



April 24, 2012

David Sheinberg, PhD
Search Committee Chair
Department of Neuroscience
Brown University

Dear Prof. Sheinberg:

I'm pleased to provide this strong recommendation for Dr. Carlos Diuk-Wasser, in his application for your open faculty position in computational neuroscience. Carlos has been working as a postdoctoral fellow here – a position spanning my lab and that of Dr. Yael Niv – for the past two years, and I have a good sense of him. Over the time he's been with us, Carlos has made the transition from a highly sophisticated computer scientist to a promising experimental psychologist and cognitive neuroscientist.

Yael and I recruited Carlos from Michael Littman's lab at Rutgers, where Carlos was finishing up his thesis on model-based reinforcement learning methods. His role in our labs has been to spearhead a joint program of research studying the potential relevance of hierarchical reinforcement learning (HRL) to neuroscience. Before Carlos's arrival, Yael and I – in collaboration with Andy Barto at U. Mass. Amherst – had published a theoretical paper laying the out the possibilities, but we were only beginning to think up empirical studies to test predictions. This was Carlos's main target. More specifically, Yael and I challenged Carlos to develop a behavioral decision-making task that would permit him to use functional neuroimaging (fMRI) to test for neural correlates of HRL.

Carlos has met this challenge, developing a 'casino' task: an extension of the standard n-armed bandit paradigm, in which experimental participants first select between two 'casinos' and then decide among 'slot machines' within the casino they chose. As simple as it sounds, it is actually quite tricky to parameterize the task so that learning can be tracked at both casino and machine levels. Carlos managed to engineer this, and more importantly managed to set things up so that prediction-error signals at the two levels of the task would be experimentally dissociable.

The resulting paradigm allows Carlos not only to compare HRL models against standard RL models based on fits to behavioral data, but also to test predictions from HRL in the scanner. Carlos has run two full fMRI studies using the casino paradigm, each of which yielded very interesting results. The first study, presented at last year's Society for Neuroscience meeting, demonstrated that the brain generates parallel, concurrent prediction errors for different levels of task structure. The second study, which Carlos has just submitted for publication, used multi-voxel pattern analysis to show that these concurrent prediction error signals actually engage different

populations of neurons in ventral striatum. (I consider this one of the most interesting and novel findings ever to have come out of my lab.)

In addition to pursuing these fMRI studies, Carlos has ported the casino task for use in a monkey neurophysiology study currently underway in an allied lab at Berkeley. Beyond this, Carlos also contributed critical elements to another fMRI study on HRL, performing Bayesian model-comparison analyses that put that project across the finish line for publication in *Neuron*. He has also been running a series of behavioral studies in my lab (some currently submitted for publication), focusing on the discovery of hierarchical structure in novel tasks; and a set of studies in Yael Niv's lab focusing on attentional learning.

These pursuits and accomplishments would be considerable in someone coming out of a graduate program in Psychology or Neuroscience; they are doubly impressive in someone coming originally from a computer-science background. Carlos has shown a natural ability to move between CS and cognitive neuroscience, absorbing new ways of thinking and working without missing the direct relevance of his already established knowledge.

Although Carlos has been in our labs for a relatively short time, he has already presented his work on HRL in multiple conference appearances, and is an author on numerous papers and chapters at different stages in the pipeline. He's been quite productive, in a position that – for a lesser mind – might have carried an extended start-up period.

Carlos has been a 'hub', both in my lab and Yael's, for discussions of computational issues. Other postdocs and graduate students here – and indeed faculty members including Yael and me – think of Carlos as a go-to guy when questions arise concerning the relevance of tools from reinforcement learning, or about how those tools work. Carlos truly now has dual citizenship: He carries with him a PI-level knowledge of machine learning, but he also speaks like a true native when discussing experimental design and data analysis. He's very good at translating among CS, psychology and neuroscience, making each accessible to denizens of the other, without dumbing things down. He has really enriched the research environment in my lab.

On top of everything, Carlos is a wonderful guy. A pleasure to interact with.

In sum, I warmly recommend Carlos to you. I hope you will let me know if there's anything more I can do to support his candidacy.

Sincerely,



Matthew Botvinick, M.D., Ph.D.
Associate Professor, Psychology and Neuroscience