

Viewpoint

Computational Social Science \neq Computer Science + Social Data


The important intersection of computer science and social science.

THIS VIEWPOINT IS about differences between computer science and social science, and their implications for *computational* social science. Spoiler alert: The punchline is simple. Despite all the hype, machine learning is not a be-all and end-all solution. We still need social scientists if we are going to use machine learning to study social phenomena in a responsible and ethical manner.

I am a machine learning researcher by training. That said, my recent work has been pretty far from traditional machine learning. Instead, my focus has been on computational social science—the study of social phenomena using digitized information and computational and statistical methods.

For example, imagine you want to know how much activity on websites such as Amazon or Netflix is caused by recommendations versus other factors. To answer this question, you might develop a statistical model for estimating causal effects from observational data such as the numbers of recommendation-based visits and numbers of total visits to individual product or movie pages over time.⁹

Alternatively, imagine you are interested in explaining when and why senators' voting patterns on particular issues deviate from what would be expected from their party affiliations and ideologies. To answer this question, you might model a set of issue-



based adjustments to each senator's ideological position using their congressional voting history and the corresponding bill text.^{4,8}

Finally, imagine you want to study the faculty hiring system in the U.S. to determine whether there is evidence of a hierarchy reflective of systematic social inequality. Here, you might model the dynamics of hiring relationships between universities over time using the placements of thousands of tenure-track faculty.³

Unsurprisingly, tackling these kinds of questions requires an interdisciplinary approach—and, indeed, computa-

tional social science sits at the intersection of computer science, statistics, and social science.

For me, shifting away from traditional machine learning and into this interdisciplinary space has meant that I have needed to think outside the algorithmic black boxes often associated with machine learning, focusing instead on the opportunities and challenges involved in developing and using machine learning methods to analyze real-world data about society.

This Viewpoint constitutes a reflection on these opportunities and challenges. I structure my discussion here

around three points—goals, models, and data—before explaining how machine learning for social science therefore differs from machine learning for other applications.

Goals

When I first started working in computational social science, I kept overhearing conversations between computer scientists and social scientists that involved sentences like, “I don’t get it—how is that even research?” And I could not understand why. But then I found this quote by Gary King and Dan Hopkins—two political scientists—that, I think, really captures the heart of this disconnect: “[C]omputer scientists may be interested in finding the needle in the haystack—such as [...] the right Web page to display from a search—but social scientists are more commonly interested in characterizing the haystack.”⁶

In other words, the conversations I kept overhearing were occurring because the goals typically pursued by computer scientists and social scientists fall into two very different categories.

The first category is prediction. Prediction is all about using observed data to reason about missing information or future, yet-to-be-observed data. To use King and Hopkins’ terminology, these are “finding the needle” tasks. In general, it is computer scientists and decision makers who are most interested in them. Sure enough, machine learning has traditionally focused on prediction tasks—such as classifying images, recognizing handwriting, and playing games like chess and Go.

The second category is explanation. Here the focus is on “why” or “how” questions—in other words, finding plausible explanations for observed data. These explanations can then be compared with established theories or previous findings, or used to generate new theories. Explanation tasks are therefore “characterizing the haystack” tasks and, in general, it is social scientists who are most interested in them. As a result, social scientists are trained to construct careful research questions with clear, testable hypotheses. For example, are women consistently excluded from long-term strategic planning in the workplace? Are government organizations more likely to comply with a public records request if they know that their peer organizations have already complied?

The goals typically pursued by computer scientists and social scientists fall into two very different categories.

Models

These different goals—prediction and explanation—lead to very different modeling approaches. In many prediction tasks, causality plays no role. The emphasis is firmly on predictive accuracy. In other words, we do not care why a model makes good predictions; we just care that it does. As a result, models for prediction seldom need to be interpretable. This means that there are few constraints on their structure. They can be arbitrarily complex black boxes that require large amounts of data to train. For example, GoogleNet, a “deep” neural network, uses 22 layers with millions of parameters to classify images into 1,000 distinct categories.¹⁰

In contrast, explanation tasks are fundamentally concerned with causality. Here, the goal is to use observed data to provide evidence in support or opposition of causal explanations. As a result, models for explanation must be interpretable. Their structure must be easily linked back to the explanation of interest and grounded in existing theoretical knowledge about the world. Many social scientists therefore use models that draw on ideas from Bayesian statistics—a natural way to express prior beliefs, represent uncertainty, and make modeling assumptions explicit.⁷

To put it differently, models for prediction are often intended to *replace* human interpretation or reasoning, whereas models for explanation are intended to *inform* or *guide* human reasoning.

Data

As well as pursuing different goals, computer scientists and social scientists typically work with different types of data. Computer scientists usually work with large-scale, digitized data-

Calendar of Events

March 5–8

HRI ‘18: ACM/IEEE International Conference on Human-Robot Interaction
Chicago, IL,
Contact: Takayuki Kanda,
Email: kanda@atr.jp

March 7–11

IUI’18: 23rd International Conference on Intelligent User Interfaces
Tokyo, Japan,
Co-Sponsored: ACM/SIG,
Contact: Shlomo Berkovsky,
Email: shlomo.berkovsky@csiro.au

March 11–15

CHIIR ‘18: Conference on Human Information Interaction and Retrieval
New Brunswick, NJ,
Sponsored: ACM/SIG,
Contact: Chirag Shah,
Email: chirags@rutgers.edu

March 15

AID ‘18: AI Decentralized 2018
Toronto, ON, Canada,
Sponsored: ACM/SIG,
Contact: Toufi Saliba,
Email: toufi@privacysell.com

March 15–16

TAU ‘18: ACM International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems
Monterey, CA,
Sponsored: ACM/SIG,
Contact: Athanasius Spyrou,
Email: tom@e-spyrou.com

March 18–21

TEI ‘18: 12th International Conference on Tangible, Embedded, and Embodied Interaction
Stockholm, Sweden,
Sponsored: ACM/SIG,
Contact: Martin Jonsson,
Email: martin.jonsson@sh.se

March 19–21

CODASPY ‘18: 7th ACM Conference on Data and Application Security and Privacy
Tempe, AZ,
Sponsored: ACM/SIG,
Contact: Gail-Joon Ahn,
Email: gahn@asu.edu

March 19–23

DATE ‘18: Design, Automation and Test in Europe
Dresden, Germany,
Contact: Jan Madsen,
Email: jama@dtu.dk

sets, often collected and made available for no particular purpose other than “machine learning research.” In contrast, social scientists often use data collected or curated in order to answer specific questions. Because this process is extremely labor intensive, these datasets have traditionally been small scale.

But—and this is one of the driving forces behind computational social science—thanks to the Internet, we now have all kinds of opportunities to obtain large-scale, digitized datasets that document a variety of social phenomena, many of which we had no way of studying previously. For example, my collaborator Bruce Desmarais and I wanted to conduct a data-driven study of local government communication networks, focusing on how political actors at the local level communicate with one another and with the general public. It turns out that most U.S. states have sunshine laws that mimic the federal Freedom of Information Act. These laws require local governments to archive textual records—including, in many states, email—and disclose them to the public upon request.

Desmarais and I therefore issued public records requests to the 100 county governments in North Carolina, requesting all non-private email messages sent and received by each county’s department managers during a randomly selected three-month time frame. Out of curiosity, we also decided to use the process of requesting these email messages as an opportunity to conduct a randomized field experiment to test whether county governments are more likely to fulfill a public records request when they are aware that their peer governments have already fulfilled the same request.

On average, we found that counties who were informed that their peers had already complied took fewer days to acknowledge our request and were more likely to actually fulfill it. And we ended up with over half a million email messages from 25 different county governments.²

Challenges

Clearly, new opportunities like this are great. But these kinds of opportunities also raise new challenges. Most conspicuously, it is very tempting to

say, “Why not use these large-scale, social datasets in combination with the powerful predictive models developed by computer scientists?” However, unlike the datasets traditionally used by computer scientists, these new datasets are often about people going about their everyday lives—their attributes, their actions, and their interactions. Not only do these datasets document social phenomena on a massive scale, they often do so at the granularity of individual people and their second-to-second behavior. As a result, they raise some complicated ethical questions regarding privacy, fairness, and accountability.

It is clear from the media that one of the things that terrifies people the most about machine learning is the use of black-box predictive models in social contexts, where it is possible to do more harm than good. There is a great deal of concern—and rightly so—that these models will reinforce existing structural biases and marginalize historically disadvantaged populations.

In addition, when datapoints are humans, error analysis takes on a whole new level of importance because errors have real-world consequences that involve people’s lives. It is not enough for a model to be 95% accurate—we need to know who is affected when there is a mistake, and in what way. For example, there is a substantial difference between a model that is 95% accurate because of noise and one that is 95% accurate because it performs perfectly for white men, but achieves only 50% accuracy when making predictions about women and minorities. Even with large datasets, there is always proportionally less data available about minorities, and statistical patterns that hold for the majority may be invalid for a given minority group. As a result, the usual machine learning objective of “good performance on average,” may be detrimental to those in a minority group.^{1,5}

Thus, when we use machine learning to reason about social phenomena—and especially when we do so to draw actionable conclusions—we have to be exceptionally careful. More so than when we use machine learning in other contexts. But here is the thing: these ethical challenges are not entirely new. Sure, they may be

new to most computer scientists, but they are not new to social scientists.

Conclusion

To me, then, this highlights an important path forward. Clearly, machine learning is incredibly useful—and, in particular, machine learning is useful for social science. But we must treat machine learning for social science very differently from the way we treat machine learning for, say, handwriting recognition or playing chess. We cannot just apply machine learning methods in a black-box fashion, as if computational social science were simply computer science plus social data. We need transparency. We need to prioritize interpretability—even in predictive contexts. We need to conduct rigorous, detailed error analyses. We need to represent uncertainty. But, most importantly, we need to work with social scientists in order to understand the ethical implications and consequences of our modeling decisions. ■

References

1. Barocas, S. and Selbst, A.D. Big data’s disparate impact. *California Law Review* 104 (2016), 671–732.
2. ben-Aaron, J. et al. Transparency by conformity: A field experiment evaluating openness in local governments. *Public Administration Review* 77, 1 (Jan. 2017), 68–77.
3. Clauaset, A., Arbesman, S., and Larremore, D.B. Systematic inequality and hierarchy in faculty hiring networks. *Science Advances* 1, 1 (Jan. 2015).
4. Gerrish, S. and Blei, D. How they vote: Issue-adjusted models of legislative behavior. In *Advances in Neural Information Processing Systems Twenty Five* (2012), 2753–2761.
5. Hardt, H. How big data is unfair; <http://bit.ly/1BBgLLr>.
6. Hopkins, D.J. and King, G. A method of automated nonparametric content analysis for social science. *American Journal of Political Science* 54, 1 (Jan. 2010), 229–247.
7. Jackman, S. *Bayesian Analysis for the Social Sciences*. Wiley, 2009.
8. Lauderdale, B.E. and Clark, T.S. Scaling politically meaningful dimensions using texts and votes. *American Journal of Political Science* 58, 3 (Mar. 2014), 754–771.
9. Sharma, A., Hofman, J., and Watts, D. Estimating the causal impact of recommendation systems from observational data. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation* (2015), 453–470.
10. Szegedy, C. et al. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2015).

Hanna Wallach (hanna@dirichlet.net) is a Senior Researcher at Microsoft Research and an Adjunct Associate Professor at the University of Massachusetts Amherst.

This article is based on an essay that appeared on Medium—see <http://bit.ly/13QIExf>. This work was supported in part by NSF grant #IIS-1320219. Any opinions, findings and conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect those of the sponsor.

Copyright held by author.