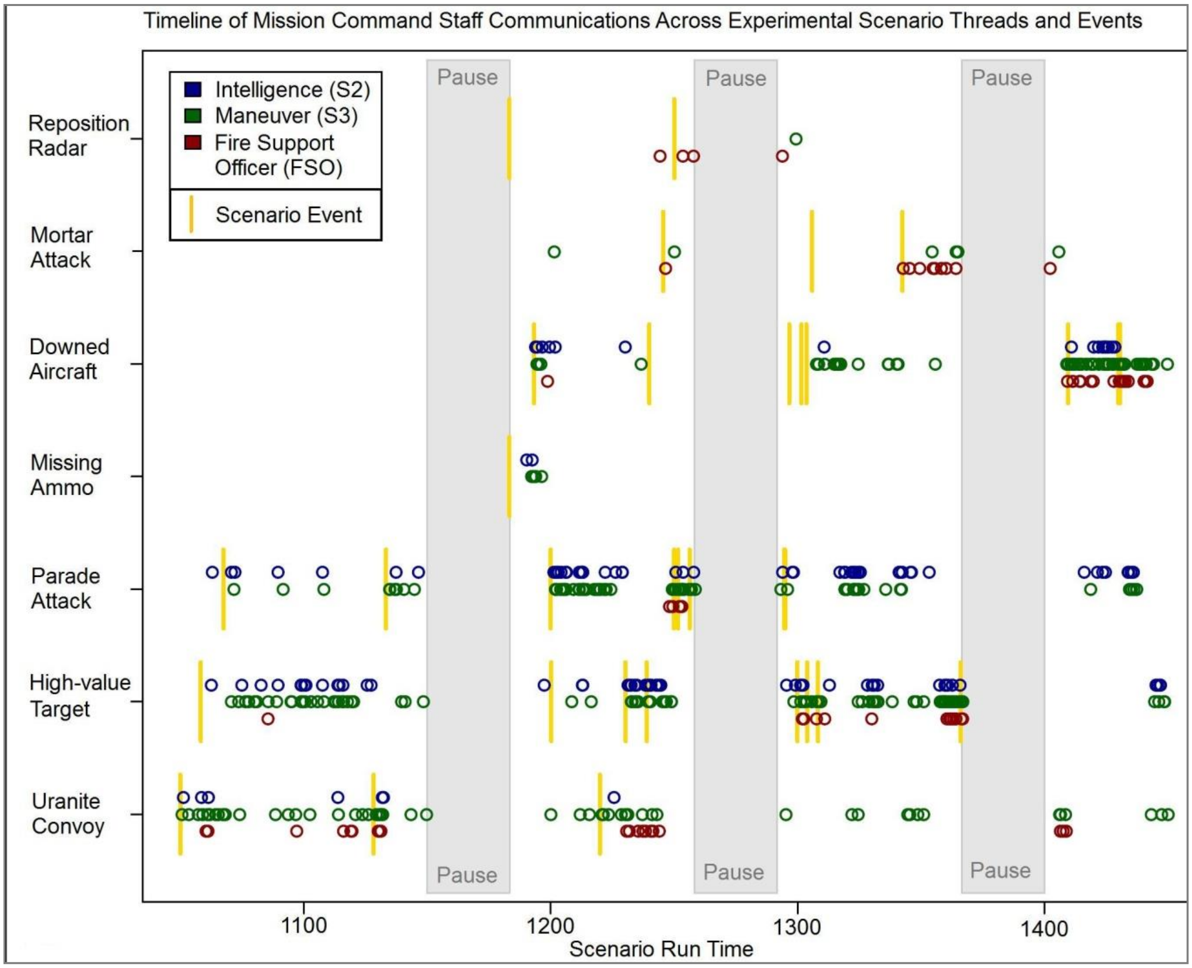

Figure 6. Timeline of chatroom inputs and voice communications (open circles) across scenario clock run-time for the Mission Command staff arrayed vertically by major scenario threads and by role position (intelligence, maneuver, and fire-support). The injects of significant scenario events are denoted by hash-marks.

In Figure 7, the activation count of one goal node from the D-A graph is depicted across the experimental scenario. This high-level goal node <*Each Threat Managed*> tracks current threats in the area of operations. As a metric, the activation profile of this high-level goal node represents the many different workflows of responding to threats. The activation count increases steadily as more and more compounding events occur in the area of operations. In representing the number of active workflows, the <*Each Threat Managed*> goal node is perhaps a good candidate for a measure of cognitive workload for the staff. Cognitive workload is recognized as an important aspect of behavior based upon studies of work performance and subjective assessments of job demands.[8,9] Deactivations or satisfaction of particular instances of the goal node <*Each Threat Managed*> are depicted as well and occur when threats are successfully managed in the area of operations. The WA logs provide a detailed accounting of exactly which threat was successfully managed at a particular time-step.

Managing and prioritizing limited ISR resources is an important Mission Command function. In Figure 8 we examine the activation count of a goal node from the D-A graph <*Threat Surveillance Adequate*> that is activated when an active threat exists in the area of operations and there are **no** ISR assets in range to cover the threat. The WA system monitors the current location of all known ISR assets and is knowledgeable of their capabilities. The <*Threat Surveillance Adequate*>goal is satisfied when an ISR asset moves in range to cover the threat, which typically requires that the staff task or re-task ISR assets. In Figure 8, the activation count of the <*Threat Surveillance Adequate*> goal is depicted across our experimental scenario in the condition without WA decision aiding. Note that for the first half of the

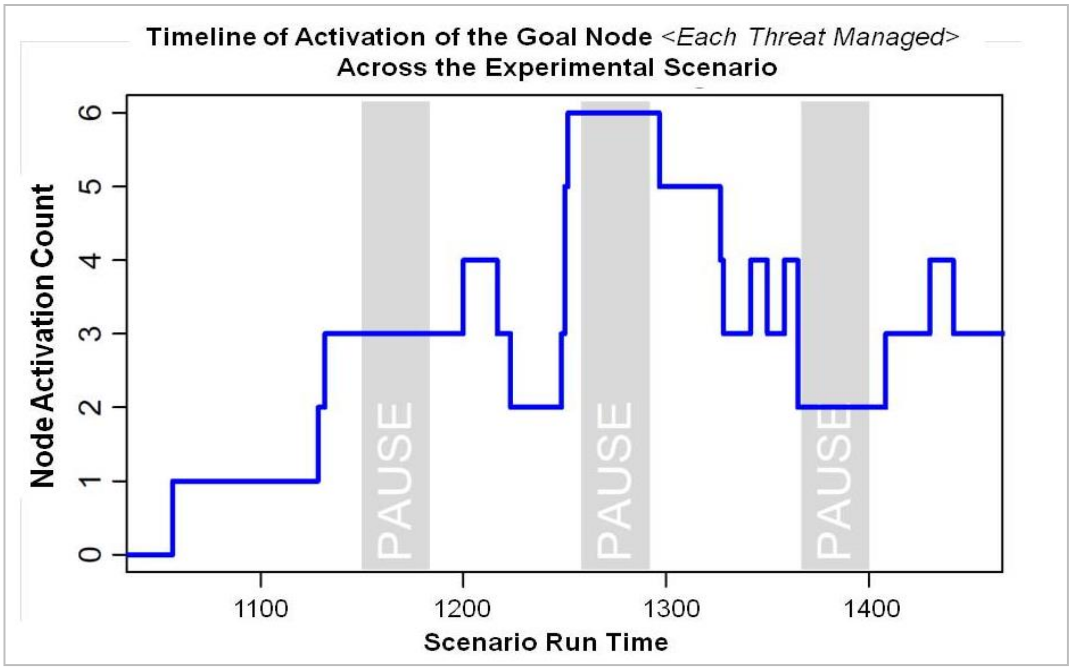

Figure 7. Timeline of the activation and deactivation (satisfaction) counts for a specific goal node <Each Threat Managed> depicted across one run of the entire experimental scenario.

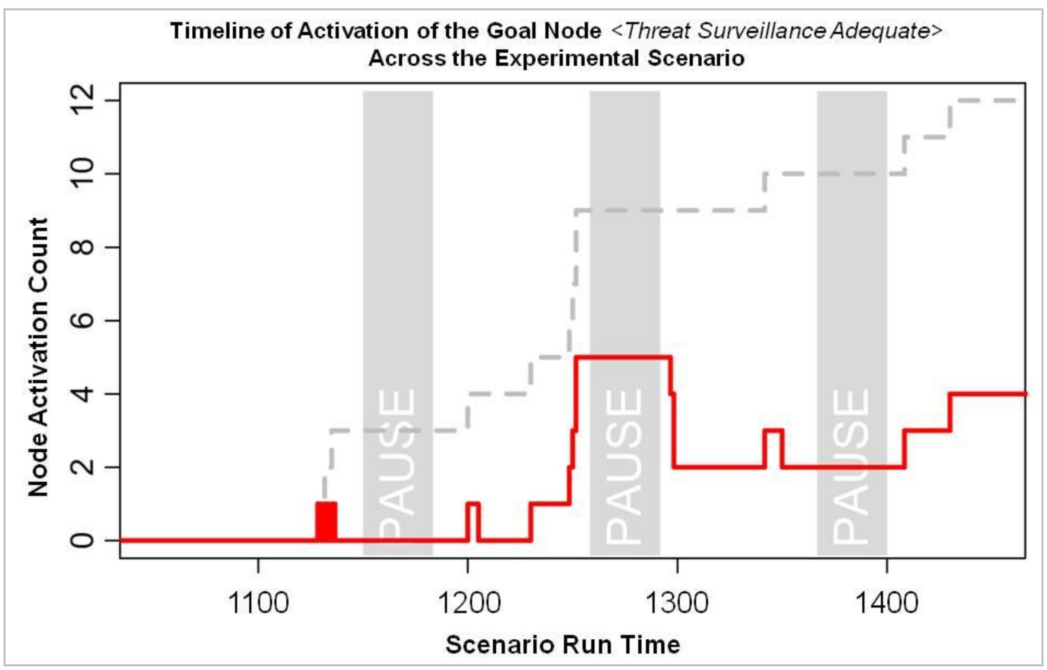

Figure 8. Timeline of the activation and deactivation (satisfaction) counts for a specific goal node <Threat Surveillance Adequate> depicted across one run of the entire experimental scenario. The <Threat Surveillance Adequate> goal node is triggered when there are no ISR assets in range of an existing threat in the area of operations. The gray dotted line represents a baseline dry run condition with no Soldier participation where ISR assets have not been tasked or re-tasked in response to scenario events.

experimental scenario (until around the 1230 time-step) the <*Threat Surveillance Adequate*> goal is efficiently satisfied, staying at zero. After this midpoint time-step, the staff is unable to satisfy all the surveillance demands in the area of operations given limited ISR resources. This provides further evidence that managing ISR assets made up a substantial portion of the staff's effort in this scenario. In fact, the demand for ISR allocation becomes nearly overwhelming in the second half of the scenario. This high level of demand makes for a good test of the WA's ability to help Soldiers manage their workload and speed up their OODA decision cycle, examined below.

4.2 Decision Timing

In order to examine the WA's effect on the speed of the OODA decision cycle, we first examined the amount of time that elapsed between the occurrence of a threat requiring ISR allocation and the tasking of an ISR asset by one of the Soldier participants. Figure 9 shows the results of this analysis – the time to task ISR assets was apparently longer and more variable in the experimental condition run without the WA than in the experimental condition with the WA. Inferential statistical tests, however, lacked sufficient statistical power given the small sample size of our experiment. The rather large variation in performance in the without WA condition implies that by comparison, the group with the WA benefitted from a structured intelligent systems approach to guide Mission Command staff decision-making.

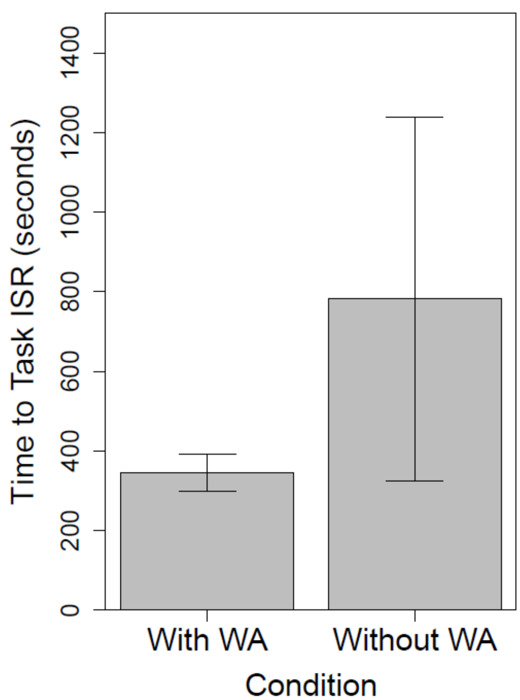

Figure 9. Time to task ISR assets after the onset of an event requiring ISR allocation in both experimental scenarios (with and without the Warfighter Associate). Error bars represent standard errors of the mean.

We also examined the amount of time that elapsed between the occurrence of a threat and the first communication about that threat from any member of the staff. This represents the front-end of the military decision-making cycle, the Observe-Orient portion of the OODA loop. The results of this analysis are shown in Figure 10. Here, the group who had the WA decision-aiding was marginally slower to communicate about threats than the group without the WA. The intuition here is that the group with the WA may have been slower in achieving situational awareness of threats. However, upon examination of what exactly was transmitted in these critical first communications, we found substantial differences between the two groups. As shown in Figure 11, the group with the WA were more likely to transmit command communications (the Decide-Act cycle of the OODA loop), whereas the group without the WA were more likely to transmit messages providing acknowledgments and sharing information with their fellow participants (the Observe-Orient cycle of the OODA loop). This result provides evidence that the WA allowed the staff to quickly achieve situational awareness and push farther into their OODA decision-making loop faster than the group without access to the

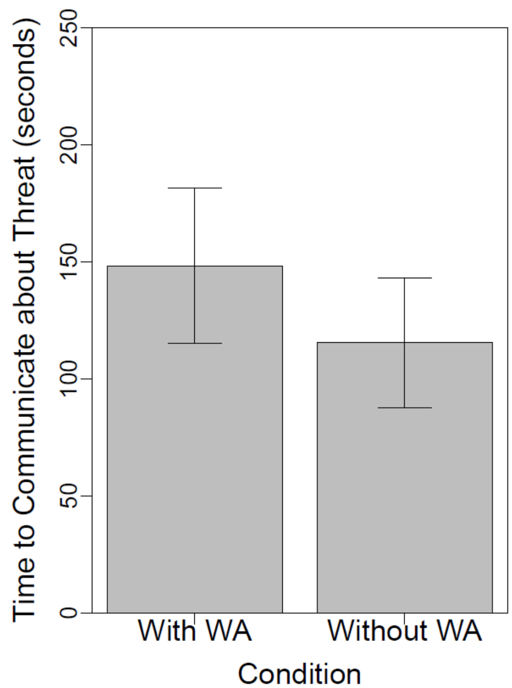

Figure 10. Time to communicate about a threat requiring ISR allocation in both experimental scenarios (with and without the Warfighter Associate). Error bars represent standard errors of the mean.

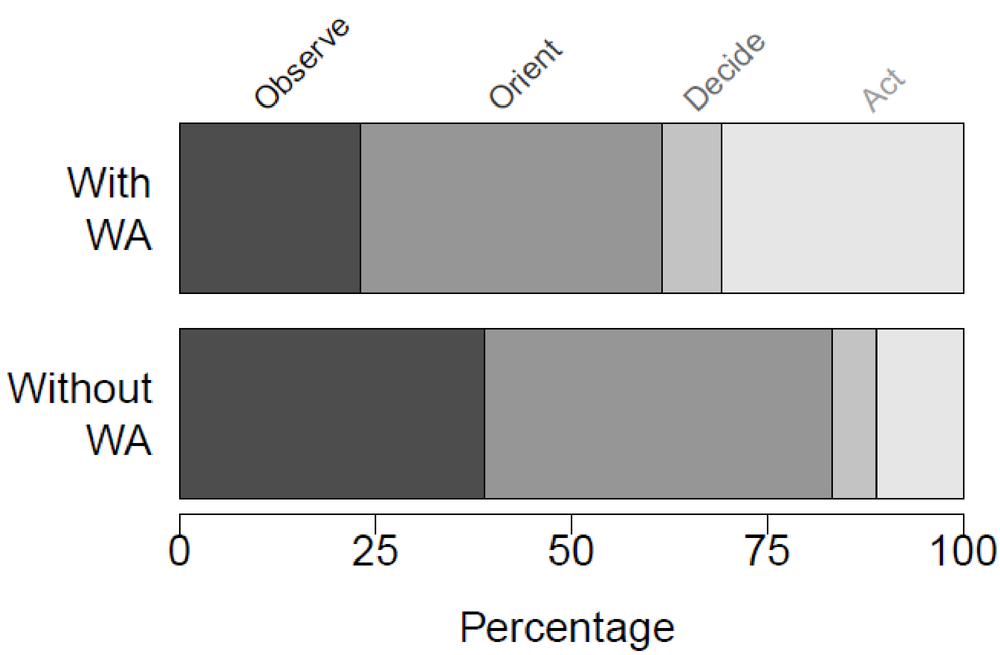

Figure 11. Percentage of first communications about threats devoted to each part of the OODA loop. The data provide evidence that the WA allowed the staff to reach more advanced stages of the OODA loop faster.

WA's decision-making. Thus, the WA provides rapid situational awareness by automatically alerting the Mission Command staff to emerging threats on the COP and providing a running summary of the current situation. Further, it is

WA's decision-making. Thus, the WA provides rapid situational awareness by automatically alerting the Mission Command staff to emerging threats on the COP and providing a running summary of the current situation. Further, it is likely that the role-specific decision-support capability (i.e. Recommendations box; see Figure 2)—to include ISR recommendations— and the critical information summary (i.e. Active PIRs box; see Figure 2) fostered Mission Command staff synchronization and a faster resource allocation and decision-making cycles.

5. CONCLUSIONS

The Warfighter Associate (WA) is an intelligent agent software system that uses doctrinally-based knowledge engineering to provide role-specific decision-support to a Mission Command staff. The WA prototype addresses a major tenet of the U.S. Office of Secretary of Defense's "data to decision" initiative[10] and the primary challenge for military commanders and their staff to shorten the cycle time from data gathering to decisions. At a staff level, our analysis of a controlled scenario-based "humans-in-the-loop" laboratory experiment demonstrated that a Mission Command staff—composed of a maneuver officer (S3), intelligence officer (S2), and fire-support officer (FSO)— with the WA prototype technology was able to cycle more fully through their military Observe-Orient-Decide-Act decision-making loop than a comparable staff without the WA prototype (baseline condition). In a focused metric analysis, the WA supported Mission Command staff was marginally slower to communicate and orient to new threats, however, they issued more commands and acted more quickly and appropriately in tasking ISR assets with the needed sensor coverage. ISR recommendations is one of the most important types of decision-aiding provided by the WA prototype. Our ISR asset selection algorithm scored each potential asset based on several factors, including time on target and sensor capabilities, to recommend a top three.

An important research objective was to develop the scientific experimental infrastructure and metric framework necessary to capture Mission Command effectiveness. Objectively measuring the effectiveness of Mission Command is a formidable Army challenge given the difficulty of capturing staff performance and complex, overlapping workflows. Military operations are inherently complex human endeavors as Army commanders and their staffs collectively face difficult, stressful, and dynamic challenges in managing battlefield operations. Warfare is chaotic and incredibly complicated, and resolving the attendant ambiguity on the battlefield is a cognitive challenge of the first order. Using the WA system, we demonstrated that it provides a metric solution by automating collection and the aggregation of data derived from field or laboratory based studies. The WA provides capabilities to improve and evaluate Mission Command performance, including difficult to quantify concepts such as cognitive workload and ISR asset management. Essentially, the knowledge representation architecture of the WA provides a state trace of activation of battlefield events and soldier workflows, represented in the Observe-Orient and Decide-Act graphs, respectfully. This constitutes a potentially important methodological breakthrough in developing objective measures of Mission Command performance that can be collected in real-time and unobtrusively. Metrics are essential to developing, evaluating, and improving soldier-system interfaces and performance; and in our case, was used to experimentally validate performance gains using a prototype WA decision-support technology.

Finally, one potential additional application of the WA is as a training technology. The alerts and decision aids offers Soldier's support in Mission Command and performance can also be evaluated and used during debriefings and after-action reviews. The WA could reduce training time, and thus training costs, while increasing Mission Command performance through situated practice in scenario-based training and in the application of doctrinal knowledge. In terms of practical utility, the WA also potentially addresses organizational and environmental challenges in providing decision-aiding to Soldiers whom are inexperienced, fatigued, or both.

REFERENCES

[1] Buchler, N., O'Neill, D., Sokoloff, S., and Bakdash, J. Z., "The Warfighter Associate: Decision-support and metrics for Mission Command," Technical Report ARL-TR-6309 (2013).

[2] Buchler, N., Marusich, L., Bakdash, J. Z., Sokoloff, S., and Hamm, R., "The Warfighter Associate: Objective and automated metrics for Mission Command," Proc. International Command and Control Research and Technology Symposium (2013).

[3] Boyd, J., "Organic design for command and control," In Discourses on Winning and Losing, unpublished slide presentation, Marine Corps University Research Archives, charts 16, 25 (1986).

[4] Anderson, J. R., Reder, L. M., and Lebiere, C., "Working memory: Activation limitations on retrieval," Cognitive Psychology 30, 221-256 (1996).

[5] Baddeley, A., "Working memory," Science 255, 556-559 (1992).

[6] Cowan, N., [Working Memory Capacity], Psychology Press, New York, NY (2005).

[7] Endsley, M., and Garland, D., [Situation Awareness: Analysis and Measurement], Lawrence Erlbaum, Mahwah, NJ (2009).

[8] Wickens, C., and Hollands, J., [Engineering Psychology and Human Performance], 3rd ed., Prentice Hall, Upper Saddle River, NJ (2000).

[9] Wierwille, W., and Eggemeier, F., "Recommendations for mental workload measurement in a test and evaluation environment," Human Factors 35, 263-281 (1993).

[10] Swan, J. and Hennig, J., "From Data to Decisions," Army Acquisition, Logistics & Technology, Jan-Mar (2012).