

2003

Learning in Dynamic Decision Making: The Recognition Process

Cleotilde Gonzalez

Carnegie Mellon University, conzalez@andrew.cmu.edu

Jose Quesada

University of Colorado at Boulder

Follow this and additional works at: <http://repository.cmu.edu/sds>

This Article is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Social and Decision Sciences by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.



Learning in Dynamic Decision Making: The Recognition Process

CLEOTILDE GONZALEZ

Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA
email: conzalez@andrew.cmu.edu

JOSE QUESADA

Institute of Cognitive Science, University of Colorado, Boulder, Boulder, CO 80309-0344, USA

Abstract

The apparent difficulty that humans experience when asked to manage dynamic complexity might be related to their inability to discriminate among familiar classes of objects (i.e., flawed recognition). In this study we examined the change in individuals' recognition ability, as measured by the change in the similarity of decisions they made when confronted repeatedly with consistent dynamic situations of varying degrees of similarity. The study generated two primary findings. First, decisions became increasingly similar with task practice, a result that suggests gradually improving discrimination by the participants. Second, the similarity was determined by the interaction of many task features rather than individual task features. The general principles highlighted by this study are applicable to dynamic situations. For example, with practice, decision makers should be able to learn to identify the time at which to intervene to achieve the maximal effect during dynamic decision making.

Keywords: cognitive representation, dynamic decision making, recognition

1. Introduction

On-line environments can reduce the lag time between decisions and their effects because such environments increase the number of times individuals can cycle around a learning loop and, presumably, facilitate improved decision making by enabling people to accumulate experience more quickly. We use the term *presumably* intentionally because little is known about how individuals rely on this accumulated experience to make decisions in dynamic situations, and even less is known about how these decisions improve with practice. Research has shown that learning to perform dynamic decision-making (DDM) tasks is difficult, mainly because of the dynamic complexity of these tasks (Diehl and Sterman, 1995). Dynamic complexity comprises time delays, feedback loops, stock and flow structures, and nonlinearities (Sterman, 2000). To improve our understanding of how individuals learn to make better decisions in a dynamic system, we investigated the human process of recognition.

Gonzalez et al. (2003) propose recognition as a pre-condition for decision making in dynamic environments. The concept of recognition—the ability to discriminate among familiar classes of objects (Simon and Langley, 1981)—is often used by researchers in

the field of naturalistic decision making (NDM), who claim that experts are accurate and quick to make decisions in ill-structured situations because they use their *experience* to recognize a situation, and then make decisions that have worked previously (Klein et al., 1993; Zsombok and Klein, 1997). When used in decision-making theories to describe the process of deciding if a specific event has occurred previously, recognition is closely related to the concept of categorization discussed in cognitive psychology (see Gonzalez et al., 2003).

In this study we examined the recognition process and its effect on the ability of individuals to recall relevant past experiences that they can use to achieve improved task performance with practice. We believe that individuals' ability to manage dynamic complexity could be improved if more were known about the skills that would enable them to take advantage of past experience. Specifically, our study focused on the following questions: (1) Do decision makers recognize situations and, if so, does decision-making performance correlate with an increase in similarity between current and past decisions? (2) Which factors influence recognition and how do these factors change as performance improves? We addressed these questions by analyzing how individuals performing a DDM task used experience to improve performance.

2. Learning in Dynamic Decision-Making Tasks

Simon and Langley (1981) defined learning as "a process that modifies a system as to improve, more or less irreversibly, its subsequent performance of the same task or of tasks drawn from the same population." DDM requires that multiple and interdependent decisions be made in an environment that changes autonomously and in response to a decision maker's actions (Brehmer, 1990). Extensive knowledge and practice are necessary for individuals to acquire efficient control of dynamic systems (Kerstholt and Raaijmakers, 1997). DDM involves the formation of feedback loops through which the results of individuals' actions define the situations that the individuals encounter in the future. Each new situation, in turn, alters future decisions (Serman, 2000). Thus, one might expect that learning and improved DDM would occur simultaneously. However, studies have shown that decision makers who learn the input and output signals required to attain control of a system do not necessarily exhibit improved performance (Diehl and Serman, 1995; Dienes and Fahey, 1995). Many individuals fail to learn even when provided outcome feedback (Diehl and Serman, 1995), and do not show good transfer abilities even when the transfer task is similar to the learning task (Gibson et al., 1997).

Although outcome feedback is integral to the learning process, decision makers may fail to use feedback because they do not *recognize* the situation that produced the outcome. During DDM, decision makers generally receive delayed feedback, which is characterized by a temporal causal disconnect between an action and its outcome (Brehmer, 1995; Serman, 1994). Therefore while completing DDM tasks, decision makers must retrieve from memory the situations and actions that may have produced a given outcome. Engaging in this process of reflection on prior actions and their apparent results may help individuals improve their DDM performance (Gibson et al., 1997).

Gonzalez et al. (2003) recently have claimed that recognition plays a key role in individuals' ability to use outcome feedback, in the form of prior decisions and their outcomes, to improve DDM performance. Called Instance-Based Learning Theory (IBLT) incorporates the psychological theories of memory and categorization, and specifies five learning mechanisms that are crucial to skill development in DDM: (1) instance-based knowledge, (2) recognition-based retrieval, (3) adaptive strategies, (4) necessity, and (5) feedback updates. Two of these learning mechanisms, instance-based knowledge and recognition-based retrieval, are relevant to the study presented here.

According to IBLT, every decision-making cycle can be described by an instance, which comprises the *situation* in which a decision is made, the *decision* made, and the expected *utility* of the decision in that situation (Situation, Decision, Utility—SDU). Recognition-based retrieval enables a decision maker to identify the instances of highest utility from memory by evaluating the *similarity* between the situation under assessment and earlier instances stored in memory. A decision reached by successful recognition-based retrieval leads to the creation of a new instance with a particular outcome, and the utility of all similar, prior instances is upgraded to reflect this most recently experienced utility. According to this theory, then, the efficient use of feedback to improve performance is dependent on a decision maker's ability to recognize similarity between current and past situations.

Similarity plays an important role in behavioral theories and, particularly, in IBLT (Gonzalez et al., 2003). Medin et al. (1995) have suggested that decision making entails the use of similarity judgments and that phenomena associated with decision making is based on similarity judgments. To further explore these claims, Gonzalez et al. (2003) tested a cognitive model of DDM by using similarity as a judgment strategy to determine the utility of decisions. According to IBLT, similarity is a heuristic that decision makers under time constraints use in the absence of complete information. Pioneering studies of the dynamics of similarity (Tversky, 1977) have defined similarity as a metric of matching and mismatching features of current and past decision-making situations. In this study we tested the hypothesis that DDM performance is related closely to the ability to recognize similar stimuli. We designed the experiments to evaluate whether decision makers in DDM systems reuse past decisions, whether increased similarity between current and past situations leads to performance improvement, whether similarity is a reliable predictor of future performance, and whether features of the task influence recognition and fluctuate during task learning.

3. Research Method

We collected data by using an on-line system that we designed to reproduce the structure and complexity of a dynamic real-world task. Systems such as this are most commonly referred to as *microworlds* (Omodei and Wearing, 1995). Study participants ran the microworld on several days while we monitored the individual decisions made during each run (trial) and evaluated overall performance. We first determined the similarity of decisions made across trials, and then analyzed these data in terms of performance improvements.

3.1. *Water Purification Plant (WPP)*

The microworld that we designed is called the Water Purification Plant (WPP). WPP is a resource management task in which participants decide how to allocate limited resources while working under deadlines for the completion of the task. As demonstrated to be important in other DDM environments, a primary characteristic of the WPP task is its requirement that participants make multiple and interdependent real-time decisions in an autonomously changing environment (Brehmer and Dorner, 1993). WPP also features the basic building blocks of a complex dynamic system: time delays, stocks and flows, and feedback processes (Sterman, 2000). A complete description of the WPP simulation and how it incorporates these fundamental characteristics of DDM is described elsewhere (Gonzalez et al., 2003), but a brief description of the task appears below.

Researchers designed WPP by using field research aimed at building decision-making support for operators in the United States Postal Service (USPS) (Lerch et al., 1997). The USPS uses a network of sorting machines to sort mail enroute to its final destination. Operators under dispatch deadlines process incoming mail by activating sorting machines. WPP is an isomorph of this real-world task. A WPP participant plays the role of a water purification plant operator. In this plant, water (mail) enters different purification tanks (sorting machines), which the operator then activates while attempting to meet a set of deadlines.

Figure 1 is a screenshot of the WPP, which contains 22 tanks with 2 pumps per tank; a set of deadlines is visible on the right of the screen. The tanks are connected by pipes that indicate the path the water traverses to reach the deadline. The set of connected tanks is called a *chain*, and the length of the chain dictates the amount of time required to pump the water out of the system. The operator must remove the water that enters the various tanks at different times as the simulation advances. The operator's goal is to distribute all the water within the allotted amount of time by activating and deactivating the pumps assigned to each tank. Only 5 pumps can be activated at any one time, and after a pump is used there is a delay of 10 simulation minutes (cleaning time) before this pump (or a different one) can be re-activated. The simulation time begins at 2 o'clock and finishes at 10 o'clock when the final deadline expires.

3.2. *Performance in WPP*

The main performance measure in the WPP simulation is the number of gallons of water (out of a total system capacity of 1080 gallons) that are not processed in time by the user. Thus the best performance is zero. To avoid confusion, we converted the number of gallons missed into a positive percentage measure: the percentage of water processed on time. Thus, larger percentages (i.e., fewer gallons missed) indicate better performance.

WPP does not have a unique solution because there are many possible ways to activate and deactivate pumps and achieve optimal performance. By conducting protocol analyses, we have identified three essential features in the WPP task: the timing of the decisions, the amount of water to process in a particular tank, and the chain length (Gonzalez, 2003).

WPP has proven to be very challenging (Gonzalez et al., 2003). In preliminary experiments, the mean first-run performance was approximately 76%. After extensive practice




Figure 1. The layout of the WPP simulation. When the screenshot was taken at simulation time 6:48 PM, the operator had missed 30 gallons of water and had five pumps activated. Water enters from outside the system and moves continuously from left to right (toward the deadlines) due to the action of activated pumps. The operator decides when to activate and deactivate pumps while the simulation time is running. Numbers have been added to tanks to reference them in text. In the actual task the tanks are not numbered.

(more than 50 times running the task), individuals have reached asymptotic performance levels of approximately 95%, after which no more improvement was observed. Therefore, although some learning occurs, there appears to be a limit for learning the WPP task. Such suboptimal performance is typical of other dynamic tasks.

3.3. Participants

Sixty-four students from local universities participated during three consecutive days (one hour per day) and received \$50 in compensation. All participants ran WPP 18 times (trials), with 6 8-minute (real time) trials per day. On the first day, participants were instructed on how to operate the WPP simulation. During training, they were informed of the task

goal and were given general instructions on how to perform the task. All participants were provided with explanations of the deadlines, simulation time, and the chains through which the water travels, and were informed that different amounts of water could be added from outside the system to any of the tanks at any time. However, participants were given no information describing the amount of water or its time of arrival. They were instructed to do their best to process all the water that appeared within the system, but no particular strategy was outlined for them. During the training period, participants practiced activating and deactivating pumps in the WPP simulation, which was running at the slowest possible pace of 30 real-time minutes per scenario. We did not allow participants to complete the full trial. Rather, we stopped the simulation before the first deadline (5 o'clock PM) so that they would not know the results of the decisions they made during training.

3.4. Instances (SDUs)

Each WPP situation is defined by the attributes of the tank (i.e., time of evaluation, amount of water, the position of the tank in the chain [the *chain value*], and deadline). Decisions required to complete the task include the activation and deactivation of the pumps associated with a tank. The utility measure used in this task is the amount of time remaining before a deadline. This remaining time is calculated as the difference between the deadline and the simulation time at the time of decision. Figure 2 shows concrete examples of instances

Trial	Type	Hr	Min	Tank	Pump	Water	Chain	Deadline
1	activate	2	6	1	3	27	3	480
1	activate	2	24	13	26	1	1	300
1	activate	2	32	0	0	15	3	300
1	activate	2	43	0	1	13	3	300
1	activate	3	1	13	26	1	1	300
1	activate	3	11	11	22	10	1	330
1	activate	3	14	15	30	30	3	480
1	activate	3	19	2	5	43	2	300
1	activate	3	29	13	26	7	1	300
1	activate	3	42	0	0	2	3	300
1	activate	3	54	2	5	23	2	300
1	activate	3	57	0	1	2	3	300
1	activate	4	4	1	2	2	3	480
1	activate	4	18	0	0	2	3	300

Figure 2. Examples of instances in the WPP task. The situation is defined by several parameters: time (Hr, Min), tank, pump, water amount, chain, and deadline. The first instance (first row) shows a decision to open pump number 3 in Tank 1 was made at 2:06, when the tank had 27 gallons of water in it. This tank has a chain of 3 and a deadline of 480 (8 o'clock).

encountered by one of the participants during the first trial. Each row in figure 2 gives the parameters of an instance at which the user decided to activate pumps. The parameters of each instance include the time at which the decision was made (in hours and minutes); the tank affected by the decision; the pump that was activated; the amount of water in the tank (in gallons); the location of the tank (e.g., Tank 13 has a chain length value of 1 because it is the last in the chain, Tank 1 has a chain length value of 3 because there are three tanks downstream); and the deadline (in minutes).

The WPP simulation stores an instance for each activation or deactivation decision. One of the main requirements for successful performance of the WPP task—and most other DDM tasks—is the ability to relate the timing of one’s own actions to the evolution of the system (Kerstholt and Raaijmakers, 1997). IBLT suggests that decision makers use recognition when confronted with a decision-making situation (e.g., Activate Tank 1?) by determining the similarity between the current situation and past instances, retrieving the utility of these past instances, and adjusting new decisions to achieve maximal utility. For example, past decisions that resulted in missed gallons after deadline expiration might require adjustment in terms of their timing. By analyzing the instances generated over many trials, it should be possible to observe the changes in decision-making performance (i.e., the learning) exhibited by individuals completing a DDM task.

3.5. Instance Similarity

Instance similarity is a measure of how closely a current situation resembles a past situation. To compute similarity, all SDUs in a trial were sorted first by tank and then by the time at which each decision was made. This sorting process produced a sequence of an individual’s activation decisions for each tank in each trial. SDUs in trial i were compared to SDUs in trial $i - 1$ (the previous trial) in sequential order according to the time of decision. Similarity was calculated by using two variables: the simulation time at which the decision was made and the amount of water in the tank at the time of decision. To enable the use of a unique measure of comparison, water was converted to time units: 1 gallon of water was pumped every 2 simulation minutes. The following formula, which was recently reported by Gonzalez et al. (2003), was used to calculate similarity:

$$\text{similarity} = 1 - (\alpha * (\text{decision}_i - \text{SDU}_{(i-1)}) + ((1 - \alpha) * (\text{accumulated similarity } \text{SDU}_{(i-1)})))$$

When calculated with this formula, similarity values range from 0 to 1; a similarity value of 1 means that the two compared instances are the same, while a value of 0 means that they are completely dissimilar. The alpha value assigns different weights to the similarity of the most recent SDUs relative to the similarity of all previous ones. Table 1 portrays the decisions made for tank 0 by a single participant, with $\alpha = .99$ (i.e., 99% weight is assigned to the SDUs of the most recent trial). The similarity value for Trial 1 is zero, because there is no previous knowledge for comparison. The similarity value in Trial 2 is calculated by comparing the decisions in Trial 2 to the decisions in Trial 1. The participant in Table 1 made three decisions for tank 0 in Trial 2, four decisions in Trial 3, etc. By incorporating

Table 1. Similarity calculations for several SDUs.

Tank	Trial	Order	Time	Water	Similarity	Similarity calculation
0	1	1	390	27	0.00	
0	2	1	172	15	0.50	$1 - ((0.99 * ((172 + 15 * 2) - (390 + 27 * 2)) / 480 + 0.01 * 0)) = 0.50$
0	2	2	226	2	0.00	
0	2	3	285	8	0.00	
0	3	1	147	10	0.92	$1 - ((0.99 * ((147 + 10 * 2) - (172 + 15 * 2)) / 480 + 0.01 * 0.5)) = 0.92$
0	3	2	222	2	0.99	$1 - ((0.99 * ((222 + 2 * 2) - (226 + 2 * 2)) / 480 + 0.01 * 0)) = 0.99$
0	3	3	248	2	0.90	$1 - ((0.99 * ((248 + 2 * 2) - (285 + 8 * 2)) / 480 + 0.01 * 0)) = 0.90$
0	3	4	269	2	0.00	

the similarity of the first SDU, the equation yields a similarity value of 0.5 for the decision in Trial 2:

$$1 - (.99 * ((172 + 15 * 2) - (390 + 27 * 2)) / 480 + .01 * 0.0).$$

Table 1 explains how we calculated the similarity value for several SDUs.

4. Results

In general, the participants' performance improved with practice. Figure 3 shows the participants' average (mean) learning curve and the mean and standard error per trial. A repeated measures analysis of performance indicated that significant learning occurred over the 18 trials ($F(17,986) = 38.212, p < .001$).

4.1. Question 1: Does Similarity between Contiguous Trials Correlate with Performance?

The SDU data set for each participant contained approximately 43 SDUs per trial, yielding a total of 49,542 instances over the 18 trials. SDU similarity was averaged per trial. Figure 4 shows changes in similarity with practice. The average similarity between trial i and $i-1$ increased with task practice as participants improved their performance. A repeated measures ANOVA indicated a significant increase in similarity with continuing practice in the task ($F(16,1008) = 69.07, p < .001$). A linear regression calculation using the average similarity as the independent variable and the average performance as the dependent variable indicated that similarity was a good predictor of performance ($F(1,63) = 14.04, p < .001, R^2 = .185, t = 3.75, p < .001$).

We also compared the use of decision similarity by good and poor performers. If good performers can recognize a decision-making situation better than poor performers can, the




Figure 3. Participants' average performance for 18 trials.




Figure 4. Average similarity over time. Trial i shows the average similarity of situations compared to trial $i - 1$.

decisions of the former should be more similar. We divided the sample of 64 participants into good and poor performers by defining good performers as those who scored at least one standard deviation above the group mean and poor performers as those who scored at least one standard deviation below the group mean. This led to the identification of two groups




Figure 5. Similarity of decisions made by good and poor performers. The decision similarity is higher for good performers.

containing 12 subjects each. If individuals depend on their past decisions to improve their current performance, then the decisions of those who perform better should show higher similarity over time than the decisions of poor performers. Figure 5 shows that decision similarity differed between good and poor performers. This difference was particularly noticeable in the earlier simulation trials. The figure indicates that good performers showed high average decision similarity from the very beginning, whereas poor performers tended to reach high average values of similarity only after several trials. This result could indicate that good performance in the task is related to the amount of time required for participants to learn to make similar decisions.

In the remainder of the paper, we assume (in accordance with Gonzalez et al., 2003) that decision makers reuse their past decisions, and that this reuse explains why similarity predicts performance. Of course, the applicability of this assumption depends on the structure of the task (i.e., the consistency of the environment). For our purposes here it should hold true because we presented people with the same scenario in each trial. In an alternative microworld with more diverse circumstances, this effect may not hold; we address this observation further in the discussion below.

4.2. Question 2: Which Factors Influence People's Similarity Judgments and How Do These Factors Change as Performance Improves?

We took two approaches in answering this question. First, we ran regression models using the similarity measured for each made decision as the dependent variable and the features of an instance as predictors. Second, we performed a two-step cluster analysis.

The regression models informed us as to what proportion of the variance in similarity observed between instances in consecutive trials was attributable to the features involved in each instance. For our experiment, the variables were tank, chain, deadline, time, and water. The easiest question to answer via this method is whether the variance in similarity can be explained by any one variable. The adjusted R^2 values for the linear models containing a single variable (see Table 2) demonstrate that individual features explain only negligible amounts of the variance in similarity.

We next ran regression models with all the features and a regression model with all the features and simple interactions. The results in Table 3 indicate that the best predictor of similarity is the model with simple interactions among the variables. The highest beta coefficient corresponded to the chain and the interaction between chain and time, suggesting that these features are the best predictors of similarity between consecutive trials.

We used cluster analysis to identify the factors that influence similarity judgments and to investigate how these factors vary as performance improves. Two-step cluster analysis is designed to explore natural groupings (clusters) for large data sets. This algorithm consists of two stages. In the first stage a sequential cluster method is applied to a large data set to compress the regions and form sub-clusters. In the second stage, the sub-clusters are organized into the desired number of clusters (Zhang, 1996). We conducted a cluster analysis using categorical and continuous variables that describe a

Table 2. Simple regression models using only one variable to predict similarities.

Variable	Adjusted R -square	F (1,49541)	p value
Chain (nominal, 3 levels)	0.001	54.150	.000
Deadline	0.001	31.823	.000
Time	0.002	105.340	.000
Tank (nominal, 23 levels)	0.000	3.607	ns
Water	0.000	.142	ns

Table 3. Multiple variable models.

Variables	Multiple R -squared	F	p value
All continuous values (deadline, time, water) No interactions	0.002	(3,49541) = 38.108	0.000
All continuous values (deadline, time, water) + chain No interactions	0.003	(4,49541) = 42.338	0.000
All continuous values (deadline, time, water) + chain + first order interactions	0.877	(10,49541) = 35231.212	0.000

Table 4. Cluster analysis results.

Trial	Group	Time	Water	Tank	Chain	Deadline
Cluster 1						
1	0	350.2	22.5	0-1-2-3-4-15-16-21-22	2-3	480-300-540
2	3	388.4	27.3	1-3-4-15-16-21-22	3	480-540
3	2	382.9	27.4	5-6-9-10-17-18-19-20	2	600-420-450
4	1	278.0	10.2	0-2-7-8-11-12-13-14	1-3-2	300-330-360-510-390-570
5	1	277.3	10.7	0-2-7-8-11-12-13-14	1-3-2	300-330-360-510-390-570
6	1	275.4	10.7	0-2-7-8-11-12-13-14	1-3-2	300-330-360-510-390-570
7	1	279.8	10.8	0-2-7-8-11-12-13-14	1-3-2	300-330-360-510-390-570
8	2	387.2	26.7	5-6-9-10-17-18-19-20	2	600-420-450
9	2	391.4	27.7	5-6-9-10-17-18-19-20	2	600-420-450
10	1	280.8	10.9	0-2-7-8-11-12-13-14	1-3-2	300-330-360-510-390-570
11	2	386.3	28.2	5-6-9-10-17-18-19-20	2	600-420-450
12	1	378.5	25.6	0-2-7-8-11-12-13-14	1-3-2	300-330-360-510-390-570
13	1	285.0	11.4	0-2-7-8-11-12-13-14	1-3-2	300-330-360-510-390-570
14	1	290.7	11.9	0-2-7-8-11-12-13-14	1-3-2	300-330-360-390-570-510
15	1	288.3	12.1	0-2-7-8-11-12-13-14	1-3-2	300-330-360-390-570-510
16	1	291.0	11.8	0-2-7-8-11-12-13-14	1-3-2	300-330-360-570-510-390
17	3	383.8	23.7	1-3-4-15-16-21-22	3	480-540
18	1	287.1	12.1	0-2-7-8-11-12-13-14	1-3-2	300-330-360-570-510-390
Cluster 2						
1	1	302.0	10.3	7-8-11-12-13-14	1	300-330-360-390-510-570
2	1	272.4	10.4	0-2-7-8-11-12-13-14	1-2-3	300-330-360-390-510-570
3	3	392.3	26.6	1-3-4-15-16-21-22	3	480-540
4	3	387.8	25.7	1-3-4-15-16-21-22	3	480-540
5	2	380.7	27.2	5-6-9-10-17-18-19-20	2	420-450-600
6	2	387.0	27.3	5-6-9-10-17-18-19-20	2	420-450-600
7	3	381.9	25.5	1-3-4-15-16-21-22	3	480-540
8	1	278.8	10.9	0-2-7-8-11-12-13-14	1-2-3	300-330-360-390-510-570
9	1	280.9	11.1	0-2-7-8-11-12-13-14	1-2-3	300-330-360-390-510-570
10	3	379.8	24.3	1-3-4-15-16-21-22	3	480-540
11	3	377.3	25.8	1-3-4-15-16-21-22	3	480-540
12	1	283.6	10.9	0-2-7-8-11-12-13-14	1-2-3	300-330-360-390-510-570
13	3	383.7	26.0	1-3-4-15-16-21-22	3	480-540
14	2	388.6	26.9	5-6-9-10-17-18-19-20	2	420-450-600
15	2	381.3	26.5	5-6-9-10-17-18-19-20	2	420-450-600
16	2	383.3	26.8	5-6-9-10-17-18-19-20	2	420-450-600
17	1	294.8	11.7	0-2-7-8-11-12-13-14	1-2-3	300-330-360-390-510-570

(Continued on next page.)

Table 4. (Continued).

Trial	Group	Time	Water	Tank	Chain	Deadline
18	3	381.6	24.6	1-3-4-15-16-21-22	3	480-540
Cluster 3						
1	2	376.0	27.0	5-6-9-10-17-18-19-20	2	420-450-600
2	2	375.0	27.6	5-6-9-10-17-18-19-20	2	420-450-600
3	1	276.5	10.1	0-2-7-8-11-12-13-14	1-2-3	300-330-360-390-510-570
4	2	379.8	26.3	5-6-9-10-17-18-19-20	2	420-450-600
5	3	381.8	24.9	1-3-4-15-16-21-22	3	480-540
6	3	382.4	25.0	1-3-4-15-16-21-22	3	480-540
7	2	389.5	27.4	5-6-9-10-17-18-19-20	2	420-450-600
8	3	385.3	25.8	1-3-4-15-16-21-22	3	480-540
9	3	380.7	25.5	1-3-4-15-16-21-22	3	480-540
10	2	387.4	28.5	5-6-9-10-17-18-19-20	2	420-450-600
11	1	283.6	11.3	0-2-7-8-11-12-13-14	1-2-3	300-330-360-390-510-570
12	2	384.1	27.4	5-6-9-10-17-18-19-20	2	420-450-600
13	2	387.1	27.3	5-6-9-10-17-18-19-20	2	420-450-600
14	3	372.1	23.8	1-3-4-15-16-21-22	3	480-540
15	3	377.1	24.1	1-3-4-15-16-21-22	3	480-540
16	3	377.8	24.0	1-3-4-15-16-21-22	3	480-540
17	2	382.2	26.1	5-6-9-10-17-18-19-20	2	420-450-600
18	2	386.5	27.1	5-6-9-10-17-18-19-20	2	420-450-600

particular decision-making situation in the WPP simulation: time, water, tank, chain, and deadline. By using the Bayesian Information Criterion (BIC; Fraley and Raftery, 1998), we determined that three were the optimum number of clusters from all the data sets. We then ran one cluster analysis per trial while specifying the number of clusters to be three.

Table 4 summarizes the results of the cluster analyses per trial. The table shows the average time, water, tanks, chains, and deadlines involved in each of the three clusters per trial. The analyses of tank, chain, and deadline parameters clearly identified three groups per cluster. Group 1 includes tanks 0, 2, 7, 8, 11, 12, 13, and 14. Group 2 includes tanks 5, 6, 9, 10, 17, 18, 19, and 20. These are the tanks involved in a two-tank chain that had to be emptied to meet the 7:00 (420 minutes), 7:30 (450 minutes), and 10:00 (600 minutes) deadlines. These tanks, which also were involved in multi-tank chains affected by early and late deadlines, were the easiest tanks to control. Group 3 includes tanks 1, 3, 4, 15, 16, 21, and 22; chain-length 3; and deadlines 8:00 (480 minutes) and 9:00 (540 minutes). Figure 6 identifies the three groups of tanks on a screenshot of the WPP simulation.

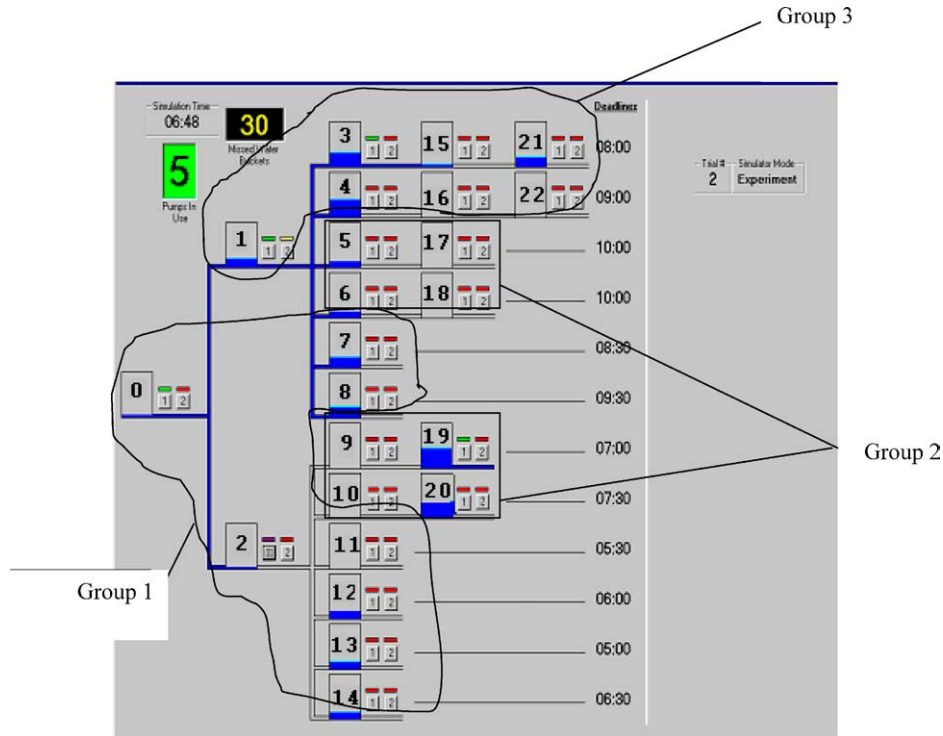




Figure 6. Three groups of Tanks identified by using a cluster analysis. Group 1: Tanks 1, 3, 4, 5, 15, 16, 21, and 22; Group 2: Tanks 5, 6, 9, 10, 17, 18, 19, and 20; Group 3: Tanks 0, 2, 7, 8, 11, 12, 13, and 14.

This cluster analysis explains the low R^2 values generated by the simple regression models and the high R^2 values observed with the model accounting for simple interactions. Similarity between two instances seems to be determined by a complicated combination of the values of different variables that change over time, rather than by single, static combinations of variables. To investigate how the variables changed over time as performance improved, we held Groups 1, 2, and 3 (as identified in the clusters) constant and looked at the average time and water over the course of 18 trials within each group. This method allowed us to ascertain if decision makers adjusted the moment of intervention and recognized the amount of water in the tanks as they learned the task. Figure 7(a) shows the average time of intervention per trial for Groups 1, 2, and 3, and figure 7(b) shows the corresponding amount of water in the tank at the time of intervention. We observed an increasing trend for both variables in group 1, a decreasing trend in Group 3, and a relatively stable trend in Group 2. These findings indicate that decision makers learned to intervene later in shorter chains (with nearer deadlines) and sooner in longer chains (with later deadlines). The corresponding amount of water also generally increased and generally decreased for Groups 1 and 3, respectively, indicating that participants changed and adjusted their recognition of a situation over time.



(a)



(b)

Figure 7. The attention to the task features changed over time. Panel (a) shows the time of intervention over time for the three groups. Panel (b) shows the average amount of water in the groups of tanks.

5. Discussion

During DDM the tendency for people to adopt an event-based, open-loop view of the dynamic system can cause misperceptions and, in so doing, can negatively affect the ability of people to learn to perform optimally (Diehl and Serman, 1995). IBLT provides a cognitive structure for these events and predicts that recognition, as measured by the similarity of situations, will determine decision makers' use of outcome feedback to make high utility decisions (Gonzalez et al., 2003).

When participating in DDM, decision makers decide if a particular event has occurred in the past and, if so, they reuse the solutions to past problems. The selection of actions that worked well in the past involves a general learning process known as the Law of Effect (Edward L. Thorndike). Researchers have used this basic learning principle to demonstrate that people learning to play games repeatedly decide to reuse the strategies that led to good outcomes in the past (Erev and Roth, 1998). In DDM, however, because of the dynamics and the complexity of the system, it is difficult to identify the strategies that led to success in the past. According to IBLT, instance-based learning is a flexible mechanism by which decision makers learn to deal with dynamic complexity (Gonzalez et al., 2003). The study described here demonstrates that SDU similarity increases with task practice and that high similarity is associated with high performance. Our findings also clearly indicate that long-term good performance is determined by the quality of the decisions made in early task practice trials. Furthermore, the initial similarity values calculated in this trial suggest that individuals make judgments based on their own unique experiences. These results explain how people may acquire control of a dynamic task, but not necessarily learn to perform it well. We have proposed previously that learning occurs not only due to the ability to successfully retrieve high utility instances, but due to the ability to adequately update this utility with feedback (Gonzalez et al., 2003). Therefore performance depends both on similarity and on the retrieval of high utility SDUs. Since feedback delays are a characteristic of DDM systems, decision makers at times may fail to attribute the consequences of decisions to the appropriate prior actions. Therefore, outcome feedback often cannot be used during DDM to update the expected utilities of SDUs to reflect experienced utilities. Under such conditions, individuals may acquire control of the system and reach a performance plateau, but will not necessarily learn how to improve performance.

Increases in similarity in DDM systems are highly dependent on the consistency of the environment—i.e., the time patterns in which exogenous events occur. In this study we observed a considerable degree of both decision similarity and learning. However, participants experienced the same exogenous events across trials. The patterns by which water arrived in the tanks was constant throughout the learning process. Such constancy is uncommon in the real world. We believe that the ability of decision makers to use past knowledge to improve performance decreases as patterns of environmental changes become more unpredictable. We currently are conducting a laboratory study to test this hypothesis.

In this study we also determined that the recognition process is influenced by a combination of task features rather than by individual features. This finding may explain why it is so difficult to identify strategies that work well in dynamic situations. Most strategies are related to particular events experienced with practice. The results from the cluster analysis

suggest that the attention that decision makers pay to different features changes over time and that with practice decision makers can adapt their actions to fulfill the task demands. This analysis also indicates that accurate timing of a behavior may require extensive knowledge of the task and of the temporal relationships among the system components, a finding that parallels the findings of other researchers (Kerstholt and Raaijmakers, 1997). The results from this portion of the study suggest that decision makers can learn to adjust the moment of intervention.

In summary, the findings generated by this study demonstrate that the recognition process contributes to the type of experience that individuals accumulate as they work to complete a dynamic task. As experience accumulates, decision makers become more selective, adjust their decisions, and thus gradually change the task features that trigger the recognition process.

Acknowledgments

This research was supported by the Advanced Decision Architectures Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory (DAAD19-01-2-0009), the Office of Naval Research (N00014-01-10677), and by Grant EIA—0121201 from the National Science Foundation.

References

- Brehmer, B. (1990), "Strategies in Real-Time, Dynamic Decision Making," in R.M. Hogarth (Ed.), *Insights in Decision Making*, University of Chicago Press, pp. 262–279.
- Brehmer, B. (1995), "Feedback Delays in Complex Dynamic Decision Tasks," in P. Frensch and J. Funke (Eds.), *Complex Problem Solving: The European Perspective*, Erlbaum, Hillsday, NJ.
- Brehmer, B. and D. Dörner (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.
- Diehl, E. and J.D. Serman (1995), "Effects of Feedback Complexity on Dynamic Decision Making," *Organizational Behavior and Human Decision Processes*, 62(2), 198–215.
- Dienes, Z. and R. Fahey (1995), "Role of Specific Instances in Controlling a Dynamic System," *Journal of Experimental Psychology: Learning, Memory and Cognition*, 21(4), 848–862.
- Erev, I. and A.E. Roth (1998), "Predicting how People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *The American Economic Review*, 848–881.
- Fraley, C. and A.E. Raftery (1998), "How Many Clusters: Which Clustering Method? Answers via Model-Based Cluster Analysis," *Computer Journal*, 4, 578–588.
- Gibson, F.P., M. Fichman, and D.C. Plaut (1997), "Learning Dynamic Decision Tasks: Computational Model and Empirical Evidence," *Organizational Behavior and Human Decision Processes*, 71(1), 1–35.
- Gonzalez, C. (2003), *Novice and Expert Verbal Protocols in Real-Time Dynamic Decision-Making*. Human Factors and Ergonomics Society, Denver.
- Gonzalez, C., F.J. Lerch and C. Lebiere (2003), "Instance-Based Learning in Real-Time Dynamic Decision Making," *Cognitive Science*, 27, 591–635.
- Kerstholt, J.H. and J.G.W. Raaijmakers (1997), "Decision Making in Dynamic Task Environments," in R. Ran- yard, R.W. Crozier, and O. Svenson (Eds.), *Decision Making: Cognitive Models and Explanations*, Routledge, New York, NY, pp. 205–217.
- Klein, G., J. Orasanu, R. Calderwood and C.E. Zsombok (Eds.) (1993), *Decision Making in Action: Models and Methods*. Ablex Publishing Corporation, Norwood, New Jersey.

- Lerch, F.J., D.B. Ballou and D.E. Harter (1997), "Using Simulation-Based Experiments for Software Requirements Engineering," *Annals of Software Engineering*, 3, 345–366.
- Medin, D.L., R.L. Goldstone and A.B. Markman (1995), "Comparison and Choice: Relations Between Similarity Processing and Decision Processing," *Psychonomic Bulletin and Review*, 2(1), 1–19.
- Simon, H.A. and P. Langley (1981), "The Central Role of Learning in Cognition," in H.A. Simon (Ed.), *Models of Thought*, Yale University Press: New Haven and London, vol. II, pp. 102–184.
- Sterman, J.D. (1994), "Learning in and About Complex Systems," *System Dynamics Review*, 10(2/3), 291–330.
- Sterman, J.D. (2000), *Business Dynamics: Systems Thinking and Modeling for a Complex World*. McGraw-Hill.
- Tversky, A. (1977), "Features of Similarity," *Psychology Review*, 84, 327–352.
- Zhang, T.R. (1996), "BIRCH: An Efficient Data Clustering Method for Very Large Databases," Paper presented at the ACM SIGMOD Conference on Management of Data, Montreal, Canada.
- Zsombok, C.E. and G. Klein (Eds.) (1997), *Naturalistic Decision Making*. Lawrence Erlbaum Associates, Inc, Mahwah, New Jersey.

Cleotilde Gonzalez is Assistant Professor of Information Systems in the Department of Social and Decision Sciences at Carnegie Mellon University. Her research focuses on the psychology of dynamic decision making. She uses behavioral, computational and brain imaging approaches to understand how people make decisions in dynamic, complex environments. She is the founder and director of the Dynamic Decision Making Laboratory at Carnegie Mellon (www.cmu.edu/ddmlab) that currently holds several post-doctoral fellows, researchers and programmers.

Jose Quesada is a Post-Doctorate fellow at the Dynamic Decision Making Laboratory at Carnegie Mellon University (CMU). Jose's work focuses on the extending the use of Latent Semantic Analysis (LSA) principles to problem solving domains including simulated fire fighting and process control as well as aircraft landing efficiency. They are all dynamic tasks, and they are all treated using corpus-based approaches to cognition. In his work at CMU, Jose plans to extend the framework to explain decision making phenomena.