8-16-1996

A Computational Model of Learning with Time Pressure in Dynamic Decision Tasks

Faison P. Gibson
Graduate School of Industrial Administration, Carnegie Mellon University, gibson+@cmu.edu

A Computational Model of Learning with Time Pressure in Dynamic Decision Tasks

Faison P. Gibson
Graduate School of Industrial Administration
Carnegie Mellon University
Pittsburgh, PA 15213
gibson+@cmu.edu
http://www.gsia.cmu.edu/andrew/gibson/web/

Introduction and Theory

In dynamic decision tasks such as managing a production process, decision makers attempt to control an evolving situation in real-time. This type of *dynamic* decision task can be distinguished from one-time decision tasks such as buying a house, by the presence of four elements: (1) The tasks require a series of decisions rather than one isolated decision; (2) The decisions are interdependent; (3) The environment changes both autonomously and as a result of decision makers' actions; (4) Decisions are goal-directed and made under time pressure, thereby reducing the decision-maker's opportunities to consider and explore options (Brehmer, 1990, 1992, 1995; Rapoport, 1975; Edwards, 1962).

Hogarth (1981) has hypothesized that an important feature of many dynamic decision environments is that they provide decision makers with feedback on the outcomes of their actions. Decision makers may use this outcome feedback to adapt their decisions on-line as they perform the task. Recently, Davis and Kottemann (1995) have shown evidence that decision makers' performance improves in a dynamic task when they are provided with continuous and cumulative outcome feedback. Sengupta and Abdel-Hamid (1993) find similar results when they provide their subjects with a decision aid that immediately models the environment's long-term response to their actions.

However, Sterman (1989a, 1989b, 1994; Paich & Sterman, 1993; Diehl & Sterman, 1995) provides and discusses evidence that subjects using outcome feedback alone fail to learn to control dynamic systems to the level that they could with simple linear heuristics. Furthermore, several researchers have obtained results supporting the general finding that subjects in dynamic tasks fail to use outcome feedback to adapt to the point where they can effectively cope with the demands of the decision environment, at least in the time-frame dictated by the experiment and with the information provided about the environment (Brehmer, 1990, 1992, 1995; Kleinmuntz, 1993; Kleinmuntz & Thomas, 1987; Richardson & Rohrbaugh, 1990). Attempts to improve subjects' performance by providing various decision aids in addition to outcome feedback have generated mixed results (e.g., Kleinmuntz & Thomas, 1987; Richardson & Rohrbaugh, 1990).

The mixture of results described in the last two paragraphs suggests that, while central to improving performance, the process by which people adapt or learn learn on-line from outcome feedback in dynamic decision tasks is only partially understood. To begin to address this issue, we have developed a computational formulation of how people might learn in these situations.

Our formulation builds on previous theoretical work in dynamic decision making by Brehmer (1990, 1992, 1995) and motor learning by Jordan and Rumelhart (1992; Jordan, 1992, in press; Jordan, Flash, & Arnon, 1994; Wolpert, Gharamani, & Jordan, 1995). It proposes two central assumptions about how decision makers learn on-line from outcome feedback in dynamic decision tasks. The first is that decision makers use outcome feedback to form two interdependent, internal submodels of the task as they participate in it. These two submodels represent knowledge, conditioned by environmental context, about: (1) How the decision maker's actions affect outcomes; (2) Which actions to take to achieve desired outcomes. Our formulation's second major assumption is that the acquisition of these two types of knowledge from outcome feedback can be simulated using on-line learning of parallel distributed processing (PDP) or neural network models. This second assumption is important because, among other reasons, it focuses the applicability of the model to situations where decision makers have little time for reflection.

Computational and Laboratory Studies

The advantage of constructing a computational formulation is that we can use it to generate testable predictions about human performance in different task manipulations and settings. We have done this using the Sugar Production Factory, a simple dynamic decision task in which subjects, using outcome feedback alone, learn to manipulate an input (typically workforce) to achieve target levels of sugar production as they participate in the task. The task has been widely used to investigate on-line learning in dynamic decision environments (Gibson, in press; Berry, 1991; Berry & Broadbent, 1984, 1987, 1988; Buchner, Funke, & Berry, 1995; Gibson & Plaut, 1995; Hayes & Broadbent, 1988; Marescaux, Luc, & Karnas, 1989; McGeorge & Burton, 1989; Squire & Frambach, 1990; Stanley, Matthews, Buss, & Kotler-Cope, 1989; Sanderson, 1990; Dienes & Fahey, 1994, 1995; Dienes, 1990; Berry & Dienes, 1993).

Figure 1: A comparison of the average learning performance across training sessions of Stanley et al.'s human subjects and our base simulation (source: Gibson & Plaut, 1995)

We have used these previous results (principally from Stanley et al., 1989, and Berry & Broadbent, 1984) as a replication test for our formulation. As illustrated in Figure 1, our formulation is able to broadly match human learning in the Sugar Production Factory over 3 training sessions of 200 trials each. Each trial represents one attempt by the learner to bring the factory to goal production by setting workforce. As used originally by Stanley et al. (1989) with their human subjects, the dependent measure is average number of trials per session out of 10 in which the learner is able to reach within plus-or-minus one level of the production goal.

However, our formulation suggests three properties of learning that we were unable to test by replicating these results:

- Learning is approximate.
- Learning is most applicable locally
- Because of this approximate, local learning, transfer of knowledge is best nearest the training region.

In a second study, which is on-going, we are using results from the first study and additional simulation results that illustrate these properties to make and test predictions concerning human learning in the Sugar Production Factory.

Conclusion

At the mini-track, I will report in detail on the findings from the computational and laboratory studies outlined above. This work has served as the theoretical basis for a field study at a large credit collections facility. As time permits, I will describe our work at this facility, further computational work based on insights gained there, and plans for future work.

A much more extensive treatment of this work is given in *Learning in Dynamic Decision Tasks: Computational Model and Empirical Evidence* which can be found at the URL at the beginning of this paper.

Acknowledgements

I would like to thank my two dissertation advisors, Mark Fichman and David C. Plaut, for their support.

References

Berry, D. C. (1991). The role of action in implicit learning. *Quarterly Journal of Experimental Psychology, 43A*(4), 881--906.

Berry, D. C., & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology, 36A*, 209--231.

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

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

Berry, D. C., & Dienes, Z. (1993). *Implict learning: Theoretical and empirical issues*. Lawrence ErlbauAssociates.

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

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

Brehmer, B. (1995). Feedback delays in complex dynamic decision tasks. In P. Frensch, & J. Funke (Eds.), *Complex problesolving: The European perspective*. Hillsdale, NJ: Lawrence ErlbauAssociates.

Buchner, A., Funke, J., & Berry, D. C. (1995). Negative correlations between control performance and verbalizeable knowledge: Indicators for implicit learning in process control tasks? *Quarterly Journal of Experimental Psychology, 48A*(1), 166--187.

Davis, F. D., & Kottemann, J. E. (1995). Determinants of decision rule use in a production planning task. *Organizational Behavior and Human Decision Processes, 64*(2), 145--157.

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. (1990). *Implicit concept formation*. PhD thesis, University of Oxford.

Dienes, Z., & Fahey, R. (1994). *The role of implicit memory in controlling a dynamic system*. Unpublished manuscript.

Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(4), 1--15.

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

Gibson, F. P. (in press). The sugar production factory---a dynamic decision task. *Journal of Computational and Mathematical Organization Theory*, 1(3).

Gibson, F. P., & Plaut, D. C. (1995). A connectionist formulation of learning in dynamic decision-making tasks. *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 512--517). Hillsdale, NJ: Lawrence ErlbauAssociates.

Hayes, N. A., & Broadbent, D. E. (1988). Two modes of learning for interactive tasks. *Cognition, 28*, 249--276.

Hogarth, R. M. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin, 90*, 197--217.

Jordan, M. I. (1992). Constrained supervised learning. *Journal of Mathematical Psychology, 36*, 396--425.

Jordan, M. I. (in press). Computational aspects of motor control and motor learning. In H. Heuer, & S. Keele (Eds.), *Handbook of perception and action: Motor skills*. New York: Academic Press.

Jordan, M. I., Flash, T., & Arnon, Y. (1994). A model of the learning of artrajectories frospatial deviations. *Journal of Cognitive Neuroscience, 6*(4), 359--376.

Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science, 16*(3), 307--354.

Kleinmuntz, D. N. (1993). Information processing and misperceptions of the implications of feedback in dynamic decision making. *System Dynamics Review, 9*(3), 223--237.

Kleinmuntz, D. N., & Thomas, J. B. (1987). The value of action and inference in dynamic decision making. *Organizational Behavior and Human Decision Processes, 39*, 341--364.

Marescaux, P.-J., Luc, F., & Karnas, G. (1989). Modes d'apprentissage selective et non-selective et connaissances acquises au controle d'un processus: 'Evaluation d'un mod`ele simul'e. *Cahiers de Psychologie Cognitive Europeenne, 9*, 239--264.

McGeorge, P., & Burton, A. M. (1989). The effect of concurrent verbalization on performance in a dynamic systems task. *British Journal of Psychology, 80*, 455--465.

Paich, M., & Sterman, J. D. (1993). Boom, bust, and failures to learn in experimental markets. *Management Science, 39*(12), 1439--1458.

Rapoport, A. (1975). Research paradigms for studying dynamic decision behavior. In D. Wendt, & C. Vlek (Eds.), *Utility, probability, and human decision making*. Dordrecht-Holland: D. Reidel Publishing Company.

Richardson, G. P., & Rohrbaugh, J. (1990). Decision making in dynamic environments: Exploring judgments in a system dynamics model-based game. In K. Borcherding, O. I. Larichev, & D. M. Messick (Eds.), *Contemporary issues in decision making* (pp. 463--472). North-Holland: Elsevier Science Publishers B. V.

Sanderson, P. M. (1990). *Implicit and explicit control of a dynamic task: Empirical and conceptual issues* (Technical Report EPRL-90-02). Urbana, IL 61801: Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign.

Sengupta, K., & Abdel-Hamid, T. K. (1993). Alternative conceptions of feedback in dynamic decision environments: An experimental investigation. *Management Science, 39*(4), 411--428.

Squire, L. R., & Frambach, M. (1990). Cognitive learning in amnesia. *Psychobiology, 18*(1), 109--117.

Stanley, W. B., Mathews, R. C., Buss, R. R., & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology*, *41A*(3), 553--577.

Sterman, J. D. (1989a). Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes*, *43*, 301--335.

Sterman, J. D. (1989b). 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*(2--3), 291--330.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). Forward dynamic models in human motor control: Psychophysical evidence. *Advances in Neural Information Processing Systems 7*. Cambridge, MA: MIT Press.