Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Levels of Analysis in Computational Social Science

Permalink

https://escholarship.org/uc/item/103299v4

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Krafft, Peter M
Griffiths, Thomas L

Publication Date

2018

Levels of Analysis in Computational Social Science

Peter M. Krafft (pkrafft@berkeley.edu)
Thomas L. Griffiths (tom_griffiths@berkeley.edu)
Department of Psychology, University of California, Berkeley
Social Science Matrix, University of California, Berkeley

Abstract

Marr's levels of analysis constitute one influential approach to the central program of cognitive science—the multilevel analysis of cognition as information processing. The distinctive aspects of Marr's framework are an emphasis on identifying the computational problems and constraints faced in cognition, and conceptual machinery to relate cognitive mechanisms to that computational level of analysis. Although related ideas have been explored in a range of social science disciplines, Marr's framework, and particularly its notion of the precise formulation of computational problems and solutions, has yet to be applied widely in social analysis. In the present work we develop a formulation of Marr's levels for social systems, provide examples of this approach, and address potential criticisms. The consequence is a computational perspective on the sociological school of structural functionalism, and an apparatus for conducting multiscale analysis of social systems.

Keywords: Computational social science; Marr's levels of analysis; structural functionalism; analytical sociology; computational social theory

Introduction

Marr (1982) famously argued that any information processing system can be analyzed at three levels, that of (1) the computational problem the system is solving; (2) the algorithm the system uses to solve that problem; and (3) how that algorithm is implemented in the "physical hardware" of the system. This decomposition offers both functional and mechanistic perspectives on information processing systems. Marr's main aim was understanding psychology and the human brain, and his levels of analysis have proven to be a useful conceptual tool for generations of cognitive scientists after him.

Considering that aspects of human cognition can be productively viewed as information processing, and that social groups consist in part of sets of people who exchange information, it is natural think about how Marr's levels of analysis could be productively applied in the analysis of social systems. An immediate difficulty occurs with a naïve application, however. Taking the disciplinary commitments of cognitive science for granted, we can clearly model social systems as distributed computer programs. Yet, there is no guarantee at all that the resulting computer programs implement any coherent distributed solution to particular computational problems. Far from all human collective behavior, or for that matter far from all combinations of synthetic intelligent agent behavior, has any functional purpose. Well-recognized examples of collective dysfunction resulting even from intelligent agents include information cascades (Bikhchandani, Hirshleifer, & Welch, 1992), phantom traffic jams (Kerner & Konhäuser, 1993), and the tragedy of the commons (Ostrom, 2015). The naïve application of Marr's levels of analysis to social systems therefore only extends as far as an algorithmic level of analysis—representing social systems as distributed computer programs. The naïve application does not necessarily extend through to the computational level, in which the combination of agent behavior, taken holistically, would have to yield coherent distributed computation at the population level. To address this difficulty, we pursue a program of identifying which social behaviors and structures can be productively conceptualized as having computational roles.

There are several existing approaches in the social sciences related to Marr's framework. The field of organization science has explicitly adopted information processing perspectives since the seminal work of Herbert Simon. Simon, who was an early thinker on the topic of information processing in the context of human behavior (Simon, 1978), also applied these ideas to organizations. In one famous passage, Simon writes: "In the post-industrial society, the central problem is not how to organize to produce efficiently ... but how to organize to make decisions, that is, to process information" (Simon, 1973, p. 269–270). Financial markets are also commonly understood as information processing systems. In a classic economics paper, Hayek states: "the economic problem of society ... is a problem of the utilization of knowledge not given to anyone in its totality" (Hayek, 1945, p. 519–520).

Marr's algorithmic and implementation levels are akin to mechanistic explanations. Mechanistic explanations are popular in the classical area of mathematical sociology, such as in Schelling's seggregation model (Schelling, 1971) or Granovetter's threshold model (Granovetter, 1978), and remain popular in the modern area of analytical sociology (Hedström & Bearman, 2009). Recent progress in the field of economics and computation highlights the algorithmic side of the mathematical notion of game-theoretic equilibrium (Daskalakis, Goldberg, & Papadimitriou, 2009). Mechanisms have also been a target of inquiry in organization science, such as in the study of transactive memory (Wegner, 1987).

These existing lenses in the social sciences fit naturally within Marr's framework, and therefore point towards a synthesis of a cognitive, information-processing view of a wide variety of social systems. At the same time, many of these classic works did not draw explicit parallels to distributed computation, or did not leverage the hierarchies of abstraction familiar to computer scientists that Marr deploys. Cognitive scientists have begun to explicitly explore the application of Marr's levels to social systems. Hutchins (1995) pioneered the application of Marr's levels to social systems in his ethnography of distributed cognition in team behavior on a naval vessel. To Hutchins, the computation performed by a naval vessel was that of navigation—calculating where you are and determining how to get where you want to go. Hutchins provided a detailed account of how this function is accomplished by the crew members and their interactions with each other and with artifacts on the ship.

Hutchins' example makes it clear that the explicit application of Marr's levels can be productive in the context of teams and organizations. However, one of the reasons this example easily fits into Marr's framework is because teams and organizations have well-defined group boundaries and have functions that are explicit in the goals of these groups. These goals then dictate the information processing challenges the groups face. An important outstanding question is to what extent Marr's approach can be applied to more loosely organized social systems that are often the subject of sociology.

The functionalist lens, used in sociology for a variety of less strictly organized and less explicitly engineered social systems, provides reason to believe that there is space for such an attempt to be fruitful. Structural or sociological functionalism—i.e., the pursuit of understanding social structures and behavior in terms of how they solve social problems—is one of the classic theoretical perspectives in sociology. Many early sociologists held views that a variety of social phenomena played functional roles in society. For instance, Spencer (1898) advocated for an equilibrium view of society and drew extensive analogies between social and biological function. Durkheim (1893) presented a functionalist argument that division of labor acts as a mechanism of social solidarity promoting a cohesive social bond. Although functionalism was and continues to be controversial in sociology (Weber, 1922; Giddens, 1984), scholars still lean on it in modern studies. Yet, unlike in organization science, explicit information processing analogies are barely ever used in sociology.

In the present work, we explicate the application of Marr's levels to loosely organized social systems, review examples of recent work that fit within this paradigm, address challenges to this approach, and explore its potential and limitations. The main benefit of Marr's approach is that computation provides an expressive language for high-level, abstract theory, while providing the conceptual machinery to relate that abstract level to mechanistic explanations. Computational social theory can therefore be precisely specified, and tested via its relation to algorithmic and behavioral descriptions. At the same time, Marr's charge to identify computational problems that information processing systems solve could provide inspiration for research questions in computational social science.

Illustrative Examples

Before more carefully defining Marr's levels for social analysis, we begin with three motivating illustrative examples.

Waiting in Line

A simple example that illustrates Marr's three levels of analysis in a social system is waiting in line. A line, for instance outside a professor's office, consists of a group of people, each standing behind another. This social behavior implements the computation of a first-in, first-out (FIFO) queue. A FIFO queue is a simple function used in computer science, for example to prioritize computer processes in the CPU. A FIFO queue takes as input a stream of entries, maintains the order of those entries, and outputs entries in that order. The representation that is used to solve the FIFO queue problem in the case of waiting in line is to maintain a linked list data structure between elements of the queue, and pop elements off the list as needed. A linked list is another data structure used computer science, in which each entry contains a "pointer" to the next element in the list. To "pop" a linked list means to remove the head element. In the example of waiting in line, the distributed algorithm that implements this linked list is for each person in the line to keep track of who is ahead of them. The physical implementation used to keep track of who is ahead of you in the line is simply to stand behind that person. Figure 1 illustrates this example.

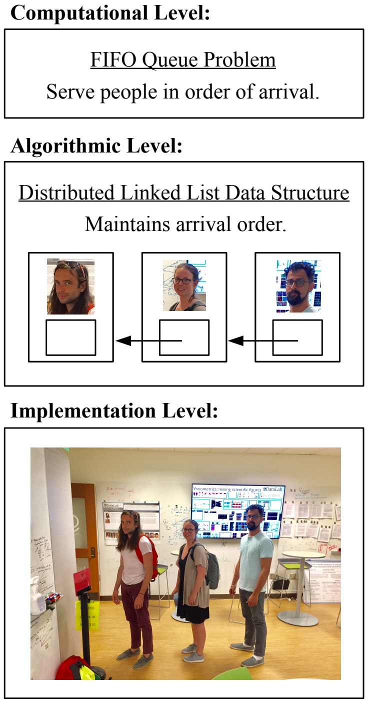

Figure 1: Marr's levels of analysis for waiting in line.

There are many failure modes to standing in line. Two people can arrive at a similar time or be in a similar position and not be certain who is ahead of whom. Some people cut in line. Sometimes lines fail to form at all or totally collapse and become disorganized crowds of people waiting. There is also cultural variation in how much importance or value people place on lines as a useful mechanism. Taking Marr's approach abstracts away these details and exposes the underlying information processing challenge at the heart of standing in line. Our ability to reason about the computational function of waiting in line can also suggest other engineered solutions. Waiting in line is not the only solution to the FIFO queue problem. For example, some delis, grocery stores, and government offices implement ticket and announcement systems that obviate the need to keep your own place in line.

Status Hierarchies

As another example, many species of animals maintain some kind of social hierarchy, influencing the interactions between animals in a group and the way that they allocate resources such as food. At a computational level, this structure can be viewed as a solution to the problem of performing resource allocation with a minimum of conflict—an alternative to a costly free-for-all whenever resources become available. At the algorithmic level, there are many kinds of algorithms that can be used to impose an ordering on a group through pairwise comparison (i.e. fights or displays)—a whole branch of theoretical computer science is devoted to questions about sorting. However, here the implementation details matter: most sorting algorithms involve maintaining some kind of global record of the current ordering of the items being sorted. In an animal group, each animal needs to maintain such a record independently. Thus it makes sense to think about distributed algorithms in which individual animals operate as computational elements. Flack and Krakauer (2011) applied exactly this approach to modeling the decisions of monkeys to engage in fights, viewing these individual decisions as forming "adaptive social circuits".

Rumors and Collective Sensemaking

As a final motivating example, a classical and now well-supported sociological theory of rumors conceptualizes rumors as a natural part of a process of collective sensemaking (Shibutani, 1966; Bordia & DiFonzo, 2004; Huang, Starbird, Orand, Stanek, & Pedersen, 2015). According to this view, people try to make sense of the world together when they find themselves in uncertain environments. A computational perspective of rumors based on this view is as functioning to communicate hypotheses about the state of the environment (Krafft, Zhou, Edwards, Starbird, & Spiro, 2017). This perspective frames rumors as oriented towards a distributed inference problem of inferring the state of the world given the evidence at hand.

Marr's Levels for Social Systems

Having motivated computational and algorithmic perspectives of social systems through our examples, we now expand upon the usage of Marr's levels of analysis in social systems.

Computational Level

The first level of analysis Marr defined is the computational level. The computational level describes the problem that an information processing system is oriented towards solving. The information processing function that the social system accomplishes may be explicit due to design or implicit, as in Merton's manifest versus latent functions (Merton, 1949). For this level of analysis to apply, the group must face some computational problem. The computational problems in our examples were implementing a FIFO queue, resource allocation, and distributed inference. Other common computational problems in social systems include aligning group member preferences and solving coordination problems (Krafft, 2018). Unlike in cognition, in which computational problems are frequently posed by the external environment, many computational problems faced by groups are endogenous. The need to coordinate is one example. The need to coordinate is an inherent result of existing as differentiated people. Another type of endogeneity is in problems that are created by history dependence. For instance, Durkheim offers that one view of the division of labor could be that by increasing our ability to create goods to relieve our increasing fatigue, division of labor functions in part to meet the needs created by its very existence (Durkheim, 1893).

The computational level of analysis is important because it allows the researcher to answer "why" questions—to understand why people behave in a certain way. In order to justify a teleological interpretation of a social function, that function should either be explicitly intended or otherwise be evidently addressing a problem that threatens the group. For instance, we can say that conventions about which side of the street to drive on exist in order to solve a coordination problem. Some coherent distributed computations do not meet this criterion of solving a computational problem associated with an intention or a need of the group. In the tragedy of the commons, rational agents are computing an equilibrium, and therefore accomplishing a computational function, but this outcome is neither intended nor meeting a need. Therefore this collective behavior cannot be productively interpreted as functional, and Marr's computational level does not apply.

Another important qualification in the social case is that the computational problem is one faced by the group, community, or society. Every individual in a group has their own problems and goals, and some behavior will be oriented towards those individual needs and not any shared needs of the group. Selfish behavior of this sort is one reason why we cannot treat all compositions of rational behavior as functional group behavior. This issue is at play in the tragedy of the commons and other social dilemmas from game theory.

The definition of social functions is also only with respect to the boundaries of the group being analyzed, and does not represent a moral judgment. Accomplishing a computational function in one group can cause problems for other groups; consider the case of one group finding a new place to build a settlement and displacing another group. In line with Weber's interpretive approach (Weber, 1922), insofar as we are aiming to understand why people are engaging in certain social behaviors, we must interpret function with respect to the values of the people in the group being analyzed.

Algorithmic Level

Marr's second level of analysis is the algorithmic level. The algorithmic level describes the way in which a computational problem associated with an information processing system is solved. An algorithm involves both the representations of information used and the transformations of those representations. In social systems, the fundamental algorithms at play are most readily conceptualized as distributed algorithms, in which multiple people are participating as agents akin to networked computer processors. This perspective of social processes as distributed algorithms is closely related to agent-based modeling (Macy & Willer, 2002), the study of social mechanisms in analytical sociology (Hedström & Bearman, 2009), and the study of natural algorithms in theoretical computer science (Chazelle, 2009).

The critical criterion for an algorithmic explanation within Marr's framework is that the behavior being examined offers a proper solution to the computational problem posed in the computational level of analysis. This criterion pushes beyond purely descriptive studies of social mechanisms, as in many agent-based or rational models, towards a formal relationship between mechanisms and social functions. In the example of waiting in line, the FIFO queue is accomplished if each person keeps their place. While deviant behavior such as line-cutting could be included in an agent-based model, a strict algorithmic analysis does not accommodate cases when some people cut in line for no reason other than their own self-interest, because this behavior undermines the correct computation of the FIFO queue.

Implementation Level

The final level of analysis Marr defines is the implementation level. An algorithm is an abstract process-level description. The implementation level explains how the algorithm is implemented through actual interaction of basic elements. In social systems, the definition of the implementation level is contingent on what elements of the system are taken as primitive. Typically, social systems are reduced to psychological processes, and connections to neural processes are left to cognitive scientists. This division leaves the implementation level to be concerned with psychological processes, details of social interaction, and contextual elements such as geography, social network structure, and artifacts in the environment as building blocks. The implementation level can be thought of as a second, lower-level algorithmic analysis.

Benefits of the Approach

There are several potential gains to be had from employing Marr's levels of analysis to understand social systems. One benefit is a deductive approach to discovering mechanisms. In analytical sociology, the discovery of mathematical descriptions for social mechanisms is often post hoc and inductive from observations. Marr's levels provide a deductive, reverse-engineering approach. In this approach, the computational problem being faced by a group is specified first, and then algorithms to solve that problem are explored. For instance, in the case of conceptualizing rumors as distributed inference, we can look to the literature on algorithms for distributed inference in search of mechanisms. To understand coordination, there is a wealth of literature in computer science on engineering distributed systems.

A second benefit is a rigorous approach to providing mathematical evidence for functionalist sociological theories. Taking the rumor example again, suppose we wanted to provide evidence that rumors function as a mechanism of collective sensemaking. Suppose we can show that a distributed inference algorithm as a behavioral model explains observed behavior better than alternative mechanisms, such as a contagion model or a thermodynamic model. The evidence for that distributed algorithm then in turn provides evidence for the functional interpretation of rumors since there is a mathematical relationship between the algorithmic model and the problem of distributed inference.

A final benefit is for design. Once a social problem is specified precisely as a computational problem, then we can do more than just understand how current social behavior might address this problem. We can also search for alternative social behaviors or structures that better solve the problem according to some criteria. The value of precise computational specification is that this search through design space can be automated. An example of this approach is in automated mechanism design (Conitzer & Sandholm, 2003).

Challenges to the Approach

There are several interrelated challenges and potential critiques of the indiscriminate application of Marr's approach to social systems. We now address what we view as the major challenges. Our responses to these challenges center around an argument that a program of Marr's approach to social systems aims to produce useful, idealized hierarchical mathematical descriptions, but should not be conducted without also paying careful attention to the specifics of the social context being studied and the political aspects of that inquiry.

Multiagent Systems versus Human Social Systems

One potential criticism that we can readily dismiss involves the difference between human social systems and artificial multiagent systems or distributed computer systems. Computer networks offer quite different affordances and constraints as compared to social systems. For instance, computers can easily communicate their entire internal states with complete precision to each other. Communication is much harder for people, but at the same time, people have a richer range of distinctive forms of communication, including symbolic and cultural systems. Social networks, the physicality of human interaction, social norms and institutions, and many other contextual factors form additional components that must be considered in the case of social systems. Although distributed computer systems and distributed social systems clearly have widely differing constraints and affordances, the mathematical language we use to describe both types of systems, the classes of algorithms that are employed, and some fraction of the computational problems each type of system faces could still be similar.

Methodological Individualism

A classic criticism of functionalism is that of methodological individualists or “reductionists”. The view of methodological individualism would assert that a group-level functionalist perspective is unnecessary for explaining the behavior of social systems. Under such a view, any group-level structure supervenes on the individual-level beliefs, intentions, plans, and behaviors. A social system therefore cannot be properly understood as having a function. To have a function means that the social system as a whole has a causal role in a broader ecosystem. But according to the reductionist view, the social system as a whole plays no causal role in the system dynamics. The only causally relevant entities are the components.

There are several responses to such a criticism. The response requiring the weakest logical commitments is that function can serve as a useful description that succinctly summarizes the behavior of the system, without making any causal claims. The usefulness of the description alone justifies a functionalist inquiry. When conceptualizing social systems “as if” they had functions allows us to better understand them, then such concepts are valuable. Another response is to assert that there can be multiple scales of causal explanations, and functions serve a causal role at an aggregate level. For instance, in a counterfactual world where people have no way to implement a FIFO queue at the professor’s office, then the operations of the group—who gets in when—would be fundamentally different. Therefore the operation of office hours can rightfully be conceptualized as having a dependence on the ability of people to implement a FIFO queue.

A related response is that certain functions are irreducible emergent properties of the social system, meaning that the behavior of the social system cannot be properly understood without understanding its emergent functions. This argument asserts that there is downward causality from emergent function to the constituents of the system. For instance, consider two competing hunting groups whose members must hunt together to be successful, and suppose the members of one group have synchronized clocks. It is the ability of the group to accomplish the task of coordination that allows the group to be more effective, not anything about the individuals in and of themselves. If the clocks did not accomplish the function of coordination, the individuals would not benefit.

Non-adaptive Functions

We argued in our definition of the computational level of analysis that the computational functions being analyzed must be beneficial to the group in order for the function to have a teleological interpretation, which is an implicit aim of Marr’s computational level. Some collective behavior implements coherent computation that is either harmful or epiphenomenal. For instance, Schelling’s (1971) segregation model showed how small individual biases could lead to a population clustering itself according to attributes such as race. Clustering or sorting could then be said to be an information processing function implemented by Schelling’s mechanism. At the same time, we might think that segregation is actually a maladaptive characteristic of a population. In a more recent example, a group of analytical sociologists presented a multiscale analysis of adolescent sexual behavior, and showed that the behavioral mechanisms of these people led to networks that tended to have structures similar to spanning trees (Bearman, Moody, & Stovel, 2004). Spanning trees are good for the sexual health of the community in some ways but bad in other ways, and thus do not serve a clear function. These analyses benefit from the same mathematical machinery that we use in the computational level analysis, but fall outside its scope in our definition.

Dysfunctional Collective Behavior

Another concern is individual behaviors that lead to incoherent or unstructured collective behavior, and individual or collective behaviors that appear functionally oriented but fail to accomplish any function. Social behavior that leads to incoherent or unstructured collective behavior, such as people going about their own individual business within their homes, simply may not have group-level structure that lends itself to illuminating interpretation via Marr’s levels. One prominent and perhaps surprising example of this sort is agent-based models. Although agent-based models can be described as computer programs, their aggregate dynamics are sometimes chaotic or unstructured. The scope of Marr’s levels therefore is not as wide as the class of all processes that can be described as distributed computer programs. Other behaviors may appear functionally oriented but are suboptimal or totally dysfunctional. Here, Marr’s approach simply may not apply if the social system does not have group-level information processing characteristics.

Sociological Critiques

A final set of threats to applying Marr’s levels to social systems are inherited from other challenges to sociological functionalism. Despite still being influential in contemporary sociology, functionalism has been criticized for abstracting and obscuring many key details of social phenomena. Conflict theorists have emphasized how functionalism diminishes the struggle of marginalized groups, the importance of revolutionary change, and the role of individual human agency in society. In a somewhat separate line of critique, Giddins’ structuration theory explores how function and structure coevolve continuously across space and time and cannot be neatly separated (Giddens, 1984).

We follow Weber (1922) in responding to these criticisms by noting firstly that a functional description can still be useful for certain ends, although non-functional aspects must be considered for a complete treatment of any system; and a functionalist analysis may still be illuminating to see how a population deviates from an idealized solution. That said, we must always keep in mind the balance between the clarity provided by abstraction and the frequent importance of the details that are abstracted away.

Conclusion

Analysis of social systems is challenging in part due to the diversity and complexity of the people in these systems and their interactions. Abstraction of social processes should be approached with caution but can help to highlight generalizable insights and underlying principles at play. In the present work we have outlined the application of Marr's levels of analysis to a broad range of social systems. This framework provides machinery for abstraction, and for relating abstract levels to lower levels of explanation.

Given our discussion of criticisms in the previous section, we can make an informed attempt at outlining the scope of Marr's approach for social systems. Marr's levels have clear utility and have been used in cases where groups have an explicit shared goal of executing an information processing task, such as in teams and organizations. We have argued that Marr's levels of analysis can also be useful in cases where people are self-organized—whether by coincidence, by intention, or through biological or cultural evolution—to execute coherent information processing. Cases that are less appropriate are when group behavior is unstructured, incoherent, or structured but not oriented towards information processing.

There are several classes of computational problems and algorithms that may be useful in deploying Marr's levels of analysis for social systems. Multiagent systems is one of the most relevant areas (Shoham & Leyton-Brown, 2008). The area of multiagent systems provides a wide variety of formalisms that are useful for both specifying problems and algorithms, including multiagent decision problems (Bernstein, Zilberstein, & Immerman, 2000) and computational models of shared cooperative activity (Grosz, Hunsberger, & Kraus, 1999). There are also many distributed algorithms outside the literature on multiagent systems, such as fault-tolerant distributed algorithms for consensus (Lynch, 1996). Distributed machine learning is another promising area to draw upon. Future work could deploy Marr's levels of analysis to further explore links between developments in these areas and the study of social systems.

Acknowledgments. Special thanks to Josh Tenenbaum, Sandy Pentland, David Krakauer, and Jessica Flack who helped stimulate an initial investigation into these ideas; to Josh for his suggestion to think about the example of waiting in line; and to Nick Logler, Steph Ballard, and Lavi Aulck for their special assistance. This work was supported by DARPA Cooperative Agreement D17AC00004.

References

Bearman, P. S., Moody, J., & Stovel, K. (2004). Chains of affection: The structure of adolescent romantic and sexual networks. *American Journal of Sociology*, *110*(1), 44–91.

Bernstein, D. S., Zilberstein, S., & Immerman, N. (2000). The complexity of decentralized control of Markov decision processes. *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 32–37.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*(5), 992–1026.

Bordia, P., & DiFonzo, N. (2004). Problem solving in social interactions on the internet: Rumor as social cognition. *Social Psychology Quarterly*, *67*(1), 33–49.

Chazelle, B. (2009). Natural algorithms. *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, 422–431.

Conitzer, V., & Sandholm, T. (2003). Applications of automated mechanism design. *UAI-03 Workshop on Bayesian Modeling Applications*.

Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, *39*(1), 195–259.

Durkheim, E. (1893). *Division of labor in society*. Paris: Alcan.

Flack, J. C., & Krakauer, D. C. (2011). Challenges for complexity measures: A perspective from social dynamics and collective social computation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *21*(3), 037108.

Giddens, A. (1984). *The constitution of society: Outline of the theory of structuration*. University of California Press.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, *83*(6), 1420–1443.

Grosz, B. J., Hunsberger, L., & Kraus, S. (1999). Planning and acting together. *AI Magazine*, *20*(4), 23.

Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, *35*(4), 519–530.

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

Huang, Y. L., Starbird, K., Orand, M., Stanek, S. A., & Pedersen, H. T. (2015). Connected through crisis: Emotional proximity and the spread of misinformation online. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 969–980.

Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

Kerner, B. S., & Konhäuser, P. (1993). Cluster effect in initially homogeneous traffic flow. *Physical Review E*, *48*(4), R2335.

Krafft, P. (2018). A simple computational theory of general collective intelligence. *Topics in Cognitive Science (topiCS) Special Issue on Computational Approaches to Social Cognition*.

Krafft, P., Zhou, K., Edwards, I., Starbird, K., & Spiro, E. S. (2017). Centralized, parallel, and distributed information processing during collective sensemaking. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2976–2987.

Lynch, N. A. (1996). *Distributed algorithms*. Morgan Kaufmann.

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

Marr, D. (1982). *Vision: A computational approach*. MIT Press.

Merton, R. K. (1949). *Social theory and social structure*. The Free Press.

Ostrom, E. (2015). *Governing the commons*. Cambridge University Press.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, *1*(2), 143–186.

Shibutani, T. (1966). *Improvised news: A sociological study of rumor*. The Bobbs-Merrill Company.

Shoham, Y., & Leyton-Brown, K. (2008). *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.

Simon, H. A. (1973). Applying information technology to organization design. *Public Administration Review*, *33*(3), 268–278.

Simon, H. A. (1978). Information-processing theory of human problem solving. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 5, pp. 271–295). Hillsdale, NJ: Lawrence Erlbaum Associates.

Spencer, H. (1898). *The principles of sociology*. New York: D. Appleton and Company.

Weber, M. (1922). *Economy and society: An outline of interpretive sociology* (Vol. 1).

Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In *Theories of group behavior* (pp. 185–208). Springer.