Integrated Critical Thinking Training and Decision Support for Tactical Anti-Air Warfare[1]

Marvin S. Cohen, Ph.D., Jared T. Freeman, Ph.D., and Bryan B. Thompson
Cognitive Technologies, Inc.
Arlington, VA 22207
703-524-4331
fax: 703-527-6417
email: mcohen@compuserve.com

Abstract

There is an opportunity to explore synergies between training and decision support in the Tactical Decision Making Under Stress (TADMUS) program. Decision making skill appears to involve a combination of recognitional and metacognitive elements. A model of decision making based on this principle has been developed, and training has been designed based on the model. The training has four components: a process for testing an assessment of a track's intent by constructing, verifying, and improving stories; a specific kind of story for intent to attack; methods for finding the assumptions underlying a story; and strategies for deciding when to examine an assessment critically and when to act immediately. Concepts for a decision support system (DSS) have been developed that facilitate these aspects of critical thinking. The DSS design provides a causal, story-based organization of information about a track, and it supports both recognitional and metacognitive strategies of decision making. Pattern matching is facilitated by color coding of supporting and conflicting cues, as well as gaps in information and expectations based on an assessment. Story building, or assumption-based reasoning, is supported by enabling users to adopt and evaluate assumptions that explain conflicting evidence and lead to a more coherent story. Experimental research will soon be underway to test the value of the DSS as a tool in critical thinking training, and the contribution of critical thinking training to performance with the DSS.

1. A Potential Synergy between Training and Decision Support

The Tactical Decision Making under Stress (TADMUS) program has involved a dual focus: training to improve decision making and team skills, on the one hand, and computer-based decision support, on the other. Although developments in these two areas have paralleled one another, up to now important potential synergies have not been exploited: First, training in relevant decision making skills might facilitate more effective use of the decision support system. Second, the decision support system might enhance the effectiveness of training, by providing a graphical interactive display for instruction, practice, and feedback. Thirdly, the importance of the relevant decision making skills might be confirmed in both training and decision support contexts. Finally, we might gain more general insight into the ways that training and decision support interact. For example, we may learn more about the features of training methods and decision support systems that promote synergy, and about how a common cognitive model can be used to guide the development of each. In effect, we might test the idea that

[1] This work was supported by Contract No. N61339-96-R-0046 with the Naval Air Warfare Center Training System Division.

training and decision support should in general be developed in a coordinated fashion.

An opportunity to implement such coordinated development and testing is now being pursued. This opportunity is based on recent research in modeling and training critical thinking strategies conducted by Cognitive Technologies, Inc., (CTI) with its subcontractor Klein Associates and sponsored by the Naval Air Warfare Center Training Systems Division (Orlando), and on the development of a prototype decision support system by Pacific Science & Engineering Group under sponsorship of the Naval Research and Development Center (San Diego).

In this report, we will describe work done at CTI on a model of critical thinking, a training strategy based on the model, and concepts for a decision support system that might facilitate critical thinking and critical thinking training.

2. A Cognitive Model of Critical Thinking

Proficient decision makers are *recognitionally skilled*: that is, they are able to recognize a large number of situations as familiar and to retrieve an appropriate response. Recent research in tactical decision making suggests that proficient decision makers are also *meta-recognitionally skilled* (Cohen, Freeman, & Wolf, 1996). In novel situations where no familiar pattern fits, proficient decision makers supplement recognition with processes that verify its results and correct problems.

Based on critical incident interviews with active-duty naval officers, we developed a framework for decision making, called the Recognition/Metacognition (R/M) model (Cohen et al., 1996). The model describes a set of critical thinking strategies that supplement recognitional processes. Structured situation models (i.e., *schemas*), often in the form of *stories* about enemy intent, causally organize information about a situation and provide a basis for metarecognitional processes. Metarecognitional processes determine when it is worthwhile to think more about a problem; identify evidence-conclusion relationships within the story; critique the story for incompleteness, conflict, and unreliability; and attempt to improve it, by collecting or retrieving new information and revising assumptions. At a somewhat more detailed level, meta-recognitional processes include:

1. Identification of evidence-conclusion relationships (or *arguments*) within the evolving situation model and plan. This is simply an implicit or explicit awareness that cue A was *observed* on this occasion, while intent to attack along with expectations of observing cue B were *inferred*. On some other occasion cue B might be observed and cue A inferred.

2. Processes of *critiquing* that identify problems in the arguments that support a conclusion (e.g., hostile intent) within the situation model or plan. Critiquing can result in the discovery of three kinds of problems: *incompleteness, unreliability*, or *conflict*. An argument is incomplete if it does not provide support either for or against a conclusion of interest (e.g., the kinematics of the track suggest only that it is a military aircraft, but say nothing about hostile intent; this conclusion is too general for deciding whether to engage). Two arguments conflict with one another if they provide support both for and against a conclusion, respectively (e.g., the heading of a track toward own ship suggests hostile intent, while its slow speed argues for routine patrol). Finally, an argument is unreliable if it provides support for, but not against, a conclusion, but the support depends on unexamined assumptions. Unreliable support may shift or vanish when its premises are further considered.

3. Processes of *correcting* that respond to these problems. Correcting steps may instigate external action, such as collecting additional data, and two kinds of internal actions, attention shifting and assumption revision, that regulate the operation of the recognitional system. Shifting the focus of attention stimulates retrieval of new, potentially relevant information in long-term memory and brings additional arguments into view for meta-recognitional critiquing. Adding or

dropping assumptions permits what-if reasoning, queries for alternative causes and effects, and adoption a single coherent model or plan. These processes in combination help to fill gaps in models or plans, resolve conflict among arguments, and search for more reliable arguments.

4. A higher-level process, called the *quick test*, which controls critiquing and correcting. Metarecognitional strategies, like other actions, are shaped in part by past experiences of success and failure. Metarecognitional processing occurs when the benefits associated with critical thinking outweigh the costs. This is likely to be the case when the costs of delay are acceptable (i.e., time is available for critical thinking), the situation is uncertain or novel (i.e., recognitional conclusions are subject to improvement), and the costs of an error in acting on the current recognitional conclusion are high. The quick test considers these three factors and, if conditions are appropriate, inhibits recognition-based responding and interposes a process of critical thinking. When these conditions are not satisfied, the quick test allows immediate action based on the current best response.

Figure 1 summarizes the relationships among these processes. It highlights the functional distinction between recognitional processes (at the top of the figure) and metacognitive ones (the shaded boxes). The recognitional level provides information to the metacognitive level, while the metacognitive level exerts control over the recognitional level. In the R/M model, metacognition monitors the recognitional processing, maintains a model or description of it (i.e., identifies arguments and problems of incompleteness, conflict, and unreliability), and modifies recognitional activity by inhibiting overt action, shifting attention, and adopting or dropping assumptions. These functional differences may or may not correspond to structural or physiological ones (see Nelson & Narens, 1994). A more detailed description of the R/M model may be found in Cohen et al. (1996).

The R / M model 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 (such as 70% chance of hostile intent, 30% chance non-hostile). Uncertainty is represented explicitly at the metacognitive level, by "annotating" the situation model or plan to highlight points of incompleteness, conflict, and unreliability. In response to specific problems of this kind, metacognitive strategies try to improve the current situation model and plan or find better ones.

Metarecognitional processing is highly dynamic and iterative. The next processing step is determined locally by the results of earlier steps, rather than by a global, fixed procedure (as in Bayesian inference). Correcting for one problem may sometimes (but not always) lead to identification and correction of another problem. For example, a gap in an argument may be filled by collecting further data or remembering previously known information, or, if these fail, by making assumptions. The resulting more specific argument may then turn out to conflict with other arguments. Such conflict may then be addressed by looking for unreliability in one of the conflicting arguments. In doing so, metarecognitional processing might shift focus from the conclusion to the grounds of the argument. This may result in retrieval of previous experiences with the source of the information that is the grounds for the conflicting argument. Such experiences may suggest that the source is not to be trusted. The conflict, which arose because of the implicit, or unconsidered, assumption that this source was accurate, is now resolved. (Alternatively, what if no relevant information were retrieved about the source? A new cycle of critiquing would identify this gap in knowledge, and it might be corrected, for example, by adopting the explicit assumption that the unfamiliar source is not trustworthy. Conflict would be eliminated, but the story now depends on the potentially unreliable assumption about this source. Attention might now be shifted to the

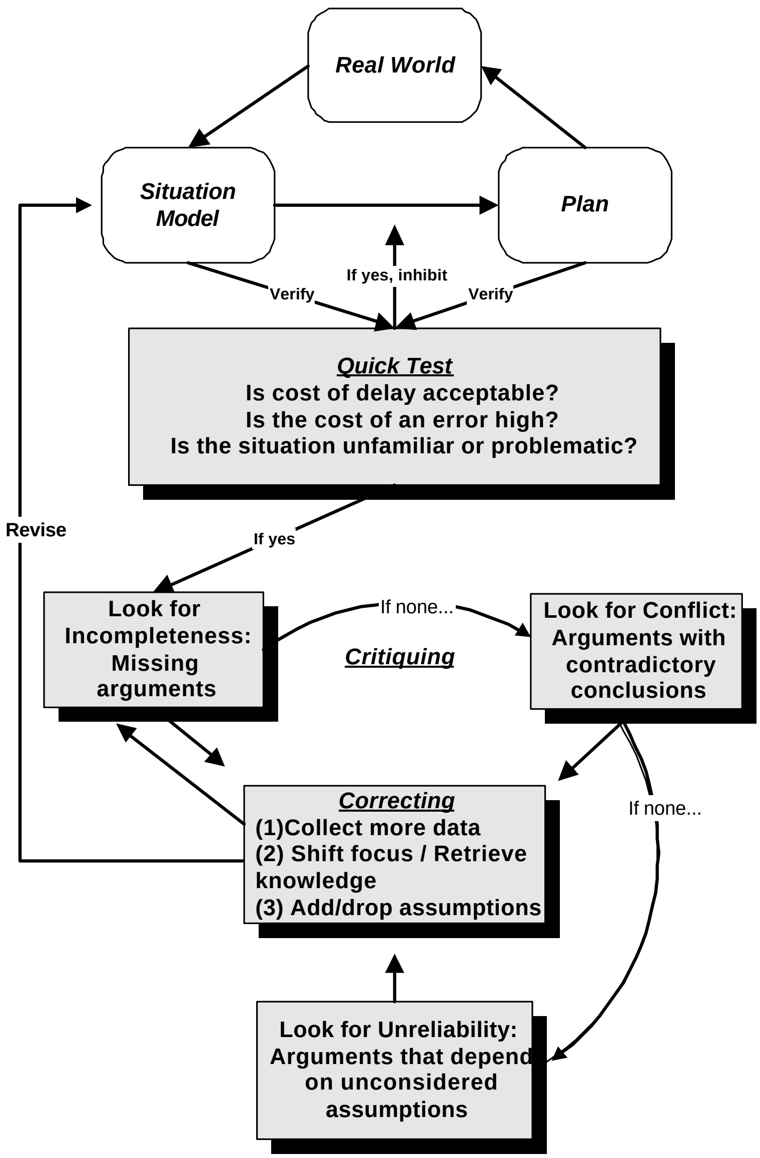

Figure 1. Components of the Recognition / Metacognition Model.

other conflicting arguments.) This process stops when the quick test indicates that the benefits of further metarecognitional actions are likely to be outweighed by the risks of delay, and that action o on the basis of the current best model or plan is called for. The output is a coherent, consistent model or plan together with an understanding of its strengths and weaknesses.

3. Critical Thinking Training

Training based on the Recognition / Metacognition model has been developed, focusing on the decision of whether or not to engage an approaching air or sea contact whose intent is unknown, under conditions of undeclared hostility. The training is based on interviews with active-duty Naval officers, in which they described experiences of this kind in the Persian Gulf, the Gulf of Sidra, and elsewhere (Kaempf, Klein, Thordsen, & Wolf, 1996). Many aspects of the training are based on differences in the way that more and less experienced officers handled similar situations.

We call the training *critical thinking* because it is designed for situations where familiar patterns or rules do not fit. For example, some features of the situation may match the standard hostile intent pattern (e.g., an aircraft turning toward own ship from a hostile country), but others do not (e.g., its speed is slower than expected) and may even match parts of another pattern (e.g., commercial airliner). The training is divided into four segments: (1) An overview of the cycle of creating, testing, and evaluating stories to improve situation understanding; (2) a particular kind of story based on hostile intent; (3) strategies for finding and correcting problems with stories; and (4) guidelines for when critical thinking is appropriate and when immediate action is necessary. In each of these segments, officers listen to a brief verbal presentation of the concepts, followed by questions and discussion. They then participate in interactive scenario-based exercises designed to provide practice in the relevant skill. Feedback during these exercises is provided by group discussion and by the instructor. The basic concepts of the four training segments are as follows:

Creating, testing, and evaluating stories. This section provides an overview of the critical thinking process, called *STEP*. When an assessment is uncertain, decision makers can take it seriously by constructing a *S*tory around it. The story includes the past and future events that would be expected if the assessment were true. Decision makers use the story to *T*est the assessment, by comparing expectations to what is known or observed. When evidence appears to conflict with the assessment, they try to patch up the story by explaining the evidence. They then *E*valuate the result; if the patched up story involves too many unreliable assumptions, they generate alternative assessments and begin the cycle again. In the meantime, they *P*lan against the possibility that their current best story is wrong. Figure 2 summarizes the *STEP* process.

Hostile-intent stories. Stories contain certain typical components. Knowledge of these components can help decision makers notice and fill gaps in the stories they construct. A particularly important kind of story is built around the assessment of hostile intent. For example, a complete hostile intent story explains why an attack is taking place against a particular target by a particular platform, and accounts for how that platform has localized the target and is arriving at a position suitable for engaging it (see Figure 3). The training teaches officers by practice and example how to discover story components and to let them guide decision makers to relevant evidence regarding assessments of intent.

Critiquing stories. After a story is constructed, decision makers step back and evaluate its plausibility. This segment of the training introduces a devil's advocate technique for uncovering hidden assumptions in a story and generating alternative interpretations of the evidence. An infallible crystal ball persistently tells the decision maker that the current assessment is wrong, despite the evidence that appears to support it, and asks for an explanation of that evidence. Regardless of how confident decision makers are in their assessments, this technique can successfully alert them to significant alternatives. It can also help them see how conflicting data could fit into a story. In each case, the technique helps decision makers expose and evaluate assumptions underlying their reading of the evidence.

When to think more. Critical thinking is not always appropriate. Unless three conditions are satisfied, the decision maker should probably act immediately: (1) The risk of delay must be acceptable. (2) The cost of an error if one acts immediately must be high. And (3) the situation must be non-routine or problematic in some way. Training focuses on the way experienced decision makers apply these criteria. For example, they tend to utilize more precise estimates of available time based on the particularities of the situation, a longer-term outlook in estimating the cost of an error, and greater sensitivity to the mismatch between the situation and any familiar pattern.

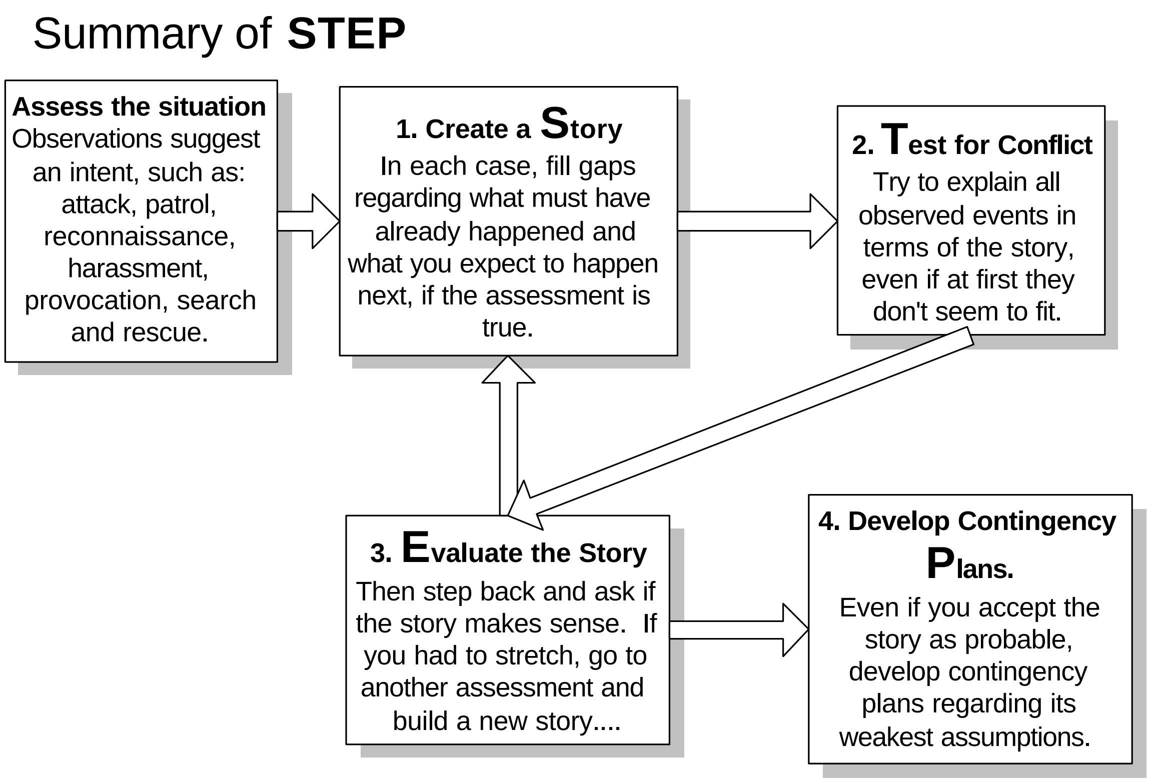

Figure 2. A cycle of four steps for critical thinking.

Critical thinking training has now been tested with active-duty officers at two Navy training facilities (Cohen, Freeman, and Thompson, in press). The evaluation examined the effects of training on critical thinking processes, its effects on performance (i.e., assessments and actions), and participants' own evaluations of the training. The results were encouraging. For example, in one study training increased the number of factors officers considered in assessing the intent of a track by 30%, increased the amount of conflicting evidence they noticed by 58%, increased the number of assumptions they identified underlying the interpretation of that evidence by 27%, and increased the number of alternative assessments they generated by 41%. Critical thinking training can also improve the accuracy of assessments. Agreement with a subject matter expert increased significantly in two out of four test scenarios in the two studies, by 79% and 35%, respectively. At the same time, the training (non-significantly) increased officers' confidence in their assessments in both studies. Subjective evaluations of the training were generally positive.

These tests strongly suggest that meta-recognitional skills can be taught effectively, that officers will use them in relatively realistic tactical situations, and that enhanced meta-recognitional skills will lead to improved performance.

4. Decision Support Concepts for Critical Thinking

Decision aiding should not force proficient decision makers to adopt a radically new approach to a problem. Instead, it should facilitate the strategies preferred by experienced decision makers while guarding against potential pitfalls or problems with those strategies. A decision aid is currently being developed by NRaD/NCCOSC that to improve tactical decision making. Certain aspects of this aid are designed to support critical thinking, and it is expected that these aspects will

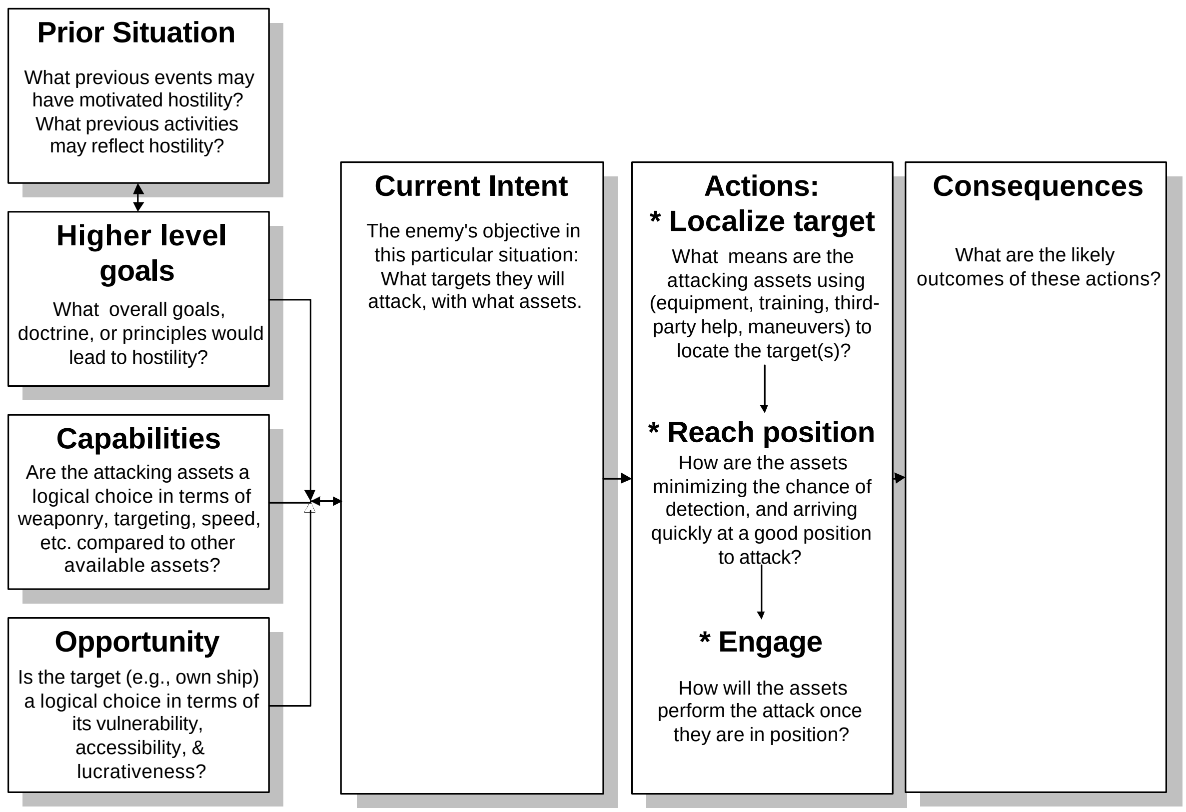

Figure 3. Elements of a hostile intent story.

be expanded in future modifications of the aid. In this section we briefly describe some of the concepts developed by CTI for a more fully developed future version of this system.

Three concepts are central for supporting critical thinking in the DSS:

1. Stories – support for a causal understanding of events through a cause-and-effect arrangement of factors that bear on the intent of a target.
2. Critiquing and correcting — support for a transition from pattern matching to story construction in novel situations, by facilitating strategies for finding and explaining conflicting evidence and for generating and evaluating alternative stories
3. The Quick Test — support for managing time and allocating attention as a function of the costs of delay, the stakes, and the degree of uncertainty

We will discuss each of these elements in turn.

4.1. Stories

The DSS will structure its display in a way that corresponds to the causal story structures discussed in critical thinking training. This will have a number of benefits purely from a DSS design perspective, in addition to compatibility with training:

(i) It will support rapid finding of information. Each type of cue (EW, prior intelligence, altitude, etc.) will have a consistent location every time the user looks for it, rather than appearing at a random location in an unstructured list.

(ii) It will facilitate noticing gaps in current information and the significance of such gaps, leading both to more effective information collection and to a better appreciation of weaknesses in the support for a hypothesis.

(iii) The structure will display information that is expected if a hypothesis about intent is true, not just information that has already been obtained. It thus guides information collection to test situation assessments and supports the identification of conflicting evidence.

(iii) The structure will place items in natural relationships to one another, supporting development of a coherent mental model rather than a mere checklist.

(iv) The structure will facilitate a big picture of the situation. Users will be encouraged to consider important issues that might otherwise be neglected, such as the reasonableness of own ship as a target, the reasonableness of the contact as an attacking platform, the localization abilities of the attacking platform).

Figure 4 is a notional display that illustrates all of these points. Information regarding a track is displayed in a causal framework corresponding to a possible, but uncertain hostile intent assessment. The display resembles the story structure in Figure 3, except that time flows in a rough sense from right to left for consistency with other DSS displays. Users select the assessment they want to evaluate in the box labeled *Intent*. An assessment consists of an intent, a platform possessing the intent, and the target of the intent. Each of these can be selected by the user from drop-down menus listing relevant intents and the platforms in the area.

In the scenario underlying this display, a F-4 has taken off from an airbase in Iran; instead of following the normal patrol route along the coast north or south of the airport, the aircraft begins to circle the airport. The circles are growing larger, bringing the aircraft closer on each cycle to own ship. On the legs of the circle in which it is facing own ship, the aircraft turns on its fire control radar. Officers in the Combat Information Center wish to evaluate the possibility that this aircraft intends to attack own ship. The display in Figure 4 depicts a story based on the assessment that track 7015 (identified as an Iranian F-4) intends to attack track 4001 (own ship, the Chosin). If this intent is true, then it should be possible to tell a plausible hostile-intent "story," i.e., it should be possible to fill in this structure in a plausible way. First, hostile intent should reflect the higher-level goals, capabilities, and opportunities of the country that owns track 7015, and second, the actions of track 7015 should be designed to achieve a successful attack.

The boxes to the right correspond to the first set of issues, i.e., the causes of the intent. *Context* refers to higher-level goals, as reflected in prior hostilities, intelligence regarding intent to attack, or presumed motives. *Capabilities* refers to the reasonableness of choosing the platform under consideration (an F-4) as a platform for attacking the Chosin. And *opportunities* refers to the reasonableness of selecting the Chosin as a target, asking, for example, if it is the most lucrative, the most accessible, or the most vulnerable of the available targets. In this scenario, there is no intelligence regarding an Iranian attack, nor have their been recent hostilities. However, an attack would be consistent with Iranian goals of retaliating against the U.S. for its presence in the Persian Gulf. The F-4 makes sense as a choice of platform to attack own ship (the Chosin). However, the selection of the Chosin as a target does not make much sense: the Essex is both more lucrative and more vulnerable to an attack. This consideration is an argument against hostile intent, though hardly definitive.

The boxes to the left of the intent box represent a second set of issues, the effects of hostile intent. These are the actions and events that should be observed in the immediate past, in the present, and in the near future if track 7015 really intends to attack. These observations are divided into three categories, corresponding to *kinematics*, electronic warfare (*EW*), and *responses* to own ship actions. The horizontal axis represents range from own ship, decreasing from right to left. The vertical black line represents the present range. The kinematics and emissions of

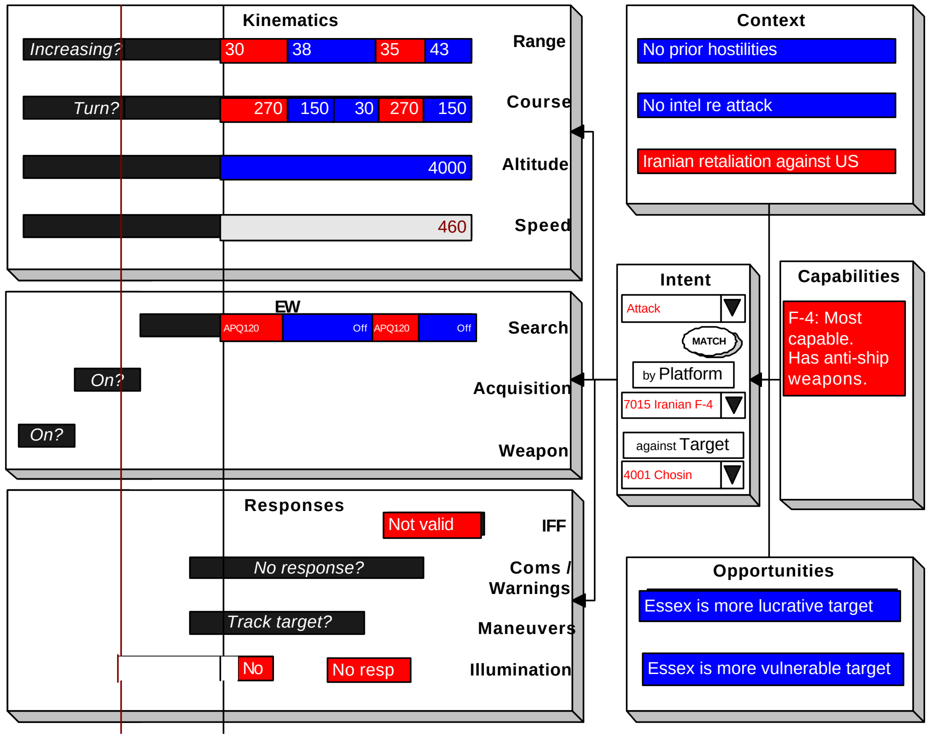

Figure 4. A causally structured display of track information. This display supports evaluation of the hypothesis that an aircraft intends to attack.

the aircraft in this scenario present a highly unusual pattern that both supports and conflicts with the hypothesis of hostile intent. On the one hand, hostile intent is supported when the aircraft is closing own ship and using fire control radar; on the other hand, the aircraft's periodic turning away from own ship and its intermediate altitude conflict with hostile intent.

In this display, information is color-coded to represent its agreement (red) or disagreement (blue) with the hypothesis of hostile intent. Items colored black represent information gaps (either present or future). For example, the display shows that the ship has not yet warned the F-4; nor has it attempted to maneuver to see if the F-4 adjusts its own course accordingly. The display also indicates expectations within the story under consideration: If this track is hostile, it should fail to respond to warnings, track own ship if it maneuvers, continue to approach own ship until reaching engagement range, fire a weapon, and then turn away.

(v) The structure will support reasoning about the assumptions underlying causal relationships. Assumptions will be collocated with the issues to which they are relevant, rather than listed randomly. For example, suppose an aircraft is flying slower than it normally would if it were hostile. On the hypothesis that it is hostile, the user has to make some assumption, such as that it is trying to localize its target visually. This assumption should be displayed in proximity to the information regarding speed.

(vi) The story structure will be integrated with the other DSS displays. For example, a Track Profile and a Response Manager window are currently related to one another through a shared range dimension (decreasing from left to right) and a moving indicator of the target's current location. They might in turn be correlated with the present display, which indicates the range at which events are expected within a hostile-intent story structure.

In sum, a story-based display would organize cues graphically according to their causal relationships. This display integrates all relevant information, places each cue in a consistent location, makes gaps and expectations salient, and provides a full picture of the situation. Such a display should provide a variety of the advantages. For example, according to Pennington & Hastie (1993), "expectations about the kinds of information necessary to make a story tell the juror when important pieces of the explanation structure are missing and when inferences must be made… story construction is a general comprehension strategy for understanding human action… The story provides an automatic index of the importance of different pieces of evidence…"

4.2. Story Building

The DSS facilitates two different decision-making strategies: pattern matching (in relatively familiar or time-stressed situations) and story-building or assumption-based reasoning (in more novel situations, where time is available), with a seamless transition from one to the other. A single story-based display supports both kinds of reasoning. Indeed, this DSS display illustrates how pattern matching and story construction strategies can be used simultaneously in the same problem.

(i) *Pattern-matching.* If evidence is displayed within a causal structure, as discussed above, the display cannot also spatially indicate which evidence supports and which evidence conflicts with a hypothesis. However, as already noted, this can be accomplished in a vivid fashion through color coding. For example, all evidence supporting the hypothesis that the contact is a threat can be shown in red, all evidence supporting non-threat in blue (see Figure 4). Missing data will be shown in back, and inconclusive or unreliable data will be shown in gray. Users will thus see at a glance the balance of supporting, conflicting, and missing or unreliable data, and (by their location in the structure) the nature and importance of each.

(ii) *Story building, or assumption-based reasoning*. The display will show the assumptions that are associated with the current interpretation of each evidence item. As long as the cue is given its normal interpretation (e.g., slow speed displayed in green to signify non-threat), the associated assumption (e.g., approach to attack is typically at high speed) is not shown at all unless the user toggles for a more detailed display, in which case it is shown in a neutral color.

Users can improve the plausibility and coherence of the story, however, by changing the assumptions associated with cues. Changed assumptions are highlighted in yellow to indicate their possible unreliability. At the same time, the color of the cue changes to reflect its new interpretation.

For example, before engaging a contact, users can test the idea that it is a threat by attempting to build a plausible and coherent threat story, and then step back to evaluate it. To do so, they explore assumptions that might explain the apparently conflicting information while keeping the assessment that the aircraft intends to attack. Figure 5 illustrates that process in the scenario of the circling F-4. The user has explained the choice of the Chosin rather than the Essex as a target by assuming that the aircraft is unaware of the presence of the Essex. The user has explained the fact that the aircraft is circling rather than directly approaching the Chosin, by assuming that it is waiting to rendezvous with another aircraft before attacking. As illustrated in Figure 5, the system provides a drop-down list of possible explanations of each cue. Such a list can remind the user of alternative interpretations of cues. Users can select from among these or create new explanations of their own. Figure 5 depicts the user in the process of examining possible explanations of the fact that the aircraft is flying at an intermediate rather than low altitude. The assumption most consistent with the story being built by the user is that the intermediate altitude is a deception to buy time.

The user can now see at a glance what assumptions a particular hypothesis implies. The user can scan the story display for assumptions (marked in yellow), and examine them to see if they make an acceptable story. By the same token, the conflicting cues that are explained by the assumptions change from blue (supporting non-threat) to gray (neutral). The user might also

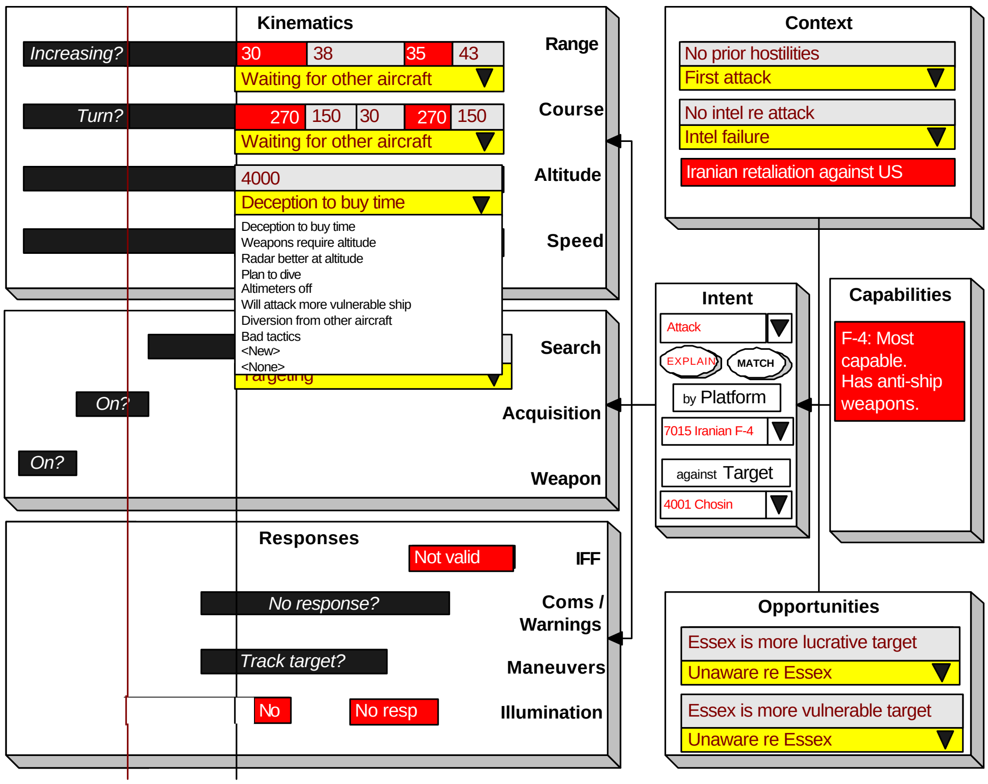

Figure 5 Assumptions have been adopted in an attempt to explain evidence that conflicts with hostile intent. The user can now evaluate this story, and possibly try another.

select the non-threat hypothesis, and see if there is a plausible story in terms of some non-threat intent (e.g., training, search and rescue, transiting, etc.).

In sum, the story-based display will support both a pattern-matching strategy and assumption-based reasoning or story construction. Officers will be able to determine at a glance the pattern of cues matching threat vs. non-threat hypotheses by virtue of color-coding (i.e., blue vs. red). Story construction is supported by displaying the assumptions underlying the interpretation of a cue. Users can change assumptions in order to create the most plausible and coherent story they can, in support of a hypothesis. They can then evaluate the hypothesis by evaluating the assumptions required by the story. Potentially unreliable assumptions are color coded (i.e., yellow) to facilitate this evaluation.

4.3. The Quick Test

The Quick Test is the process by which officers decide either to think more about a problem or to take prompt action. In situations that involve immediate threats or a high density of potential threats, critical thinking may be neither feasible nor wise. However, in the face of a potential but uncertain threat, where the cost of an error (e.g., downing a civilian aircraft) is high, and where time is available, critical thinking may be both possible and worthwhile. An important function of the DSS will be to provide guidance to users regarding when critical thinking is appropriate.

The DSS provides a series of icons representing current contacts in order of priority for attention. By clicking on (or pointing to) an icon, the user receives a more detailed display of information regarding the selected track. In preliminary algorithms that we have developed, the priority order is based on the relative values of the actions that it is currently most desirable to take with regard to a track. Such actions include warning, maneuvering, monitoring, setting up alerts and tripwires, illuminating, and engaging. The value of these actions is based on outcomes like avoiding a mistaken engagement (e.g., by warning), and successfully executing an engagement if necessary (e.g., by setting alerts and tripwires). Critical thinking is simply another action added to the list of available actions regarding a track. Its value increases with the uncertainty of the decision whether or not to engage a given track, the potential costs of being wrong (i.e., mistakenly engaging a non-hostile aircraft or mistakenly not engaging a hostile one), and the time available before the contact becomes an immediate threat.

Typically, for tracks that represent immediate threats, the most desirable actions will involve such actions as warning, illuminating, or engaging. These immediate threats will be ranked higher in priority than tracks whose preferred action is critical thinking. When such immediate threats are not present, however, critical thinking may itself take the highest priority. The DSS will use color coding to mark icons for tracks when critical thinking is the most desirable action for that track.

Currently existing aspects of the DSS will provide backup information relevant to the Quick Test. A Track Profile display of a contact's weapons' range, and a Response Manager display of appropriate actions in regard to a track, provide a basis for estimating available time before risk to own ship becomes unacceptable. The story-based display described above shows the degree of uncertainty regarding a target's intent by highlighting conflicting cues (red vs. blue), missing information (black areas), inconclusive or unreliable information (gray areas), and assumptions (yellow). The costs of a mistaken engagement are suggested by the alternative non-threat assessments that seem plausible (e.g., commercial air, military patrol).

5. Experimental Test

An experimental test of the interaction of training and the current decision support system is scheduled for this spring and summer. The primary objectives of the experimental research are (1) to examine whether training helps users make more effective use of the DSS, and (2) to examine whether instruction, practice, and feedback supported by the DSS enhances critical thinking training. Secondary objectives include (3) testing the relevance and importance of critical thinking skills as conveyed either through training, through decision support, or through both, and (4) improved understanding of how decision support and training can be coordinated.

TAO-qualified active-duty Naval officers will participate in the study. Three treatment conditions will be administered to different groups of participants and compared: (1) critical thinking training utilizing the DSS, (2) familiarization and use of the DSS only, and (3) critical thinking training with current displays rather than the DSS. In the first condition, the DSS will be utilized in the training to illustrate critical thinking concepts, for classroom demonstrations, and for practice. Changes in critical thinking skills from pretest to posttest will be evaluated as a function of these treatment conditions. The results should shed light on the potential value of the DSS as a training tool, and on the potential contribution of critical thinking training to more effective use of the DSS.

6. Conclusion

The research reported here illustrates the transition of a cognitive model of decision making, first to training, and second to decision support system design. The effectiveness of the training has been successfully tested. It increases the frequency of critical thinking strategies and appears to improve both the accuracy of assessments and the appropriateness of actions. We anticipate that support for critical thinking by a real-time decision aid will also improve both critical thinking strategies and performance. The most significant gains, however, are expected when critical thinking training is combined with appropriate decision support.

References

[Cohen, Freeman, & Wolf, 1996] Cohen, M.S., Freeman, J.T., & Wolf, Steve. Meta-Recognition in Time Stressed Decision Making: Recognizing, Critiquing, and Correcting. *Human Factors, 38* (2), pp. 206-219. 1996.

[Cohen & Freeman, in press] Cohen, Marvin S. & Freeman, Jared T. Critical thinking skills in tactical decision making: A model and a training strategy. In Cannon-Bowers, Janis A. & Salas, Eduardo (Eds.), *Decision making Under Stress: Implications for Training and Simulation*. In press.

Kaempf, G. L., Klein, G., Thordsen, M. L., & Wolf, S. Decision making in complex command-and-control environments. *Human Factors, 38* (2), pp. 206-219.

[Nelson & Narens, 1994] Nelson, T.O., & Narens, L., (1994). Why investigate metacognition? In J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition* (pp. 1-25). The MIT Press, Cambridge, MA.

[Pennington & Hastie, 1993] Pennington, N. & Hastie, R. The story model for juror decision making. In Hastie, R. (Ed.). *Inside the Juror: The Psychology of Juror Decision Making*. Cambridge University Press, 1993.