AlphaACT™
ADVANCED CRISIS TECHNOLOGY

The AlphaACT Decision Support System for Emergency Responders

An T. Oskarsson & C. Reed Hodgin

To appear in

The Proceedings of the 54th Annual Meeting of the Human Factors and Ergonomics Society

September/October 2010

AlphaTRAC™
CRISIS SOLUTIONS THAT SAVE LIVES

The AlphaACT Decision Support System for Emergency Responders

An T. Oskarsson & C. Reed Hodgin
AlphaTRAC, Inc.

We first present cognitive psychological theories relevant to crisis decision making. Next we describe the AlphaACT™ (Alpha Advanced Crisis Technology) decision support system, a software application designed to support a range of military and civilian emergency responders. The essential components of the system are as follows: (a) a user interface that walks the decision maker through a multi-step decision process inspired by the recognition-primed decision model and case-based reasoning theory; (b) a pattern recognition engine that prompts the user for diagnostic information and retrieves similar cases; and (c) a community-wide shared knowledge base of cases that grows as the system is used. AlphaACT's objective is to train and enable responders in crisis situations to think like experienced decision makers, and to quickly build their store of available experiences. The first AlphaACT application under development will support key decisions made by first responders managing a hazardous materials emergency.

BACKGROUND

Human error has been found to be responsible for 70% of aviation accidents (Sniezak, Wilkins, Wadington, & Baumann, 2002). During crisis and emergency situations, good decision making is critical because the primary goal is to prevent or mitigate extremely negative consequences. Events are often unexpected, life-threatening, and occur under conditions of great urgency, stress, instability, and uncertainty.

An approach to decision research that is well aligned with crisis decision-making is naturalistic decision-making (NDM) (Lipshitz, Klein, Orasanu, & Salas, 2001). NDM focuses on people who use their experience to make decisions in real-world contexts, and aims to examine how successful decision makers size up the decision problem and make reasonable decisions that are compatible with the situation. According to this view, experienced decision makers can skillfully use their experience and prior knowledge to assess a situation and appraise decision options — a process sometimes colloquially referred to as "intuition" — and quickly settle on a course of action (Klein, 2003). The NDM perspective emphasizes the importance of studying complex decision making under conditions of uncertainty, time pressure, and stress. Crisis decision making involves all of these factors, often to an extreme degree, and under rapidly changing conditions.

Recognition-primed Decisions (RPD)

The most cited and prototypical NDM model is the RPD model (Klein, 1993; 1998). It is considered by some to be the most appropriate framework for describing how proficient decision makers operate, especially under conditions of time pressure and uncertainty (Lipshitz, et al. 2001). According to RPD, often the goal of the decision maker is to rapidly find and select the first reasonable and workable solution. We choose here to focus on the variations of the RPD model that can involve less typical situations. In our view, RPD essentially propose the following major steps in the decision process:

- Characterize the decision problem and diagnose the situation
- Recognize in memory a similar situation
- Achieve a better understanding of the current situation through comparison with the recognized situation (considering cues, goals, expectations, & actions)
- Mentally simulate the indicated course of action to gauge whether it will succeed, making modifications if needed

Originally based on a cognitive task analysis of firefighters, replications with other groups (e.g., design engineers, offshore oil installation managers, commercial aviation pilots, British army officers) have found that RPD strategies were used by experienced decision makers in 80-95% of cases (Klein, 1993). Research on judgment heuristics lends indirect support for RPD. For example, recognizing cases based on similarity is consistent with the representativeness heuristic, a mental rule-of-thumb, whereby judgments depend on the degree of assumed similarity between cases (Kahneman & Tversky, 1973).

Beliefs in "the adaptive decision maker" and "fast & frugal heuristics" (Payne, Bettman, & Johnson, 1993; Todd & Gigerenzer, 2001) support the notion that suitable and efficient decisions can often be made based on the consideration of minimal information (e.g., Take-the-Best heuristic, recognition heuristic). In satisficing choice strategies, options are considered one at a time and searching stops when an option is considered to be "good enough." Along these lines, the mechanism underlying sequential sampling process models assumes that, rather than taking a predetermined quantity of information, sampling of each option occurs until evidence sufficient to favor one option over the other has accumulated (Busemeyer & Townsend, 1993). One can argue that the measure of "rational" decision making depends on the extent to which the decision works well in our environment, not whether it adheres to a set of formalisms.

Case-based Reasoning (CBR)

The theory of CBR resembles the basic process underlying the RPD model (Kolodner, 1993, 1997; Riesbeck & Schank, 1989). CBR focuses on the mind's ability to apply analogs to solve real-world problems. More often than not, solutions are not constructed from scratch; instead, previous experience brings to mind old problems that suggest possible solutions to new problems. "History repeats itself," and solutions that worked in one situation are likely to be applicable in similar situations.

Support for CBR can be found in psychology and artificial intelligence literature (Kolodner, 1993; Noh, Lee, Kim, Lee, & Kim, 2000; see Bloomfield & Moulton, 2008, for an applied example). CBR can be found in the world around us (e.g., a doctor's diagnosis based on a prior patient case, a lawyer preparing arguments based on legal precedents, a mechanic fixing an engine by recalling a car with similar symptoms), and in everyday personal problem solving. CBR involves:

- Retrieval: Selection of a similar source case from memory
- Adaptation: Revision of the proposed solution, if needed
- Learning: Retention of the solution to form a new case; update experience

Successful CBR requires recognizing the applicability of an old situation to a new one. Knowledge and experience relate to the number and richness of cases in the decision maker's memory, and to the ability to encode into memory the case information that later will facilitate efficient retrieval of appropriate source cases. According to Ross (1986, 1989), the difference between novices and experts is that novices have fewer relevant experiences to draw upon and do not encode cases and details as well, making it difficult to retrieve appropriate source cases.

CBR retains the value of specific case details. Through induction processes, CBR methods can acquire knowledge with ease—making CBR especially useful when knowledge is incomplete and information is limited. This kind of structured analogical reasoning tends to be a good approach for rich, complex domains in which there are many qualitatively different ways to generalize a case. Critics of CBR argue that their theories rely too much on observational methods and anecdotal evidence. However, all inductive reasoning in circumstances involving scarce data is less amenable to statistical generalizations and regression methods. This is especially relevant to crisis decision making, since crises are by definition infrequent occurrences.

THE AlphaACT SYSTEM

One common objective for a decision support system is to help users analyze more and more information in order to come up with an optimal decision. The authors of this paper took this approach in the past, when we assisted U.S. Department of Energy (DOE) emergency responders in preparing for and dealing with hazardous material emergencies - and our efforts were met with resistance. We were told that decisions are typically made with limited or wrong information, and that often a "90% solution now is better than a 100% solution later". First responders were wary of succumbing to "paralysis by analysis", and becoming ineffectual when lives could be lost if one did not act fast.

It became apparent that the most effective decision makers were the ones who were consistently able to quickly generate successful solutions based on a few pieces of information, by drawing on their prior experiences with similar events. We observed that their cognitive processes and decision strategies were compatible with the fundamentals of RPD and CBR theory. We therefore took an alternate approach to designing the AlphaACT™ system, and focused on helping decision makers use a sufficient amount of data to diagnose the situation and quickly arrive at a workable solution.

This led us to our primary objectives for AlphaACT. Firstly, we aim to support decision makers in following a decision process that is guided by the RPD framework. Secondly, the system helps decision makers efficiently use and look for information that is important for recognition of appropriate cases. Finally, the system builds and expands the decision makers' knowledge store of available cases. Accomplishing these objectives would result in a system that can train user to think like experienced decision makers. In the next sections we describe the essential components of AlphaACT: A pattern *recognition engine* that searches for the key unknown information and similar cases, an expansive and growing *knowledge base* of cases, and a user interface walking the user through our *decision process*.

Recognition Engine

Research in biological neural systems has led to attempts to simulate their functionality as cognitive architecture in computer software. Practical applications have resulted in artificial neural network systems with great computational capacity. Connectionist and neural networks can process information in parallel and in a nonlinear and distributed fashion, and they have a capacity to adapt with new data (Glockner & Betch, 2008; Newell & Broder, 2008).

AlphaACT uses a simple neural network, or "recognition engine", to quickly and efficiently search through large relational knowledge bases and recognize patterns in data. It is a multi-layer fully connected network that uses a sigmoid symmetric activation function and is trained using the Rprop training algorithm (described by Igel & Husken, 2000). With each new input submitted by the user, the recognition engine determines case matches from the knowledge base and the next piece of key data for situation diagnosis. This enables a question & answer interview interface, dynamically identifying the next question for which an answer is likely to lead to convergence onto matching cases in the knowledge base. Even if the user does not know the answers, the questions teach the user to focus on and gather the critical information.

Other potential advantages of the recognition engine include an ability to fill in missing data, abstract a prototypical case from partial data, and detect "poor" data that is inconsistent with the overall pattern of activation. Furthermore, it can periodically re-train and update its weights when new cases, cues, and decisions are entered into the database, so that the engine's algorithm improves with use.

Knowledge Base

A knowledge base of cases serves as the foundation of the AlphaACT system. The knowledge base is initially created using a streamlined version of the applied cognitive task analysis method (Militello & Hutton, 1998), whereby knowledge is elicited during structured interviews with subject matter experts. Training materials are also a source of information.

In later stages of database development, the system employs an innovative approach to knowledge acquisition in which decision makers can add new cases, cues, and decisions into the knowledge base through a semi-automated process. As the system is used by a community of first responders for training and after-action reviews, competitive "gaming", and field response, the centralized knowledge base expands. This increases the amount of "experience" from which a user can draw. In addition, over time, the relative predictive validity of event information in determining decisions improves as the system's recognition engine "re-trains" on updated databases and takes into account the collective experience and input from a community of users. Description of these functions is beyond the scope of this paper; however, a shared and expansive knowledge base is vital to AlphaACT's utility.

Decision Process

The AlphaACT system is not a strict representation of the RPD model. Rather, it is *inspired* by RPD and CBR, and intended to facilitate the fundamental cognitive processes underlying these models. To this end we structured the decision process into the following steps: Characterize, Recognize, Analyze, Customize, Dramatize, and Utilize. We next describe the process in some detail and explain how each step is supported in the AlphaACT HAZMAT prototype.

AlphaACT FOR HAZMAT EMERGENCIES

Whenever a chemical is accidentally released into the environment, first responders need to decide within minutes whether to declare a state of emergency, the alert level, whether to evacuate and/or shelter people in place, the best safe route for evacuation, and the physical distance to which protective actions should be taken. The incident commander determines where to set up the command post and how to secure the area. For our first application, we chose to focus on these key decisions that occur during the initial stages of a hazmat emergency.

Step 1: Characterize. The decision maker receives initial event information and sizes up the situation. AlphaACT prompts the decision maker to provide known event information that is important for retrieving an appropriate case from the database. In AlphaACT HAZMAT, the user can use a scroll-down menu to input information about the context (e.g., weather, wind direction, terrain) and the "target" (e.g., physical state of the substance, number and type of containers). Users are able to click on links and scroll over terms to view pictures or textual explanations. See Figure 1 for an example screen.

Step 2: Recognize. The decision maker further assesses the situation, answering key questions that facilitate efficient recognition of possible matching cases. AlphaACT uses a question & answer interview interface to ask the user for observations and cues that are likely to diagnose the event. This is done in a "smart" order, as each question presented is based on what information has been entered thus far. The recognition engine quickly weights and combines cues to retrieve cases and diagnostic questions in the database. Answers are processed and recognition is iteratively updated.

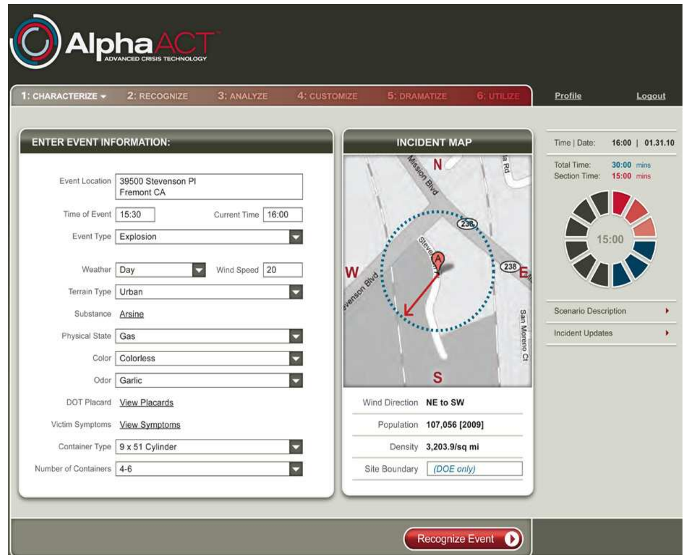

Figure 1. Example screen from the Characterize decision step.

Figure 2 depicts an example Recognize screen in AlphaACT HAZMAT. Here the system deems that the question "What is the estimated container size" is likely to help convergence towards matching cases. After an answer is submitted, the recognition engine selects a new question to ask (e.g., "Is anyone nearby experiencing nausea or vomiting?") The next probable questions can also be displayed, so that the user can plan to obtain the indicated information if they so choose. In Figure 2, multiple matching cases are presented so that the user – especially novices - can see the range of possible protective action distances and event classification. The user can continue to answer questions or select a case (or cases) for further analysis.

Step 3: Analyze. The decision maker examines the case(s) selected from the database in more detail. He or she can perform side-by-side comparisons between the selected case and the present emergency to determine if it is good enough to work from. The user takes into consideration differences in the particulars of the current situation, as well as their present needs and goals. For example, in the HAZMAT application, the user is able to read a description of the past case and access other supplementary information, such as a material safety data sheet that shows the recommended personal protective equipment, first aid information, and chemical incompatibilities. This information can serve to generate expectancies, goals, and actions.

Step 4: Customize. This step includes features that allow the user to "customize" the selected case when the case selected from the database differs from the present emergency. In AlphaACT HAZMAT, an analytical model can predict the resulting difference in the recommended decision. In other words, in this step users can run "what if" analyses to see how changes in information might affect their decisions.

Step 5: Dramatize. The decision maker plays out the decision mentally (or with the team) to assess the feasibility of performing the decision actions. In this step, he or she uses checklists, visual displays, and other tools to conduct mental rehearsal of the recommended decision, evaluate its viability and appropriateness for the present circumstances, and make modifications if necessary. To illustrate, in Figure 3, users are able to “drag and drop” the command post and traffic control points onto a map, a list of high-value targets are highlighted on the map, and a protective action checklist is provided.

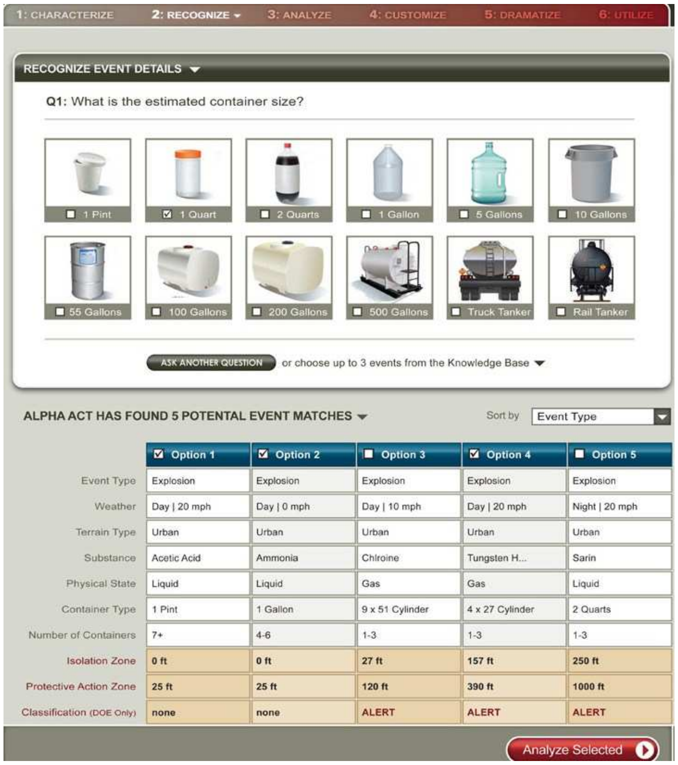

Figure 2. Example screen from the Recognize decision step.

Step 6: Utilize. The decision maker executes the course of action and the system assists in communication and implementation of the decision. Resources and features (such as the ability to share, print, and email maps, reports, and forms) are available. Afterwards, the new case and the cues that prompted recognition, as well as the actual decision and its outcome, are entered into the central database of cases.

AlphaACT HAZMAT will be able to operate in training as well as in response mode. The training version presents a scenario at the beginning of a session, and the scenario “unfolds” as the user moves through the decision process. New information is intermittently “injected” into the scenario in real time. Various types of performance feedback are also provided; Indicators of the user’s decision success, skill level, and statistics (e.g., average decision time) are displayed. The scenario’s critical information is explained, and the user’s inputs and decision process are critiqued.

Feedback from Chemical Emergency Responders

We presented the AlphaACT concept and solicited feedback at workshops involving eight DOE sites in 2009. Participants included incident commanders, emergency directors, fire chiefs, security officers, and technical staff.

Several common themes emerged. Firstly, there was consensus that the AlphaACT decision process is an approach used by DOE emergency decision makers. Participants believed that AlphaACT could bring significant advances in the support of emergency responders, and that it should be extended beyond hazmat to include other materials, such as biological, radiological, nuclear, and explosive. Use of the tool to promote knowledge sharing across the DOE complex would be a substantial benefit, providing “experience compression” to the novices and “experience expansion” to the experts.

A pilot study was also conducted to gather evaluation data for a prototype of AlphaACT HAZMAT. Seven inexperienced (with no emergency response experience) and five experienced (with an average of 26 years experience) participants read a brief scenario about a hypothetical HAZMAT emergency. They imagined they were the Incident Commander and listed the information they would consider to help them size up the situation. After playing with AlphaACT for about 30 minutes, participants completed another size-up questionnaire with a different scenario. Lastly, they completed a satisfaction survey.

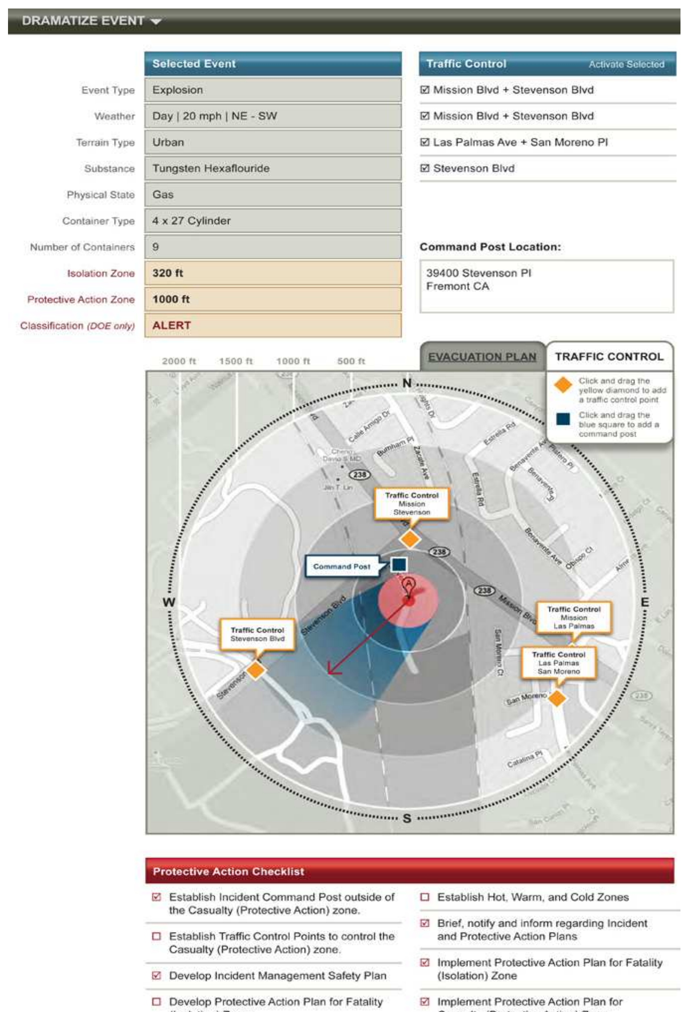

Figure 3. Example screen from the Dramatize decision step.

Survey results. Participants rated their overall satisfaction with AlphaACT, its decision process, and its Q&A interview. They also rated its ease of use, utility, their willingness to use

AlphaACT, and whether they believed it would help them make good decisions. The mean ratings were all favorable, ranging between 4.0 and 4.4 on a scale of 1 to 5, and all were significantly above the neutral point of 3 ($p < .05$). Likely due to the small sample, the only difference found between the inexperienced and experienced was that the inexperienced participants reported being more willing to use AlphaACT ($M = 4.8$ vs. 4.0, respectively, $p < .05$).

Questionnaire results. To test whether participants learned to consider information critical for sizing up a HAZMAT emergency situation, we analyzed whether participants listed more key items of information *after* they played with AlphaACT HAZMAT, compared to *before* they played with it. On average, 8.1 pieces of critical information were listed *after* playing with AlphaACT HAZMAT, while 5.1 items were listed *before* playing with the system. This difference was statistically significant ($M = 3.0$, $p < .05$). No experience effects or scenario order effects were found.

SUMMARY & DISCUSSION

A number of psychological theories support the core principles of the AlphaACT system. This case-based system is designed to emulate how successful decision makers think in crisis situations, and trains users to follow a decision process inspired by CBR theory and the RPD model.

A theoretical issue to consider is whether the RPD model should be treated as a prescriptive model. Although it may accurately describe how experts operate when making emergency decisions, some may argue that a decision aid should still aim to support normative and rational decision making. There can be skepticism about a system that relies on experts' cognitive processes, because of work by (for example) Meehl (1954) and Swets, Dawes, and Monahan (2000), which convincingly explain how actuarial judgments rendered by statistical models tend to outperform clinical judgment of experts. These issues highlight concerns that some have with the NDM approach (Bazerman, 2001).

However, AlphaACT applies optimizing principles to produce decision options, utilizing statistical algorithms to process information and recognize patterns in the data. To some extent the system "mechanizes" situation assessment and recognition in a way that mitigates possible biases and problems with expert validity. Although the user experiences a subjective feeling of satisficing, the system's recognition engine maximizes the search and retrieval of matching cases and question prompts, based on the user's input.

A pilot study in which people evaluated the AlphaACT HAZMAT prototype yielded promising results. The preliminary findings showed that the system was favorably evaluated by users and suggested that familiarity with AlphaACT HAZMAT could lead to rapid learning. Further studies examining the system's effectiveness are warranted.

AlphaACT's utility will be further tested when it is applied to other domains. Prototypes serving a spectrum of emergency responder communities are under development. For example, a combat patrol application trains the detection of improvised explosive devices and supports response to a complex attack. Our hope is that AlphaACT will lead to better decision outcomes in emergency situations, and to significant advances in crisis decision making.

REFERENCES

Bazerman, M. (2001). The study of 'real' decision making. Journal of Behavioral Decision Making, 14, 353-384.

Bloomfield, L.P., & Moulton, A. (2008). MIT Cascon System for Analyzing International Conflict, accessed from http://web.mit.edu/cascon/ on April 16, 2008.

Busemeyer, J.R. & Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432-459.

Glockner, A., & Betsch, T. (2008). Modeling option & strategy choices with connectionist networks: Towards an integrative model of automatic & deliberate decision making. *Judgment & Decision Making*, 3(3), 215-228.

Igel, C., & Husken, M. (2000). Improving the Rprop learning algorithm. Proceedings of the Second International ICSC Symposium on Neural Computation (eds. H. Bothe and R. Rojas), 115–121, ICSC Academic Press.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237-251.

Klein, G.A. (1993). A recognition-primed decision model of rapid decision making. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsambok (Eds). *Decision Making in Action: Models & Methods*. Ablex Publishing, Norwood, CT.

Klein, G.A. (1998). *Sources of Power: How People Make Decisions*. MIT Press, Cambridge, Massachusetts.

Klein, G.A. (2003). *Intuition at Work: Why Developing Your Gut Instincts Will Make You Better at What You Do.* Doubleday Business, New York, NY.

Kolodner, J. (1993). *Case-based Reasoning*. Morgan Kaufmann Publishers, San Mateo, CA.

Kolodner, J. (1997). Educational implications of analogy. *American Psychologist*, 52(1), 57-66.

Lipshitz, R., Klein, G.A., Orasanu, J., & Salas, E. (2001). Taking stock of naturalistic decision making, *Journal of Behavioral Decision Making*, 14, 331-352.

Meehl, P.E. (1954). *Clinical versus statistical prediction: A theoretical analysis and a review of the evidence.* Minneapolis, MN: University of Minnesota Press.

Militello, L.G., & Hutton, R.J.B. (1998). Applied cognitive analysis (ACTA): A practitioner's toolkit for understanding cognitive task demands. *Ergonomics, Special Issue, 41*(11) 1618-1641.

Noh, J.B., Lee, K.C., Kim, J. K., Lee, J.K., & Kim, S.H. (2000). A case-based reasoning approach to cognitive map-driven tacit knowledge management. *Expert Systems with Applications*, 19, 249-259.

Newell, B., & Broder, A. (2008) Cognitive processes, models & metaphors in decision research. *Judgment & Decision Making*, 3(3), 195-204.

Payne, J., Bettman, J., & Johnson, E. (1993). *The Adaptive Decision Maker*. Cambridge University Press, New York, NY.

Riesbeck, C., & Schank, R. (1989). *Inside Case-based Reasoning*. Lawrence Erlbaum Associates, Northvale, NJ.

Ross, B. (1986). Remindings in learning: Objects & tools. In S. Vosniadou & A. Ortony (Eds.), *Similarity, Analogy, & Thought*. Cambridge University Press, New York, NY.

Ross, B. (1989). Some psychology results on case-based reasoning. In K.J. Hammond (Ed.), *Proceedings: Second Workshop on Case-based Reasoning* (DARPA), Morgan Kaufmann Publishers, San Mateo, CA.

Swets, J.A., Dawes, R.M., and Monahan, J. (2000). Better decisions through science, *Scientific American, 283*, 4, 70-75.

Sniezak, J., Wilkins, D.C., Wadington, P.L., & Baumann, M.R. (2002). Training for crisis decision making: Psychological issues and computer-based solutions. *Journal of Management Information Systems*, 18(4), 147-168.

Todd, P., & Gigerenzer, G. (2001). Putting naturalistic decision making into the adaptive toolbox. *Journal of Behavioral Decision Making*, 14, 353-384.