UNIVERSITY of ROCHESTER

DEPARTMENT OF
BRAIN AND COGNITIVE SCIENCES

Alexandre Pouget, Ph.D
Full Professor
Brain and Cognitive Sciences
Center for Visual Science
Meliora Hall
University of Rochester
Rochester, NY 14627-0268
Phone: (585) 275-0760
Fax: (585) 442-9216
e-mail: alex@bcs.rochester.edu

April 19th, 2012

To whom this may concern,

This is a letter of recommendation for Xaq Pitkow, who has recently applied for a faculty position in your department. Xaq is among the top three postdocs who have joined my lab over the years, at the same level as Wei Ji Ma or Jeff Beck, and he is among the top 10% of the people I know in computational neuroscience. I'm not saying this lightly: I'm not sure I belong to the top 10% of my field... Since the very first day he joined the lab, Xaq has been more of a PI than a postdoc. His knowledge of neuroscience and computational neuroscience is well above what I expect from postdocs and more in line with the best of my colleagues. I think it's the first time in my career I had a postdoc challenge me at every lab meeting with incisive and penetrating comments on virtually any topics. I have absolutely no doubt that Xaq has a very bright future ahead of him and will undoubtedly be one of the major computational neuroscientists of his generation.

I first encountered Xaq's work when my lab discussed a paper he published in PLoS Biology. This was a study based on his thesis research that explained how neural networks can maintain high visual acuity even when eye movements smear the image across the retina. The approach he took was a particularly elegant application of Bayesian inference in the brain, and earned rare high praise from our group. I later met Xaq in person when he came to a seminar I gave at Columbia. He asked a number of insightful questions, and we continued our discussion enthusiastically for over an hour afterwards. He and I had very similar views on the most fundamental questions and promising avenues for understanding neural computation. I was therefore very pleased when he contacted me about the possibility of joining my group in Rochester.

As I have appreciated since he joined my lab, Xaq comes to neuroscience with a valuable interdisciplinary perspective. He trained in physics and biophysics at Princeton and Harvard, and has impressive mathematical skills and a great intuition for abstract mathematical concepts. There is a lot of overlap between concepts in physics and statistics but a significant barrier in language and formalism; Xaq knows both and can translate fluidly between them, which is handy in a field like computational neuroscience that benefits from both. For someone with his theoretical skills, he is unusual in having also worked in neuroscience doing electrophysiology and psychophysics experiments and analyzing data from real neurons. This experience lets him speak fluently with experimentalists, and keeps his theorizing very grounded in achievable data.

Recently Xaq presented to our department a study on optimal coding in the retina that is an elaboration of some of his thesis work. One of the textbook examples of a quantitatively successful theory in neuroscience is Atick's efficient coding theory. It says that center-surround receptive fields in the retina remove correlations in images to efficiently transmit information to the brain, and it predicts the shape of contrast sensitivity curves with impressive accuracy. However, when Xaq looked carefully at the details, he found that although the predictions match the psychophysical measurements well, they had not been rigorously tested in the retina. When he recorded from retinal neurons and analyzed their responses, he found that the linear receptive field failed to fully remove the spatial correlations. These results overturn Atick's famous theory. On the other hand, Xaq also showed that nonlinear retinal processing did remove most of the remaining correlations, and did so in a fashion that was nearly optimal according to a generalization of efficient coding theory that he developed. A paper on these results is working its way through a third round of reviews at Nature Neuroscience, where I hope it will shortly be accepted because of its considerable interest to the wider neuroscience community.

As a postdoc Xaq has been working on how the visual system can extract objects and their properties from a retinal image. This is a daunting task even for computer vision, let alone when trying to reconstruct how the brain does it. Xaq has made several significant contributions to this larger problem.

One of his contributions was in image statistics. This topic plays a crucial role in Bayesian inference, as stimulus statistics determine how uncertain evidence should be interpreted. These days, most vision research uses model images constructed by adding features together. Even though this is unrealistic, it is done because the equations are solvable. Xaq studied a harder but more realistic model that has distinct objects, and was able to compute joint probability distributions for all the relevant variables in the model images. Such probabilities are necessary for any model of Bayesian inference, and so Xaq's model is now one of the few that can be used as a quantitative benchmark for visual inference. He published this single-author study last year in the Journal of Vision, and I fully expect it will prove very useful in the field. Already, for example, Jack Gallant at Berkeley has been corresponding with Xaq about how to apply his model to decode human brain activity, and Tony Movshon's group at NYU has been recording cortical responses to stimuli Xaq designed to test for certain neural selectivities that his model predicts to be highly informative.

To apply these results to the bigger question of how neurons can extract object properties, Xaq turned to a classic model of inference, known as the Boltzmann machine in computer science and the Ising model in physics. Researchers have identified several ways of learning and performing inference in this model, but there is still no consensus about which, if any, best describes activity in the brain. Xaq developed a neural implementation of one computer inference algorithm known as belief propagation, and showed that its properties are compatible with those seen in some real neurons. When this neural network has recurrent connections tuned to his image model, its dynamics enhance the encoding of object boundaries compared to the feedforward-only network. Others have hypothesized that lateral connections served to enhance contours but Xaq was able to derive this behavior from a few simple principles and quantify how strong the effect should be. He is preparing a manuscript on these results, which he has already presented at the Computational and Systems Neuroscience conference. One very nice thing about Xaq's approach is that it can immediately be applied to inference in other perceptual domains if their prior statistics can be calculated. It then gives testable predictions about how neurons should be connected and how this connectivity should alter neural responses. There are many brain regions for which neuroscience presently has no good theories about what lateral circuitry does, and Xaq's model can provide much-needed predictions. It will be exciting to see whether these are borne out by experiments.

Incidentally, his work also makes an important contribution to computer science. Even though belief propagation has been used in a huge range of applications and extensively studied over the past two decades, Xaq found a critical flaw, unnoticed by specialists, that guaranteed poor results in a great number of applications. He then found a novel way to patch up the flaw for orders-of-magnitude better accuracy. These striking results have recently been accepted as a paper at the highly selective conference, Neural Information Processing Systems, which draws researchers around the world from neuroscience and machine learning.

Although the type of neural representation that Xaq has been using has an esteemed pedigree in physics and machine learning, previous work in my lab showed that this model does not capture certain established properties of real neurons, most notably their contrast-invariant tuning curves. We have argued instead that probabilistic population codes (PPCs) better reflect how neural populations encode uncertainty. Xaq pointed out that these representational differences hinge on whether a population of neurons encodes probabilities of a single stimulus or of multiple stimuli, a question that is still hotly debated in the field. Xaq and I have been working to combine the best aspects of these two models, and he has already found a natural way of interconnecting two neural populations that implements a simple form of the contextual integration that he has studied. He is presently refining this model and expanding it further to neural coding in larger populations with more complex statistical dependencies.

More recently, Xaq also contributed to a paper in which we argued that the main cause of behavioral variability is not neuronal noise, as assumed by most models, but suboptimal inference. We had been working on this paper for 6 years when Xaq joined my laboratory. We thought we had an important point to make but, despite numerous attempts, we never felt like we managed to make a compelling case. Xaq is the one who put us back on track. He pointed several mistakes and suggested several additions to the paper (such as the last two figures), which turned this paper around. At the time, he was not supposed to be a co-author, but his contribution turned out to be critical which is why I asked him to join us on the author list. At this point, it is fair to say that he has contributed as much to this project as anybody else, and I’d hate to speculate beyond that as I may not like what I would conclude…

To summarize, I strongly suspect that Xaq is by far the best candidate on the market this year, or one of the very best on any year. I should also mention that he is a pleasure to have around: open to new ideas, always willing to provide feedback, constantly taking initiative and always on the look-out for new approaches.

Don’t hesitate to contact me if you have any further questions.

Best,

Alexandre Pouget
Full Professor