Adaptive Technologies for Training and Education

Edited by
PAULA J. DURLACH
U.S. Army Research Institute
ALAN M. LESGOLD
University of Pittsburgh

CAMBRIDGE
UNIVERSITY PRESS

CHAPTER 8

Training Decisions from Experience with Decision-Making Games

Cleotilde Gonzalez

The research reported here is relevant to human decision making in disaster, emergency, and generally dynamic and changing conditions. In emergency conditions, both external resources and cognitive capabilities are limited. A decision maker must manage his/her cognitive load, that of others in the situation, and allocate the often limited resources in a short period of time. The main job of a decision maker in an emergency situation is to allocate resources wisely where the situation changes constantly, outcomes are delayed, decisions are interdependent, priorities change over time, and there is high uncertainty of outcomes.

These types of decision-making conditions are part of a field of research called Dynamic Decision Making (DDM) (Brehmer, 1990; Edwards, 1962; Rapoport, 1990). Some classical examples of DDM include, among others: firefighting resource allocation and management in real time; triage decisions in a medical emergency room; 911 operators determining relative urgency and deploying resources; and supply-chain management. In general, DDM often involves a dynamic allocation of limited resources in real time. Disaster and emergency responses are examples of decision making in dynamic conditions. For instance in an emergency room, doctors must allocate limited hospital resources (i.e., nurses, doctors, beds) in the presence of a large number of patients in need of help. The rate and timing of the inflow of patients is unknown, and physicians must make assignments in real time. These allocations are interdependent, perhaps producing suboptimal decisions like when a mildly injured patient is given priority over future unknown severely injured patients.

A common paradigm in the behavioral decision sciences has been that of *decisions from description* (DFD, Hertwig, Barron, Weber, & Erev, 2004). A human is given access to descriptive information, often including probabilities and outcomes, and asked to make a choice based on the conditions described. This is also a common paradigm in the real world. For example, to make an investment decision, people often read brochures with the different pros and cons of the alternatives while considering the risk of the different alternatives. These

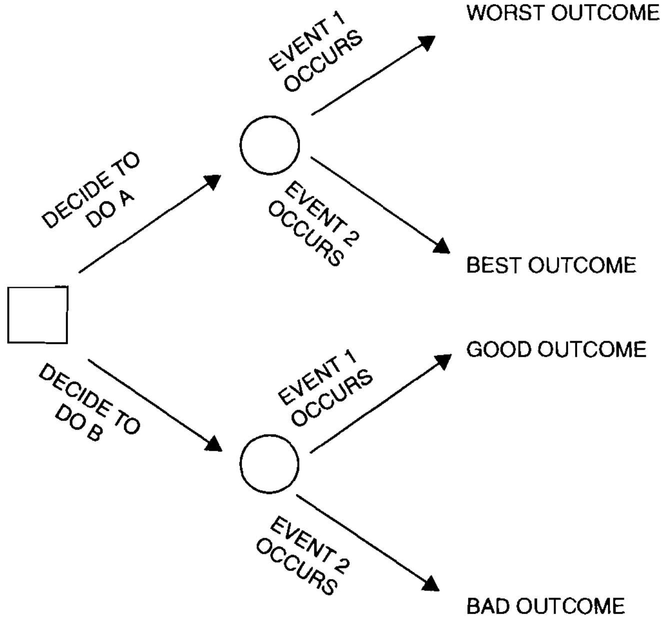

Figure 8.1. Example of a common problem structure studied in behavioral decision-making research.

type of decisions are one-shot (open-loop, no-feedback) decisions, illustrated in Figure 8.1, where people are expected to act "rationally" and select the alternative that results in the best outcome.

In emergency and DDM conditions, DFD are quite uncommon. One rarely finds well-defined alternatives; one rarely understands or is able to correctly predict all possible outcomes; and often high uncertainty prevents one from calculating or estimating any probabilities involved in the occurrence of those alternatives (Gonzalez, Lerch, & Lebiere, 2003; Klein, 1989). Traditional behavioral decision research often makes predictions based on "rational theory," the simple and intuitive idea that people search for the optimal outcome (often monetary outcome) in decision situations. Drawing conclusions from controlled experiments in very simplified decision-making situations has multiple advantages. It allows us to identify clear strategies and measures of human decision making, and thus make theoretical progress by drawing generic inferences about how people make decisions. However, this approach has some limitations, notably their applicability in disaster, emergency, and dynamic environments in general. Thus, the many years invested in studying DFD in the behavioral decision sciences have resulted in practically no concrete guideline for application to train people that deal with emergency, disaster, and DDM in general.

In contrast, Naturalistic Decision Making is a field that has been concerned with studying decision making in the "wild," in contexts in which proficient decision makers draw conclusions from realistic cases and scenarios that are relevant to their experience and knowledge (Lipshitz, Klein, Orasanu, & Salas, 2001). Those who study naturalistic decision making are confronted with serious challenges. They often study large groups and real decision makers in complex decision situations. Real-world decision makers typically confront many uncertainties about the available options; they have inadequate information about their options; they rarely know the likely costs and benefits, and the value trade-offs they entail. Although one could expect decision makers in the real world to have clear goals and to promote those goals with their decisions, the reality is that decision making in real-world situations is seldom rational, and in fact it

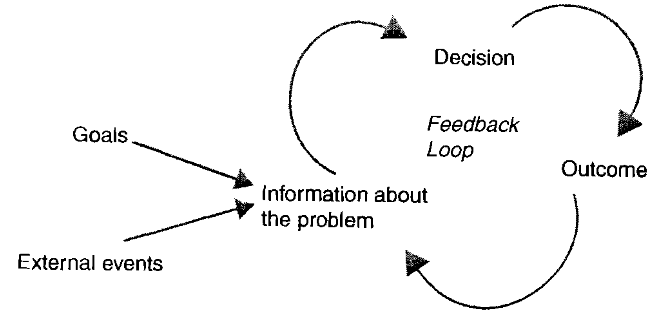

Figure 8.2. A closed-loop view of decision making.

is often hard to understand what constitutes rational choice under such conditions. Thus, these studies have multiple limitations. Realistic studies demonstrate only particular examples of decision-making situations from which general predictions and inferences are hard to derive.

Recently in the behavioral decision sciences, there has been some progress to study decision making in "less ideal" conditions, in situations where there are repeated decisions and uncertainty about the outcomes and probabilities of events. This paradigm, named Decisions from Experience (DFE), has revealed that many assumptions of the perceptions of risk and human preferences in one-shot static decisions do not hold for DFE (Hertwig et al., 2004). Although this work brings a very optimistic view of the application of behavioral decision sciences to DDM situations, there are many years of research ahead to clearly understand all the aspects involved in these tasks in a systematic way (Lejarraga, Dutt, & Gonzalez, in press).

I conceptualize DDM as a control process; a closed-loop learning process where decisions are influenced by goals, external events, and previous decisions. Thus in my view, decision making is a learning process where decisions are made from experience and are feedback-dependent (feedback here is the association that a human makes between actions and their outcomes) (see Figure 8.2). In this paradigm, alternatives are not presented at the same time, but rather unfold over time, and decision making is a learning loop: Decisions depend on previous choices as well as on external events and conditions.

DFE is, very likely, the only method by which decisions can be made in dynamic conditions. In fact, a recent study demonstrates that as the complexity of a problem increases, people prefer to make decisions from experience rather than interpreting the given probabilities and outcomes of a one-shot decision (Lejarraga, 2010). My research has focused on the study of DFE in DDM situations. I have used quite diverse research approaches, different from those used in the behavioral decision sciences and those used in naturalistic decision making, including laboratory experiments with complex interactive decision-making games, and computational cognitive modeling. I believe that these approaches improve the degree of convergence between the traditional behavioral decision sciences experiments and the naturalistic decision-making studies. In what follows I summarize what I have learned in the past years from using these two approaches and present the practical and concrete lessons for application to training decision makers that deal with emergency, disaster, and dynamic environments in general.

Learning in DDM

Research on learning in and about dynamic systems indicate that humans remain suboptimal decision makers even after extended practice and after being given unlimited time and performance incentives (Diehl & Sterman, 1995; Sterman, 1994); that is, humans do not always and frequently do not improve their decisions from experience (Brehmer, 1980). One main impediment to learning in dynamic tasks is the difficulty in processing feedback, particularly delayed feedback (Brehmer, 1992; Sterman, 1989). However, many other difficulties have also been documented, including our abilities to deal with time constraints, high workload, and limitations in our inherent cognitive abilities (Gonzalez, 2004, 2005a; Gonzalez, Thomas, & Vanyukov, 2005).

Various accounts have been proposed regarding how a human learns in dynamic systems (for summaries of these, see Busemeyer, 2002; Gonzalez, 2005b). One theory is that specific instances are used to control dynamic systems (Dienes & Fahey, 1995). This learning model was based on two cognitive mechanisms that compete every time someone encounters a decision-making situation: an algorithm and a set of context-action exemplars. The algorithm is a general heuristic or rule that one uses in a novel situation; the context-action exemplars are discrete representations of knowledge that are called "instances," a name derived from Logan's (1988) instance theory of automatization. In this model, an implicit assumption is that a decision maker stores actions and their outcomes together in memory and retrieves them on the basis of their similarity to subsequently encountered situations.

Another theory of learning is proposed by the connectionist approach, in which decision making is built from interconnected units (Gibson, Fichman, & Plaut, 1997). This model of learning in DDM is based on the control theory approach proposed by Brehmer (1990) and was implemented computationally via neural networks. This theory assumes that decision makers use outcome feedback to form two submodels: the *judgment submodel* that represents how the decision maker's actions affect outcomes, and the *choice submodel* that represents which actions are taken to achieve desired outcomes. The judgment submodel learns by minimizing the differences between the outcomes it predicts and the outcomes received from feedback, whereas the *choice submodel* learns by minimizing the differences between the alternatives predicted by the judgment model and the alternatives actually selected. This model provides a good account of individuals' learning in dynamic situations and their transfer of knowledge to novel situations (Gibson, 2000).

A third theory is the Instance-Based Learning Theory (IBLT) (Gonzalez & Lebiere, 2005; Gonzalez et al., 2003). IBLT was developed to reproduce decision-making behavior in dynamic tasks. IBLT characterizes learning by storing in memory a sequence of action-outcome links produced by experienced events through a feedback-loop process of human and environment interactions. This process increases knowledge and allows decisions to improve as experience accumulates in memory. IBLT assumes the following components. *Instances* are examples of choices that are stored in memory. Each instance contains cues about the situation in which a decision was made, the decision itself, and the subsequent outcome. Situational cues are relevant in dynamic environments because situations are continuous and variable, and not all experiences are informative for future choice situations. Learning resides in the *Activation* (e.g., frequency and recency) of experienced choices and outcomes. IBLT assumes that the instances experienced by the decision maker are activated in memory as a function of their occurrence. More recent and frequent instances are more active in memory than less recent and less frequent ones. Choice situations are never equivalent in dynamic tasks as environments change over time. Thus, past experiences are not necessarily directly applicable in new conditions. A *similarity* rule, defined on situational cues, is specified to evaluate the resemblance of previous situations with respect to the current situation being evaluated. Finally, IBLT uses *blending* as a mechanism to average the value of different observed outcomes in previous similar situations (Lebiere, 1999). The value of an option is the addition of the subjective value of each possible outcome weighted by its subjective likelihood.

The three learning models summarized earlier incorporate at least two common characteristics; all three models take into account the need for two forms of learning: *explicit* (i.e., decision making based on rules of action) and *implicit* (i.e., decision making based on context-based knowledge and recognition). There is some evidence that individuals who have completed a dynamic task are not always aware of the task structure (i.e., their knowledge is implicit), which suggests that the knowledge they acquired was not in the form of rules about how the system works (Dienes & Fahey, 1995). Often, individuals performing DDM tasks are unable to describe the key elements of the task or verbalize the ways in which they make decisions (Berry & Broadbent, 1987, 1988). Second, these models rely on a similarity process that determines the applicability of accumulated experiences to familiar situations. Research in analogical reasoning has demonstrated the high relevance of analogy to learning and decision-making processes (Kurtz, Miao, & Gentner, 2001; Medin, Goldstone, & Markman, 1995).

In summary, there are well-documented difficulties when humans make decisions in dynamic systems. Humans remain suboptimal or learn very slowly, often due to feedback delays, time constraints, and the cognitive workload required by these environments. To be able to understand and improve training protocols and guidelines, one needs to first understand how humans make decisions in these tasks. Fortunately, the similarities across the most prominent theories of learning in DDM help converge in some issues. All models agree in that humans learn facts, cause-and-effect knowledge related to the context, and none of the models present the main form of learning as being structural knowledge or rules. Also, all of the models agree on the relevance of some form of recognition of familiar patterns from past experience; that decisions are made from experience by retrieving a solution from similar situations in past experience.

Because IBLT is the basis for my proposed guidelines to training decisions from experience with decision-making games, we explain IBLT's principles and its formulations in more detail next.

Instance-Based Learning Theory

An instance in IBLT consists of environmental cues (the situation), the set of actions that are applicable to the situation (the decision), and the evaluation of the goodness of a decision in a particular situation (the utility) (see Figure 8.3). Thus, the accumulation of instances involves storing situation-decision-utility (SDU) triplets in memory. Figure 8.4 presents the generic IBLT process by which decisions are made in an interactive environment, consisting of the Recognition, Judgment, Choice, Execution, and Feedback steps. When faced with a particular decision situation, people are likely to retrieve similar SDUs (SDUs with similar situations) from memory (Recognition step). In a typical situation (situation similar to past SDUs), the expected utility of an action is calculated by combining the utility of similar instances retrieved from memory (a procedure called Blending). In atypical situations, however, people fall back on heuristics in their evaluation of expected utility of the action. The evaluation of the expected utility of a decision in a situation is done in the Judgment step. Alternative actions are evaluated sequentially, and after each evaluation, the decision of whether or not more alternative should be evaluated is determined by a *necessity mechanism*. Necessity may be subjectively determined by the decision maker's own preferences or by exogenous factors such as lack of time or changes in the environmental conditions. The alternative with the highest utility among the evaluated alternatives is then selected (the Choice step) and executed (the Execution step), changing the environment and noting which SDU was executed in memory. Once a decision has been made, the outcome of the decision is

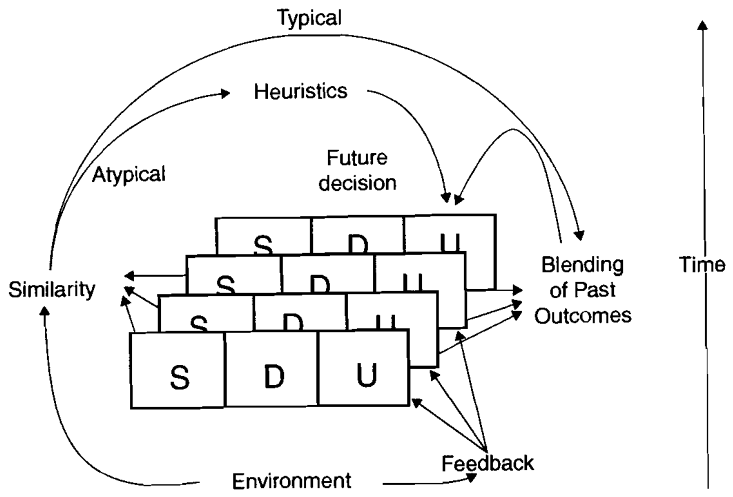

Figure 8.3. Instance-based learning. Instances (SDUs) accumulate over time. At a decision point, a situation is compared to past instances, and based on its similarity to past instances, these are reused and blended to determine the expected utility of the current decision situation. Good instances are reinforced through feedback from the environment.

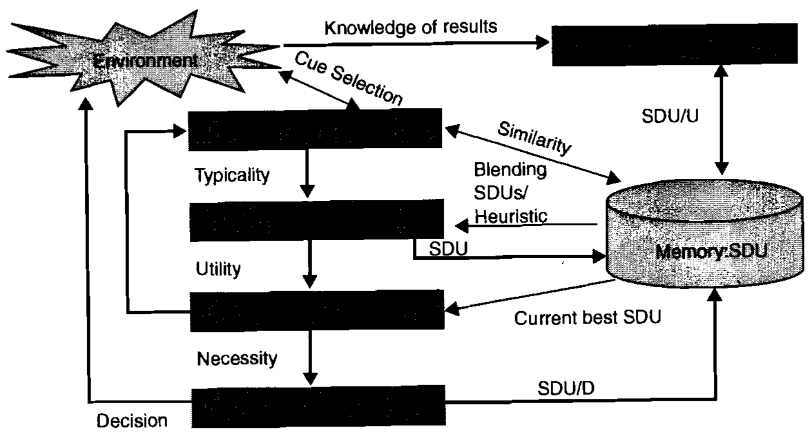

Figure 8.4. The IBLT process.

used as feedback to modify the utility value of the original SDUs (Feedback step).

The computational implementation of the IBLT relies on several mechanisms proposed by the ACT-R cognitive architecture (Anderson & Lebiere, 1998), notably the Activation mechanism. Activation in ACT-R is a value assigned to each "chunk" (i.e., instance) that reflects the estimate of how likely the chunk would be retrieved and the speed of retrieval for the chunk. The activation of a chunk reflects the frequency and recency of use of the chunk and the degree to which the chunk matches a context (i.e., the extent to which a given chunk is similar to previously presented chunks in each chunk position). ACT-R architectural mechanisms underlie and constrain the various steps of the IBLT process. Learned instances and their accessibility provide for the accumulation of instances in memory; similarity-based retrieval from memory drives the recognition process, and so on. We have used the computational implementation of IBLT to confirm and predict human performance in many tasks. IBLT's

computational models' data often agree with human data (e.g., Gonzalez & Lebiere, 2005; Gonzalez et al., 2003; Martin, Gonzalez, & Lebiere, 2004).

Training and Decision-Making Games

Given the similarity of the different theories of learning in DDM and the known difficulties of human learning in these environments, the question we address here is: What are the recommendations that can be drawn from these theories to address the learning difficulties in DDM tasks? Because each DDM situation is unique and the reusability of good past instances from memory depends on the similarity between the current situation and those stored in memory, disaster and emergency situations present a challenge to the improvement of decision performance over time: There are not many similar instances stored in our memories. In addition, the unpredictability of events in dynamic situations makes it difficult to determine the timing of major events, with possibly severe consequences. Thus, our main concern in this section is to present guidelines and suggestions from the IBLT to help individuals become alert and able to perform as best as possible in disaster and emergency situations.

In my view, an essential way to achieve successful training for dynamic tasks is the use of Decision-Making Games (DMGames). DMGames are graphical models (abstractions of reality) used for experimentation with human decision makers. My concept of DMGames has evolved from that of Microworlds, a term commonly used in the DDM field (Brehmer & Dörner, 1993; Gonzalez et al., 2005; Turkle, 1984), and the more recent developments of serious games and serious games initiatives (Cannon-Bowers & Bowers, 2010). Many disciplines are now adopting simulations and games in research, including engineering (Foss & Eikass, 2006), business and management (Zantow, Knowlton, & Sharp, 2005), medicine (Bradley, 2006), and political science (Kelle, 2008; Mintz, Geva, Redd, & Carnes, 1997). Microworlds were developed to study DDM, and through the years, technological advancements have allowed for the development of more graphical and interactive tools for research, which are also more fun. DMGames may incorporate temporal dependencies among system states (i.e., dynamics), feedback delays, nonlinear relationships among system variables, and uncertainty (in the form of exogenous disturbances). They are interactive and allow repeated, interrelated decisions. They also may incorporate external events and time pressure. Thus, DMGames are essential to compress time and space in the laboratory setting. They reproduce difficult learning in conditions with rare, novel events and unpredictable timing, such as in disaster and emergency situations. DMGames may speed up learning and help people acquire the instances they cannot acquire from the real world. DMGames may help a human acquire the skills needed to be alert and become adaptable in the real world.

The training recommendations that follow are based on the use of DMGames in laboratory experiments where we have manipulated experience (type of instances stored) and the dynamic conditions on which decisions are made, such as timing, workload, and feedback delays. All recommendations come from IBLT and the empirical work done through the years to test IBLT's predictions on ways to speed up learning and facilitate prompt adaptation to novel and rare situations.

Slow Is Fast

IBLT recommends that *slow is fast* when it comes to adapting to time-constrained environments (Gonzalez, 2004). In a dynamic resource allocation task like disaster and emergency situations, it has been demonstrated that individuals trained on a task at a slower pace were able to adapt more successfully to greater time constraints, compared to those who only trained under high time constraints, regardless of exceedingly large number of practice sessions given to those trained under time constraints. Thus, a few slow practice sessions were more beneficial than a larger number of fast practice sessions because they enable people to acquire more complex and useful knowledge.

Thirty-three graduate and undergraduate college students recruited from local universities were randomly assigned to either the fast or slow condition group. The Water Purification Plant simulation™ was used for this study. The goal in this task is to distribute all the water in the system on time and under time constraints by activating and deactivating pumps. The environment is opaque, so the user is uncertain about some key variable values. For example, water appears in the system according to a scenario defined by the experimenter and unknown to the user. The environment changes both autonomously and in response to the user's decisions. Because a maximum of five pumps can be activated at any one time, the decision maker's actions are interrelated. This task translates directly to disaster and emergency situations because it involves time pressure, limited resources, incomplete knowledge about the situation, unexpected events, and the need to coordinate efforts to meet the demands. Participants did the task over three consecutive days. Under the fast condition, each simulation trial lasted eight minutes. Participants under this condition completed eighteen trials over the three-day period (six trials/day). Under the slow condition, each simulation trial on the first two days lasted twenty-four minutes (two trials/day), whereas each trial on the last day lasted eight minutes (six trials). For all participants, the first two days were training days and the last day was the test day. The results show that slow training led to better performance than fast training on day three with fast performance for both groups.

IBLT explains this finding in several ways. First, learning at a slower pace results in more alternatives being considered within each recognition-choice cycle. More instances are considered and stored in memory during each evaluation of possible courses of action, given more time within each cycle. Further, a greater chance of finding an optimal action exist, given that more alternatives are being evaluated. Second, when an individual is trained at a slower pace and then asked to perform at a faster pace, she can retrieve and rely on the larger and possibly more diverse set of instances in memory, which may be applied during a time-constrained condition. In contrast, when someone is trained under high time constraints, there is no chance of getting more instances representing different alternatives in memory during training. Performance under time constraints is thus limited to the possible retrieval of a few, selected sets of instances in memory, leading to poorer performance.

Less Workload Helps Adaptation

Similar to the results on time constraints, it has been found that individuals who trained on a task under low workload were able to perform more accurately during transfer with a high workload than those trained under high-workload conditions all along (Gonzalez, 2005a).

Fifty-one students were recruited and assigned to conditions that differed in the amount of workload (number of simultaneous tasks performed at the same time) during training. In the high-workload condition, participants had to complete the same Water Purification Plant simulation™ as in Gonzalez (2004), but also had to simultaneously perform two additional, independent tasks. This group was contrasted to a low-workload condition in which participants performed the task at the same pace but with no additional independent tasks. Participants ran the DMGame on three consecutive days. The first two days were the training days, during which participants worked under one of the workload conditions, and a third day, during which all participants performed the same DMGame under workload.

Similar to the results in Gonzalez (2004), the findings indicate that high task workload during training hindered performance and transfer compared to low-workload training. Thus, these two studies demonstrated that it is not a good idea to train individuals in conditions "close to the real conditions" when it comes to workload and time constraints. Slower pacing and low workload are best during training for people to perform well in fast and high-workload tasks at test. Once again, from the IBLT perspective, this study demonstrates the effects of workload on the recognition-choice cycle, which under workload inhibits the generation of instances that can then be reused at test.

Heterogeneity Helps Adaptation

In several studies together with my colleagues (Brunstein & Gonzalez, in press; Gonzalez & Madhavan, 2011; Gonzalez & Quesada, 2003; Madhavan & Gonzalez, 2010), I have found that the variation in the situations that people confront during training influences how fast and how well they learn to adapt to novel and unexpected situations, with higher variation leading to better transfer. For example, Gonzalez and Quesada (2003) demonstrated the influence of the similarity of past decisions on future decisions. In this study, decisions became increasingly similar with task practice, but the similarity depended on the interaction of many task features rather than by any single task features. This study demonstrated the relevance of similarity of situations in the IBLT process.

In another study, Madhavan and Gonzalez (2010) used a luggage-screening task to investigate the effects of similarity of experiences during the learning process. In such a task, each piece of luggage could have distractors and targets. Targets could be one of several types: knives, guns, glass objects, liquids, and so on. This task resembles disaster and emergency situations in its key features. For example, disaster responders have to discriminate between patients who would profit most from treatment based on their symptoms and available resources (targets) and patients who would not (distractors). As a dynamic decision-making task, the condition of a patient might change over time, resulting in a different category membership. My work with Brunstein and Madhavan has demonstrated that diverse and heterogeneous conditions of training lead to better adaptation to novel and unexpected situations. For example, training conditions in which the targets change (targets can be targets in some trials and distractors in others) or training conditions in which targets are drawn from diverse categories result in better adaptation to novel and rare situations, compared to training in consistent conditions (in which the targets set is constant and are items from a narrow set of categories).

Heterogeneity in IBLT increases the chance that a similar instance will be retrieved in future choices. The heterogeneity is defined by the multidimensional space of the cues that form the situation of the instances. The more diverse those instances are, the better the chances will be of finding similarity to past decisions when a novel instance is confronted. Thus, the greater the possibility of reusing well-matching stored instances instead of merely general heuristics.

Feedforward Helps Adaptation

I have found that knowledge of results is not enough for improving learning and adaptation (Gonzalez, 2005b). The provision of outcome feedback was an inferior way of trying to aid learning and performance compared to the viewing of a highly skilled decision maker (e.g., a form of feedforward). Participants were trained in one of four groups during the first day and transfer to the same DDM conditions in the second day. During training, participants were assigned to one of the following four conditions. Control participants only received feedback about the outcome in the task. Process feedback participants received outcome feedback broken down in multiple steps and pieces. Participants in the self-exemplar group ran one trial under the control condition and then viewed a replay of their own performance. Finally, participants in the expert-exemplar condition ran one trial under the control condition and then replayed the trial of a highly skilled participant. On the second day, all participants were asked to perform the DDM task without any feedback (the detailed feedback, self-exemplar, and expert-exemplar aids were removed).

Results showed that the performance of all groups improved across trials. However, in the early stages of learning, people with the detailed-outcome feedback condition actually showed poorer learning than the control participants. The self-exemplar and the expert-exemplar conditions showed superior performance compared to the detailed-outcome feedback condition; however, when the aid was removed in the testing period, the only treatment that outperformed the control condition was the expert-exemplar group. The expert-exemplar group began to outperform the other groups midway through the training trials, and this superior performance continued throughout the testing period, even after the intervention of seeing the expert performance had been removed.

The results of the detailed feedback are similar to those found in Gonzalez (2005b). Additional detailed feedback increased the workload during training. Because this task was a real-time task, the time left to evaluate more alternatives was used instead in evaluating the detailed feedback. Thus, although presumably the utility of the instances would increase, the number of instances decreased compared to a condition with no detailed feedback. This is clear in the first sessions of the task, in which performance decreased for the self-exemplar group. The effects of self-exemplar and expert-exemplar are also clear from IBLT. In both conditions, people are acquiring more instances without the need of executing the task. The self-exemplar instances, however, are less effective than the expert-exemplar instances. Reflecting on one's own performance only reinforces poor instances, whereas reflecting on expert's performance reinforces instances with higher utility.

Summary and Conclusions

Behavioral scientists, particularly those with interest in human decision making, should pay more attention to the process and skills needed to make decisions in dynamic situations. The majority of research in decision sciences has focused on one-shot (open-loop) decisions that present people with description of hypothetical problems of choice. Very little has been done in the decision sciences to understand the formation of alternatives in the first place and the effects of "closing the loop"; that is, making decisions from experience by reusing past decisions and outcomes.

My research has focused on the study of decisions from experience in dynamic situations. My belief is that experience is the most likely method by which people make decisions in dynamic conditions. I see decision making in dynamic tasks as a learning loop, where decisions are made from experience and are feedback-dependent. In this chapter, I explained current research of learning theories that are directly relevant to DDM. I concentrated particularly on the explanation of IBLT, a theory that has been implemented computationally and from which several training guidelines have been derived. I have illustrated how research on DDM using DMGames and the IBLT can be used to generate practical and concrete training principles. The factors of DDM that are of special relevance for emergency training are time constraints, workload, similarity or diversity of experiences, and types of feedback. For these factors, we found what best prepares people for uncertain and novel situations is to train under conditions that foster skill acquisition and a deeper understanding of the situations confronted.

Some examples of the principles for emergency training derived from DDM research are the diversity of practice and the *slow is fast* principles. Greater diversity of instances during training helped individuals detect unknown and novel targets more accurately at transfer than those trained with a consistent set of targets. In addition, slower training led to better performance than fast training. Thus, training for high-time-pressure tasks is more effective if performed in a slower path before releasing learners to realistic conditions. In laboratory studies, we have demonstrated how these guidelines of training result in better performance and adaptation to unexpected and uncertain conditions. Thus, I suggest that these guidelines could be used to design training protocols for decision makers that have to be prepared to deal with unexpected conditions and possible emergencies in their daily activities.

Acknowledgments

This is a summary of research performed through many years, and thus I owe recognition to a large number of organizations that have trusted and supported my work, including the Army Research Laboratory, the Office of Naval Research, the National Science Foundation, and the Army Research Office, among others. I also owe recognition to doctoral students, postdoctoral fellows, and other research staff of my dynamic decision making laboratory. To all of them, thank you.

References

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Berry, D. C., & Broadbent, D. E. (1987). The combination of explicit and implicit learning processes in task control. *Psychological Research, 49*(1), 7–15.

Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology, 79*, 251–272.

Bradley, P. (2006). The history of simulation in medical education and possible future directions. *Medical Education, 40*, 254–262.

Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica, 45*, 223–241.

Brehmer, B. (1990). Strategies in real-time, dynamic decision making. In R. M. Hogarth (Ed.), *Insights in decision making* (pp. 262–279). Chicago: University of Chicago Press.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica, 81*(3), 211–241.

Brehmer, B., & Dörner, D. (1993). Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior, 9*(2–3), 171–184.

Brunstein, A., & Gonzalez, C. (in press). Preparing for novelty with diverse training. *Applied Cognitive Psychology.*

Busemeyer, J. R. (2002). Dynamic decision making. In N. J. Smelser & P. B. Baltes (Eds.), *International encyclopedia of the social and behavioral sciences, Vol. 6* (pp. 3903–3908). Oxford: Elsevier Press.

Cannon-Bowers, J., & Bowers, C. (2010). *Serious game design and development technologies for training and learning.* Hershey, PA: IGI Global.

Diehl, E., & Sterman, J. D. (1995). Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes, 62*(2), 198–215.

Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory and Cognition, 21*(4), 848–862.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors, 4*, 59–73.

Foss, B. A., & Eikaas, T. I. (2006). Game play in engineering education: Concept and experimental results. *International Journal of Engineering Education, 22*(5), 1043–1052.

Gibson, F. P. (2000). Feedback delays: How can decision makers learn not to buy a new car every time the garage is empty? *Organizational Behavior & Human Decision Processes, 83*(1), 141–166.

Gibson, F. P., Fichman, M., & Plaut, D. C. (1997). Learning in dynamic decision tasks: Computational model and empirical evidence. *Organizational Behavior and Human Decision Processes, 71*(1), 1–35.

Gonzalez, C. (2004). Learning to make decisions in dynamic environments: Effects of time constraints and cognitive abilities. *Human Factors, 46*(3), 449–460.

Gonzalez, C. (2005a). The relationship between task workload and cognitive abilities in dynamic decision making. *Human Factors, 47*(1), 92–101.

Gonzalez, C. (2005b). Decision support for real-time dynamic decision making tasks. *Organizational Behavior & Human Decision Processes, 96*, 142–154.

Gonzalez, C., & Lebiere, C. (2005). Instance-based cognitive models of decision making. In D. Zizzo & A. Courakis (Eds.), *Transfer of* *knowledge in economic decision-making*. New York: Macmillan (Palgrave Macmillan).

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*(4), 591–635.

Gonzalez, C. & Madhavan, P. (2011). Diversity during training enhances detection of novel stimuli. *Journal of Cognitive Psychology, 23*(3), 342–350.

Gonzalez, C., & Quesada, J. (2003). Learning in dynamic decision making: The recognition process. *Computational and Mathematical Organization Theory, 9*(4), 287–304.

Gonzalez, C., Thomas, R. P., & Vanyukov, P. (2005). The relationships between cognitive ability and dynamic decision making. *Intelligence, 33*(2), 169–186.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15*(8), 534–539.

Kelle, A. (2008). Experiential learning in an arms control simulation. *PSOnline, April*, 379–385.

Klein, G. A. (1989). Recognition-primed decisions. In W. B. Rouse (Ed.), *Advances in Man-Machine Systems Research, Vol. 5* (pp. 47–92). Greenwich, CT: JAI Press.

Kurtz, K. J., Miao, C., & Gentner, D. (2001). Learning by analogical bootstrapping. *The Journal of Learning Sciences, 10*(4), 417–446.

Lebiere, C. (1999). *Blending: An ACT-R mechanism for aggregate retrievals*. Paper presented at the Sixth Annual ACT-R Workshop at George Mason University.

Lejarraga, T. (2010). When experience is better than description: Time delays and complexity. *Journal of Behavioral Decision Making, 23*(1), 100–116.

Lejarraga, T., Dutt, V., & Gonzalez, C. (in press). Instance-based learning in repeated binary choice. *Journal of Behavioral Decision Making*.

Lipshitz, R., Klein, G., Orasanu, J., & Salas, E. (2001). Taking stock of naturalistic decision making. *Journal of Behavioral Decision Making, 14*(5), 331–352.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95*(4), 492–527.

Madhavan, P. & Gonzalez, C. (2010). The relationship between stimulus-response mappings and the detection of novel stimuli in a simulated luggage screening task. *Theoretical Issues in Ergonomics Science, 11*(5), 461–473.

Martin, M. K., Gonzalez, C., & Lebiere, C. (2004). Learning to make decisions in dynamic environments: ACT-R plays the beer game. In M. C. Lovett, C. D. Schunn, C. Lebiere, & P. Munro (Eds.), *Proceedings of the Sixth International Conference on Cognitive Modeling, Vol. 420* (pp. 178–183). Pittsburgh, PA: Carnegie Mellon University/University of Pittsburgh; Lawrence Erlbaum Associates Publishers.

Medin, D. L., Goldstone, R. L., & Markman, A. B. (1995). Comparison and choice: Relations between similarity processing and decision processing. *Psychonomic Bulletin and Review, 2*(1), 1–19.

Mintz, A., Geva, N., Redd, S. B., & Carnes, A. (1997). The effect of dynamic and static choice sets on political decision making: An analysis using the decision board platform. *The American Political Science Review, 91*(3), 553–566.

Rapoport, A. (1990). *The meaning of the built environment: A nonverbal communication approach*. Tucson: University of Arizona Press.

Sterman, J. D. (1989). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science, 35*(3), 321–339.

Sterman, J. D. (1994). Learning in and about complex systems. *System Dynamics Review, 10*, 291–330.

Turkle, S. (1984). *The second self: Computers and the human spirit*. London: Granada.

Zantow, K., Knowlton, D. S., & Sharp, D. C. (2005). More than fun and games: Reconsidering the virtues of strategic management simulations. *Academy of Management Learning & Education, 4*(4), 451–458.

CHAPTER 9

Adaptive Tutoring Technologies and Ill-Defined Domains

Collin Lynch, Kevin D. Ashley, Niels Pinkwart, and Vincent Aleven

Introduction

Consider the following problems:

Pw: Given a fleet of trucks $F = \{f1, f2, \ldots, fn\}$ and a shipment of supplies S packed in N crates each weighing $\leq T$ tons, how long would it take to ship the boxes from Bagram AFB to a forward operating base using one of N predefined routes?

Pi: Arrange a logistical distribution system that will provide a sufficient supply of medical and military stores for firebases and patrol units in Helmand province.

Both are standard logistical problems that any hapless staff officer might face on a given day. The former, however, is quite *well-defined* in that the problem is fairly constrained, the relevant knowledge is clear and available to the solver, and the answer is unambiguous. It could easily be approached using standard production rules as employed in GPS (Newell & Simon, 1995) or with a pattern-matching or schema-driven approach (Greeno, 1976).

The latter problem, by contrast, is *ill-defined*. Much of the necessary information is unavailable or underspecified: What supplies are required? How often must they be replenished? Who is available to make the shipments? And so on. To solve the problem, our hapless staffer will be forced to articulate, and answer, a number of related questions, frame and define missing criteria (e.g., what is "sufficient"), and be ready to defend their decisions after the fact. In short, he or she must *recharacterize* the problem to solve it, and that process will define the solution. In doing so, our officer will need to draw on a wide range of information, from the shelf-life of medical supplies, to the status of roads in Helmand, to the local governor's taste in cigars.

Our focus in this chapter is on ill-defined problems and domains. What does it mean for a problem or domain to be ill-defined and what are the implications for adaptive training technologies? We begin by clarifying what we mean by "problem,"