RUTGERS

Department of Computer Science
School of Arts and Sciences
Rutgers, the State University of New Jersey
110 Frelinghuysen Road
Piscataway, NJ 08854-8019

Rutgers Laboratory for Real Life
Reinforcement Learning (RL^3)
732-445-2001 ext 5587
Fax: 732-445-3546
mlittman@cs.rutgers.edu
http://www.cs.rutgers.edu/∼mlittman/

February 10, 2012

Hiring Committee:

Carlos Diuk joined my lab Spring 2004 and completed his dissertation research with me Summer 2010. Over the years, we worked on several projects applying learning technology to decision-making problems and have collaborated on a number of exploratory projects with colleagues in the Rutgers Center for Cognitive Science. I've found him to be consistently industrious, knowledgeable, and a pleasure to work with.

Our first project together was on simple learning problems in non-stationary environments. The main idea of the project was to detect and adapt to sudden shifts in environmental conditions. We focused primarily on the disk-spin-up problem—predicting the next time a laptop's disk will be accessed. If there will be a substantial delay, the disk should be spun down to conserve energy. If not, the disk should remain on to avoid expending the energy needed up spin it up again. Carlos' role was in collecting new data, finding sources of existing data, researching methods for "change detection", and implementing and evaluating the resulting algorithms. We found that, indeed, accurate predictions could be learned and would result in energy savings. However, due to the inherent noise in human behavior, the savings were small compared to much simpler schemes. Nevertheless, the fact that the most successful algorithm we identified came from the cognitive science literature on rodent learning made the project very interesting for both us.

Our next project was in the area of hierarchical reinforcement learning. The basic premise of this research area is that behavior learning can be accelerated through the introduction of explicit subgoals and other "behavior within a behavior" structures. We compared several existing algorithms to a new approach we developed. We found that, contrary to conventional wisdom, hierarchical structures were not critical to speed up the *acquisition* of good behaviors from experience. They did play a key *computational* role, however, speeding up the "mental" search for good behaviors. The work was published in the Autonomous Agents and Multiagent Systems (AAMAS 2006) conference.

Building on this foundation, the focus of Carlos' dissertation work was on how the selection of a task representation mediates between learning efficiency and expressibility. In a paper published with Alex Strehl (AAAI 2007), we solved the open problem of learning a Bayesian-network structure for transition dynamics. We provided theoretical guarantees and experimental validation of our algorithm. Carlos helped conceptualize the algorithm and implemented and ran all the experiments. He took the initiative and assembled a team of fellow students to follow up on this work with a more efficient algorithm and a second—robotic—application of the approach. This paper was published in the International Conference on Machine Learning 2009.

His final project at Rutgers extended our representation-learning techniques to handle environ-

ments with object-based dynamics. The basic idea here is that many higher level processes in human cognition appear to center around notions of objects of various types and predictions of how these object types behave and interact. In contrast, the representations studied in the reinforcement-learning community are considerably impoverished, in part due to their mathematical simplicity and generality. We showed how a class of algorithms could be extended to work with an object-based representation and applied it in several natural examples including an Atari video game testbed. Results were published in ICML 2008 and in his dissertation.

After graduating, Carlos joined a cognitive neuroscience lab at Princeton where he has been working with a team of psychologists to devise experiments that elucidate how humans solve reinforcement-learning tasks. His expertise in algorithm design has been of great value to the team and he has been quite effective at the kind of cross-disciplinary work that is sorely needed by both fields. He's also served as a liaison between our groups and has set up a series of joint lab meetings that have been very interesting for all involved. I've been delighted by his growth and progress. His efforts have led to co-authoring several exciting cognitive neuroscience papers. I believe he would be particularly effective at Riverside connecting Psychologists with Computer Scientists such as my colleague Prof. Christian Shelton.

Throughout his time at Rutgers, Carlos consistently proved to be a reliable partner in the creation and evaluation of novel computational approaches. He is mature and competent and was a vital contributor to the lab. I'd be remiss if I didn't note his amazing administrative skills, demonstrated in his role running the Graduate Student Association for all of Rutgers University. I strongly support his application!

Sincerely,

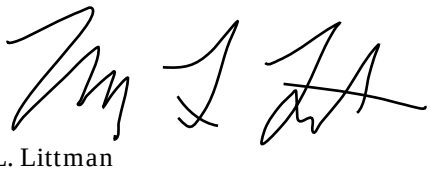

Michael L. Littman
Professor and Chair
Department of Computer Science