April 24th, 2012

Dear search committee,

Please find attached my application for the position of Assistant Professor position in the Department of Neuroscience. I specialize in learning and decision-making from a computational perspective. In particular, I have a strong theoretical and experimental background in reinforcement learning from my PhD at the Rutgers University Computer Science Department and my three years of postdoctoral research at the Neuroscience Institute and Psychology Department at Princeton University.

In my research I use computational modeling, behavioral experiments, EEG and fMRI to understand how experience shapes our representations about the world and how this knowledge enables planning of actions that maximize reward. I believe your advertised position is right in my area of expertise and I'm excited by the possibility of establishing strong collaborations within the Brown Brain Science community. For example, I foresee immediate lines of synergy with ongoing research at Brown's Department of Cognitive, Linguistic and Psychological Sciences (in particular with researchers like Michael Frank and David Badre). I believe this position is ideally suited to my background and research interest in reinforcement learning, as detailed in my accompanying statements.

Enclosed please find my curriculum vitae, a detailed research statement, a brief teaching statement and three sample papers. I have chosen papers that, although in some cases being drafts under review, I believe are the best illustrations of my current and future research directions.

Please do not hesitate to contact me at cdiuk@princeton.edu or 732-648-3042 if you request any additional information.

Sincerely,

Carlos Diuk
Psychology Department and Neuroscience Institute
Princeton University

CARLOS G. DIUK-WASSER

Department of Psychology and Neuroscience Institute
Princeton University, Princeton, NJ 08540, USA
cdiuk@princeton.edu
http://www.princeton.edu/~cdiuk
(+1) 732 648 3042

POSITIONS

POSTDOCTORAL RESEARCH ASSOCIATE with Yael Niv and Matthew Botvinick (2009-)
Department of Psychology and Neuroscience Institute, Princeton University.

VISITING PROFESSOR – UNIVERSITY OF BUENOS AIRES (*Summer 2012*)

EDUCATION

2010 PHD IN COMPUTER SCIENCE, Rutgers University-New Brunswick.
Dissertation: "An Object-oriented Representation for Efficient Reinforcement Learning". Advisor: Michael Littman.

2003 LICENCIATURA IN COMPUTER SCIENCE (6-year degree), University of Buenos Aires.
Thesis: "A Computational Tool for the Reconstruction of Genealogies".

GRANTS & AWARDS

2012 Fellow at 2nd LA School for Education, Cognitive and Neural Sciences, Argentina.
2011 Scalable Data Analytics for A Smarter Planet Grant (IBM)
Scientific Collaborator. Project: OLPC@School, Incorporating Information Technology in the Teaching and Learning Cycle.
2008 Travel Award. *International Conference on Machine Learning. Helsinki, Finland.*
2006 Travel Award. *Autonomous Agents and Multiagent Systems. Hakodate, Japan.*

PUBLICATIONS

Available pre-prints

- Carlos Diuk, Yael Niv and Matthew Botvinick
Two simultaneous but separable prediction errors in human ventral striatum. *Submitted for review.*
- Carlos Diuk, Natalia Córdova, Yael Niv and Matthew Botvinick
Discovering hierarchical task structure. *Submitted for review.*
- Carlos Diuk, Anna Schapiro, Natalia Córdova, Yael Niv and Matthew Botvinick
Divide and conquer: hierarchical reinforcement learning and task decomposition in humans. *Invited chapter in "Intrinsically Motivated Cumulative Learning in Natural and Artificial Systems" book. To appear.*
- Carlos Diuk, Diego F. Slezak, Iván Raskovsky, Mariano Sigman and Guillermo Cecchi
A quantitative philology of introspection. *Under review.*

Journals and Peer-reviewed Proceedings

- José J.F. Ribas-Fernandes, Alec Solway, Carlos Diuk, Joseph T. McGuire, Andrew G. Barto, Yael Niv and Matthew M. Botvinick (2011)
  A Neural Signature of Hierarchical Reinforcement Learning
  Neuron, 71 (2), 370-379.

- Thomas J. Walsh, Kaushik Subramanian, Michael L. Littman and Carlos Diuk (2010)
  Generalizing Apprenticeship Learning across Hypothesis Classes
  Proceedings of the Twenty-Seventh International Conference on Machine Learning (ICML-10), Haifa, Israel.

- Carlos Diuk, Lihong Li and Bethany R. Leffler (2009)
  The Adaptive k-Meteorologists Problem and Its Application to Structure Learning and Feature Selection in Reinforcement Learning
  Proceedings of the Twenty-Sixth International Conference on Machine Learning (ICML-09), Montreal, Canada.

- Thomas Walsh, Istvan Szita, Carlos Diuk and Michael Littman (2009)
  Exploring compact reinforcement-learning representations with linear regression
  Proceedings of the Twenty-Fifth Conference in Uncertainty in Artificial Intelligence (UAI-09), Montreal, Canada.

- Carlos Diuk, Andre Cohen and Michael Littman (2008)
  An Object-Oriented Representation for Efficient Reinforcement Learning
  Proceedings of the Twenty-Fifth International Conference on Machine Learning (ICML-08), Helsinki, Finland.

- Carlos Diuk and Michael Littman (2008)
  Hierarchical Reinforcement Learning
  Encyclopedia of Artificial Intelligence, IGI Global.

- Alexander L. Strehl, Carlos Diuk and Michael L. Littman (2007)
  Efficient Structure Learning in Factored-state MDPs
  Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, Vancouver, Canada.

- Carlos Diuk, Alexander Strehl and Michael Littman (2006)
  A Hierarchical Approach to Efficient Reinforcement Learning in Deterministic Domains"
  Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems. Hakodate, Japan.

- Carlos Diuk and Enrique Tándeter (2002)
  Computer tools for reconstructing a genealogy

International Journal of History and Computing, 12 (3), 329-346.

Peer-reviewed full-paper Workshops

- Ivan Raskovsky, Diego Fernández Slezak, Carlos Diuk and Guillermo Cecchi (2010)
  The emergence of the modern concept of introspection: a quantitative linguistic analysis
  North American Association for Computational Linguistics Young Researchers Workshop. Los Angeles, CA, USA.

- John Mark Agosta, Carlos Diuk, Jaideep Chandrashekar and Carl Livadas (2007)
  An adaptive anomaly detector for worm detection
  Second Workshop on Systems and Machine Learning. Cambridge, MA, USA.

- Carlos Diuk, Michael L. Littman, Alexander L. Strehl (2006)
  Efficient exploration and learning of structure in factored-state MDPs
  Towards a New Reinforcement Learning Workshop (NIPS-06). Vancouver, Canada.

- Carlos Diuk, Michael L. Littman, and Alexander L. Strehl (2005)
  A Hierarchical Approach to Efficient Reinforcement Learning in Factored State Spaces
  Rich Representations for Reinforcement Learning Workshop (ICML-05). Bonn, Germany.

Abstracts, Workshop Presentations and Posters

- Carlos Diuk, Matthew Botvinick and Yael Niv (2011)
  Two coincident but separable reward prediction errors in human ventral striatum
  Society of Neuroscience Meeting 2011. Washington, DC, USA.

- Carlos Diuk, José Ribas-Fernandes, Natalia Córdova, Matthew Botvinick and Yael Niv (2011)
  Hierarchical behavior and the brain: a reinforcement learning perspective.
  Computational and Systems Neuroscience (Cosyne-2011). Salt Lake City, Utah.

- Carlos Diuk, Matthew Botvinick, Andrew Barto and Yael Niv (2010)
  Hierarchical Reinforcement Learning: An fMRI Study of learning in a two-level gambling task
  Society of Neuroscience Meeting 2010. San Diego, CA, USA.

- Carlos Diuk (2009)
  Tools for learning reinforcement learning representations
  Fourth Barbados Workshop on Reinforcement Learning. Holetown, Barbados.

- Carlos Diuk (2008)
  Object-oriented MDPs
  Third Barbados Workshop on Reinforcement Learning. Holetown, Barbados.

- Thomas J. Walsh, Carlos Diuk and Michael Littman (2006)

Using Classifiers to Transfer Knowledge
New York Academy of Sciences Machine Learning Symposium. New York, NY, USA.

SELECTED TALKS

2011 RUTGERS UNIVERSITY - NEW BRUNSWICK.
2011 PRINCETON UNIVERSITY. *Neuroscience of Social Decision Making invited speaker.*
2008 UNIVERSITY OF MASSACHUSSETTS - AMHERST. *Invited speaker.*
2008 UNIVERSITAT POMPEU FABRA (Barcelona, Spain).
2006 AAMAS HIERARCHICAL AGENTS WORKSHOP. *Invited speaker.*
2005 UNIVERSIDAD DE BUENOS AIRES (Argentina).

OTHER RESEARCH EXPERIENCE

2007 GOOGLE (NY). *Graduate Summer Intern in the Personalized Search group.*
2006 INTEL RESEARCH. *Graduate Research Intern in the Distributed Detection and Inference group.*
2004 YALE UNIVERSITY. *Developer of web infrastructure for research projects in the Dept. of Public Health.*
1998-03 UNIVERSITY OF BUENOS AIRES. *Research Assistant for Enrique Tándeter, developing algorithms for genealogy reconstruction from historical records.*

TEACHING

2010 COMPUTER SCIENCE WINTER SCHOOL (ECI 2010) - UNIVERSITY OF BUENOS AIRES. *Invited to teach the course "Reinforcement Learning: Theory and Applications in Robotics, Games and Neuroscience."*
2008 RUTGERS UNIVERSITY – NEW BRUNSWICK. *Co-organizer of "Bayesian Reinforcement Learning" seminar.*
2003-04 RUTGERS UNIVERSITY – NEW BRUNSWICK. *Teaching Assistant for senior undergraduate class "Design & Analysis of Algorithms". Worked with Professors Martin Farach-Colton, S. Muthukrishnan and Leonid Kachiyan.*
1999-02 UNIVERSITY OF BUENOS AIRES. *Teaching Assistant and Instructor for first-year course "Algortihms and Data Structures 1". Position obtained in annual contests.*

MENTORING

- Advisor for Anthony Lenton's MSc in Data Mining at U. of Buenos Aires (ongoing)
- Co-advisor (with Yael Niv) for Dominic Kao's MSc in Computer Science at Princeton University (ongoing).
- Co-advisor (with Guillermo Cecchi and Diego Fernandez-Slezak) for Iván Raskovsky's Licenciatura in Computer Science at the University of Buenos Aires (ongoing).
- Advisor for Pablo Rodriguez Zivic's Licenciatura in Computer Science at the University of Buenos Aires (2010).

- Mentor of a group of high-school students of New Jersey Governor's School Program of Engineering and Technology (Summer 2009).

SERVICE AND PROFESSIONAL ACTIVITIES

- Journal Reviewer:
  Journal of Neuroscience, Cognition, *Neurocomputing*, *Trends in Cognitive Sciences (TICS), Journal of Artificial Intelligence Research (JAIR), Adaptive Behavior, Artificial Intelligence*, *Journal of Autonomous Agents and Multiagent Systems (JAAMAS), Frontiers in Decision Neuroscience*
- Conference Reviewer / Program Committe member:
  Neural Information Processing Systems (NIPS, multiple)
  International Conference on Machine Learning (ICML 2010, 2011, 2012)
  European Conference in Machine Learning (ECML, multiple)
  European Workshop on Reinforcement Learning (EWRL 2011, 2012)
  Association for the Advancement of Artificial Intelligence (AAAI, multiple)
  International Joint Conference in Artificial Intelligence (IJCAI 2011)
  Artificial Intelligence and Statistics (AI-Stats 2011)
  Intn'l Conf. on Pattern Recognition Apps and Methods (ICPRAM 2012, 2013)
  International Conference on Robotics and Automation (ICRA 2011)
- Program Committees:
  AAAI-11 Workshop on Lifelong Learning from Sensorimotor Experience
  ICML/UAI/COLT-09 Workshop on Abstraction in Reinforcement Learning
  AAAI Video Competition (2008, 2009, 2010)
  ECAI-08 Workshop on Machine Learning and Planning
- Grant evaluator for Netherlands Organisation for Scientific Research (NWO) – 2012.
- Grant evaluator for PICT 2010 (Argentina's Agency for Science and Technology Promotion) and UBACyT 2012 (University of Buenos Aires) grant programs.
- Chair – Technical Committee of the Reinforcement Learning Competition (2009)
- Judge in Princeton Undergraduate Research Symposium 2011.
- President of the Graduate Student Association of Rutgers University-New Brunswick (2006-07, re-elected for 2007-08).
- President of the Computer Science Graduate Student Society of Rutgers University-New Brunswick (2004-05).
- Student member of the Computer Science Governing Council at University of Buenos Aires (elected for three consecutive periods, 1998-2001).

OTHER PROFESSIONAL EXPERIENCE

2003 Project Manager at Tecnonexo (USA office).
2000-03 Independent IT Consultant for Unilever, Aguas Argentinas and Banco Río.
1998-00 Project Leader at Lemma Informatics for Movicom/Bellsouth Argentina and the Public Library System of the City of Buenos Aires.

STATEMENT OF RESEARCH INTERESTS – CARLOS DIUK

My research focuses on the neural and computational processes underlying learning and decision-making. In particular, I am interested in human behavior in the real world, where a wealth of streaming stimuli needs to be arranged into coherent representations that can be used to plan actions that maximize reward and minimize punishment. Significant progress has been made in recent years in understanding simple decision-making, both computationally and at the implementation level. My research goal is to leverage my background in computational reinforcement learning to expand our understanding of learning and decision-making in complex, real-life situations like the ones faced by humans every day. I start from rigorous computational models of learning in tasks with complex, hierarchical structure to inform the design of behavioral and neuroimaging studies.

Reinforcement Learning with Hierarchical and Object-based Representations

In recent years, the computational framework of reinforcement learning (RL) has revolutionized our understanding of how animals and humans learn to choose actions and obtain rewards. Importantly, RL theory has suggested that dopaminergic signals to the basal ganglia report a reward prediction error signal that is crucial for the well-established principle of learning via error-correction. However, RL models are fundamentally limited when applied to real-world situations. Specifically, computational RL suffers from what is known as the *curse of dimensionality*: as more environmental variables are considered, there is an exponential explosion in the size of the problem's representation, and thus in the search for its solution. Given this curse, a question emerges: how do animals and humans manage to adapt so well to their complex environments? I hypothesize that the solution involves employing a combination of abstract representations, task decompositions, smart ways of exploring the environment, and a toolset of learning heuristics. My working hypothesis is that the RL system in the brain works in a hierarchical fashion, decomposing problems into simpler sub-problems, and abstracting away irrelevant aspects of the environment appropriately for each sub-problem. My research strategy involves using these computational ideas to design behavioral and fMRI experiments to understand if and how the brain implements such mechanisms.

My research in this area started on the computational side of RL during my doctoral studies in the Rutgers Laboratory for Real-Life Reinforcement Learning, under the supervision of Michael Littman. The focus of my doctoral research was on hierarchical and object-oriented ways of representing RL problems. In my earlier work, I had shown how an artificial agent provided with an appropriate decomposition of a task into a set of simple sub-tasks could exploit this structure to significantly speed up learning. Later on, inspired by ideas about how humans parse an unknown domain and efficiently explore it, in my dissertation I showed how representing problems in terms of objects and their interactions enables the design of new RL algorithms that learn orders-of-magnitude faster than prior approaches. As an example, I studied how an artificial agent using my object-based representation would very efficiently learn the rules of a videogame by exploring how unknown objects behave and interact with each other.

My research at Rutgers University provided me with a very strong background in machine learning, in hierarchical RL, and in the problems of exploration and state abstraction. I complemented this computational training with considerable exposure to cognitive science through my work and collaboration with Randy Gallistel at the Rutgers Center for Cognitive Science. To continue this line of cross-fertilization between fields, in 2009, I joined the Neuroscience Institute and the Psychology Department at Princeton University, as a member of the laboratories of Yael Niv and Matthew Botvinick. The main aim of my postdoctoral research is to understand the

behavioral and neural underpinnings of hierarchical RL, following up on pioneering work by my mentors (Botvinick, Niv & Barto, 2009).

Hierarchical Reinforcement Learning in the Brain

The computational framework of hierarchical reinforcement learning (HRL) proposes that tasks should be broken down into simpler sub-tasks in order to tackle the curse of dimensionality. Existing approaches in the field, including some that I had developed, provide a set of well-defined predictions regarding the mechanisms underlying HRL. For example, HRL posits that an agent learning hierarchically must contain a representation for sub-tasks, and commit to efficiently perform them even though reaching sub-goals provides no primary reinforcement. HRL further stipulates that learning has to occur, often simultaneously, at multiple levels of the hierarchical representation. My postdoctoral work tested whether these predictions hold true for human learning.

Towards this end, I designed and ran fMRI experiments in a complex hierarchical task and showed that the human brain produces multiple simultaneous prediction-error signals, corresponding to the evaluation of performance at different levels of a hierarchy (Diuk et al., submitted). These signals are likely the basis for the simultaneous learning on multiple hierarchy levels that I observed behaviorally. This striking result is a breakthrough in two ways: first, it establishes that there are (at least) two teaching signals operating in the brain at once, contrary to prior belief that there is only one global, dopaminergic teaching signal. Second, it further establishes the relevance of the theoretical framework of HRL to human cognition. Additionally, using fMRI and EEG, we have been able to observe a reward prediction error signal related to subjective sub-goals (which are not externally rewarded) (Ribas-Fernandes et al., 2011). These error signals appear within structures previously implicated in RL, further attesting to an implementation of HRL in the human brain. Currently, I am exploring how humans identify useful sub-goals in tasks and exploit them to not only perform better, but also to potentially build collections of reusable skills on which to draw upon in future tasks (Diuk et al., submitted).

My work so far indicates that the RL system in the brain is indeed working hierarchically and lets us glimpse into how humans might be scaling up their learning capabilities to the challenges of the real world. My immediate plan is to continue research along these lines, performing behavioral and, when available, fMRI experiments to test specific hypotheses concerning how human learners construct and exploit hierarchical task structure. Specifically, a few questions inspired by my latest work guide my planned agenda for the next few years. First, what constitutes a “good” decomposition of a task into component sub-tasks, and how can it be produced from experience? (And complementary to that, what can we learn from how humans parse complex tasks in order to design better decision-making algorithms?) Second, how do humans leverage skills and solutions acquired in one task to better perform in a new one? How does experience shape human representations and how are representational generalizations brought to bear in order to achieve complex, high-level transfer learning? Third, given an abstract, hierarchical task representation, what is the best way to explore a new domain?

To achieve these aims, I see my research agenda as being necessarily interdisciplinary, putting computational techniques in direct contact with neural and behavioral experimental data. I believe my strong background in computational RL, my training in model-based analysis of empirical trial-by-trial data, and my neuroimaging experience put me in a unique position to achieve my long-term research goals.

TEACHING STATEMENT – CARLOS DIUK

I have a considerable background in teaching and mentoring and, in general, I find teaching enjoyable, rewarding and extremely educating. Although my teaching interests broadly parallel my research ones, I am always eager and happy to teach new topics beyond my area of expertise, as a challenge and a learning opportunity for myself.

Early on in my career, as an undergraduate at the University of Buenos Aires, I started collaborating as an assistant instructor for the first-year Algorithms course, giving a few lectures per semester on functional and imperative programming, and helping to run the programming labs. Soon, I started participating in the University's annual contests by which teaching positions were assigned, obtaining a position in all three occasions in which I participated, and ending up in top place at my last one. This experience of more than three years, starting when I was only a second-year undergraduate, was incredibly educational for me. Lecturing for a class I had only recently taken, in front of more than a hundred students, gave me confidence, a love of teaching and a deep knowledge of the material that profoundly shaped my future training years.

On my first year as a PhD student at Rutgers University, and based on my prior experience, I was assigned to teach an advanced, senior-level course in Algorithms. This was my first time lecturing in English and to a different student population than the one I was used to, which significantly enriched my teaching experience. During my weekly recitation sections I gave short lectures reviewing that week's material, and then presented a few exercises that we solved together in class, to make sure my students were prepared for upcoming homework assignments or term exams.

During the rest of my doctorate, while being funded through grants, I sought further opportunities to teach and mentor students. During the summer of 2008, I volunteered and was assigned to mentor a group of high-school students, part of the New Jersey Governor's School Program that selects some of the State's top students and provides for a two-month research experience at Rutgers University. Once again, I was exposed to a very different student population and found it extremely enjoyable. We met on a weekly basis for a three-hour session in which I taught them about reinforcement learning, brainstormed about their summer project, and monitored their progress. As their project they programmed a set of reinforcement learning agents that learned how to play a simple video game.

Last year I was invited to teach a one-week, intensive 15-hour course at the University of Buenos Aires' CS Winter School. I taught reinforcement learning theory and algorithms, focused on its connections to psychology and neuroscience. After that course, I was approached by a number of students wanting me to mentor their senior and Master-level theses, and I have been the main advisor or co-advisor for three of them so far on subjects as diverse as music perception, computational linguistics and planning and decision-making in elementary school children. I have been invited to go back to the University of Buenos Aires this Summer as a Visiting Professor. Recently, I have also co-adviced an MSc student at

Princeton, designing behavioral experiments to try to understand what makes certain videogames more engaging than others. In all cases in which I am mentoring students, I closely monitor their progress while at the same time letting them explore subjects of their interest independently. As shown by the diversity of their research topics, I think I have found a successful strategy that lets my student approach advanced theoretical material from their own particular interests.

Based on my experience, at an undergraduate level I would be most comfortable teaching courses on a diversity of topics in cognitive and computational neuroscience, but would be open and eager to teach a variety of courses. At a graduate level, I would be very interested in teaching both basic courses in Reinforcement Learning or Neuroimaging, as well as designing advanced interdisciplinary courses and seminars ranging from computational learning theory to animal and human behavior. Finally, I believe that an important aspect of my academic career will be the individual training and mentoring of students in my lab. I take this role very seriously, and plan to devote time and effort towards training my students both in terms of knowledge and tools for their specific topics of research, as well as in terms of skills that will assist them in pursuing careers as independent researchers themselves.