HANDBOOK OF COMPUTATIONAL SOCIAL SCIENCE, VOLUME 1

Theory, Case Studies and Ethics

Edited by Uwe Engel, Anabel Quan-Haase, Sunny Xun Liu, and Lars Lyberg

First published 2022

ISBN: 978-0-367-45653-5 (hbk)
ISBN: 978-0-367-45652-8 (pbk)
ISBN: 978-1-003-02458-3 (ebk)

3

ANALYTICAL SOCIOLOGY AMIDST A COMPUTATIONAL SOCIAL SCIENCE REVOLUTION

Benjamin F. Jarvis, Marc Keuschnigg, and Peter Hedström

DOI: 10.4324/9781003024583-4

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

3

ANALYTICAL SOCIOLOGY AMIDST A COMPUTATIONAL SOCIAL SCIENCE REVOLUTION

Benjamin F. Jarvis, Marc Keuschnigg, and Peter Hedström

Acknowledgments: We thank Martin Arvidsson, Miriam Hurtado Bodell, Dorottya Kisfalusi, Etienne Ollion, and Christian Steglich for valuable comments. We are grateful for financial support through the Swedish Research Council (2016–01987, 2018–05170, and 445–2013–7681) and Formas (2018–00269).

Introduction

The emerging fields of computational social science (CSS) and analytical sociology (AS) have been developing in parallel over the last two decades. The development of CSS has been largely technological, driven by enormous increases in the power and availability of computational tools, the increasing digitization of contemporary life, and the corresponding broadening and deepening of data describing that life (Lazer et al., 2009, 2020). This explosion in computing capabilities and social data has inspired an interdisciplinary assortment of social and computational scientists to grapple with and mine these data for new insights into the social world (Conte et al., 2012; Watts, 2013; Ledford, 2020). The development of AS, meanwhile, has been spurred by an underlying substantive and theoretical interest in producing mechanistic explanations of collective dynamics (Hedström & Ylikoski, 2010; Keuschnigg, Lovsjö, & Hedström, 2018). By collective dynamics, we mean the emergence and transformation of system-level properties of social collectivities. AS is premised on the idea that the explanation of collective dynamics can only be achieved by understanding the social mechanisms that bring those dynamics about. This means analyzing the activities of individual actors; uncovering the social, institutional, and environmental contexts and cues that influence their actions; and demonstrating how their interdependent behaviors accumulate into macro-level social patterns and collective change (Hedström & Bearman, 2009). Drawing on this underlying philosophy, analytical sociologists have sought to achieve deeper understandings of collective phenomena such as inequality, segregation, market success, and political change.

As CSS and AS have matured, obvious connections between them have emerged. Both are keenly interested in complexity and emergence, whereby "the behavior of entities at one 'scale' of reality is not easily traced to the properties of the entities at the scale below" (Watts, 2013, p. 6; see also Anderson, 1972). And both build models that view emergence as driven by the interactions of socially embedded, interdependent actors. CSS and AS have also come to share

DOI: 10.4324/9781003024583-4

some methodological tools, including simulation approaches like agent-based modeling. But despite these connections, there remains a distinct gap between these two fields. While analytical sociologists have often used computer simulations for theoretical purposes, it is mainly within the last decade that they have begun to adopt other CSS techniques and research designs, like computational text analysis and online experiments, for empirical work. Even then, AS is often playing catch-up with the most cutting-edge CSS tools. Meanwhile, CSS research, while often making gestures to the importance of scale and emergence in social systems, all too often falls short in offering empirical confirmation of postulated behaviors at lower scales and in demonstrating how these behaviors aggregate up to real, observed patterns at larger scales.

The purpose of this chapter is to describe and strengthen the nascent connections between CSS and AS. In the process, we consider how CSS methods and perspectives have been used to augment and advance the analytical sociology project and how an analytical sociology perspective can improve the practice of CSS. We begin by offering a primer on the analytical sociology tradition. We then highlight how different CSS approaches, including empirically calibrated agent-based models, large-scale online experiments, broad and deep digital datasets, and natural language processing, fit into AS research designs. Along the way, we highlight several points for analytical sociologists to consider when incorporating CSS approaches into their research. This includes assessing the scale at which potential CSS techniques provide research insights, discerning how these tools enhance the explanatory power of AS research designs, and determining whether and how to synthesize CSS methods in AS research. Finally, we suggest how an AS perspective can guide the development of CSS and improve CSS research designs in ways that will uncover and elucidate key social mechanisms that drive collective dynamics.

Analytical sociology: a primer

This section presents core features of analytical sociology. At the same time, we distinguish its intellectual project from traditional quantitative sociology. Of course, the exact details of what analytical sociology is and how it should be practiced are not entirely settled (see, e.g., Hedström, 2005; Hedström & Bearman, 2009; Hedström & Ylikoski, 2010; Manzo, 2010, 2014; León-Medina, 2017). With that said, there is agreement that the key aspects of AS are a commitment to clarity and precision in developing mechanism-based explanations, a concern with bottom-up social processes, and an interest in achieving realism in sociological theory building.

How social systems come to evince particular patterns and how those patterns change over time or vary across contexts are main concerns for many sociologists. Much of the quantitative sociological tradition has examined these long-standing interests by focusing on "factors" rather than the interdependent behaviors of "actors" (Macy & Willer, 2002). Largely, although not exclusively, this has meant applying statistical models that identify and describe associations, or even causal relationships, between time-ordered social variables, categories, or events but with insufficient attention paid to the concrete activities, relations, and patterns of mutual influence among the social actors involved in key social processes (Sørensen, 1998). What is more, micro-level data, typically collected and analyzed based on a premise of statistical independence, came to dominate empirical social analysis (Boelaert & Ollion, 2018). This tradition has yielded at times revelatory insights into the social world, but it has rarely provided scientifically satisfying answers to questions about *how* certain social phenomena come about or *how* associations emerge between certain variables, categories, or events.

Analytical sociology distinguishes itself from this tradition by addressing these *how* questions. Analytical sociologists intend to produce mechanism-based explanations of social phenomena and their dynamics (Hedström, 2005; Hedström & Ylikoski, 2010). By a mechanism-based explanation for a social phenomenon, we mean the identification of "entities and activities organized in such a way that they are responsible for the phenomenon" (Illari & Williamson, 2012, p. 120; see also, Glennan & Illari, 2017). From the AS perspective, it is insufficient to establish a predictive relationship, even a causal one, between antecedent and subsequent variables, categories, and events. Proper explanation of a social phenomenon is achieved by specifying who the relevant actors are, theorizing the social behaviors and relations that are expected to bring about the phenomenon, and, critically, demonstrating how these "nuts and bolts [and] cogs and wheels" produce the social phenomenon in question (Elster, 1989, p. 3).

Mechanism-based explanation is premised on an ideal-typical distinction between the micro and macro levels, or scales, of a social system (Coleman, 1986). At the macro level are observed aggregate social patterns such as inequality, network density, segregation, population prevalence, and so on. These are the social phenomena to be explained. At the micro level are the actors who make up the social system and their inter-relations. In many research cases, the actors are individual persons, but in other cases, those actors may be aggregations, like families, firms, or political parties. AS researchers usually subscribe to some version of methodological individualism (Coleman, 1986, 1994; Udehn, 2002). As such, AS assigns explanatory primacy to these micro-level entities, because they – not variables, categories, or events – are capable of social action. This means that AS explanations of social processes necessarily proceed from the bottom up.

The micro and macro levels of the system are coupled. Changes in macro properties may redound to social actors at the micro level by, for example, shaping their opportunity structures, altering decision rules, shifting incentives, or channeling information. Changes in the local circumstances or thought processes of actors can lead to behaviors that, in turn, influence macro-level social patterns. This is especially likely when people are interconnected such that the actions of some will influence the actions of others, leading to domino effects. One classic example of this kind of micro-macro link is the Schelling segregation model, which shows how a population that tolerates intergroup mixing nonetheless achieves a state of near complete segregation through a process of neighborhood "tipping" induced by chains of interdependent moves (Schelling, 1971, 1978).

Finally, analytical sociologists are keenly interested in achieving realism in their research. AS aims to explain social dynamics in the real world, not in abstract, fictive social systems. Few analytical sociologists would be content to demonstrate that one particular mechanism *might* be implicated in the production of an aggregate social phenomenon. Instead, the goal is to identify actually operating mechanisms in empirically observed social systems and to demonstrate how these mechanisms, taken together, bring about the social phenomenon to be explained. This contrasts with the tendency in traditional quantitative approaches in the social sciences, most notably perhaps in economics (Friedman, 1953), to elide the often messy, heterogeneous cognitive and social processes guiding individual behavior in favor of reductive, analytically tractable models. To meet the empirical prerequisites of realistic theory building, analytical sociologists are increasingly turning to CSS methods.

New tools for sociology

Up until recently, most quantitative sociology has been limited to the conventional statistical analysis of survey data. The survey-research model and its attendant analytical tools largely rely on the sampling of independent observations. This reliance on statistical independence has often precluded the analysis of large-scale social phenomena produced by interconnected actors (Coleman, 1986). This has driven a wedge between sociological theory, primarily interested in larger-scale phenomena emerging from networked social systems, and quantitative practice, which has often narrowed its ambitions to the prediction of individual-level outcomes as functions of various temporally prior psychological and sociodemographic attributes. In short, the survey-research paradigm has steered quantitative sociology's research aims, shifting attention from the "social processes . . . shaping the system's behavior to psychological and demographic processes shaping individual behavior" (Coleman, 1986, p. 1315).

Thanks to the rise of CSS, the data and tools for realigning sociology's empirical attention with its long-standing theoretical interests are either presently or nearly at hand. Producing mechanism-based explanations of system-level properties and dynamics is predicated on an ability to access or collect datasets describing the interactions of thousands or even millions of individuals and analyzing those data to uncover the complex processes that may drive social change or maintain stability. The observations in these data are necessarily statistically dependent. With abundant computing resources and the increasing digitization of social life, not only is it now possible to gather these data, but also it is now within the realm of possibility to properly analyze them and the interdependent behaviors they describe. This has brought sociology to the cusp of transcending the survey-research paradigm and truly putting quantitative tools to work in service to sociological theorizing.

In the remainder of this section, we discuss CSS tools that analytical sociologists are using to build and investigate social theories. We discuss some of the current applications of these tools, provide suggestions for how these can be applied fruitfully in research, and suggest some potential opportunities for combining these tools. We place a particular emphasis on empirically calibrated agent-based models which have played a crucial role in developing the AS paradigm over the past two decades. Conventional agent-based models (ABMs) have occupied a niche space in sociology for several decades now (Schelling, 1971; Sakoda, 1971), but increasing computing power and programming sophistication are making it possible to apply ABMs in ways that recreate, in simulations, real-world populations and social processes. We then highlight the role of large-scale online experiments for sociological theory building. We subsequently turn to the increasing availability of broad and deep digital datasets covering networked populations and their digital traces, like those generated on social media platforms, and the potential for these datasets to revolutionize observational studies. Finally, we discuss computational text analysis, which is opening up vast, text-based troves of data describing both collective and individual sentiments, ideologies, and bodies of knowledge. We argue that CSS approaches to data collection and analysis are well suited to sociological interests in explaining the dynamics of systems populated by interdependent actors, and we consider ways in which these different approaches overlap and might be combined.

Empirically calibrated agent-based models

Empirically calibrated agent-based models (ECABMs) are becoming a key tool for demonstrating how micro-level behaviors produce macro-level outcomes not only in theory but also in real empirical settings (Bruch & Atwell, 2015). In some cases, ECABMs remain largely theoretical tools that attempt to overcome the limitations of more abstract, conventional ABMs. Building ABMs has typically involved constructing hermetically sealed, digital social worlds in which micro behaviors and the structure and form of agent interactions are entirely fabricated by the model implementer. Arbitrary assumptions coded into these models throw their real-world relevance into question. "Low-dimensional realism" ECABMs (Bruch & Atwell, 2015, p. 187), in contrast, tether some portion of the model to real data, substituting empirical findings for arbitrary assumptions. Ideally, this can yield theoretical models that are more pertinent to understanding processes in existing social systems. ECABMs are also used to achieve "high-dimensional realism" (Bruch & Atwell, 2015, p. 187). Such models integrate data about individual behaviors and aggregate outcomes so as to reproduce, as closely as possible, the collective dynamics in the observed system. The highly calibrated models can then be used to perform *in silico* experiments to judge how an intervention – possibly a policy intervention – affects a target macro-level outcome. High-dimensional-realism models are of particular interest when experimental intervention in a social system is precluded by ethical dilemmas, exorbitant costs, or practical feasibility. Programming this class of ECBMs involves thoroughly grounding the agent behaviors, characteristics, and structural positions in empirical data, the more the better. This yields greater realism but potentially at the expense of unmanageable complexity that may render opaque the ultimate mechanisms producing a macro-level phenomenon (León-Medina, 2017). And programming detailed ECABMs is no small feat: because the models are populated by many thousands, if not millions, of agents, each making dozens, hundreds, or even thousands of decisions, special care must be made to program the models efficiently and to optimize for distributed computing platforms (Deissenberg, van der Hoog, & Dawid, 2008; Collier & North, 2013). Both the low-dimensional and high-dimensional realism approaches to ECABMs have their place, and both can align with the AS desire to achieve realistic, mechanism-based explanations.

Up until now, empirically calibrated models have rarely been used in sociology. Some prominent exceptions include studies of segregation (Bruch & Mare, 2006; Xie & Zhou, 2012; Bruch, 2014), network dynamics (Snijders, 2001; Snijders, van de Bunt, & Steglich, 2010; Stadtfeld, 2018), and noncontagious disease spread (Liu & Bearman, 2015). Many applications of ECABMs have fallen on the "low-dimensional realism" side of the ECABM complexity scale. But in fields such as epidemiology, urban planning, natural resource management, and computational economics, ECABMs of the "high-dimensional realism" variety have gained wider acceptance, in part because of their utility in performing policy analyses and generating predictions (e.g., Zhang & Vorobeychik, 2019). Two prominent examples are the UrbanSIM model (Waddell, 2002) of urban land use, utilized by numerous city-planning departments, and the Global Scale Agent Model of disease transmission developed by the Brookings Institution (Epstein, 2009; Parker & Epstein, 2011) and used to study pandemics.

There are two main techniques for incorporating empirical data into ABMs: either by using micro-level data about human decision-making to calibrate agent behaviors or by fitting the full ABM model to distributional data describing the target population. Calibrating the micro-level behaviors of agents is typically done by fitting statistical models or applying machine learning algorithms to data about the relevant behaviors – such as consumption, mobility, or tie formation – to extract the parameters or decision rules that guide action (Bruch & Atwell, 2015). The identified micro-level behaviors and their parameters are then directly programmed into the ABM. In the analytical-sociology tradition, this micro-level calibration approach has been deployed most notably in studies of segregation (Bruch & Mare, 2006; Xie & Zhou, 2012; Bruch, 2014; Jarvis, 2015).

When data about micro-level behaviors are lacking or incomplete, an alternative is to directly fit an ABM to empirical data describing a self-contained social system. To do this sort of fitting, an analyst must make assumptions about micro-level agent behaviors, specify a parameterized micro-level model that links agent contexts and attributes to their behaviors, and then tune the micro-level parameters until they reproduce, in simulation, relevant macro-level statistics describing the observed system. Most prominently in sociology and network science, this approach has spawned stochastic actor-oriented models (SAOMs) for network analysis (Snijders et al., 2010). The SAOM approach to network analysis uses longitudinal network data in conjunction with assumptions about the micro-level tie-formation process to generate estimates of behavioral parameters determining tie formation and dissolution. The estimates are produced such that simulated networks match selected network-level statistics observed in the real data (e.g., degree distributions) but not to exactly reproduce the observed network. Not only can SAOMs be used to study network evolution, but also they can uncover how social influence leads to changes in distributions of node-level outcomes and behaviors, like academic achievement, alcohol consumption, and tobacco use (Steglich, Snijders, & Pearson, 2010; although see Daza & Kreuger, 2019; Ragan, Osgood, Ramirez, Moody, & Gest, 2019). For example, adams and Schaefer (2016) use this approach to examine how smoking behavior is both a cause and consequence of friendship ties in schools and to understand the implications of this endogeneity for school-level smoking prevalence. Their study illustrates the power of the SAOM approach: because the same machinery is used for both estimation and simulation, the models can account for endogenous micro-level processes during the parameter estimation stage while simultaneously providing tools for answering "what if" questions about the trajectories for the analyzed social system as a whole.

Both kinds of ECABMs – those using the independent calibration approach and those, like SAOMs, calibrated directly to data about a social system – achieve the micro-macro link in mechanistic explanation by offering a platform for performing counterfactual experiments. During counterfactual experiments, an analyst manipulates features of the model, that is, the mechanistic cogs and wheels, and examines the impact of those manipulations on aggregate outputs (Manzo, 2011; Marshall & Galea, 2015). With empirical calibration comes not only an ability to assess whether a given mechanism or set of mechanisms is capable of producing some macro-social pattern in a broad sense, as is the case with conventional ABMs, but also to analyze – using precise, quantitative measures – whether and to what degree the mechanisms in the model produce the empirically observed phenomena. This ability to simulate the social world under different conditions is one way to unpack how a social mechanism works and sort out which behaviors and structural components are needed to produce a given macro outcome. Importantly, data derived from these simulations can be used to explore whether and how those components interact as social processes unfold.

Performing a counterfactual experiment with ABMs involves either adjusting behavioral parameters and rules or altering the attributes, resources, or social contexts of agents populating the system. To begin, one can ask whether suppressing, enhancing, or otherwise modifying an aspect of the micro-level behaviors leads to changes in observed macro-level outcomes. This is typically done by modifying parameters attached to those behaviors (e.g., Snijders & Steglich, 2013). In the segregation literature, this can involve changing ethnic preferences of one or more groups and examining what different levels or patterns of segregation are realized as a result (e.g., Jarvis, 2015). In the case of network dynamics, one can alter popularity and transmission effects related to certain behaviors and trace the resulting effects on behavioral adoption (Schaefer, adams, & Haas, 2013; Lakon, Hipp, Wang, Butts, & Jose, 2015; Fujimoto, Snijders, & Valente, 2018). An alternate counterfactual approach is to modify the agent population structure, covariate distributions, or social relations without altering behavioral rules or parameters. This approach can be used to answer questions like: How does network structure shape the diffusion of social behaviors (adams & Schaefer, 2016)? How does between- and within-group inequality affect patterns of segregation (Bruch, 2014)? And how does adding or switching network ties influence mobility and segregation (Arvidsson, Collet, & Hedström, 2021)?

Both counterfactual approaches have their advantages and disadvantages. The first approach of adjusting agents' behavioral parameters is useful for understanding whether and to what degree particular behaviors contribute to macro outcomes. For example, by manipulating group preferences in a model of segregation, it should be possible to attribute the degree of segregation attained at equilibrium to the preferences of one group or another or to the interactions between these preferences. This sort of attributional exercise is of some intellectual value but may have less practical value in terms of guiding policy making or in understanding the dynamics of actually existing social systems. This is because directly modifying underlying behavioral parameters like preferences or beliefs may be tremendously difficult in practice. The second approach – adjusting variables representing relations, distributions, or other social structures – is potentially more promising for understanding policy interventions and the behavior of real social systems. This is because these interventions are more plausible – they take people as they are rather than as we might like them to be – and consider how people would behave with slightly different resources, incentives, or social influences.

Deciphering why and how a particular ABM model produces a particular macro phenomenon under some counterfactual conditions but not others is a mounting challenge in the AS and ABM community. There are increasing calls to open up ABM "black boxes" to uncover the precise chains of behavior that bring about the macro outcomes of interest (León-Medina, 2017). There are no easy solutions to these problems, but other CSS techniques may offer avenues for further exploration. Pattern recognition approaches, like sequence analysis (Cornwell, 2015; Ritschard & Studer, 2018), might be fruitfully applied not only to empirical data but also to the outputs of agent-based models to understand the sequences of events that account for macro-level outcomes that differ between *in silico* counterfactual experiments. Using these approaches might allow AS to move beyond identifying the relative importance of different "cogs and wheels" that bring macro outcomes about and towards a fuller account of how these parts fit together in causal chains of transformative social action.

Large-scale web-based experiments

Increasingly, sociologists make use of web-based experiments not only to study behavior among participant groups that are more diverse than those sustained by traditional laboratory pools (e.g., Bader & Keuschnigg, 2020; Schaub, Gereke, & Baldassarri, 2020) but also, importantly, to elicit behavior of larger groups of interacting population members (Salganik & Watts, 2009; Centola, 2018). This type of experiment – introduced to the analytical sociology toolbox in 2006 by a study of music downloading in artificial cultural markets (Salganik, Dodds, & Watts, 2006) – focuses on the social processes arising from behavioral interdependencies, and it tackles questions about the "social production" of macro phenomena that cannot be addressed using the more individualistic perspectives taken in small-group experiments. This design breaks with the older experimental traditions by randomizing participants into separate "social systems", or "multiple worlds", that each serve as a unit of analysis.

During the experiment, entire miniature social systems are populated by real users, with each miniature system varying in the conditions guiding social interactions and user choices. By manipulating conditions in theoretically informed ways, it can be determined which interactive mechanisms produce system-level outcomes. Interpreting each social system, rather than each experimental subject, as a unit of analysis shifts the focus from individual action towards macro phenomena. Participants in an experimental run learn about the others' decisions either through direct interaction – such as on a social media-like platform that signals network contacts' adoptions of certain behaviors (Centola, 2010, 2011) – or through statistics that summarize the aggregate behaviors of other participants – such as a popularity ranking derived from past participants' choices (Salganik et al., 2006; Macy, Deri, Ruch, & Tong, 2019). Experimental runs start on an all-to-all network and are left alone to evolve endogenously, or they start from predefined worlds, for example, by placing participants on a network with a certain number of homophilic contacts. Each run then resembles a realization of a social process in a population of interconnected individuals, and it provides controlled data on the social interactions that lead to a specific collective outcome.

Observing larger groups of interacting individuals in such "macrosociological" experiments (Hedström, 2006) acknowledges that collective properties, such as status hierarchies, the diffusion of products and ideas, the strength of social norms, or network segregation, are not defined by the micro-level characteristics of the population members alone. As argued previously, many macro-level outcomes do not result from a linear aggregation of individual preferences and individualistic choices but often depend on critical masses and tipping points of contingent behaviors. Understanding how such social dynamics unfold is important because they have the capacity to construct highly path-dependent, and arbitrary, realities. In Salganik et al.'s prominent study, where participants could listen to and download songs of unknown artists in parallel "cultural markets", it became highly arbitrary which artists were most often downloaded once participants were exposed to a popularity ranking summarizing other participants' choices, and different artists led the popularity rankings in the parallel social systems (Salganik et al., 2006; see also Macy et al., 2019 on the polarization of arbitrary policy stances across political affiliations in the United States).

There is a potentially strong link here between ECABMs and large-scale web experiments. ABMs are increasingly used by experiment designers to target mechanisms for experimental manipulation, generate hypotheses, and extrapolate experimental results (e.g., Frey & van de Rijt, 2020; Stein, Keuschnigg, & van de Rijt, 2021). Typically, experiments are designed to collect detailed information about participants' micro-level behaviors. These data can then be used to (1) fit micro-level behavioral models that yield estimates of key behavioral parameters and (2) generate macro-level predictions using ECABMs that capture the empirically observed micro behaviors. The predictions can be produced as an internal consistency check to confirm that the experimentally observed micro behaviors indeed produce the macro phenomenon as hypothesized. The behavioral parameters estimated from experimental data might also be used to propose additional "virtual" interventions in the experimental platform or to scale up findings to understand the system behaviors among larger populations (e.g., Analytis, Stojic, & Moussaïd, 2015).

Digital trace data

Many sociological research questions defy randomized experimentation, and there are often other reasons, not least the transportability of empirical findings to the real world, to instead rely on observational data collected in real social environments. In the last decade, digitization has provided a plethora of sociologically relevant observational data from sources such as social-media platforms, online retail sites, mobile apps, administrative records, and historical archives. A particularly active research field investigates the behavioral mechanisms (e.g., confirmation bias) and situational mechanisms (e.g., network segregation) underlying political polarization on platforms such as Twitter and Facebook (e.g., Bakshy, Messing, & Adamic, 2015; Boutyline & Willer, 2017). These studies make particular use of the relational data provided by users following and maintaining ties to other individuals (e.g., friends, politicians) and organizations (e.g., media outlets, political parties). Analyses of co-following graphs on Twitter (Shi, Mast, Weber, Kellum, & Macy, 2017) further reveal that political polarization extends to other domains of social life, most notably cultural consumption, with conservative and liberal Twitter users (as indicated by their following of Republican or Democratic members of Congress) following different musical artists, restaurants, sports teams, and universities.

Typically, observational data from digital traces are not only wide – capturing, in the extreme, entire populations – but also deep in terms of granularity and the sheer number of variables available. The high dimensionality of many online datasets allows researchers to construct new measures of latent characteristics and to describe local social environments in detail, including network structures, homophily levels, and information flows (Golder & Macy, 2014; Bruch & Feinberg, 2017; Salganik, 2018; Edelmann, Wolff, Montagne, & Bail, 2020). These types of data differ in both size and kind from those typically used in the social sciences, and this has prompted interest in new methods emanating from the vibrant field of machine learning. Social scientists are using machine learning to distill new measures of hard-to-quantify constructs and to refine methods of causal inference from observational data, in contrast to the mainly predictive uses of machine learning in computer science and the emerging field of data science (Grimmer, 2015; Molina & Garip, 2019). Legewie and Schaeffer (2016), for example, study millions of geo-coded service and complaint calls made to New York City's 311 service to better understand neighborhood conflict in ethnically diverse settings. The authors use computer vision algorithms applied to census data to locate boundaries between racially and ethnically dissimilar areas of the city. They find that poorly defined rather than crisp and polarized boundaries between ethnic and racial groups act as drivers of neighborhood conflict.

In a recent study using Spotify data, Arvidsson, Hedström, and Keuschnigg (2020) estimate the causal effect that exposure to new music through a friend has on individuals' listening behavior. The user data contain information about who follows whom as well as musical tastes, characterized by fine-grained digital traces of users' listening habits. The difficulty of isolating social influence from confounding factors – homophily among interconnected individuals (e.g., friends like the same music) and common exposure to external stimuli (e.g., the like-minded receive similar algorithmic recommendations) – makes such analyses of social influence challenging (Aral, Muchnik, & Sundararajan, 2009; Shalizi & Thomas, 2011). The Spotify study uses the granular data on users' tastes to substantially improve the statistical matching of "treated" and "untreated" users and arrives at estimates of peer influence with substantially reduced confounding biases. Rather than matching on sociodemographic variables that correlate with adoption behavior, their procedure allows matching directly on music taste, a main driver of adoption behavior, tie formation, and exposure to music outside of Spotify. Importantly, the matching is performed after pre-processing individuals' playlist data. Playlists follow strong thematic patterns such that song co-occurrences in playlists indicate musical similarities. Inferring relational structures from co-occurrences is a central task in natural language processing. The study uses probabilistic topic models (Blei, Ng, & Jordan, 2003; see discussion in the following section) to map artists and their songs onto genres. The topic model allocates artists with different probabilities to different topics, capturing graded memberships in musical genres (Hannan et al., 2019). Replacing millions of songs contained in the individuals' playlists by their inferred genre preferences, the study arrives at a lower-dimensional representation of users' music tastes. The pre-processing makes sparse and granular data tractable for traditional matching models, such as propensity score matching and coarsened exact matching (Stuart, 2010), which have been developed for relatively low-dimensional settings where observations markedly outnumber variables.

Computational text analysis

For much of their disciplinary history, quantitative social scientists have found it difficult to study text. Lexical approaches that rely on word counts came to dominate, and with them attempts to infer meaning from word frequencies (Lasswell, Lerner, & de Sola Pool, 1952; Riffe et al., 2014). This tradition received renewed attention with the increasing availability of digitized text and the vast number of documents that can be searched in automated ways (Michel et al., 2011; Lorenz-Spreen, Mønsted, Hövel, & Lehmann, 2019). Related lexicon-based methods have also been used in analyses of sentiment expressed in text which, with the surge in human-annotated as well as machine-generated sentiment lexica, have been substantially refined in recent years (e.g., Pang & Lee, 2008; Pennebaker, Boyd, Jordan, & Blackburn, 2015). A particularly interesting design combines sentiment measures (e.g., from Twitter posts interpreted as "social sensors" over time) with a natural experiment (an exogenous variation of environmental conditions) such as policy changes or terrorist attacks in order to draw conclusions on their causal impact on public sentiment (Flores, 2017; Garcia & Rimé, 2019).

Advances in machine learning revolutionized the use of text as data (Grimmer & Stewart, 2013; DiMaggio, 2015; Mohr, Wagner-Pacifici, & Breiger, 2015; Evans & Aceves, 2016). Combined with ever larger corpora of digitized text, these tools offer new ways to measure what people think, feel, and talk about – on the level of a whole society. Traditional approaches to revealing such patterns of social behavior through text analysis required qualitative deep reading and time-consuming hand-coding which restricted analyses to small- and medium-sized collections of text (Franzosi, 2010; Riffe, Lacy, Fico, Lacy, & Fico, 2014). The abundance of text data now available makes the limited scalability of traditional approaches more apparent (Mützel, 2015), and the use of keywords to restrict analyses to important parts of a corpus, for example, has been shown to lead to biased results (King, Lam, & Roberts, 2017). At the same time, procedures based on hand- or automated-coding that rest on predefined classifications have been criticized as ill equipped to identify underlying cultural meanings and the context of social text (Biernacki, 2012; Guo, Vargo, Pan, Ding, & Ishwar, 2016). Fortunately, a collection of machine-learning methods developed for analyzing vast quantities of text data – often subsumed under the *natural language processing* label – have emerged in the last two decades (Jurafsky & Martin, 2009; Hirschberg & Manning, 2015).

Natural language processing offers new ways to describe both the macro-level properties of social systems and the characteristics of individual actors. In terms of macro properties, computational text analysis offers a lens through which to view relationships between words and, perhaps most interestingly for sociologists, the shared understandings of terms prevalent in a given population. On the micro level, text-analytic tools provide new measures of individuals' beliefs, sentiments, and tastes which in the past had to be collected using costly and difficult-to-administer survey instruments.

There is a classic distinction between supervised and unsupervised methods of computational text analysis. Supervised machine-learning algorithms can extrapolate hand-coded annotations from small subsets of documents to vast digitized corpora, making them available to large-scale quantitative analyses. Training computers to classify content opens up avenues for a broader and more representative study of vast bodies of digitized text, and "[i]nstead of restricting ourselves to collecting the best small pieces of information that can stand in for the textual whole. . ., contemporary technologies give us the ability to instead consider a textual corpus in its full hermeneutic complexity and nuance" (Mohr, Wagner-Pacifici, & Breiger, 2015, p. 3).

Unsupervised machine-learning algorithms for the analysis of text require no prior coding but recognize words that frequently occur together. A prominent approach is probabilistic topic modeling (Blei et al., 2003) which distills themes or categories from texts. Each corpus consists of multiple such "topics", and each document (e.g., online post, newspaper article, political speech) exhibits these topics in different proportions. From the viewpoint of the social sciences, topic models can reveal the "hidden thematic structure in large collections of documents"

(DiMaggio, Nag, & Blei, 2013, p. 577). Although topic models capture documents as "bags-of-words", ignoring syntax and word location, their coding of a corpus into meaningful categories often yields plausible readings of the texts, demonstrating what DiMaggio et al. (2013) describe as high levels of "substantive interpretability". The text-analytic concept of studying word co-occurrences connects to sociological ideas about how people create meaning and make sense of the social world by relating terms to other words and concepts (Goffman, 1974; DiMaggio, 1997; Mohr et al., 2020). Word co-occurrences are thought to capture such sociocultural associations, providing indications of schemata of interpretation and cultural frames more generally (Bail, 2014; DiMaggio, 2015). Correspondingly, topic models are increasingly used in sociology to operationalize relationality and meaning structures (e.g., DiMaggio et al., 2013; Lindgren, 2017; Nelson, 2020).

Another class of unsupervised algorithms, word embedding models (Mikolov, Yih, & Zweig, 2013; Pennington, Socher, & Manning, 2014), accounts for semantic structures by letting a sliding window pass through documents, recording the frequency with which words occur in a narrow context of other words. Word embeddings capture relations between words as distances between vectors in a high-dimensional space, allowing the inference of social meanings attached to words based on their positioning relative to other words. Because the dimensions of the identified vectors are largely uninterpretable, however, it is often unclear why words are predicted to be related. Important methodological developments thus focus on the interpretability of word embeddings. Hurtado Bodell, Arvidsson, and Magnusson (2019) and Kozlowski et al. (2019), for example, propose novel methodologies to study the meaning of individual words in relation to predetermined dimensions of interest (e.g., sentiment, gender, social status). An alternative to word embeddings builds on the older approach of co-occurrence networks where words – as nodes – are linked to their nearest neighbors, and a target term will associate with different close words over time, capturing the fluidity of meaning (Leskovec, Backstrom, & Kleinberg, 2009; Rule, Cointet, & Bearman, 2015; Bail, 2016). Such methodological developments strengthen the applicability of natural language processing to social science research questions, and they can deepen our understanding of how and why the cultural associations of words change over time.

It is also worthwhile to consider text-analytic models as a non-parametric way to conduct research using complex categorical data that are not typically thought of as corpora. The world is filled with detailed, linguistically delineated categories such as ethnicities, occupations, industries, and music genres. In many cases, the universe of categories runs into the dozens, if not hundreds or thousands. Sociologists often want to describe aggregations of people (e.g., neighborhoods, firm employees) or activities (e.g., cultural consumption, opinion expression) according to their compositions along these different categorical axes. In conventional statistical analyses with finite samples, it is rarely feasible to use all of the detailed categorical information at one's disposal, and often substantial simplification is necessary. Models derived from natural language processing can be used to induce, in a non-parametric way, salient regularities in compositional data rendered in full detail. To refer to a previous example, Arvidsson et al. (2020) apply topic models to Spotify playlists (i.e., documents) and their constituent artists (i.e., words). This use of topic models yielded condensed but nuanced descriptions of users' listening habits, which could then be employed in a causal analysis of social influence. Similar approaches could be taken to describe, for example, the ethnic mix of neighborhoods, the industrial mix of cities, or the mix of educational credentials in firms. Text-analytic models like topic models are especially appealing because they explicitly assume that any "document" (e.g., playlist, neighborhood, firm) is a mixture of topics. This allows for ambiguity and mixtures in classification, unlike many clustering methods which algorithmically assign analytical units to single categories.

While text-analytic models, on their own, stand to contribute novel insights for analytical sociologists, even greater insight may be unlocked by combining them with other CSS methods. The possibility of combining text analytic methods with ABMs is one largely unexplored avenue of research. Most ABMs, for practical reasons, postulate agents that think, perceive, and act in terms of numeric quantities or stark categorical distinctions. But language – its production and comprehension – is the basis for many human interactions, whether face to face or online. Combining ABMs with text analysis is an opportunity to leverage the dual interpretive and generative nature of some text-analytic models, like topic models, to understand social dynamics that are mediated by language. Agents in ABMs could be programmed to receive and interpret messages based on a text-analytic model and generate new messages based on the same discursive model. They could then take other actions, such as forming or dropping network ties or engaging in mobility, based on interpretations or classifications of messages received (e.g., Mordatch & Abbeel, 2018; Karell & Freedman, 2020). These kinds of agent-based models could be particularly helpful for modeling social processes in digital trace data, understanding not just the evolution of network ties but also, potentially, the evolution of the discourse itself.

Discussion

The dual growth in the power of computational tools and availability of digital trace data has pushed quantitative sociology to the brink of a new "watershed" moment, similar in significance to the introduction of representative population surveys and the tools for their statistical analysis in the middle of the twentieth century (Coleman, 1986; McFarland, Lewis, & Goldberg, 2016). Now the question is whether sociologists will take advantage of these new tools and evolve their empirical practice to live up to sociology's theoretical ambitions. Doing so requires transcending the tendency to limit quantitative analysis to the prediction of individual-level outcomes using psychological and sociodemographic variables contained in survey data. The new tools of computational social science present sociologists generally, and analytical sociologists in particular, with the chance to identify the complex social mechanisms that cause macro-level social patterns to emerge from the behaviors of networked, micro-level actors.

In this chapter, we have examined CSS methods that offer promise for analytical sociology's aim to understand and explain collective dynamics. In particular, we have shown how empirically calibrated agent-based models have made it possible to perform *in silico* counterfactual experiments in cases where *in situ* experimentation is virtually impossible, thereby identifying micro mechanisms that generate macro-level phenomena. We have discussed how large-scale web-based experiments make it possible to treat social systems, rather than individuals, as units of analysis in experimental tests of social mechanisms. And we have discussed how digital trace data, in combination with computational techniques of dimensionality reduction, including the tools of natural language processing, are opening up new data frontiers for quantifying difficult to measure concepts and observing related micro-level behaviors.

Our presentation has perhaps given the impression that the intellectual avenues connecting CSS and analytical sociology run one way: first, computational scholars in CSS create methods for their own purposes, and then sociologists adapt those methods to fit their substantive and theoretical interests. This is not our intent. We believe that AS has important contributions to make to CSS. Primarily, we believe an AS influence would lead CSS scholars to shed a preoccupation with producing aggregate-level descriptions of digital trace, text, and other "big" data that lack explanatory depth. We also believe that an AS influence would encourage CSS scholars to divert energy away from producing black-box predictive models and in the direction of developing tools to identify social mechanisms and assess their influence on macro-level social phenomena.

CSS techniques have greatly increased the ability of social scientists to collect data about complex social systems, to detect empirical regularities across these social systems, and to make predictions about individual and macro outcomes (Watts, 2013; Salganik, 2018). However, CSS scholars have too often rested at providing evidence of an empirical regularity and using a highly abstract mathematical or simulation model to propose a simple mechanism capable of producing that regularity. This theoretical work often lacks explanatory depth because it ignores empirical evidence at the micro level or gives little consideration to other mechanisms capable of producing similar empirical patterns. Examples of this tendency include the literature on urban scaling (Bettencourt, Lobo, Helbing, Kühnert, & West, 2007; Bettencourt, 2013) and research on the prevalence and emergence of power law distributions and scale-free networks (Barabási & Albert, 1999; Mitzenmacher, 2004). To summarize the critiques of these literatures (e.g., Stumpf & Porter, 2012; Keuschnigg, Mutgan, & Hedström, 2019), exclusive reliance on macro-level data to attribute a given regularity to a particular micro-level behavior is difficult, all the more so when multiple mechanisms are theoretically implicated (Young, 2009).

The antidote to the neglect of micro mechanisms is not to abandon explanatory ambition and resign ourselves to constructing complex prediction algorithms. Certainly, data scientists and other CSS researchers have made strides in making precise predictions for human behaviors, and the new innovations are increasingly finding their way into the social sciences (Molina & Garip, 2019; Edelmann et al., 2020; although see Salganik et al., 2020). However, we should recognize that the digital machinery used to, say, accurately predict a Netflix user's movie ratings or predict epidemics as a function of internet search terms is not necessarily conducive to understanding how aggregate social patterns like the distribution of box office revenues or the spread of global pandemics come about. For one, the digital machines used for prediction are typically black boxes. In the worst case, this raises serious questions about their reliability and reproducibility (Hutson, 2018). But even in the best case, these black boxes make it difficult to connect a prediction's inputs to its outputs, and in so doing may even hamper our ability to identify social mechanisms (Boelaert & Ollion, 2018; Wolbring, 2020). Machine learning models and artificial intelligence algorithms may provide more accurate predictions of human action than conventional statistical models, but ambiguity in linking this action to particular motivations, cognitive biases, or social influences poses a challenge for connecting micro behaviors to macro outcomes.

One potential solution is to employ theory to avoid the explanatory pitfalls that come from relying on aggregated digital data or micro-level predictive models. CSS should give fuller consideration to the many micro mechanisms that can produce a macro pattern of interest and explicitly investigate how these mechanisms combine and interact. Failure to articulate and demonstrate the social mechanisms driving a particular social phenomenon creates doubt about the practical implications of CSS research and hamstrings the research community's ability to generalize its findings. We are left with shallower understanding and greater uncertainty about effective policy responses to remedy perceived social problems (Martin & Sell, 1979; Deaton, 2010; Hedström & Ylikoski, 2010). A theory-driven approach that strives to identify mechanisms would improve the causal analysis of social systems, yielding insights that can be ported to other research cases covering different social domains.

A concern with social mechanisms should translate into a good-faith effort to incorporate mechanistic thinking into CSS research designs and analytic techniques. This means more than paying lip service to social mechanisms or relying on post-hoc, common-sense explanations of empirical findings (Kalter & Kroneberg, 2014; Watts, 2014). Instead, it requires CSS scholars to think more carefully about applying existing methods, or developing new tools, to explicitly locate and elucidate generalizable social mechanisms. In some cases, this may require thinking more carefully about heterogeneity and the importance of variation and randomness to the processes under study (Macy & Tsvetkova, 2015). To address this, machine learning algorithms might be deployed not only to produce precise predictions but also to sort through dense empirical data to identify behavioral "ecologies", making behavioral variation and its role in system-level dynamics the object of study (Arthur, 1994; Molina & Garip, 2019).

In other cases, understanding how a mechanism works requires identifying sequences of interrelated actions that bring about large-scale phenomena of interest. The point here is to see how the mechanistic cogs and wheels fit together and set each other in motion. To take an example, Schelling's (1971) work on segregation is still appreciated today not only because it connected micro-level behaviors to equilibrium patterns of segregation but also because Schelling delved into his model's unfolding micro-level dynamics to understand how chains of mobility created cascades towards segregation (Hegselmann, 2017). In Schelling's case, the micro-level dynamics were accessible because he acted as the computational engine: Schelling implemented his model by manually moving physical pieces around on a board. But more complex models, implemented *in silico*, may offer resistance when researchers attempt to pick out the chains of events that precipitate the emergence of macro properties. CSS techniques may help here. Pattern recognition algorithms designed for ordered events, like sequence analysis (Cornwell, 2015; Ritschard & Studer, 2018), could be strategically applied to real and simulated data alike to identify sequences of interconnected actions that generate macro outcomes. Similar care in examining micro-level dynamics can also be applied to the analysis of text data. Adopting a longitudinal perspective that explicitly acknowledges the relational and temporally contingent nature of discourse can provide insight into how cross-temporal social influence works and can even be used to imagine counterfactual patterns of discourse (Gerow, Hu, Boyd-Graber, Blei, & Evans, 2018). In general, there remains ample room for technical innovation in this space. As our generative models become more complex, social scientists will likely need complementary computational tools to help with unpacking the dynamic processes connecting micro behaviors to macro outcomes in these models.

CSS scholars do not have to act alone in introducing mechanistic thinking into their research. They can invite analytical sociologists, and social scientists more generally, to join their research projects from inception. And social scientists should consider returning the favor. Encouraging greater collaboration between AS and CSS is perhaps the most likely way to improve social explanation and technological practice in both (Watts, 2013; Subrahmanian & Kumar, 2017). However, extensive interdisciplinary collaboration between analytical sociologists and computational social scientists remains elusive. Perhaps a philosophical disconnect stemming from very different objects of research in the originating disciplines – physical systems rather than social systems – is the main stumbling block. CSS researchers, who often hail from computer science, statistics, and physics, typically emphasize predictive power rather than mechanistic explanation, whereas accurate prediction is a lesser concern, if it is even a realistic possibility, in the social sciences (Lieberson & Lynn, 2002; Salganik et al., 2020). It is also possible that more practical obstacles related to publication strategies, career expectations, and target audiences are impeding collaboration. These barriers separating CSS from analytical sociology will only be overcome with concerted effort. CSS scholars and analytical sociologists alike must invest in interdisciplinary activities and outlets, like journals and conferences, to build the lasting professional relationships that can cement ties between these disciplines. New models of research, communication, and publication that satisfy the intellectual and career needs of both CSS and AS researchers will need to be forged through dialogue and, eventually, collaboration. By building a shared intellectual community, the potential of CSS and analytical sociology to produce profound insights into the social world can be more fully realized.

References

adams, jimi, & Schaefer, D. R. (2016). How initial prevalence moderates network-based smoking change: Estimating contextual effects with stochastic actor-based models. *Journal of Health and Social Behavior*. https://doi.org/10.1177/0022146515627848

Analytis, P. P., Stojic, H., & Moussaïd, M. (2015). *The collective dynamics of sequential search in markets for cultural products* (SFI Working Paper No. 2015–06–023). Sante Fe Institute.

Anderson, P. W. (1972). More is different. *Science*, *177*(4047), 393–396. https://doi.org/10.1126/science.177.4047.393

Aral, S., Muchnik, L., & Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, *106*(51), 21544–21549. https://doi.org/10.1073/pnas.0908800106

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Review*, *84*(2), 406–411.

Arvidsson, M., Collet, F., & Hedström, P. (2021). The Trojan-horse mechanism: How networks reduce gender segregation. *Science Advances*, 7(16), eabf6730. https://doi.org/10.1126/sciadv.abf6730

Arvidsson, M., Hedström, P., & Keuschnigg, M. (2020). *Social influence and the emergence of the unexpected: How novel music spreads on Spotify* (Working Paper). Norrköping: Institute for Analytical Sociology.

Bader, F., & Keuschnigg, M. (2020). Bounded solidarity in cross-national encounters: Individuals share more with others from poor countries but trust them less. *Sociological Science*, 7, 415–432. https://doi.org/10.15195/v7.a17

Bail, C. A. (2014). The cultural environment: Measuring culture with big data. *Theory and Society*, *43*(3), 465–482. https://doi.org/10.1007/s11186-014-9216-5

Bail, C. A. (2016). Combining natural language processing and network analysis to examine how advocacy organizations stimulate conversation on social media. *Proceedings of the National Academy of Sciences*, *113*(42), 11823–11828. https://doi.org/10.1073/pnas.1607151113

Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, *348*(6239), 1130–1132. https://doi.org/10.1126/science.aaa1160

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509–512. https://doi.org/10.1126/science.286.5439.509

Bettencourt, L. M. A. (2013). The origins of scaling in cities. *Science*, *340*(6139), 1438–1441. https://doi.org/10.1126/science.1235823

Bettencourt, L. M. A., Lobo, J., Helbing, D., Kühnert, C., & West, G. B. (2007). Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the National Academy of Sciences*, *104*(17), 7301–7306. https://doi.org/10.1073/pnas.0610172104

Biernacki, R. (2012). *Reinventing evidence in social inquiry: Decoding facts and variables*. Palgrave Macmillan. https://doi.org/10.1057/9781137007285

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, *3*(Jan), 993–1022.

Boelaert, J., & Ollion, É. (2018). The great regression. *Revue Française de Sociologie*, *59*(3), 475–506. https://doi.org/10.3917/rfs.593.0475

Boutyline, A., & Willer, R. (2017). The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, *38*(3), 551–569. https://doi.org/10.1111/pops.12337

Bruch, E. E. (2014). How population structure shapes neighborhood segregation. *American Journal of Sociology*, *119*(5), 1221–1278. https://doi.org/10.1086/675411

Bruch, E. E., & Atwell, J. (2015). Agent-based models in empirical social research. *Sociological Methods & Research*, *44*(2), 186–221. https://doi.org/10.1177/0049124113506405

Bruch, E. E., & Feinberg, F. (2017). Decision-making processes in social contexts. *Annual Review of Sociology*, *43*(1), 207–227. https://doi.org/10.1146/annurev-soc-060116-053622

Bruch, E. E., & Mare, R. D. (2006). Neighborhood choice and neighborhood change. *American Journal of Sociology*, *112*, 667–709. https://doi.org/10.1086/507856

Centola, D. (2010). The spread of behavior in an online social network experiment. *Science*, *329*(5996), 1194–1197. https://doi.org/10.1126/science.1185231

Centola, D. (2011). An experimental study of homophily in the adoption of health behavior. *Science*, *334*(6060), 1269–1272. https://doi.org/10.1126/science.1207055

Centola, D. (2018). *How behavior spreads*. Princeton University Press.

Coleman, J. S. (1986). Social theory, social research, and a theory of action. *American Journal of Sociology*, *91*(6), 1309–1335. https://doi.org/10.1086/228423

Coleman, J. S. (1994). *Foundations of social theory*. Belknap Press of Harvard University Press.

Collier, N., & North, M. (2013). Parallel agent-based simulation with repast for high performance computing. *Simulation*, *89*(10), 1215–1235. https://doi.org/10.1177/0037549712462620

Conte, R., Gilbert, N., Bonelli, G., Cioffi-Revilla, C., Deffuant, G., Kertesz, J., . . . Helbing, D. (2012). Manifesto of computational social science. *The European Physical Journal Special Topics*, *214*(1), 325–346. https://doi.org/10.1140/epjst/e2012-01697-8

Cornwell, B. (2015). *Social sequence analysis: Methods and applications*. Cambridge University Press. https://doi.org/10.1017/CBO9781316212530

Daza, S., & Kreuger, L. K. (2019). Agent-based models for assessing complex statistical models: An example evaluating selection and social influence estimates from SIENA. *Sociological Methods & Research*. https://doi.org/10.1177/0049124119826147

Deaton, A. (2010). Instruments, randomization, and learning about development. *Journal of Economic Literature*, *48*(2), 424–455. https://doi.org/10.1257/jel.48.2.424

Deissenberg, C., van der Hoog, S., & Dawid, H. (2008). EURACE: A massively parallel agent-based model of the European economy. *Applied Mathematics and Computation*, *204*(2), 541–552. https://doi.org/10.1016/j.amc.2008.05.116

DiMaggio, P. (1997). Culture and cognition. *Annual Review of Sociology*, *23*(1), 263–287. https://doi.org/10.1146/annurev.soc.23.1.263

DiMaggio, P. (2015). Adapting computational text analysis to social science (and vice versa). *Big Data & Society*, *2*(2), Article 2053951715602908. https://doi.org/10.1177/2053951715602908

DiMaggio, P., Nag, M., & Blei, D. (2013). Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of U.S. government arts funding. *Poetics*, *41*(6), 570–606. https://doi.org/10.1016/j.poetic.2013.08.004

Edelmann, A., Wolff, T., Montagne, D., & Bail, C. A. (2020). Computational social science and sociology. *Annual Review of Sociology*, *46*(1), 61–81. https://doi.org/10.1146/annurev-soc-121919-054621

Elster, J. (1989). *Nuts and bolts for the social sciences*. Cambridge University Press. https://doi.org/10.1017/CBO9780511812255

Epstein, J. M. (2009). Modelling to contain pandemics. *Nature*, *460*(7256), 687–687. https://doi.org/10.1038/460687a

Evans, J. A., & Aceves, P. (2016). Machine translation: Mining text for social theory. *Annual Review of Sociology*, *42*(1), 21–50. https://doi.org/10.1146/annurev-soc-081715-074206

Flores, R. D. (2017). Do anti-immigrant laws shape public sentiment? A study of Arizona's SB 1070 using Twitter data. *American Journal of Sociology*, *123*(2), 333–384. https://doi.org/10.1086/692983

Franzosi, R. (2010). *Quantitative narrative analysis*. Sage.

Frey, V., & van de Rijt, A. (2000). Social influence undermines the wisdom of the crowd in sequential decision making. *Management Science*. https://doi.org/10.1287/mnsc.2020.3713

Friedman, M. (1953). The methodology of positive economics. In M. Friedman (Ed.), *Essays in positive economics* (pp. 3–43). University of Chicago Press.

Fujimoto, K., Snijders, T. A. B., & Valente, T. W. (2018). Multivariate dynamics of one-mode and two-mode networks: Explaining similarity in sports participation among friends. *Network Science*, *6*(3), 370–395. https://doi.org/10.1017/nws.2018.11

Garcia, D., & Rimé, B. (2019). Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychological Science*, *30*(4), 617–628. https://doi.org/10.1177/0956797619831964

Gerow, A., Hu, Y., Boyd-Graber, J., Blei, D. M., & Evans, J. A. (2018). Measuring discursive influence across scholarship. *Proceedings of the National Academy of Sciences*, *115*(13), 3308–3313. https://doi.org/10.1073/pnas.1719792115

Glennan, S., & Illari, P. (2017). Introduction: Mechanisms and mechanical philosophy. In S. Glennan & P. Illari (Eds.), *The Routledge Handbook of Mechanisms and Mechanical Philosophy* (pp. 1–7). Taylor & Francis. https://doi.org/10.1007/978-3-030-10707-9_1

Goffman, E. (1974). *Frame analysis: An essay on the organization of experience*. Harvard University Press.

Golder, S. A., & Macy, M. W. (2014). Digital footprints: Opportunities and challenges for online social research. *Annual Review of Sociology*, *40*(1), 129–152. https://doi.org/10.1146/annurev-soc-071913-043145

Grimmer, J. (2015). We are all social scientists now: How big data, machine learning, and causal inference work together. *PS: Political Science & Politics*, *48*(1), 80–83. https://doi.org/10.1017/S1049096514001784

Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, *21*(3), 267–297. https://doi.org/10.1093/pan/mps028

Guo, L., Vargo, C. J., Pan, Z., Ding, W., & Ishwar, P. (2016). Big social data analytics in journalism and mass communication: Comparing dictionary-based text analysis and unsupervised topic modeling. *Journalism & Mass Communication Quarterly*, *93*(2), 332–359. https://doi.org/10.1177/1077699016639231

Hannan, M. T., Mens, G. L., Hsu, G., Kovács, B., Negro, G., Pólos, L., Pontikes, E., & Sharkey, A. J. (2019). *Concepts and categories: Foundations for sociological and cultural analysis*. Columbia University Press.

Hedström, P. (2005). *Dissecting the social: On the principles of analytical sociology*. Cambridge University Press. https://doi.org/10.1017/CBO9780511488801

Hedström, P. (2006). Experimental macro sociology: Predicting the next best seller. *Science*, *311*(5762), 786–787. https://doi.org/10.1126/science.1124707

Hedström, P., & Bearman, P. (Eds.). (2009). *The Oxford handbook of analytical sociology*. Oxford University Press.

Hedström, P., & Ylikoski, P. (2010). Causal mechanisms in the social sciences. *Annual Review of Sociology*, *36*(1), 49–67. https://doi.org/10.1146/annurev.soc.012809.102632

Hegselmann, R. (2017). Thomas C. Schelling and James M. Sakoda: The intellectual, technical, and social history of a model. *Journal of Artificial Societies and Social Simulation*, *20*(3), 15. https://doi.org/10.18564/jasss.3511

Hirschberg, J., & Manning, C. D. (2015). Advances in natural language processing. *Science*, *349*(6245), 261–266. https://doi.org/10.1126/science.aaa8685

Hurtado Bodell, M., Arvidsson, M., & Magnusson, M. (2019). Interpretable word embeddings via informative priors. In K. Inui, J. Jiang, X. Wan (Eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 6323–6329). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1661

Hutson, M. (2018). Artificial intelligence faces reproducibility crisis. *Science*, *359*(6377), 725–726. https://doi.org/10.1126/science.359.6377.725

Illari, P. M., & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science*, *2*(1), 119–135. https://doi.org/10.1007/s13194-011-0038-2

Jarvis, B. F. (2015). *Matching residential mobility to raced-based neighborhood preferences in Los Angeles: Implications for racial residential segregation*. University of California Press. http://escholarship.org/uc/item/9174793s

Jurafsky, D., & Martin, J. H. (2009). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*. Prentice Hall.

Kalter, F., & Kroneberg, C. (2014). Between mechanism talk and mechanism cult: New emphases in explanatory sociology and empirical research. *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, *66*(1), 91–115. https://doi.org/10.1007/s11577-014-0272-7

Karell, D., & Freedman, M. (2020). Sociocultural mechanisms of conflict: Combining topic and stochastic actor-oriented models in an analysis of Afghanistan, 1979–2001. *Poetics*, *78*, 101403. https://doi.org/10.1016/j.poetic.2019.101403

Keuschnigg, M., Lovsjö, N., & Hedström, P. (2018). Analytical sociology and computational social science. *Journal of Computational Social Science*, *1*(1), 3–14. https://doi.org/10.1007/s42001-017-0006-5

Keuschnigg, M., Mutgan, S., & Hedström, P. (2019). Urban scaling and the regional divide. *Science Advances*, *5*(1), eaav0042. https://doi.org/10.1126/sciadv.aav0042

King, G., Lam, P., & Roberts, M. E. (2017). Computer-assisted keyword and document set discovery from unstructured text. *American Journal of Political Science*, *61*(4), 971–988. https://doi.org/10.1111/ajps.12291

Kozlowski, A. C., Taddy, M., & Evans, J. A. (2019). The geometry of culture: Analyzing the meanings of class through word embeddings. *American Sociological Review*, *84*(5), 905–949. https://doi.org/10.1177/0003122419877135

Lakon, C. M., Hipp, J. R., Wang, C., Butts, C. T., & Jose, R. (2015). Simulating dynamic network models and adolescent smoking: The impact of varying peer influence and peer selection. *American Journal of Public Health*, *105*(12), 2438–2448. https://doi.org/10.2105/AJPH.2015.302789

Lasswell, H. D., Lerner, D., & de Sola Pool, I. (1952). *The comparative study of symbols: An introduction*. Stanford University Press.

Lazer, D. M. J., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., . . . Alstyne, M. V. (2009). Computational social science. *Science*, *323*(5915), 721–723. https://doi.org/10.1126/science.1167742

Lazer, D. M. J., Pentland, A., Watts, D. J., Aral, S., Athey, S., Contractor, N., . . . Wagner, C. (2020). Computational social science: Obstacles and opportunities. *Science*, *369*(6507), 1060–1062. https://doi.org/10.1126/science.aaz8170

Ledford, H. (2020). How Facebook, Twitter and other data troves are revolutionizing social science. *Nature*, *582*(7812), 328–330. https://doi.org/10.1038/d41586-020-01747-1

Legewie, J., & Schaeffer, M. (2016). Contested boundaries: Explaining where ethnocidal diversity provokes neighborhood conflict. *American Journal of Sociology*, *122*(1), 125–161. https://doi.org/10.1086/686942

León-Medina, F. J. (2017). Analytical sociology and agent-based modeling: Is generative sufficiency sufficient? *Sociological Theory*, *35*(3), 157–178. https://doi.org/10.1177/0735275117725766

Leskovec, J., Backstrom, L., & Kleinberg, J. (2009). Meme-tracking and the dynamics of the news cycle. In J. Elder, F. S. Fogelman, P. Flach, & M. Zaki (Eds.), *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 497–506). Association for Computing Machinery. https://doi.org/10.1145/1557019.1557077

Lieberson, S., & Lynn, F. B. (2002). Barking up the wrong branch: Scientific alternatives to the current model of sociological science. *Annual Review of Sociology*, *28*(1), 1–19. https://doi.org/10.1146/annurev.soc.28.110601.141122

Lindgren, S. (2017). *Digital media and society*. Sage.

Liu, K., & Bearman, P. S. (2015). Focal points, endogenous processes, and exogenous shocks in the autism epidemic. *Sociological Methods & Research*, *44*(2), 272–305. https://doi.org/10.1177/0049124112460369

Lorenz-Spreen, P., Mønsted, B. M., Hövel, P., & Lehmann, S. (2019). Accelerating dynamics of collective attention. *Nature Communications*, *10*(1), Article 1759. https://doi.org/10.1038/s41467-019-09311-w

Macy, M. W., Deri, S., Ruch, A., & Tong, N. (2019). Opinion cascades and the unpredictability of partisan polarization. *Science Advances, 5*(8), eaax0754. https://doi.org/10.1126/sciadv.aax0754

Macy, M. W., & Tsvetkova, M. (2015). The signal importance of noise. *Sociological Methods & Research*, *44*(2), 306–328. https://doi.org/10.1177/0049124113508093

Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, *28*(1), 143–166. https://doi.org/10.1146/annurev.soc.28.110601.141117

Manzo, G. (2010). Analytical sociology and its critics. *European Journal of Sociology*, *51*(1), 129–170. https://doi.org/10.1017/S0003975610000056

Manzo, G. (2011). Relative deprivation in silico: Agent-based models and causality in analytical sociology. In P. Demeulenaere (Ed.), *Analytical sociology and social mechanisms* (pp. 266–308). Cambridge University Press. https://doi.org/10.1017/CBO9780511921315.014

Manzo, G. (2014). Data, generative models, and mechanisms: More on the principles of analytical sociology. In G. Manzo (Ed.), *Analytical sociology* (pp. 1–52). John Wiley & Sons. https://doi.org/10.1002/9781118762707.ch01

Marshall, B. D. L., & Galea, S. (2015). Formalizing the role of agent-based modeling in causal inference and epidemiology. *American Journal of Epidemiology*, *181*(2), 92–99. https://doi.org/10.1093/aje/kwu274

Martin, M. W., & Sell, J. (1979). The role of the experiment in the social sciences. *Sociological Quarterly*, *20*(4), 581–590. https://doi.org/10.1111/j.1533-8525.1979.tb01237.x

McFarland, D. A., Lewis, K., & Goldberg, A. (2016). Sociology in the era of big data: The ascent of forensic social science. *American Sociologist*, *47*(1), 12–35. https://doi.org/10.1007/s12108-015-9291-8

Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Team, T. G. B., . . . Aiden, E. L. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, *331*(6014), 176–182. https://doi.org/10.1126/science.1199644

Mikolov, T., Yih, W., & Zweig, G. (2013). Linguistic regularities in continuous space word representations. In L. Vanderwende, H. Daumé III, & K. Kirchhoff (Eds.), *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 746–751). Association for Computational Linguistics. https://aclanthology.org/N13-1090/

Mitzenmacher, M. (2004). A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, *1*(2), 1385. https://doi.org/10.1080/15427951.2004.10129088

Mohr, J. W., Bail, C. A., Frye, M., Lena, J. C., Lizardo, O., McDonnell, T. E., . . . Wherry, F. F. (2020). *Measuring culture*. Columbia University Press.

Mohr, J. W., Wagner-Pacifici, R., & Breiger, R. L. (2015). Toward a computational hermeneutics. *Big Data & Society*. https://doi.org/10.1177/2053951715613809

Molina, M., & Garip, F. (2019). Machine learning for sociology. *Annual Review of Sociology*, *45*(1), 27–45. https://doi.org/10.1146/annurev-soc-073117-041106

Mordatch, I., & Abbeel, P. (2018). Emergence of grounded compositional language in multi-agent populations. *ArXiv:1703.04908*. http://arxiv.org/abs/1703.04908

Mützel, S. (2015). Facing big data: Making sociology relevant. *Big Data & Society*, *2*(2), Article 2053951715599179. https://doi.org/10.1177/2053951715599179

Nelson, L. K. (2020). Computational grounded theory: A methodological framework. *Sociological Methods & Research*, *49*(1), 3–42. https://doi.org/10.1177/0049124117729703

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, *2*(1–2), 1–135. https://doi.org/10.1561/1500000011

Parker, J., & Epstein, J. M. (2011). A distributed platform for global-scale agent-based models of disease transmission. *ACM Transactions on Modeling and Computer Simulation*, *22*(1), Article 2 https://doi.org/10.1145/2043635.2043637

Pennebaker, J. W., Boyd, R. L., Jordan, K., & Blackburn, K. (2015). *The development and psychometric properties of LIWC2015*. University of Texas. https://repositories.lib.utexas.edu/handle/2152/31333

Pennington, J., Socher, R., & Manning, C. (2014). GloVe: Global vectors for word representation. In A. Moschitti, B. Pang, & W. Daelemans (Eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1532–1543). Association for Computational Linguistics. https://doi.org/10.3115/v1/D14-1162

Ragan, D. T., Osgood, D. W., Ramirez, N. G., Moody, J., & Gest, S. D. (2019). A comparison of peer influence estimates from SIENA stochastic actor-based models and from conventional regression approaches. *Sociological Methods & Research*. https://doi.org/10.1177/0049124119852369

Riffe, D., Lacy, S., Fico, F., Lacy, S., & Fico, F. (2014). *Analyzing media messages: Using quantitative content analysis in research*. Routledge. https://doi.org/10.4324/9780203551691

Ritschard, G., & Studer, M. (Eds.). (2018). *Sequence analysis and related approaches: Innovative methods and applications*. Springer.

Rule, A., Cointet, J.-P., & Bearman, P. S. (2015). Lexical shifts, substantive changes, and continuity in State of the Union discourse, 1790–2014. *Proceedings of the National Academy of Sciences*, *112*(35), 10837–10844. https://doi.org/10.1073/pnas.1512221112

Sakoda, J. M. (1971). The checkerboard model of social interaction. *Journal of Mathematical Sociology*, *1*(1), 119–132. https://doi.org/10.1080/0022250X.1971.9989791

Salganik, M. J. (2018). *Bit by bit: Social research in the digital age*. Princeton University Press.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*(5762), 854–856. https://doi.org/10.1126/science.1121066

Salganik, M. J., Lundberg, I., Kindel, A. T., Ahearn, C. E., Al-Ghoneim, K., Almaatouq, A., . . . McLanahan, S. (2020). Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, *117*(15), 8398–8403. https://doi.org/10.1073/pnas.1915006117

Salganik, M. J., & Watts, D. J. (2009). Web-based experiments for the study of collective social dynamics in cultural markets. *Topics in Cognitive Science*, *1*(3), 439–468. https://doi.org/10.1111/j.1756-8765.2009.01030.x

Schaefer, D. R., adams, jimi, & Haas, S. A. (2013). Social networks and smoking: Exploring the effects of peer influence and smoker popularity through simulations. *Health Education & Behavior*, *40*(1_suppl), 24S–32S. https://doi.org/10.1177/1090198113493091

Schaub, M., Gereke, J., & Baldassarri, D. (2020). Does poverty undermine cooperation in multiethnic settings? Evidence from a cooperative investment experiment. *Journal of Experimental Political Science*, 7(1), 72–74. https://doi.org/10.1017/XPS.2019.19

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, *1*(2), 143–186. https://doi.org/10.1080/0022250X.1971.9989794

Schelling, T. C. (1978). *Micromotives and macrobehavior*. W. W. Norton & Company.

Shalizi, C. R., & Thomas, A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research*, *40*(2), 211–239. https://doi.org/10.1177/0049124111404820

Shi, Y., Mast, K., Weber, I., Kellum, A., & Macy, M. (2017). Cultural fault lines and political polarization. In P. Fox, D. McGuinness, L. Poirer, P. Boldi, & K. Kinder-Kurlanda (Eds.), *Proceedings of the 2017*

ACM on Web Science Conference (pp. 213–217). Association for Computing Machinery. https://doi.org/10.1145/3091478.3091520

Snijders, T. A. B. (2001). The statistical evaluation of social network dynamics. *Sociological Methodology*, *31*(1), 361–395. https://doi.org/10.1111/0081-1750.00099

Snijders, T. A. B., & Steglich, C. (2013). Representing micro-macro linkages by actor-based dynamic network models: *Sociological Methods & Research*, *44*(2), 222–271. https://doi.org/10.1177/0049124113494573

Snijders, T. A. B., van de Bunt, G. G., & Steglich, C. (2010). Introduction to stochastic actor-based models for network dynamics. *Social Networks*, *32*(1), 44–60. https://doi.org/10.1016/j.socnet.2009.02.004

Sørensen, A. B. (1998). Theoretical mechanisms and the empirical study of social processes. In P. Hedström & R. Swedberg (Eds.), *Social mechanisms: An analytical approach to social theory* (pp. 238–266). Cambridge University Press. https://doi.org/10.1017/CBO9780511663901.010

Stadtfeld, C. (2018). The micro–macro link in social networks. In *Emerging trends in the social and behavioral sciences* (pp. 1–15). John Wiley & Sons, Inc. https://doi.org/10.1002/9781118900772.etrds0463

Steglich, C., Snijders, T. A. B., & Pearson, M. (2010). Dynamic networks and behavior: Separating selection from influence. *Sociological Methodology*, *40*(1), 329–393. https://doi.org/10.1111/j.1467-9531.2010.01225.x

Stein, J., Keuschnigg, M., & van de Rijt, A. (2021). *Network segregation and the propagation of misinformation* (Working Paper) Florence: European University Institute.

Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science*, *25*(1), 1–21. https://doi.org/10.1214/09-STS313

Stumpf, M. P. H., & Porter, M. A. (2012). Critical truths about power laws. *Science*, *335*(6069), 665–666. https://doi.org/10.1126/science.1216142

Subrahmanian, V. S., & Kumar, S. (2017). Predicting human behavior: The next frontiers. *Science*, *355*(6324), 489–489. https://doi.org/10.1126/science.aam7032

Udehn, L. (2002). *Methodological individualism: Background, history and meaning*. Routledge.

Waddell, P. (2002). UrbanSim: Modeling urban development for land use, transportation, and environmental planning. *Journal of the American Planning Association*, *68*(3), 297–314. https://doi.org/10.1080/01944360208976274

Watts, D. J. (2013). Computational social science: Exciting progress and future directions. *The Bridge: Linking Engineering and Society*, *43*(4). https://nae.edu/106118/Computational-Social-Science-Exciting-Progress-and-Future-Directions

Watts, D. J. (2014). Common sense and sociological explanations. *American Journal of Sociology*, *120*(2), 313–351. https://doi.org/10.1086/678271

Wolbring, T. (2020). The digital revolution in the social sciences: Five theses about big data and other recent methodological innovations from an analytical sociologist. *Soziale Welt, Special Issue*, *23*, 52–64.

Xie, Y., & Zhou, X. (2012). Modeling individual-level heterogeneity in racial residential segregation. *Proceedings of the National Academy of Sciences*, *109*(29), 11646–11651. https://doi.org/10.1073/pnas.1202218109

Young, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American Economic Review*, *99*(5), 1899–1924. https://doi.org/10.1257/aer.99.5.1899

Zhang, H., & Vorobeychik, Y. (2019). Empirically grounded agent-based models of innovation diffusion: A critical review. *Artificial Intelligence Review*, *52*(1), 707–741. https://doi.org/10.1007/s10462-017-9577-z