

Computational social science and social computing

Winter Mason · Jennifer Wortman Vaughan ·
Hanna Wallach

Received: 7 November 2013 / Accepted: 15 November 2013 / Published online: 27 November 2013
© The Author(s) 2013

Computational social science is an emerging research area at the intersection of computer science, statistics, and the social sciences, in which novel computational methods are used to answer questions about society. The field is inherently collaborative: social scientists provide vital context and insight into pertinent research questions, data sources, and acquisition methods, while statisticians and computer scientists contribute expertise in developing mathematical models and computational tools. New, large-scale sources of demographic, behavioral, and network data from the Internet, sensor networks, and crowdsourcing systems augment more traditional data sources to form the heart of this nascent discipline, along with recent advances in machine learning, statistics, social network analysis, and natural language processing.

The related research area of *social computing* deals with the mechanisms through which people interact with computational systems, examining questions such as how and why people contribute user-generated content and how to design systems that better enable them to do so. Examples of social computing systems include prediction markets, crowdsourcing markets, product review sites, and collaboratively edited wikis, all of which encapsulate some notion of aggregating crowd wisdom, beliefs, or ideas—albeit in different ways. Like computational social science, social computing blends techniques from machine learning and statistics with ideas from the social sciences. For example, the economics literature on incentive design has been especially influential.

W. Mason
Facebook, 335 Madison Ave, New York, NY, 10017, USA
e-mail: m@winteram.com

J.W. Vaughan (✉)
Microsoft Research, 641 Avenue of the Americas, 7th Floor, New York, NY, 10011, USA
e-mail: jenn@microsoft.com

H. Wallach
School of Computer Science, University of Massachusetts Amherst, 140 Governors Drive, Amherst,
MA, 01003, USA
e-mail: wallach@cs.umass.edu

The First and Second NIPS Workshops on Computational Social Science and the Wisdom of Crowds, held in 2010 and 2011 respectively, were established with the goal of providing a forum for workshop attendees with diverse backgrounds to meet, interact, share ideas, establish new collaborations, and inform the wider machine learning community about current research in computational social science and social computing. These workshops brought together experts from economics, political science, psychology, sociology, machine learning, statistics, and beyond, connecting researchers with common interests but different backgrounds, viewpoints, and cultural norms.

The idea for this special issue grew out of discussions at these workshops. The six featured articles were selected from seventeen submissions, each of which was carefully reviewed by at least three experts. All candidates for acceptance were asked to make thorough revisions which were evaluated by the editors and in some cases sent back to the initial reviewers. The resulting six articles provide a taste of mature research in computational social science and social computing. Their authors include psychologists, cognitive scientists, linguists, and communications scholars, as well as computer scientists.

It is well-known that in many forecasting scenarios, averaging the forecasts of a set of individuals yields a collective prediction that is more accurate than the majority of the individuals' forecasts—the so-called “wisdom of crowds” effect (Surowiecki 2004). However, there are smarter ways of aggregating forecasts than simple averaging. In “Forecast Aggregation via Recalibration,” Turner et al. present a comparison of models and algorithms for simultaneously calibrating and aggregating probabilistic forecasts about future events provided by experts who may exhibit systematic biases, such as overestimating the likelihood of rare events. The authors empirically examine whether calibration techniques should be applied before or after averaging, whether averaging should be done in probability space or log-odds space, and whether using hierarchical models improves the accuracy of forecasts. This research came about as part of the IARPA ACE program, a US government-sponsored program focused on the development of new techniques for combining the judgments of intelligence analysts, but is applicable anywhere that diverse expert forecasts are available.

For many situations, such as when estimating the opinions of a group of individuals, it is desirable to learn a classification model but there is no underlying ground truth. In other cases, it is easier or less costly to obtain noisy labels from a group of annotators (e.g., crowdsourcing sites) than obtaining the ground truth labels. In “Learning from Multiple Annotators with Varying Expertise,” Yan and colleagues present a model that estimates correct labels for data sets in which a subset of instances are annotated with noisy labels. They utilize a probabilistic model that assumes an annotator's reliability with respect to a given label depends on the input data, capturing potential variability in the annotator's expertise. As a result, the model is able to estimate annotators' reliability at the same time as estimating the most likely label. The authors demonstrate the superiority of their method on a wide range of standard classification tasks in addition to tasks that exemplify the unique contribution of their work, namely facilitating the prediction of labels for new data, the estimation of annotators' expertise given the input data, and identification of “spammer” behavior in annotators.

For decades sociologists have collected longitudinal social network data on small groups, such as children at a school or employees in a business, using a series of network surveys collected in waves. Inevitably, however, there is some failure to collect data on all nodes and all edges in all waves due to various types of attrition. In “Imputation of Missing Links and Attributes in Longitudinal Social Surveys,” Ouzienko and Obradovic develop a model based on the exponential random graph model to infer such missing data. By extending previous work, the authors are able to simultaneously infer edges and node attributes of a network

recorded over multiple time periods. Unlike prior work that focused on only one or two such features, such as link imputation over time or attribute and link imputation in a static network, their technique uses information from all three features to infer missing data. The authors compare their method with several baselines and show that their method has higher accuracy in both link and attribute inference, thereby demonstrating clear utility for social scientists conducting longitudinal network surveys.

In “Manifestations of User Personality in Website Choice and Behavior on Online Social Networks,” Kosinski and colleagues analyze a very large data set that relates people’s web browsing behavior, as well as their actions and profile information on Facebook with a validated measure of individual personality. Exploring the relationship between personality and online activity can reveal relationships that may be useful for online marketers seeking to identify personality-based segments of the population. These online behaviors can also be used as predictors for the many offline behaviors known to be associated with particular personalities. With the largest study relating online behaviors and personality to date, Kosinski et al. are able to distinguish between contrasting results from previous work that had significantly smaller sample sizes. They also use a regression model to assess how well personality can be predicted from Facebook activity, showing moderate but reasonable accuracy.

The task of identifying influential speakers in conversations is of interest to researchers in many different social science disciplines, including political science, sociology, and psychology. One widely-recognized mechanism for establishing influence in conversations is to control the topic of discussion by introducing new topics. In “Modeling Topic Control to Detect Influence in Conversations using Nonparametric Topic Models,” Nguyen et al. present a new computational model for quantitatively characterizing individuals’ tendencies to exercise control over the topic of conversation and, therefore, their influence. One of the key benefits of this model is that it infers topics, topic usage patterns, topic shifts, and individuals’ topic control in an automated, unsupervised fashion, thereby reducing annotation and training burdens on end users. The authors validate their model using real world conversations from a variety of domains, including work meetings, online discussions, and political debates. Via a range of quantitative and qualitative analyses, the authors demonstrate that their model outperforms other methods at topic segmentation, capturing topic shifts, and topic-based detection of influencers.

Many social processes involve or produce large quantities of unstructured text data. Understanding such document collections, which often are too large for any single human to read, remains an important challenge for researchers in the social sciences. Statistical topic models provide an automated way to explore the high-level themes represented in such collections and identify documents of particular interest. Despite increasing interest from social scientists, adoption of topic models into researchers’ work flows has been slow, in part because of the level of technical expertise needed to ensure such models reflect domain-specific prior knowledge and expectations. In “Interactive Topic Modeling,” Hu et al. present a novel framework for iteratively encoding users’ feedback into a topic model. By allowing users to provide the information needed to create the kinds of topics they expect, this interactive topic modeling framework facilitates large-scale analyses of document collections that are driven by the actual needs and interests of social scientists, rather than the algorithmic assumptions of computer scientists. To validate their framework, the authors undertake an extensive examination of the ways in which interactive topic modeling supports navigation of and engagement with new document collections, ranging from general interest tasks to exploration of a legislative data set focused on immigration and other political policies.

The six papers that form this special issue cover a diverse set of topics, ranging from predicting personality to identifying influencers, from labeling through crowdsourcing to

providing tools for inferring topics from large document collections. Despite their differences, these papers all share a common thread of using, and in some cases, advancing state-of-the-art machine learning methods to better understand social processes. We hope that by presenting these papers together in this special issue, their commonalities are highlighted thereby provoking further research that leverages machine learning methods to advance the fields of computational social science and social computing.

Acknowledgements We would like to thank everyone who participated in the First and Second NIPS Workshops on Computational Social Science and the Wisdom of Crowds, all of the authors who submitted articles to this special issue, and all of the referees who took the time to thoroughly review those submissions and provide valuable feedback.

References

Surowiecki, J. (2004). *The wisdom of crowds: why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. Doubleday.