TRAINING THE NATURALISTIC DECISION MAKER [1]

by

Marvin S. Cohen, Jared T. Freeman, and Bryan B. Thompson

Cognitive Technologies, Inc.
4200 Lorcom Lane
Arlington, VA 22207
(703) 524-4331

1 March, 1995

1 This research was funded by Contract No. MDA903-92-C-0053 from the Army Research Institute. We are grateful to Jon Fallesen of the Fort Leavenworth Field Unit for his insightful assistance at all stages of the project. We also thank Martin Tolcott, Leonard Adelman, Terry Bresnick, Freeman Marvin, and Gen. Charles Ottstot (US Army Ret.) for their invaluable help.

The naturalistic approach to training decision making takes as its starting point the way real decision makers make decisions. Unlike normative approaches that are based on decision theory or other formal models, naturalistic training is in no rush to replace ordinary decision processes with methods that may appear to have theoretical value, but which are qualitatively different from the methods that proficient decision makers use. If a house is basically sound, it can be better to renovate than to tear down and rebuild from scratch.

This chapter addresses a specific set of skills that are utilized by experienced decision makers in novel situations. While naturalistic models of decision making have tended to emphasize recognitional processes, recognition is inadequate when no familiar pattern fits the current situation. Although recognition is at the heart of proficient decision making, other processes may often be crucial for success. For example, Klein (1993) discusses mental simulation of a recognized option. We will argue that many of these processes, which verify the results of recognition and improve situation understanding in novel situations, are meta-recognitional in function.

Training metacognitive skills may be applicable across a variety of domains as part of a naturalistic approach to improved decision making. We will briefly discuss a framework for naturalistic decision making, called the Recognition/Metacognition (R/M) model (Cohen, Adelman, Tolcott, Bresnick, &

Marvin, 1993; Cohen, Freeman, & Wolf, 1994), that highlights meta-recognitional skills. We will then turn to our research on training meta-recognitional skills in Army battlefield situation assessment.

META-RECOGNITION IN DECISION MAKING

Meta-recognition is a cluster of skills that support and go beyond the re-cognitional 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. More formally, according to Nelson & Narens (1994), metacognition involves splitting cognitive processes into at least two levels, whose relationship is characterized by the direction of information flow and control. The meta-level monitors the object-level, maintains a model or description of the object-level, and modifies object-level activity, but not vice versa.

In the R/M model, the object-level consists of recognitional processes that activate schemas in response to internal and external cues. It also includes a situation model that integrates these recognitional schemas under the influence of meta-level control. The meta-level includes (1) processes of critiquing that identify problems with the recognitional schemas and the evolving situation model and (2) processes of correcting that may instigate additional observation, additional information retrieval, and/or reinterpretation of cues in order to produce a more satisfactory situation model and plan. These two processes are controlled in turn by (3) a higher-level process called the quick test, which considers the available time, costs of an error, and degree of uncertainty or novelty.

The following example is based on think-aloud problem solving sessions with active duty Army officers who were presented with a battlefield scenario. A division plans officer 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, city Y, is in the south; it has the best supplies in the south; and the best roads are in the south. The planner concludes that the attack will be in the south.

The normal, recognitional meaning of each cue (prior success, a lucrative goal, supplies, and roads) 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, the planner will not consider the issue further. However, when the stakes are high, time is available, and the situation is not completely routine, he may not be content with this initial recognitional response; he may critique it.

Critiquing can result in the discovery of three kinds of problems with an assessment: incompleteness, unreliability, or conflict. An assessment is incomplete if key elements of a situation model or plan based on the assessment are missing. In order to identify incompleteness, the recognitional meanings of the cues must be embedded within a story structure. Story structures depict causal and intentional relations among events and have characteristic sets of components (Pennington & Hastie, 1993). In particular, the main components of stories concerned with assessments of enemy intent are goals, capabilities, and opportunities (which elicit) the intent to attack at a particular place and time (which leads to) actions (which result in) consequences. For example, an officer might conclude that the enemy's intent to attack in the south was adopted because of higher-level goals such as capturing city Y and exploiting prior success in the south, superior capabilities in the south by virtue of better supplies, and superior opportunity via better roads. Future actions that would be expected include removing obstacles in the relevant sector, massing artillery, and moving up troops.

In our example, the officer looks for an argument supporting the conclusion that the enemy will attack in the south based on each component of the story structure. He finds the story to be incomplete because none of the enemy actions expected to occur prior to an attack have yet been observed. More subtly, the story may also be incomplete because the officer has not fully considered the factor of capability. What about the relative strength of artillery, armor, and leadership in the north versus the south? Moreover, he has not fully considered the factor of accessibility. What about mountain or river crossings required in a southern versus a northern attack? Correcting steps may generate the information required to complete this story by directing the retrieval of prior knowledge, the collection of new observations or analyses, or the revision of assumptions.

Another function of critiquing is to find conflict, new arguments whose conclusions contradict the conclusions of existing arguments. In our example, the officer's further consideration of enemy capabilities produced an assessment that both troop strength and leadership were superior in the north. The normal, recognitional meanings of these assessments are that the enemy intends to attack in the north. Moreover, fleshing out the accessibility component of the story produced another conflicting argument: The northern forces had superior river crossing skills, making the northern route easier on the whole.

Critiquing can also expose unreliability in a situation model or plan. Understanding and planning is unreliable if the argument from evidence to conclusion, or from goals to action, is conditioned on doubtful assumptions. For example, taken by itself, 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. Unreliability is different from conflict, because here critiquing neutralizes the argument for attack in the south based on troop movements but does not provide an argument against attack in the south.

Critiquing and correcting for one problem may lead to the creation and detection of other problems. In this example, efforts to create a complete story led to discovery of the conflict between better capabilities and accessibility in the north versus more plausible goals in the south. The officer resolved this conflict by rejecting the normal, recognitional meaning of the evidence favoring attack in the north. He 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 1 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.

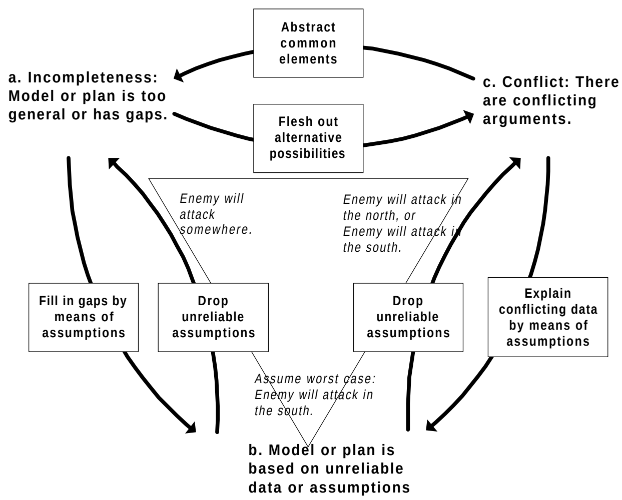

Figure 1. Ways in which correcting steps can lead to new problems.

Our analysis of 34 critical incident interviews with Army command staff suggests an important feature of naturalistic decision making related to Figure 1. 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 will hopefully expose it and trigger efforts to construct a new story.

The officer in our example continued to monitor the situation, but new observations conflicted with his hypothesis of a main attack in the south. Each time, however, the officer generated a modification of the story that accommodated the new evidence. These modifications involved assumptions that neutralized the impact of the conflicting argument. For example, the enemy's deep interdiction efforts destroyed southern bridges, making it difficult for the enemy to advance along the expected attack route. The officer rejected the normal, recognitional interpretation of this observation, which supports attack in the north. He explained the observation by supposing that the enemy might have more bridging equipment than previously known, or that the enemy was more concerned about reinforcements of our forces in the south than rapidity of advance. Another piece of potentially conflicting evidence involved movement of enemy artillery to the north. The officer generated possible explanations of this as well: The artillery might be moved at the last minute to the south; it might have a longer range than supposed; it might have moved north by mistake.

In creating a coherent and complete story, the officer appears to be guilty of the confirmation bias, reinterpreting conflicting data to make it consistent with a favored hypothesis. But we are not convinced that this "bias" is always an error. As discussed in Cohen (1993), it can also be an error to abandon a hypothesis in the face of one apparently conflicting observation. It is possible that the normal meaning of the conflicting observation is not correct. It can be perfectly reasonable for decision makers to use conflict with a larger or more trusted set of arguments as a prompt to look for alternative explanations of conflicting cues.

While meta-recognitional processes help trigger the search for alternative interpretations of cues, they also guard against excesses. They monitor for unreliable assumptions required by explaining away too much. For example, in order to hold onto the view that the main attack will be in the south, the officer must now accept at least some of the assumptions mentioned above, or other equivalent ones: for example, that a diversionary attack is planned for the north, that the enemy has better bridging equipment than expected, and that the artillery will be moved prior to the attack. He might be expected to become increasingly uncomfortable as the number of conflicting observations grows, along with the unsubstantiated assumptions required to explain them. Each new explanation is like stretching a spring. At some point, the extended spring snaps back. At this point he drops the assumptions, and decides to focus seriously on the alternative possibility, a main attack in the north. He then explores a new situation model.

TRAINING META-RECOGNITION

The R/M model describes a set of skills that supplement pattern recognition in novel situations. These skills include identifying key assessments and the recognitional support for them, checking stories and plans based on those assessments for completeness, noticing conflicts among the recognitional meanings of cues, elaborating stories to explain a conflicting cue rather than simply disregarding it, sensitivity to problems of unreliability in explaining away too much conflicting data, attempting to generate alternative coherent stories to account for data, and a sensitivity to available time, stakes, and novelty that regulates the use of these techniques.

These skills are neither as domain-specific as simple pattern recognition, nor as general-purpose as analytical methods. Like analytical tools, meta-recognitional skills may be applicable with minor adaptations across a wide range of domains. Unlike analytical skills, however, their use requires a relatively strong base of familiarity in a domain. They build upon the knowledge embedded in recognitional skills, but do not by any means replace it.

We have prepared and tested a training program on meta-recognition skills in two domains: Army command staff battlefield situation assessment and Navy ship-based anti-air warfare. In both cases, the existing training method consists of several segments, each of which contains (a) reading material, (b) some brief explanations and examples by the instructor, and (c) classroom discussion and interactive exercises. Examples for exercises and discussion are drawn from trainees' own experiences, from the pre-test scenario that is run prior to the training, and from modified case studies in the training materials.

The Army training has two major segments. One focuses on situations in which the decision maker feels relatively certain of his or her conclusions, and focuses on critiquing and correcting for unreliability. The other focuses on critiquing and correcting for conflicting observations. (The Navy training addresses critiquing and correcting for incompleteness and the quick test.) We will very briefly convey some of the flavor of each of the two Army segments.

<u>Handling "Certainty"</u>. We begin the discussion by asking officers for a personal experience in which they felt completely certain of some assessment. We then show how that "certainty" could be questioned. This method forces the officers to generate an alternative story that covers all the evidence. In doing so, it exposes assumptions underlying the story that they currently accept, and helps them evaluate the story for reliability. These assumptions can be evaluated and, if

time and stakes warrant, can be checked. Appropriate correcting steps can be taken when weaknesses in the story are found. In the end, even if officers retain the original story, their confidence in it will have been earned.

The following example of "certainty" was volunteered in one class. A battalion officer, facing an enemy across the river, predicted that they would cross the river at point X. Point X was relatively close to the enemy's present position, the river at point X was relatively shallow, and a combination of vegetation and terrain there would provide concealment. He recommended concentration of friendly forces in the vicinity of point X.

The method for finding hidden assumptions consists of four steps:

1. Select a critical assessment, no matter how confident you are that it is true (e.g., that the enemy will cross the river at point X).
2. Imagine that a perfect intelligence source, such as a crystal ball, tells you that this assessment is wrong.
3. Explain how this assessment could be wrong.
4. The crystal ball now tells you that your explanation is wrong and sends you back to step 3.

After each new exception was mentioned, the crystal ball told the trainee, "No, that's not the reason why the assessment is wrong. Come up with another explanation." In this particular case, the crystal ball elicited a number of ways this “certain” assessment might fail: (i) The enemy might anticipate that our force will be at point X and decide not to cross there. (ii) The enemy might detect the movement of our force to point X and decide not to cross there. (iii) There are good crossing sites that we missed. (iv) The enemy doesn’t know how good a location point X is. (v) The enemy doesn’t have any river crossing assets. He can’t cross the river at all. (vi) The enemy’s river crossing assets are so good that he can cross elsewhere. (vii) The enemy has a large enough force that they can accept casualties in crossing elsewhere. (viii) The enemy’s objectives are not what we thought. He doesn’t need to cross the river. (ix) The enemy will use air assault forces to get across the river.

Usually, trainees are surprised at the quantity and the plausibility of the exceptions that the crystal ball elicits from them. They now realize that the original assessment rested on the assumption that none of these exceptions was true. However, the existence of these possible exceptions is not adequate cause to abandon that assessment. The next step is to evaluate the exceptions. Each one should be considered, at least briefly. The class is asked how they would handle each one. Some possible exceptions may be implausible, for example, that the enemy can afford large casualties. Some can be tested by data collection or by requesting additional intelligence, for example, that the enemy has superior river crossing assets. Other exceptions may motivate a change in plans to make them less likely. For example, to avoid anticipation or detection of our forces at point X, we might position our forces elsewhere, then move to point X later. Other exceptions may cause adjustments in planning to handle them in case they turn out true. For example, we might place reserves on paths behind the river in case we missed some sites or the enemy missed point X. Exceptions may also cause the adoption of a contingency plan. For example, if the enemy's objective turns out to be on the other side of the river, we might prepare to cross the river ourselves. Finally, some exceptions might have to be accepted as known risks, for example, if the enemy uses air assault.

Handling Conflicting Data. In the second unit of training, we ask officers to describe personal experiences in which they were surprised; for example, the enemy attacked in an unexpected sector. We then ask if any cues or indicators had been observed that, in hindsight at least, could have served as a warning. Typically, such cues are clearly remembered, but were disregarded at the time.

A common response to observations that conflict with a previous conclusion is to disregard or discount them. Another response to conflict, which may be equally bad, is to lose confidence and immediately abandon the original assessment. An unexpected event means that situation understanding is imperfect, but the fault may not lie in the original assessment. It may lie in an incorrect interpretation of the new event. In situations where no patterns fits all the data, the correct assessment must involve some "explaining away" of conflicting data.

In this training segment, officers learn to monitor or critique for conflicting evidence and learn how to handle conflicting observations when they occur. When conflict is detected, they begin by modifying the current story to explain the surprising events in terms of their original assessment. They then evaluate the reliability of the resulting story. If the explanations prove to be implausible, officers alter the assessment itself and create a new story. The procedure consists of these steps:

1. Notice unexpected events.
2. Explain an unexpected event in terms of your current assessment. If there have been previous unexpected events, try to find the simplest reliable explanation covering all of them.
3. Evaluate the reliability of your account of all the unexpected events.
4. If the explanation is not reliable, change your assessment and return to step 1.

The crystal ball technique can be useful in generating explanations of conflicting data. Now, however, the crystal ball tells you that the original assessment is correct, despite the conflicting observation, and asks you to explain how this could be. The crystal ball rejects each explanation the trainees generate and asks them to find another. For example, the destruction of southern bridges might be consistent with a main attack in the south, if: (i) The bridge was destroyed to prevent our troops from being reinforced; (ii) the enemy has better bridging equipment than we thought; (iii) destruction of the bridge was a mistake; (iv) it was part of a deception; or (v) the bridge was destroyed by our own troops rather than by the enemy. The other surprising event, observation of artillery moving north, can be explained similarly by elaborating the story based on the original assessment.

A list of this kind may never be exhaustive. Nevertheless, it provides an understanding of the kinds of ways in which the current assessment could still be true despite a conflicting observation. To hold onto the assessment, it is not necessary to know which, if any, of these explanations is the case. But some such story elaboration must be true if the original assessment is to be maintained. Thus, the original assessment is no more reliable than the best of these explanations. These explanations may also point to ways that the assessment can be tested.

The training proceeds to illustrate and discuss the dangers of explaining away too many conflicting cues. We have compared this process of adopting new assumptions to the gradual stretching of a spring until finally it snaps back, the assumptions are dropped, and a new assessment is explored. The spring stretches less if explanations of new conflicting events do not introduce new assumptions. The most reliable overall account must consist of individually plausible explanations, but it should also be simple. It should not require a large number of separate assumptions. For example, an account of events that relies heavily on enemy mistakes to explain all of the unexpected events is maximally simple, but not necessarily reliable. Nevertheless, it may be worthwhile to look for an explanation that accounts for a variety of unexpected observations in the same way, such as a single unreliable source or a single enemy strategic concept.

EXPERIMENTAL TESTS

The training methods have been tested in a pretest/posttest control design with active duty Army officers. We have space here only for the most cursory summary of the results. Participants varied in rank from Captain to Lieutenant Colonel. Twenty-nine received 90 minutes of training, while 8 controls were engaged in a discussion of traditional situation assessment techniques. All participants were given a pretest and posttest consisting of a military scenario followed by a set of information updates pertaining to the scenario. Each update consisted of two parts. The first part presented observations concerning enemy activity, terrain, attitudes of the local populace, international diplomatic developments, and other issues and gave an assessment based on those observations. The second part gave a new set of observations, which supported or disconfirmed the original assessment. Participants were asked to evaluate the assessment twice: first, after receiving the original observations, then after receiving the new ones.

Trained participants tended to consider more factors in their evaluations of the assessments than did controls ($F_{1,31}$=3.67, p=.065). More importantly, training increased the value of the factors that officers considered ($F_{3,27}$=3.166, p=.04). Training caused participants to consider fewer factors that were neutral (or irrelevant) with respect to the assessment and to notice more factors that conflicted with the assessment. Moreover, the increase in the number of factors that trained officers considered did not occur at the expense of quality. On the contrary, the issues presented by trained participants were judged in a blind rating by a subject matter expert to be of higher quality on average than those presented by control participants ($F_{1,31}$=3.255, p=.081).

CONCLUSION

Based on critical incident interviews in two domains, the R/M framework appears to capture the way experienced decision makers test and improve the results of recognition. Moreover, it explains how critical thinking can be an effective component of situation assessment without the need for inflexible analytical methods. For example, the R/M model offers a very different approach to situation assessment than Bayesian decision theory. In the latter, every possible implication of every observation must be identified and quantified at the beginning of an analysis. To avoid intractability, the assessment of the impact of a given piece of evidence must be independent of the context of other evidence. This requires a difficult counterfactual mind set: What would this piece of evidence mean if I didn't know about all the other things that have occurred and if I had never thought about them? The product of all this work is a set of probabilities over statements rather than a coherent picture of the situation.

According to the R/M framework, by contrast, proficient experienced decision makers work with evolving situation models or stories. They approach new data in terms of these models, while at the same time looking for gaps and conflicts and being prepared for surprises. When an unexpected or conflicting event occurs, they elaborate the model to take it into account. But they maintain an awareness of their elaborative efforts and stay alert to the danger of going too far. Alternative stories are investigated when the current story has been stretched too thin (as in the training for handling conflict) or even if a story appears sound, when time is available and the stakes are high (as in the training for handling "certainty").

Based on our very preliminary tests, these meta-recognitional skills can be enhanced by training. The result may be more effective situation assessment across a very wide range of decision domains.

REFERENCES

Baker, L. (1985). How do we know when we don’t understand? Standards for evaluating text comprehension. In D.L. Forrest-Pressley, G.E. MacKinnon, &

Waller, T.G. (Eds.), *Metacognition, cognition, and human performance*. NY: Academic Press, Inc., 155-205.

Cohen, M.S. (1993). 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.

Cohen, M.S., Freeman, J.T., & Wolf, S. (1994). Meta-cognition in time-stressed decision making: Recognizing, critiquing, and correcting. Submitted to the *Journal of the Human Factors and Ergonomics Society*.

Cohen, M.S., Adelman, L., Tolcott, M.A., Bresnick, T.A., & Marvin, F.F. (1993). *A cognitive framework for battlefield commanders' situation assessment* (Technical Report 93-1). Arlington, VA: Cognitive Technologies, Inc.

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.

Nelson, T.O., & Narens, L (1994). Why investigate metacognition? In J. Metcalfe & A.P. Shimamura (Eds.), *Metacognition: Knowing about Knowing*. Cambridge, MA: The MIT Press.

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.