Dynamic Decision Making: Implications for Recommender System Design

Cleotilde Gonzalez
Dynamic Decision Making Laboratory
Department of Social and Decision Sciences
Carnegie Mellon University
coty@cmu.edu

We make decisions in increasingly complex, high-risk, and dynamic environments that evolve over time in unpredictable ways, and the options that we have available in our daily decisions has exponentially increased. For example, when shopping in a store the item diversity on the shelves are large, menus in the restaurants offer large variety, books to select from in the bookstore are large, etc. We are living a choice explosion era. Even more dramatic is the choice explosion in the cyber-world. Given no physical storage restrictions, the options to choose from in the cyber world is immense. More than ever before, these situations challenge our cognitive abilities to process information and to make accurate decisions. How do we choose from this large diversity of options and how do we decide which ones best match our preferences? In the physical world, we may get advice from people we know: experts, friends and family, or we may get help and support from technology such as while driving relying on a GPS. In the cyber world, we now rely on recommender systems that help to filter the large amounts of information and to reduce possible decision options by predicting preferences of a decision maker and offering best possible alternatives.

Recommender systems vary in their approach and ways in which individual preferences are collected and the way in which information and alternatives are filtered for particular users. However, ultimately, all recommender systems aim at predicting human preferences and choice and the essence of every recommender system is the human decision making process. Furthermore, because human preferences are not static, recommender algorithms must be dynamic and adaptable to changes. Often preferences are constructed through past experience (choices and outcomes observed in the past) and through explicit information provided. These characteristics suggest that human preferences are dynamic and contingent to the decision environment. I suggest that Dynamic Decision Making (DDM) research may help to build recommender systems that learn and adapt recommendations dynamically and to a particular user's experience, to maximize benefits and overall utility of her choices. I present a conceptual framework for dynamic decision making that is different from the traditional view of choice in the behavioral sciences, summarize main behavioral results obtained from experimental studies in dynamic situations; and summarize a theory and a computational model that has demonstrated accuracy in predicting human choice in a large diversity of tasks, which may provide an initial departure point for improving recommender algorithms.

1 Dynamic Decision Making defined

In contrast to a popular static view of decision making, DDM characterizes choice as a closed-loop process representing the interaction between the environment and a decision maker (Forrester, 1961; Rapoport, 1975; Sterman, 1989; Gonzalez, 2012, 2013). The figure below conceptualizes this idea: A decision maker perceives information from the environment and transforms that information to find and create alternatives, to build preferences, and to evaluate options that lead to a choice. An action is taken which changes the environment, and feedback from the action is processed in a way that one may reuse past decisions in future actions (Gonzalez et al., 2003; Gonzalez, 2012).

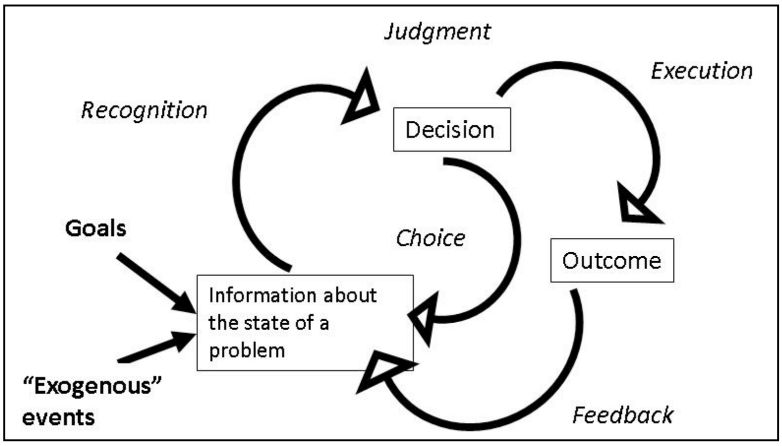

The essential element of DDM is a series of choices taken over time to achieve some overall goal. Decision making may be dynamic at different degrees, according to additional characteristics such as: 1) choices are interdependent so that later decisions are contingent to earlier actions; 2) the environment changes both spontaneously and as a consequence of earlier actions; and 3) decisions need to be made in real-time (Edwards, 1962; Rapoport, 1975; Brehmer, 1992; Hogarth, 1981; Gonzalez et al., 2003; Kerstholt & Raaijmakers, 1997). Under this view, DDM is a learning process where alternatives unfold over time, decisions depend on previous choices and on external events and conditions, and decisions are made from experience and based on feedback.

2 Main behavioral results from psychological experiments in DDM

Decision making has been studied in complex dynamic environments using "microworlds," simulation systems representing a realistic situation and context (Brehmer, 1993; Funke, 1988; Omodei, Wearing, McLennan & Hansen, 2001; Gonzalez, Vanyukov & Martin, 2005; Frensch & Funke, 1995). Experiments with microworlds have identified common human errors committed when working with complex tasks (Brehmer & Dorner, 1993; Dorner, 1987), including the processes and problems of dealing with feedback delays, types of feedback and feedback specificity (Brehmer, 1990; Gonzalez, 2005; Sterman, 1989); time constraints (Kerstholt, 1994; Gonzalez, 2004); cognitive workload (Gonzalez, 2005b); and the relationships between cognitive abilities and performance (Gonzalez, Thomas & Vanyukov, 2005; Rigas, Carling & Brehmer, 2002). Findings suggest that in dynamic situations, learning from only outcome feedback is slow and generally ineffective (Gonzalez, 2005a). Instead, reflecting on an expert's performance improved dynamic decision more effectively. Relatedly, we have found that while learning a dynamic resource allocation task, one needs to learn slowly and slow learning results in best subsequent performance under high-stress and time pressure conditions (Gonzalez, 2005b). A similar demonstration showed that individuals who learn under low cognitive workload are able to perform more accurately in a transfer task while under high workload (Gonzalez, 2004). Another important insight relates to how to make the best use of our tendency to rely on context-specific instances in order to improve adaptation to novel situations. Our results suggest that an effective way to do so is through instance diversity. The diversity of instances is defined by the attributes in each situation. When individuals are trained in multiple, diverse situations (e.g., many categories), they have been found to adapt more successfully to novel conditions compared to when they are exposed to less diverse conditions (Brunstein & Gonzalez, 2011; Gonzalez & Madhavan, 2011).

DDM has also been examined in extreme simplifications of dynamic tasks. For example, recent developments in decision sciences provide new insights into our understanding of DDM. This is a shift of attention from one-shot decisions in which all information is provided to the decision maker (probabilities and outcomes are explicit) to repeated decisions in which no information at all is given requiring that decisions are made from experience (Hertwig et al., 2004; Barron & Erev, 2003; Erev & Barron, 2005; Hertwig & Erev, 2009). Three main insights have emerged from our research in these simplified paradigms. First, is conditional reinforcement: people increasingly select actions that led to best outcomes in similar past experiences (Gonzalez & Dutt, 2012; Erev & Barron, 2005); second, reduced exploration: people decrease exploration of options over time in consistent environments (Gonzalez & Dutt, 2011); third: recommenders may act as distractions for humans' own exploration and search for best value, although they tend to abandon imperfect recommenders with practice (Harman, ODonovan, Abdelzaher & Gonzalez, 2014; Harman & Gonzalez, in preparation).

3 Instance-Based Learning Theory and computational models

Instance-Based Learning Theory (IBLT) was developed to explain human decision making behavior in dynamic tasks (Gonzalez et al., 2003). IBLT characterizes learning in dynamic tasks by storing "instances" in memory as a result of having experienced decision making events. These instances are representations of three elements: a situation (S), which is defined by a set of attributes or cues; a decision

(D), which corresponds to the action taken in situation S; and a utility or value (U), which is expected or received for making a decision D in situation S. IBLT proposes a generic decision making process through which SDU instances are built, retrieved, evaluated, and reinforced (see detailed description of this process in Gonzalez et al., 2003); with the steps consisting of: recognition (similarity-based retrieval of past instances), judgment (evaluation of the expected utility of a decision in a situation through experience or heuristics), choice (decision on when to stop information search and select the optimal current alternative), execution (implementation of the decision selected), and feedback (update of the utility of decision instances according to feedback) (see Figure above). The decision process of IBLT is determined by a set of learning mechanisms needed at different stages, including: Blending (the aggregated weighted value of alternatives involving the instance's utility weighted by its probability of retrieval), Necessity (the decision to continue or stop exploration of the environment), and Feedback (the selection of instances to be reinforced and the proportion by which the utility of these instances is reinforced). To test theories of human behavior and IBLT in particular, we use computational models: representations of some or all aspects of a theory as it applies to a particular task or context. Many IBL models have been developed in a wide variety of dynamic decision making tasks including: dynamically-complex tasks (Gonzalez & Lebiere, 2005; Martin, Gonzalez, & Lebiere, 2004), training paradigms of simple and complex tasks (Gonzalez, Best, Healy, Kole, & Bourne, 2011; Gonzalez & Dutt, 2010), simple stimulus-response practice and skill acquisition tasks (Dutt, Yamaguchi, Gonzalez, & Proctor, 2009), and repeated binary-choice tasks (Lebiere, Gonzalez, & Martin, 2007; Lejarraga et al., 2012) among others. A recent IBL model has shown generalization across multiple tasks, and accurate predictions of human choice (Gonzalez & Dutt, 2011; Lejarraga et. al., 2012; Gonzalez, Dutt, & Lejarraga, 2011). Current work involves the use of this model to predict the dynamics of trust in recommendations which have been found in behavioral studies (Harman et al., 2014).

4 Conclusion

In conclusion, dynamic decision making research may help to inform and improve the construction of recommender systems that learn and adapt their recommendations dynamically, to users' experience and to maximize benefits and overall utility from their choices.

References

1. Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. Journal of Behavioral Decision Making, 16(3), 215-233. doi: 10.1002/bdm.443
2. 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.

3. Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. Acta Psychologica, 81(3), 211-241. doi: 10.1016/0001-6918(92)90019-A
4. 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. doi: 10.1016/0747-5632(93)90005-D
5. Brunstein, A., & Gonzalez, C. (2011). Preparing for novelty with diverse training. Applied Cognitive Psychology, 25(5), 682-691. doi: 10.1002/acp.1739
6. Dörner, D. (1987). On the difficulties people have in dealing with complexity. In J. Rasmussen, K. Duncan & J. Leplat (Eds.), New Technology and Human Error (pp. 97-109). Chichester: John Wiley & Sons Ltd.
7. Dutt, V., Yamaguchi, M., Gonzalez, C., & Proctor, R. W. (2009). An instance-based learning model of stimulus-response compatibility effects in mixed location-relevant and location-irrelevant tasks. In A. Howes, D. Peebles & R. Cooper (Eds.), 9th International Conference on Cognitive Modeling – ICCM2009. Manchester, UK.
8. Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. Human Factors, 4(2), 59-73. doi: 10.1177/001872086200400201
9. Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. Psychological Review, 112(4), 912-931. doi: 10.1037/0033-295X.112.4.912
10. Forrester, J. W. (1961). Industrial dynamics. Waltham, MA: Pegasus Communications.
11. Frensch, P. A., & Funke, J. (Eds.). (1995). Complex problem solving: The European perspective. Hillsdale, NJ: Lawrence Erlbaum.
12. Funke, J. (1988). Using simulation to study complex problem solving. Simulation & Games, 19(3), 277-303. doi: 10.1177/0037550088193003
13. Gonzalez, C. (2004). Learning to make decisions in dynamic environments: Effects of time constraints and cognitive abilities. Human Factors, 46(3), 449-460. doi: 10.1518/hfes.46.3.449.50395
14. Gonzalez, C. (2005a). Decision support for real-time dynamic decision making tasks. Organizational Behavior and Human Decision Processes, 96(2), 142-154. doi: 10.1016/j.obhdp.2004.11.002
15. Gonzalez, C. (2005b). The relationship between task workload and cognitive abilities in dynamic decision making. Human Factors, 47(1), 92-101. doi: 10.1518/0018720053653767
16. Gonzalez, C. (2012). Training decisions from experience with decision making games. In P. Durlach & A. M. Lesgold (Eds.), Adaptive technologies for training and education (pp. 167-178). New York: Cambridge University Press.
17. Gonzalez, C. (2013). The boundaries of Instance-based Learning Theory for explaining decisions from experience. In V. S. Pammi & N. Srinivasan (Eds.), Progress in brain research (Vol. 202, pp. 73-98). Amsterdam, Netherlands: Elsevier.
18. Gonzalez, C., Best, B. J., Healy, A. F., Kole, J. A., & Bourne, L. E., Jr. (2011). A cognitive modeling account of simultaneous learning and fatigue effects. Journal of Cognitive Systems Research, 12(1), 19-32. doi: 10.1016/j.cogsys.2010.06.004
19. Gonzalez, C., & Dutt, V. (2010). Instance-based learning models of training. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 54(27), 2319-2323. doi: 10.1177/154193121005402721
20. Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating decisions from experience in sampling and repeated choice paradigms. Psychological Review, 118(4), 523-551. doi: 10.1037/a0024558
21. Gonzalez, C., & Dutt, V. (2012). Refuting data aggregation arguments and how the IBL model stands criticism: A reply to Hills and Hertwig (2012). Psychological Review, 119(4), 893-898. doi: 10.1037/a0029445

22. Gonzalez, C., Dutt, V., & Lejarraga, T. (2011). A loser can be a winner: Comparison of two instance-based learning models in a market entry competition. Games, 2(1), 136-162. doi: 10.3390/g2010136
23. 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 (pp. 148-165). New York: Macmillan (Palgrave Macmillan).
24. Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. Cognitive Science, 27(4), 591-635. doi: 10.1016/S0364-0213(03)00031-4
25. Gonzalez, C., & Madhavan, P. (2011). Diversity during training enhances detection of novel stimuli. Journal of Cognitive Psychology, 23(3), 342-350. doi: 10.1080/20445911.2011.507187
26. Gonzalez, C., Vanyukov, P., & Martin, M. K. (2005). The use of microworlds to study dynamic decision making. Computers in Human Behavior, 21(2), 273-286. doi: 10.1016/j.chb.2004.02.014
27. Harman, J., O'Donovan, J., Abdelzaher, T., & Gonzalez, C. (2014). Dynamics of human trust in recommender systems. Paper to be presented at the RecSys'14, Foster City, CA.
28. Harman, J. L., & Gonzalez, C. (in preparation). Allais from experience: The process of reducing Allais reversals in repeated decisions. Unpublished manuscript in preparation.
29. 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. doi: 10.1111/j.0956-7976.2004.00715.x
30. Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. Trends in Cognitive Sciences, 13(12), 517-523. doi: 10.1016/j.tics.2009.09.004
31. Hogarth, R. M. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. Psychological Bulletin, 90(2), 197-217. doi: 10.1037/0033-2909.90.2.197
32. Kerstholt, J. H. (1994). The effect of time pressure on decision-making behaviour in a dynamic task environment. Acta Psychologica, 86(1), 89-104. doi: 10.1016/0001-6918(94)90013-2
33. Kerstholt, J. H., & Raaijmakers, J. G. W. (1997). Decision making in dynamic task environments. In R. Ranyard, W. R. Crozier & O. Svenson (Eds.), Decision making: Cognitive models and explanations (pp. 205-217). London: Routledge.
34. Lebiere, C., Gonzalez, C., & Martin, M. (2007). Instance-based decision making model of repeated binary choice. Paper presented at the 8th International Conference on Cognitive Modeling, Oxford, UK.
35. Lejarraga, T., Dutt, V., & Gonzalez, C. (2012). Instance-based learning: A general model of repeated binary choice. Journal of Behavioral Decision Making, 25(2), 143-153. doi: 10.1002/bdm.722
36. 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 (pp. 178-183). Pittsburgh, PA: Lawrence Erlbaum.
37. Omodei, M. M., Wearing, A. J., McLennan, J., & Hansen, J. (2001). Human decision making in complex systems Interim summary report: Research agreement #2 (1998-2000). Melbourne: The Defence Science Technology Organisation-Information Technology Division & The University of Melbourne.
38. Rapoport, A. (1975). Research paradigms for studying dynamic decision behavior. In D. Wendt & C. Vlek (Eds.), Utility, probability, and human decision making (Vol. 11, pp. 349-375). Dordrecht, Netherlands: Reidel.

39. Rigas, G., Carling, E., & Brehmer, B. (2002). Reliability and validity of performance measures in microworlds. Intelligence, 30(5), 463-480. doi: 10.1016/S0160-2896(02)00121-6
40. Sterman, J. D. (1989). Misperceptions of feedback in dynamic decision making. Organizational Behavior and Human Decision Processes, 43(3), 301-335. doi: 10.1016/0749-5978(89)90041-1
41. Teodorescu, Kinneret, and Ido Erev. (2014). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. Psychological science. DOI: 10.1177/0956797614543022.