My primary research interest is to develop and use statistical machine learning methods to identify influences on human behavior.

I seek to develop computational methods to identify influences on human behavior in settings such as voting, health-care, education, media platforms, and online markets; my research is concerned with all aspects of these computational approaches, from developing statistical machine learning models of human behavior to building tools to explore and interpret model results. It is not sufficient to simply apply existing methods to problems in these domains; identifying influence in these settings requires domain knowledge to tailor methods to peculiarities of the data (e.g., non-random missingness) and processes of interest. Thus, I develop custom statistical machine learning models and inference algorithms to assist in answering questions about behavioral influence; I augment this work by building tools to increase the accessibility of the resulting model fits. I have worked extensively with unstructured text and logged user actions, and I am more broadly interested in disentangling composite signals, including feedback loops. I anticipate that, in addition to its academic significance, this work will impact practices in industry and government.

General Methodology I rely on Bayesian latent variable models, which are well-suited to exploratory data analysis because variables can map to intuitive concepts such as the "topic" of a document or the "influence" of one person on another. These variables represent assumptions about some hidden structure that was involved in the creation of the data—we do not directly observe the "topics" of a document, only the resulting words. Given a joint probability model of latent and observed variables, the central computational task is to compute or estimate the posterior distribution of the latent variables, given the observed data. The goal of exploratory tools is to then translate this posterior distribution into a visualization, browser, or navigator that is accessible to an investigator; this allows the model to be a lens through which to view the data. My research involves every stage of this process: first, working with domain experts to develop models of human behavior, then deriving, implementing, and applying scalable inference algorithms to estimate the posterior given real-world data, and finally building visualizations of the results to help interpret, validate, and critique the original models.

Unstructured Text Written text is a rich and abundant source of data that tells us about human behavior, relationships, and influences. Probabilistic topic models discover the underlying themes in collection of documents; these themes can be used to summarize, organize, explore, and analyze the corpus.

Topic models, however, are high-level statistical tools—a user must scrutinize numerical distributions to understand their results. To make these results accessible, David Blei and I developed a method for visualizing topic models [1]. Our method creates a navigator of the documents, allowing users to explore the hidden structure that a topic model discovers and understand the collection in new ways. This work included the release of open-source software; the method and its accompany software have been influential in shaping the exploration of topic models and in their application to a variety of domains. More recently, I have worked with Brandon Stewart and others to develop a new open-source tool to not only visualize, but also manually aggregate topics to be used for measurement in the social sciences [4].

In addition to understanding the general themes in large corpora, historians, political scientists, and journalists often wish to identify significant events that influence individuals and agencies. Hanna Wallach, David Blei, and I developed methods to help historians identify possible events from diplomatic messages or similarly structured text (such as email) [7, 6]. We built on topic modeling to distinguish between topics that describe "business-as-usual" and event topics that deviate from these patterns during particular periods of time. We developed a scalable variational inference algorithm for this model, as well as a visualization pipeline, and released open-source software. As an example application, we analyzed over two million diplomatic messages from the 1970s as provided by Matt Connelly's History Lab at Columbia; this analysis highlighted historical events of which our expert historian had previously been unaware.

Logged User Actions User actions, such as clicks on web posts, can also help us understand influences on behavior. Algorithmic recommendation systems use this data to uncover latent "preferences" for items and form personal recommendations based on the activity of others with similar tastes.

With David Blei and Tina Eliassi-Rad, I developed the *social Poisson factorization* (SPF) recommendation model [2]. Prior work represents users only in terms of general preferences; these models do not capture that a user may like an item because her friend likes that item. SPF models both signals, discovering both latent preferences and unobserved influence between pairs of connected users; these learned parameters can then be used to explore data. We developed scalable algorithms for analyzing data with SPF and demonstrated that it outperforms competing methods on six real-world datasets.

With Mike Gartrell, Jake Hofman, and others, I explored how group settings influence users by performing a large-scale study of television viewing habits [3]. Our analysis revealed how engagement in group viewing varies by viewer and content type, and how viewing patterns shift across various group contexts. We then constructed a simple model of how individual preferences are combined in group settings.

These recommendation systems have the potential to influence product consumption, individuals' perceptions of the world, and life-altering decisions. Deployed systems are often part of a feedback loop: models are evaluated or trained with observed data that are confounded because users are already exposed to algorithmic recommendations. With Brandon Stewart and Barbara Engelhardt, I used simulations to show that using confounded data may disadvantage algorithms in training, bias held-out evaluation, and amplify homogenization of user behavior without gains in utility [5]. I plan to continue characterizing this influence with user studies, formalize how to account for it, and build new recommendation systems that address these dynamics. This work is supported by the Princeton Center for Information Technology Policy (CITP) Interdisciplinary Seed Grant Program.

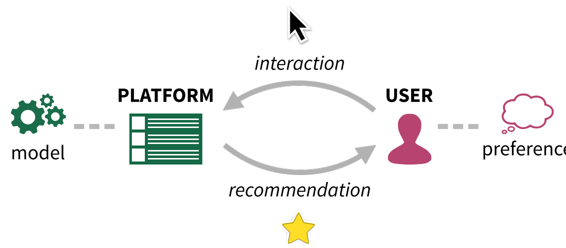

The feedback loop between user behavior and algorithmic recommendation systems. Confounding occurs when the platform model attempts to capture user preferences using observational data without accounting for past recommendations. User preferences then influence both recommendations and interactions, obfuscating the causal impact of recommendations on user behavior.

Disentangling composite signals One of the main themes of my research is to develop machine learning methods to find patterns in user behavior data; this allows us to disentangle different signals that explain that behavior and attribute weights to each source of influence. I have described my work to identify different kinds of topics in text data, as well as my research to understand various mechanisms that influence user actions including social networks, group viewing, and algorithmic recommendations.

I have also considered the broad problem of modeling collections of convolved observations, work done with Young-suk Lee, Archit Verma, and Babara Engelhardt. In this paradigm, each feature of an observation is the sum of particles that originate from distinct factors. This structure exists in data from many disciplines, including political voting among different demographics, RNA gene expression across cell types, fMRI scans of neuron activity, and financial investment across stocks and investors. We are currently working on a manuscript to present *generalized nonparametric deconvolution models*, a family of Bayesian nonparametric models for data with this structure, and study its performance on data from a variety of domains. These models learn 1) the features of global factors shared among all observations and the number and global proportions of these factors; 2) for each observation, the proportion of particles that belong to each factor; and 3) the features of observation-specific (or local) factors for each observation. While the first two objectives are fulfilled by existing models, the final objective, which we call deconvolution, is unique to our model, and allows us to ask scientific questions that are difficult to address with other models. This research is supported by an appointment to the Intelligence Community Postdoctoral Research Fellowship Program at Princeton

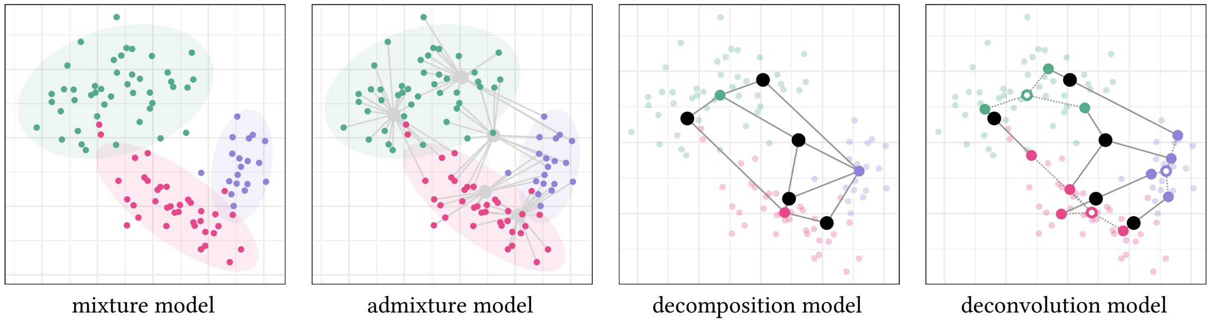

Illustrations of multiple latent variable models. ***Mixture models*** *assign each observation to one cluster or factor.* ***Admixture models*** *represent groups of observations, each with its own mixture of shared factors.* ***Decomposition models*** *decompose observations into constituent parts by representing observations as a product between group representations and factor features.* ***Deconvolution models*** *(my work) similarly decompose, or deconvolve, observations into constituent parts, but also capture group-specific (or local) fluctuations in factor features.*

University, administered by Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence.

Future Directions My research centers on developing custom statistical machine learning models and inference algorithms to identify behavioral influence; I also build tools to explore the resulting model fits. To date, I have worked to understand these signals of influence using unstructured text and logged user actions, both sparse data; I have also developed methods to learn from dense convolved, or aggregate, data. I plan to continue to build off of my expertise in modeling influences on human behavior to contribute advances in machine learning, visualization, and computational social science. I will investigate other sources of influence on human behavior in application areas such as voting, censorship, social network dynamics, education, and health-care. I will continue to explore the complex influence effects of algorithms on user behavior, including recommendation systems, search algorithms, text correction and prediction, and automatic personal assistants. In pursuing these research interests, I anticipate collaborating with faculty and students interested in machine learning, visualization, human-computer interaction, sociology, and political science.

References

[1] Allison J B Chaney and David M Blei. Visualizing topic models. In *ICWSM*, 2012.

[2] Allison J.B. Chaney, David M. Blei, and Tina Eliassi-Rad. A probabilistic model for using social networks in personalized item recommendation. In *RecSys*, RecSys '15, pages 43–50, New York, NY, USA, 2015. ACM.

[3] Allison J.B. Chaney, Mike Gartrell, Jake M. Hofman, John Guiver, Noam Koenigstein, Pushmeet Kohli, and Ulrich Paquet. A large-scale exploration of group viewing patterns. In *Proceedings of the 2014 ACM International Conference on Interactive Experiences for TV and Online Video*, TVX '14, pages 31–38, New York, NY, USA, 2014. ACM.

[4] Allison J.B. Chaney, Yuki Shiraito, and Brandon M. Stewart. The power of aggregation in topic models used for measurement. *Text as Data*, 2017.

[5] Allison J.B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. *arXiv preprint arXiv:1710.11214*, 2017.

[6] Allison J.B. Chaney, Hanna Wallach, and David M. Blei. Who, what, when, where, and why? a computational approach to understanding historical events using state department cables. *Text as Data*, 2015.

[7] Allison J.B. Chaney, Hanna Wallach, Matthew Connelly, and David M. Blei. Detecting and characterizing events. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1142–1152, 2016.