Bayesian Poisson Tensor Decomposition for International Relations*

Hanna Wallach

Abstract: Like their inhabitants, countries interact with one another: they consult, negotiate, trade, threaten, and fight. These interactions are seldom uncoordinated. Rather, they are connected by a fabric of overlapping communities, such as security coalitions, treaties, trade cartels, and military alliances. A single country can belong to multiple communities, reflecting its many overlapping identities, and can engage in both within- and between-community interactions, depending on the capacity in which it is acting. In this talk, I will introduce two tensor decomposition models for modeling interaction events of the form "country i took action a toward country j at time t." The first model (Bayesian Poisson CP decomposition) discovers coherent threads of events, characterized by sender countries, receiver countries, action types, and time steps; the second model (Bayesian Poisson Tucker decomposition) discovers latent country–community memberships, including the number of latent communities, as well as directed community–community interaction networks that are specific to "topics" of similar action types. I will demonstrate that these models infer interpretable latent structures that conform to and inform our knowledge of international relations. Many existing models for discrete data (such as networks and text) are special cases of these models, including infinite relational models, stochastic block models, and latent Dirichlet allocation. As a result, Bayesian Poisson tensor decomposition is a general framework for analyzing and understanding discrete data sets in the social sciences.

Bio: Hanna Wallach is a Senior Researcher at Microsoft Research New York City and an Adjunct Associate Professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst. She is also a member of UMass's Computational Social Science Institute. Hanna develops machine learning methods for analyzing the structure, content, and dynamics of social processes. Her work is inherently interdisciplinary: she collaborates with political scientists, sociologists, and journalists to understand how organizations work by analyzing publicly available interaction data, such as email networks, document collections, press releases, meeting transcripts, and news articles. To complement this agenda, she also studies issues of fairness, accountability, and transparency as they relate to machine learning. Hanna's research has had broad impact in machine learning, natural language processing, and computational social science. In 2010, her work on infinite belief networks won the best paper award at the Artificial Intelligence and Statistics conference; in 2014, she was named one of Glamour magazine's "35 Women Under 35 Who Are Changing the Tech Industry"; in 2015, she was elected to the International Machine Learning Society's Board of Trustees; and in 2016, she was named co-winner of the 2016 Borg Early Career Award. She is the recipient of several National Science Foundation grants, an Intelligence Advanced Research Projects Activity grant, and a grant from the Office of Juvenile Justice and Delinquency Prevention. Hanna is committed to increasing diversity and has worked for over a decade to address the underrepresentation of women in computing. She co-founded two projects—the first of their kind—to increase women's involvement in free and open source software development: Debian Women and the GNOME Women's Summer Outreach Program. She also co-founded the annual Women in Machine Learning Workshop, which is now in its eleventh year. Hanna holds a BA in computer science from the University of Cambridge, an MSc in cognitive science and machine learning from the University of Edinburgh, and a PhD in machine learning from the University of Cambridge.

*Machine Learning External Seminar, Gatsby Unit, September 7, 2016.