

Final Performance Report
AFOSR Grant #FA9550-07-1-0178

“Category Learning by Clustering
with Extension to Dynamic Environments”

Program Manager: Dr. Jun Zhang

Bradley C. Love
1 University Station A8000
The University of Texas at Austin
Austin, TX 78712

brad_love@mail.utexas.edu

(tel) 512-232-5732

(fax) 512-471-5935

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 03 MAY 2010	2. REPORT TYPE Final	3. DATES COVERED 01-03-2007 to 30-06-2010	
4. TITLE AND SUBTITLE CATEGORY LEARNING BY CLUSTERING WITH EXTENSION TO		5a. CONTRACT NUMBER FA9550-07-1-0178	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Bradley Love		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Bradley C. Love, The University of Texas at Austin, 1 University Station A8000, Austin, TX, 78712-0187		8. PERFORMING ORGANIZATION REPORT NUMBER ; AFRL-OSR-VA-TR-2011-0227	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/Dr. Jun Zhang, 4015 Wilson Blvd. Room 713, Arlington, VA, 22203		10. SPONSOR/MONITOR'S ACRONYM(S)	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-OSR-VA-TR-2011-0227	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited			
13. SUPPLEMENTARY NOTES			
14. ABSTRACT This project focuses on how humans master new categories by learning from examples with extension to dynamic environments. Decision making tends to take place in dynamic environments in which successive decisions are contingent on one another, and in which the rewards associated with actions can be delayed, yet most tasks that have been studied in the laboratory are broken up into brief, independent trials (e.g., classification of a stimulus) in which responses are determined only by the immediate context and have no bearing on future states of the task environment. Thus, this project narrows the gap between the range of mental processes typically addressed by cognitive scientists and the mental processes that underlie performance in Air Force relevant activities. We find that people's performance profiles are generally consistent with modern reinforcement learning models. For example, including perceptual information that disambiguates a person's current state within a task improves performance. Additionally, consistent with model-based predictions, people appear to hill climb on reward gradient, as opposed to globally optimize performance and show other suboptimal behavior, such as poorer performance under certain circumstance when given more information about response options.			
15. SUBJECT TERMS cognitive modeling, behavioral research, training, learning			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	
			18. NUMBER OF PAGES 6
19a. NAME OF RESPONSIBLE PERSON			

Abstract

This project focuses on how humans master new categories by learning from examples with extension to dynamic environments. Decision making tends to take place in dynamic environments in which successive decisions are contingent on one another, and in which the rewards associated with actions can be delayed, yet most tasks that have been studied in the laboratory are broken up into brief, independent trials (e.g., classification of a stimulus) in which responses are determined only by the immediate context and have no bearing on future states of the task environment. Thus, this project narrows the gap between the range of mental processes typically addressed by cognitive scientists and the mental processes that underlie performance in Air Force relevant activities. We find that people's performance profiles are generally consistent with modern reinforcement learning models. For example, including perceptual information that disambiguates a person's current state within a task improves performance. Additionally, consistent with model-based predictions, people appear to hill climb on reward gradient, as opposed to globally optimize performance, and show other suboptimal behavior, such as poorer performance under certain circumstance when given more information about response options.

Project Overview

In this project, the PI and his collaborators have made progress in understanding human category learning and have extended this work to dynamic decision making environments. Below, findings from this project are briefly reviewed. Following this review, doctoral students who have graduated during this project are listed, as our project publications.

Todd Gureckis and the PI have published a number of articles that develop the sequential learning aspects of the project. In the Cognitive Science article, we conduct a formal model comparison of simple recurrent and buffer networks and find that the simpler buffer networks do a better job of characterizing human learning and sequential performance. Surprisingly, there has been little previous fine grain evaluation of sequential learning models. We derived predictions from our buffer network and found a strong linear (through time) constraint on human sequential learning that is not present in human category learning.

In two papers, one published in the Journal of Mathematical Psychology and the other in Cognition, we explore human learning and decision making in a dynamic environment in which short- and long-term rewards are put in conflict. We find that people can learn to make long-term responses when state cues are present that de-alias underlying system states and allow for generalization of rewards to yet unexplored states. In noisy environments, we find that noise on state cues is much more detrimental to human and model performance than is equivalent noise on rewards, even though rewards define the learning problem. In fact, moderate levels of noise on rewards can be beneficial in that it encourages exploration in a task in which humans and models under explore.

We use simple reinforcement learning models to derive our study designs and characterize our results.

Three other papers have been published exploring human learning and decision making when short- and long-term rewards are in conflict. In a paper published in *Psychonomic Bulletin & Review*, we examined whether state cues make people more rational or just more sensitive to the gradient of reward as our models predict. By comparing performance when reward curves are close or far apart, we found that state cues led people to be more sensitive to reward gradient, not more rational. People still climbed toward states with increasing rewards even when doing so was not optimal. In a *Judgment and Decision Making* paper, we found (as reinforcement learning models predict) that giving additional information about forgone rewards (i.e., information about the choice option not selected) lowers performance (i.e., people meliorate and choose the short-term option). Finally, in a *Journal of Experimental Psychology: Learning, Memory, & Cognition* paper, we manipulate people's motivational focus and find a systematic effect on people's exploration strategies. In particular, people are more streaky (i.e., explore systematically by making a number of identical responses consecutively) when in a regulatory fit motivational state.

In two papers (a *Memory & Cognition* and *Psychological Science* paper), we find that people's estimation of category mean and variance is consistent with error-driven learning models that make sequential updates. In the *Psychological Science* paper, we find that people's conceptions of categories distort away from contrasting categories. The mechanisms we explore in these papers can explain high-level idealization effects.

Finally, in a second *Memory & Cognition* paper, we find evidence for two pathways for stimulus encoding. We borrow theoretical ideas from the object recognition literature. We find that one pathway that experts use is holistic and whereas the second pathway is more part-based or discrete. This latter pathway requires effortful processing to decompose and analyze stimulus parts. Although many researchers have explored the possibility that there are multiple learning systems in the brain, fewer have explored the possibility that visual stimuli can be encoded in multiple formats.

A final journal article most closely related to the proposed work is the Maddox et al. contribution. In that paper, rule-based and information-integration category learning were compared under minimal and full feedback conditions. Rule-based category structures are those for which the optimal rule is verbalizable. Information-integration category structures are those for which the optimal rule is not verbalizable. With minimal feedback subjects are told whether their response was correct or incorrect, but are not informed of the correct category assignment. With full feedback subjects are informed of the correctness of their response and are also informed of the correct category assignment. An examination of the distinct neural circuits that subserve rule-based and information-integration category learning leads to the counterintuitive prediction that full feedback should facilitate rule-based learning but should also hinder information integration learning. These predictions held. The results were

modeled by a reinforcement learning system and a Bayesian hypothesis testing system whose outputs were combined by a gating mechanism. The reinforcement learning systems processing of only feedback valence was explained by making recourse to additional dynamic tasks it subserves, like motor control and the kinds of problems considered in the aforementioned Gureckis and Love papers.

Dissertations Supported By Grant

Tyler Davis (2010). Declarative Category Learning System. University of Texas at Austin.

Marc T. Tomlinson (2010). Building BRIDGES: Combining analogy and category learning to learn relation-based categories. University of Texas at Austin.

References

Peer Reviewed Articles

Otto, A.R., Gureckis, T.M., Markman, A.B., & Love, B.C. (in press). Regulatory Fit and Systematic Exploration in a Dynamic Decision-Making Environment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.

Otto, A.R., & Love, B.C. (2010). You Don't Want To Know What You're Missing: When Information about Forgone Rewards Impedes Dynamic Decision Making. *Judgment and Decision Making*, 5, 1-10.

Davis, T., & Love, B.C. (2010). Memory for Category Information is Idealized through Contrast with Competing Options. *Psychological Science*, 21, 234-242.

Gureckis, T. M., & Love, B. C. (2010). Direct Associations or Internal Transformations? Exploring the Mechanisms Underlying Sequential Learning Behavior. *Cognitive Science*, 34, 10-50.

Gureckis, T.M., & Love, B.C. (2009). Short Term Gains, Long Term Pains: Reinforcement Learning in Dynamic Environments. *Cognition*, 113, 293-313.

Otto, A.R., Gureckis, T.M., Markman, A.B., & Love, B.C. (2009). Navigating through Abstract Decision Spaces: Evaluating the Role of State Generalization in a Dynamic Decision-Making Task. *Psychonomic Bulletin & Review*, 16, 957-963.

Davis, T., & Love, B.C. (2009). Anticipatory Emotions in Decision Tasks: Covert Markers of Value or Attentional Processes? *Cognition*, 112, 195-200.

- Davis, T., Love, B.C., & Maddox, W.T. (2009). Two Pathways to Stimulus Encoding in Category Learning? *Memory & Cognition* 37, 394-413
- Gureckis, T. M., & Love, B. C. (2009). Learning in Noise: Dynamic Decision-Making in a Variable Environment. *Journal of Mathematical Psychology*, 150, 180-193.
- Sakamoto, Y., Jones, M., & Love, B. C. (2008). Putting the Psychology Back into Psychological Models: Mechanistic vs. Rational Approaches. *Memory & Cognition*, 36, 1057-1065.
- Maddox, W. T., Love, B. C., Glass, B. D., & Filoteo, J. V. (2008). When more is less: Feedback effects in perceptual category learning. *Cognition*, 108, 578-589.

Peer Reviewed Proceedings

- Tomlinson, M.T., Howe, M., Love, B.C. (2009). Seeing the world through an expert's eyes: Context-aware display as a training companion. *Proceedings of HCI International*, LNAI 5638, 668-677.
- Love, B. C., Jones, M., Tomlinson, M.T., & Howe, M. (2009). Learning to Predict Information Needs: Context-Aware Display as a Cognitive Aid and an Assessment Tool. *Proceedings of The ACM SIGCHI Conf. on Human Factors in Computing Systems (CHI 2009)*, 1351-1360.
- Sakamoto, Y., & Love, B.C. (2009). You Only Had to Ask Me Once: Long-term Retention Requires Direct Queries During Learning. *Proceedings of the Annual Meeting of Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Otto, A.R., Gureckis, T.M., Markman, A.B., & Love, B.C. (2009). When Things Get Worse before they Get Better: Regulatory Fit and Average-Reward Learning in a Dynamic Decision-Making Environment. . *Proceedings of the Annual Meeting of Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Love, B. C., Jones, M., Tomlinson, M.T., & Howe, M. (2008). Predicting information needs: Adaptive display in dynamic environments. *Proceedings of the Annual Meeting of Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Davis, T., Love, B.C., & Maddox, W.T. (2008). How Goals Shape Category Acquisition: The Role of Contrasting Categories. *Proceedings of the Annual Meeting of Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.

Gureckis, T. M., & Love, B. C. (2007). Behaviorism Reborn? Statistical Learning as Simple Conditioning. *Proceedings of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.

Publications - Other

Love, B. C., & Tomlinson, M. (in press). Rule-based vs. similarity-based concept learning. In Denis Mareschal, Paul Quinn, Stephen Lea (Eds.), *The Emergence of Uniquely Human Concepts?*

Love, B.C. (invited encyclopedia entry). Category Learning, Computational Perspectives. In Hal Pashler, *Encyclopedia of the Mind*.

Love, B.C. (invited encyclopedia entry). Bayesian Learning. In Norbert Seel, *Encyclopedia of the Sciences of Learning*.

Tomlinson, M.T., & Love, B. C. (2008). Monkey see, monkey do: Learning relations through concrete examples. *Behavioral and Brain Sciences*, 31, 150-151.

Love, B.C., Tomlinson, M., & Gureckis, T.M. (2008). The concrete substrates of abstract rule use. In B.H. Ross, *The Psychology of Learning and Motivation*, 49, 167-207.