

AI on the Battlefield: an Experimental Exploration

Robert Rasch

Battle Command Battle Lab
ATTN: ATZL-FD-BC
415 Sherman
Ft. Leavenworth, KS 66027
raschr@leavenworth.army.mil

Alexander Kott

BBN Technologies
11 Stanwix St., Suite 1410
Pittsburgh, Pa 15222
akott@bbn.com

Kenneth D. Forbus

Northwestern University
CS Department
1890 Maple Avenue
Evanston, IL 60201
forbus@northwestern.edu

Abstract

The US Army Battle Command Battle Lab conducted an experiment with the ICCES system -- an integrated decision aid for performing several critical steps of a US Army Brigade Military Decision Making Process: from capturing a high-level Course of Action to producing a detailed analysis and plan of tasks. The system integrated several available technologies based largely on AI techniques, ranging from qualitative spatial interpretation of course-of-action diagrams to interleaved adversarial planning and scheduling. The experiment dispelled concerns about potential negative impacts of such tools on the creative aspects of the art of war, showed a potential for dramatic time savings in the MDMP process, and confirmed the maturity and suitability of the technologies for near-future deployment.

Decision Making at a Brigade Command Post

A US Army Brigade includes an impressive range of assets and capabilities: thousands of professional soldiers and officers, hundreds of combat and support vehicles, helicopters, sophisticated intelligence and communication equipment and specialists, artillery and missiles, engineers, medical units, repair shops, and much more. In a battle, these assets may perform hundreds of complex tasks of multiple types: collection of intelligence, movements, direct and indirect fires, construction of roads, bridges, and obstacles, transportation and handling of supplies, managing civilian population, command and control, and so on.

Detailed planning of a military operation -- whether a battle with an enemy or a peacekeeping operation -- requires an intensive effort of highly trained professionals, the Brigade planning staff. To accomplish this effort, the Army teaches and uses a methodologically rigorous process called the Military Decision Making Process (MDMP) (Department of the Army 1997).

The process is typically performed by a primary staff of 4-5 officers, typically ranging in ranks from captains to lieutenant colonels, with the support of a considerable sized subordinate staff, over a period of several hours. The

physical environment often consists of a tent extended from the back of one or several Army's light trucks, or armored command and control vehicles, folding tables, maps hung on the walls of the tent and covered with acetate sheets on which the officers draw symbols of units and arrows of movements.

To describe the process for the purposes of this paper, let's focus on a few salient aspects of it. The input for their effort comes usually from the unit Commander in the form of the commander's intent, concept of operation and desired end-state for the operation-- a high-level specification of the operation. This information is then used to develop COA (Course of Action) sketches and statements. In effect, such sketches and statements comprise a set of high-level actions, goals, and sequencing, referring largely to movements and objectives of the friendly forces, e.g., "Task Force Arrow attacks along axis Bull to complete the destruction of the 2nd Red Battalion." With this input, working as a team for several hours (typically ranging from 2 to 8 hours), the members of the planning staff examine the most critical elements of the friendly COAs in minute detail. The process involves planning and scheduling of the detailed tasks required to accomplish the specified COA; allocation of tasks to the diverse forces comprising the Brigade; assignment of suitable locations and routes; estimates of friendly and enemy battle losses (attrition); predictions of enemy actions or reactions, etc. This latter process is referred to as the wargaming process.

The outcome of the process is usually recorded in a synchronization matrix format (FM 101-5 1997), a type of Gantt chart. Time periods constitute the columns and functional classes of actions, such as the Battlefield Operating Systems (BOS), are the rows (see Fig. 3). Examples of BOS include Maneuver, Combat Service Support (e.g., logistics), Military Intelligence, etc. The content of this plan, recorded largely in the cells of the matrix, includes the tasks and actions of the multiple sub-units and assets of the friendly force; their objectives and manner of execution, expected timing, dependencies and synchronization; routes and locations; availability of supplies, combat losses, enemy situation and actions, etc.

How Decision Aids Can Help MDMP

It is easy to see a number of areas in which dramatic improvements are desired and might be affected by a

judicious introduction of computer aids (Wass de Czege and Biever 2001).

Currently, manual products cannot be reused downstream in the process. Multiple re-entry of information, hand-jamming, and just the fact of creating multiple overlays take time. Even when the products are produced in an ostensibly “electronic” format, e.g., a PowerPoint presentation, it is not a true digitization – it lacks semantic content and cannot be readily reused in the downstream processes and tools. Could a better set of tools, capable of capturing the semantics of the digital information, address this deficiency?

Remaining essentially manual, the current process is time and manpower consuming (Bohman 1999, Paparone 2001). Much of this consumption of man-hours is directed toward computational tasks such as logistics consumption, time/space analysis, etc., which could be at least in theory allocated to a computer aid.

The time demands of the manual process force the staff into drastically limiting the number and diversity of options they are able to explore and analyze (Banner 1997). Perhaps, an intelligent computer aid could explore a greater range of options, enabling the staff to analyze more possible options in the same amount of time, or possibly conducting deeper analysis of the same number of options that the current process allows.

The dichotomy of planning and execution remains pervasive. The gulf separating the two is unacceptably wide, and could be explained at least in part by the fact that today’s planning process is far too slow to be merged effectively into execution decision-making. If computer aids make fast, real-time planning and re-planning possible, would it enable a major step toward the integration of planning and execution?

The Army’s corporate knowledge continuously evolves, and the rate of this evolution and adaptation has increased under the pressure of multiple factors: new military-political realities, the threat of asymmetric warfare, and the rise in operations other than conventional war, to name just a few. The effective mechanisms for capture and transmission of such knowledge are elusive. Could it be that computer decision aids (which by necessity must contain some of the warrior’s knowledge, continuously updated) can become one of such mechanisms?

Fighting by the Book and by Numbers?

In spite of potential benefits of decision aids in MDMP, their roles, limitations and concept of operations in military environments are justifiably open to a number of serious questions and concerns. These questions and concerns include:

Will they inhibit agility and dynamics of command, forcing greater reliance on slow and bureaucratic processes, command-by-plan and reduced latitude afforded to the tactical commanders?

Will such computer aids impose extensive training and specialization requirements, turning warriors into narrow-focused computer tool operators?

Will they encourage excessive fixation on analytical aspects of command, by the book and by numbers? And detract from intuitive, adaptive, art-like aspects of the military command decision making?

Will they engender undue dependence of future commanders and staff on technology that may be vulnerable in a combat environment? After all, isn’t it often said with a great justification that “a map with a hole in it is still a map, but a computer with a hole in it is a doorstop?”

Will it make the plans and actions more predictable to the enemy?

The Motivation and Focus of the Experimental Exploration

Some of these questions and arguments can be clarified, if not necessarily answered with opportunities and promise of experimental investigation. That was the rationale behind the Integrated Course of Action Critiquing and Elaboration System (ICCES) experiment, conducted by the Battle Command Battle Laboratory – Leavenworth (BCBL-L) as a result of a TRADOC sponsored Concept Experimentation Program (CEP) during the Government fiscal year 2000. In this experiment, several promising and representative prototype technologies were inexpensively “lashed together” to produce a necessarily crude but sufficiently useable suite of decision aids. The resulting ICCES system was then placed in the hands of several Army officers in controlled experiments. The key question was: can such tools provide value to Army decision-making?

For the purposes of the ICCES experiment we focused on the course of action planning and analysis steps of the MDMP: from documenting and communicating a high-level COA to producing a detailed analysis and plan of tasks. A highly creative step of inventing a high-level COA, currently explored by a number of researchers (Hayes and Schlabach 1998, Atkin et al. 1999, Tate et al. 2001, Kewley and Embrecht) was left outside the scope of this effort. To further circumscribe the scope of the experiment (subject as always to budgetary constraints) we focused on the planning process at the Army Brigade echelon.

The Experimental Rig

To provide computer-aided support to the selected aspects of MDMP, we identified several existing, advanced R&D prototype tools, modified them lightly and integrated them loosely and inexpensively into a conceptually seamless suite of decision aids (Fig. 1). The resulting “rig” offered a basis for conducting practical experiments structured around the key tasks of the staff process.

COA Creator, developed by the Qualitative Reasoning Group at Northwestern University, is a tool that allows a user to sketch a COA into the computer (Ferguson et al. 2000). Although superficially similar to familiar drawing tools like MS PowerPoint, COACreator is fundamentally different in that there are semantic knowledge based

representations stored into the computer for each item added to the COA sketch. Additionally, this tool uses an “overlay” approach to graphics which allows the user to switch graphics on and off easily in a fashion which is analogous to taking acetate graphics on and off a map, which is the current practice. Finally, the system is doctrinally based, to allow the military user to work in a domain environment that is familiar to him. The tool is also speech enabled, but for the purpose of the ICCES experiment, the users used drag-and-drop functionalities instead.

The COA statement tool, a product of Alphatech, Inc., was modified under the ICCES experiment to allow the staff planner to enter the COA statement. Unlike a word processor that captures words but not the semantics of the text entered, this tool presented the users with an interface that allowed them to produce natural language sentences to construct their COA statement. Additionally, this system was linked to the sketch tool to, in a sense, “auto-fill”

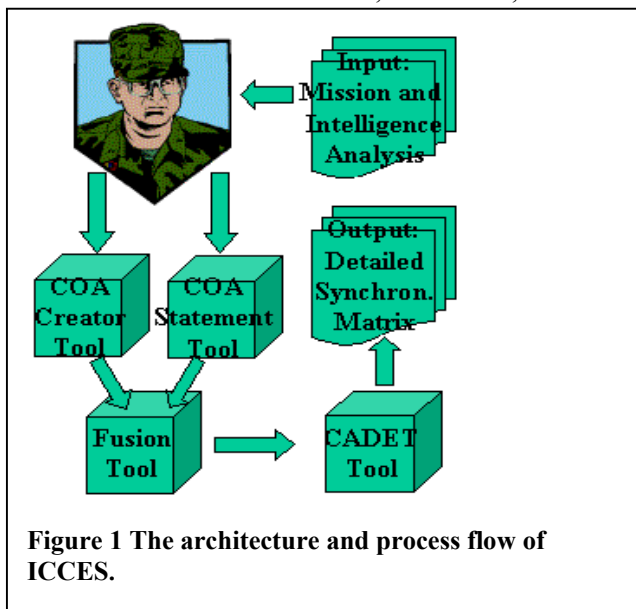


Figure 1 The architecture and process flow of ICCES.

portions of the COA statement that could be derived automatically from the sketch (e.g. units, tasks, etc.). Some examples of the sentences that can be constructed with the system are:

Close: TF 1-8, a mechanized infantry task force (Supporting Effort 1) attacks to destroys REDARCAVBN2 in order to prevent REDINBN17 and REDARCAVBN2 from engages in offensive operations. Fires: Fires will suppress OBJ CUB, then suppress OBJ ROYALS, then suppress OBJ BRAVES, then suppress OBJ BREWERS.

The sketch and statement of today’s staff process during COA development reflect different aspects of the course of action. Although they go hand-in-hand, they each contain some unique information that cannot be gleaned from the other (“purpose” for example, cannot be easily inferred from a COA sketch but is usually clearly defined in the statement). These two distinct aspects also reflected themselves in the fact that we had to use two different tools

– COACreator and the COA statement tool – that capture the content of the COA from two distinct prospective. Thus, we needed a mechanism that could merge the digital representations of sketch and statement in a unified product – that was the task of the Fusion Engine. It was developed by Teknowledge, Inc. and generated a single information file from the two separate sources as well as eliminating inconsistencies between the information. Additionally, Teknowledge built an XML translator to translate the knowledge fragments into the XML schema needed for the next system in the experiment.

Once the digital representation of the sketch and statement information is properly fused and translated, it goes to the next tool called CADET, developed by Carnegie Group, Inc. (now owned by BBN Technologies) as an SBIR program, under the sponsorship of CECOM RDEC. This tool transforms the sketch and statement into a detailed plan/schedule of the operation. CADET expands friendly tasks, determines the necessary supporting relations, allocates/schedules tasks to friendly assets, takes into account dependencies between tasks and availability of assets, predicts enemy actions and reactions, devises friendly counter-actions, estimates paths of movements, timing requirements, logistics consumption, attrition and risk. The resulting digital product can be then displayed in a number of different forms – as a traditional synchronization matrix or as an animated map. Although the resulting plan still requires careful review and editing by the planning cell officers, it was our expectation that it could serve as a good starting point for further analysis by the officers, and potentially a major time saver. A detailed discussion of CADET is found in (Kott et al. 2002, Ground, Kott and Budd 2002).

Once the COA is truly digitized, a tool like CADET can automatically (or with human guidance) perform the detailed planning, including the traditionally time-consuming tasks such as time-distance analysis, logistics calculations, and potential attrition calculations for the plan.

These tools were linked together mainly via file transfer, crudely but sufficient for exercising a carefully controlled experiment. The overall “rig” supported a logical concept of operation for the end-users, a group of staff planners:

- enter the COA sketch into the COA Creator, discuss and modify (e.g., Fig. 2);
- enter the COA statement into the Statement Tool, discuss and modify;
- review the detailed plan produced by CADET (e.g., Fig. 3), modify it as desired or return to the sketch and statement to produce a new or modified COA;
- use the detailed plan product to generate the OPLAN/OPORD.

Potentially, the entire process could be accomplished in a few minutes (minus the manual generation of the textual OPLAN/OPORD). However, only experiments could determine whether it would work at all.

The Experiment

The experiment was conducted over a 3-day period and involved eight Army officers (majors and lieutenant colonels) at BCBL-L facilities in Fort Leavenworth, Kansas. All the subjects were from combat arms branches and had a variety of tactical experience, with 11-23 years of Active Service. None of the users had prior technical backgrounds, but all possessed basic computer skills with MS Office products like PowerPoint, Word, etc.

The first day of the experiment consisted of training all the officers on how to use the system. The training consisted of walking the users through a complete scenario of COA development (sketch and statement) and COA expansion within the ICCES system. The training occurred over a 4-hour period and included a description of each system within the experiment, and then a sample COA was developed by the instructor with the students following along on their own machine. Given limited resources, the users worked in pairs during training, but were each given opportunities to manipulate the software. Observers noted the users performances, and at the end of the training, the users were broken into two roughly equivalent groups of four based on their tactical skills/experience and their demonstrated technical skills during the training.

On Day Two of the experiment, each group of four officers conducted the MDMP process given a tactical scenario. One group (control group) was to use the traditional, manual process. The other group was to use the ICCES system to conduct their planning. Each group received the same plan and briefing from their simulated higher headquarters, and both groups were allowed to ask questions in order to ensure their understanding of the plan (similar to how military units request additional information in order to ensure their understanding of orders from higher). Once the groups were confident in their understanding of the high-level plan and their requirements, they were allowed to organize and conduct their planning activities. Each group was informed that their deliverable products at the end of the day were three COA sketches and statements, and one COA synchronization matrix that reflected the one COA they had chosen internally as their “best” COA with a level of detail that would allow execution of the plan. The groups were not given a specific time limit to complete their planning, but observers monitored times for post-experiment analysis.

Day Three of the experiment would involve the same procedures as Day Two, but the roles of the groups would reverse. The control group assumed the role of the automated group and vice-versa. The scenario was slightly different, but similar enough to be comparable with the Day Two scenario with regards to complexity of the plan, etc. Both groups were tasked to provide the same products as generated in Day 2 for the new scenario in their new roles (automated or manual).

Although the experiment would provide valuable data and insights into the issues of focus, there are several

considerations that prevent us from claiming statistical relevance to our results. First, given the limited time availability of the user groups, we were unable to conduct enough iterations of the experiment to provide statistically valid results. Second, although the groups were broken out in order to attempt to achieve parity with regards to their tactical and technical abilities, human factors such as personalities could not be completely accounted for. Finally, by switching roles of the groups from Day Two to Day Three in the experiment, we introduced several other uncontrollable variables into the experiment, such as team

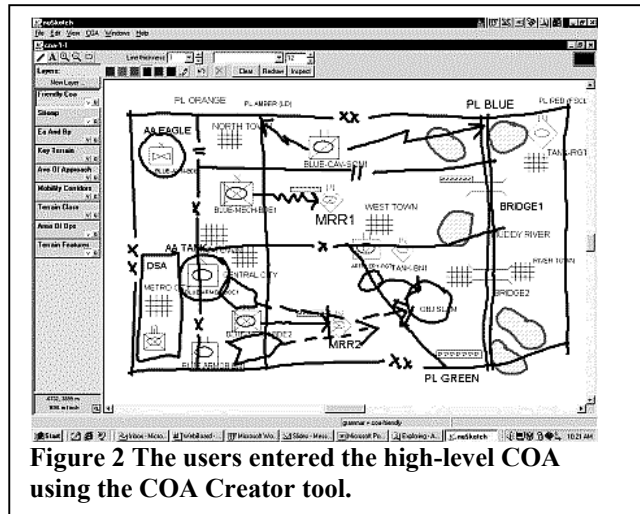


Figure 2 The users entered the high-level COA using the COA Creator tool.

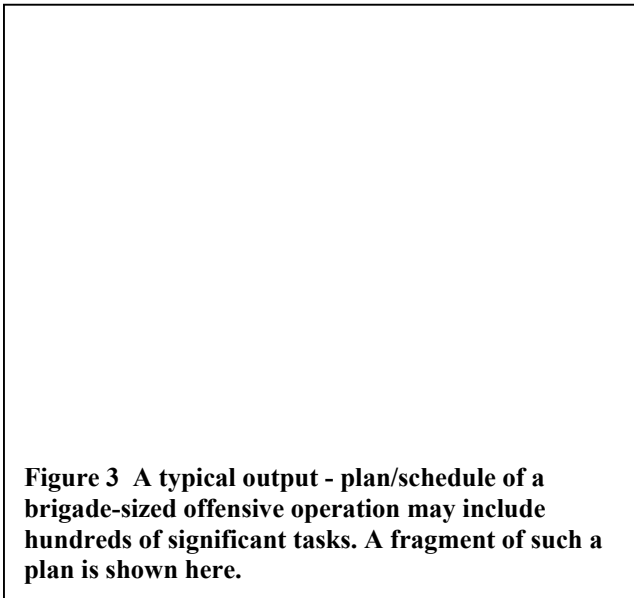
building within the groups and the ability of the initial control group to retain their training from Day One to Day Three with regards to manipulating the software.

Observations and Lessons Learned

Significant observations began in the training phase. In spite of a very modest time allocated to the training session, users did not exhibit any hesitation or difficulties in operating the system that could be attributed to a need for additional training. This was all the more notable in view of the fact that most of the training focused on the workarounds necessitated by the limited integration of the system. E.g., we had to train the users how to pass files between the components of the system, how to avoid crash-prone situation, etc. None of this should be necessary in a mature, fully-developed system. Even with this overhead, we were able to complete the training session in four hours. Without the overhead, we estimate that the training could be accomplished in less than an hour.

A key factor allowing the low training requirements and rapid, easy learning curve was the use of a sketch-based multimodal interface. The nuSketch approach to multimodal interfaces (Forbus et.al. 2001, 2002) used in the COA Creator, like other multimodal interface systems such as QuickSet (Cohen et.al 1997), exploits the naturalness of drawing and visual annotations for communicating with software. While QuickSet has shown itself to be very useful, the nuSketch approach had several advantages for this task over QuickSet. The QuickSet

such as QuickSet (Cohen et.al 1997), exploits the naturalness of drawing and visual annotations for communicating with software. While QuickSet has shown itself to be very useful, the nuSketch approach had several advantages for this task over QuickSet. The QuickSet approach focuses on providing recognition services as interfaces to legacy computer systems; its “smarts” are in statistical recognizers for visual symbols, speech and natural language understanding, and integrating these information sources into commands for the underlying software system. By contrast, nuSketch-based systems focus on rich conceptual understanding of the domain, spatial reasoning about the user’s ink, and clever interface design instead of recognition. These differences were important for this task in several ways. First, the conceptual understanding of the domain used in the COA creator provided the representational framework needed for CADET to do its work. Second, extensive pre-training of speech vocabularies and grammars was not needed, as it would be with QuickSet or any system using existing speech recognition technology¹. Instead, officers used the software equivalent of push-buttons (organized in layer-specific glyph bars) to indicate the intended meaning of their ink as they drew. This allows them to draw complex shapes (which cannot be handled by today’s statistical recognition technologies) and deal with interruptions such



as conversations with fellow officers (which would cause problems for most multimodal interfaces, which interpret

¹ While speech recognition is useful in many applications, today’s technology has severe limitations for battlefield use, including sensitivity to environmental noise and operator stress, user-specific training of the software, and training operators to work with limited vocabularies and grammars. Technology advances will change this over time, but it is worth being wary about near-future applications of speech recognition in battlefield systems.

pauses or lifting the pen as a signal that what the person is drawing is finished).

At the output side of the system, the users generally did not express dissatisfaction with the quality of the planning products generated with the ICCES system. The group that used ICCES elected to make only a small number of changes to the automatically generated product i.e., the highly detailed synchronization matrix (Fig. 3). Only about 10-15 % of the entries in the matrix were manually edited, indicating that the users were in agreement with the overwhelming majority of the plan produced by ICCES. After the editing, the products compared favorably with the products produced by the control group. For example, the COA’s produced by both groups when analyzed/wargamed either through the ICCES system or manually, all resulted in roughly the same estimates for time to complete the mission and overall attrition of friendly forces. This observation was later confirmed by a different experiment reported in (Kott, Ground and Budd 2002) with regard to the CADET module of ICCES, where a larger number of test cases and multiple unbiased judges were used to compare the products of manual and the computerized processes.

Further, there was no evidence that the computer-assisted process resulted in less imaginative, more cook-book type products. This can be simply explained by the fact that the overall COA inputted into the system came directly from the user and was not constrained in any way by the software. By allowing the user to free-hand a COA sketch into the system, there was complete freedom of tactical actions for the user.

On the other hand, there were discouraging observations with regard to the presentation aspect of the products. Synchronization matrix is an accepted way of recording the results of COA analysis. However, in the ICCES experiment we found that the users had difficulties comprehending the synchronization matrix generated by the computer tool, even though it was presented in a very conventional, presumably familiar manner. Perhaps, the synchronization matrix functions well only as a mechanism for short-hand recording of one’s own mental process. However, the same synchronization matrix is not nearly as useful when used to present the results of someone else’s, e.g., a computer tool’s, reasoning process. In effect, the synchronization matrix serves as a textual representation of a visual process. The problem was further exacerbated by the fact that ICCES-generated matrices were unusually detailed and therefore large, making it difficult for the users to navigate within this large volume of information. It appeared a system like ICCES requires a qualitative simulation/animation capability to visually present the expanded plan to the user.

Another factor contributing to the low training requirements appeared to be the intentionally simple, straightforward process flow and the concept of user-system interactions. These consisted of the sequence of steps outlined earlier in section "The Experimental Rig," and the users readily accepted them as natural and

consistent with their prior training and experience in the manual MDMP process.

In fact, the users displayed preferences for further simplification of the process. For example, the users stated that they would prefer a single process of entering sketch and statement, rather than the two sequential steps that they had to perform in ICCES. Their desire for a simple and straightforward concept of operations was further demonstrated in their use of the COA Statement tool – they consistently looked and asked for one, simple way to enter the statement, and shied away from the rich, flexible, but necessarily complex approach offered by the tool. We will return to this issue in the conclusions.

Consistent with the preference for a simple concept of operations were the users' requests for a mechanism that would allow them to perform easy modifications and iterations within the process. In particular, the users wanted to make changes in the synchronization matrix produced in CADET and see it automatically reflected in the COA sketch and statement. Although such a capability is technically feasible to develop, the ICCES system did not have it at the time of the experiment.

Of the greatest practical significance were those observations that confirmed a potential for major reduction in time and manpower required for performing a typical MDMP cycle. In particular, the COA analysis/wargaming process could potentially be shortened by several hours. Significantly, the savings were realized primarily in the downstream tasks, particularly in the step that generates a detailed plan/schedule of the operation. This is hardly surprising. The upstream processes of capturing the digitized information, such as was done in our experiment in the COA Creator tool, may not be any less time-consuming than a manual counterpart. However, once the information is captured in a digitized form, great time-savings accrue in the downstream processes. To put it differently, none of the ICCES components alone can deliver time-savings; but a system of such components can.

Conclusions

This study suggests evidence for several important conclusions. To start with, AI techniques can be used to create natural sketch-based interfaces that domain experts can use with little training. The nuSketch approach to multimodal interfaces, with its emphasis on visual understanding of the user's ink tied to a rich conceptual understanding of the domain, provides a practical method of expressing COA sketches with today's technology. One important limitation noted was the desire expressed for a single tool that captures COA sketches and statements simultaneously. This approach could be extended to provide a unified map-based interface to do both tasks, but it would require additional research to apply advances in natural language understanding and dialogue management to supply this capability. Another research opportunity is to use the rich representations in the COA Creator as inputs to other support tools, such as critiquers and pattern completion (for hypotheses about Red intent) and access to

previous plans via analogy (Forbus, 2001). Such capabilities could be incrementally added to near-future deployed systems as they became available, given a stable semantic framework.

With a semantically-rich input provided by a tool like the COA Creator, techniques of tightly interleaved adversarial planning and scheduling, such as those applied in the CADET component of ICCES, can be used to create thoroughly detailed plans comparable in quality to manually generated, but dramatically faster. Decision aids that combine both natural COA sketch interfaces and a full-functionality COA expansion mechanism can indeed contribute greatly into dramatic increase in speed and agility of the staff planning process, potentially bringing it into an integrated execution-planning cycle.

With regard to the concerns that decision aids of such nature might adversely impact the creative aspects of art of war, the experiment illustrated the fact that such a suite of computer aids is merely a tool. No tool is a substitute for training, doctrine and personal qualities of the decision makers, and regardless of tools, it is ultimately up to the decision-makers to define their approach and style of decision making. Although tools do lead to changes in the details and form of the process, the experiment offered no evidence that the substantive aspects of the decision making processes will be either inhibited or dictated by any such tools. Different commanders and staffs, with different styles, will use such tools to leverage their own preferences and strengths.

Currently, there is no evidence that a tool like ICCES would in any way increase predictability of the plans, or to encourage cook-book approach. To the contrary, because it allows the staff planners to explore rapidly a broader range of possible COAs, including those that are more unconventional and out-of-the-box, there is a potential for such tools to encourage greater ingenuity, creativity and adaptivity.

Overall, the experiment suggests that the ICCES concept is a practical paradigm for a planning staff decision aid, with near-future deployment potential. Staff officers would have such a tool available on a rugged, light-weight, highly portable device, such as a PDA, linked to other such devices via tactical internet. The decision-aid tool, in keeping with ICCES lessons, would be tightly integrated, capable of producing complex operational plans and orders rapidly and with minimal manual input, with simple, straightforward and natural operation concept, easy to learn and easy to use even in stressful field conditions. An officer would use it routinely to perform planning of tactical operations, to collaborate with others while on the move and dispersed over the battlefield, to issue operational plans and orders, and to monitor and modify the plans as the operation is executed and the situation evolves. Furthermore, as demonstrated in ICCES, the current AI technology is not far from being directly transitioned into a practical tool of such nature.

Acknowledgements

The US Army CEP program, sponsored by TRADOC, provided the funding for this work. LTC's Eugene Stockel, John Duquette, and Brooks Lyles from BCBL-L provided valuable advice and guidance. Mr. Murray Burke (DARPA) made invaluable contributions to the preparation and execution of the experiment. Several of the tools used in ICCES were developed with support by DARPA through the High Performance Knowledge Bases and Command Post of the Future programs.

References

- Atkin, M. S., Westbrook, D. L., and Cohen, P. R. 1999. Capture the Flag: Military simulation meets computer games. In Proceedings of AAAI Spring Symposium Series on AI and Computer Games, 1-5.
- Banner, G. T. 1997. Decision Making - A Better Way, Military Review, Headquarters, Department of the Army, vol.LXXVII, no.5.
- Bohman, W. E. 1999. STAFFSIM, An Interactive Simulation for Rapid, Real Time Course of Action Analysis by U.S. Army Brigade Staffs, Thesis, Naval Postgraduate School, Monterey, CA.
- Department of the Army 1997. *Field Manual 101-5. Staff Organization and Operations*, Headquarters, Washington, D.C.
- Cohen, P. R., Johnston, M., McGee, D., Oviatt, S., Pittman, J., Smith, I., Chen, L., and Clow, J. 1997. QuickSet: Multimodal interaction for distributed applications. In Proceedings of the Fifth Annual International Multimodal Conference (Multimedia '97), Seattle, WA, ACM Press, pp 31-40.
- Ferguson, R.W., Rasch, R.A., Turmel, W., and Forbus, K.D. 2000. Qualitative Spatial Interpretation of Course-of-Action Diagrams. In Proceedings of the 14th International Workshop on Qualitative Reasoning. Morelia, Mexico.
- Forbus, K., Ferguson, R. and Usher, J. 2001. Towards a computational model of sketching. In Proceedings of IUI'01, January 14-17, 2001, Santa Fe, New Mexico.
- Forbus, K. 2001. Exploring analogy in the large. In Gentner, D., Holyoak, K., and Kokinov, B. eds., *Analogy: Perspectives from Cognitive Science*. MIT Press.
- Forbus, K. and Usher, J. 2002. Sketching for knowledge capture: A progress report. In Proceedings of IUI'02, San Francisco, California.
- Ground, L., Kott, A., and Budd, R. 2002. A Knowledge-Based Tool for Planning of Military Operations: the Coalition Perspective. In Proceedings of the Second International Conference on Knowledge Systems for Coalition Operations, Toulouse, France.
- Hayes, C. C., and Schlabach, J. L. 1998. FOX-GA: A Planning Support Tool for assisting Military Planners in a Dynamic and Uncertain Environment. In Integrating Planning, Scheduling and Execution in Dynamic and Uncertain Environments, R. Bergmann and A. Kott, eds., AAAI Press, Madison, Wisconsin.
- Kewley, R. and Embrecht, M. 1998 Fuzzy-Genetic Decision Optimization for Positioning of Military Combat Units. In Proceedings SMC'98, pp. 3658 – 3664, 1998 IEEE International Conference on Systems, Man, and Cybernetics, October 11-14, La Jolla, California
- Kott, A., Ground, L., Budd, R., Rebbapragada, L., and Langston, J. 2002. Toward Practical Knowledge-Based Tools for Battle Planning and Scheduling. In Proceedings of IAAI 2002 Conference, Edmonton, Alberta, Canada.
- Paparone, C. R. 2001. US Army Decisionmaking: Past, Present and Future, Military Review, Headquarters, Department of the Army, vol.LXXXI, no.4.
- Tate, A., Levine, J., Jarvis, P., and Dalton, J. 2000. Using AI Planning Technology for Army Small Unit Operations. In Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems.
- Wass de Czege, H. and Biever, J. D. 2001. Six Compelling Ideas On the Road to a Future Army, Army Magazine, Vol.51, No.2, pp. 43-48..