Rapid Capturing of Battlefield Mental Models

Technical Report 95-3

Marvin S. Cohen, Bryan B. Thompson, Leonard Adelman,
Terry A. Bresnick, Martin A. Tolcott, and Jared T. Freeman

Cognitive Technologies, Inc.
4200 Lorcom Lane
Arlington, VA 22207

United States Army Research Institute
Fort Leavenworth Field Unit
Fort Leavenworth, Kansas 66027-0347

6 October 1995

ACKNOWLEDGMENTS

We are grateful to Dr. Joan Silver for her invaluable advice and support in this project. We are also grateful to Dr. Rex Michel for providing us with the Arisle scenario materials for use in our demonstration system. This work builds on earlier research for ARI on battlefield situation assessment, and we thank Dr. Jon Fallesen for his guidance and support in that effort.

RAPID CAPTURING OF BATTLEFIELD MENTAL MODELS

EXECUTIVE SUMMARY

Requirements:

Methods for rapidly representing battlefield situation knowledge would have utility in several important applications. Displays and decision aids might be more effective if designed to reflect the knowledge that users draw on in specific situations. Training might be more effective if designed to support a smooth trajectory of internal models leading up to expertise. Team coordination might be enhanced by methods for developing and communicating a shared situation understanding.

Yet there are significant obstacles in the way of real-time capturing of battlefield mental models. Such obstacles include improving the efficiency and overcoming the biases inherent in knowledge elicitation techniques, graphically representing knowledge that is not easily verbalized, and incorporating uncertainty into concrete visualizable representations.

Procedure:

The present report pursues theoretical, empirical, and practical issues in the design of a real-time mental model capturing system. A theoretical framework for understanding different types of mental models is developed, based on recent research in cognitive psychology. Special attention is paid to the problem of representing uncertainty in the context of such models.

We analyze critical incident interviews and think-aloud problem-solving sessions with active-duty officers to identify a set of mental models utilized by command staff in Army tactical battlefield operations and planning. Based on further analysis of the same interview and problem-solving data, we look for features of the environment, the context, or the individual officer that predict the type of mental model that is used.

Finally, we draw on these results to develop an initial proof-of-concept mental-model-capturing system. The system is implemented in NeXSTEP on a Pentium desktop computer, and utilizes a combination of off-the-shelf and customized software components.

Findings:

We describe a framework for situation understanding that includes several qualitatively different types of mental models (pattern-matching, interpretative, and generative), and a set of metacognitive processes that monitor mental models for problems of uncertainty and adopt corrective strategies when problems are found.

Based on coding and analysis of interview and problem-solving data, a set of five key mental model structures was identified that officers consistently use to organize their understanding of battlefield situations. These structures pertain to *intent* (enemy or friendly), *principles and methods* (e.g., for attacking or defending), *action execution* (temporal and causal relations among actions), *rate of movement*, and *evidence interpretation*. We also identified a set of metacognitive processes that influence the construction, elaboration, and modification of

mental models: generating alternative possible causes and effects, noticing surprising events, and revising the model to explain unexpected events.

We tested the ability to predict when different mental model structures and metacognitive processes will be used. Significant predictive relationships were found involving three kinds of variables: environmental (e.g., mission, type of unit), personal (e.g., amount of experience), and contextual (e.g., other structures that have been used).

We designed and implemented a proof-of-concept mental model capturing system. The system provides users with flexible tools for creating graphical structures of the kind identified as important in battlefield situation understanding. Users can annotate mental models to indicate gaps, unreliability (e.g., alternative causes and effects), or conflicting implications in such structures. Information of any kind (e.g., from intelligence estimates, orders, the commander's estimate, message traffic, spot reports, maps, etc.) can be inserted into the model or linked to its components.

The system dynamically adjusts its advice based on the environment, user, and immediate context (drawing on the empirical findings alluded to above). The system examines the models under construction and finds and recommends information that is relevant to the model (to fill gaps, and to confirm or disconfirm the model). The system also recommends ways to elaborate and complete the on-going model. Finally, the system will recommend structures and components of structures that may help interpret a particular piece of information.

Utilization of Findings:

A mental model capturing system will be of value as a computer-based aid for situation assessment, planning, and operations; as a team aid, to support dissemination, integration, and collaborative development of shared mental models; as a research tool, to investigate the knowledge structures and processing strategies utilized under different circumstances; and in training, both to diagnose levels of skill and to transfer appropriate knowledge structures and processing skills. These results should therefore be of interest to instructors, designers of C^2 systems, C^2 analysts, and researchers. The present results provide a foundation for a mental model capturing system through empirical analysis of mental models, metacognitive strategies, and the variables that predict their use, and by the development of a proof-of-concept system.

RAPID CAPTURING OF BATTLEFIELD MENTAL MODELS

CONTENTS

TABLES

FIGURES

INTRODUCTION

<u>The Problem</u>

Knowledge has moved to center stage in recent theoretical and empirical studies of expertise. In domains ranging from chess to physics to computer programming and medicine, more proficient problem solving appears distinguished not by general-purpose, analytical strategies, but by the accumulation of effectively structured knowledge and strategies for exploiting it (e.g., Chase and Simon, 1973; Larkin, 1983; Chi, Feltovich, & Glaser, 1981).

Recent developments in Army doctrine have heightened the importance of knowledge and knowledge structure in battlefield decision making (Department of the Army, 1993). One example is the concept of *battlespace*. A commander's battlespace is a three-dimensional moving volume that contains anything relevant to the commander's planning or operations. Unlike the traditional *area of interest*, battlespace is not handed down by higher command, but is defined by each commander. It reflects the commander's ability to visualize relevant events at an appropriate level of detail, far enough into the future, and in a large enough volume of space. In short, the commander must use knowledge not only to interpret the situation, but to define what counts as a relevant part of the situation. A second important development in recent doctrine is the decreasing emphasis on enumerating and comparing alternative courses of action. In circumstances of time stress, commanders might use "abbreviated" methods, in which only a single course of action is proposed and assessed. The effectiveness of that course of action will clearly depend on the validity of the commanders' understanding of the situation, and that in turn will depend on the knowledge the commander applies.

The problem is, how can we know — quickly and accurately — what someone else knows. The practical implications of research on rapidly diagnosing knowledge are potentially enormous. *Displays and decision aids* might be more effective if designed to reflect the knowledge structures that users bring to bear, or attempt to construct, in specific situations (e.g., Cohen, 1987, 1993c; Noble, 1993). Such aids might prompt users regarding problems, such as incompleteness, conflict, or unreliability, within their situation models and plans. *Training* might be more effective if it is based on accurate diagnosis of a trainees' current ability to construct situation models, and if training is designed to support a smooth trajectory of internal models leading up to expertise (e.g., Koedinger & Anderson, 1993; Cohen, Freeman, Wolf, Millitello, & Klein, 1994). Instruction, practice, and feedback would focus on the ability to identify the most useful types of model for a task and to use those models effectively. *Team coordination* might be enhanced by a better understanding of the mental models required to operate in a team or distributed team setting, including models of the relevant parts of the overall task and plan, models of the relevant parts of the overall organizational structure, and models of the resources, personnel, information sources, capabilities, attitudes and habits of other team members or other teams (e.g., Cannon-Bowers, Salas, & Converse, 1991). Shared situation understanding would also be enhanced by methods for critiquing and iterative modification of mental models.

Unfortunately, the introduction of knowledge modeling into these applications has been relatively slow. Knowledge may not only be the most important element in expertise; it may also be the most difficult to analyze and assess. There are several different kinds of difficulties to be overcome:

- Elicitation techniques vary from those that are both close to the natural task and non-directive (e.g., observation of actual performance) to those that are both highly artificial and highly directive (e.g., structured interviews that attempt to directly elicit general rules). Between these extremes lie techniques that are somewhat unnatural and somewhat directive (e.g., think-aloud solving of simulated problems, critical incident interviews). Other methods are more natural in some respects (e.g., requiring less verbalization) but at the cost of failing to resemble the original task at all (e.g., concept mapping, problem sorting, multi-dimensional scaling). In general, the more efficient a technique is at extracting specifically targeted knowledge, the more artificial and/or directive it is, and thus the more likely it is either to misrepresent or to influence the processes it purports to capture (Ericsson & Simon, 1993; Nisbett & Wilson, 1977). The cost for avoiding bias is time.
- Knowledge that is accumulated during direct experience in a domain may not be easily verbalized. Anderson (1982) argues that explicit knowledge is converted through experience into implicit knowledge that is not accessible to consciousness. Moreover, knowledge may be represented in a variety of different internal and external formats (e.g., spatial imagery, verbal or propositional, numerical, procedural).
- Different individuals may differ significantly in the knowledge they posses, in the way their knowledge is represented, and in the strategies they apply in using it.
- Standard normative methods for representing and handling uncertainty (e.g., by Bayes' rule) do not seem to fit the way real-world decision makers handle uncertainty (Kahneman, Slovic, & Tversky, 1982; Cohen, 1993b). Yet there is no well-developed alternative approach to representing uncertainty in mental models.

We have explored an approach that responds to each of these challenges.

A Cognitive Approach

Can methods that are both fast and non-obtrusive be developed for representing situation knowledge in naturalistic battlefield environments? We think the tradeoff between speed and bias in knowledge elicitation can be mitigated by separating the knowledge capturing process into three stages:

1. The major work of eliciting the model structures that are characteristic of the domain and strategies for using them is accomplished beforehand. Stage 1 can be done slowly and accurately by methods such as critical incident interviews and think-aloud problem solving.
2. An interactive computer-based system permits its users to create appropriate situation models and to quickly and flexibly map the data flow of intelligence, reports, and orders into them. Stage 2 must be rapid enough to capture the current state of understanding during an evolving tactical situation, and to support the kind of real-time knowledge assessment required for applications in team coordination, aiding, and training. Since this process is facilitated by the generic structures and strategies captured in the earlier, more accurate stage 1, biasing effects are minimized.

3. A component of the interactive computer-based system monitors the user's activity, and revises the body of generic structures and strategies in response to the performance of an individual user. Stage 3 thus provides a further hedge against the biasing effects of stage 2, by means of a more slowly evolving refinement and individualization of the database of structures and strategies.

People are not particularly accurate when asked to articulate how they perform a task or why they perform it in one way rather than another (Nisbett & Wilson, 1977). Fortunately, none of the three stages of knowledge capturing requires that they do so. Neither the critical incident method nor think-aloud problem solving ask for general rules. Participants need only say or draw what is (or was) on their mind in an actual incident; in problem-solving, they may directly demonstrate a procedure. The interactive capturing process in stage 2 can organize maps, photographs, sketches, or other representations as easily as words. The mental model structure itself is represented graphically. Finally, stage 3 is based on observation of performance in stage 2, and does not require direct verbalization.

Individuals will surely differ in how they represent situations, solve problems, and make decisions. Real-time capturing in stage 2 will be faster to the degree that the structures and strategies elicited beforehand fit the needs of the current user. But if they do not, the user is free to create any structure the user wishes. The price for using this flexibility is, of course, time, but that price is only temporary. The stage 3 tracking process will eventually produce a more customized set of representations and strategies for use in stage 2, and stage 2 will tend to become both faster and more accurate over time. Yet another solution is to carry out the stage 1 elicitation process on each individual user, so that the system is tailored to their particular representations and strategies from the start. This requires that additional time be invested up front.

Another apparent dilemma in using mental models involves the representation of uncertainty. Most approaches have been abstract and mathematical, such as Bayesian statistics, Shafer-Dempster belief functions, and fuzzy logic. These systems do not provide a concrete, visualizable representation of the current real-world situation. Actual decision makers, moreover, do not appear to generate, quantify, and aggregate all possible hypotheses as these models require (e.g., Kahneman, Slovic, & Tversky, 1982; Cohen, 1993b). There seems to be a tradeoff between representations that incorporate uncertainty at the expense of visualizability, and visualizable representations that ignore uncertainty.

We believe, on the contrary, that processes for detecting and addressing uncertainty provide one of the most important ways that visualizable mental models evolve. A variety of cognitive strategies are utilized to monitor for problems of incompleteness, unreliability, and conflict in the current situation model; other strategies correct problems that are found by collecting and retrieving new data, modifying the model, or abandoning it in favor of another. A key feature of the model structures that we develop, therefore, is their incorporation of uncertainty into visualizable situation representations. They do so by overlaying the primary situation representation with a set of qualitative annotations, representing alternative causes and effects, arguments, and rebuttals.

In sum, three steps lead to a system for rapid capturing of mental models: First, to develop a background database of knowledge structures and structure components along with rules for their combination, modification, and use; second, to develop computer-based interactive techniques for matching and assembling elements from the structural database into real-time representations that track a decision maker's momentary mental model; and third, to provide for the refinement and individualization of the background database as a function of the model-building choices of the user. The interactive capturing of mental models builds on a repertoire of structures that becomes increasingly tailored to specific situations, but is always capable of modification by the user. The real-time representations will graphically organize information that is originally represented in a variety of different formats. Moreover, by explicitly representing uncertainty, such models will support the processes by which mental models are dynamically verified and revised in evolving situations. This approach is both theoretically and empirically grounded: It draws on cognitive research in knowledge representation and executive strategies and at the same time relies on extensive interviews with experienced active-duty officers.

Phase I research provides a feasibility demonstration of key components of this concept. This report describes the results in four logically related steps:

1. *Understanding Mental Models*. We describe a theoretical framework for understanding mental models and the cognitive strategies that help construct and verify them, under dynamic and uncertain conditions.

2. *Eliciting Mental Models*. We characterize a set of mental models utilized by command staff in Army tactical battlefield operations and planning, based on a series of interviews with active duty officers. The goal is to define a relatively small set of canonical structures that can be combined and elaborated in a mental model capturing system, to generate a rich variety of real-world situation representations.

3. *Predicting Mental Models*. Based on the interview data, we look for features of the context, the mission, or the individual officer that predict the type of mental model that is used. These findings permit a mental model capturing system to facilitate the construction of models that are most likely to be used in particular circumstances.

4. *Capturing Mental Models*. We describe an initial proof-of-concept mental-model-capturing system. The system permits users to construct and update situation-specific mental models from a database of structures and structural components, and from a digitized flow of information including such items as the intelligence estimate, commander's estimate, messages, spot reports, orders, and so on. The system illustrates the potential for predictive support of mental model capturing. It suggests candidate structures based on a perusal of the available information, suggests candidate structural completions based on the structures thus far created by the user, and guides the user to relevant information based on the structures that the user is building.

UNDERSTANDING MENTAL MODELS

It will be useful to refer to situation-specific structures as *mental models*, and to distinguish them from more permanent underlying knowledge that is used to generate them. But what is the nature of a theoretically well-founded and empirically justified notion of mental model? Are different formats of knowledge representation best suited for different types of knowledge? Are different types of mental models used to interpret battlefield situations?

There are a number of competing theories of mental models and, more generally, of how situation knowledge is represented (e.g., Rasmussen, 1979; Young, 1983; Gentner & Stevens, 1983; Rouse and Morris, 1986; Rogers, Rutherford, & Bibby, 1993). We can begin to clarify some of the differences among these theories, and move toward a synthesis, by considering three levels at which experts might respond to a problem, as a function of its degree of novelty in relation to their existing knowledge:

1. Intuitive mental models (pattern-matching): The decision maker directly recognizes the situation as familiar or typical. The situation has activated a schema comprising either a past case that resembles the present situation or an aggregation of many such cases. The decision maker performs a response associated with the schema. This level of decision making or problem-solving is reflected in Schank & Abelson's (1977) notion of a *script*, as well as in Noble's (1993) concept of *schema*.

2. Interpretative mental models: The situation is not fully captured by any single pre-existing schema, but it partially matches several. The decision maker *constructs* a cognitive model of the environment by combining pre-existing representations corresponding to pieces of the situation. The decision maker mentally tests a candidate action within the constructed model (or draws some other desired conclusion from the model). This level of decision making is addressed by Johnson-Laird's (1983) notion of a mental model, by Schank's (1982) *Memory Organization Packets (MOPs),* and by Pennington & Hastie's (1993) causal stories.

3. Generative mental models: The decision maker cannot construct a cognitive model as in 2 because there are gaps in experience of relevant event sequences. The decision maker generates expected sequences of events by drawing on deeper qualitative or quantitative knowledge of the relevant objects and forces and their causal relationships. The consequences of an action might be mentally simulated by considering the mechanisms of its interaction with objects or agents in the environment. This level of decision making suggests the notion of a device model (e.g., deKleer and Brown, 1981).

In the remainder of this section, we explore these three kinds of models and their implications for cognitive strategies in battlefield decision making. The main focus will be on interpretative models and on the strategies that verify and improve them. These strategies make mental models dynamic. They detect problems, in particular different kinds of uncertainty, and expand, modify, or elaborate the models to correct them. In doing so, they draw on both pattern recognition and generative models. This discussion provides the theoretical basis for rapid capturing of evolving situation models.

Pattern-Matching Models

In artificial intelligence, general-purpose methods (such as means-ends analysis) have proven insufficient for the solution of real-world problems, especially where the space of possible solutions to be searched is very large and time is limited. The second generation of artificial intelligence systems, expert systems, draw their power instead from the incorporation of large quantities of domain-specific knowledge, enabling them to drastically reduce or eliminate search (Feigenbaum, 1977). This knowledge is often embedded in production rules, whose antecedents represent domain-specific patterns and whose consequents represent associated actions. Pattern-based knowledge may also be represented in frames or schemas, corresponding to familiar types of objects, situations, or activities.

Similar developments occurred in the psychological modeling of human problem solving. DeGroot (1965) and Chase & Simon (1973) concluded that chessmasters are able to recognize and identify a large number of familiar structural patterns. Anderson (1982) presented a theory of the development of cognitive skill as a replacement of explicit data structures by implicit procedural methods for representing knowledge, through repetition of a task. Larkin, McDermott, Simon, & Simon (1980) found support for this notion in a study comparing expert and novice performance on physics problems, Novices applied general problem-solving methods (e.g., working backward from the unknown to the given quantities) to explicit representations of the domain knowledge (e.g., relevant physical equations). Experts, by contrast, were faster, worked forward from the given to the desired quantities, and usually verbalized only numerical results rather than the equations themselves. Berry and Broadbent (1984) and Koedinger & Anderson (1993) argued that procedural or pattern-based knowledge may arise directly from experience in a domain, rather than by compiling pre-existing declarative knowledge.

Klein (1989, 1993) and his associates have proposed a model of decision making based (in part) on recognitional processes that closely parallels the above work in problem solving. Retrospective accounts by urban fireground commanders (FGCs) revealed no evidence for traditional analytical methods of generating and comparing options. Instead, the FGCs relied on their ability to recognize and appropriately classify a situation. Once they knew that it was “that” type of case, they usually also knew the typical way of reacting to it. In the simplest case, called *rapid Recognition-Primed Decision Making (RPD)*, the situation is recognized as familiar and a typical action is implemented (Klein, 1989; Klein, Calderwood, and MacGregor, 1989; Klein, Calderwood, and Clinton-Cirocco, 1986).

Pattern-based models are sometimes represented as relatively unstructured lists of cues and associated responses. According to Klein, for example, familiar cues activate patterns containing expectancies, goals, and typical responses. Schank and Abelson's (1977) notion of a script is a collection of scenes (or activities) in pursuit of a general set of goals.

Interpretative Mental Models

There is evidence that a simple pattern-matching approach to decisions is inadequate:

1. A rapid recognition-based approach does not explain how decision makers handle uncertainty, novelty, or unexpected situations. In recent research (Serfaty, 1993), experienced Army planners were compared to novice planners. The experienced planners did not perceive more similarities with prior situations, did not generate plans more rapidly, tended to see the situation as more complex, were less confident in their solutions, and felt the need for more time. Similarly, Shanteau (1992) provide evidence that experienced decision makers are more likely to acknowledge uncertainty than novices. Clearly, something more than simple pattern matching is responsible for this behavior.

2. Situation models are often constructed over time, not simply retrieved whole. There may be no pre-existing pattern that fits the current situation. Research by Schank (1982), Pennington and Hastie (1988), and others suggests that understanding a situation is not instantaneous, but may require the activation and combination of multiple structures in many steps. Schank (1982) abandoned the original notion of a script as a fixed, self-contained sequence of events, and adopted the more abstract and flexible notion of a MOP (memory organization packet), which combines with other MOPs to represent a particular situation. The construction of situation models is controlled by reference to goals rather than being completely automatic. In

reading, people monitor for comprehension by asking questions and attempting to predict subsequent text; the level of effort, and the amount of comprehension they consider adequate, is determined by the purpose of reading (Collins, Brown, and Larkin, 1980). Goals even influence the level of specificity at which familiar objects are verbally classified (Cruse, 1977).

3. Plans and actions are often *designed* on the basis of an evolving situation representation, not *retrieved* after situation assessment is finished. Situation assessment and action selection are often intertwined rather than separate steps. Decision makers evaluate their current understanding of the situation to discover constraints that it imposes on action (Voss, Wolf, Lawrence, & Engle, 1991). If actions are sufficiently constrained to accomplish the task, the situation model is deemed adequate; if not, the situation model must be further elaborated until a satisfactory plan is developed (Cohen, Adelman, Tolcott, Bresnick, & Marvin, 1993).

A third generation of theories may now be emerging in cognitive science, advancing beyond both sterile general-purpose models and overly narrow pattern recognition approaches (Holyoak, 1991). Like the second-generation theories, these theories recognize the importance of the knowledge base experts use to solve problems in their domains. However, there is a much greater attention to the processes and strategies experts employ to effectively construct and apply that knowledge. Thus, these theories account for both routine and adaptive expertise, i.e., the ability to recognize and respond swiftly to familiar situations, as well as to handle novel and uncertain situations well.

The *Recognition / Metacognition (R/M)* theory (Cohen, 1993a) is a theory of this kind that emphasizes: (i) the active construction of situation models and plans over time by a process of activating and combining existing knowledge; and (ii) metacognitive skills for monitoring and regulating that process. Both of these components enhance the ability of decision makers to deal with the novel and uncertain situations. In the R/M model, knowledge is specific to a domain, but varies in its level of abstraction. In a given situation, an individual may have access to short term memory of the current situation, long-term memory of similar situations, and long-term memory of abstract classes of such situations.

Abstract organization of knowledge is associated with effective problem solving. In physics, Chi et al. (1981) found that experts and novices differ in the way they sort problems by similarity. Novices categorize problems by "surface structure," i.e., superficial features such as type of apparatus, while experts rely on basic principles of physics (e.g., conservation of energy) and generic solution techniques associated with such principles. Similar differences between experts and novices in algebra are reported by Shoenfeld and Herrmann (1982), and in computer programming by Weiser and Shertz (1983). In concurrent think-aloud protocols for physics (Larkin, 1981) and geometry (Greeno, 1983), the entities mentioned in expert as compared with novice protocols tend to be technical rather than familiar, and to be closely tied to fundamental laws.

According to Pennington and Hastie (1992) comprehension of trial evidence by jurors is a constructive process, in which the jurors create explanatory causal models of the available facts in the form of stories, and thereby give meaning to the data. Stories also enable jurors to identify gaps where important pieces of the explanatory structure are missing and where inferences might be necessary. The R/M model posits a similar process, in which abstract structural knowledge is used to organize information, and then the result is subjected to repeated evaluation and modification. As decision makers become familiar with a domain, they acquire knowledge about the types of events and relationships among events that are relevant in particular situations. In new situations of the same kind, decision makers use this generic knowledge to integrate the new information. In particular, structural knowledge consisting of causal and intentional relations between events is used to construct narrative story structures. The main components of a story episode, according to Pennington and Hastie (1992), are *initiating events* (which elicit) *goals* (which motivate) *actions* (which result in) *consequences*. Pennington and Hastie suggest that story construction is a general comprehension strategy for understanding human action. We have identified numerous examples of stories of this kind in interviews with command and G-3 staff.

A story episode is an example of an *interpretative mental model*. Interpretative models represent events, objects, or properties within a structure that describes relationships among entities of the relevant types. The relations may be causal, temporal, is-a-part-of, or is-a-kind-of. Such structures provide an abstract framework within which vastly different contents can be organized, understood, and evaluated.

In the R/M model, metacognition monitors and regulates the process of using abstract knowledge to construct situation interpretations. Two aspects of this metacognitive skill emerge clearly from expert/novice studies. First, experts use metacognitive skills to test the quality of their problem solutions. Larkin, et al. (1980) noted that physics experts utilize the abstract physical representation of a problem to verify the correctness of their method and result, e.g., by checking whether all forces are balanced, whether all entities in the diagram are related to givens in the problem, etc. Patel & Groen (1991) made similar observations of expert physicians' verifying their diagnoses.

Second, experts use metacognitive skill is to facilitate the discovery of problem solutions and to improve problem solutions when they are found to be inadequate. For example, Larkin, et al. (1980) found that physics experts often construct and examine a sketch of the superficial objects and relations in a physics problem in order to determine the next step: if the depicted system is familiar, the expert may proceed directly to the equations required for solution. If the system is still unfamiliar, the expert constructs an idealized representation (i.e., a free-body diagram), which is then used in the generation of solution equations. Chi, Glaser, and Rees (1982) found that experts returned to, and refined, the initial gross representation throughout the course of the problem.

In the R/M model, the basic level of cognition is *recognitional*, including processes that activate assessments in response to internal and external cues. For example, the presence of enemy objectives in a sector may be recognized as a cue regarding an intent to attack in that sector. A correlation of forces that is favorable to the enemy in that sector will be similarly recognized and reinforce the assessment of intent to attack. Assessments may in turn activate abstract structures, i.e., interpretative models, that organize actual and potential information into a situation model or plan.

Meta-recognition is a cluster of skills that go beyond the recognitional processes in situation assessment. They are analogous in many ways to the meta-comprehension skills that proficient readers use when they construct a mental model based on the information in a text. For example, according to Baker (1985), skilled readers test and evaluate the current state of their ongoing comprehension, and they adopt a variety of strategies for correcting problems that are found, such as inconsistency or gaps in their understanding.

The R/M model includes three meta-recognitional processes: the quick test, critiquing, and correcting. Taken together, these processes help determine what kinds of mental models decision makers will construct on a given occasion. The quick test is a gate-keeping function that determines (1) whether to engage in critiquing and correcting processes that might improve situation understanding or (2) whether the current level of understanding can (or must) suffice. The quick test considers the available time, the costs of an error, and the degree of uncertainty or novelty in the situation. This can be a relatively explicit and conscious process, or a form of recognitional processing at a higher level that is rapid and virtually automatic. The quick test may reveal that understanding is problematic, but not pinpoint the source of difficulty. In this case, the next step is to ferret out problems.

A key aspect of critiquing is identification of evidence-conclusion relationships, or *arguments*, within the evolving situation model or plan. This is simply an implicit or explicit awareness that fact A served as *evidence* on this occasion (e.g., the presence of an enemy objective in the sector) while intent to attack at a particular place and time along with expectations of observing other events B, C, and D served as *conclusions* (e.g., the predictions that the enemy will move up artillery, remove obstacles, and concentrate forces in the sector). On some other occasion B, C, or D might be serve as grounds and A might be inferred or predicted. Critiquing strategies may be used to identify three kinds of problems with the arguments underlying a situation model or plan: incompleteness, conflict, or unreliable assumptions. A model or plan is incomplete if expected information is missing, i.e., some information has not been obtained that would ordinarily provide an argument for or against an assessment of interest (such as location of the attack). Conflict exists when there are arguments that support contradictory assessments. An argument is unreliable if it depends on implausible assumptions.

When a specific problem with the model or plan is identified through critiquing, the third major function of metacognition is enlisted: facilitating the construction of an improved model or plan. This *correction* stage may involve gathering more information or spurring recall; it may also involve adopting or revising assumptions to create a more plausible situation interpretation. Each cycle of corrective activity restructures the problem so that critiquing processes can again be applied. The process stops if and when the newly represented model is satisfactory— or the quick test concludes that time has run out or the importance of the problem (relative to other priorities) has been reduced. Figure 1 summarizes these processes.

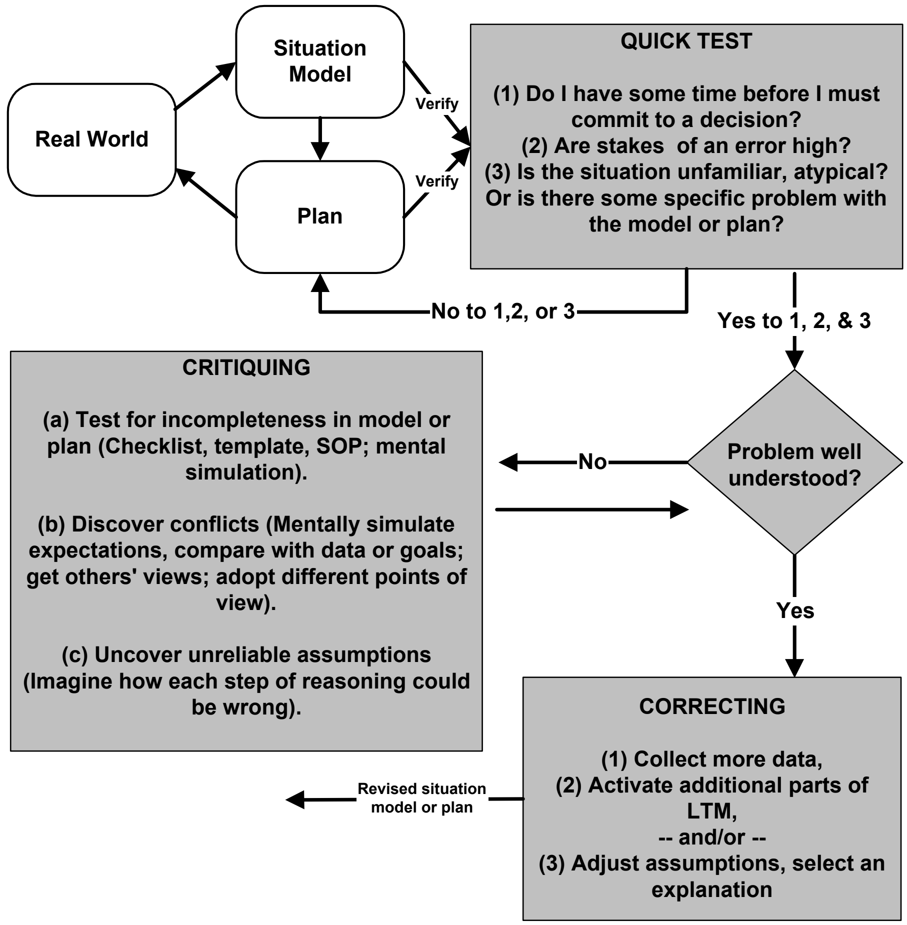

Figure 1. Component processes in the Recognition / Metacognition model. Unshaded boxes represent recognitional processes. Shaded boxes represent meta-recognitional processes.

Interpretative models lend themselves to meta-recognitional critiquing and correcting because of the way they explicitly represent structure, e.g., causal relations among events. For example, in critiquing the incompleteness of a model, an officer looks for information about potential *causes* and *effects* of the phenomenon of interest. In critiquing the reliability of the hypothesis of enemy intent to attack in a particular place at a particular time, an officer might look for *alternative causes* of observed actions, and *alternative effects* of enemy goals.

The explicit representation of causal relations, in a two-dimensional spatial format, permits a more perspicuous, *analog* representation of the situation. There is considerable discussion and debate in the research community regarding the nature of (and the need for) a distinction between “analog” and “propositional” representations (e.g., Pylyshin, 1979; Kosslyn, 1994, 1980; Rumelhart and Norman, 1985; Shepard, 1975; Metzler and Shepard, 1974). Analog representations, we believe, play a key role in the process of rapidly combining multiple structures into a single representation to support reasoning about a situation that has not previously been explicitly represented.

According to Johnson-Laird(1983), what distinguishes a mental model from other forms of knowledge representation is the close structural isomorphism between the model and the state of affairs it represents. This isomorphism involves two properties: There is a one-to-one correspondence between components of the representation and components of the situation which it represents; and the representing relations have the same inherent constraints as the represented relations (cf., Rumelhart and Norman, 1985).

When these properties are satisfied, numerous related facts can be combined with one another in the same representation, and the implications of the combined facts can then simply be *recognized* or *read off* the representation without the need for logical deduction or for the separate explicit statement of each implied fact that is characteristic of logical or probabilistic models. The simplest example is a map, which uses symbols and spatial relations to represent objects and their spatial relations. If an officer learns from one observer that howitzer A is west of tank B, and then learns from another observer that tank B is west of howitzer C, the officer knows that howitzer A is west of howitzer C; the relation “to the left of” among the symbols for A, B, and C on the map preserves the same transitivity as “to the west of.” (To perform this inference deductively in a propositional system would require general rules stating the transitivity of “to the west of” and all other such relations.) The relations represented by mental models need not be spatial, but may, for example, be temporal, causal, kinematic or dynamic (i.e., a continuously changing representation), or conceptual; the relations used in the representation may be, but need not be, the same type as the relations being represented, although they must be isomorphic. Causally structured mental models can also yield new inferences that are “read off” the representation, without the need for laborious deductive reasoning.

Mental model representations, because they mirror what they represent, more directly trigger the schemas that embody the user's underlying knowledge of a domain. According to Johnson-Laird (1983), they enable decision makers to “experience events by proxy.” Johnson-Laird speculates that people use schemas to construct mental models, which they manipulate to understand and anticipate phenomena, activate new schemas, and to select and control the execution of actions. Mental model representations of this sort thus maximize decision maker understanding and effectiveness.

As a result of its representational directness, there is no simple way to represent uncertainty in analog mental models. Let us consider maps again and suppose, for example, that we know that howitzer C is west of howitzer A and tank B is west of howitzer A. How can this be represented in an analogical model. Since the information does not specify the relationship between B and C, there are two possibilities:

A B C

or

A C B

We cannot have a model with a direct mapping to the state of affairs it represents when we do not know what that state of affairs is.

The strict requirement of isomorphism of analog models can be relaxed in various ways to represent uncertainty. Each approach has both advantages and drawbacks:

- Utilization of more imprecise models — e.g.,

A D

Here a coarser mapping onto the real state of affairs is preserved by lumping B and C together as a single aggregated object D. This may be a viable approach, unless decision making requires that the relative locations of B and C be known. There is a parallel solution in the case of causal interpretative models. For example, in a story about enemy attack where it is not known whether the enemy will attack in the north or the south (as in Figure 11), we might indicate that the intent is to attack somewhere in a broader region encompassing both northern and the southern sectors. If this is insufficient to support planning (e.g., due to limitations on friendly defensive assets), then meta-recognitional critiquing will find this model *incomplete*.

- Use of multiple models — e.g.,

  A B C
  A C B

  While each model taken separately is isomorphic to a possible state of affairs, the disjunction of the two, which represents the uncertainty, is not. An additional problem is the potential combinatorial explosion as new information, and new indeterminacies, are added. Generation and evaluation of all possible mental models quickly exceeds human (and computer) capabilities. One or two important alternative models may, however, be considered. In the R/M framework, decision makers construct and evaluate an alternative model when the current model depends on questionable, or *unreliable,* assumptions.

- Adoption of one model by assumption, with subsequent revision if necessary — e.g.,

  A B C

  This is a common solution: for example, planning against the worst case, or assuming that a source of information is reliable (or unreliable) until proven otherwise. Human reasoning is often assumption-based: a single model is tentatively adopted, subject to later revision (Cohen, 1989; Johnson-Laird, 1983; Doyle, 1979; Reiter, 1980). The R/M model provides for a form of assumption-based reasoning in which assumptions are adopted, evaluated, and revised when they lead to *conflict*. Dynamic adjustment in response to feedback and new information replaces exhaustive up-front analysis. The danger in assumption-based reasoning is that we may lose track of (or be unaware of) assumptions and feel an unwarranted sense of certainty.

- Adding supplementary notation to indicate uncertainty — e.g.,

  A [B C]

  where the box indicates that the relative locations of B and C are assumed rather than reliably known. As in the previous option, spatial relations in the map continue to be isomorphic to the spatial relations in the world. At the same time, however, assumptions are explicitly noted, so they can be reevaluated in case of conflict. In the context of the R/M model, supplementary notation is used to represent metacognitive awareness of uncertainty in the analog situation model. We have already illustrated two complementary types of notation for this purpose. First, a form of "ghosted" representation (i.e., dotted lines and boxes) represents alternative possible causes and effects in the context of a particular mental model (see, for example, Figure 11 and

Figure 21). Second, when causal relations are not explicit, mental model components can be associated with *rebuttals* indicating that particular assumptions may be unreliable (see, for example, Figure 13 and Figure 22). These devices permit a thorough exploration of a particular concrete mental model, while minimizing the danger of overlooking or forgetting its shortcomings.

By contrast with analogical models, normative approaches represent uncertainty by mathematically aggregating the possibilities (such as 70% chance of attacking in the north, 30% chance of attacking in south). Normative models thus provide an abstract level of representation that corresponds to no actually realizable state of affairs. For decision making in the context of uncertainty about facts, an *expected value* is computed for each option: i.e., a weighted average of the possibilities, in which the probabilities assigned to each possible outcome serve as the weights (Raiffa, 1968). For uncertainty about values or goals, a *multiattribute utility* score is computed; i.e., a weighted average of the scores on different evaluative dimensions, in which measures of the relative importance of differences on each dimension serve as the weights (Keeney and Raiffa, 1976). Abstractions such as these can play little role in an officer's mental models of the situation since they are averages rather than real or even possible events. They cannot be visualized, anticipated, planned for, or even understood.

The R/M model addresses the question of how uncertainty can be handled in the context of analog situation models. It explains how experienced decision makers are able to exploit their experience in a domain and at the same time handle uncertainty and novelty. They construct and manipulate concrete, visualizable models of the situation, not abstract aggregations. At the same time, uncertainty is represented explicitly at the *metacognitive* level, in terms of incompleteness, conflict, and unreliability attributed to components of the first-level analog model. In response to specific types of uncertainty, metacognitive correcting strategies try to improve the first-level model or find a better one.

Generative mental models

In some particularly novel situations, combining existing structures into a single, more complex model is not sufficient to generate needed inferences or responses. In these cases, mental models based on deeper causal knowledge is required. Another tradition in the mental model literature has focused on human understanding of machines (e.g., deKleer and Brown, 1981) or the knowledge required to control physical processes (e.g., Bainbridge, 1992).

According to deKleer and Brown (1981), for example, a mental model consists, first, of a *device topology*, i.e,, a set of well-understood components, a set of well-understood *conduits* (connections by means of which components may causally affect one another), and a specification of which components are connected with which by conduits. Thus each component has a set of states it can be in, and a set of rules determining how its state will change as a function of changes in the values of conduit attributes. We can understand deKleer and Brown's notion of a "device topology" as a system of frames that represent objects and their properties, including slots to contain rules describing their causal behavior. The objects send "messages" to one another representing cause-effect relationships and triggering state changes in one another.

The process of using device topologies to generate predictions is called by deKleer and Brown (1981) *envisioning*. Envisioning is a process of propagation whereby one starts with a single input state (e.g., a candidate hypothesis about enemy intent or a friendly action option), then examines the nearby components to observe its effects, examines the nearby components of those components, and so on. Envisioning resembles mental simulation, except that it need not proceed in strict temporal order. It results in a "causal model," i.e., a dependency graph of causes and effects; e.g., if I do x, y happens; as a result, I do z, and w happens, and so on. In other words, envisioning converts a representation in terms of interacting objects (a generative mental model) into a representation of temporally and causally related events (an interpretative mental model).

A device topology is constructed by linking together familiar components (objects with their properties and rules) in novel combinations. To this extent, it is similar to the combinations of interpretative mental models that we considered in the previous section (such as the combined *intent-to-attack*, *evidence-interpretation*, and *command* structure). But interpretative mental models represent a *single, concretely realized situation*. A device topology, by contrast, is a system which has the potential for generating many different sequences of events. It provides a description of the *underlying rules or mechanisms* that give rise to different event sequences in different circumstances. Instead of relationships among actual events, this kind of model represents consistent or lawful relationships among object characteristics or variables.

For example, an interpretative model might contain the event sequence "bring tanks to the river bank — cross river — assemble on other side — ..." A generative mental model, by contrast, might contain qualitative or quantitative rules describing how the chance of successfully crossing depends on the slope of the bank and the weight and speed of the tank. An interpretative *intent-to-attack* model might include the *consequence*, "enemy penetrates friendly line," while a generative model would contain rules that determine the chances of penetration as a function of the relative numbers and types of forces, terrain, morale, momentum, and so on. In our interviews, we have found that this model can take the form of a modal ratio (e.g., 3-to-1 for defense), a set of causal factors that explain the ratio, and a set of circumstances that cause adjustments in the ratio because they causally interact with the underlying factors. As a final example, a generative mental model might predict that unless terrain obstacles are reduced by engineers, an avenue of approach will not accommodate sufficient troops to ensure successful penetration.

The decision to construct a generative mental model is one possible response to the metacognitive process of critiquing one's on-going situation understanding for incompleteness, unreliability, and conflict. As we just saw, a generative mental model can fill in details or help resolve uncertainties not specified by the more abstract *intent-to-attack* structure. Yet it is subject to the quick test step in the R/M model; i.e., generative models will not be developed unless the problem is truly novel, time is available, and stakes are sufficiently high.

Once they are built, generative mental models, like interpretative structures, are themselves subject to metacognitive critiquing and correcting. Difficulties may arise because knowledge of the component objects and their qualitative behavioral rules may be insufficient to determine the behavior of the system (e.g., if I do x, y might happen but z might also happen). When this is the case, deKleer and Brown propose that envisioning eliminates the ambiguity by making assumptions. Such assumptions may concern the existence of causally relevant but unobserved attributes, the temporal order of events, the satisfaction of rule conditions, or precise attribute values. In the R/M model, metacognitive processes play a role in keeping track of such assumptions and revising them subsequently if actually observed events conflict with the events predicted by the model.

ELICITING MENTAL MODELS

The previous section provided a framework for understanding different types of mental models and the processes, both recognitional and meta-recognitional, that influence their construction and evolution over time. In this section, we turn to an empirical investigation of the types of mental models that are required to represent tactical battlefield knowledge. Our primary focus is brigade and division level operations, as conducted by the G-3 staff, G-3, executive officer (XO), and commander.

This investigation of mental models is both partial and preliminary. We are limited by the size of the sample (23 officers and 25 incidents), the time available for the interviews (about half a day each), and the background and experience of these particular interviewees.

Method

Data. The primary source for identification of mental models was a set of transcripts of critical incident interviews and think-aloud problem-solving sessions that we conducted in a related project on modeling battlefield situation assessment skills.[1] A total of 33 interviews and problem-solving sessions were conducted with active duty officers. These officers were located at Fort Stewart, Hunter Army Airfield, Fort Leavenworth, Fort Ord, and Fort Riley.

We evaluated the 33 sessions for appropriateness to the goals of this project. Ten of the interviews (those at Fort Riley) were rejected because the brevity of the interview period (about one hour) did not yield sufficiently rich material to permit inferences about situation models. The 23 sessions that we utilized each involved a half-day interview. Five of these involved officers who had held positions at the division level. Nine of the officers had held positions only as high as the brigade level. Seven of the officers had held positions only at the battalion level and two only at the regimental level. These latter were included because the incidents described were of interest, and appeared to shed light on processes shared with division and brigade operations staff. All individuals in the selected sessions served as G3's, Assistant G3's, XO's, or S3's, with the exception of one, who was a Fire Support Officer (FSO). Table 1 summarizes information on the participants.

Table 1 also provides summary information on the elicited critical incidents and think-aloud scenario. There were 25 incidents in all (two incidents each were elicited from two of the participants). Of these 25 incidents, four involved combat in Desert Storm and two involved duty during the Los Angeles riots. Thirteen involved either command post exercises (cpx) or field exercises (field x). Six incidents involved a think-aloud problem-solving session.

The think-aloud scenario involved a defensive mission in open (desert) terrain, with the participants asked to play the role of the G-3 of a US corps. The 19 critical incidents elicited from the participants, however, varied in their missions, terrain, and the character of the participant's unit. Eleven of the incidents involved attack missions (with emphasis either on seizing territory or destroying enemy troops), while 8 involved defense. Six of the incidents involved open terrain, 9 involved closed terrain (such as mountains or dense vegetation), while 4 involved urban warfare. The interviewee participated in a heavy unit in 4 of the incidents, a light

---

[1] Contract No. MDA 903-92-C-0053 with the Army Research Institute, Fort Leavenworth Field Unit, Dr. Jon Fallesen COTR.

unit in 5 incidents, and a specialized unit in 10 incidents (5 aviation, 2 engineer, 2 artillery, and 1 fire support).

An additional source of data is a set of interviews with intelligence officers conducted in earlier research on knowledge elicitation techniques.[2] Although we draw some examples from these interviews, they are not included in the statistical summaries of results.

Analysis. The critical incident transcripts were analyzed in two stages.

Coding of significant cognitive events. In this step, cognitive events were extracted from the critical incident descriptions and subjected to a preliminary analysis. This involved:

1. Identifying interesting cognitive events described in the transcript
2. Correlating material regarding the same event that appears in different parts of the transcript, for example, when the interviewer or interviewee revisits a topic previously discussed
3. Developing an approximate time-line of the selected events
4. Characterizing each event as (A) action-related (an action, concept of an action, option, or action-related problem) or (B) situation assessment-related (a belief, concept, hypothesis, or question).
5. Identifying the sources, reasons, underlying knowledge, or doubts regarding the action or belief.

The output of this stage was a table for each critical incident, in which the rows represented cognitive events in temporal order, and the columns indicated whether the event was belief- or action-related, its content, and the reasons, underlying knowledge, or doubts pertaining to it.

Development of situation models. The next stage of analysis involved the construction of a set of situation representations underlying the significant cognitive events identified in the previous step. Construction of these representations involved:

1. Identifying clusters of cognitive events (beliefs or actions) that are linked in the analysis of step 1. Cognitive events are linked when they occur in a chain or network of reasons. In such a network, one event provides a source, reason, underlying knowledge, or reason to doubt another event.

---

[2] Contract No. MDA 903-86-C-0383 (Decision Science Consortium, Inc.) with the Army Research Institute.

Table 1
Interviews used for the analysis of mental models

Participant					Incident			
#	**Echelon**	**Position**	**Rank**	**Months as 3, XO, etc**	**Mission**	**Terrain**	**Own unit**	**Context**
1	Div	G3 plans	Maj	8	attack/territory	open	heavy	cpx
1					defend	closed	heavy	cpx
2	Bde	S3 plans/ops	Maj	55	defend	closed	heavy	cpx
3	Bde	XO	Maj	35	attack/troops	open	heavy	DesStorm
4	Div	XO	LTC	21	attack/territory	open	arty	DesStorm
5	Div	Ass't S3	Maj	35	attack/troops	open	arty	DesStorm
6	Bde	S3	Maj	38	attack	open	egr	cpx
7	Bde	XO	Maj	11	defend	closed	egr	cpx
8	Bde	XO	LTC	12	defend	closed	aviation	field x
9	Bn	S3	Maj	32	attack/troops	open	aviation	DesStorm
10	Bde	XO/S3	LTC	53	defend	open	heavy	scenario
11	Bn	S3	Maj	31	defend	open	heavy	scenario
12	Div	G3	LTC	64	defend	open	heavy	scenario
13	Bn	S3	Maj	19	defend	open	heavy	scenario
14	Bde	Ass't G3	LTC	59	defend	open	heavy	scenario
15	Div	Ass't G3	LTC	33	defend	open	heavy	scenario
16	Bn	XO	Maj	33	attack/troops	closed	aviation	field x
17	Regt	S3	Maj	16	defend	closed	light	field x
18	Bn	S3	Maj	12	attack/troops	closed	aviation	cpx
19	Bn	S3	Maj	0	attack/territory	closed	light	field x
20	Bde	S3	Maj	14	attack/territory	urban	light	field x
21	Regt	FSO	Maj	39	attack/territory	urban	fire sup.	field x
22	Bde	S3	LTC	15	defend	urban	aviation	LosAngls
23	Bn	XO	Maj		defend	closed	light	field x
23					defend	urban	light	LosAngls

2. Constructing connected graphs to represent the links among actions and beliefs.
3. Identifying category labels for beliefs and actions. For example, a belief might be labeled as pertaining to enemy *intent* or enemy *forces*. An action might be labeled as *preparatory* to attack or *positioning forces*.
4. Preliminary identification of generic structure in the set of connected labeled graphs. This involves finding recurrent patterns of linkages among category labels across the critical incidents. For example, a number of critical incidents included a similar pattern of linkages among cognitive events categorized as pertaining to enemy *goals*, *forces*, *opportunity*, *intent*, *preparatory* actions, *positioning forces*, and so on. These patterns were identified as instances of a generic structure, enemy *intent-to-attack*., even though there are minor differences among them (e.g., they may focus on different categories of activities). Other structures might be grouped together under the heading of *command* structures, *evidence-interpretation* structures, and so on.
5. Refined grouping and classification of the situation representations which have been constructed. This step involves determination of the boundaries, labeling, and type of the structures based on co-occurrence patterns, frequency of use, and meta-recognitional behavior.

Step 5 involves is to some degree a matter of judgment, guided by the pattern of data within the 25 incidents.

The boundary issue refers to the question of whether a candidate structure is a single structure or a combination created by joining two or more structures. This must be decided in part based on the pattern of correlations of occurrences across incidents among the elements of the structure. The elements within a single structure occur together more frequently than elements from different structures.[3] Another indicator that a structure is a hybrid is that its occurrence is associated with meta-recognitional strategies for handling conflict and unreliability. Such strategies can prompt the elaboration of an initial structure by joining with other structures, in order to find explanations for conflicting data or to expose and test assumptions in an argument.[4]

Decisions about the boundaries of structures interact with decisions about the appropriate labeling of structural components. Two structures that share a common element may be joined together, and their labels may be replaced by more specific labels. The result may itself occur frequently and spontaneously enough to deserve recognition as a distinguishable structure in its

---

[3] For example, Figure 20 would have been regarded as a single structure rather than the combination of several different types of structures, if its components never occurred independently. On the contrary, however, we have frequently observed the elements of *evidence-interpretation* structures in the absence of elements from the *intent-to-attack* and *command* structures; similarly, we have observed elements of each of the latter two structures in the absence of elements of the others. Co-occurrence patterns suggest that Figure 20 is a hybrid structure.

[4] The elements in Figure 20, for example, were activated in response to a simulated conflict of evidence, in which the officer had to construct explanations of an unexpected event (Figure 21 and Figure 22). It is unlikely that the entire structure was called to mind at once, since new explanations were produced only as the officer was told that earlier explanations were wrong.

own right. In effect, at the more specific level of labeling, the correlation pattern among elements is broader.[5]

Finally, there are also important methodological and theoretical issues in identifying different types of structures. A variety of discriminators are used in the identification of candidate knowledge structures. The initial indicators involve the types of entities and relations referred to in the interviewee's remarks. For example, an interpretative mental model is suggested by discussion of events, their temporal sequence, and/or their causal relations. More specific interpretative models, such as the *intent* structure, are suggested by contents appropriate to their slots, e.g., enemy *intent* or *opportunity*. Frames are indicated by discussions of spatially related objects. A generative mental model is suggested by the presence of a frame, quantitative or qualitative attributes of the objects in the frame, and causal relationships among the attributes (e.g., the number of troops versus the width of an avenue of approach determining rate of insertion of an enemy force; the slope of the river bank versus the weight of vehicles determining passability).

The output of our analysis is a relatively small set of canonical mental model structures. These structures function like the grammatical devices of a language. Instantiated and combined with one another in appropriate ways, they are capable of representing a significant portion of the situational understanding that was manifest in the critical incident interviews and problem-solving sessions.

Results

The results will be discussed in two sections: first, generic mental model structures, and second, meta-recognitional processes that are used to evaluate and modify those structures.

Mental model structures. Table 2 provides a list of the generic mental model structures that were identified and the number of critical incidents and problem-solving sessions (out of a total of 25) in which at least one instance of the relevant structure occurred.

Table 2
Mental model structures and the number of incidents containing at least one example

Model structure	No. of incidents	% of incidents (n=25)
Intent	20	80
Principles and methods	9	36
Action execution	13	52
Rate of movement	7	28
Evidence interpretation	4	16

We will discuss each of these in turn.

[5] For example, we will discuss the *intent* structure below, with components for *goals, forces, opportunity, intent,* and *activities*. But more specific and more elaborated structures based on this structure also occur frequently, i.e., *intent-to-attack* and *intent-to-defend* structures. The slots of these structures are more specific (e.g., *intent to attack* or *intent to defend*), and the *activities* slot has been expanded by joining this structure with the *principles and methods* structure. The latter supplies a more detailed set of actions appropriate to the more specific intent. There could be a still more specific structure corresponding to an enemy *intent-to-ambush* model, and so on.

Intent structures provide a framework for understanding human action in terms of its purposes. As shown in Figure 2, it includes the *intent* itself along with its possible causes and effects. The causes or prior context of an intent include *goals* that the decision maker wishes to pursue and that the intent may help achieve, *forces* (in the tactical setting) that may be capable of achieving the goal, and *opportunities* to use the forces at a relevant time and place to achieve the goal. *Goals* include higher-level intents, and *activities* may involve setting up new subgoals; thus, the *intent* structure is implicitly hierarchical. We coded the presence of an *intent* structure only if there was explicit mention of *intent, goals, forces,* and *opportunity*. We did not require explicit mention of *prior activities*, *activities* undertaken to implement the intent, or their *consequences*.

Variants of *intent* structures may be used to model enemy or friendly intent and may apply to either attack or defense. These are obtained from the more general intent structure simply by specializing the slots (e.g., enemy *goals*, enemy *opportunity*, enemy *intent* to attack). These specialized variants also contain a more detailed account of the *activities* undertaken to implement the specified *intent*.

The enemy *intent-to-attack* structure, for example, is used to generate or evaluate hypotheses about where, when, or if the enemy plans to attack. The central element in this structure is the *intent* of the enemy, defined as the assets that are to be used and the expected time and place of attack. Among the causes of *intent* to attack are the *goals* of the enemy (including peculiarities of doctrine, strategy, historical practice, or personality of the commander); the *forces* that it has available (including their strength, composition, and disposition in relation to the friendly forces); and *opportunity* (including the relations in time and space of its forces to objectives and friendly forces, with respect to mobility, terrain, and weather). Among the effects of *intent* are the *activities* undertaken to implement the attack. These activities are of four general kinds: *preparatory activities* (e.g., reconnaissance, logistics activity, engineering activity such as removing obstacles or building bridges, cross-attaching units, preparing air defense, etc.), *positioning forces* (with sufficient concentration and strength to achieve break through), *reducing opposing forces* (including deception, diversionary attacks, artillery barrage, jamming, etc.) and *launching the attack*. These actions, finally, have expected *consequences*, which typically involve the satisfaction of the original *intent* and *goals*.

Decision makers first try to fill in the slots in this structure (guided by metacognitive judgments of incompleteness); then they evaluate the resulting story in terms of conflict (with other observations or with other lines of reasoning about what the enemy is likely to do) and unreliable assumptions (e.g., are there alternative explanations or predictions that are consistent with the contents that have been placed in the slots?). If a hypothesis regarding enemy intent to attack is to be accepted, these slots must be plausibly filled in — i.e., the decision maker must be able to tell a convincing story. If the structure remains incomplete, conflicting, or unreliable, the decision maker may try to generate alternative stories to account for the data.

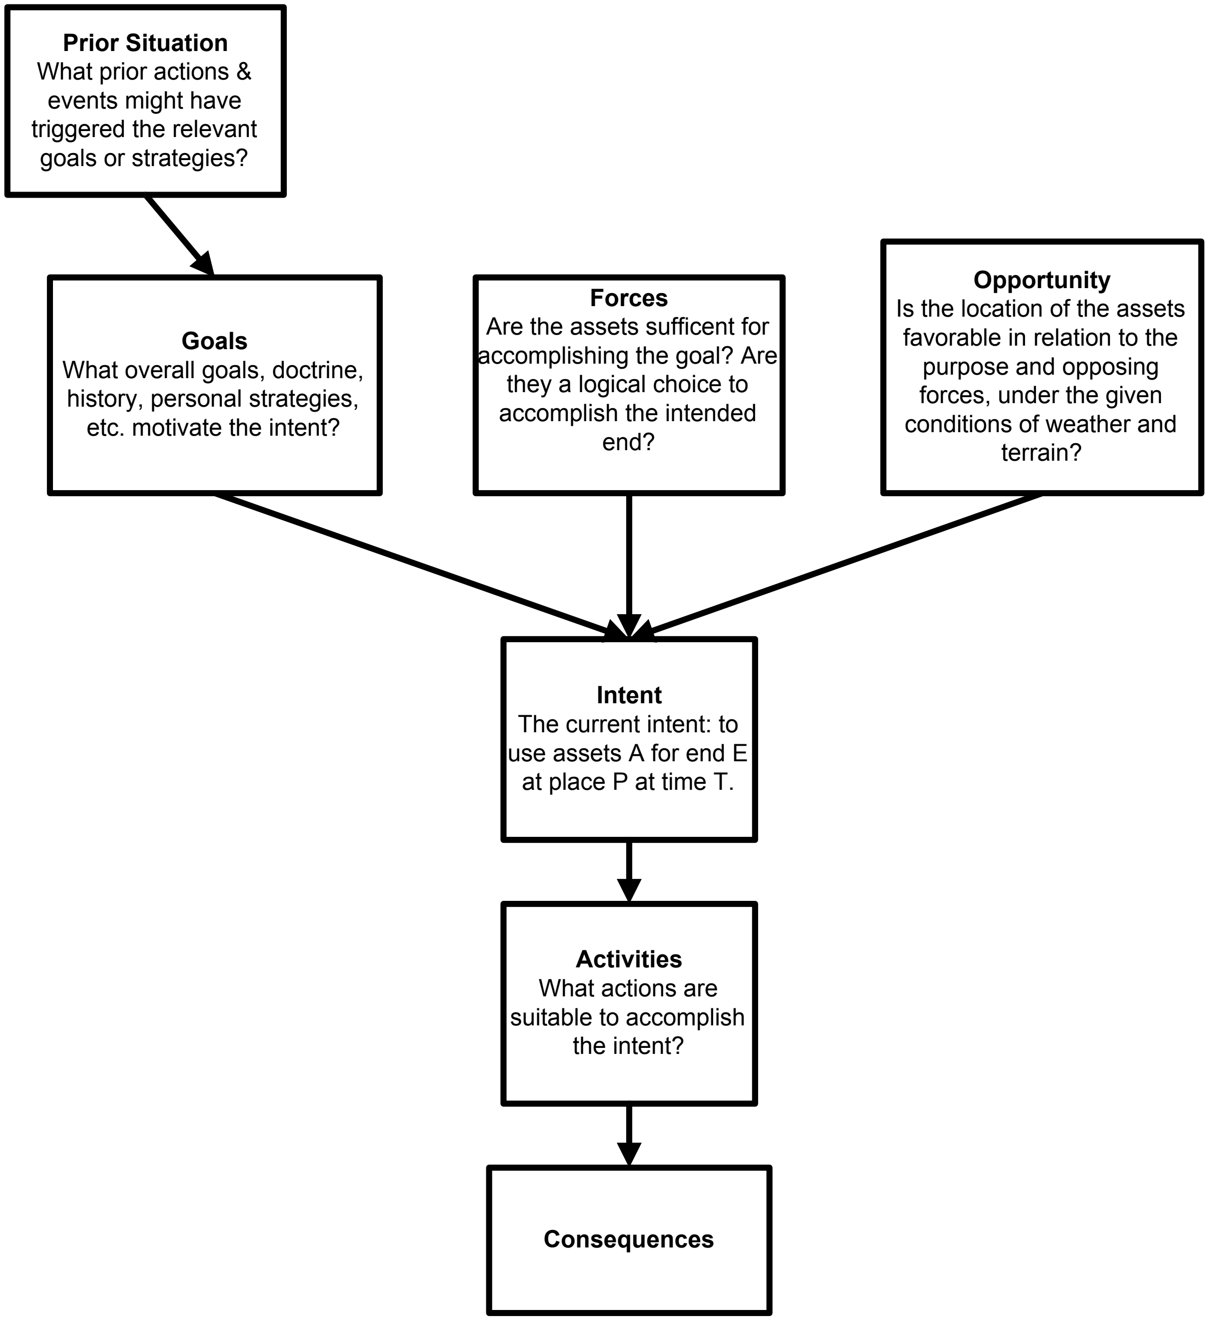

Figure 2. *Intent* model.

As an example, consider a division plans officer who is trying to predict the location of an enemy attack. The enemy has had the greatest success in the south, which the enemy is likely to want to exploit; its most likely goal, Frankfurt, is in the south; it has the best supplies in the south; and the best roads are in the south. Moreover, the enemy has been observed moving troops to the south. The planner concludes that the attack will be in the south.

Figure 3 shows how all this information fits together with an *intent-to-attack* story. It shows the causal relationships among elements of the prior context (location of objectives, prior success, supplies, mobility), intent to attack, observed activities (moving troops), and predicted activities (moving artillery).

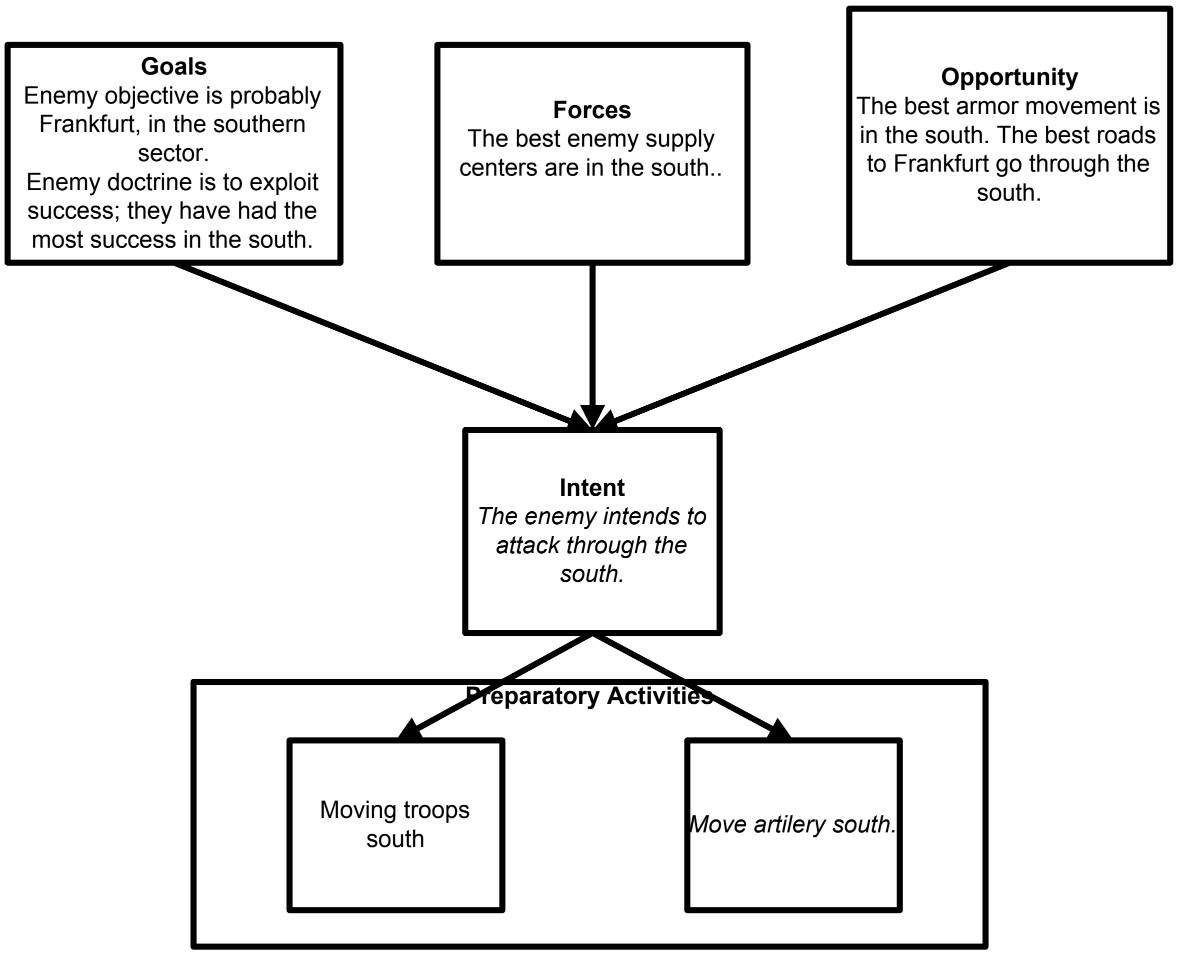

Figure 3. *Intent-to-attack* story at an early stage in an experimental scenario.

Officers used very similar structures for understanding and predicting enemy intent to defend, and for planning friendly intent to attack or defend. Interestingly, three incidents involved use of the friendly structure to predict (rather than plan) friendly intent to attack. In these cases, artillery and aviation officers used it to predict the support requirements of friendly maneuver forces so that they could anticipate where and when support would be needed. An enemy and a friendly *intent* structure may also be linked together causally. This can occur either because friendly actions are expected to influence the enemy, or because predictions regarding enemy actions influence friendly plans. (We shall return to linked structures below.) Table 3 shows the relative frequency with which these variants of intent structures were used.

Table 3
Number of incidents in which there was at least one example of a particular variant of the *intent* structure

Model variant	No. of incidents	% of incidents (n=25)
Enemy intent to attack	11	44
Enemy intent to defend	3	12
Friendly intent to attack	8	32
Friendly intent to defend	8	32
Linked enemy & friendly	8	32

Intent structures can also be distinguished according to whether their use was proactive, predictive, or reactive (Cohen, et. al., 1993). A proactive strategy attempts to influence the intent of the enemy for friendly advantage, by changing the enemy's perception of forces or opportunity. A predictive strategy attempts to predict the enemy's intent ahead of time based on an understanding of enemy goals, and the enemy's perception of forces and opportunities. A reactive strategy, finally, attempts to infer the enemy's intent after the fact, from the actions the enemy has already taken to implement it. Table 4 shows the frequencies of these different uses.

Table 4
Number of incidents in which there was at least one example of a particular use of the *intent* structure

Model use	No. of incidents	% of incidents (n=25)
Predictive	14	56
Proactive	17	68
Reactive	3	12

Proactive and predictive strategies often involve linked enemy and friendly *intent* structures. The strategies differ in whether the enemy or friendly *intent* structure has causal priority. In a proactive strategy, where the friendly plan is to influence enemy intent, friendly *activities* causally affect enemy perceptions of *forces* or *opportunities*. In a predictive strategy, where predictions regarding enemy intent influence friendly plans, enemy *activities* provide *opportunities* for friendly intent. Figure 4 and Figure 5 show how enemy and friendly intent structures may be linked for predictive or proactive uses, respectively.[6]

[6] Proactive and predictive uses of *intent* structures, however, do not necessarily entail use of linked enemy and friendly *intent* structures. One of the two structures might only be partially present, and thus fail to be coded according to our criteria.

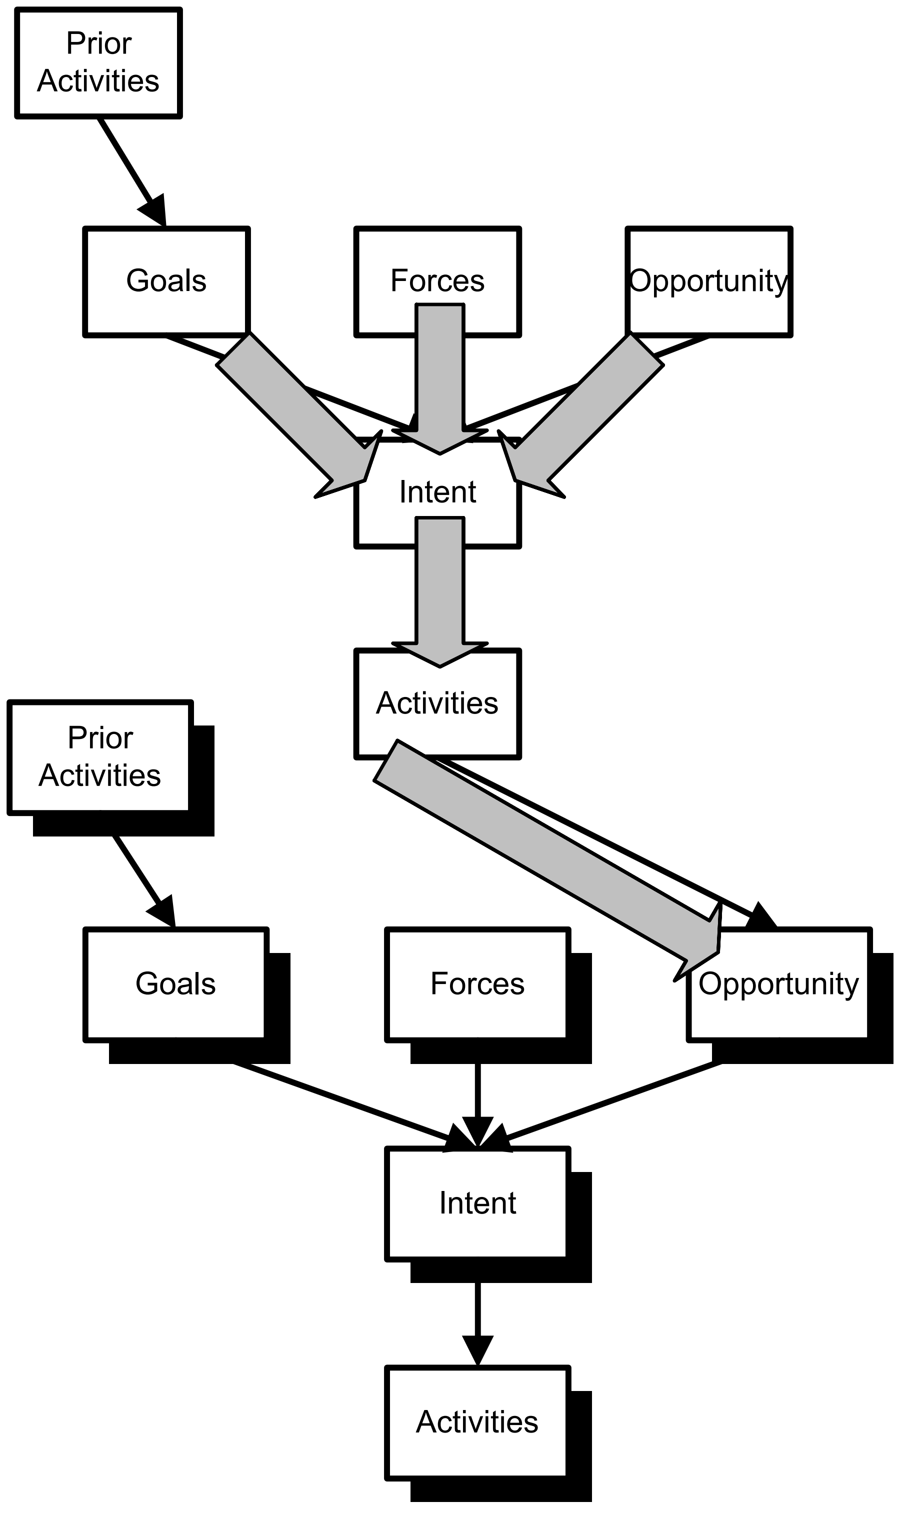

Figure 4. Predictive use of linked enemy and friendly *intent* structures. The enemy story (at the top) supplies an opportunity for the friendly plan (at bottom). Large shaded arrows represent arguments, i.e., the sequence in which ideas influenced one another.

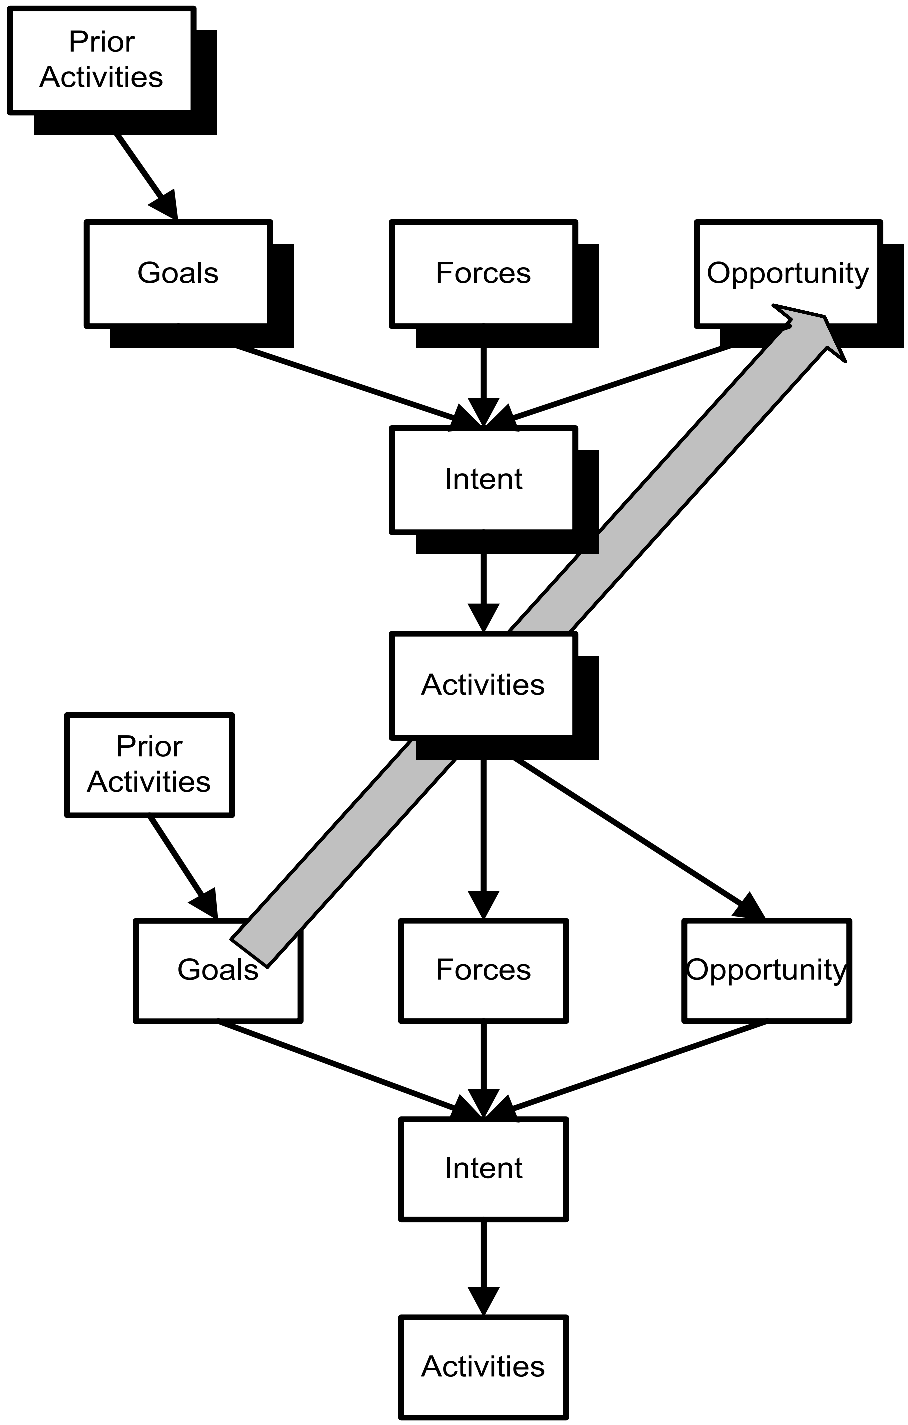

Figure 5. Proactive use of linked enemy and friendly *intent* structures. Knowledge of enemy goals or habits in the enemy story provides an opportunity that is exploited in the friendly plan. The friendly plan in turn influences the enemy's real or perceived forces and opportunities in order to shape their intent. The large shaded arrow represents an argument, i.e., the influence of knowledge about enemy goals on the perception of a friendly opportunity.

Despite its importance, the *intent* structure is only one example of structures that commanders and command staff use to interpret situations. Types of activities are not specified in the general *intent* structure, since the activities actually addressed are highly variable from situation to situation. The specification of actions draws on knowledge of how particular actions tend to achieve particular goals under particular circumstances. A representation of such knowledge is contained in the *principles and methods* structure (Figure 6). This is a hierarchical structure which decomposes goals into subgoals, subgoals into lower-level subgoals, and so on. For example, there are two broad classes of methods (or subgoals) when the intent is to attack: to increase the capabilities of own forces and to reduce the capabilities of the opposing forces. Each of these can be decomposed further. For example, one can reduce the capabilities of opposing forces by drawing them off (in regard to either the time or place of attack) or by directly weakening them (e.g., by artillery, air, or chemical attack). We coded the presence of a *principles & methods* structure only when higher-level goals were made explicit, and when multi-faceted tactics were adopted, e.g., actions were planned to concentrate and strengthen own forces, and draw off enemy forces.

Principles & methods structures can be used to flesh out the *activities* slot in an intent structure. The *intent* structure is joined with the relevant portion of the *principle & methods* structure by equating its *intent* slot with a goal somewhere in the *principles & methods* structure; the subgoals under that goal are then incorporated into the *intent* structure as *activities*.

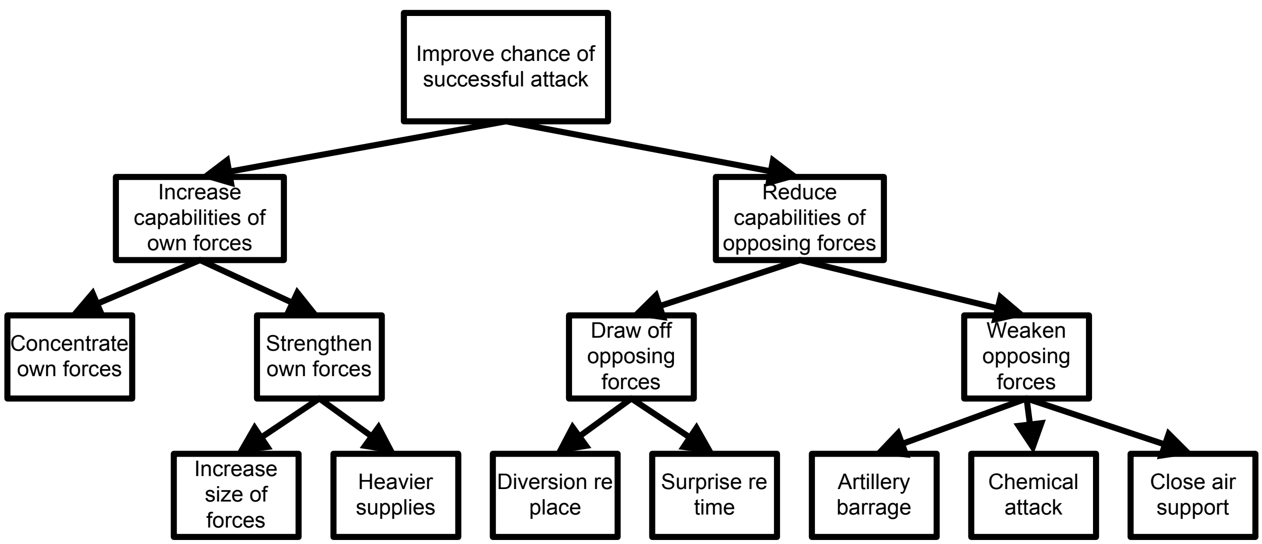

Figure 6. *Principles and methods* structure for intent to attack.

Another important class of models are *action execution* structures. These structures show in more detail how specific actions depend for their success on one another (e.g., action A must be done before action B) or upon time (e.g., action A must be initiated 1 hour before the planned time of attack). *Action execution* structures can be used to flesh out *activities* in *intent* structures and bottom-level subgoals in *principles and methods* structures. Action execution structures were scored only when officers made explicit the dependence of two or more actions on each other and/or on time.

Figure 7 is an example of an *action execution* structure representing enemy *command*. It contains a series of events: Higher echelon unit makes decision — higher echelon unit communicates command to lower echelon unit — lower echelon unit understands command — lower echelon unit carries out command. Such a structure can be used in planning how to disrupt enemy command and control. The key decision maker in a sector is identified and then methods are considered to disrupt that decision maker's effectiveness or impact, by interdicting the command and control process at one or more of these key stages.

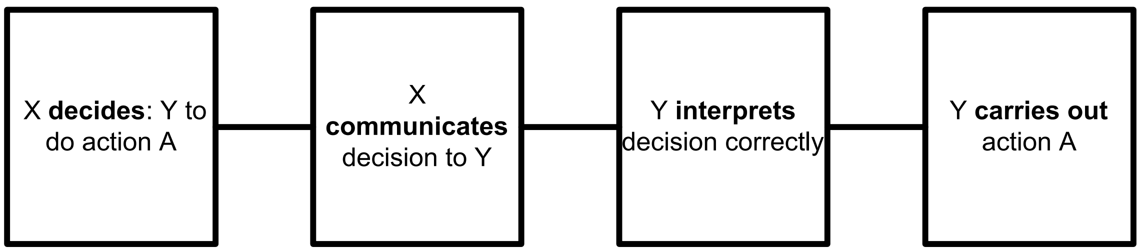

Figure 7. *Command* structure.

Another structure, called the *evidence-interpretation* structure, represents the transmission of information rather than commands (Figure 8). It supports reasoning about the plausibility of one's own situation understanding, by representing the causal stages that supposedly led from the real-world fact of interest to one's conclusion. If these links are "tight," then one's conclusion is valid. The components of this structure include: accurate observation of the event — honestly and correctly reporting one's belief about the event, including correct translation and transmission of the report — valid analysis of the report — and forming a conclusion.

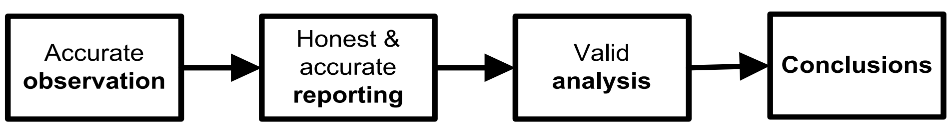

Figure 8. *Evidence-interpretation* structure.

The 25 incidents provided evidence for one type of generative structure, for predicting *rate of movement*. This structure was used in 7 incidents. As shown in Figure 9, it begins with a nominal rate of movement for the relevant type of force or vehicle and specifies how it should be adjusted up or down as a function of several variables pertaining to the enemy and terrain. This model was used in a variety of ways: to determine whether or not a given avenue of approach was feasible, to predict the time it would take if the avenue of approach were used, and to explain an actual delay either by the enemy or by friendly troops.

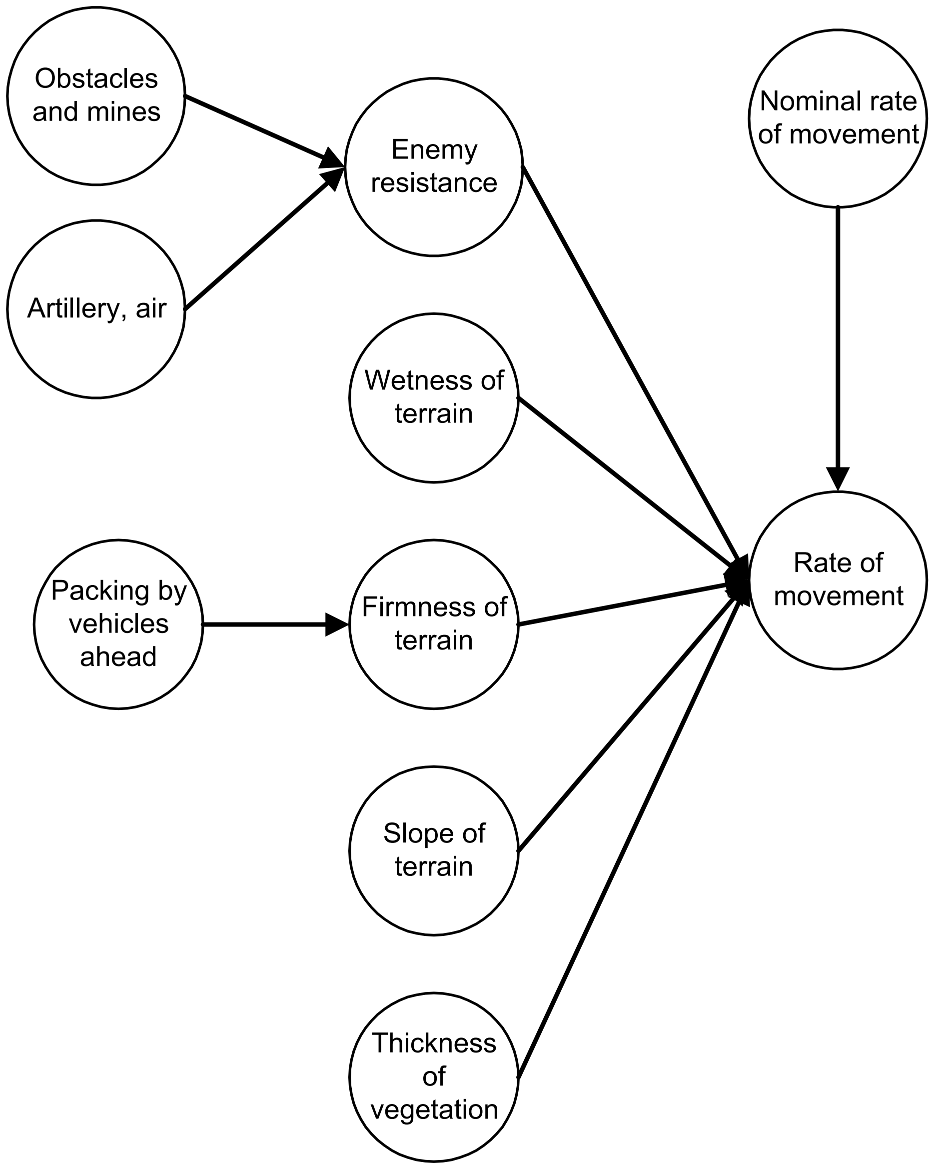

Figure 9. Generative model for predicting rate of movement as a function of different variables.

Meta-recognitional processes. Meta-recognitional processes include strategies that verify the completeness of a model, examine the reliability and consistency of the inferences it contains, monitor for observations that conflict with the model, and attempt to handle surprises by revising the model. We did not code instances of critiquing the completeness of models (although participants in the interview and problem-solving sessions frequently mentioned problems of missing data). We did score other meta-recognitional processes:

- critiquing the model for unreliability and/or conflict, i.e., generating alternative possible causes and effects within the model and looking for information that might support them
- noticing surprising or unexpected events when they occur (a form of critiquing for conflict)

- correcting conflict, by revising the model so that surprising events are explained

Table 5 shows the frequency with which each type of process occurred in the 25 incidents and problem-solving sessions. 19 of the 25 incidents involved some exploration of alternatives to the current story, in order to test its reliability. Thirteen of the incidents involved some sort of surprise.

Table 5
Number of incidents containing at least one example of a meta-recognitional process

Meta-recognitional process	No. of incidents	% of incidents (n=25)
Alternative cause/effect	19	76
Surprise	13	52
Explain surprise	7	28

The following figures, based on a think-aloud scenario, illustrate how officers can represent and deal with uncertainty using these meta-recognitional processes.

In Figure 3, the normal, recognitional meaning of each cue (prior success, a lucrative goal, supplies, roads, and moving troops) is to expect attack in the sector associated with the cue. If time is limited or the consequences of being wrong about the location of attack are not great, planners will not consider the issue further. However, when the stakes are high, time is available, and the situation is not completely routine, planners may not be content with the model based on these initial recognitional responses; they may critique it.

A first step in critiquing a model is to understood the *argument* relationships that underlie it. The causal structure in Figure 3 does not correspond to the order in which the components were activated in the course of information processing. Figure 10 shows the arguments that reflect the flow of recognition or reasoning from grounds to conclusions on this particular occasion. This flow is based on causal relationships, but may follow them in any direction. For example, recognition or reasoning may flow from knowledge of a cause to inference of an effect (e.g., inferring intent to attack from the presence of enemy objectives in the sector), or it may flow from knowledge of an effect to inference of its cause (e.g., inferring intent to attack from setting up logistics bases, movement of forces, and removal of obstacles). It can also involve predicting one effect from another observed effect of a common cause. Figure 10 shows the five arguments used to infer intent to attack in the south, in the context of the causal *intent-to-attack* structure. It also shows the argument from intent to attack in the south to a prediction that artillery will be moved south.

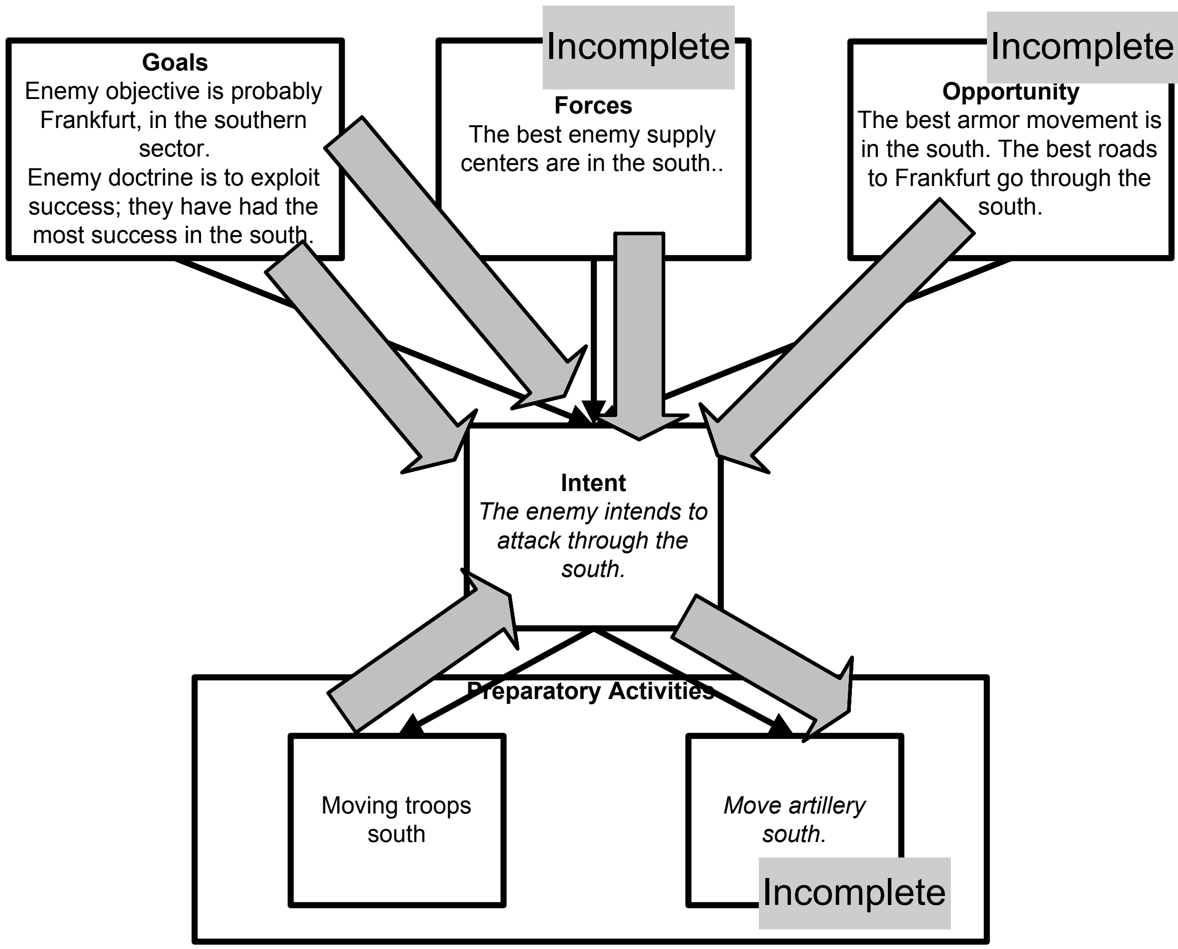

Figure 10. Large shaded arrows represent arguments in the original *intent-to-attack* story. Four arguments support attack in the south, and one argument yields a prediction regarding artillery movements based on attack in the south. Italicized items are inferred, i.e., they are the conclusions of arguments.

The first step in critiquing is to search for incompleteness in a situation model or plan. Mental model structures help in identifying such gaps. Figure 10 shows, for example, that the story is incomplete because many of the actions expected to occur prior to an attack in the south (such as moving artillery south) have not been observed. These gaps can be filled by tasking intelligence assets to the appropriate areas. The story is also incomplete because the officer has not yet fully considered all the factors that pertain to forces or opportunity. The best enemy supply centers are in the south, but what about the relative strength of the maneuver forces in the south? Similarly, roads are better in the south, but an attack will require river crossing, and the officer has not yet considered the relative river-crossing capabilities of enemy forces in the north and south.

Another function of critiquing is to find conflict, new arguments whose conclusions contradict the conclusions of existing arguments. Filling in the missing information can lead to the discovery of conflict. In this example, further analysis of enemy forces led to the conclusion that forces were stronger in the north, and that the northern commander was more experienced. Both of these provide arguments for an attack in the north, conflicting with the earlier arguments in favor of attack in the south. Thus, an alternative effect of considering the strength and disposition of forces is that the enemy will decide to attack in the north, represented by a dotted line in Figure 11. Similarly, with respect to opportunity, the northern forces had better river crossing equipment. Thus, an alternative effect of considering opportunity is an intent to attack in the north. Finally, artillery was observed moving to the north rather than the south. Intent to attack in the north is a possible cause of these observations. These alternative cause-effect relationships are shown by the dotted lines in Figure 11. The conflicting arguments that they generate are made explicit in Figure 12.

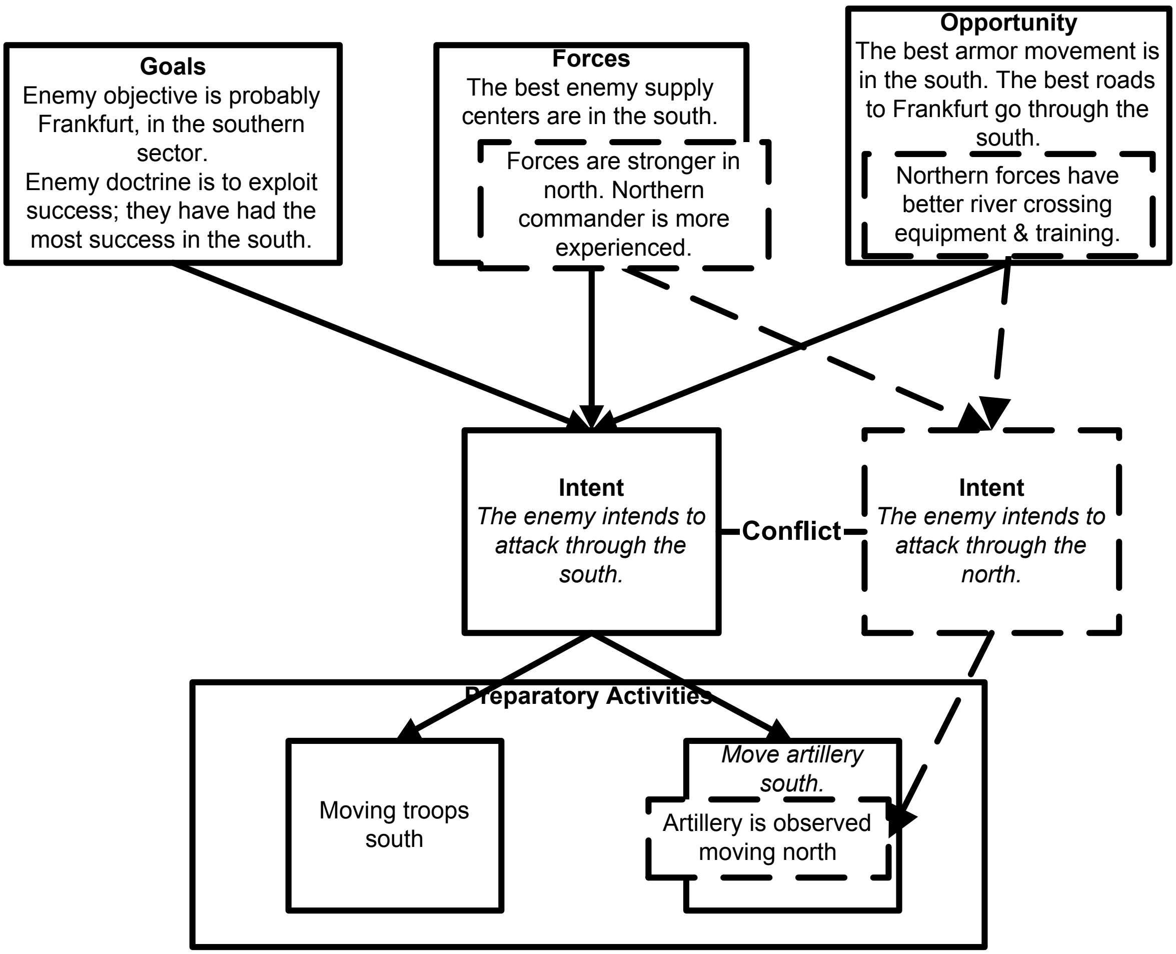

Figure 11. Fleshing out incomplete components in the story model led to awareness of conflict with assessment of intent to attack in the south. Dotted arrows and boxes represent alternative cause-effect relationships from those in the original story.

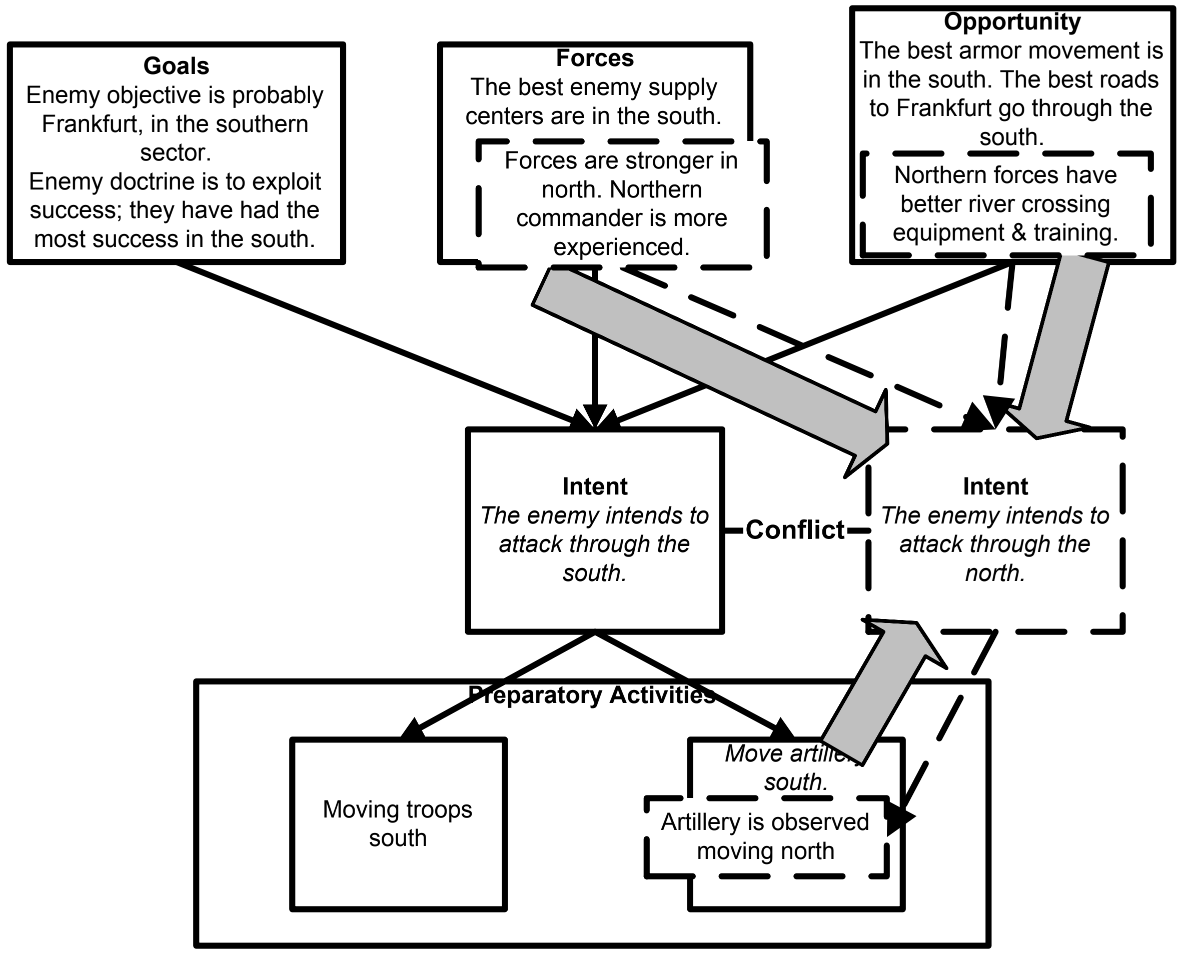

Figure 12. Large shaded arrows represent new arguments against attack in the south.

The original assessment of attack in the south is now faced with three conflicting arguments (as shown in Figure 12). Correcting steps for conflict can include retrieving or collecting new information that tips the balance in favor of one of the conflicting conclusions. Such information is not always available, however. Correcting steps may also include a reassessment of the conflicting arguments themselves. If two set of arguments point in contradictory directions, the premises in one set of arguments or the other *must* be wrong. Conflict can be resolved by adjusting the assumptions underlying one or the other of the argument sets. The officer must then evaluate the plausibility of the patched up story.

Figure 13 shows how conflict with the assessment of attack in the south can be resolved. To do so, the normal recognitional meaning of the observations supporting attack in the north must be rejected. The arguments representing these meanings are critiqued for unreliability. Understanding and planning is unreliable if the argument from evidence to conclusion, or from goals to action, is conditioned on doubtful assumptions. The argument based on the movement of artillery, for example, is based on the assumption that the artillery's range is insufficient for use from the north against the southern sector. Rejecting this assumption provides a possible rebuttal to the argument for attack in the north based on artillery movements. The argument based on the northern movement of artillery is neutralized if the artillery has a longer range than expected. Similarly, the officer generates rebuttals for all the other conflicting arguments. The arguments based on forces and opportunities are both neutralized if the northern forces and roads are to be used by the enemy in a diversionary attack.

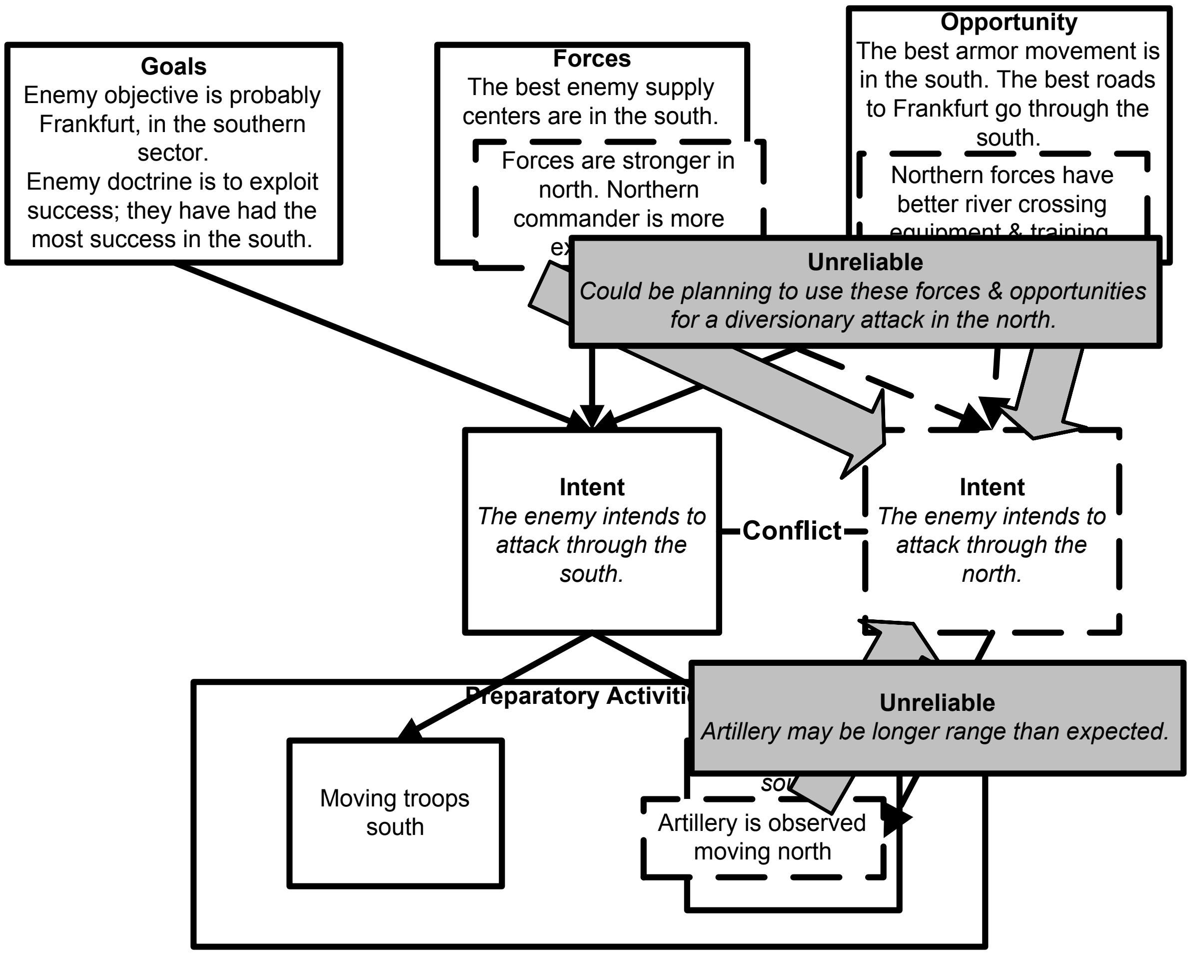

Figure 13. Assumptions required to resolve conflict with the assessment of attack in the south. Large shaded arrows represent arguments that support attack in the north. Shaded boxes overlaid on those arrows represent rebuttals, i.e., ways the conflicting arguments could fail. If these assumptions are true, the conflicting evidence is unreliable.

The original story, based on attack in the south, now is seen to depend on two assumptions (as shown in Figure 13). Since these assumptions are not totally reliable, it may be worthwhile taking seriously the alternative story, that attack will be in the north. This story, too, is subject to processes of critiquing and correcting. First, for example, the officers must fill in gaps, such as identifying a plausible objective that would be achieved by a northern attack. This story also faces conflicting arguments, namely, the considerations that led to the original view that attack would be in the south. These original arguments can now critiqued for unreliability, to see if a coherent and plausible story can be constructed around the hypothesis of attack in the north.

Troop movement toward the south is an unreliable indicator of attack in the south since there may be even more troops moving north, or the enemy may intend to move the observed troops north at the last minute. Figure 14 shows that an alternative cause of moving troops to the south is intent to attack in the north — if we add the assumption that even more troops will be moved north, or that the southern troops will be shifted north at the last minute. These assumptions represent rebuttals to the argument for attack in the south based on troop movements to the south.

To the degree that the exceptions are plausible, the argument for intent to attack in south based on troop movements is unreliable. Unreliability is different from conflict, however, because here critiquing can at best neutralize the argument for attack in the south based on troop movements, but does not provide an argument *against* attack in the south. In Figure 14 moving troops south does not provide an argument *for* attack in the north; it is merely compatible with it.

Similarly, the officers looked for rebuttals for the argument based on the presence of an objective in the south. They speculate that there may be an objective of which they are unaware in the north, or that the enemy plans a lengthier route to Frankfurt through the north.

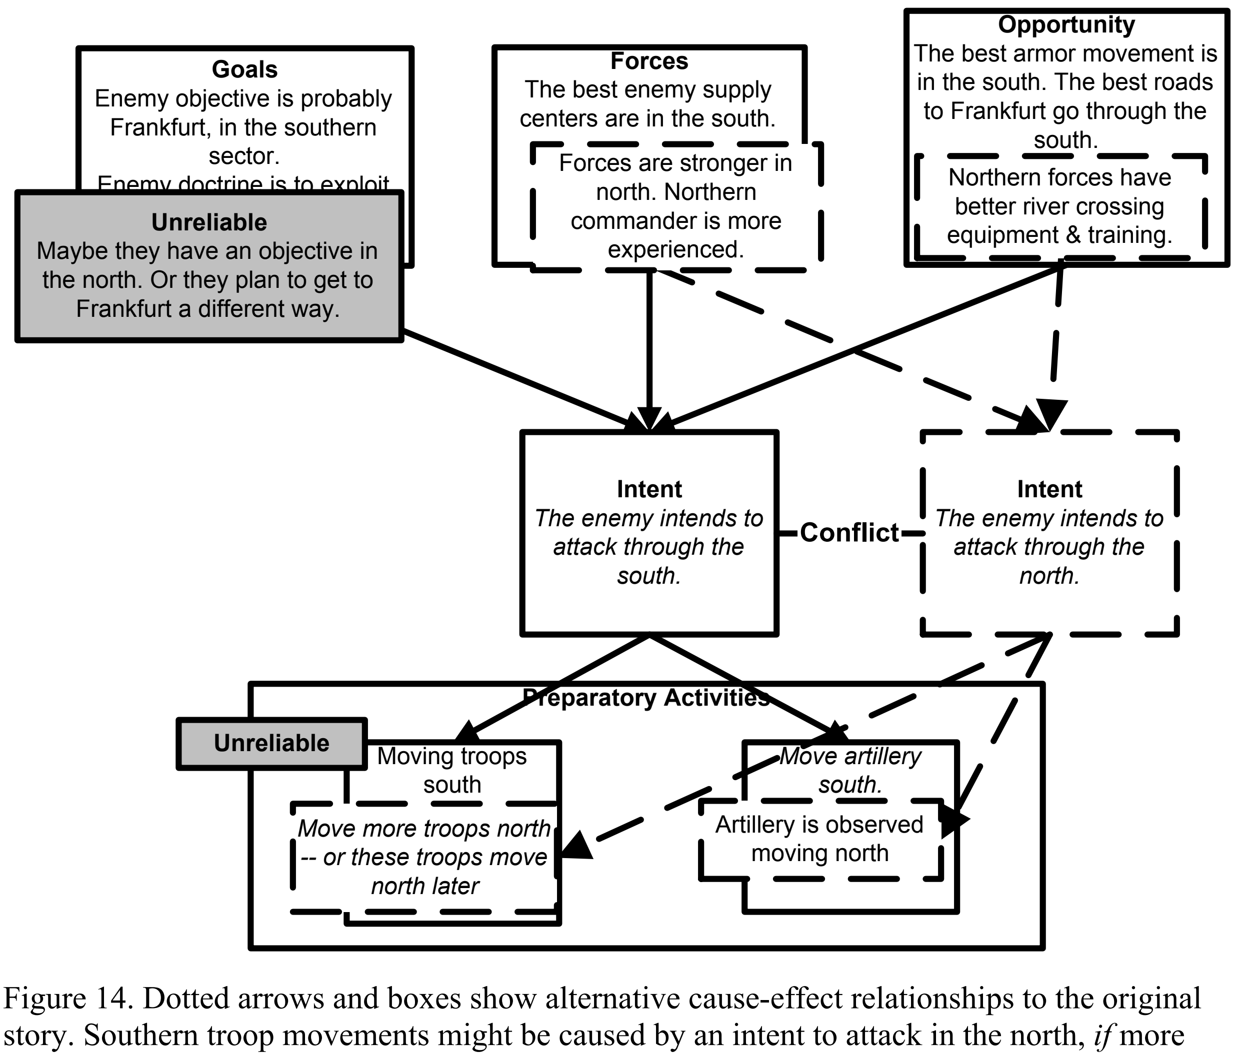

Figure 14. Dotted arrows and boxes show alternative cause-effect relationships to the original story. Southern troop movements might be caused by an intent to attack in the north, *if* more troops are being moved north, or if the southern troops will be moved north later. Troop movements in the south are compatible with, but do not provide an argument for, attack in the north.

The officer now has two coherent stories, one for attack in the south represented in Figure 13, and one for attack in the north represented in Figure 14. Because the evidence does not perfectly fit either hypothesis, assumptions were required in both cases, to explain the conflicting data. Each story will only be as plausible as the assumptions required to flesh it out and make it consistent. The officers were not satisfied with the attack-in-the-north story, because they regarded the assumptions about an unknown objective or roundabout route as especially implausible. Moreover, in reevaluating the attack-in-the-south story, they found a simpler and more plausible version, which required only a single assumption (that the enemy intends a diversion) to explain all the conflicting arguments. The final version of the attack-in-the-south story is shown in Figure 15. Compared to the original attack in the south story of Figure 10, movement of artillery is no longer regarded as preparatory for the main attack, but as part of a strategy to weaken opposing forces by diversion. Meta-recognitional critiquing has thus led to a restructuring of the situation model.

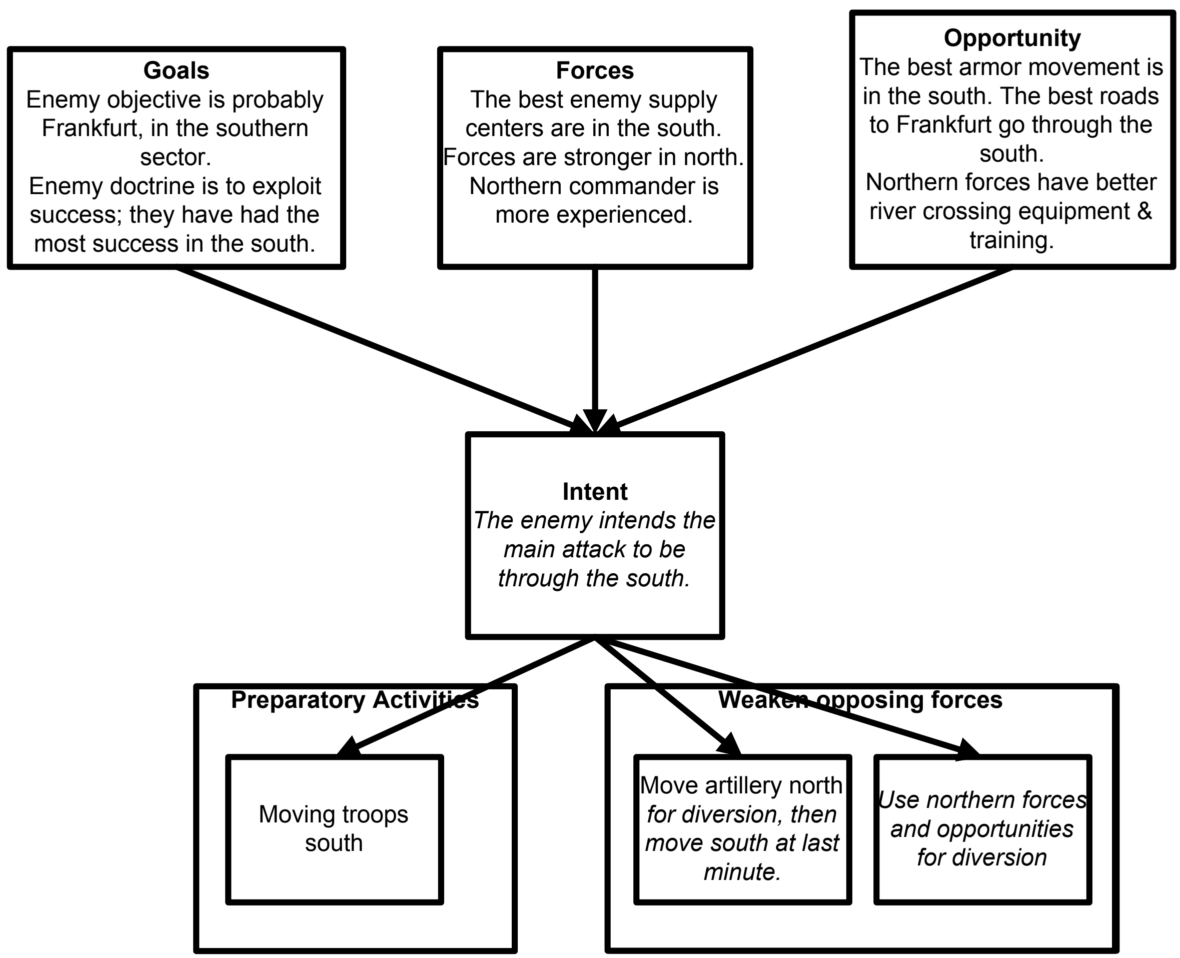

Figure 15. A coherent story based on main attack in the south. It requires assumptions regarding use of northern forces and opportunities, and artillery, as part of a diversion.

Mental models evolve both in their content and structure through iterative phases of critiquing and correcting. In the course of meta-recognitional processing, critiquing and correcting for one problem may lead to the creation and detection of other problems. In the above example, efforts to create a *complete* story based on attack in the south led to discovery of the *conflict* between superior forces and river-crossing in the north versus more plausible goals, better supplies, and better mobility in the south. The officers resolved this conflict by rejecting the normal, recognitional meaning of the evidence favoring attack in the north. They generated an alternative interpretation of these same data, that the main attack will be in the south but that a diversionary attack is planned for the north. This resolution of the conflict, however, opened the door to a new problem: *unreliability* of the assumption about a diversionary attack in the north.

Figure 16 summarizes how steps of critiquing and correcting can be linked in the R/M framework. The three types of problems explored by critiquing are shown as three points on a triangle, representing model incompleteness, unreliable assumptions in arguments for the key assessment (e.g., intent to attack in the south) or in rebuttals of arguments against the key assessment, and the existence of conflicting arguments that contradict the key assessment. The arrows showing transitions from one corner of the triangle to another represent correcting steps. It is these correcting steps that may sometimes, but not always, produce new problems. For example, correcting incompleteness in the situation model by retrieving or collecting data or by making assumptions can lead either to unreliable arguments or to conflict with other arguments. Resolving conflict by critiquing a conflicting argument can lead to unreliable assumptions in rebuttals. Dropping or replacing unreliable assumptions can restore the original problems of incompleteness or conflict. These new problems may then be detected and addressed in a subsequent iteration of critiquing.

Our analysis of critical incident interviews with Army command staff suggests an important feature of naturalistic decision making related to Figure 16. Proficient decision makers first try to fill gaps and explain conflict, and only then assess the reliability of assumptions. Thus they tend to advance from the upper right and left corners of the triangle down to the bottom, converting problems of incompleteness and conflict into problems of unreliability. In short, they try to construct *complete and coherent situation models*. They do this if possible by means of newly collected or retrieved information, but if necessary by adopting assumptions. Success in filling gaps and resolving conflict does not mean that decision makers accept the resulting situation model. But it does tell them what they must believe *if* they were to accept it. This process facilitates evaluation of a model by reducing all considerations to a single common currency: the reliability of its assumptions. If unreliability is too great, a new cycle of critiquing may expose it and trigger efforts to construct a new story.

Critiquing often reveal alternative causes and effects in other model structures besides *intent* structures. Figure 17 shows how an *action execution* structure representing a command sequence can be used to explain why an expected event has not occurred by identifying points where the expected sequence could have broken down. The chain of arguments that runs from cause to effect (i.e., from X's decision to Y's action) can be derailed at any point by an alternative possible effect.

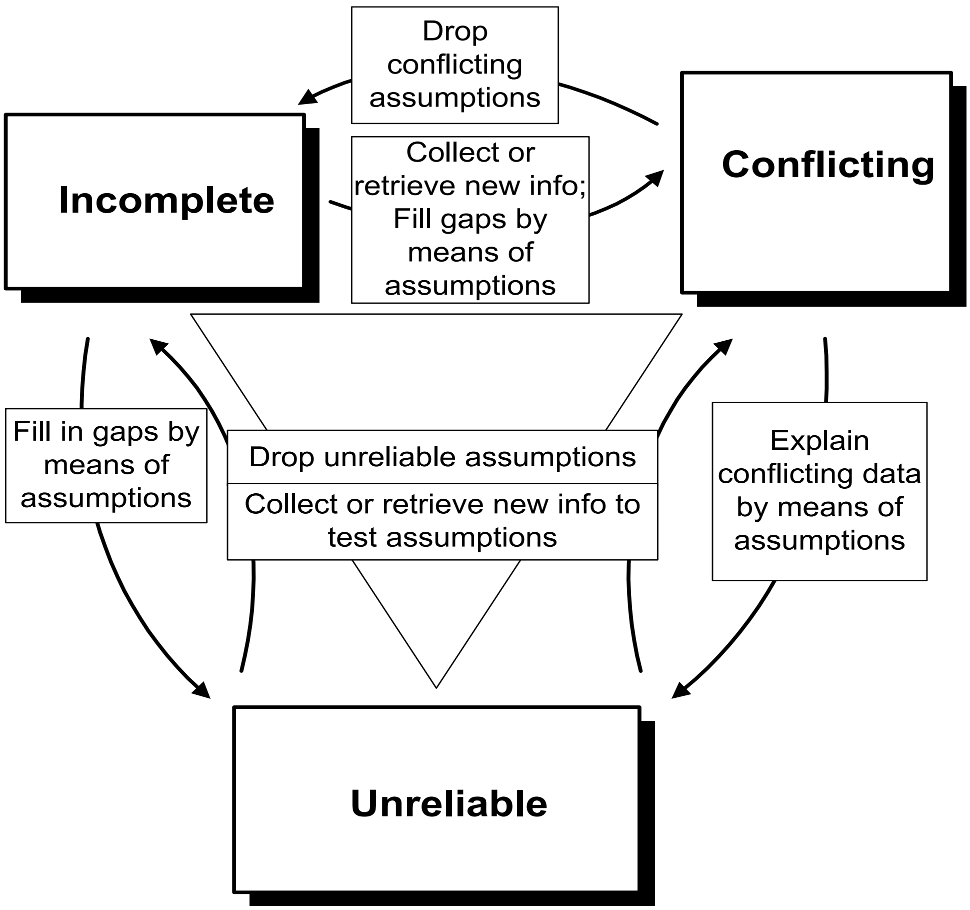

Figure 16. Ways in which correcting steps can lead to new problems in meta-recognitional processing.

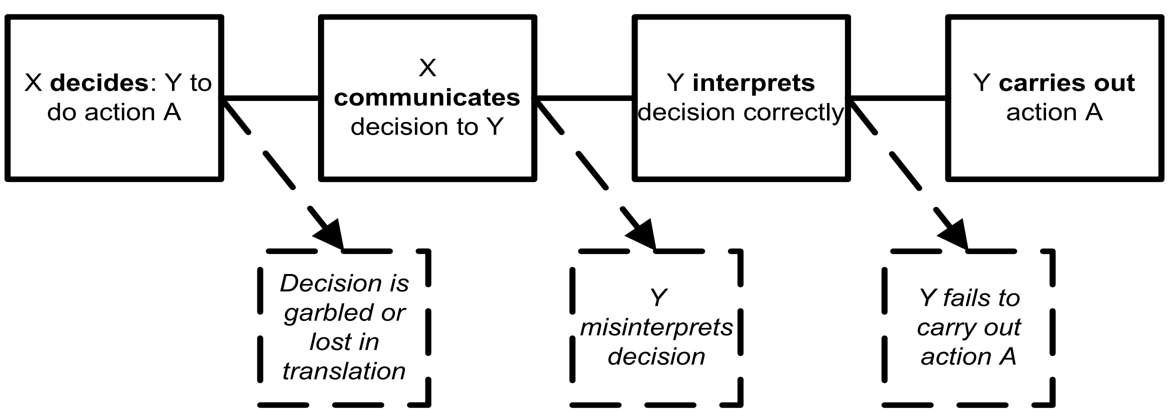

Figure 17. *Command* structure with illustrative alternative effects. Since the argument flows from causes (on the left) to effects (on the right), it can be derailed by any of these alternative effects.

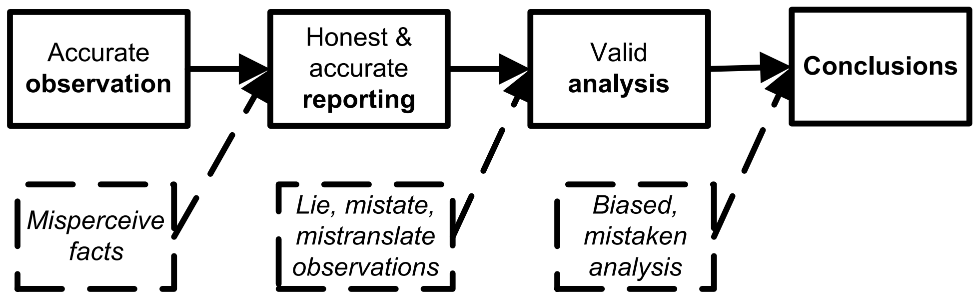

Figure 18. *Evidence-interpretation* structure with illustrative alternative causes. Since the argument goes from effects (the conclusion on the right) to causes (the accurate observation on the left), it can be derailed by any of these alternative causes.

Similarly, evidence-interpretation structures can be used to show how a conclusion could be wrong. Suppose a report is received that a vehicle belonging to a particular enemy commander has been seen in sector x. This would normally be taken as evidence that the associated enemy unit plans to attack in sector x, because moving the command post into a sector fits into the *preparatory activities* slot in the enemy *intent-to-attack* structure. But was that really the general's car? To critique this conclusion, an officer might represent all the links in the causal chain from the observation of the event to the conclusion of the analysis. These links might, for example, include: the car is spotted and its specific markings are noted by a member of the indigenous population — the person who spotted the car reports it honestly to friendly intelligence officers — the report is translated correctly — the translated report is transmitted successfully to HQ — the car markings are accurately correlated with prior intelligence about General Y's car. The chain of arguments runs from effect to cause (i.e., was the intelligence conclusion really caused by the general's car being there, or by something else), and can be derailed at any point by an alternative cause, as shown in Figure 18.[7]

Meta-recognitional processes, as we have seen, can lead to the elaboration and modification of situation models. One way a model can be elaborated is by combining it with other models, which flesh out the details of some of its components. In the following example, an intelligence officer received reports that a follow-on enemy army was at least 72 hours from a position suitable to support an attack; his assessment was that the follow-on army was required for success in an attack in his sector; he inferred that the main attack by the front-echelon army was at least 72 hours away. Figure 19 shows the *intent* structure for this situation.

In this example we combined the *critical incident* technique for knowledge elicitation with a technique we call *conflict resolution*. The officer was asked to image that an infallible crystal ball said that the attack would not be in 72 hours, and to explain how that could be. Each time the officer produced an explanation, the crystal ball told him it was wrong, and asked for another explanation. This method simulates the situation in which expectations are violated by events, and the officer tries to elaborate the current situation model to explain why.

---

[7] Even if it is the general's car, of course, the inference regarding attack could still be wrong: Someone other than General Y might now be using the car (perhaps for deception); General Y might be there but the command post somewhere else; and so on. Action-execution structures handle these possibilities.

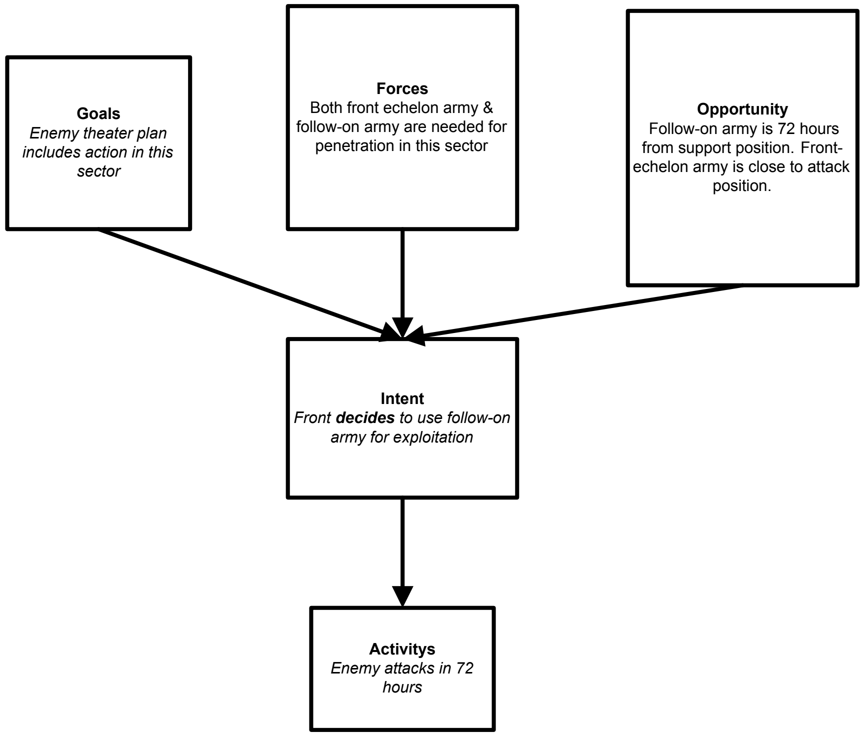

Figure 19. Story structure early in a critical incident elicitation.

The result of the conflict resolution method was a more elaborate mental model, shown in Figure 20. In this structure, the starting point of this reasoning (estimated distance of follow-on army) and the ending point of the reasoning (expected time of enemy attack) are linked by a causal chain that combines several of the knowledge structures discussed above. The link between the estimate of distance and observation of the actual distance is an example of an *evidence-interpretation* structure. A second *evidence-interpretation* structure occurs in the link between observation of the actual distance and the *enemy's* estimate of the distance. In other words, the actual distance of the army is the causal origin of two chains of events — one, consisting in scout reports, ELINT readings, etc. that results in the US estimate of the distance; and another, consisting in courier and radio reports from the army itself that results in the enemy front command's knowledge of where its own army is. The enemy's estimate of the distance in turn fills the *opportunity* slot in an enemy *intent-to-attack* structure. Other important components of this *intent-to-attack* structure are the required *forces* to penetrate in the relevant sector, hence, the need for the follow-on army, and the theater echelon *goal* of action in this sector. One of the *preparatory activities* in the *intent-to-attack* structure is a command to the front-echelon army to attack at a particular time. This command is linked to the actual attack by the front-echelon army via a *command* structure. Thus, four knowledge structures are unified into a single novel representation of the situation, by equating nodes in one structure with relevant nodes in another.

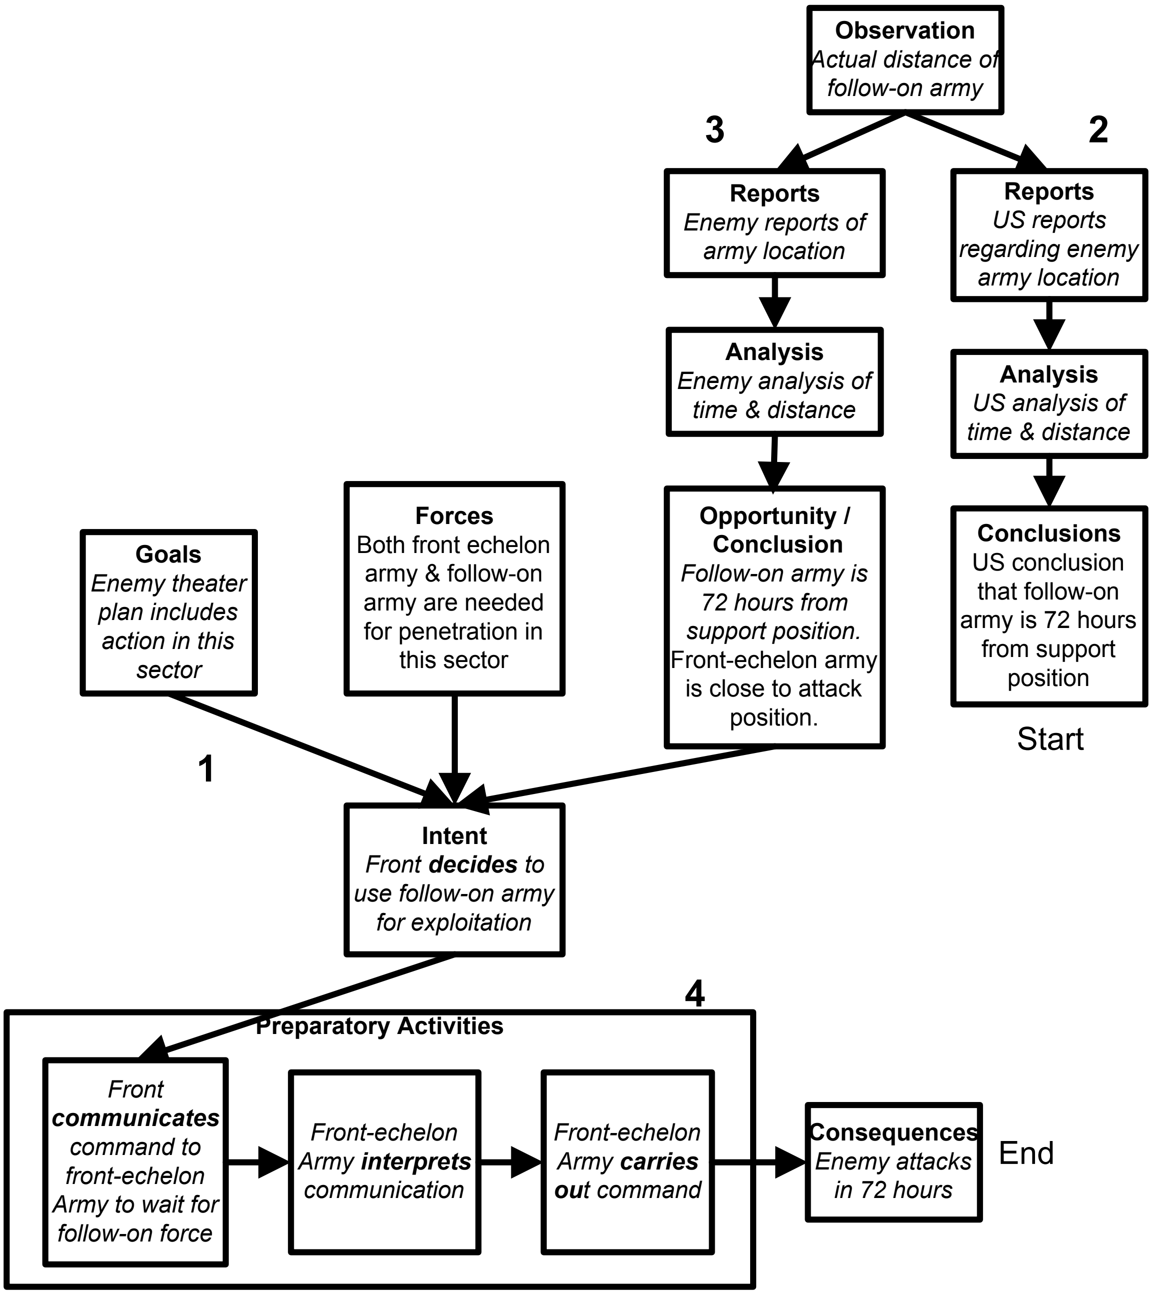

Figure 20. Combination of (1) an *intent-to-attack* structure, (2 & 3) two *evidence-interpretation* structures, and (4) a *command* structure. Note that the *observation* node is part of both *evidence-interpretation* structures; the *intent* node in the *intent-to-attack* structure also serves as the *decide* node in the *command* structure; and the remainder of the *command* structure is a set of *preparatory activities* within the *intent-to-attack* structure. The chain of argument goes from the node labeled "start" to the node labeled "end."

Mental models will not always be elaborated in this much detail. It is not usually necessary to spell out the steps of evidence interpretation involved in arriving at a conclusion about distance, or to consider all the stages of enemy command communications. Here again, the role of meta-recognitional processes is critical. The original incident involved a front-echelon enemy army apparently moving into position for attack, with an enemy follow-on army estimated to be 72 hours away from its ideal support position. The initial recognitional response to these cues was a judgment that the enemy attack was about 72 hours away and, probably, activation of an *intent-to-attack* structure. There is no reason to suppose that the *evidence-interpretation* and *command* structures represented in Figure 20 were active in the officer's mind at this time.

The structure in Figure 20 was activated by the officer as part of the effort to explain conflict. (Each time an explanation was generated, the officer was told that it was wrong, and asked to generate another.) The R/M framework predicts that conflict can be resolved by discovering and revising the assumptions underlying one of the conflicting arguments. In this case, the officer was forced to elaborate the situation model underlying his original conclusion in order to find the relevant assumptions. The explanations generated by the officer pertain to virtually every link in the combined structure. Figure 21 shows the explanations generated by the officer that pertain to the two *evidence-interpretation* structures. Figure 22 shows the explanations that pertain to the *intent-to-attack* and *command* structures. By the end of this exercise, therefore, the officer had generated something resembling the structure shown in Figure 20. The argument is now much more complex, and runs along the causal links that connect the start and end nodes. The unified representation thus provides a way of reading off new causal relationships that span across different structures (e.g., between possible erroneous identification of enemy units by our intel and enemy time of attack).

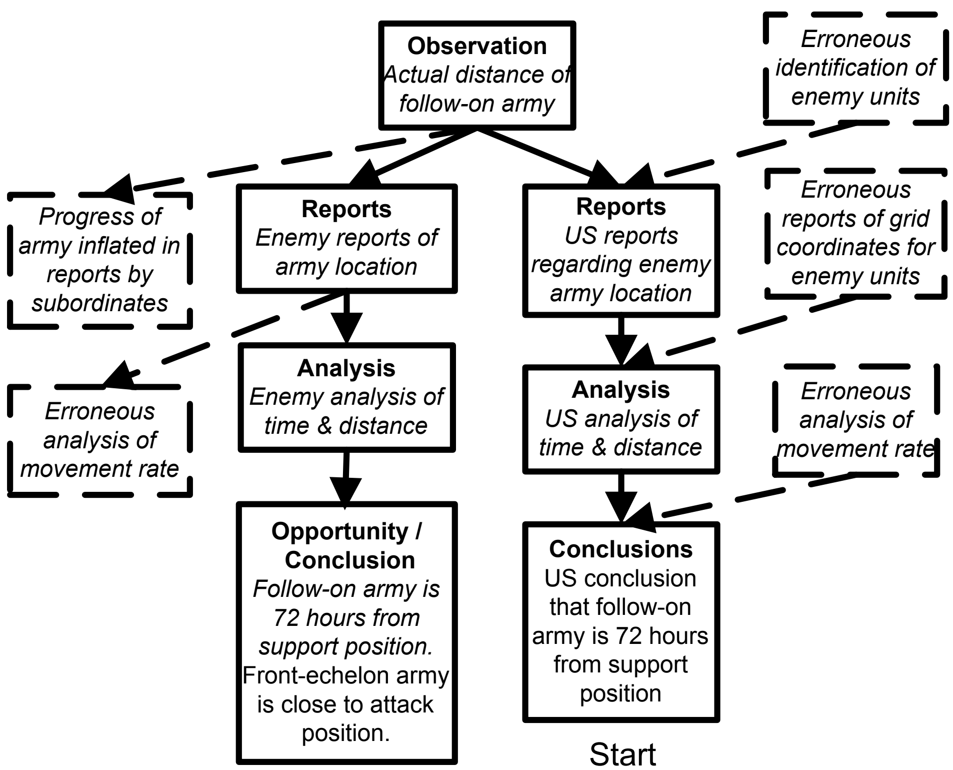

Figure 21. Explanations of conflict generated by the officer pertaining to the two *evidence-interpretation* structures. Dotted arrows and boxes represent alternative causes or alternative effects of the information in the original story.

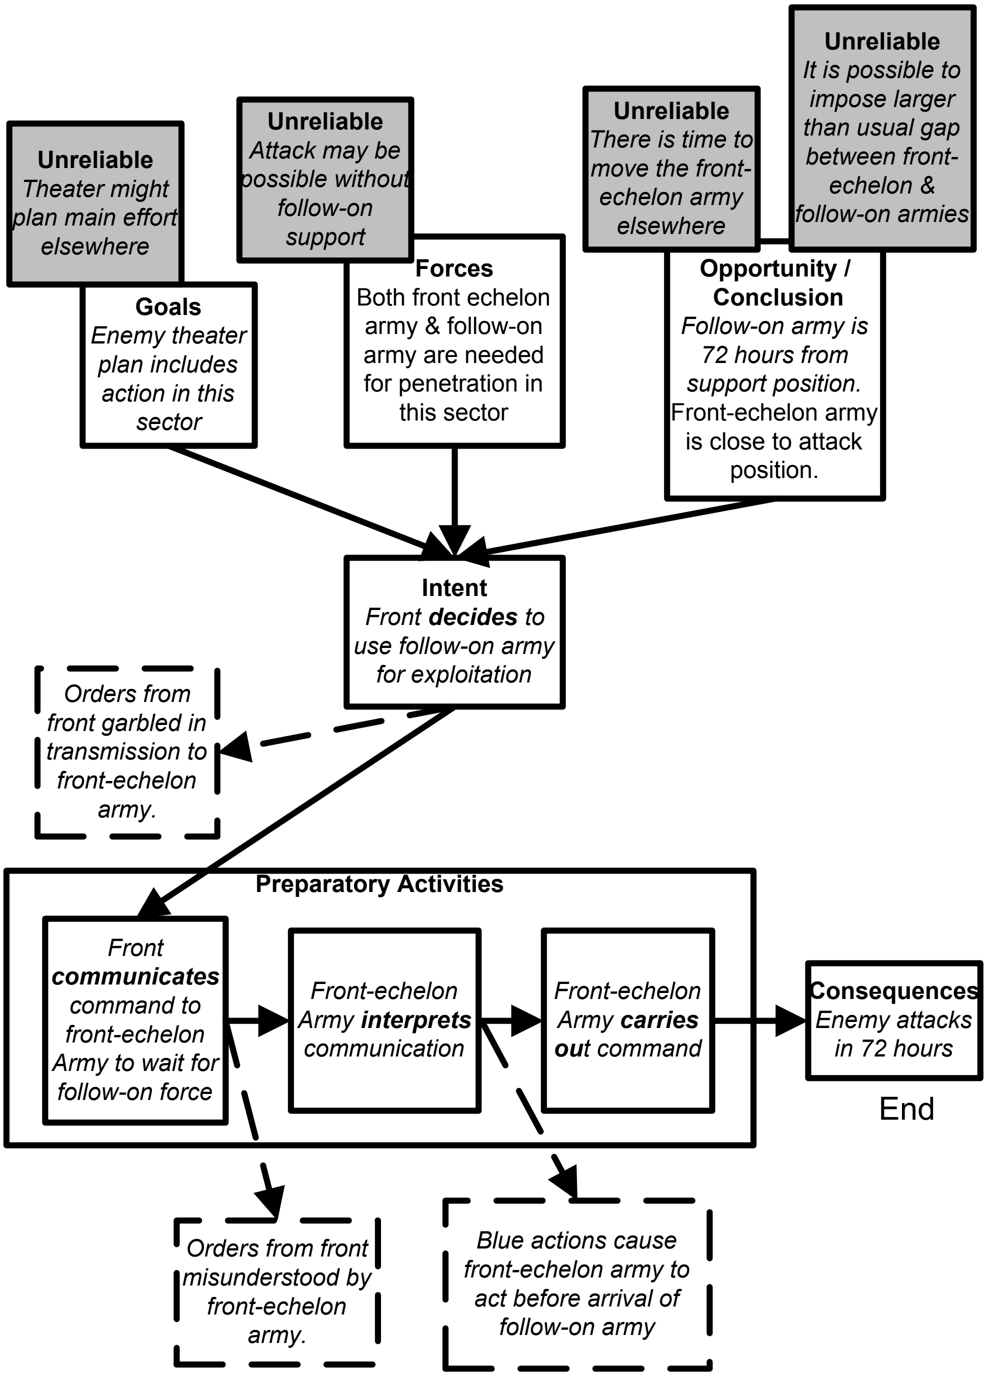

Figure 22. Explanations of conflict generated by the officer pertaining to the *intent-to-attack* and *command* structures. Dotted arrows and boxes represent alternative causes or alternative effects. Shaded boxes represent rebuttals, or exception conditions.

Theoretical Summary: Varieties of knowledge. We have found evidence for several different types of knowledge representation that might be used to support a decision makers interpretation of a situation. Other types of representation have been implicit in our discussion. An understanding of situation understanding requires, at least, the following kinds of representation: (1) learned recognitional patterns (associations among events and actions), (2) features that define selectional constraints on information that can be inserted into a slot in a mental model, (3) interpretative mental models (abstract causal/temporal structures of events and actions), (4) meta-recognitional procedures or strategies that inspect and modify mental models, and (5) generative mental models (combinations of objects with associated rules for their behavior). Figure 23 graphically illustrates how these different types of knowledge interact in situation understanding.

In the R/M model, recognition-based responses play a significant theoretical role. They activate interpretative structures and help supply their contents. For example, at the beginning of the incident addressed by Figure 3, the initial cues (enemy success in the south, enemy doctrine to exploit success, and the presence of an objective in the southern sector) recognitionally supported an assessment that the enemy intends to attack in the south (Figure 23, item 1). These cues also triggered an *intent-to-attack* structure (Figure 23, item 3). This recognitional process requires a matching of the cues to selectional constraints associated with slots in the model structure (Figure 23, item 2). Success in the south matches the features associated with the *prior activities* slot of the *intent-to-attack* structure, the presence of an objective in the south matches features associated with the *goal* slot, and the assessment of intent to attack in the south matches the features associated with the *intent* slot. These cues and assessments thus became the contents of the *prior activities, goal*, and *intent* slots, respectively. Similarly, in the incident addressed by Figure 19, the initial cues (the front-echelon army close to the FEBA and the follow-on army 72 hours away) recognitionally triggered an *intent-to-attack* structure and supported an assessment that the enemy intends to attack in that sector in 72 hours. The initial cues match the features associated with the *opportunity* slot of the *intent-to-attack* structure (the *opportunity* slot concerns the location of forces with respect to an objective of attack), and the assessment of intent to attack in 72 hours matches the features associated with the *intent* slot. *Semantic networks* are convenient for representing arbitrary sets of features and the categorical relationships among events or objects that govern their inheritance.

Slot constraints thus help activate and fill in the contents of an interpretative structure, but they do not guarantee that the contents tell a coherent story. Recognitional knowledge is required to ensure that the story makes sense. For example, the intent to attack in 3 hours, or in 100 hours, would also satisfy the constraints associated with the *intent* structure in the example of Figure 19. Recognitional *schemas* are required to support the recognition that the observed cues are associated with attack in 72 hours. As another example, cues that enemy theater objectives were in sector B would satisfy the constraints for the *goals* slot in the *intent-to-attack* structure, but would not support the assessment that the enemy intended to attack in sector A. There must be a schema indicating that the presence of an objective in a sector is associated with attack in that sector. Schemas of this sort embody a large store of accumulated recognitional knowledge, and such schemas underlie the construction of mental models for particular situations. A particular situational model, such as Figure 13, typically integrates a large number of recognitional products that have not occurred together previously. The interpretative structure highlights missing information and clarifies the causal relationships in the information that has been obtained; it thus supports meta-recognitional reasoning about incompleteness and about alternate causes and effects.

A key contribution to situation understanding is made by meta-recognitional skills (Figure 23, item 4). Such skills contribute to filling in interpretative structures and combining different types of structures through critiquing and correcting for incompleteness. They can lead to the generation of new or revised structures through critiquing and correcting for conflict and unreliability. Meta-recognitional concepts contribute in two ways to the present study: (1) They supplement the first-level situation model representation with higher-level notations, indicating alternative causes and effects, arguments, and rebuttals. These notations permit us to keep the first-level representation reasonably direct and simple, while keeping track of potential problems. (2) An understanding of meta-recognitional strategies may provide insight into when different types of first-level situation models will be called upon. The ability to anticipate the kinds of mental models that officers will utilize in different situations may be a key factor in the development of techniques for rapid capturing of their real-time situation understanding.

Finally, generative mental models also contribute to filling in the contents of interpretative mental models, in cases where recognitional associations are inadequate (Figure 23, item 5) . For example, the location of an army on an avenue of approach toward an objective may satisfy the constraints of the *opportunity* slot of an *intent-to-attack* structure. But restrictions on the passability of the terrain may make it unclear whether the army can arrive at the FEBA in time to make a difference in the battle at the expected time and place. A generative model, incorporating the size of the avenue of approach, the size of the army, and its rate of movement over the relevant terrain, may be required to determine whether the presence of the army in fact supports the intent to attack at a particular time and place. The construction of generative mental models presupposes something like *frame* representations of objects, with slots for their relevant attributes and for rules representing their qualitative or quantitative causal interactions.

Though by no means complete, this discussion provides a preliminary list of the underlying knowledge that is used to interpret situations: (1) A repertoire of learned recognitional associations, (2) semantic networks that describe the features and categorical relationships of objects and events (used to define selectional constraints on slots in mental models), (3) a repertoire of abstract interpretative mental models (such as *intent-to-attack, action-execution, evidence- interpretation,* and others), (4) meta-recognitional strategies and procedures, and (5) frames representing spatially related objects, and rules representing the qualitative or quantitative causal interactions among objects (used together to construct generative mental models).

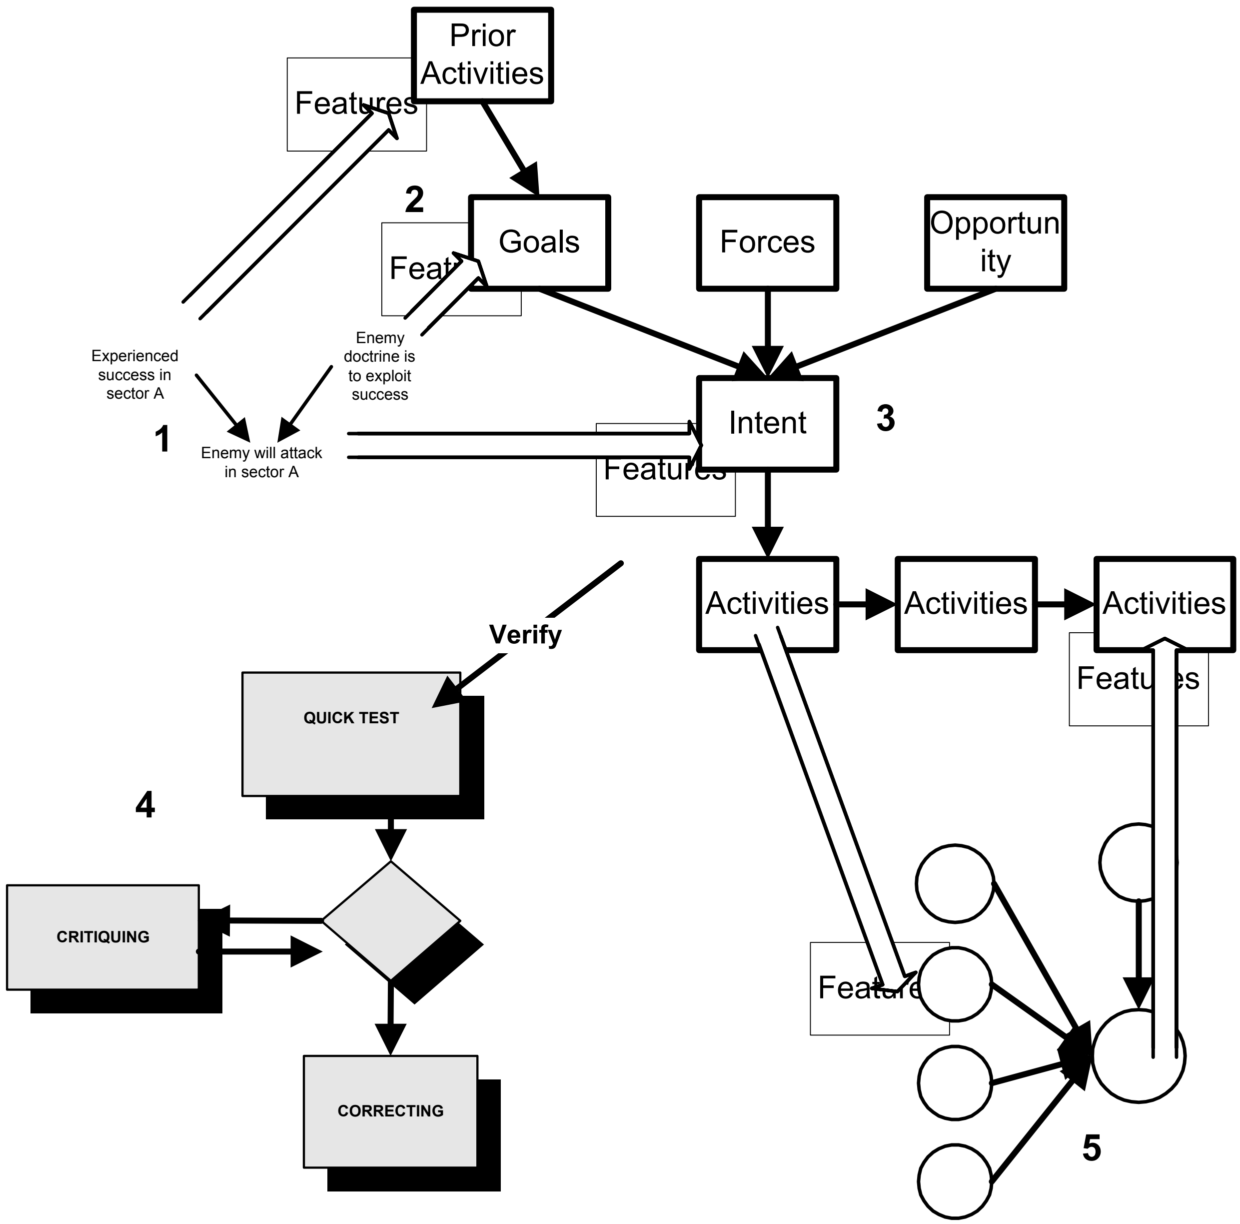

Figure 23. Different types of knowledge that interact in situation understanding: (1) Recognitional patterns, (2) semantic features, (3) interpretative models, (4) meta-recognitional strategies, (5) generative models.

PREDICTING MENTAL MODELS

We have identified a set of mental model structures that were consistently used by officers in understanding tactical situations. A mental-model capturing system might provide users with flexible tools for creating structures of these and other kinds. A more rapid model capturing capability might provide, in addition, a palette of prebuilt shapes that are most commonly used, and which users might select, reject, or modify as they choose. Still more efficiency would be achieved if the palette were dynamically tailored to the specific environment, user, and situation.

We hypothesize that mental models can be characterized in terms of features that predict when they are most likely to be used. Such features pertain to: the environments in which the structure is likely to be used, (2) the people by whom the structure is likely to be used, and (3) the immediate context of concerns to which it is likely to be relevant. The following are illustrative features in these three categories:

1 Environmental

- 1.1 Mission type (e.g., attack, defend)
- 1.2 Terrain type (e.g., open, closed, urban)
- 1.3 Own unit type (e.g., heavy, light, specialized)
- 1.4 Outlook (e.g., favorable vs. unfavorable force ratio)
- 1.5 Clarity of situation (e.g., presence of incomplete, conflicting, or unreliable evidence)
- 1.6 Time available for planning

2 Personal

- 2.1 Rank
- 2.2 Position (e.g., Commander, XO, G-3, Ass't G-3)
- 2.3 Amount of experience in the relevant position
- 2.4 Task (e.g., creating the commander's estimate, selecting a course of action, monitoring execution of plan, updating the plan)
- 2.5 Area of specialization (e.g., operations, planning, artillery, engineers, aviation, fire support)

3 Contextual

- 3.1 Types of entities focused on (e.g., objects, events, variables, attributes)
- 3.2 Types of relations focused on (e.g., temporal, causal, goal/subgoal, is-a-kind-of, is-a-part-of, is-similar-to, is-a-function-of)
- 3.3 Topics focused on (e.g., goals, forces, intent, concentrating own forces, weakening opposing forces, communicating an order, analyzing intelligence reports, wetness of terrain, thickness of vegetation)
- 3.4 Other structures being used, to which a structure is frequently joined (e.g., *intent, principles & methods, action execution, rate of movement*)

Environmental and personal features permit anticipation of a structure's use in an appropriate class of situations before it is actually used. In making these predictions, the R/M model is helpful, particularly the quick test for determining when a particular cognitive activity is worthwhile as reflected in items 1.4, 1.5, and 1.6 above. For example, the *intent-to-attack* structure might be used by a commander or G-3 in a defensive mission, when there is uncertainty regarding enemy intent, there is time to attempt to resolve it, and the costs of mistaking the enemy's intent are high because the enemy has an advantage in forces. The *evidence-interpretation* structure might be more likely to be used by the G-2 or a member of the G-2 staff, under conditions when an observation conflicts with other indicators or is inconsistent with prior expectations regarding the conclusion, and when the conclusion is important for mission success.

Contextual features, on the other hand, require concurrent behavioral evidence or inputs from the officers. Types of entities and types of relations distinguish the basic types of representation from one another, e.g., interpretative mental models (events — time, causality), generative mental models (objects, variables — causality), semantic network (objects, attributes — is-a-kind-of), causal rule (variables — causality), or frames (objects — part-whole). Content topics narrow the identification of a structure to subtypes that have appropriately labeled slots, e.g., the *intent-to-attack* structure versus the *evidence-interpretation* structure. Finally, once some mental model structures are identified, others become more likely. For example, the *principles & methods* structure is likely in conjunction with an *intent* structure, to flesh out the activities that might achieve a particular intent.

A real-time mental model capturing system can use features of the environment, person, and immediate context in conjunction, screening out irrelevant structures by means of environmental and personal features, and then selecting the most likely structures from those remaining based on the immediate context. The resulting structures, finally, are offered as options to users which they may accept, reject, or modify as they choose.

In the following analysis, we tested the feasibility of predicting mental model use, based on environmental, personal, and contextual features.

Method

Independent variables. Independent variables were a subset of the more exhaustive list above. We tested the predictive value of variables from each of the three major categories (environmental, personal, and contextual):

1 Environmental

- Mission type (e.g., attack, defend)
- Own forces type (e.g., heavy, light, specialized)
- Terrain type (e.g., open, closed, urban)

2 Personal

- Amount of experience in relevant positions

3 Contextual

- Other structures used (e.g., *intent, principles & methods, action execution, evidence interpretation, speed of movement*)

Dependent variables. The dependent variable was the occurrence or non-occurrence within a critical incident or problem-solving session of at least one instance of a specified mental model or meta-recognitional process. The procedure for identifying these structures was discussed in the previous section (Eliciting Mental Models). The following mental models structures, variants, and uses were examined:

- *Enemy intent*
- *Friendly intent*

- *Friendly* plus *enemy intent* linked
- Proactive use of models
- Predictive use of models
- Reactive use of models
- *Principles & methods*
- *Action execution*
- *Rate of movement*

In addition, the following meta-recognitional processes were examined:

- Generating alternative causes and effects (i.e., critiquing for unreliability or conflict)
- Detecting surprising events (i.e., critiquing for conflict)
- Explaining surprising events (i.e., correcting conflict)

Analysis. The analysis proceeded in three phases:

1. The aim of this step was to find a preliminary set of predictive relationships. We conducted separate *chi-square* tests of association for each combination of a qualitative independent variable and a dependent variable. For quantitative dependent variables, such as months of experience, we did a regression of the dependent variable on the independent variable. To investigate the value of one structure for predicting the use of another, we did *chi-square* tests of association among all pairs of dependent variables. It would have been impossible to test all interactions among all the variables in this step. Thus, step 1 was an initial screening for potentially significant variables.
2. This step began with the variables that seemed promising based in step 1, and explored their effects in combination with one another. An analysis of variance was constructed for each dependent variable, to test interactions among the effects of the independent variables. The ANOVA incorporated only the dependent variables that were associated with the independent variable in step 1 at a level of $p < .20$ or better.[8] It was a factorial between-subjects design, including tests for all interactions among the dependent variables. (If all interactions could not be tested due to sparseness of data in one or more cells, a series of simpler ANOVA's was conducted, combining as many of the dependent variables with one another as possible.)
3. A simplified regression model was constructed for each dependent variable. The model included dependent variables that were significant at a level of $p < .15$ in step

[8] By including varaiables that achieved a level of $p < .20$ or better in step 1, we allowed for the possibility that some variables that have only a weak predictive relationship on their own might be more predictive in combination with other variables.

2, and interactions that were significant at a level of $p < .10$ in step 2.[9] (No interactions were in fact significant at that level.)

We do not contend that a final model, capable of predicting mental model usage, can be extracted from this sample of 23 interviews and problem-solving sessions. A more reasonable goal is to demonstrate the feasibility of these methods for developing a fuller and more adequate predictive model, and to build an initial picture of what that model would look like.

Results

Results for mental model structures and meta-recognitional processes will be reported in two parts: First, the associations discovered in the first stage of analysis will be reported, in a discussion that is organized by independent variable. In other words, we will ask what a given independent variable, such as mission, unit type, terrain, and so forth, tells us about the mental models and processes that are associated with it. Second, the simplified regression models will be reported by dependent variable. In this part, we ask the converse question, how can we predict when a mental model or meta-recognitional process will be used?

Predictive associations. *Mission.* Type of mission had a significant effect on the occurrence of different types of *intent* structures. The use of enemy *intent* structures in defensive missions (71%) was almost double the use of enemy *intent* structure in attack missions (36%) ($\chi^2_1 = 3.074$; $p = .080$). There was an even stronger effect of mission on the occurrence of linked enemy and friendly intent structures: The rate of use of combined structures increased from 9% in attack missions to 50% in defense missions ($\chi^2_1 = 4.738$; $p = .030$). There was no effect of mission of the use of friendly *intent* structures; rates of modeling own intent were virtually identical in attack and defense.

Type of mission also had an effect on the use of *intent* structures as proactive (attempting to influence enemy intent), predictive (attempting to predict enemy intent ahead of time), and reactive (attempting to understand enemy intent after the fact, based on actions taken to implement it). Reactive uses were more likely on attack (27%) than on defense (0%) ($\chi^2_1 = 4.339$; $p = .037$). Conversely, there was a non-significant hint that proactive uses might be more likely on defense (43%) than attack (18%) ($\chi^2_1 = 1.724$; $p = .189$). Predictive uses were virtually the same in the two types of mission.

Mission had no effect on *principles & methods, action execution,* or *rate of movement* structures.

Overall, these results are plausible. They suggest that officers focus more on the enemy in defensive missions. These are the situations where the enemy has the initiative, often because the balance of forces does not yet permit a US attack (e.g., US forces are still arriving in theater). Proactive strategies (for example, deceiving the enemy regarding our weakest points in order to lure him into a trap) also make sense under disadvantageous circumstances. Conversely, on the attack, less effort goes into modeling the enemy ahead of time. When a surprise occurs, however,

---

[9] Abelson (1995) echoes an interesting suggestion of Tukey's that results between .05 and .15 be reported as *leaning* in the indicated direction. We decided to report such trends, along with the relevant test statistic and *p* values, and to let the reader judge.

it may trigger a modeling effort to identify and correct erroneous assumptions regarding the enemy or own forces.

The influence of mission on meta-recognitional processes supports the notion that uncertainty plays a more prominent role in attack. Significantly more surprising events were reported during attack missions (82%) than during defensive missions (21%) ($\chi^2_1 = 9.000$; $p = .003$). There was also a non-significant trend toward more exploration of alternative possible causes and effects (e.g., alternative hypotheses about enemy intent, or contingencies in a friendly plan) in attack (91%) than defense (64%) ($\chi^2_1 = 2.394$; $p = .122$). There was no effect of mission on tendency to explain surprising events.

Unit type. Unit type had a strong effect on the use of enemy *intent* structures. Specialized units (e.g., engineers, artillery, aviation, fire support) modeled the enemy only 10% of the time, while heavy and light units modeled the enemy 90% and 80% of the time, respectively ($\chi^2_1 = 14.448$; $p = .001$). There was a similar but non-significant trend for the use of combined enemy and friendly intent structures. Specialized units used combined structures 10% of the time, while heavy and light units used them 50% and 40% of the time, respectively ($\chi^2_1 = 3.860$; $p = .145$).[10] There was no effect of unit type on modeling friendly *intent*.

Unit type also had an effect on how intent structures were used. Proactive uses (to influence enemy intent) were far more likely for heavy units (70%) than for light or specialized units (0% and 10% respectively) ($\chi^2_1 = 11.213$; $p = .004$). There was no effect of unit type on predictive or reactive uses of mental models.

In addition, heavy units were more likely to use *principles & methods* structures (70%) than light or specialized units (20% and 10%, respectively) ($\chi^2_1 = 8.507$; $p = .014$). There was no effect of unit type on *action execution* or *rate of movement* structures.

These results are on the whole plausible. The role of specialized units is in general to support friendly maneuver forces. They try to model friendly intent in order to anticipate the needs of friendly commanders for engineering, artillery, or aviation support. It is less worthwhile to attempt to directly model enemy intent.

Heavy units may be more likely to be proactive because they are more able to influence enemy intent (e.g., by positioning) than light or specialized units. The greater likelihood of heavy units to use *principles & methods* structures is part of the same picture: They are more likely to adopt multi-faceted tactics, e.g., that involve concentrating own forces, and diverting and weakening enemy forces.

Meta-recognitional results add another piece to this picture. Heavy units are less likely to encounter surprises (20%) than light or specialized units (60% and 80%, respectively) ($\chi^2_1 = 5.769$; $p = .056$). There was no difference among types of units in exploration of alternative causes and effects or in tendency to explain surprises.

Terrain. Terrain had little effect on either mental model structures or on meta-recognitional strategies. Indeed, there is only one marginally significant result, but it is a plausible one: The use of a generative model to estimate *rate of movement* was more frequent in

[10] There was only one instance of modeling enemy or friend among the ten cases involving specialized units, and it was a combined structure.

closed terrain such as mountains or jungle (56%) than in open terrain such as deserts (25%), and never occurred in urban terrain (0%) ($\chi^2_1 = 4.448$; $p = .108$). In addition, there was a non-significant tendency for reactive uses of *intent* structures in open terrain ($\chi^2_1 = 3.693$; $p = .158$).

Experience. Experience was defined as the number of months served in tactical operations or planning positions, such as commander, XO, G-3, S-3, or Assistant G-3 or S-3. Among the officers who participated in the critical incident interviews and problem solving sessions, the mean experience was 30 months, and the median was 21. The minimum was 0, while the maximum was 64 months or 5.33 years. For purposes of reporting results, we will divide the officers into two groups: those above the median (more experienced) and those below (less experienced).

The most significant effect of experience was on use of the *principle & methods* structure. 50% of the more experienced officers used multifaceted tactics with explicitly identified higher-level goals, as required by this structure, while only 23% of the less experienced officers did so ($F_{1,23} = 4.042$; $p = .056$).

There was a tendency for more experienced officers to model friendly *intent* more often (75%) than less experienced officers (24%) ($F_{1,23} = 2.802$; $p = .108$). Coupled with this was a non-significant hint that more experienced officers linked enemy and friendly structures more often (42%) than less experienced officers (23%) ($F_{1,23} = 1.842$; $p = .188$). There was a related, non-significant hint that more experienced officers were more likely to be proactive (42%) than less experienced officers (23%) ($F_{1,23} = 1.992$; $p = .172$).

These effects and trends paint a consistent picture of the influence of experience. As reflected in their greater use of friendly *intent* structures, more experienced officers are more explicit about their own *goals, forces,* and *opportunities*. They are also more thorough in exploring the subgoals, or different classes of activities, that they can use to achieve their goals, as reflected in their use of the *principles & methods* structure. The *principles & methods* structure in particular supports a greater likelihood of being proactive, i.e., using tactics that will influence rather than simply predict or react to, the enemy.

Experience had no effect on meta-recognitional strategies of generating alternative causes and effects, or explaining surprises. However, it did have a significant effect on the number of surprises that occurred. Less experienced officers were surprised in 69% of their incidents, while more experienced officers were surprised only 25% of the time ($F_{1,23} = 7.801$; $p = .010$).

Other structures. Use of certain mental model structures was associated with use of certain other mental model structures. Table 6 shows the associations among structures and processes that cleared our screening threshold.[11] A convenient way to summarize these relationships is in a multidimensional similarity space. Figure 24 is the result of multidimensional scaling applied to a correlation matrix of all the model structures and meta-recognitional processes (except explaining surprise).

It is tempting to interpret the horizontal axis of Figure 24 as running from proactive processes on the right, through predictive processes in the center, to reactive processes at the left.

---

[11] There were no significant relationships involving either action execution or generating alternative causes and effects, so they are omitted from the table. Explaining surprises was not tested, since it applies only to a subset of the incidents and problem-solving sessions.

In any case, this division corresponds to distinguishable clusters of associated structures and processes.

At the right of the Figure 24 is an interconnected group consisting of proactive uses of structures, *principles & methods*, enemy *intent* structures, and linked enemy-friendly *intent* structures. As both Table 6 and Figure 24 show, the *principles & methods* structure was closely associated with proactive use of *intent* structures. This is not entirely surprising, since one of the criteria for the presence of *principles & methods* is consideration of a multifaceted tactics, e.g., both improving own forces and weakening enemy forces. Proactive use of structures is more closely associated with models of enemy *intent* than with models of friendly *intent*.

At the top center of Figure 24 is a cluster of three interconnected structures and processes: friendly *intent* structures, predictive use of structures, and generative models of *rate of movement*. Models of friendly *intent* are more closely associated with predictive uses of structures than with proactive uses of structures — the mirror image of the association of enemy *intent* models with proactive uses. Interestingly, predictive use of structures is not significantly associated with complete enemy *intent* models. (One reason for this may be the use of friendly structures by specialized units to predict friendly support requirements.) Also of interest is a related point: Officers who construct enemy *intent* structures are more likely also to construct a linked friendly-enemy *intent* structure, than officers who construct a friendly *intent* structure. In other words, friendly *intent* structures are more common in the absence of enemy *intent* structures than vice versa.

Finally, at the far left of the diagram are two nodes connected with one another and with nothing else: encountering surprises and reactive use of intent structures. Reactive structures are used to infer intent from actions the enemy has taken, rather than trying to influence intent or predict it.

Summary of predictive associations. We have found significant and near-significant associations between the use of mental model structures, on the one hand, and three kinds of variables, on the other: environmental variables (mission and unit type), a personal variable (amount of experience), and a contextual variable (use of other structures). Table 7 summarizes these relationships.

Table 6
Associations among different structures and processes

	enemy	friend	enemy+ friend	predictive	proactive	reactive
enemy+ friend	(χ^2_1 = 9.244; *p* = .002)	(χ^2_1 = 6.618; *p* = .010)				
predictive		(χ^2_1 = 6.512; *p* = .011)				
proactive	(χ^2_1 = 4.738; *p* = .030)	(χ^2_1 = 2.820; *p* = .093)				
principles &methods	(χ^2_1 = 2.707; *p* = .100)		(χ^2_1 = 3.586; *p* = .058)		(χ^2_1 = 7.767; *p* = .005)	
movement		(χ^2_1 = 2.820; *p* = .093)		(χ^2_1 = 4.738; *p* = .030)		
surprises					(χ^2_1 = 2.493; *p* = .114)	(χ^2_1 = 3.693; *p* = .055)

Predictive models. The associations that we have just discussed do not take into account either correlations among independent variables (such as unit type and terrain) or interactions among them in their predictive effects. The effects of some features may be artifacts of their correlations with other features. Alternatively, features may be associated with different kinds of effects in combination with other features than when considered in isolation; the effect of a variable may be masked by the failure to consider an interaction. We addressed these issues by constructing a regression model for each dependent variable. Each regression model started with the independent variables found in the previous section to have a marginally significant or better association with the relevant dependent variable (Table 6). The initial model also contained all the interaction effects among independent variables, insofar as the sample size permitted. The next step was to test for interaction effects among independent variables, and to drop interaction terms where they were not significant ($p<.10$). The final step was to drop main effects when they did not achieve significance in the regression model.

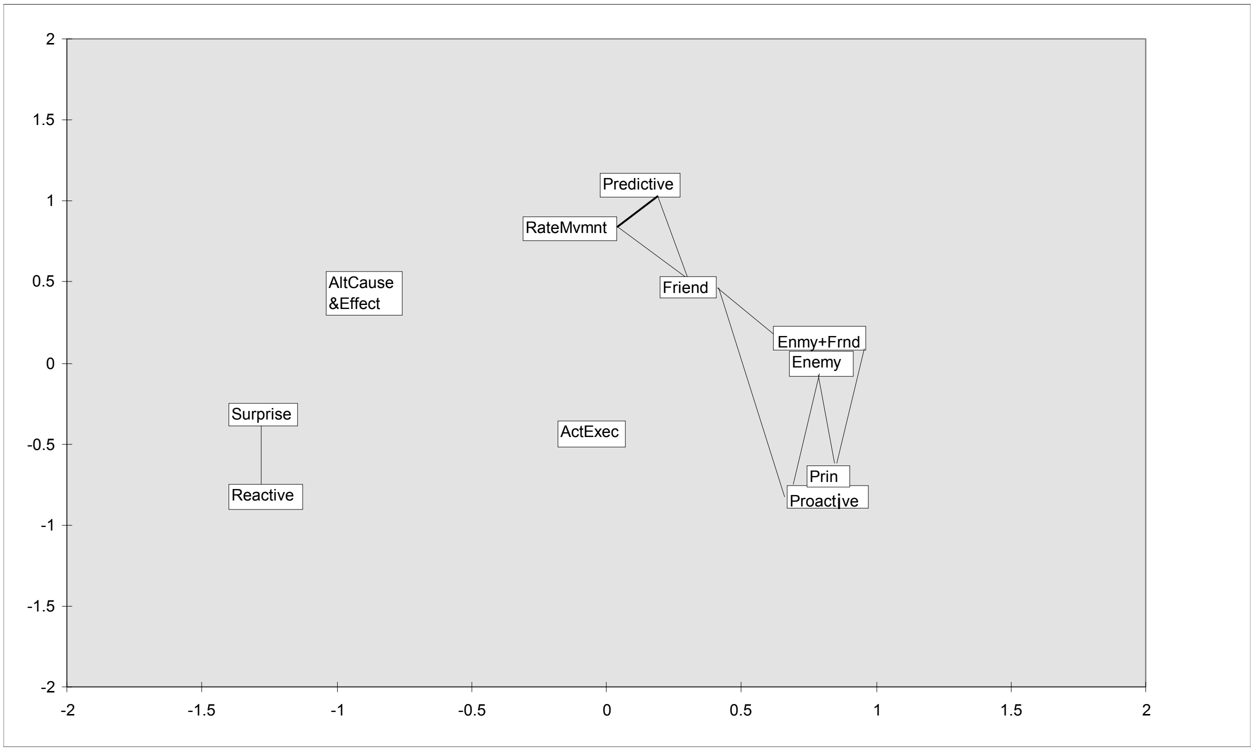

Figure 24. Multidimensional space in which distances represent a two-dimensional Euclidean best-fit to pairwise correlations among mental model structures and/or processes. Lines connect structures and/or processes for which the statistical association is significant, except when the boxes are adjacent (see Table 6).

The initial palette of structures offered by a mental-model capturing system will not be able to draw on information about the structures already built by the user. It will rely on environmental variables, such as mission, terrain, and unit type, and personal variables, such as experience. Thus, a set of models was developed for each dependent variable based on environmental and personal variables only. A second set of models was then developed which also incorporated contextual variables, i.e., information about the on-going situation understanding process.

Table 8 shows the final regression models in the absence of contextual information. Table 9 shows the final regression models with contextual information included. Each table shows the features that were included in the model for each dependent variable, along with the regression coefficient for that variable and its level of significance. The percentage of variance accounted for by the model is also shown. The estimated "probability" of observing the dependent variable in a given incident can be calculated by adding the constant to the sum of the terms representing the effects of the variables (i.e., the coefficient multiplied by the variable).[12]

[12] Coding of variables was as follows. Mission: attack = 0, defense = 1. Unit type: specialized = 0, light = 1, heavy = 2. Terrain (in predicting reactive use of models): urban = 0, closed = 1, open = 2. Terrain (in predicting rate of movement): urban = 0, open or closed = 1. Experience was measured in months. Structural variables were coded as: not present in an incident = 0, present in an incident = 1.

Table 7
Summary of associations, showing the value associated with a high level of the dependent variable, and the significance level in parentheses

Structure	Environmental	Personal	Contextual
Enemy intent	unit: not-specialized (.001) mission: defense (.080)		structures: proactive (.030), principles & methods (.100)
Friendly intent		experience: high (.108)	structures: predictive (.011), proactive (.093), movement (.093)
Enemy+Friendly	mission: defense (.030) unit: not specialized (.145)	experience: high (.188)	structures: enemy (.002), friendly (.010), principles & methods (.058)
Proactive	unit: heavy (.004) mission: defense (.189)	experience: high (.172)	structures: principles & methods (.005), enemy (.030), friendly (.093)
Predictive			structures: friendly (.011), movement (.030)
Reactive	mission: attack (.037) terrain: open (.158)		processes: surprises (.055)
Principles&Methods	unit: heavy (.014)	experience: high (.056)	structures: proactive (.005), enemy + friendly (.058), enemy (.100)
Rate of movement	terrain: not urban (.108)		structures: predictive (.030), friendly (.093)
Generate alternative causes & effects	mission: attack (.122)		
Encounter surprises	mission: attack (.003) unit: not heavy (.056)	experience: low (.010)	structures: reactive (.055), not proactive (.114)

Table 8
Predictive models for each dependent variable, without contextual information. Each independent variable is shown in the following format: *coefficient * name of variable (level of significance)*. Left column also shows percent of the variance accounted for by the model.

Dependent variable	Constant	Environmental	Personal
Enemy intent **52%**	.16	.4 * unit (.000)	
Friendly intent **11%**	.39		.009 * experience (.108)
Enemy+Friendly **19%**	.091	.41 * mission (.030)	
Proactive **33%**	-.19	.30 * unit (.003)	
Predictive **0%**			
Reactive **28%**	.078	-.26 * mission (.039) .14 * terrain (.091)	
Principles&Method **31%**	.06	.30 *unit (.004)	
Rate of movement **9%**	.00	.38 * terrain (.146)	
Generate alternative causes & effects **10%**	.91	-.27 * mission (.132)	
Encounter surprises **49%**	1.063	-.51 * mission (.004)	-.01 * experience (.025)

A comparison of the two tables shows that the addition of contextual information contributes substantial predictive value for some dependent variables. This is particularly true for modeling friendly *intent*, linked friendly and enemy *intent*, predictive use of models, *principles & methods*, and *rate of movement*. The average amount of variance accounted for without contextual information was 24%; with contextual information the average rose to 40%.

Table 9
Predictive models for each dependent variable, including contextual information. Each independent variable is shown in the following format: *coefficient * name of variable (level of significance)*. Left column also shows percent of the variance accounted for by the model.

Dependent variable	Constant	Environmental	Personal	Contextual
Enemy intent **52%**	.16 +	.4 * unit (.000)		
Friendly intent **49%**	.13 +			.62 * predictive (.001) .51 * proactive (.004)
Enemy+Friendly **61%**	-.33 +	.26 * mission (.068)		.41 * enemy (.008) .43 * friendly (.005)
Proactive **44%**	-.19 +	.30 * unit (.002)		.33 * friendly (.046)
Predictive **34%**	.19 +			.42 *friendly (.011) .32 * movement (.030)
Reactive **38%**	-.35 +	.23 * terrain (.008)		.35 * surprises (.005)
Principles&Method **40%**	.53 +	.19 *unit (.093)		.36 * proactive (.096)
Rate of movement **27%**	-.20 +	.35 * terrain (.144)		.40 * predictive (.031)
Generate alternative causes & effects **10%**	.91 +	-.27 * mission (.132)		
Encounter surprises **49%**	1.063 +	-.51 * mission (.004)	-.01 * experience (.025)	

CAPTURING MENTAL MODELS

Overview

A system for rapid capturing of battlefield mental models can build on the results reported in the previous two sections. First, it can provide general tools for creating any interpretative or generative mental model, and for annotating problems such as incompleteness, conflict, and unreliable assumptions in such models. Second, it can specifically facilitate the construction of mental models like the ones that officers actually used in a variety of critical incidents and problem solving sessions. Thirdly, it can respond sensitively and dynamically to information about the environment, user, and context to facilitate the more specific structures that are likely to be appropriate. Finally, it can monitor the structures that a particular user creates and modify its dynamic responses to more closely match that user's personal style.

As part of the Phase I effort, CTI has developed a proof-of-concept prototype that demonstrates a number of these ideas for rapidly representing and analyzing mental models. The prototype has been applied to a scenario developed by the Army Research Institute, called Arisle.[13] The system, however, is not limited to that application; a wide variety of situations may be readily depicted and explored. Although the system is built primarily from off-the-shelf software, many of its components are fully functional sketches of the capabilities we intend in the Phase II design.

This prototype is intended to show how mental models can be constructed interactively. The user works with a palette of shapes contained in a Structure window, which represent component structures of mental models, and one or more models that are created in a Workspace window. By dragging objects, such as *intent* or *action-execution* structures, from the Structure window palette and dropping them into the Workspace, users are able to quickly build up mental models that correspond to critical aspects of their situation understanding. Other objects in the Structure window palette, such as *assumptions* or *rebuttals*, enable the user to elaborate these mental models in ways that highlight the unreliability, incompleteness, or conflict that may be present in the model.

Mental models can be used to organize and interpret a vast quantity of information. Such information arrives in familiar forms, such as intelligence estimates and reports, mission statements, commander's guidance, commander's estimate, frag orders, spot reports from units in the field, and so on. This information flow is available electronically to users in the Estimate window, in a hierarchical outline with hyper-media links to supporting documents. Information from the Estimate window can be inserted into mental models that the user creates in the Workspace. Users may insert text, maps, diagrams, or other models directly into the model under construction. To simplify the visual display, they can type summary statements in the displayed portion of the model, and insert supporting material into *backing* boxes that are available by hot button from the structure itself.

The model capturing process may proceed in either a top-down or a bottom-up fashion. In the top-down case, users begin constructing a model that reflects their current understanding of the situation. They then seek information that fleshes out, and confirms or disconfirms, different aspects of the model. In the bottom-up case, users begin with a flow of information (i.e., intelligence estimates and reports, mission statements, commander's guidance, commander's estimate, frag orders, and so on). They then seek appropriate ways to organize that information within a set of one or more mental models.

The proof-of-concept system supports both top-down and bottom-up strategies, with processes that we call *advisory* functions. For top-down processing, the system examines the models under construction in the Workspace, and finds and recommends material in the Estimate

---

[13] This scenario, developed by Dr. Rex Michel, provided much of the intelligence and command estimate material that we utilized in the demonstration system.

window that is likely to be relevant to it. Also based on its perusal of the Workspace, the system recommends ways to elaborate and complete the on-going model by means of elements from the Structure window. For bottom-up processing, the system recommends models in the Structure window that seem appropriate to the information that the user is considering in the Estimate window.

We envision a variety of potential applications for a fully developed version of this system:

- As an individual decision aid, it could support processes of situation understanding, planning, and operations by command staff. The advisory functions would facilitate the construction of appropriate mental models, and guard against pitfalls such as overlooking conflicting evidence, relying on unsubstantiated assumptions, or failing to consider important factors.
- As a team decision aid, it could support the communication of shared situation understanding and plans. Users might take mental models they have constructed in the Workspace, and insert them into documents being prepared for distribution in the Estimate window (e.g., the commander's estimate). Members of the team to whom such documents are circulated might use the tools provided by the system to mark up the models with their own concerns regarding potentially unreliable assumptions, gaps, or conflicts, or to produce alternative models for consideration if they think it necessary.
- As a training aid, the system could be used both to diagnose situation understanding skills and to improve them. With advisory functions turned off, the system could monitor the success of trainees in building appropriate mental models, filling gaps, noticing crucial information, handling conflicts, and so on. (Such diagnostic functions could also be used for personnel selection.) With the advisory functions turned on, the system could assist trainees in learning the structures and processing skills that underlie successful performance.
- As a research and knowledge engineering tool the system could be used to elicit and model the knowledge and processing strategies utilized in different domains.

Main components of the system

For the user, the system consists of three main windows for interaction. These are the Estimate window, the Structure window, and the Workspace window. The Estimate window makes available to the user a wide variety of documents, such as the Intelligence Estimate or the Commander's Estimate, in electronic form. The Workspace is the area in which the user constructs visual diagrams that represent mental models for the selected task. The Structure window is a collection of both basic and complex shapes. These shapes may be dragged by the user from the Structure window into the Workspace. These operations are discussed in some detail below.

Estimate window. The Estimate window provides the user an interface to electronic documents pertaining to the current situation. In the demonstration system, electronic copies of the Intelligence Estimate and the Commander's Estimate for the Arisle scenario are made available. The Estimate window organizes these documents in a collapsible outline form. The documents have been divided into subsections following the standard form for the documents. The text of any subsection associated with a heading in the outline is visible at the bottom of the Estimate window, as shown in Figure 25 and Figure 27. Headings with embedded content are represented by dark gray circles if collapsed, and by dark gray rings if expanded. Headings without embedded content are always represented by light gray rings. Finally, some headings are associated with light-blue lightening bolts on the far right hand side of the Estimate window, e.g., Figure 26. These icons represent hyper-media links to other documents. The demonstration system has examples of such links that lead to detailed area maps, the auxiliary *Status of Forces* document, and some mental models of the Arisle situation.

Structure window (Palette). The Structure window simplifies the task of constructing visual representations of mental models. Various kinds of pre-defined structures are combined into a palette of both simple and complex shapes (see Figure 28). The user may very easily place copies of these shapes in the Workspace window using a *drag-and-drop* mouse technique. Once in the Workspace window, the user is free to resize these shapes, or their components, to enter text, to interconnect shapes, and so on. See the Workspace window, below, for more detail on the available manipulations.

The Structure window displayed in Figure 28 is a typical palette of sample shapes. It consists of a set of shapes for constructing causal models, and a set of shapes for annotating those causal models with arguments, marking assumptions, stating rebuttals, and identifying conflict. In the palette, shapes are typically displayed at a reduced scale such that the text contents are not wholly visible. We have found that the user can readily identify the desired shape by its basic visual structure in the absence of its textual content. For special symbols, such as maps, the different symbols can be explicitly labeled, e.g., “Map of Arisle”, so that they can be more readily identified. Some of the shapes on the palette have been expanded in Figure 29 for clarity.

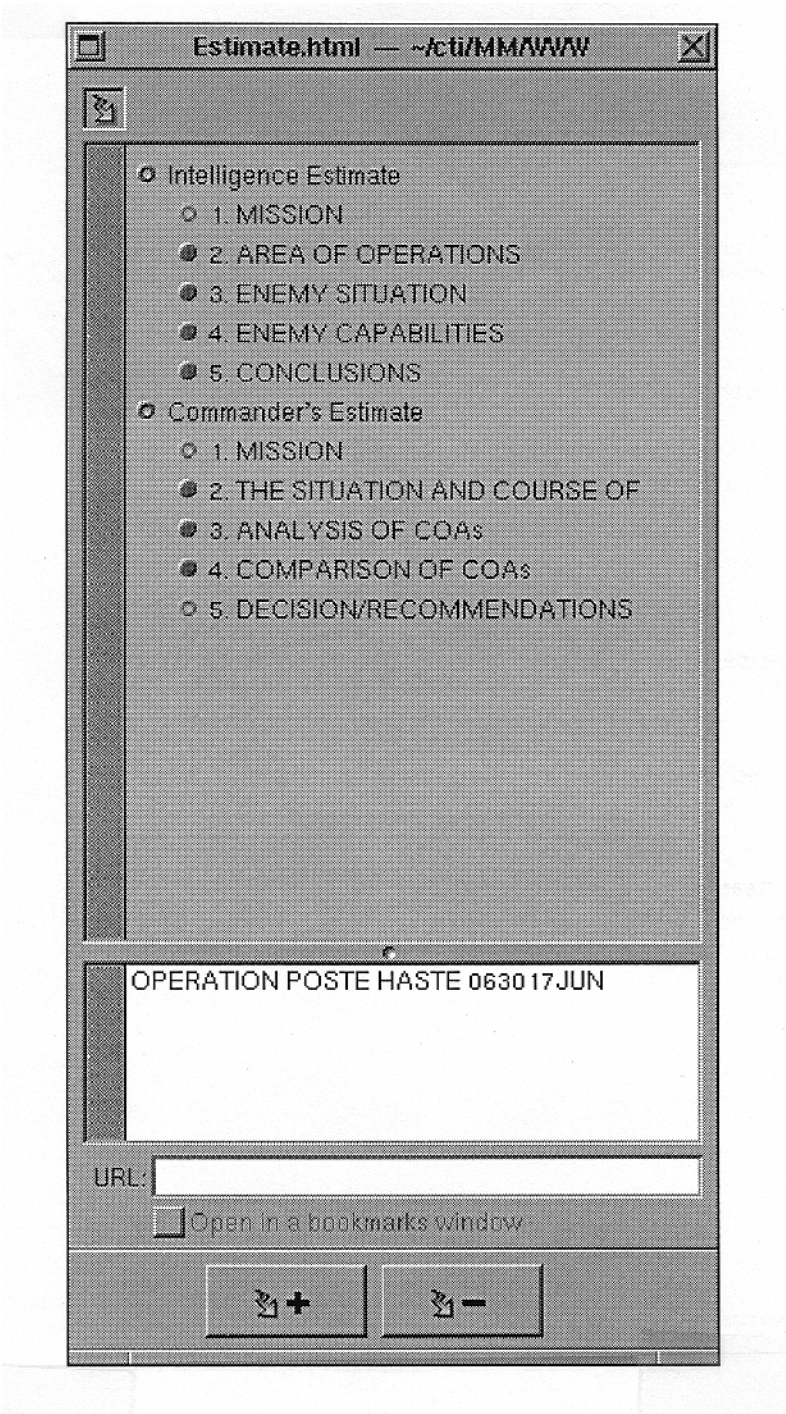

Figure 25. Estimate window.

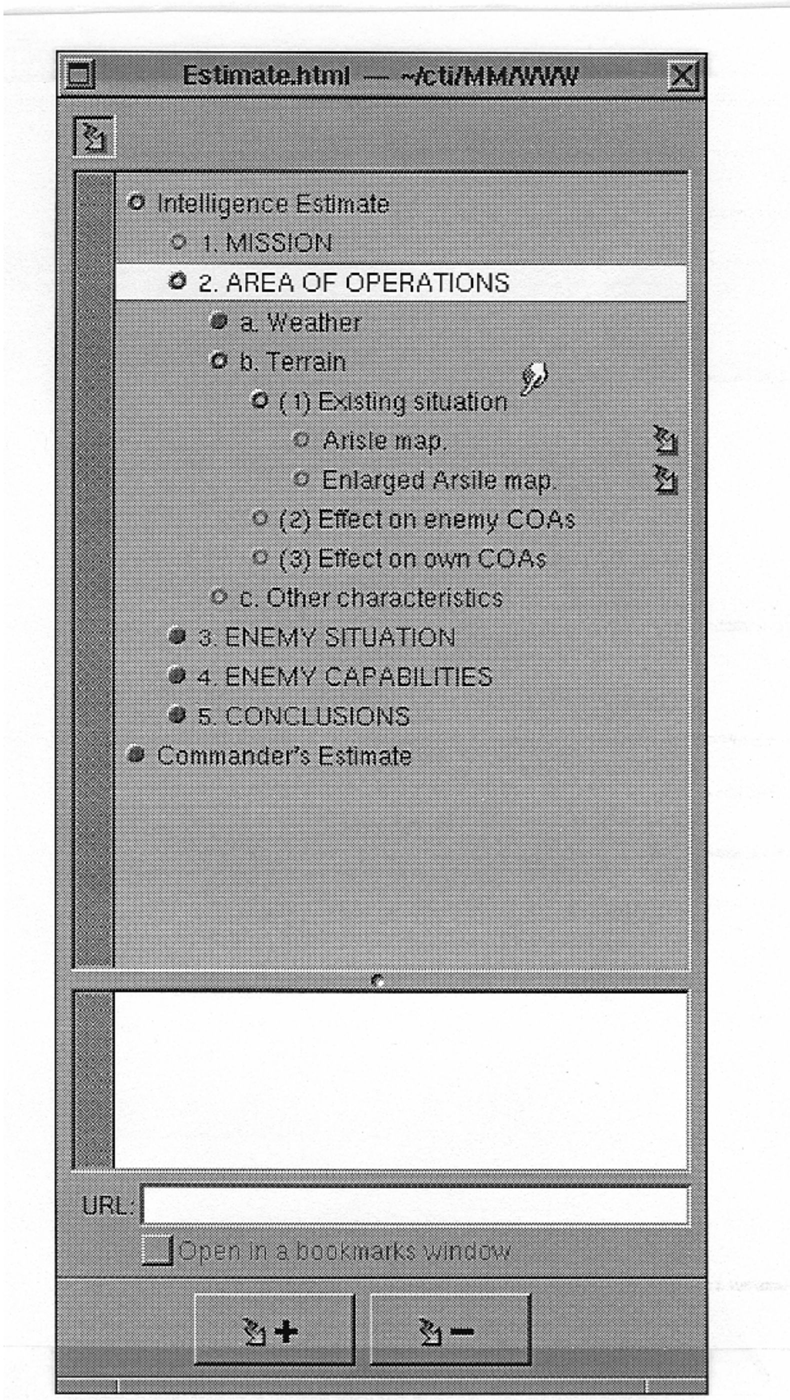

Figure 26. Estimate window, expanded.

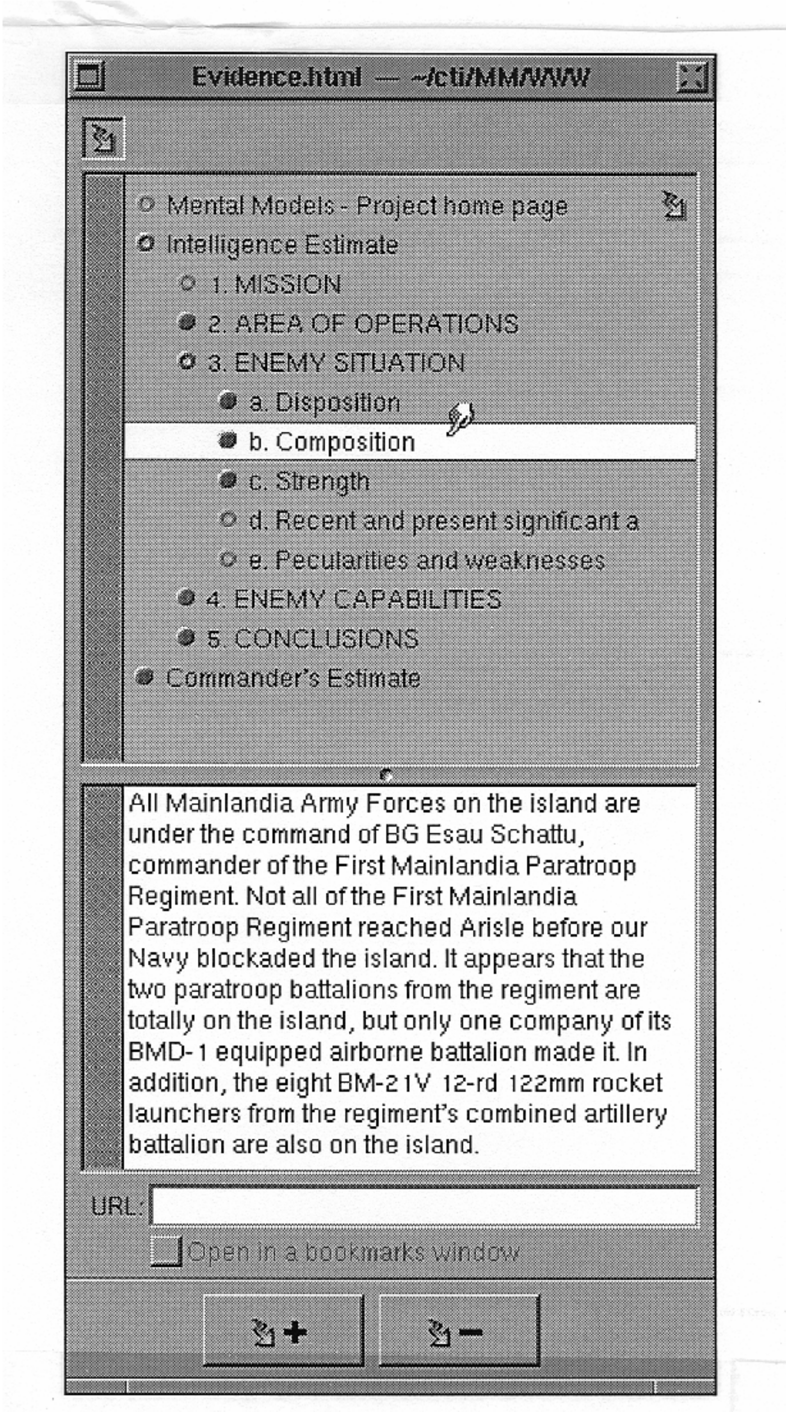

Figure 27. Estimate window, with detailed text.

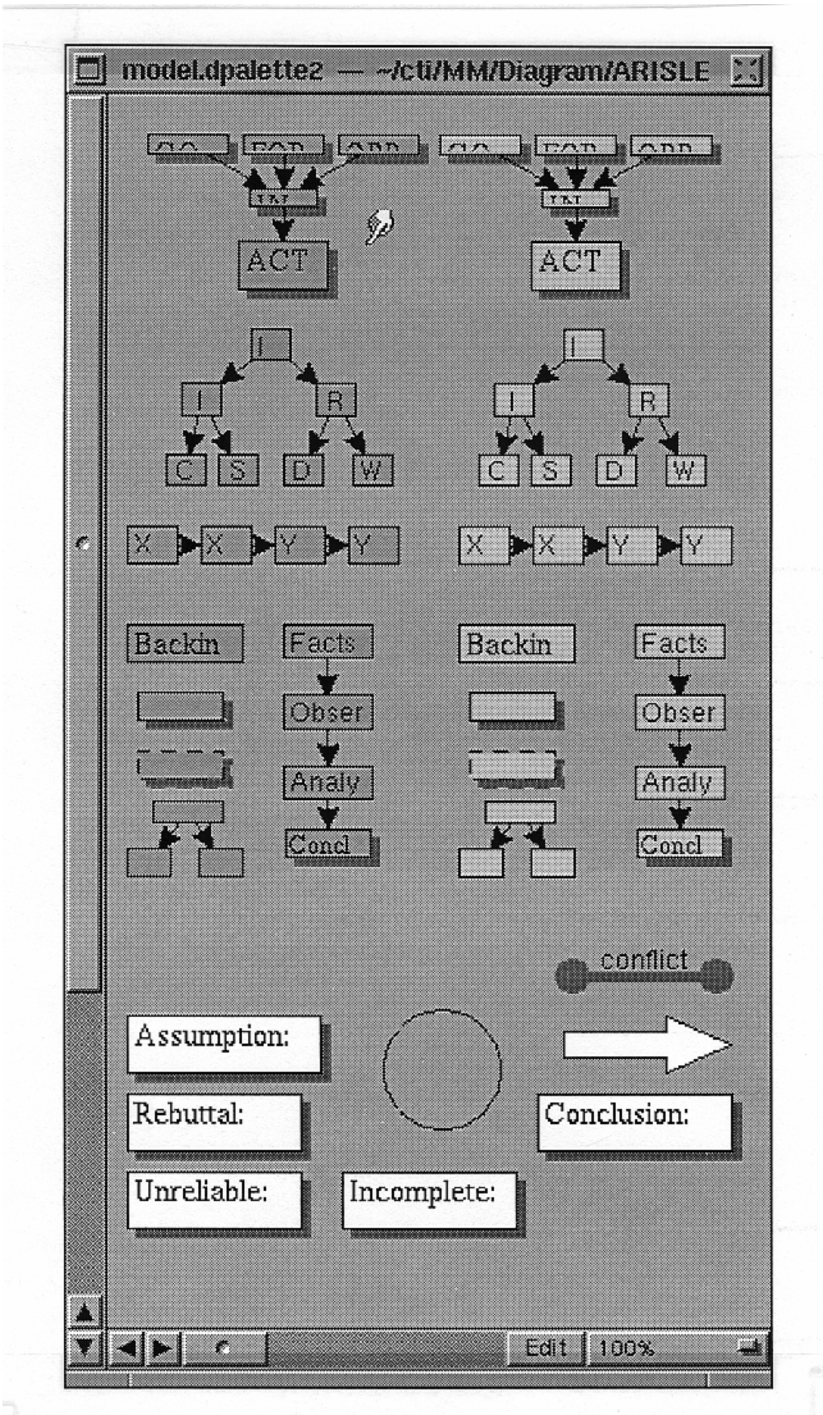

Figure 28. Structure window, with palette of shapes.

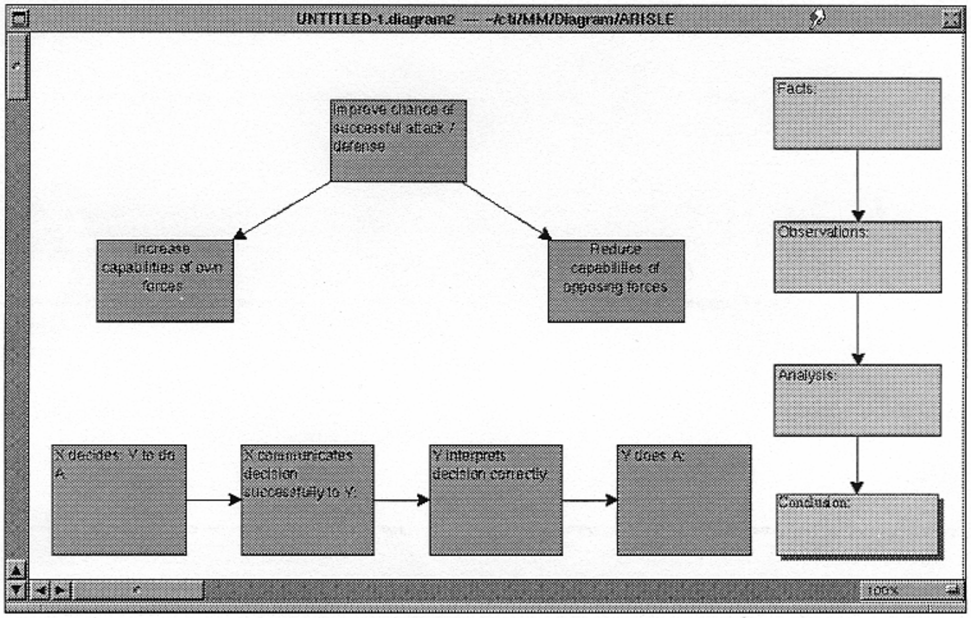

Figure 29. Sample of structures from the palette.

Two sets of shapes are provided for building causal models, one colored red (to represent our understanding of the enemy's thinking) and the other colored blue (to represent our own planning). These are found in parallel columns in the top two-thirds of the palette in Figure 28. The following templates are contained in each column, moving from top to bottom:

Intent. These are basic structures for understanding the enemy's intent, in order to influence it, predict it, or explain surprises. It is also important for friendly planning. The structure contains nodes for *goals, forces, opportunity, intent,* and *activities*. Each of these can be expanded by linking it to other structures, or to backing.

Principles & methods. This structure is useful for fleshing out the activities box in the *intent* structure. It contains a set of goals, subgoals, and actions to achieve them. The template itself is provided as a set of overlapping hierarchical templates, permitting users to determine what level of detail they wish to model. Whichever node the user selects to drag into the Workspace, only that node and the nodes immediately below it actually appear in the Workspace.

Command. This is just one example of a large number of *action execution* structures. They represent the temporal and causal constraints among actions, which require them to be performed in a particular sequence or at particular times. The command structure represents the sequence of events from making a decision, to communicating it, to implementing it.

Evidence interpretation. This structure supports the process of questioning the reliability of a conclusion. It spells out the steps of observation, reporting, and reasoning involved in arriving at the conclusion.

Backing. A backing box is a versatile element that can be linked to any component of any model structure. It represents the information, assumptions, and reasoning underlying the conclusions stated in that component. Thus, information from the Estimate window can be placed in the backing box for the appropriate component. In addition, the backing box can itself contain models of any kind. Thus, it can represent the reasoning underlying a conclusion at a greater level of detail than the top-level model. The backing box for a component can be retrieved in order to review the information and reasoning behind a conclusion in the model.

Generic components. The palette contains unlabeled boxes that can be dragged to the Workspace and linked in anyway the user wishes to existing models, or combined into new models of any kind.

Alternative causes and effects. This box can be dragged into the Workspace and linked anywhere into an existing model, to represent alternative possible causes (e.g., of an enemy action) or alternative possible effects (e.g., of a friendly action). It is distinguished by other causal boxes by use of dashed lines.

Meta-recognitional elements. In the bottom third of the palette are shapes that annotate a causal mental model to highlight problems of uncertainty. These components help the user mark *assumptions* that underlie a conclusion in the model*, rebuttals* or counterarguments to conclusions in the model, *unreliable* assumptions in the model, *incomplete* components in the model, i.e., gaps in the information that has been considered or accounted for, *conflict* between components of a model or between two models, and *conclusions* that are drawn from the model after considering the assumptions, rebuttals, unreliabilities, and conflicts. An arrow shape is provided for marking arguments within the model, i.e., the direction of reasoning, which may be from cause to effect, from effect to cause, or from effect to effect. The circle is actually four separate 90-degree arcs that can be used for grouping common causes or effects by passing the arc through their links. The arc is a traditional notation to represent conjunction, i.e., the requirement that all the causes be present.

Editing capabilities provided with the Structure window permit the creation of new shapes or structures as desired. The Structure window and Workspace can display any arbitrary image. For example, a map could be defined as a symbol on a map palette (in reduced scale). By dragging a particular map symbol into the Workspace, the user could create a scaleable map of that terrain. The system could also represent a palette of different military units and resources. These, in turn, could be used to annotate the map. Symbols for units could be arranged on the terrain using the standard drag-and-drop technique and rearranged freely. Unit movements indicators and synchronization information could be easily added.

The contents of the Structure window are currently static, although the user is free to edit its contents. In a future version, part of the Structure window would be dedicated to contents that are dynamically adapted to the current situation and user. Such dynamic adjustment might include both the mental model structures presently addressed by the system and the maps and units relevant to the commander's current battlespace.

Workspace window. The Workspace window serves as a canvas on which users arrange the components drawn from the Structure palette into mental models that reflect their current situation assessment. Suppose, for example, a user wants to explore the possibility that the enemy has a particular *intent.* The process might begin by dragging a structure for red *intent* into the Workspace window. The user would type the hypothetical intent into the *intent* box. The user would then look for a way to make sense of that intent by constructing a complete story, including *goals*, *forces*, and *opportunity* (as perceived by the enemy) that might lead to that intent, on the one hand, and the observable *activities* that the intent might lead to, on the other. The user thus tries to fill in the various components of the *intent* model, i.e., *goals, forces, opportunity, intent,* and *activities*. The user would draw information from the Estimate window concerning enemy strength, composition, disposition, likely COAs, etc., that is relevant to fleshing out and testing this story.

The full text of this information may be placed in *backing* boxes using a simple *copy-and-paste* operation, while summary statements are placed in the appropriate components of the story structure. *Backing* is marked in the Workspace by a blue diamond placed next to the component for which it provides backing. (These blue diamonds mark hyper-media links to backing boxes, which in turn contain text, models, maps, etc.) The *backing* may then be retrieved at a later time by clicking on the blue diamond.

Alternatively, users might not start with any particular hypothesis about enemy intent in mind. For example, they may be stimulated to wonder about intent by information they discover in the Estimate window. They may then start building an *intent* structure. Users can use that structure as a guide to other information that might be relevant, find it in the Estimate window, and use it to flesh out the *intent* structure. As they build the *intent* structure, they is likely to generate ideas regarding what the enemy is likely to do. Typically, situation understanding will be an iterative mix of these top-down and bottom-up processes. (We shall see how system advisory functions support both kinds of processes at each step, e.g., finding information in the Estimate window that is relevant to a structure, and prompting to build appropriate structures based on current structures and information in the Estimate window.)

In the demonstration scenario, Mainlandia has invaded the tiny island of Arisle; at the same time, an indigenous terrorist group called the Noclas, which is allied with Mainlandia, has seized foreign hostages; and the US is deploying forces to restore the island's independence.[14] In Figure 30, we imagine that the US (blue) commander has built an *intent* structure to try to capture the thinking of the President of Mainlandia. Consideration of *forces* suggests that Mainlandia will not be able militarily to withstand a US invasion; but it has the potential to consolidate the seizure of the island diplomatically at a meeting of a regional alliance called FOCOP. In considering the *opportunity* box of this structure, the blue commander realizes that Mainlandia truly has a window of opportunity: To achieve a diplomatic victory in the time available before US forces can be brought to the island. This insight suggests a hypothesis about enemy *intent*: To delay the allied military until diplomatic victory is achieved at FOCOP. This *intent* in turn leads to a prediction regarding diplomatic and military *activities*.

Figure 31 and Figure 32 show information from the Estimate window that the blue commander has used as *backing* for the *forces* component of the mental model in Figure 30. As these two figures show, when fleshing out a structure, users may find information that does not support the story they are building. These may be incorporated into backing boxes as *rebuttals* or *unreliable* assumptions; they may then be reasoned about, and *conclusions* drawn. (The user may even decide to build an alternative model, which *conflicts* with the original one.) In this example, the blue commander concludes (tentatively) that Mainlandia will not in fact be able to delay the US invasion to any significant degree. Nonetheless, Mainlandia still has a brief period to pursue

---

[14] The scenario was developed and made available to us by Dr. Rex Michel of the Army Research Institute, Fort Leavenworth Field Unit.

diplomatic goals before a US invasion can begin. And they may be relying on the hostages as a means of delaying the US further.

This model of the President's intent has important implications for the US military effort: It places pressure on the blue commander to devise and execute a plan to quickly regain control of Arisle, to forestall diplomatic defeat. The blue commander begins developing a blue *intent* structure in parallel with a model of red *intent*. A mental model of blue *intent*, and some associated backing, is shown in Figure 35 and Figure 36. This is a predictive use of the two *intent* structures: The prediction that the enemy *intent* is to achieve a diplomatic victory provides a time constraint (i.e., an *opportunity*) for the US plan. Figure 35 contains additional backing for the blue *intent* structure: hyper-media links to a map of Arisle (Figure 37) and a COA analysis table (Figure 38).

In Figure 33, the blue commander carries the modeling of red *intent* to a finer level of detail, moving from the President's overall objectives to the *intent* of the military commander of the Mainlandia forces. One of the means at the red commander's disposal (i.e., part of the enemy *forces*) are the hostages. Figure 34 shows the *backing* underlying this red capability. The blue commander has also annotated this *backing* with rebuttals that reflect doubts that Mainlandia can truly control the Noclas or use the hostages effectively. This reasoning, too, had an influence on details of the blue plan, as reflected in the blue *intent* structure. (The blue commander decided to use Special Forces both to attack enemy air defense and to free the hostages.)

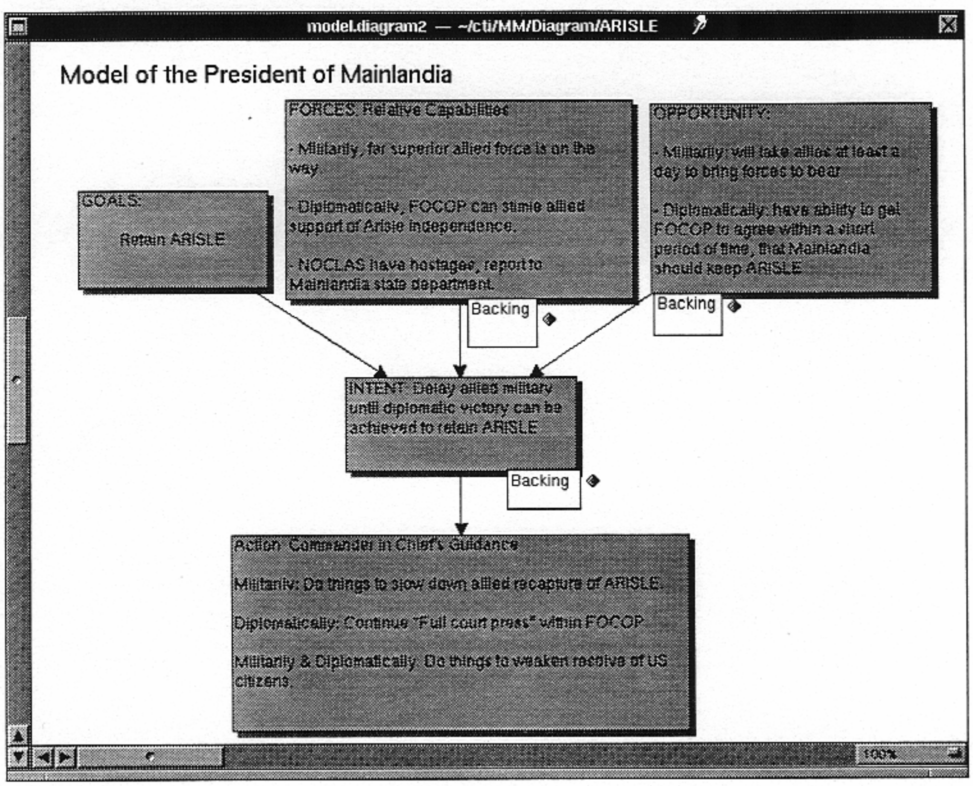

Figure 30. Intent structure representing the possible thinking of the President of Mainlandia.

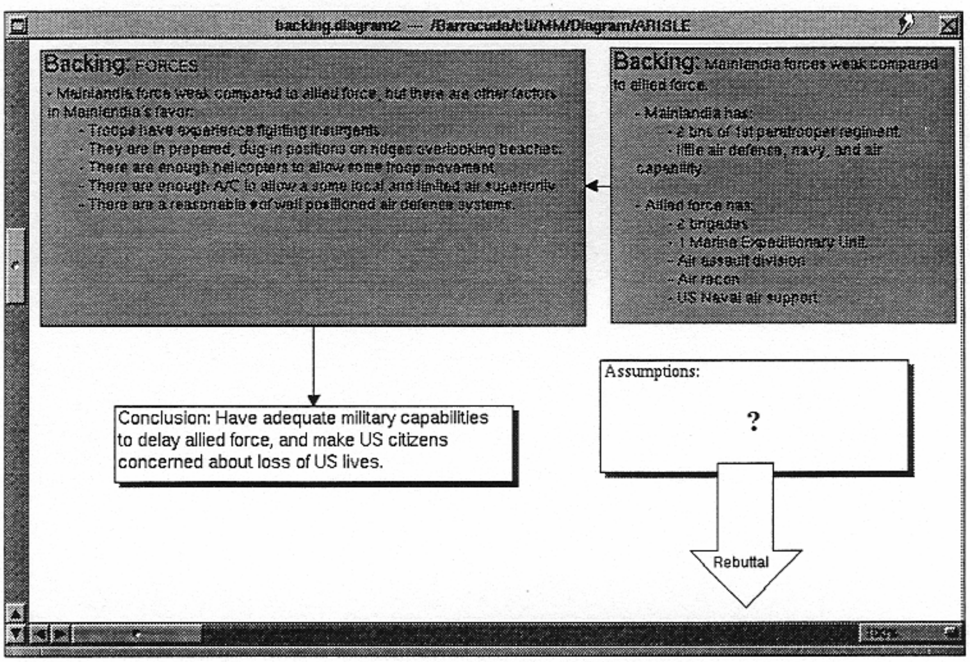

Figure 31. Backing for *forces* component in red model (President of Mainlandia).

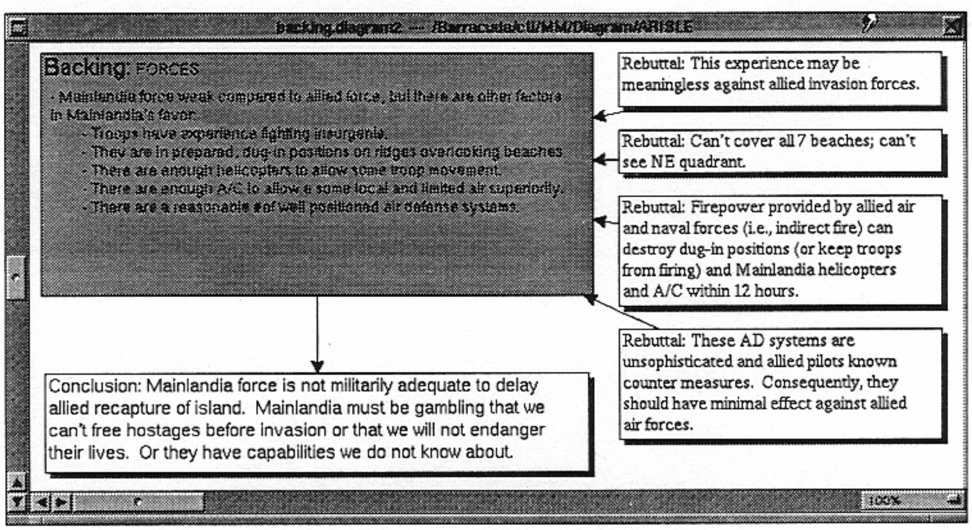

Figure 32. Backing for *forces* component in red model, with rebuttals (President of Mainlandia).

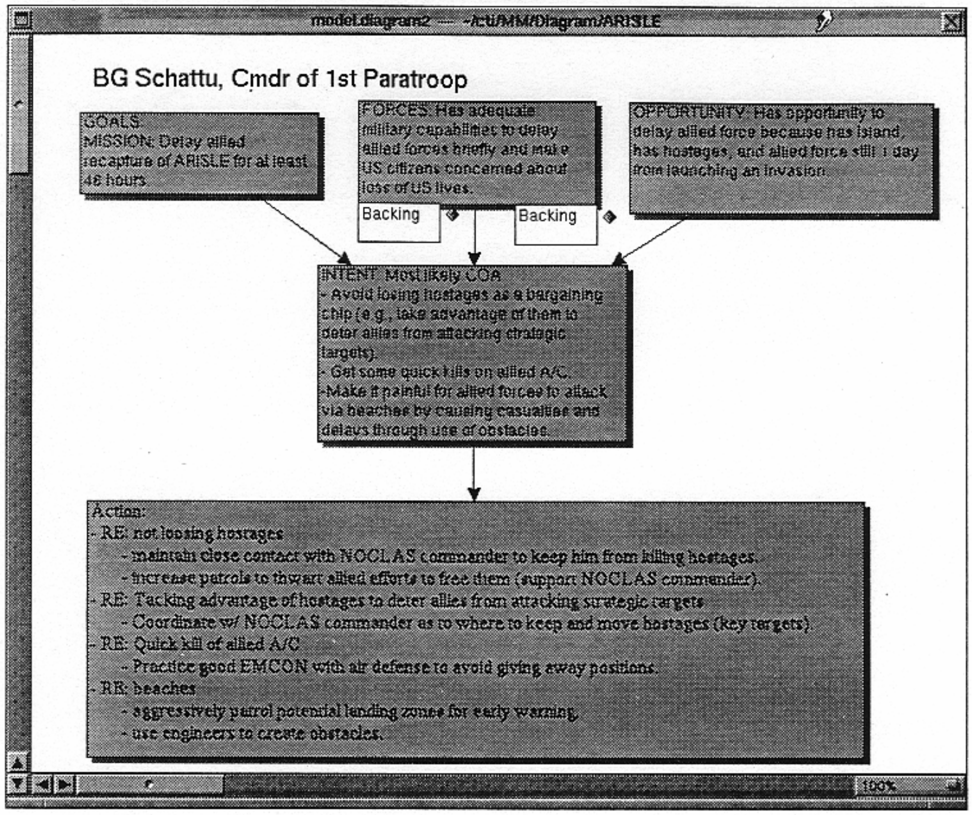

Figure 33. *Intent* model for red commander (General Shattu).

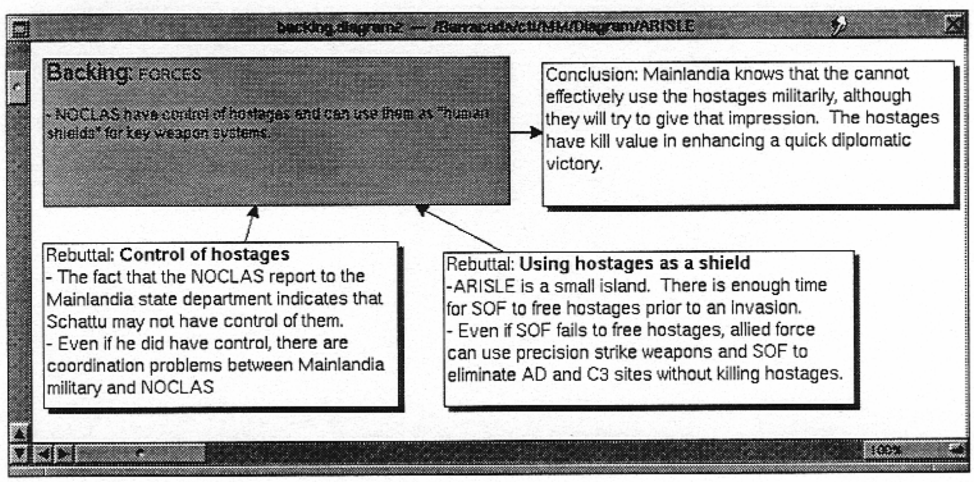

Figure 34. Backing for *forces* component in red model (General Shattu).

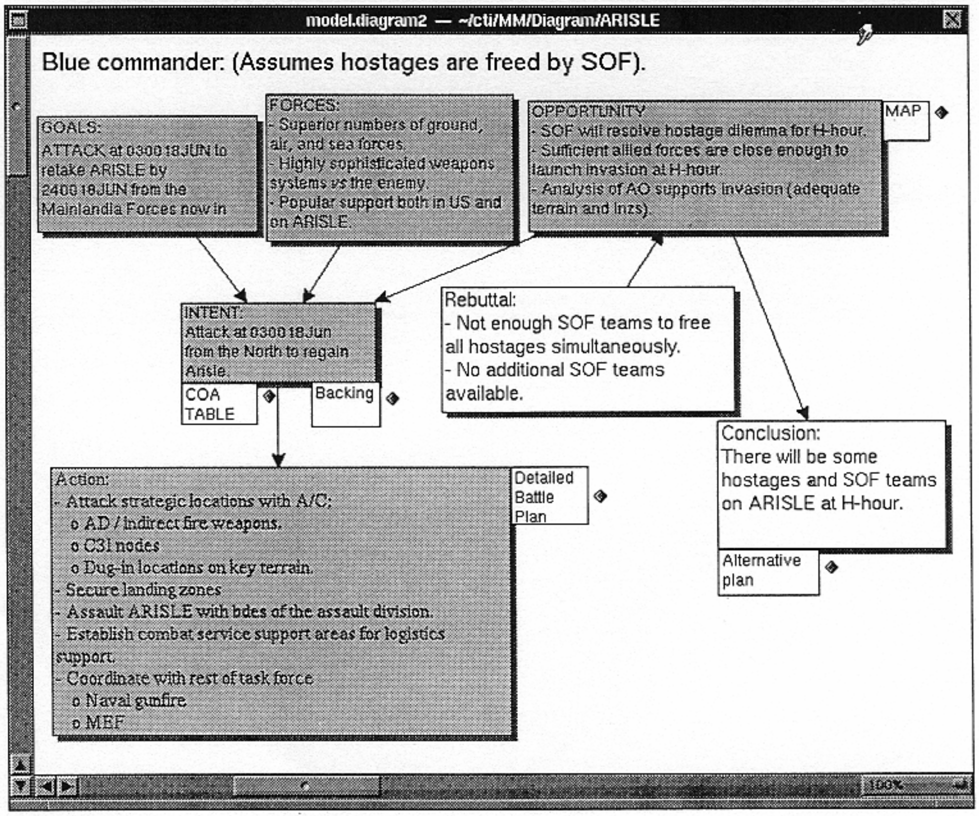

Figure 35. *Intent* model for blue commander.

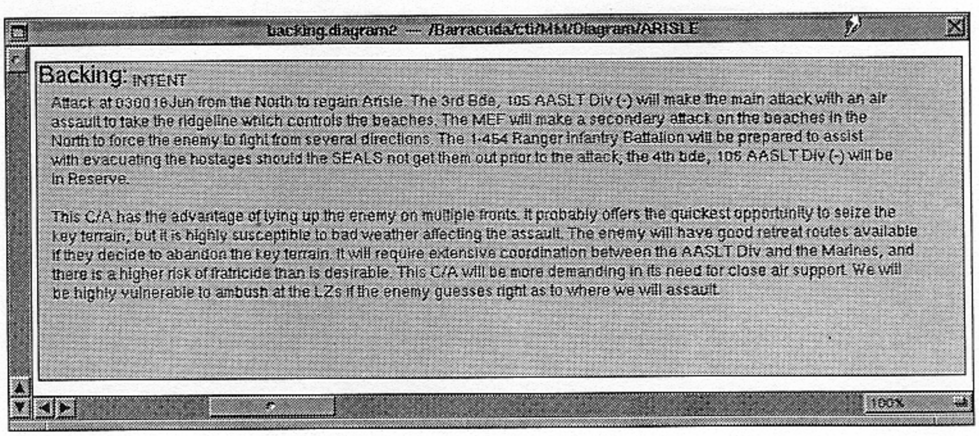

Figure 36. Backing for *intent* component in blue model.

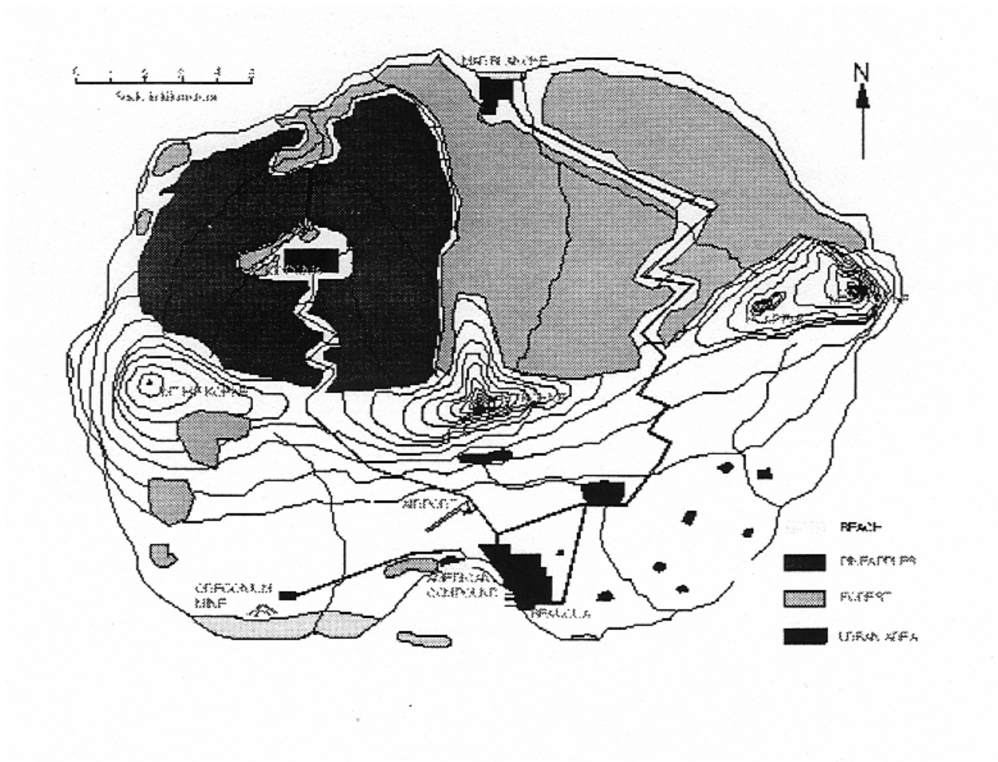

Figure 37. Map of Arisle.

table.rtf — /Barracuda/cti/MM/Documents

Comparison of own courses of action:
Significant Factors

C/A	*Dispositions*	*Weather/ terrain*	*Supporting Attack*	*Obstacles*
C/A 1	Hits main enemy strength from a single direction	Beach assault vulnerable to bad weather	AASLT Div can sspt MEF	Highly vulnerable to obstacles on beach
C/A 2	Hits main enemy strength from multiple directions	Both assaults vulnerable to bad weather	Depends heavily on success of supporting attack	Somewhat vulnerable to obstacles on beach
C/A 3	Hits main enemy strength from multiple directions	Both assaults vulnerable to bad weather	Only needs one of the attacks to succeed	Highly vulnerable to obstacles on beach
FAVORS:	C/A 2,3	C/A 1	C/A 3	C/A 2

Figure 38. Course of action analysis table

Advisory Functions

In this section, we turn to advanced features of the demonstration system, which support the rapid capturing of user mental models. These features are a crucial aspect of the system's capabilities in virtually all its potential applications. They are designed to speed up the process by which a user creates and maintains a model of an evolving situation. The three a*dvisory* functions are:

Recommend backing: This function monitors the content of the models being constructed by the user in the Workspace, and recommends relevant information from the Estimate window. Such information may confirm or disconfirm the model being constructed.

Candidate completions: This function monitors both the abstract form and the content of the models being constructed. It recommends additional structures to elaborate and complete the structures being built, to fill gaps in the overall model of the situation.

Recommend placement: This function monitors the information attended to by the user in the Estimate window and recommends structures and components of structures that might be appropriate for modeling the specified information. [15]

In the demonstration system, these functions are based on the results of the empirical analyses reported above, regarding the mental models that officers utilize and the environmental, personal, and contextual features that predict when they will be used. Later versions of the system would use this as a starting point, but might permit an evolution of the advisory function based on modeling the actual performance of an individual user.

As presently designed, these functions operate only on the specific request of users. That is, in building a structure, they may request recommendations regarding relevant information in the Estimate window, or candidate completions of the structure. In perusing the Estimate window, they may request recommendations regarding appropriate structures in which to place the material they are reading. For some purposes, e.g, training, it may be desirable to have these functions operate in a more automatic fashion. For purposes of diagnosing cognitive skills, it may be desirable to turn the functions off. In a decision aiding context, it may be desirable to allow users to toggle the advisory mode of the system on and off, rather than having to request it on each specific occasion.

These functions provide advice or guidance for tasks that users could perform by themselves. The functions make performance more efficient by intelligently filtering the range of choices, so that the most likely options are most salient. They highlight or otherwise emphasize likely options (i.e., the information most likely to be relevant as *backing*, the structures most likely to be useful in the given circumstances) based on the environment and the structure and content of the mental model being developed. At no time, however, do these functions reduce flexibility. They never remove or even reduce the accessibility of the broader range of options available to the user.

From structure to evidence: Recommend Backing. This advisory function automatically queries the available electronic documents, e.g., the Intelligence Estimate and the

---

[15] Of these three advice functions, the first two were implemented in the Phase I effort, and the third should be a simple extension of the existing mechanisms. The functions relating to episodic memory were not implemented for reasons that are detailed more fully in the Implementation section below.

Commander’s Estimate, based on the content of the structures in the Workspace. In the present implementation, users specify the portion of the current structures that they wish to serve as the basis for the query.

The Recommend Backing window provides a simple interface to sophisticated text analysis algorithms. These algorithms treat the components of the reports in the Estimate window as a collection of many small documents (i.e., one document for each sub-section of each report). The algorithms treat user-designated text in the Workspace (e.g., the summary statement of intent) as a query. These algorithms then search the reports and provide the user with a ordered list of likely sections in each report that contain material bearing on the subject matter of the query.

An example is shown in Figure 39 and Figure 40. The user has filled in the *goals, forces,* and *opportunity* components of the story structure, and in Figure 39 has asked for advice regarding additional relevant material in the Estimate window. The user has specified that the summary material in all three boxes serve as the basis for the query. This is a rather broad query, encompassing material from three different components of the *intent* structure. As a result, the matching sections detected by the system in Figure 40 are also broad in scope. In fact, they pertain directly to the possible *intent* of the enemy, since this is influenced by all three of the designated components. The user can now use the retrieved information as *backing* for the *intent* component of the structure, and to fill in predicted *activities*.

Queries can, of course, be more specific than this. Had the user designated only the summary material in the *forces* box, for example, the user would have received a far more specific selection of information on enemy air defense and on the Noclas terrorists.

From structure to structure: Candidate Completions. This advisory function focuses on the structures currently present in the user's mental model as represented in the Workspace. Using the results from our analysis of the critical incident interviews and problem-solving sessions, this function suggests structural completions or extensions of the user's current mental model. For example, if the user has developed a model of blue *intent*, but has not fleshed out the *activities* components of the model, then the system might recommend structures from the appropriate level of the *principles and methods* hierarchy. If the user has developed a model of red *intent*, but has not linked it to a model of blue *intent*, then the system might suggest that the user model blue *intent* as well.

These suggestions take the form of highlighting that is dynamically generated and applied to appropriate templates in the Structure window (see Figure 41). The user is then free to drag appropriate structures from this dynamic palette into the Workspace. Once in the Workspace, the structures may be elaborated or modified as usual. Connections with existing structures may be made that explore and seek to exploit inter-dependencies, for example, between blue actions and red planning.

In the example of Figure 41, the user has only created a simple high-level model of red *intent*. The advisory function has highlighted two templates in the Structure window representing Candidate Completions. One is that the user flesh out the activities portion of the red intent model with a red *principles & methods* structure. The other is that the user begin a model of blue intent, linked to the red *intent* structure.

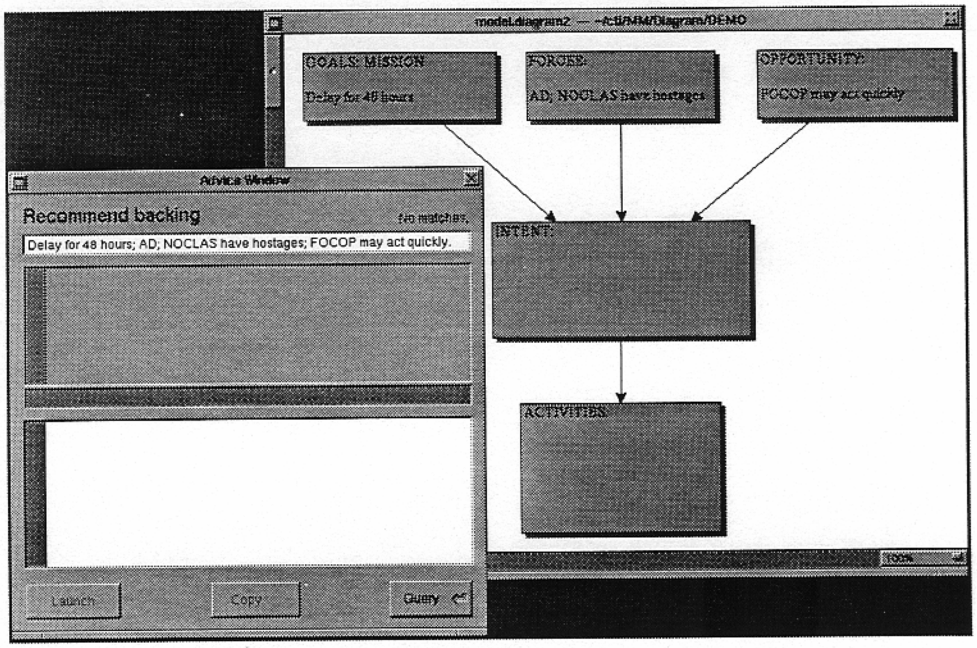

Figure 39. Beginning of Recommend Backing process.

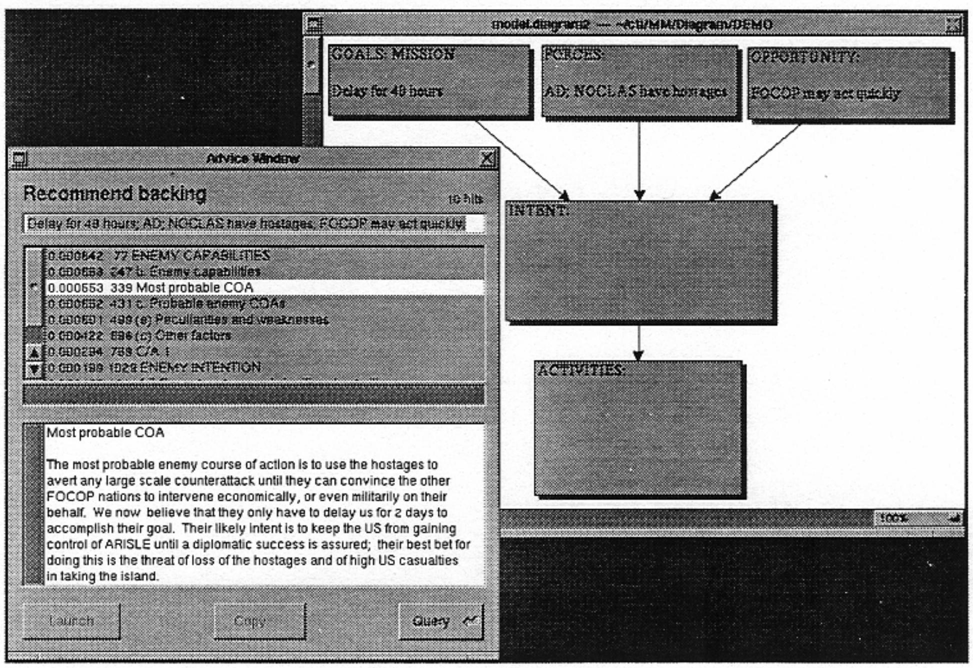

Figure 40. Recommend backing results.

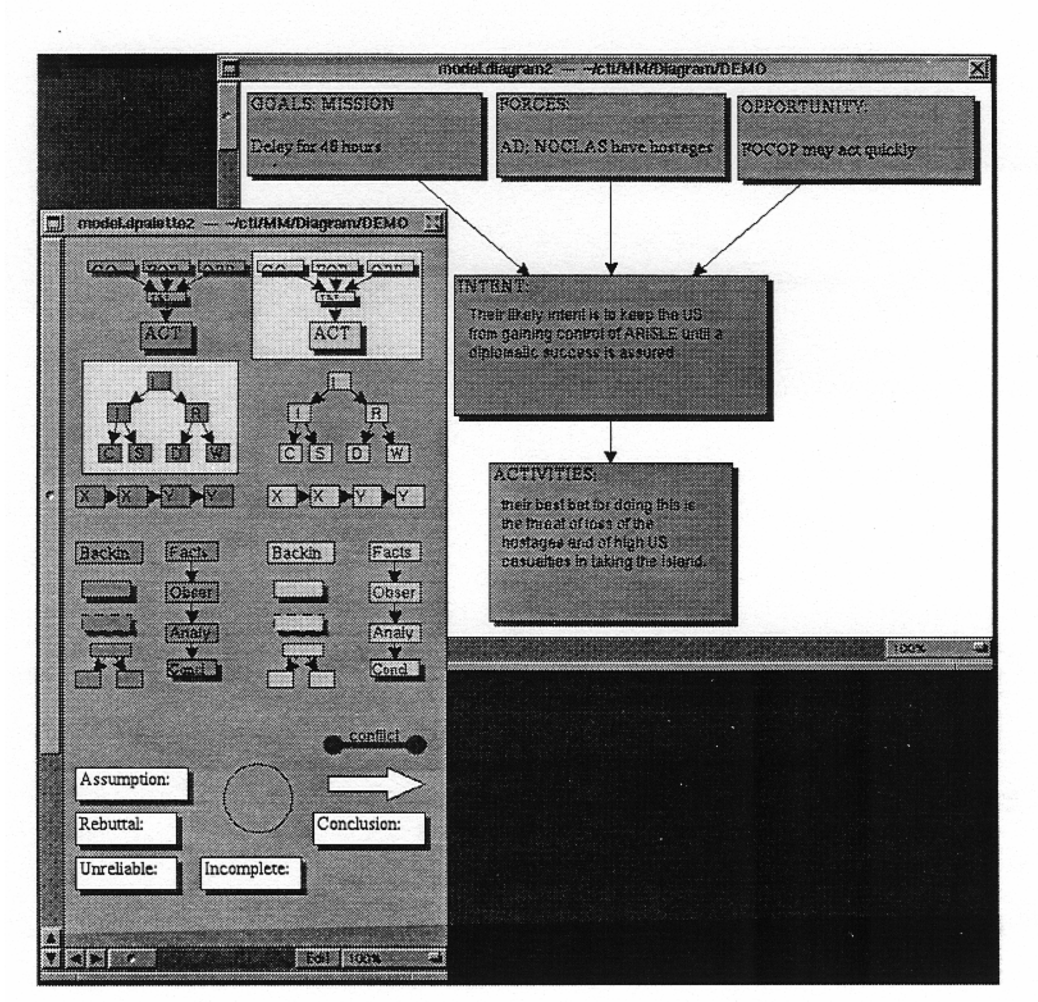

Figure 41. Candidate completions.

From evidence to structure: Recommend Placement. This advisory function would facilitate the rapid generation of mental models by suggesting structures that are relevant to selected text from the Estimate window (or from the Recommend Backing advisory function). When the user designates material in the Estimate window, this function would identify appropriate structural templates (such as *intent* or *principles and methods*) and appropriate components of those templates (such as *forces*). If the structures and components are already present in the Workspace, the advisory function would highlight them. A simplified interface mechanism would then permit direct transfer of the information from the Estimate window into those structural components of the current mental model (or into relevant *backing*). If the appropriate structures are not yet present in the current mental model or if there is no current mental model, the function would highlight appropriate structures and components in the Structure window. These could then be dragged into the Workspace and the material inserted into them.

For example, when working with a mental model of the *Red Air Defense* plan and reviewing the Intelligence Estimate in the Estimate window, the user could use this function to have the system suggest structural placements for such selected information as *Air Defense*

Emplacement, Attached Forces, or *NOCLAS control of hostages*. All three of these components potentially play a role in shaping Red's air defense plan.

Implementation

The Phase I demonstration system represents the integration of a number of systems of software, some commercial, some public domain, and some specially written for the tasks of integrating these components and providing sample algorithms. Custom algorithms were required for the analysis of the structure and content of diagrammatic representations of mental models in support of the advisory functions. A rapid prototyping approach was dictated by the short task performance time and the level of effort. Therefore, numerous rapid iterations of development and review shaped the final form of the demonstration system in combination with the limits of the available off-the-shelf software solutions. Yet, the present design is well suited to extension and expansion. Much of the material developed during this effort would remain in a future version of the system, including the document parsing tools, the integration of the text analysis tools, and the mechanisms for identifying and analyzing the mental models that are represented in Workspace documents.

Hardware, operating system and development environment. Implemented within NeXTSTEP, the system runs on a DEC Pentium-90 machine using a high-resolution color display and 32 MB of memory. The prototype utilizes the commercially available package *Diagram!,* from Lighthouse Designs, and the NeXTSTEP Developer's tools. Non-commercial components include the freely available *OmniWeb*, a World Wide Web HTML browser akin to *Mosaic*, which is also provided by Lighthouse Design on a limited licensing basis. The text analysis and semantic matching tools were developed by Bryan Thompson while at the National Science Foundation (NSF) and are based on both traditional full text analysis methods from information retrieval and on co-term analysis. The latter has been used by both the NSF and the Office of Navy Research for analysis of trends in the development of science and technology. The structural analysis algorithms are implemented in Prolog, using *wamcc*, a Prolog to C compiler and interpreter. The algorithms that monitor for changes in the mental models, which coordinate and integrate these components, and which provide the user interface for the content based advisory mechanisms were developed in the course of the Phase I effort and are written in Objective-C using the NeXTSTEP Developer's toolkit. Various small components are written as *Perl* scripts serving utility roles, such as preprocessing the HTML documents (i.e., the Intelligence Estimate and the Commander's Estimate) for input to the text analysis algorithms and massaging the output of the information retrieval queries.

Software modules and how they interact. While the demonstration system is running, three separate graphical user interface (GUI) applications are operating at the same time. These are: *OmniWeb*, which provides the Estimate window; *Diagram!*, which provides the Structure and Workspace windows; and a specially developed application that provides the advisory functions and the "glue" that integrates these separate components. This glue application is continually monitoring the file system for updates in the saved versions of the mental models that the user is building. Therefore, whenever the user saves a Workspace document, i.e., records the current state of the mental model being developed to the hard disk file system, this application notes the change and initiates a number of activities. First, the changed *Diagram!* file(s) are parsed into an internal representation by a *yacc* grammar. Next the Prolog term database associated with the changed file(s) is re-written to reflect the current contents of the mental models in the Workspace. The next time the user requests, e.g., the Candidate Completions advisory function, the *wamcc* Prolog interpreter is invoked. It analyzes the term database, identifying the locations, content, and interconnections among the shape primitives extracted from the *Diagram!* document, and identifies instances of known configurations of mental models, e.g., red or blue *intent* structures. This information may be used to generate candidate completions or to provide the text analysis routines with detailed context for a Recommend Backing query.

The mental models, as constructed by the user, are "live" diagrams whose structure and content are actively analyzed and interpreted. That content drives the various advisory functions. The principle constraint is that the diagrammatic representations obey a flexible and extensible visual "grammar" for expressing mental models. The basic primitives of this visual grammar include relationships such as *encapsulates, points-to, color-of,* and *title-of*. Specific mental model structures, such as an *intent* structure, are identified on the basis of such primitive relationships and the algorithms for identifying such structures are easily extensible.

How a future system might differ

The most severe limits in the current integration mechanism are: (1) the lack of direct perception by the demonstration system of the user's current selection in the Estimate window and model Workspace; (2) the inability to programatically select items within the Estimate window in response to the computations of the advisory algorithms; and (3) the inability to programatically update currently visible Diagram! documents, such as the Structure window (palette) and the Model workspace -- the user must request such updates manually.

Fortunately, Lighthouse Designs, the makers of *Diagram!* and *OmniWeb*, will be releasing in the fourth quarter of 1995 a software library consisting of the classes and methods that they use in implementing their commercial applications, especially *Diagram!*. Working with this library, it should be relatively easy to alleviate the above mentioned shortcomings of the present system design and advance rapidly towards a more fully featured system that is focused, in a cost effective fashion, on the tasks of the rapid representation and analysis of mental models, as sketched in a functional form during the Phase I design and implementation effort. Such a system would be well positioned for use as both a research tool and as an aid to training or evaluation.

Differences in basic features. The use of a prepared class library, such as mentioned above, will permit a substantially smoother user interface that is tailored specifically to the task of representing and analyzing mental models. Many somewhat awkward features of the present user interface, such as: the need to explicitly resize text bounding boxes in *Diagram!* to insure that all text is visible to the user; the difficulty in using the current hyper-media link creation mechanism for embedding backing into mental models in the Workspace (*Diagram!*); placing links to components of a mental model into the Estimate window (*OmniWeb*); or the visual separation between the Estimate window and the Recommend Backing advisory function window would be easily addressed using this software library. In addition, the library will permit us to readily track the user's actions and develop algorithms that model the user's behavior at a cognitive level. Finally, the NeXTSTEP development environment provides support at many levels for distributed and group computing. These facilities could be readily exploited to support many different kinds of sharing of mental models, including: sharing at different organization levels or between forces; sharing at different places and / or times; and the real-time shared development, critiquing, and correcting of mental models.

Differences in advanced features. It is within the advanced features of the system, i.e., advisory mechanisms and displays of the user interface, that the most change would be directly perceived as a result of the Phase II effort. Due to certain limitations in the ability to programatically integrate *Diagram!* and *OmniWeb* with the other parts of the user interface, as detailed above, these features in the proof-of-concept system are necessarily not quite all we would like or envision. The next version of this system will demonstrate far greater dynamic response to the user's developing mental models, including episodic memory of the user's actions, and provide more finely tuned feedback and guidance in a number of ways.

Modeling and tracking the user. There are a number of very good reasons to want to track and model the user. Analysis of the user's patterns of behavior, at the appropriate levels of granularity, can be an invaluable part of a research tool or training aid. By providing a better basis for diagnosis, and hence improvements in the advisory functions of the user interface, it can facilitate more rapid elicitation of mental models from, and representation of mental models by, the user. In fact, many of the advisory functions rely on an accurate sense of the user's current focus to provide proper context for recommending backing, suggesting structural completions, and placing evidence into the current situation model. It is principally in this manner that the demonstration system is most limited and these facilities would be extended in a Phase II design.

Finally, an estimate of the contents of the user's episodic memory is key to many predictions concerning what structure and content the user will be needing, or returning to, next. Truly dynamic palettes in the Structure window and sensitive and continuous filtering of electronic reports and messages both require a sophisticated sense of the user's current attentional focus and reasonable hypotheses concerning the structures currently active in the user's memory. CTI has modeled the user's cognitive strategies, in the context of research tools, training, and computerized decision aids.

It is important to distinguish this cognitive tracking capability from the more input-output oriented tracking of many other tools. Other tools provide facilities for recording traces of mouse movements and keystrokes, and these traces may be preserved for analysis or used to automate "macros" that remember or facilitate very specific sequences of actions at the level of specific user interface actions. We believe that the modeling of the user must occur at a higher granularity. Cognitive events must be inferred and tracked on the basis of operations by the user, and a model must be developed and maintained of the cognitive processes involved in developing, criticizing, and improving situation assessments.

Additional advisory functions. As noted above, users can link information in the Estimate window to the mental model they are constructing, either by placing the information directly in the structure or by placing it in backing boxes that are linked to components of the structure. Currently, the Recommend Backing function helps users find information that might be relevant to a structure they are building. There are a variety of ways in which this function could be supplemented. For example, the system might mark information in the Estimate window that has not been used in any part of the user's mental model. This information might be graded according to its potential relevance to the user's current concerns, as evidenced by the structures that the user is building. In addition, the system might highlight information in the Estimate window that appears to conflict with the conclusions in the user's current set of models.

Alerts of these kinds could be presented graphically in the form of gauges displayed next to headings in the Estimate window's scrollable outline. Such gauges might display a small horizontal meter depicting the degree of relevance of information in a section of the document that the user has not considered. Color might be used, e.g., bright red, to alert the user when the information is not only relevant, but appears to conflict with current conclusions. Alerts of this kind resemble devices developed for decision aids by CTI in a variety of military contexts (e.g., Cohen, Laskey, & Tolcott, 1987; Cohen, Thompson, & Chinnis, 1985).

A future system could provide more comprehensive integration with Army doctrine and methods. One way to accomplish this would be a graphical representation of the flow of procedures relevant to the current mission phase and task. Users might indicate to the system their current location in this flow simply by selecting it from the graphical representation. The Estimate window and Structure window might dynamically adapt as a function of this information.

An additional set of advanced features pertains to mechanisms for communicating and integrating mental models of physically or organizationally non-proximate users, such as integrating RECON reports with the estimate of the higher level ground component commander. Such issues will be addressed more fully in Task 4.

Differences in implementation methods and resources. The current algorithms for Recommend Backing utilize the textual content of designated mental models components in the Workspace to formulate a query to the Estimate window text. They do not, however, utilize information about the structural components and structural relations within which this textual content appears. By contrast, the algorithms for Candidate Completions do analyze the structures of the mental models.

The algorithms that support Candidate Completions analyze mental models as they are developed by the user in the Workspace. They operate by parsing the *Diagram!* documents as they are modified, maintaining a database of Prolog terms that encodes the simple entities, e.g., *shape, line, point*, and their properties, e.g., *filled, color, arrowhead, dashed,* found in those documents. Working with this term database, the unification and backtracking mechanisms of Prolog are exploited to identify where the user has generated instances of different kinds of mental model representations. Incomplete instances of mental models are identified and used to inform the advisory displays in the user interface. For example, the user might be alerted when issues pertaining to enemy *intent* have not been addressed, such as *opportunity*, or when *activities* have not been spelled out in appropriate detail, e.g., by use of a s *principles & methods* hierarchy,

Mental model structures, however, may be incomplete in content as well as structure. The Recommend Backing function might thus draw on structural as well as textual information. Predictions concerning relevant new material may be drawn through various kinds of contextualized content analysis. For example, the user selecting the summary statement in an *intent* box as the basis for a query, could indicate whether the retrieved material will be placed in a *backing* box or a *rebuttal* box. The current algorithms for content analysis do not exploit structural context, since it is not available to them due to their limited integration with *Diagram!*

and *OmniWeb*. However, these algorithms could be extended to be aware of, and utilize, such context.

Sensitivity to structural context can be complemented by the addition of semantic and pragmatic filtering of the Estimate window material, to insure that the user was only presented with truly relevant material. Such filtering, for example, would be required to support distinctions between searches for *backing* and searches for *rebuttals*. It is practical in the present application due to the relatively small size of the documents, such as the Intelligence Estimate, that are being processed, when compared to traditional information retrieval domains. Further, these documents are already broken up into sections, by the multi-level outline format, that result in a larger number of highly focused sub-documents. This results in increased accuracy in the predictions of relevance and simplifies the task of applying post-retrieval semantic and pragmatic filtering mechanisms.

Just as the Recommend Backing function can be enhanced by adding structural sensitivity, so the Candidate Completions function can be enhanced by adding semantic text analysis. The semantic text analysis algorithms could identify when the content of two structural components in different structures permit their identification. For example, they might determine that the *intent* component in an *intent* structure corresponds to a particular level on the *principles & methods* hierarchy, and suggest fleshing out the *activities* components of the intent structure accordingly. As another example, they might determine when the content of a *forces* component in an enemy *intent* structure corresponds to the *consequence* of an *activity* in a friendly *intent* structure, thus indicating a proactive friendly plan to influence the enemy's intent by affecting the enemy's forces or perception of forces. Advisory functions of this sort would go beyond the current Candidate Completions functionality, which focus on structure alone. By helping users causally link different structures, they would promote the iterative evolution of improved plans. For example, in the Arisle scenario, "Hostages move with AD" was a conclusion reached in an *intent* model of the Red commander. This would support the generation of a new model by the Blue commander in which the SOF forces are used to attack the AD installations (in place of indirect fire weapons) as well as rescue the hostages, resulting in a more effective use of Blue resources.

Shared Mental Models. There are at least three broad senses in which mental models may be shared. They are:

- dissemination of mental models;
- integrating mental models; and
- collaborative development, critiquing, and correcting of mental models.

Dissemination of models. By the dissemination of models we mean simply the delivery of all or part of the situation model to fellow staff members, subordinates, and superiors. The media for the model could include paper and/or electronic publication. This is the simplest concept of sharing, and may amount to as little as enclosing the diagrammatic representations of mental models, for example, of different COAs or conditional responses to possible enemy counteractions, within the five paragraph mission statement developed by the commander and the command staff. However, this concept of sharing can also extend to supplying the different, and in some ways richer, information captured by such mental models to other levels of command, e.g., higher command, or to other armed forces at own and higher levels that are working together in a coordinated action. While many of the key issues in the dissemination of mental models relate to doctrine and communications channels, rather than technical issues, this does lead naturally into the second concern — integrating mental models.

Integrating models. By integrating mental models, we mean combining mental models originated at other levels, e.g., within logistics vs operations, or within other forces, with one's own mental models in such a way as to leverage the richness of the situation estimates developed by people with differing responsibilities and strengths. In the context of a group, this includes by extension integrating models developed by the individuals of the groups as well as those extracted from members of other groups. For example, the commander, through the command staff, may want to review the mental models developed by RECON. At the same time, RECON, close air support missions, etc., may benefit from the sense of a broader perspective depicted in mental models developed by higher command.

To integrate models carries two meanings. Where the models from two team members are in fact non-overlapping components of a larger model (e.g., concepts concerning *capabilities* and *opportunities*, respectively), integration amounts to joining one model to the other, or pasting both into a larger structure. Where team members have generated competing models (e.g., of enemy intent), however, integration involves identifying equivalent elements of the competing models and flagging differences. This is a complex task. However, the mapping of similarities can be performed using search engines similar to those already constructed to locate appropriate backing for a summary statement.

The difficulties involved in integrating mental models at disparate levels of command and function are intriguing. It is technically simple to disseminate the mental models; however that is not enough. We must provide useful tools for exploring mental models developed by other individuals and groups, who often have differing concerns, and help users to exploit those models in their own planning, operations, etc. Some of the tools developed in the Phase I demonstration system might be appropriately extended to this task, including Recommend Backing and Candidate Completions. Consider a close air support mission, which must have detailed knowledge concerning friendly movements and composition in order to increase certainty of target identification and engagement decisions. Aviation officers would benefit from being able to anticipate various contingencies of friendly planning with respect to own and enemy movements and engagements. For example, they might be more proactive in allocating their resources to different sectors of the battle as the situation evolves. Improved understanding of the key uncertainties in the commander's plan and mental model of the enemy could facilitate such proactive strategies.

Collaborative development, critiquing, and correcting of models. By collaborative construction and editing of models, we mean simultaneous work on a single model by two or more individuals. This may mean that the model is built, its arguments critiqued, and search requests and military orders generated concurrently by several members of a group, such as those responsible for logistics vs planning or operations.

Sequential work on the same model and parallel work on different mental models, whether generated by individuals or by different groups, would be supported by the considerations outlined above. Collaborative development of mental models, and collaborative critiquing and correcting, by contrast, are within the realm of computer supported cooperative workgroups. Typical paradigms include visual blackboard models where all users share access to the same Workspace. Often a separate window is provided for a common communications dialog, and sometimes subgroups may be formed with their own dialog windows. Various extensions include: sharing audio or visual contact with other users; the ability to allocate responsibility for modifications made to part of the common model; the ability to replicate part or all of the common model in a private Workspace, and to copy changes from the private Workspace back into the common Workspace; and the ability to perceive at full scale the section of the common or private Workspace on which any user is currently working. Using these various facilities, a large number of different group dynamics may be formed for working on common or related problems.

However, sharing a Workspace does not imply that the users need to share other aspects of the design. Consider the Phase I demonstration system. The users might each be working with different Structure window palettes (at their own choosing) and focusing on elaboration of different aspects of the common model. The tracking and modeling of each user would also proceed separately, and could bring relevant information to each user within the scope of their current cognitive focus. This separate cognitive modeling also opens the door to providing each user with feedback on what cognitive tasks the other users are concerned with. In turn, this might lead to new mechanisms for coordinating group activity and dynamically allocating problems to subgroups.

The development environment used for the Phase I demonstration provides support at a number of levels for each of these kinds of sharing: dissemination, integration, and cooperative workgroups.

CONCLUSION

We have attempted to lay theoretical, empirical, and practical foundations for a rapid mental model capturing system.

From the theoretical point of view, we explored different types of mental models that are utilized to organize situation understanding. Pattern-based or recognitional knowledge involves familiar associations of cues, goals, expectations, and responses. Such knowledge is inadequate by itself in novel or uncertain situations, but it does provide the starting point for more elaborate processing. Recognitional knowledge can lead to the activation of interpretative mental models, which organize events by means of causal or temporal relationships. A variety of meta-recognitional processes are used to critique interpretative models, to discover gaps, unreliability (i.e., alternative possible causes or effects), and conflicting implications of evidence. Meta-recognitional processes activate strategies for correcting such problems when they are found. Such corrective strategies may involve amplifying, elaborating, revising, or rejecting the situation model that is under construction. A final type of mental model structure, the generative, involves causal relationships among variables, and may be used to fill gaps or resolve conflicts in interpretative models.

From the empirical point of view, we subjected a body of data to analysis in order to discover the actual mental model structures and meta-recognitional processes that officers used in a variety of realistic situations. The situations were derived from 23 critical incident interviews and think-aloud problem solving sessions. Five mental model structures were found to be particularly prominent: a model of enemy or friendly *intent*, with components for *goals, opportunity, forces, intent, activities,* and *consequences*; a model of *principles & methods* for accomplishing an attack or defense, with component substructures for enhancing own forces and diminishing enemy forces; a set of *action execution* structures, which depict interdependencies among events in the timing of their execution; a generative *rate of movement* structure, which permits the prediction of speed from variables describing slope, wetness, vegetation, equipment weight, etc.; and an *evidence interpretation* structure, which describes the steps that are required in arriving at a conclusion from an observation. We also discovered a set of meta-recognitional strategies that were consistently used to critique and to improve these mental models: i.e., generating alternative possible causes and effects in order to test the reliability of a model, noticing events that are unexpected in terms of a model, and revising the model to explain the unexpected events. After revising the model, officers evaluate the result and may consider alternative models.

The next step in our empirical analysis was to test whether the occurrence of mental models and meta-recognitional strategies could be predicted. We investigated two environmental variables (mission and unit type), one personal variable (amount of experience in relevant positions), and a set of contextual variables (the specific structures that have already been built to account for the situation). We found a number of significant and plausible associations. For example, models of enemy *intent* were more complete in defensive than in offensive situations. Officers in heavy units were more likely to create proactive models, in which the friendly *intent* was to influence enemy *intent*, than officers in light or specialized units. Heavy units were also more likely to adopt the multi-faceted tactics represented in *principles & methods* structures. Personal experience also led to an increase in the use of *principles & methods* structures. Models of friendly *intent* are typically used predictively, and in conjunction with *rate of movement* models. Enemy *intent* models are typically used proactively, in conjunction with *principles & methods* and models of friendly *intent*.

This research has had a practical fruition: We built a proof-of-concept mental model capturing system, based on both the cognitive theory and on the empirical results. The system provides a palette of pre-built shapes reflecting the model structures that officers consistently use. It also provides a set of advisory functions that facilitate different knowledge structures and processing strategies as a function of the immediate context (and potentially, as a function of the environment and user as well). At the same time, the system remains highly flexible. Users can create any structures they desire, and link them to any content.

The three main components of the system are the Structure window, the Workspace, and the Estimate window. The Structure window contains a palette of general and specialized shapes, as well as meta-recognitional annotation devices to mark different kinds of uncertainty in mental models. The user can drag and drop these shapes into the Workspace, in order to construct a mental model. Shapes in the Workspace can be linked, deleted, moved, modified, sized, and labeled as the user wishes. Model building can both begin and end with familiar types of information in the Estimate window. The Estimate window contains a hyper-media linked representation, in collapsible outline form, of data such as intelligence estimates, the commander's estimate, orders, messages, spot reports, and so on. Users can paste information from the Estimate window into structures being built in the Workspace, or into backing boxes that are linked to such structures. In a completed version of this system, the user would be able to go in the reverse direction as well, e.g., to paste a completed mental model from the Workspace into the commander's estimate, an order, or a message.

Three advisory functions have been designed (and two implemented) for the proof-of-concept system. Two of the functions monitor the Workspace. Recommend Backing tells the user what information from the Estimate window appears relevant for filling gaps in the model that is being built and for testing its validity. Candidate Completions tells the user what other shapes from the Structure window are likely to be useful in elaborating and expanding the current mental model. The third advisory function works in the opposite direction. It monitors the information being perused in the Estimate window and recommends shapes in the Structure window that appear most appropriate for interpreting it.

The system as it now stands is hardly complete. An important additional function would be the ability to monitor a user's performance and dynamically adjust the palette of shapes and the advisory functions to more nearly match that user's knowledge and processing strategies. Many other capabilities can be added to make the system practical as a team aid, in the dissemination, integration, and collaborative development of shared mental models. As it now stands, the mental model capturing system promises to become a versatile tool, with potential applications not only in research, but in evaluation, in training, and as a real-time operational aid for both individuals and teams.

REFERENCES

Abelson, R.P. (1995) Statistics as principled argument. Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J.R. (1982). Acquisition of cognitive skill. Psychological Review, 89(4), 369-406.

Baker, L. (1985). How do we know when we don't understand? Standards for evaluating text comprehension. In D.L. Forrest-Pressley, et al. (Eds.), Metacognition, cognition, and human performance. Vol I: Theoretical perspectives. NY: Academic Press, Inc.

Bainbridge, L. (1992). Mental models in cognitive skill: The example of industrial process operation. In Y. Rogers, A. Rutherford, & P.A. Bibby (Eds.), Models in the mind: Theory, perspective and application. NY: Academic Press.

Berry, D.C., & Broadbent, D.E. (1984). On the relationship between task performance and associated verbalizable knowledge. Quarterly Journal of Experimental Psychology, 36(A), 209-231.

Cannon-Bowers, J.A., Salas, E., & Converse, S. (1991). Shared mental models in expert team decision making. In N.J. Castellan, Jr. (Ed.), Individual and group decision making. Hillsdale, NJ: Lawrence Erlbaum Associates, 221-246.

Chase, W.G., & Simon, H.A. (1973). The mind's eye in chess. In W.G. Chase (Ed.), Visual information processing. NY: Academic Press.

Chi, M., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Chi, M., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R.S. Steinberg (Ed.), Advances in the psychology of human intelligence (Vol. I).Hillsdale, NJ: Lawrence Erlbaum Associates.

Cohen, M.S. (1987, April). Mental models, uncertainty, and in-flight threat responses by air force pilots. Proceedings of the Fourth Biennial Symposium on Aviation Psychology. Columbus, Ohio.

Cohen, M.S. (1989, August). Decision making "biases" and support for assumption-based higher-order reasoning. Proceedings of the Fifth Workshop on Uncertainty and AI. Windsor, Ontario.

Cohen, M.S. (1993a). Metacognitive strategies in support of recognition. Proceedings, Human Factors and Ergonomics Society, 37th Annual Meeting, Seattle, WA.

Cohen, M.S (1993b). The naturalistic basis of decision biases. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsambok (Eds.), Decision making in action: Models and methods. Norwood, NJ: Ablex Publishing Corp, 51-99.

Cohen, M.S. (1993c). The bottom line: naturalistic decision aiding. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsambok (Eds.), Decision making in action: Models and methods. Norwood, NJ: Ablex Publishing Corp, 265-269.

Cohen, M.S., Adelman, L., Tolcott, M.A., Bresnick, T.A., & Marvin, F. (1993). A cognitive framework for battlefield commanders' situation assessment. Arlington, VA: Cognitive Technologies, Inc.

Cohen, M.S., Freeman, J., Wolf, S.P., Millitello, L., & Klein, G.A. (1994). Recognizing, critiquing, and correcting: A cognitive basis for training anti-air warfare decisions. Arlington, VA: Cognitive Technologies, Inc.

Collins, A., Brown, J.S., & Larkin, K.M. (1980). Inference in text understanding. In R.J. Spiro, B.C. Bruce, & W.F. Brewer (Eds.), Theoretical issues in reading comprehension (pp. 385-407). Hillsdale, NJ: Lawrence Erlbaum Associates.

Cruse, D.A. (1977). The pragmatics of lexical specificity. Journal of Linguistics, 13, 153-164.

DeGroot, A.D. (1965/1978). Thought and choice in chess. (2nd Ed.). NY: Mouton.

deKleer, J., & Brown, J.S. (1981). Mental models of physical mechanisms and their acquisition. In J.R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, NJ: Lawrence Erlbaum Associates.

Doyle, J. (1979). A truth maintenance system. Artificial Intelligence, 12(3), 231-262.

Department of the Army (1993, June). FM 100-5 operations. Washington, DC: Headquarters.

Ericsson, K.A., & Simon, H.A. (1993). Protocol analysis. Cambridge, MA: The MIT Press.

Feigenbaum, E.A. (1977). The art of artificial intelligence: Theme and case studies in knowledge engineering. Proc. IJCAI-77, 1014-1029.

Gentner, D. & Stevens, A.L. (1983). Mental models. Hillsdale, NJ: Lawrence Erlbaum Associates.

Greeno, J.G. (1983). Conceptual entities. In D. Gentner & A.L. Stevens (Eds.), Mental models. Hillsdale, NJ: Lawrence Erlbaum Associates, 227-252.

Holyoak, K.J. (1991). Symbolic connectionism: Toward third-generation theories of expertise. In K.A. Ericsson & J. Smith (Eds.), Toward a general theory of expertise. Cambridge: Cambridge University Press.

Johnson-Laird, P.N. (1983). Mental models. Cambridge, MA: Harvard University Press.

Kahneman, Slovic, & Tversky (1982). Judgment under uncertainty: Heuristics and biases. NY: Cambridge University Press.

Keeney, R., & Raiffa, H. (1976). Decisions with multiple objectives: Preferences and value tradeoffs. NY: John Wiley & Sons.

Klein, G.A. (1989). Recognition-primed decisions. In W.B. Rouse (Ed.), Advances in man-machine system research (Vol. 5). Greenwich, CT: JAI Press, 47-92.

Klein, G.A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsambok (Eds.), Decision making in action: Models and methods. Norwood, NJ: Ablex Publishing Corp.

Klein, G.A., Calderwood, R., & MacGregor, D. (1989). Critical decision method for eliciting knowledge. IEEE Systems, Man, and Cybernetics, 19(3), 462-472.

Klein, G.A., Calderwood, R., & Clinton-Cirocco, A. (1986). Rapid decision making on the fire ground. Proceedings of the Human Factors Society 30th Annual Meeting, 1, 576-580.

Koedinger, K.R., & Anderson, J.R. (1993). Reifying implicit planning in geometry: Guidelines for model-based inteligent tutoring system design. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kosslyn, S.M. (1980). Image and mind. Cambridge, MA: Harvard University Press.

Kosslyn, S.M. (1994). Image and brain. Cambridge, MA: The MIT Press.

Larkin, J.H. (1981). Enriching formal knowledge: A model for learning to solve textbook physics problems. In J.R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, NJ: Lawrence Erlbaum Associates.

Larkin, J.H. (1983). The role of problem representation in physics. In D. Gentner & A.L. Stevens (Eds.), Mental models. Hillsdale, NJ: Lawrence Erlbaum Associates.

Larkin, J.H., McDermott, J., Simon, D.P., & Simon, H.A. (1980). Models of competence in solving physics problems. Cognitive Science, 4, 317-348.

Metzler, J., & Shepard, R.N. (1974). Transformational studies of the internal representation of three-dimensional objects. In R. Solso (Ed.), Theories in cognitive psychology: The Loyola symposium. Hillsdale, NJ: Lawrence Erlbaum Associates.

Nisbett, R.E., & Wilson, T.D. (1977). Telling more than we can know: Verbal reports on mental processes. Psychological Review, 84, 231-259.

Noble, D. (1993). A model to support development of situation assessment aids. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsambok (Eds.), Decision making in action: Models and methods. Norwood, NJ: Ablex Publishing Corp.

Patel, V.L., & Groen, G.J. (1991). The general and specific nature of medical expertise: A critical look. In K.A. Ericsson & J. Smith (Eds.), Toward a general theory of expertise. Cambridge: Cambridge University Press

Pennington, N., & Hastie, R. (1988). Explanation-based decision making: Effects of memory structure on judgment. Journal of Experimental Psychology: Learning, Memory, & Cognition. 14(3), 521-533.

Pennington, N., & Hastie, R. (1992). Explaining the evidence: Tests of the story model for juror decision making. Journal of Personality anda Social Psychology. 62, 189-206.

Pennington, N., & Hastie, R. (1993). A theory of explanation-based decision making. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsambok (Eds.), Decision making in action: Models and methods. Norwood, NJ: Ablex Publishing Corp.

Pylyshin, P.W. (1979). Imagery theory: Not mysterious--just wrong. Behavioral and Brain Sciences, 2, 561-563.

Raiffa, H. (1968). Decision analysis: Introductory lectures on choices under uncertainty. Reading, MA: Addison-Wesley.

Rasmussen, J. (1979). On the structure of knowledge--A morphology of mental models in a man-machine context (Report No. M-2192). Roskilde, Denmark: Riso National Laboratory.

Reiter, R. (1980). A logic for default reasoning. Artificial Intelligence, 13, 81-132.

Rogers, Y., Rutherford, A., & Bibby, P.A. (1993). Models in the mind: Theory, persepctive and application. NY: Academic Press.

Rouse, W.B. & Morris, N.M. (1986). On looking into the black box: Prospects and limits in the search for mental models. Psychological Bulletin, 100(3), 349-363.

Rumelhart, D.E., & Norman, D.A. (1985). Representation of knowledge. In A.M. Aitkenhead & J.M. Slack (Eds.), Issues in cognitive modeling. Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, R.C. (1982). Dynamic memory. Cambridge: Cambridge University Press.

Schank, R.C., & Abelson, R.P. (1977). Scripts, plans, goals, and understanding: An inquiry into human knowledge structures. Hillsdale, NJ: Lawrence Erlbaum Associates.

Serfaty, D. (1993, February 4-5). Hypotheses and recent findings on command decision-making expertise. Presentation at U.S. Army Research Institute Workshop on Developing Expertise in Command Decision-Making, Fort Leavenworth, Kansas.

Shanteau, J. (1992). The psychology of experts: An alternative view. In G. Wright & F. Bolger (Eds.), Expertise and decision support. NY: Plenum Press, 11-25.

Shepard, R.N. (1975). Form, formation, and transformation of internal representations. In R. Solso (Ed.), Information processing and cognition: The Loyola symposium. Hillsdale, NJ: Lawrence Erlbaum Associates, 87-122.

Schoenfeld, A.H., & Herrmann, D.J. (1982). Problem perception and knowledge structure in expert and novice mathematical problems solvers. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8, 484-494.

Voss, J.F., Wolf, C.R., Lawrence, J.A., & Engle, R.A. (1991). From representation to decision: An analysis of problem solving in international relations. In R.J. Sternberg & P.A. French (Eds.), Complex problem solving: Principles and mechanisms. Hillsdale, NJ: Lawrence Erlbaum Associates.

Weiser, M., & Shertz, J. (1983). Programming problem representation in novice and expert programmers. International Journal of Man-Machine Studies, 19, 391-398.

Young (1983). Surrogates and mappings: Two kinds of conceptual models for interactive devices. In J.R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, NJ: Lawrence Erlbaum Associates.