

Dr. Liam Paninski
Associate Professor
Department of Statistics
Center for Theoretical Neuroscience
Columbia University
1255 Amsterdam Ave
New York, NY 10027
Phone: 212-851-2166
e-mail: liam@stat.columbia.edu
<http://www.stat.columbia.edu/~liam>

April 29, 2012

Dear Professors:

It is my pleasure to strongly recommend Dr. Ed Lalor. I worked with Ed for about a year or so on what was essentially a side project for him, but I was truly impressed by his abilities in that short time. I have not followed his work closely in the last couple years, but it is clear that he has done quite a bit of work on a variety of topics, and seems to be off to a strong start as a new assistant professor in Dublin.

Ed contacted me in 2007; he was interested in working part-time with me on a project in order to gain some hands-on experience in computational neuroscience. Ed's main interests were in EEG analyses of the human visual and auditory systems, where he has done some very nice work; my work focuses more on the analysis of data from single neurons or populations of well-isolated neurons, and Ed's goal was to develop some expertise in bridging these two domains.

I was familiar with some of Ed's earlier nice work on brain-computer interfaces, so I was quite happy to work with him. I was working on a project that seemed to fit well with his interests, and he picked it up and ran with it. The question is very simple: how much information about image velocity is contained in the noisy activity of a population of retinal ganglion cells? Recent advances in large-scale multielectrode recordings and in multiple-spike-train analytical methods present us with an exciting opportunity to quantitatively connect the output of a fully-observed neural population (in the primate retina) with behavioral output (in simple psychophysical set-

tings), and Ed's project was a part of this longer-term goal.

To attack these questions properly, we needed to construct an optimal decoder, and this is where Ed was really impressive. Optimal decoding of these spike trains requires statistically-sophisticated likelihood-based techniques from point process theory. Ed immediately understood these concepts, and seemingly effortlessly wrote some (fairly involved) computer code to implement the retinal model and perform the decoding. He was aided in the high-dimensional optimal Bayesian decoding by Yashar Ahmadian, a talented postdoc here; they worked together very well. I was startled by Ed's efficiency in getting things to work and arriving at some nice scientific results, especially given that this was a very part-time project for him (a couple hours a week). We wrote up our results for a JOSA special issue on ideal observer analysis; Ed was an efficient, clear, and easy writer.

So in short, I think Ed has already been quite successful as an independent researcher. He picks up new ideas quickly; he is eager to expand his research skills in order to solve the scientific problems he's interested in; he works efficiently and is a good manager of the many demands on his time; and he is focused on producing high-quality scientific work. In addition, his research interests and strengths fit quite well with the description of the advertised position: his focus is on developing computational approaches to analyze neural information processing, with an emphasis on diagnosis and study of psychiatric disorders. I think he'd play an important role in bridging the engineering, clinical, and basic neuroscience communities at Brown, and I hope you'll take a close look at him.

Please do not hesitate to contact me with any further questions.

Sincerely,

Liam Paninski