No cover image available

Oxford Handbook of Engaged Methodological Pluralism in Political Science

(In Progress)

Janet M. Box-Steffensmeier (ed.) et al.

https://doi.org/10.1093/oxfordhb/9780192868282.001.0001

Published online: 23 October 2023 **Published in print:** 15 October 2024 **Online ISBN:** 9780191964220 **Print ISBN:** 9780192868282

Search in this book

CHAPTER

A Matter of Perspective: Computational Social Science and Researcher Choice

Sarah Shugars

https://doi.org/10.1093/oxfordhb/9780192868282.013.46

Published: 18 July 2024

Abstract

We are living in a golden age for social science—a time when researchers have both unprecedented access to human data and the computational power to interrogate those data. Innovation has proceeded with remarkable speed; new methods regularly deliver novel insights with increasingly improved accuracy. These data and tools have greatly impacted the social sciences by allowing researchers to ask new questions in meaningful, new ways and to revisit classic questions with renewed rigor. While computational capacity is no longer the insurmountable barrier it once was, the opportunity presented by this methodological innovation underscores the urgent need to address another long-standing limitation of scientific advancement: the relative homogeneity of the research community. Methodological pluralism—diversity of both methods and methodologists—is essential for ensuring the collective creativity necessary to fully leverage the tools and data available to researchers. How we ask and answer questions matters. How we conceive of questions, what data and tools we bring to those questions, and how we interpret our findings are all intimately interconnected with researcher perspectives and identities. If social science ultimately aims to understand the human experience, it cannot do so without full consideration for the diversity of that experience. This paper therefore details recent innovations in computational social science and machine learning, highlighting the connections between researcher perspectives, methodological choices, and data interpretation. With a particular focus on text and network methods, we illustrate the role of researcher choice in shaping scientific advancement and demonstrate the need for human interpretation of algorithmic output. By interrogating the role of human researchers in scientific discovery we showcase the value of methodological pluralism and argue that diversity is an essential and necessary condition for scientific advancement.

Keywords: computational social science, text analysis, network analysis, interpretation, methodological pluralism

Subject: Political Methodology, Politics

Series: Oxford Handbooks

Collection: Oxford Handbooks Online

Introduction

Modern computing power, coupled with unprecedented access to vast troves of human data, has brought about a computational revolution. Falling alternately under the banner of "machine learning," "artificial intelligence," "data science," or "computational social science," the rise of "big data" and increasingly sophisticated algorithms for analyzing that data has transformed, and continues to transform, the ways in which we can study, understand, and even predict human behavior. These powerful new approaches have proven themselves to be valuable tools for examining social science questions, but they are not, unfortunately, panaceas capable of magically transforming all our scientific inquiries into perfectly accurate answers. Just like the methods that have come before them, computational approaches have their flaws and limitations. Specifically, this joint revolution in data access and algorithmic capacity is accompanied by important limitations in measurement and interpretation. As discussed in detail in the section "The Humanity of Science", these limitations are driven both by the "readymade" (Salganik 2019) nature of the data, and the high-dimensional approach of many algorithms. Furthermore, as many researchers have argued (D'ignazio and Klein 2020; Guyan 2022), both limitations highlight the essential role that human validation and interpretation play in accurately gathering meaning from algorithmic results.

While computational social science has perhaps most starkly demonstrated the seriousness of errors that can arise from disregarding human judgement, science has *always* been a distinctly human endeavor. Human scientists must intentionally select from the infinite universe of possible questions, make myriad choices about how to ask those questions, and ultimately use their human judgement in order to best interpret the answers. Scientific inquiry cannot be separated from the perspectives of the individual humans who ask questions and seek answers. Rather, scientific advancement relies on iteration and deliberation, with researchers sharing methods, models, and results in the collaborative pursuit of the truth. In other words, scientific inquiry aims to achieve the deliberative ideal, with researchers finding and evaluating arguments so as to "convince others and be convinced when it is appropriate" (Mercier and Landemore 2012).

Pluralism has long been an essential principle of deliberation (Mansbridge 2015; Habermas 1984), with democratic theorists arguing that solutions to society's most complex challenges can *only* be achieved through sharing knowledge and insights across a range of perspectives (Neblo 2015; Levine 2015). This dedication to pluralism acknowledges that the challenges a person sees and the solutions they envision are both shaped by their personal experience and perspective of the world. No single person, no matter how smart, capable, or well-informed, can possibly know and synthesize all the necessary information on their own; addressing collective challenges requires a collective effort. Indeed, the very concept of social privilege emphasizes this reality, as some people have the "privilege" to not engage with, or even be aware of, social issues that seriously impact others (McIntosh 1988; Black and Stone 2005).

Scientific inquiry similarly relies upon methodological pluralism—diversity of both methods and methodologists—in order to earnestly pursue the richest understanding of our world. Any single analysis only presents a piece of the puzzle, and any single intellectual tradition necessarily approaches that puzzle from a given vantage point. Scientific advancement therefore requires scholars to ask different questions, to pursue those questions in different ways, and to engage thoughtfully with each other in good-faith, deliberative exchange. It is not enough for scholars to merely be tolerant of other ways of thinking about and approaching research questions, they must pro-actively *engage* in pluralistic research practices—seeking out, taking seriously, and learning from diverse perspectives (Box-Steffensmeier 2022). Through this process of collective inquiry—of respectful skepticism and methodological pluralism—we can collectively move the needle of human understanding.

Of course, it is important to acknowledge here that truly valuing pluralism and elevating underrepresented perspectives are some of the biggest barriers to actually achieving the deliberative ideal (Sanders 1997). Imbalances in who has power, whose work is valued, and whose experiences are believed, all severely hamper and can actively impede true collaboration and collective advancement (D'ignazio and Klein 2020). In the context of scientific inquiry, systems of faculty hiring, tenure and promotion, and grant-making have all been shown to systematically push women, gender minorities, and people of color out of the scientific community (Morgan et al. 2021; Way et al. 2017; Clauset, Arbesman, and Larremore 2015; Mitchell and Martin 2018). While this historic challenge has long undermined efforts for scientific advancement, modern computational methods have both highlighted the need for increased methodological pluralism and raised the stakes of failing to engage diverse voices and populations in a collective process of scientific inquiry.

The Humanity of Science

While researchers now enjoy access to incredible quantities of data and computational power, these new approaches are not without their limitations. Specifically, both the data and methods employed in modern computational analysis rely heavily on researcher choice and interpretation. This human element of scientific inquiry cannot be separated from the humans who engage in that inquiry, as individuals inevitably bring their perspectives and experience to the scientific endeavor. Furthermore, as modern life is increasingly shaped by the algorithmic analysis of passively collected data, the stakes of this process have moved beyond the generation of new knowledge to actively shape and impact people's lives.

Modern data science shapes what information people see (Noble 2018; O'Neil 2016), their ability to qualify for loans and other resources (Eubanks 2018), and how they will be treated by judges and other officials (Angwin et al. 2016; Benjamin 2019). More fundamentally, the very collection of data into necessarily standardized categories results in the computational erasure of anyone who does not fit neatly into those categories—such as racial and gender minorities, including individuals who are interracial or intersex (Hanna et al. 2020; Guyan 2022; Shugars et al. 2024). While the potentially negative effects of such standardization far pre-date computational methods (Scott 2008), the datasets now used for training algorithms present an important new form of social infrastructure—one that is not only informed by existing social norms, but which actively *reinforces* and reifies those norms as reflective of a scientific truth (Scheuerman, Hanna, and Denton 2021; Buolamwini and Gebru 2018).

In other words, the application of computational methods to human subject data comes with critical ethical considerations—not only for the individual subjects whose data is used, but for society as a whole (Hu 2017; Shugars 2024). In a data-driven world, algorithms not only shape and influence our lives, but potentially dictate our very sense of reality. If the human judgement which goes into designing these algorithms fails to properly account for the full range of human experience, it not only limits the scope of scientific inquiry but potentially has significant and harmful impacts on real people's lives. For example, if researchers construct a training dataset for image classification that doesn't include sufficient training samples from dark-skinned individuals, it will result in technologies that cannot recognize, respond to, or work for people with darker skin (Buolamwini and Gebru 2018).

Given that the "big data" revolution is driven both by data availability and algorithmic innovation, the rest of this chapter elaborates on the role of human judgement along each of these dimensions. Particular attention is given to highlighting the value of methodological pluralism in advancing both scientific inquiry and social good.

Readymade Data: Conceptualization and Operationalization

Borrowing language from the art world, Salganik (2019) differentiates between readymade and custom-made data. The traditional scientific process imagines scholars developing a research question and then carefully conceiving of and collecting data in pursuit of that question. Such custom-made data certainly has flaws, but this approach comes with the benefit that researchers were involved in every step of conceptualizing and operationalizing the measures. They know exactly how their data was constructed and had the ability to make that process as controlled and homogeneous as possible. Readymade data, on the other hand, comes to researchers having been created for some other, typically non-research, purpose. For example, much of the “big data” available today involves passively collected trace data—credit card records, social media interactions, and phone geo-location are all examples of readymade data. While bringing benefits in size and scope, this readymade data presents researchers with notably more uncertainty than custom-made data. While researchers typically have some sense of how this data was generated, they often don’t have access to nuanced details of data generation, and they may not be aware of underlying heterogeneity in this process.

For example, consider custom-made and readymade approaches to measuring public opinion. In a custom-made survey, researchers can choose the exact questions they ask and test specific wording in order to ensure consistent interpretation. They can determine the format in which respondents will take the survey and, if needed, provide detailed training to enumerators in order to control how respondents hear a question. There are, of course, limitations to survey-based measures, but researchers have a well-developed collective sense of these limitations and have direct knowledge of the exact questions and procedures which generated the data they are working with.

The ready availability of social media data, on the other hand, presents an attractive alternative to this approach. Every day, millions of people publicly volunteer and record their opinions on a range of topics, creating a resource that holds the potential to capture and reflect public opinion in near real time. Yet, this readymade data comes with a higher-level of underlying uncertainty than custom-made survey data. Demographic and other covariate information about *who* is expressing an opinion is often unavailable, and content can even be generated by bots or other non-human actors such as news outlets or corporations.

Furthermore, this data does not capture a random sample of the population, but rather reflects a sample with unknown bias shaped by the underlying user base of a platform, the popularity of a user’s opinion, and even the content that is produced by others. As a public forum, social media analysis is valuable in its own right (Jackson, Bailey, and Foucault Welles 2020; Brock 2012; Shugars et al. 2021), but when attempting to operationalize this data as if it were public opinion data, researchers need to be very thoughtful about exactly what they are measuring and how they are measuring it (Joseph et al. 2021; McGregor 2019). Readymade social media data may be indicative of public opinion, but its analysis requires a vast number of methodological choices, each of which may be debated by reasonable scholars, and each of which might vary based on the specific corpus or research question (Joseph et al. 2021).

Successfully having a scientific debate about each of these measures and how they ought to be interpreted requires methodological pluralism. It requires a diversity of methodological approaches, as scholars from different intellectual traditions may see and approach the challenge in different ways. And it further requires a diversity of researcher identities, as the personal experience of individual researchers may influence the scope of options and opportunities a researcher sees.

For example, imagine a study that involves a demographic analysis of participants. While it is common for researchers to operationalize demographic identities into discrete, separable buckets, it may be important for researchers to actively conceptualize those identities through an intersectional lens which acknowledges demographic permeability. In coining the term “intersectionality,” Crenshaw (1989) was responding to a scholarly and legal shortcoming in the existing conceptualization of demographic identities. She illustrated the issue with a U.S. Supreme Court case in which five Black women sued General Motors for employment discrimination. The court sided with General Motors, finding that the company did not discriminate on the basis of race—having hired many Black men—and did not discriminate on the basis of gender—with many White, female employees. This single-axis framework of discrimination in which all people of a given gender or race were expected to share the same experience ultimately erased the unique oppression faced by the Black women plaintiffs (Crenshaw 1989). Their experience was not because they were Black, and it was not because they women. Rather, they were discriminated against precisely because of their intersectional identity as Black women.

In principle, anyone could have observed that race and gender are overlapping categories and then identified the obvious flaw in treating them as separate. Yet, it took a Black woman—someone who had the direct experience of being doubly marginalized in this way—to demonstrate this conceptual error. It simply did not occur to scholars who did not directly experience that erasure that there was a problem with this framework (McIntosh 1988). In current studies, researchers are constantly making decisions about selection and aggregation, both in terms of demographics as well other characteristics. What decisions are most appropriate depends heavily on the specific data and research questions, but debating and discussing those choices in order to identify the optimal solutions truly does require a diversity of scholarly and individual perspectives.

Algorithms: Latent Features in High-Dimensional Space

One of the greatest strengths of computational approaches is the ability to algorithmically identify patterns in high-dimensional space. Indeed, this is the core feature at the heart of many computational algorithms and one of the driving benefits of working with "big data." In these settings, "high dimensional" typically means that a single observation is associated with many features or attributes, each of which can be understood as describing a dimension of that observation.

For example, a collection of m political speeches containing a total of n unique words could be understood as an $m \times n$ document-term matrix in which every row indicates a unique speech, and every column indicates a unique word. Each row of this matrix would then provide a vector which counts how many times each word occurs within a given speech. This same vector could be further understood as describing this speech within an n-dimensional word space. Similarly, the columns of this matrix would describe word frequencies across documents, providing a vector representation of each word in m-dimensional document space. Identifying latent patterns in this $m \times n$ matrix would be an incredibly difficult task for humans, but it's a rather trivial task for computers.

Indeed, this is the core intuition behind topic modeling (Blei, Ng, and Jordan 2003; Roberts et al. 2014), a method which—given a number of topics, k—transforms this high-dimensional $m \times n$ document-term matrix into two lower-dimensional matrices. The resulting $m \times k$ matrix describes each document's loading onto k topics—essentially transforming the vector representation of each document from high-dimensional word space (n dimensions) into lower-dimensional topic space (k dimensions). The same process provides a $k \times n$ matrix in which each of the k topics describes a distribution over n words. Intuitively, this process captures the fact that "a topic" is essentially a collection of words and "a document" is a collection of topics. In this sense, the "topics" are latent features of the documents which can be inferred based on shared word frequencies, or vector representations of documents in high-dimensional "word space."

Topic modeling has been shown to be an incredibly powerful approach for understanding and categorizing documents. Yet, the reliance on computationally inferring latent features from this high-dimensional space points to the essential role that human judgement plays in interpreting algorithmic output. As Schmidt (2012) details, a key challenge in the realm of topic modeling is that humans cannot meaningfully interpret "a distribution over words." For example, in the standard topic modeling approach of Latent Dirichlet Allocation (LDA) (Blei, Ng, and Jordan 2003), a topic is represented by a vector of length n which describes that topic as a distribution over every word which occurs across all modeled documents. While mathematically sound, this is not a human-interpretable construct—neither a researcher nor a reviewer can fully understand exactly what a topic "means."

For this reason, the coherence and validity of a topic model is often determined anecdotally by manually examining the first ten words in each topic distribution. If those top words appear to be cohesive, the model is argued to capture meaningful latent dimensions of the text. This is premised from the fact that these "top words" will have the strongest loadings, i.e., will do the most to describe the topic. However, any given word typically occurs infrequently across a large corpus, meaning that even the highest-frequency words for a topic will typically have low-relative frequency. In other words, those "top words" may not actually be that meaningful or representative of the topic as a whole, leading to misinterpretation of results (Schmidt 2012). Furthermore, the substantive output of a topic model can be highly sensitive to pre-processing decisions made by researchers (Denny and Spirling 2018) and nuanced topics have little chance of being picked up by these models at all. For more discussion of the limitations of topic modeling, please see Justin Gross' chapter in this volume. Throughout the topic modeling process researchers must make decisions—often with little conceptual guidance—about which words or punctuation are meaningful. Importantly, those choices can then have substantive effects on algorithmic output.

A key challenge here is that the successful production of algorithmic output is not sufficient to guarantee that the resulting topics are meaningful—the algorithm will run and suggest topics whether those "topics" have underlying meaning or not. This puts an incredibly high burden on researchers who must determine whether the latent, topical dimensions identified by LDA are meaningful and, if so, interpret that meaning. While individual researchers must therefore aim to be extremely cautious and careful in algorithmic interpretation, this reliance on human judgement more broadly points to the need for engaged pluralism within the methodological community. Some words may have specific or nuanced meaning in a given context which may not always be apparent to researchers who are not deeply embedded within that context. Reasonable people can and should disagree about how to interpret meaning in these latent topical dimensions, but doing that with full scientific rigor requires a deliberative community which embraces pluralism and mutual respect (Mansbridge 2015).

Importantly, this interpretability challenge is pervasive in computational approaches and is inherent to the very algorithms which have transformed our methodological understanding. In network analysis, community detection is another example of how latent features are inferred from high-dimensional data in a way that requires human interpretation and validation (Peel, Larremore, and Clauset 2017).

In these approaches, "communities" are latent features which can be inferred based on the structure of network connections. The network itself is a high-dimensional object, describing relationships between a set of n items, or nodes. For example, imagine a co-sponsorship network among a group of n members of congress. In this network, any two congress people—Representative i and Representative j—are connected if they have co-sponsored a bill together. This network can then be understood as an $n \times n$ matrix capturing the connections, or edges, between every pair of congress people. Similar to the document-term matrix of the political speech example, we can understand each row and column of a network's matrix representation as describing a node in n-dimensional space, where n is the total number of nodes in the network. For example, row i could be interpreted as a length n vector describing connections from node i to every node j. Similarly, column j could be understood as a length n vector of connections from node j to every node i. In the co-sponsorship example, this matrix would be symmetric—if Representative i co-sponsors with Representative j then Representative j necessarily also co-sponsors with Representative i. However, the general matrix format can accommodate directed relationships as well—for example, you might have a network of political donations in which connections indicate that node *i* made a donation to node *j*'s campaign. In either example, a network of connections among a group of *n* actors can be formalized as an $n \times n$ matrix which provides a high-dimensional representation of a network, i.e., a representation in *n* dimensions.

The task of community detection, then, can be understood as identifying nodes which tend to have similar connections. This can then produce a lower-dimensional, $c \times c$ representation of the network, in which *c* defines the number of communities and each node is assigned to at least one community. For example, in a co-sponsorship network based on the U.S. Congress, we may find two primary communities—one comprised of Democrats who tend to co-sponsor with other Democrats, and the other comprised of Republicans tending to co-sponsors with other Republicans. While there are number of ways in which to identify these latent communities, one common intuition is that the ties within a community should be denser than the ties between communities (Fortunato 2010; Leskovec, Lang, and Mahoney 2010).

Such a rule can be implemented in a brute-force manner: check every possible number of communities, $1 \leq c \leq n$, and for each value of *c* determine the best way to assign nodes to communities such that there are more edges within communities than between communities. Across all these possible partitions, the one which best exemplifies this with-community and between-community intuition is taken to be the true community structure.

Similar to the political speech example, interpreting and validating output from a community detection algorithm is highly reliant on human judgement. Again, the mere algorithmic identification of "community structure" is not enough—human researchers must determine whether the identified structure is meaningful and, if so, must interpret that meaning. Given an $n \times n$ network matrix, these methods *will* return proposed community labels—even for randomly generated networks in which there are definitionally no true communities. This reliance on human interpretation again highlights the need for pluralism in these methodological approaches—an algorithm can tell a researcher whether or not two nodes share similar connections, but only a human can determine whether that similarity is meaningful. Researchers must therefore have deep, substantive insight into the content of the networks they are studying and must engage with a range of perspectives in order to robustly interpret meaning from these high-dimensional constructs.

There is great potential in the high-dimensional representation of data and great value in examining lower-dimensional projections of that data. Modern computational power has only recently made these algorithms tractable, but they have already transformed the ways in which we can analyze and interpret the incredible data researchers now have access to. In doing this, however, it's essential to remember that human subjectivity and judgement are an integral part of these methods. Relying so heavily on human interpretation heightens the need for pluralism and for spaces in which researchers can respectfully discuss, debate, and deliberate about a range of possible approaches and interpretations.

Discussion

We are, indeed, living in a golden age for social science. The joint revolution in data accessibility and computational power has irrevocably transformed the ways in which we can study social phenomena—creating space to ask entirely new questions in exciting new ways. However, the great power of computational social science also comes with great responsibility. The methods and approaches we're using have implications not only for our own research, but have direct and serious impacts on real people's lives. Methodological pluralism has always been important, but the stakes today are incredibly high.

Both the readymade nature of typical "big data" and the core algorithmic approach of inferring latent features from high dimensional space make modern computational methods arguably even more reliant on human interpretation than previous methodological approaches. Human judgement has always been essential for scientific advancement, and that, in turn, requires deliberative scholarly spaces. Deliberative scholars have long embraced pluralism as a core principle because no single person can ever be perfectly correct one hundred percent of the time. Strong democracies require citizens who can work together to co-create new solutions to pressing challenges, and political science requires a diverse body of scholars who can do the same.

Our human identities cannot be separated from our scholarly drives. What questions we think to ask, how we ask those questions, and how we interpret our findings are all shaped by our past knowledge and experience. We cannot ever be truly impartial because we can only ever think like us. Creating engaged pluralistic spaces in which scholars with different training and perspectives can supportively and productively share their unique insights is essential not only for advancing our science, but for ensuring equity in how that science is applied.

References

Angwin, J., J. Larson, S. Mattu, and L. Kirchner. 2016. “Machine Bias.” *ProPublica*, available online at, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing/.

Benjamin, R. 2019. “Race after Technology: Abolitionist Tools for the New Jim Code.” *Social Forces*.

Black, L. L., and D. Stone. 2005. “Expanding the Definition of Privilege: The Concept of Social Privilege.” *Journal of Multicultural Counseling and Development* 33 (4), 243–55.
Google Scholar WorldCat

Blei, D. M., A. Y. Ng, and M. I. Jordan. 2003. “Latent Dirichlet Allocation.” *The Journal of Machine Learning Research* 3, 993–1022.
Google Scholar WorldCat

Box-Steffensmeier, J. M. 2022. “Engaged Pluralism: The Importance of Commitment.” *Perspectives on Politics* 20 (1), 9–21.
Google Scholar WorldCat

Brock, A. 2012. “From the Blackhand Side: Twitter as a Cultural Conversation.” *Journal of Broadcasting & Electronic Media* 56 (4), 529–49.
Google Scholar WorldCat

Buolamwini, J., and T. Gebru. 2018. “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification.” In *Conference on Fairness, Accountability and Transparency*, PMLR 81, 77–91.

Clauset, A., S. Arbesman, and D. B. Larremore. 2015. “Systematic Inequality and Hierarchy in Faculty Hiring Networks.” *Science Advances* 1 (1), e1400005.
Google Scholar WorldCat

Crenshaw, K. 1989. “Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics.” *U. Chi. Legal F.* 1(8), 139.
Google Scholar WorldCat

Denny, M. J., and A. Spirling. 2018. “Text Preprocessing for Unsupervised Learning: Why It Matters, When It Misleads, and What to Do About it.” *Political Analysis* 26(2), 168–89.
Google Scholar WorldCat

D’ignazio, C., and L. F. Klein. 2020. *Data Feminism*. MIT Press. Cambridge, MA.
Google Scholar Google Preview WorldCat COPAC

Eubanks, V. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin’s Press. New York, NY.
Google Scholar Google Preview WorldCat COPAC

Fortunato, S. 2010. “Community Detection in Graphs.” *Physics Reports* 486 (3–5), 75–174.
Google Scholar WorldCat

Guyan, K. 2022. *Queer Data: Using Gender, Sex and Sexuality Data for Action*. Bloomsbury Publishing. London, UK.
Google Scholar Google Preview WorldCat COPAC

Habermas, J. 1984. *The Theory of Communicative Action*. Boston: Beacon Press.
Google Scholar Google Preview WorldCat COPAC

Hanna, A., E. Denton, A. Smart, and J. Smith-Loud. 2020. “Towards a Critical Race Methodology in Algorithmic Fairness.” In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 501–12.

Hu, M. 2017. “Algorithmic Jim Crow.” *Fordham L. Rev*. 86, 633.
Google Scholar WorldCat

Jackson, S. J., M. Bailey, and B. Foucault Welles. 2020. *#HashtagActivism: Networks of Race and Gender Justice*. MIT Press. Cambridge, MA.
Google Scholar Google Preview WorldCat COPAC

Joseph, K., S. Shugars, R. Gallagher, J. Green, A. Q. Math´e, Z. An, and D. Lazer. 2021. “(Mis)Alignment Between Stance Expressed in Social Media Data and Public Opinion Surveys.” In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP).

Leskovec, J., K. J. Lang, and M. Mahoney. 2010. “Empirical Comparison of Algorithms for Network Community Detection.” In Proceedings of the 19th International Conference on World Wide Web, 631–40.

Levine, P. 2015. *We Are the Ones We Have Been Waiting For: The Promise of Civic Renewal in America*. Oxford University Press.
Google Scholar Google Preview WorldCat COPAC

Mansbridge, J. 2015. “A Minimalist Definition of Deliberation.” *Deliberation and Development: Rethinking the Role of Voice and Collective Action in Unequal Societies*, Heller, Patrick, and Vijayendra Rao, eds. 27–50. World Bank Publications.
Google Scholar Google Preview WorldCat COPAC

McGregor, S. C. 2019. “Social Media as Public Opinion: How Journalists Use Social Media To Represent Public Opinion.” *Journalism* 20 (8), 1070–86.
Google Scholar WorldCat

McIntosh, P. 1988. “White Privilege: Unpacking the Invisible Knapsack.” Peace and Freedom Magazine, July/August, 10–12.

Mercier, H., and H. Landemore. 2012. “Reasoning Is for Arguing: Understanding the Successes and Failures of Deliberation.” *Political Psychology* 33 (2), 243–58.
Google Scholar WorldCat

Mitchell, K. M., and J. Martin. 2018. “Gender Bias in Student Evaluations.” *PS: Political Science & Politics* 51(3), 648–52.
Google Scholar WorldCat

Morgan, A. C., S. F. Way, M. J. Hoefer, D. B. Larremore, M. Galesic, and A. Clauset. 2021. “The Unequal Impact of Parenthood in Academia.” *Science Advances* 7 (9), eabd1996.
Google Scholar WorldCat

Neblo, M. A. 2015. *Deliberative Democracy between Theory and Practice*. Oxford University Press.
Google Scholar Google Preview WorldCat COPAC

Noble, S. U. 2018. *Algorithms of Oppression*. New York University Press.
Google Scholar Google Preview WorldCat COPAC

O’Neil, C. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Books. New York, NY.
Google Scholar Google Preview WorldCat COPAC

Peel, L., D. B. Larremore, and A. Clauset. 2017. “The Ground Truth about Metadata and Community Detection in Networks.” *Science Advances* 3 (5), e1602548.
Google Scholar WorldCat

Roberts, M. E., B. M. Stewart, D. Tingley, C. Lucas, J. Leder-Luis, S. K. Gadarian, Albertson, B., et al. 2014. “Structural Topic Models for Open-Ended Survey Responses.” *American Journal of Political Science* 58 (4), 1064–82.
Google Scholar WorldCat

Salganik, M. J. 2019. *Bit by Bit: Social Research in the Digital Age*. Princeton University Press.
Google Scholar   Google Preview   WorldCat   COPAC

Sanders, L. M. 1997. "Against Deliberation." *Political Theory* 25(3), 347–76.
Google Scholar   WorldCat

Scheuerman, M. K., A. Hanna, and E. Denton. 2021. "Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development." *Proc. ACM Hum.-Comput. Interact*. 5 (CSCW2), 1–37.
Google Scholar   WorldCat

Schmidt, B. M. 2012. "Words Alone: Dismantling Topic Models in the Humanities." *Journal of Digital Humanities* 2 (1), 49–65.
Google Scholar   WorldCat

Scott, J. C. 2008. *Seeing Like a State*. Yale University Press.
Google Scholar   Google Preview   WorldCat   COPAC

Shugars, S. 2024. "Critical Computational Social Science." *EPJ Data Science* 13(1), 1–10.
Google Scholar   WorldCat

Shugars, S., Quintana-Mathé, A., Lange, R., & Lazer, D. 2024. "Categorizing the Non-Categorical: The Challenges of Studying Gendered Phenomena Online." *Journal of Computer-Mediated Communication*, 29(1), 1–12.
Google Scholar   WorldCat

Shugars, S., A. Gitomer, S. McCabe, R. J. Gallagher, K. Joseph, N. Grinberg, Doroshenko, L., et al. 2021. "Pandemics, Protests, and Publics: Demographic Activity and Engagement on Twitter in 2020." *Journal of Quantitative Description: Digital Media*, 1(1), 1–68.
Google Scholar   WorldCat

Way, S. F., A. C. Morgan, A. Clauset, and D. B. Larremore. 2017. "The Misleading Narrative of the Canonical Faculty Productivity Trajectory." *Proceedings of the National Academy of Sciences* 114 (44), E9216–E23.
Google Scholar   WorldCat