Peter Krafft (PhD Applicant): Research Statement

Bayesian statistics provides both a powerful set of tools for quantifying our scientific uncertainty about neural, cognitive, and social processes, and a reasonable model for the way these processes deal with the uncertainty they themselves face. My fundamental interest is in understanding and advancing the capabilities of Bayesian inference as a scientific tool and as a model of inductive inference in human systems.

Prof. Andrew Barto, a pioneer in computational neuroscience, sparked my initial interest in these areas. I met Prof. Barto in my first year of university when I was actively seeking out researchers in neuroscience and cognitive science. In one of our meetings, he suggested that I read Dayan and Abbott's "Theoretical Neuroscience", which inspired me to think more deeply about computational models of learning. Prof. Barto also encouraged me to move forward with my plan to found a neuroscience club on my campus, which helped sustain my involvement with the local neuroscience community throughout my undergraduate program.[1]

Meanwhile, Prof. Andrew Cohen had invited me to be a research assistant in his Perception and Cognition Lab. In this lab I learned about experimental design in cognitive psychology, and I became aware of the role of Bayesian statistics in cognitive modeling. Intrigued by this initial exposure, I undertook a two-semester graduate-level mathematical statistics course with Prof. Michael Lavine in my second year at university. I learned about the practice of Bayesian statistics through these classes, as well as through the two research projects I did with Prof. Lavine—one on modeling the light that reaches the floor of Harvard Forest and one on diagnosing an estimation issue in a model from spatial statistics.[2] In particular I learned about how to design Bayesian models, how to use Monte Carlo methods to infer posterior distributions, and how to diagnose unexpected properties of complicated models both through close consideration of their mathematical forms and through simulation.

I combined my interests in cognitive science and Bayesian statistics in my undergraduate thesis advised by Prof. Andrew Barto and his student, Dr. George Konidaris. This project involved the recently developed deep belief network model, which uses Bayesian inference to extract hierarchically organized latent features from data. We hypothesized that a deep belief network would be able to represent the concept of *shape* in the game of Go by using its lowest level features to capture meaningful configurations of pieces, and its higher levels to capture common interactions between those configurations. Go is notoriously difficult because of its large game tree and because of the complex interactions that occur between pieces on the board, so our hope was that we could improve a Go-playing computer program by giving it access to the network's learned features. Although the results of my experiments were inconclusive, the project was a valuable introduction to machine learning for me. In this project I had to read and understand several recent technical research papers with only the high-level guidance of my advisors, since neither Prof. Barto nor Dr. Konidaris had any prior experience with deep belief networks.

During my MS program, I decided to focus on finding an application relevant to my interests that necessitated a new statistical model. With Prof. Hanna Wallach (a computer scientist and an expert in Bayesian statistics) and Prof. Bruce Desmarais (a political scientist and

[1] The website, `bcrc.bio.umass.edu/neuroscienceclub/`, lists more information and recent activities.

[2] For full papers from these and other projects, please see `www.cs.umass.edu/~pkrafft/papers/`.

an expert in network analysis), I am using the email records from the county department managers of New Hanover County, North Carolina to discover communication patterns that distinguish particularly effective managers. This research poses a major modeling challenge. Email data contain both text attributes (the subject and body of each email) and network attributes (the author and recipients of each email), but very little work exists on joint text and network modeling. By extending a model from the network literature to a framework that is appropriate for email data, and combining it with a standard probabilistic Bayesian model of text, I developed a new model that can identify with whom each manager communicates with respect to particular topics. This model will allow us to discover whether managers with broader communication strategies are more effective than managers with more targeted strategies.[3] The project has given me valuable experience modeling a complex system using Bayesian statistics.

Despite my recent work in computational social science, I would like to return to my initial training in neuroscience for my PhD. In particular, I am interested in using and developing techniques from Bayesian statistics and machine learning to understand human behavior and neural processes. The Gatsby Unit is ideal for this goal since it combines world-class researchers in machine learning and neuroscience. I am particularly interested in Prof. Dayan's work in the new field of computational psychiatry, which has potential to help psychiatric patients and to contribute to our understanding of neuroscience and psychology. For example, "A Bayesian formulation of behavioral control" by Huys and Dayan illustrated how Bayesian statistics in the context of control can help us understand depression and anxiety. Another recent paper, "Bayesian modeling of Jumping-to-Conclusions bias in delusional patients" by Moutoussis et al., explored how delusional participants in their study deviated more from the optimal behavior specified by Bayesian statistics than healthy participants did. Both of these projects address the type of problems that I would like to research during my PhD. However, I would also be interested in developing more sophisticated Bayesian methods that would be useful for modeling neural processes and that would be of interest to the machine learning community. When I met Prof. Yee Whye Teh at NIPS this year, we briefly discussed the feasibility of such a joint project. One area where this type of work seems possible would be in modeling neural spike train data.

I am well-prepared for research on these types of problems. I have a thorough knowledge of Bayesian statistics from my undergraduate and MS programs, and I have a background in neuroscience from the classes I took as an undergraduate, the research talks I attended at the neuroscience club I founded, the research assistantship I had in Prof. Cohen's Perception and Cognition Lab, and the experiences I had with Prof. Barto and his Autonomous Learning Lab. My research interests are closely in line with those of the Gatsby Unit, and obtaining a PhD from UCL would be an important step in achieving my ultimate goals of becoming a research academic and making fundamental contributions to Bayesian statistics and to our understanding of human systems.

[3] I am currently preparing a technical report describing our preliminary results. We plan to submit the completed work to the next International Conference on Machine Learning.