

Peter M Krafft
University of Washington
Information School, Box 352840
Mary Gates Hall, Suite 370
Seattle, WA 98195-2840
+1 617 960 7045
pmkrafft@uw.edu

May 7th, 2019

Academic Office
Somerville College
Woodstock Road
Oxford OX2 6HD

Dear Selection Committee,

I am delighted to be writing to apply for the Fulford Non-Stipendiary Junior Research Fellowship at Somerville College. I completed my PhD in Computer Science at the Massachusetts Institute of Technology in 2017, and am currently a Moore/Sloan & WRF Innovation in Data Science Postdoctoral Fellow at the University of Washington. My position at Oxford starting this June will be as a Senior Research Fellow at the Oxford Internet Institute, with five years of full-time departmental funding, from June 3rd 2019 to June 2nd 2024. My primary area of expertise is in the field of social data science, which touches upon data science, machine learning, artificial intelligence (AI), and computational social science. As an interdisciplinary and civically oriented researcher, I am confident that the Junior Research Fellowship would provide me an excellent opportunity to develop as a researcher and teacher in this multifaceted area, and to undertake outreach that will benefit the College, the University, and the broader public.

The research I will be conducting at Oxford centers around AI and computer science as communities of practice. The fields of AI and data science continue to grow at an outstanding pace. Yet these communities, even with this growth, still have persistent deep issues with inclusion, diversity, and open-mindedness to further-afield disciplines such as the humanities and social sciences. My ongoing and future work revolves around a set of projects-as-interventions aimed at redressing these problems in AI and data science, especially in their application to social domains. Broadly, my projects fall along two lines: (1) working from inside data and computer science communities developing methods and technologies that bridge to disciplines in the humanities and social sciences, such as by translating concepts from those disciplines in terms accessible to computer scientists, and (2) working outside the AI and data science communities to make tools for broadening AI and data literacy, and to broaden access to AI and data science resources.

Examples of specific projects in these two lines of research include:

(1) Social Considerations for Computer Science: One ongoing project in this line, associated with a \$2,000,000 (£1,500,000) DARPA grant I co-PI, examines the notion of ground truth in computational modeling of social systems. One question we ask in this work is around how the training of researchers impacts the structure of models they build and the subsets of data they prioritize analyzing. This

research is meant to respond to the perception of objectivity in computational work and help computational researchers recognize how their personal histories and positions influence their analyses.

(2) Social Considerations for Computer Science: Another ongoing project in my first line of research aims at understanding how to promote democracy and participatory design in online platforms. In this collaboration with researchers from information science and cognitive science, we apply Elinor Ostrom's Institutional Analysis and Development framework to three varied cases of online platforms, in each case recommending ways in which developing democratic constitutions for those platforms could promote sustainability and resilience. This work could help address important societal issues such as tech companies' handling of data privacy and problematic social media content like fake news.

(3) Broadening Access: In the second line of my research, one ongoing project focuses on AI literacy for a policy audience. Following work that collaborators and I have done identifying substantial gaps between expert, policy-maker, and lay understandings of AI, my collaborators and I are developing an AI literacy toolkit aimed at helping non-specialists recognize and reason about AI and algorithmic technologies. This toolkit will include (1) an interactive web application that illustrates the relationship between how machine learning models are trained and adverse social impacts; (2) a technical questionnaire for policymakers and non-experts to identify algorithmic systems and their attributes; and (3) a stepwise evaluation procedure for surfacing the social context of a given system, its technical failure modes (i.e., potential for not working correctly, such as false positives), and its social failure modes (i.e. its potential for discrimination when working correctly). Several policy-makers and policy advocates have expressed enthusiasm to us about having a toolkit like the one we are designing to support their work. There is substantial need for this kind of translational work to increase AI literacy so that effective policies can be drafted.

These projects and the lines of my research they represent will help further raise awareness of the importance of inclusion and diversity in social data science, AI, and computational social science, while also creating opportunities to bring underrepresented voices into these areas and while also paving the way towards technologies that embed these values within them.

In addition to making interventions through research, I conduct organizing and outreach work oriented towards the same goals of inclusion, diversity, and open-mindedness in computer science. I am currently organizing an MIT workshop on social responsibility in computing that will connect MIT's new college of computing with nonprofits and community organizers. I also organize events and trainings in AI and data literacy. A successful recent event drew thousands of people interested online, and hundreds of participants to our university location. One upcoming training I am leading will be a knowledge share with policy advocates in California. For this event, I will utilize an early prototype of the AI literacy toolkit I am developing in my research.

Several excellent researchers at Somerville inspire me, including at the senior level Stephen Roberts in machine learning, and Mari Mikkola in philosophy. The open-mindedness and inclusiveness of the College also provide an excellent atmosphere for my own research aligned with those values. I am excited about the opportunity to join the diverse and thriving community at Somerville.

Thank you for your time and consideration.

Sincerely,
Peter M Krafft