Social Network Analysis for Social Neuroscientists

Elisa C. Baek^{1,2}, Mason A. Porter² & Carolyn Parkinson^{1,3}

¹Department of Psychology, University of California, Los Angeles

²Department of Mathematics, University of California, Los Angeles

^{1,3}Brain Research Institute, University of California, Los Angeles

Author Note

This work was supported by NSF grant SBE-1835239 and by the NSF SBE Postdoctoral Research Fellowship grant 1911783.

Correspondence concerning this article should be addressed to Elisa C. Baek,
Department of Psychology, University of California, Los Angeles, Los Angeles, CA 90095.

Contact: elisabaek@ucla.edu

Abstract

Although social neuroscience is concerned with understanding how the brain interacts with its social environment, prevailing research in the field has primarily considered the human brain in isolation, deprived of its rich social context. Emerging work in social neuroscience that leverages tools from network analysis has begun to pursue this issue, advancing knowledge of how the human brain influences and is influenced by the structures of its social environment. In this paper, we provide an overview of key theory and methods in network analysis (especially for social systems) as an introduction for social neuroscientists who are interested in relating individual cognition to the structures of an individual's social environments. We also highlight some exciting new work as examples of how to productively use these tools to investigate questions of relevance to social neuroscientists. We include tutorials to help with practical implementation of the concepts that we discuss. We conclude by highlighting the broad range of exciting research opportunities for social neuroscientists who are interested in using network analysis to study social systems.

Social Network Analysis for Social Neuroscientists

Humans are social beings and are immersed in intricate social structures. Social interactions and relationships play important roles in healthy human development and function (House, Landis, & Umberson, 1988; Seeman, 1996; Uchino, 2006), and it is thought that the need to navigate complicated social interactions for survival advantage has contributed to the evolution of the human brain (Dunbar, 2008). Nevertheless, the vast majority of work in social neuroscience has studied individual cognition in isolation, deprived of its rich social context. As demonstrated recently (Morelli, Leong, Carlson, Kullar, & Zaki, 2018; O'Donnell, Bayer, Cascio, & Falk, 2017; Parkinson, Kleinbaum, & Wheatley, 2017, 2018; Zerubavel, Bearman, Weber, & Ochsner, 2015), social neuroscientists can leverage tools from network analysis to characterize the structure of individuals' social worlds to improve understanding of how individual brains shape and are shaped by their social networks (Weaverdyck & Parkinson, 2018). In this paper, we provide an overview of key theory and methods in network analysis (especially for social systems) and discuss practical examples of these ideas to highlight how network analysis can be useful for social neuroscientists who are interested in relating individual cognition to the structure of individuals' social environments.

Key Concepts of Network Analysis for Social Systems

In this section, we introduce some key concepts of network analysis that are particularly relevant for understanding social systems.

Nodes and Edges

Suppose that we want to characterize how people are connected to one another in a small town or an academic program. What might we want to know? We may first wish to identify the individuals who are in the network. We represent individuals in a network (i.e., "graph") by *nodes*, which are often called "vertices" in mathematics and "actors" in the context of social systems (see Figure 1. For introductions to networks, see Wasserman & Faust, 1994 for a sociological perspective, Kolaczyk, 2009 for a statistical perspective, and Newman, 2018 for a

physical-science perspective. In our hypothetical examples, a node may represent an inhabitant of a town or a student in an academic program. We may next wish to understand who is connected to whom in a network. Taking into account the connections between individuals is what differentiates studying a group (a collection of nodes) from a network (which also encodes the connections). We represent these connections between individuals in a network by edges (which are also often called “ties” or “links”). Edges can encode different relationships, depending on the questions of interest. For instance, edges can represent friendship (e.g., in academic cohorts; Parkinson et al., 2017, or in student organizations; Zerubavel et al., 2015) or professional relationships (e.g., in sports teams; Grund, 2012, or in private firms; Zaheer & Bell, 2005). One can define such relationships in terms of subjective reports (e.g., of who likes whom; Zerubavel et al., 2015, or who trusts whom in the network; Morelli et al., 2018) or in terms of the frequency of particular kinds of interactions or communications (e.g., physical encounters per day; Read, Eames, & Edmunds, 2008, or exchange of e-mails; Wuchty & Uzzi, 2011). Edges can also represent other phenomena, such as parts of a transportation route (e.g., movement through ports by cargo ships; Kaluza, Kölzsch, Gastner, & Blasius, 2010), or shared attributes (e.g., attendance at the same social events; Davis, Gardner, & Gardner, 1941) or common behavior patterns (e.g., voting similarities; Waugh, Pei, Fowler, Mucha, & Porter, 2009) between individuals.

In some cases, it is important to consider the directions of edges. For example, in a friendship network, we might draw an edge from node A to node B if actor A reports “liking” actor B; however, although it may be awkward, it is possible that actor B may not “like” actor A. One can represent such relationships with directed edges, with an arrow that points from one node to another (in the hypothetical example, from actor A to actor B). In other cases, edges are undirected, either because the criterion that is used to define them is undirected in nature (e.g., edges based only on shared attributes are inherently undirected) or because it can sometimes be pragmatic to consider edges as undirected. For example, a researcher may choose to

consider an undirected “friendship” tie between actors A and B if they both reported liking one another to impose a more stringent definition of friendship and/or if the researcher wishes to relate these data to other undirected data, such as interpersonal similarities. It is also sometimes desirable to consider weights of edges to represent the strengths of relationships. For example, one can encode interaction frequency with edges that are weighted by the number of interactions (during some time period) between two actors. In other cases, edges are unweighted, either because one obtains them in a way that is unweighted by nature (e.g., edges that encode whether a relationship exists or does not exist), or because there is a compelling reason to consider edges as unweighted. For example, to characterize only meaningful relationships, one may choose to use an edge that represents a relationship between two people if and only if it meets or exceeds a minimum threshold on the number of interactions.

In summary, edges in a network can be directed or undirected, and they can be either weighted or unweighted. These choices depend both on how data are acquired (e.g., by asking questions that produce binary or continuous responses) and on how they are encoded in a network (e.g., decisions to threshold and binarize continuous responses). There are advantages and disadvantages to using directed and weighted edges, rather than using edges that are symmetric and unweighted. Although directed and weighted edges can provide additional information about the nature of a relationship between two nodes, they can also complicate analysis. As we discuss in the following sections, considering edges as directed and/or weighted can complicate the characterization of various network measures and affects associated inferences. (Some methods also do not work in such more complicated cases; Newman, 2018.) Consequently, researchers should carefully consider these factors when deciding how to mathematically represent a social network. Other complications are that a network can include multiple types of edges (“multiplex networks”) and the edges in a network can change over time (“temporal networks”). We discuss these issues later in this manuscript (see the section on “Multilayer Networks”), and they are reviewed in detail elsewhere (Aleta &

Moreno, 2019; Kivela et al., 2014). It is important for researchers to consider whether these and other elaborations are appropriate to include in a given study.

Sociocentric Networks versus Egocentric Networks

One can study networks either by considering a sociocentric network (also called a “complete network”; Marsden, 2002; Newman, 2018) or by taking an egocentric (i.e., “ego-network”) approach (Crossley et al., 2015). A sociocentric network approach encapsulates a complete picture of who is connected¹ with whom in a network. One can construct a sociocentric social network by asking each person in a network about those with whom they are connected directly using a desired type of connection (depending on the question of interest). For instance, we might survey all of the members of a sports team to characterize a friendship network by asking who their friends are, or we might survey all of the residents of a dorm and ask who people turn to for emotional support. Recent work in social neuroscience that leverages tools from network science has often used a sociocentric network approach to characterize relatively small, bounded networks. Bounded networks (which are also called “closed networks”) are networks in which the boundaries of a network are clearly defined. In the strictest adherence to the definition of a bounded network, the boundaries of a social network are known perfectly, because individuals reside in a restrictive physical environment, such as a remote island (Brent, Ruiz-Lambides, & Platt, 2017), or because individuals are assigned to isolated social groups (Sallet et al., 2011). It is difficult to obtain perfectly bounded networks in humans, but recent work in social neuroscience has characterized relatively bounded networks, such as academic

¹ In the present paper, we use the term “connected” to indicate that two individuals have a relationship with one another. We use the term “connected directly” to indicate that two individuals are connected with a distance of 1 (i.e., they are “adjacent” to one another in a network). Our use of the term “connected directly” is synonymous with “direct ties” and the mathematical definition of “adjacent.” We also use the term “connected indirectly” to indicate that two individuals do not have a direct relationship with one another, but they each have relationships through one or more third-parties (e.g., through one or more mutual friends). We use the term “connected” throughout the paper, because we expect this terminology to be intuitive to our target audience in conveying our intended meaning. It is important not to confuse our use of the word “connected” with the use of the term in describing graphs or components of graphs (rather than individual nodes) in graph theory. The latter usage of “connected” refers to the idea that a path exists between every pair of nodes in a graph or in a component of a graph.

programs, dorms, and clubs (Morelli et al., 2018; Parkinson et al., 2017, 2018; Zerubavel et al., 2015); it has then collected neuroimaging data from some of the members of a relatively bounded network to relate neural processing to social network measures. Such an approach demonstrates one useful way to study individual cognition in the context of a broader social environment.

It is often insightful to study social networks using an ego-network approach. An ego-network is a network based on an individual (“ego”) and their friends (“alters”). One can construct ego-networks in a few different ways. If one possesses data on an entire bounded network, one can use it to extract “objective” ego networks that consist of just one individual and their friends. In such cases, where one obtains ego networks as part of a study that also characterizes sociocentric networks, researchers may also be interested in comparing an individual’s perceptions of a network to actual characteristics of the network. Such a comparison can lead to interesting questions about how people think about their relationships and relate to the social world around them through “cognitive social structures” (Krackhardt, 1987). In this case, one can construct “subjective” ego networks by asking individuals (“egos”) to complete a questionnaire about the people (“alters”) to whom they are connected¹ directly and whether these people are also connected directly to one another. For instance, one can survey a single member of a sports team to ask who their friends are and which of their friends are also friends with one another. Similarly, one can ask a single member of a first-year-student dorm which residents they turn to for emotional support and which residents turn to one another for emotional support. Although it is relatively uncommon to obtain data on individuals’ perceptions of relationships between third parties in situations when one already has characterized a sociocentric network to which those individuals (and their alters) belong, such an approach provides a useful way to explore questions about individuals’ perceptions of their networks and characteristics of a sociocentric network.

It is most common to obtain and characterize ego networks independently, without possessing information about an associated sociocentric network. In this situation, one typically characterizes ego networks through questionnaires that ask one individual (the “ego”) about the people (the “alters”) to whom they are connected directly and, in some cases, whether those people are connected directly to one another (similar to how one can obtain “subjective” ego networks to compare with sociocentric networks, as we discussed above). In cases where obtaining a sociocentric network is infeasible or inconvenient, employing an ego-network approach alone can be very useful, despite the inherent limitation that ego networks do not provide a complete picture of an entire sociocentric network. This limits the type of inferences that it is possible to draw from such subjective ego networks. For instance, if one finds that individual differences in network position² are associated with a behavioral or neural outcome, it would be unclear whether this relationship is due to actual differences in network position or differences in individuals’ perceptions of their network position (e.g., how many friends people think that they have versus how many friends they actually have). Despite their limitations, a key advantage of ego networks over sociocentric networks is that it is much easier to collect the former, and one can conveniently add them to a study simply through administering questionnaires to individuals in isolation. Indeed, several exciting new insights in social neuroscience have resulted from the use of ego-network approaches. For example, it has been demonstrated that network size measures that were identified from self-reporting and Facebook ego networks are associated with structural and functional differences in brain regions (Hampton, Unger, Von Der Heide, & Olson, 2016; Heide, Vyas, & Olson, 2014), and individual

² In the present paper, we use “network position” as a general term to refer to features that are related to an individual’s location in a social network (e.g., with whom they are close in social ties) and their node-level characteristics (e.g., centrality measures that quantify the influence of an individual in relation to other individuals). It is important not to confuse our use of this terminology with the more specific use of “network position” in relation to positional analysis (Wasserman & Faust, 1994).

differences in network position that were identified through Facebook ego networks have been associated with brain activity during a social-influence task (O'Donnell et al., 2017).

It is worth noting that, by default, an ego network is a 1-ego network, which consists of an ego's alters (i.e., people with whom the ego is connected directly) and the edges between the alters. In this way, a 1-ego network consists of the nodes and edges that are in an ego's personal social network (Crossley et al., 2015; Jeub, Balachandran, Porter, Mucha, & Mahoney, 2015). When mathematically analyzing 1-ego networks, one often does not include the direct connections between the ego and the alter, as one instead concentrates on the direct connections that exist between the alters. In the present paper, when we write "ego networks", we refer specifically to 1-ego networks, unless we specify otherwise. One can go further than an ego's 1-neighborhood by obtaining information about the alters' additional connections, beyond just those who have direct ties with an ego (e.g., by also obtaining the ego networks of each of the ego's alters). This yields a 2-ego network, which gives information about the nodes of distance 2 or less from an ego (e.g., "friends of friends" of the ego). We discuss social distance in detail later in the paper. One can iterate this process further to obtain k -ego networks (i.e., about all nodes within distance k of an ego), allowing one to obtain information about larger social structures in which an ego is immersed. A benefit of k -ego networks is that they provide more information about the broader social contexts of an individual than 1-ego networks, although it is often more cumbersome to obtain them in practice.

Mathematical Representation of Networks

One can represent a network mathematically using an adjacency matrix. An adjacency matrix $\mathbf{A}$ of a network is an $n \times n$ matrix (where n is the number of nodes) with elements A_{ij} . (More complicated network structures, such as multilayer networks, have more complicated adjacency structures; Kivela et al., 2014.) In an undirected and unweighted network, A_{ij} is 1 if there is an edge between nodes i and j , and it is 0 if there is no edge between nodes i and j . Because $A_{ij} = A_{ji}$, an adjacency matrix of an undirected network is symmetric (see Figure 1).

One can also represent a network using an edge list, which is a list of node pairs that are connected directly by edges (see Figure 1).

Figure 1.

Approaches to study and mathematically represent social networks. **(a–c)** In a sociocentric approach, one characterizes relationships between all members of a bounded social network. **(a)** A graphical representation of an undirected, unweighted sociocentric network that

represents friendships between members of a bounded community. The colored circles are nodes (also called vertices), which represent individuals in the social network. The lines between the nodes are edges, which represent friendships or some other relationship between individuals. **(b)** One can also represent networks with an edge list, which is a list of all direct connections between nodes. **(c)** It is also common to represent an n -node network with an adjacency matrix $\mathbf{A}$ of size $n \times n$ (with $n = 10$ in this example) whose elements A_{ij} encode the edges (both their existence and their weights) between each node pair (i, j) in a network. In an undirected, unweighted network, such as the depicted one, an associated adjacency matrix is symmetric. For example, the edge between Nick and Jen yields a 1 in the associated element of an adjacency matrix. **(d–f)** In an ego-network approach, one characterizes relationships in a network from an ego's point of view. Suppose that we obtain information about the same social network as the one in the left column from interviewing only Mike, a single member of the network. This gives us Mike's ego network. We draw solid lines to represent Mike's responses about his direct friendships and dotted lines from Mike's responses about whether his friends are also friends with one another. Comparing the graph from the sociocentric and ego-network approaches illustrates that the latter is missing information about several existing edges between nodes (e.g., between Nick and Elena, Nick and Jen, and so on). We also see this in the ego network's associated **(e)** edge list and **(f)** adjacency matrix.

Social Distance

Consider two strangers who are meeting for the first time. After speaking with one another for a while, they may be surprised to find that they have an acquaintance in common and then marvel at how small the world seems to be. Anecdotal evidence suggests that many people have had this sort of experience, reflecting the “small-world phenomenon” or, in other words, the idea that people in general are connected to each other by relatively short chains of relationships (Newman, 2018). Many people have an intuitive sense of the small-world phenomenon, but one may wonder how “small” the world really is (i.e., how close together, in terms of social ties, people actually are). In his pioneering studies of social distance, social psychologist Stanley Milgram and his colleagues sought to test this question (Milgram, 1967, 1969). In these experiments, Milgram and colleagues recruited participants who were living in cities in the Midwest of the United States and instructed them that their goal was to send a package (which included an official-looking letter and a stack of cards that were meant to track each person in the chain) to reach a target individual who lived in Massachusetts. If they did not

personally know the person on a first-name basis, they were instructed to forward the package to one of their direct connections who they thought was likely to be closer to the target. Milgram and his colleagues found that, on average, it took six steps for the packages (among those that completed their journey) to reach the target individual. This finding has been popularized in pop culture as “six degrees of separation,” expressing the idea that any two people in the world are separated by six or fewer social connections. More recently, scholars have continued to examine the small-world phenomenon through algorithmic frameworks (Kleinberg, 2000, 2011) and experiments like those of Milgram and colleagues, but using media such as e-mail (Dodds et al., 2003). The mean distance between individuals has also been studied in online social networks, such as Facebook (Ugander, Karrer, Backstrom, & Marlow, 2011). In this section, we provide an overview of concepts and methods for calculating social distance. We also discuss how they may be useful for questions of interest to social neuroscientists.

Given a network of nodes and edges, we can calculate a distance between two nodes (e.g., how far actor A is from actor B in a network). There are several ways of calculating distances in a network. The simplest is geodesic distance, which is the smallest number of edges that one needs to traverse to connect two nodes in a network. In other words, it is a shortest path. Two nodes can be connected either by direct ties (e.g., “friends” in a friendship network, with a distance of 1, as they are “adjacent” in the network) or by indirect ties (e.g., “friends of friends”, which yields a distance of 2, “friends of friends of friends”, which yields a distance of 3, and so on). The numerical values of social distance lead to different sociological inferences, depending on context. For instance, consider the friendship network of a new first-year cohort at a university. Suppose that actor A and actor B in this network are separated by a social distance of 4 (e.g., “friends of friends of friends of friends”). We may be interested in interpreting the absence of friendship between the two actors based on the social distance of 4. For instance, perhaps the two actors are distant from one another in friendship ties because they do not have much in common with each other. However, we would make different

inferences from this social distance of 4 depending on whether the two individuals live in dorms on opposite sides of the university campus or on the same floor of the same dorm. In the first scenario, the two individuals may be distant from one another in friendship ties due to a lack of opportunity to interact (and not necessarily due to a lack of compatibility or common interests). By contrast, the two individuals in the second scenario likely have had opportunities to interact, but they are not friends, so a lack of common interests or compatibility may be a more plausible explanation for the large social distance between them. Missing data can also complicate the interpretation of social distance, as missing ties can lead to an overestimation of distance between two individuals in a network. For example, in this scenario, if we are missing data from an individual in the network who is friends with both individuals (but we know that these two individuals are definitely not friends with each other), the actual distance between the two individuals is 2, rather than 4. As we have just demonstrated, when we are interested in drawing inferences from social distance, it is advantageous to choose networks that are bounded (so that we do not miss indirect connections between individuals, as this may lead to overestimation of some social distances) and in which we can be confident that opportunities to interact are relatively equally distributed across pairs of people in the network (to constrain interpretations of the potential causes of the relative distances between people). This can avoid potential complications in interpreting social distances. However, the reason that actors are distant from each other may not matter as much in other situations, such as if we are interested in the spread of information or behavior. When considering which network measures to use, researchers should ensure that they use methods and tools that are appropriate for their questions of interest.

Recent neuroimaging work suggests both that the human brain tracks the social distance between oneself and familiar others and that people spontaneously retrieve information about others' social-network positions when viewing their faces (Morelli et al., 2018; Parkinson et al., 2017; Zerubavel et al., 2015). This spontaneous retrieval of social-network knowledge

when encountering familiar others may help people respond appropriately when they interact with different people. Moreover, there is evidence that the brain not only tracks information about social-network position, but also that it influences and is influenced by the social networks in which it is embedded. For example, friendship is associated with similarity of neural responses to naturalistic stimuli. More specifically, recent work found that participants tended to have more similar time series of neural responses to audiovisual movies to people with whom they were connected directly (i.e., friends) than to people with whom they were only connected indirectly (e.g., friends of friends), with neural similarity decreasing with increasing social distance (Parkinson et al., 2018). This suggests that (1) people process information about the world in similar ways to those who are close to them in terms of social ties and that (2) individual brains may shape, and be shaped by, other brains that surround them. Such results demonstrate that one can leverage tools from network analysis to advance understanding of how individual brains represent and process the world around them.

Distance in Weighted Networks. Thus far, we have focused our discussion on geodesic distance, which is the simplest way of computing distance and is used often in unweighted networks. Computing distance in weighted networks is more complicated, and there are many ways to do it. A comprehensive discussion is beyond the scope of this paper, but see Cherkassky, Goldberg, and Radzik (1996) for a detailed presentation on shortest paths in weighted networks. A common way to calculate distance in weighted networks is to convert pairwise weights to costs and then to use Dijkstra's shortest-path-first algorithm (Dijkstra, 1959; Newman, 2001). Dijkstra's algorithm works by finding a path of "least resistance" between two nodes, where the "resistance" is the cost of traversing a path between two nodes. In a weighted network, the simplest choice for the cost of a tie between two nodes is the inverse of the tie's weight, where larger weights represent stronger ties and associated lower costs. For instance, if the weight of the edge between Felix and Sam is 2 and the weight of the edge between Sam and Dave is 5, the associated costs are $1/2$ and $1/5$, respectively (see Figure 2). For indirect

connections between two nodes (i.e., paths that require at least two edges), one calculates cost as the sum of the costs of the direct ties between nodes. For example, if Felix and Dave are connected through the edge Felix $\rightarrow$ Sam of weight 2 and the edge Sam $\rightarrow$ Dave of weight 5, then the cost of the 2-step path Felix $\rightarrow$ Sam $\rightarrow$ Dave is $1/2 + 1/5 = 0.7$. In many situations, Dijkstra's algorithm may identify two nodes that are connected only indirectly as "closer" than two nodes that are connected directly. In the example above, Felix and Dave are not connected directly and have a cost (i.e., distance) of 0.7. Consider another pair of nodes, Sarah and Sam, who are connected directly with an edge of weight 1. This yields a cost of 1 for their edge, because $1/1 = 1$. In Dijkstra's algorithm, Felix and Dave are considered to be closer to each other than Sarah and Sam are to each other, even though Felix and Dave are not connected directly. This is an important implication of what measures to consider, as many social neuroscientists may want to consider directly connected people (e.g., friends in a friendship network) as closer than indirectly connected people, a premise that fits well with the types of applications and research questions that are common in the field. For researchers who are interested in understanding the spread of phenomena (e.g., information or disease) in networks, Dijkstra's algorithm may give a helpful estimate of distance because, for example, information is more likely to spread faster through edges that represent very frequent interactions than through ones that represent infrequent interactions. As our discussion demonstrates, we encourage researchers to be driven by their research questions when making decisions about which network measures to use. We also encourage them to be attuned to the details of methods before applying and drawing inferences from them.

Figure 2.

Example of a Weighted Network and Dijkstra's Algorithm

Example of computing distances between nodes in a weighted network using Dijkstra's shortest-path-first algorithm. Dijkstra's algorithm uses distance based on the costs of the edges (i.e., distances between directly connected nodes), and it sums the costs of direct connections to determine the costs of indirect connections. Given a weighted network, one first needs to turn those weights into costs, and the simplest choice is to suppose that the cost of traversing an edge between two nodes is given by the inverse of the edge weight. In situations in which two nodes are connected only indirectly (e.g., Felix and Dave), one calculates the cost of traversing the edges between the two nodes by summing the costs of edges that connect those nodes. When there are multiple ways to travel between two indirectly connected nodes, Dijkstra's algorithm chooses a path of least resistance (i.e., lowest cost) to calculate distance. In this example, connecting Felix and Dave through Felix $\rightarrow$ Sam and Sam $\rightarrow$ Dave ($1/2 + 1/5 = 0.7$) is less costly than connecting Felix and Dave through Felix $\rightarrow$ Sarah $\rightarrow$ Sam $\rightarrow$ Dave ($1 + 1 + 1/5 = 2.2$). In many situations, the distance between indirectly connected nodes (e.g., Felix and Dave) is shorter than the distance between directly connected nodes (e.g., Sarah and Sam).

Centrality

It is often of interest to characterize the importance of actors (or of edges between them) in a social network (Newman, 2018; Wasserman & Faust, 1994). For instance, we may wish to know who is well-connected or popular in an academic program. The concept of "centrality" in network analysis is helpful for examining such questions. There are myriad variants of centrality;

we discuss some of the most common types in network analysis of social structures, with a focus on methods for calculating these centrality measures in unweighted, undirected networks. We also point to some resources for discussions of variations of these measures in weighted and directed networks. See Bringmann et al. (2019) for important caveats about studying and interpreting centralities in networks.

Degree Centrality. Degree centrality (also known simply as “degree”) is the number of edges that are attached to a node, so it is the number of direct connections of a person in a social network (see Figure 3). Another way, which is also useful for conceptualizing many other types of centrality, to think about degree is in terms of “walks” across edges in a network. Consider a robot who is walking around a social network. Given an undirected and unweighted network, we calculate the degree centrality of a node by taking the number of different ways that a robot can reach that node through a walk length of 1 (i.e., from a directly connected neighbor). Although degree centrality is a simple concept to grasp intuitively without the need to illustrate it with a walking robot, we include this description because it is helpful for comparing degree centrality with other centrality measures. There are various generalizations of degree centrality when incorporating direction and weights of edges, and we discuss some of these generalizations later in this manuscript. See the section “Consideration of Direction and Weights in Centrality Measures”.

Eigenvector Centrality. Although degree is a useful measure of centrality, it counts the number of connections of a node without taking into account the *quality* of those connections. Consider a townspeople who does not have many friends but is friends with the mayor, who has a large degree (and hence is well-connected in that respect). Although that townspeople has few friends, they may have more influence in the town than an individual with many friends with small degrees. Eigenvector centrality takes this type of influence into account, providing one way (see Figure 3) to capture how well-connected a node is to other well-connected nodes (in the context of social networks, how well-connected an individual is to well-connected

individuals; Bonacich, 1972). One calculates the eigenvector centralities of the nodes in a “connected” (in the graph-theoretic sense) network as the components of the leading eigenvector of the network’s adjacency matrix **A**. (As we described earlier in this manuscript¹, a network is “connected” in this sense if, for all pairs of actors, there is a walk between these actors. A directed network in which any node can be reached by a path that emanates from any other node is called “strongly connected”.) Another way to visualize the idea behind eigenvector centrality is through a random walk. Suppose that a robot goes on an infinitely long random walk through a network. The eigenvector centrality of a node is proportional to the frequency of visits to that node by the robot during its infinitely long walk in the network. The robot visits a node with a large value of eigenvector centrality more often than a node with a small value, because the former node’s direct neighbors are well-connected to other nodes in the network. Using this idea, one can derive the formula for eigenvector centrality using a random walk, and different variants of random walks lead to different types of eigenvector-based centralities. See Masuda, Porter, and Lambiotte (2017) for details.

Eigenvector centrality has been associated with various social and health-relevant phenomena in humans—including happiness (Fowler & Christakis, 2008), body weight (Christakis & Fowler, 2007), and job retention (Ballinger, Cross, & Holtom, 2016)—and with reproductive success in animals (Brent, 2015), suggesting that indirect ties (e.g., friends of friends, friends of friends of friends, and so on) may influence an individual’s well-being and behavior (and vice versa). People may be more likely to be aware of who is well-connected to well-connected others than they are of who has a lot of friends. For instance, in a large high school, people may be keenly aware of which individuals are popular in a popular group, but they may be less aware of which individuals in a less important group have many friends. This knowledge of who is well-connected to well-connected others has important implications. Mistreating an individual who is well-connected with well-connected ties may be risky, as the individual may be defended by their friends and their friends of friends, whereas mistreating an

individual who is not well-connected may have minimal consequences, given their limited influence (Ellwardt, Labianca, & Wittek, 2012; Salmivalli, Lagerspetz, Björkqvist, Österman, & Kaukiainen, 1996). In light of these scenarios, eigenvector centrality may be particularly useful when studying how people perceive social status in a network and how these perceptions shape people's behavior.

PageRank. We also briefly discuss PageRank, a variation of eigenvector centrality (that incorporates a chance for our walking robot to “teleport” to random nodes in a network in addition to traversing the network's edges) that has been used most famously to rank search results on the World Wide Web based on ordering the importances of web pages. A node (e.g., a web page) tends to be central according to PageRank if it has large in-degree (e.g., many other web pages point to it) and the incoming edges are from nodes that themselves have a large in-degree (e.g., the web pages that point to it have a lot of other web pages that point to them). PageRank takes into consideration both the direction and the weights of edges, and one construes a web page to be important if many other important web pages link to it. Suppose that a robot is randomly surfing the web, randomly walking from one web page (i.e., node) to another through directed edges (i.e., hyperlinks that point from one web page to another), and randomly “teleporting” to other web pages by opening a new browser window. One can calculate the PageRank centrality of a web page by examining how often the robot visits the web page, including through teleportation, if it surfs the web forever (Masuda et al., 2017). Although PageRank is associated most famously with ranking web pages, it has also been applied to investigate questions in a variety of different fields, including ranking the influence of Twitter users, ranking academic journals and doctoral programs, and finding correlated genes and proteins. For a review of PageRank, see Gleich (2015).

Diffusion centrality. Diffusion centrality, a type of centrality that generalizes both eigenvector centrality and Katz centrality (another notion of importance that is based on a walk on a network; Newman 2018), captures an individual's centrality with respect to a simple

spreading process on a network (Banerjee, Chandrasekhar, Duflo, & Jackson, 2013).

Calculating diffusion centrality may be useful for social neuroscientists who are interested in characterizing how central individuals are in their ability to spread items (such as information) in a dissemination process. Prior work has suggested that people are accurate at identifying others who are good at spreading information in a social network and that these estimates are positively and significantly correlated with diffusion centrality (Banerjee, Chandrasekhar, Duflo, & Jackson, 2014).

Betweenness Centrality. Another popular type of centrality is geodesic betweenness centrality, which measures the extent to which shortest paths (or, in generalizations of betweenness, other types of short paths) between pairs of nodes pass through a node. Suppose that a robot is traversing a network and takes a shortest path between each pair of nodes. One can calculate betweenness centrality of a node by tracking the number of times that the robot passes through the node to connect each pair of nodes (see Figure 3). It is common to interpret betweenness centrality as a measure of brokerage, because it captures some information about the extent to which a node connects distally connected nodes (Wasserman & Faust, 1994). For instance, an individual with a large betweenness centrality may have a high capacity for brokerage, because more of their friends have to go through them to communicate with one another. However, it is necessary to be cautious when interpreting betweenness as a measure of brokerage, as many different factors in network structure (including ones that are unrelated to a given individual) can strongly influence the value of betweenness, yet it is standard to interpret brokerage as a node-level characteristic (Everett & Valente, 2016). As with any network measure, context plays an important role in drawing inferences from betweenness centrality. In large networks, for instance, an individual may not be well-connected (as quantified, e.g., by a small degree centrality) and not well-connected to well-connected others (as quantified, e.g., by a small eigenvector centrality), but they may still have large betweenness. This individual may be in the periphery of multiple groups of friends; although

they may broker information between groups of otherwise unconnected nodes (e.g., two friendship groups), they may not be very influential in each of the individual groups. Another possibility is that individuals may have large betweenness centrality if they are connected directly (i.e., adjacent) to nodes that are brokers, even if they are not much of a broker themselves. If a researcher is interested in characterizing individual differences in socio-behavioral tendencies that are related to brokerage (e.g., how often people introduce their friends to one another), it may be useful (and sometimes more relevant) to calculate local network measures (such as local clustering coefficient; Watts & Strogatz, 1998 and Burt's constraint measure; Burt, 2004). It is also worth noting that many centrality measures, including betweenness centrality, are not robust to noise in data (e.g., missing edges), so one needs to pay careful to these issues when using these measures (Bringmann et al., 2019; Everett & Valente, 2016).

Consideration of Direction and Weights in Centrality Measures. In directed networks, each node has both an in-degree centrality (the number of edges that point to it) and an out-degree centrality (the number of edges that emanate from it). Depending on the question of interest, it may be appropriate to calculate versions of centrality measures (such as those that we discussed earlier) for networks with directions and/or weights. In some cases, generalizations are straightforward; for example, generalizing betweenness centrality to directed networks only requires restricting the node pairs (i.e., origin–destination pairs) that one considers, and one can directly generalize eigenvector centrality to weighted and directed networks by defining it based on random walks or as the leading eigenvector of an adjacency matrix. Moreover, PageRank is formulated specifically for directed networks and generalizes to weighted networks in the same way as eigenvector centrality. Other centralities entail more difficulties; for example, once one decides how to transform from edge weights to edge costs (i.e., edge distances), it becomes straightforward to generalize betweenness centrality to weighted networks (because one now knows how to calculate distances), but deciding what

function to use (e.g., inverting the weights or doing something else) to obtain distances in the first place involves an arbitrary decision that can severely impact the interpretation of the betweenness centrality values.

In a friendship network, one may be interested in the number of people with whom an individual says they are friends (i.e., out-degree); the number of people who say that they are friends with an individual (i.e., in-degree); counting any type of edge, regardless of the direction; or counting an edge if and only if it is mutually reported (i.e., “reciprocal”). Any of these choices can be useful, depending on the question of interest, and it is important to select appropriate measures that best fit a researcher’s question. For instance, if we seek to identify the most popular people in an academic program, it may be relevant to use in-degree centrality. One way to quantify popularity is by calculating (unweighted) in-degree centrality (e.g., by counting the number of people who say they like the individual using a binary survey question or by thresholding a continuous liking rating to create an unweighted edge) or through weighted in-degree (i.e., “in-strength”) centrality (e.g., by summing continuous liking ratings that an individual receives from different people; Zerubavel et al., 2015). If we are interested in understanding the spread of sexually transmitted diseases, however, we may not care about the direction of the ties and opt instead to calculate degree using undirected, unweighted edges (based, e.g., on the number of sexual partners of an individual, counting any edge between two actors even if it is not mutual; Christley et al., 2005). However, incorporating directions and/or weights can become complicated for various centrality measures (both mathematically and with respect to the substantive interpretation of centrality values), and a detailed review is beyond the scope of the present paper. See Wang, Hernandez, & Van Mieghem (2008) and White & Borgatti (1994) for helpful discussions.

Recent Examples. The centrality measures that we discussed above can be useful for investigating links between individual cognition and individuals’ broader social context. Recent research on centrality measures has further advanced the understanding of individual cognition

in rich social environments. For instance, some work has suggested that individual brains spontaneously encode and track network features of others, including eigenvector centrality (Parkinson et al., 2017), brokerage (Parkinson et al., 2017), and weighted in-degree centrality (Zerubavel et al., 2015). Furthermore, individual differences in betweenness centrality have been associated with individual differences in recruitment of brain regions during social influence (O'Donnell et al., 2017). Work on nonhuman primates has reported that having a larger degree centrality (and hence a larger number of connections with other individuals in a bounded social network) causally increases gray matter and resting-state functional connectivity in regions of the brain that are thought to be important in social functioning (Sallet et al., 2011). Although these examples highlight ways in which network analysis can advance understanding of individual cognition, it is necessary to be cautious when drawing broad inferences across such studies, given the heterogeneity of studies in design and specific choices when calculating network measures. Moreover, even the same (or very similar) network measure can represent very different phenomena, depending on the context of a study. For example, degree centrality encoded the potential number of social contacts in Sallet et al. (2011), but it encoded how much a person is liked in Zerubavel et al. (2015). We also note that the former paper used undirected, unweighted degree centrality, and the latter one used directed, weighted degree centralities. In many situations, it is likely that results that use different centrality measures—even ones that may seem very different from each other—may be picking up some shared information. We advise researchers to carefully keep these considerations in mind when aggregating findings across studies and in forming hypotheses for future studies.

Figure 3.

A few common measures of centrality. This adapted version of Krackhardt's kite graph (Krackhardt, 1990) demonstrates several variants of centrality. **(a)** An example friendship network, with each node labeled with the name of an individual. **(b–d)** Variations of the same network, with the nodes resized to reflect the size of a particular centrality measure. **(b)** Degree centrality is the number of other nodes to which a particular node is connected directly. Mike has a degree centrality of 7, the largest value in the network. **(c)** Eigenvector centrality captures how well-connected a node is to well-connected others. Although Elena, Dan, and Sam all have the same degree centrality (of 3), Sam has a much smaller eigenvector centrality value, as his friendships are with relatively poorly connected individuals. **(d)** Betweenness centrality captures the extent to which a node lies on shortest paths between pairs of nodes. Sam has the largest betweenness centrality, because he connects many nodes in the network that otherwise would be on disconnected components of the network.

Community Structure and Other Large-Scale Network Structures

Given a network, it is often insightful to study its large-scale structural patterns. Consider

your own social network of friends. How might you organize the individuals in your social network? One intuitive way may be to categorize your friends into groups, such as friends from high school, classmates from graduate school, teammates from a recreational sports league, fellow cosplayers, and so on. Similarly, many researchers are interested in understanding how nodes in a social network congregate into groups (Porter, Onnela, & Mucha, 2009). They are also often interested in other large-scale patterns, such as core versus peripheral groups (Csermely, London, Wu, & Uzzi, 2013; Rombach, Porter, Fowler, & Mucha, 2017), the roles and positions of individuals in a network (Rossi & Ahmed, 2015; Wasserman & Faust, 1994), and so on. In the present section, we focus on the idea of algorithmically detecting tightly-knit sets of nodes called “communities”, and we briefly discuss other large-scale network structures in the section on “Future Directions”.

The most prominently studied type of large-scale structure in a network is “community structure”, in which (in idealized form) dense sets of nodes are connected sparsely to other dense sets of nodes (Newman, 2018; Porter et al., 2009). Observing the clustered structure of a network of an academic program can provide insight into the features by which people organize into friendship groups (e.g., by mutual interests or academic subdisciplines). Furthermore, in a very large network, finding dense communities of nodes in an algorithmic way may allow one to break down the network into smaller, manageable subsets. However, how do we identify sets of nodes that form a community in a network? There are numerous methods to detect communities in networks, including both sociocentric (i.e., complete-network) approaches and egocentric approaches. A “community” is a set of nodes that are densely connected with one another but sparsely connected with other communities of nodes. For example, a community (e.g., a group of friends) in a network of an academic program consists of nodes (i.e., individuals) who are densely connected with one another (directly through edges) but more sparsely connected with other nodes (i.e., other individuals in the network). Although the notion of communities (and related notions, such as cohesive groups; Wasserman & Faust, 1994) in a

network is intuitively appealing, it is very challenging to precisely define what it means to be “densely connected” and “sparsely connected” (Fortunato & Hric, 2016). One common approach to detect communities is modularity maximization, in which one seeks a partition of a network that maximizes “modularity”, an objective function that quantifies the extent to which nodes in a community connect with one another in comparison to some baseline (Newman, 2006). Another popular approach is statistical inference of communities (and other large-scale network structures) using stochastic block models (Fortunato & Hric, 2016; Peixoto, 2017). There are numerous other algorithms to identify communities (with new ones published very frequently), but a review of these methods is beyond the scope of this paper. See Porter et al. (2009) for a friendly introduction to community structure and Fortunato & Hric (2016) for a recent review.

Multilayer Networks

Thus far, we have discussed single-layer (i.e., “monolayer”) networks, as we have concentrated on networks with a single type of node in which the nodes are connected to each other with a single type of tie. Mathematically, one represents a monolayer network as a graph (Newman, 2018). However, real networks are typically more complicated, as they encode multiple types of relationships (sometimes between multiple types of nodes) and interactions that change over time. Multilayer network analysis, an emerging area of network science, allows the study of richer network representations to further understanding of how different elements that comprise the social world interact with each another. A multilayer network consists of a set of layers that each have their own network of nodes and edges, along with interlayer edges that connect nodes from different layers. For a detailed review on multilayer networks, see Kivela et al. (2014). For a recent survey, see Aleta and Moreno (2019). For a review of multilayer networks in the context of animal behavior, see Finn et al. (2019)

As we indicated previously, individuals (i.e., nodes) in a social network can have many different types of relationships (i.e., edges). For instance, nodes that encode all of the individuals in a closed network (e.g., a town) can be connected to each other with edges that

represent different types of relationships, such as friendship, professional ties, recreational relationships, and so on. One can simultaneously encode all of these relationships in a multilayer network, with each type of relationship in a different layer. In our town example, each layer includes the same nodes (e.g., every member of a town), although this need not be true in general, but there are different types of edges in different layers (e.g., with layer 1 encoding friendships, layer 2 encoding professional ties, and layer 3 encoding recreational relationships) (see Figure 4). We also suppose that all interlayer edges in this example are between instantiations of the same individual in different layers. This type of multilayer network, in which different layers encode different types of relationships and interlayer connections exist only between the corresponding node across layers, is a type of multilayer network called a “multiplex” network.

Multilayer networks can include different types of nodes and/or different types of nodes in different layers. For example, consider the online social networks of an individual. An individual may use Facebook to connect with friends but LinkedIn for professional ties. If we encode connections in these social media in a multilayer network, with the individual’s Facebook network in layer 1 and LinkedIn network in layer 2, different nodes exist in each layer and some edges may cross the two layers (e.g., nodes that communicate across the two platforms). Multilayer networks can also encode more complicated types of interactions. For instance, one layer may consist of friendships, with nodes encoding people and edges encoding friendships; and a second layer may consist of a network of restaurants in town, with nodes encoding restaurants and edges encoding culinary collaborations (see Figure 4). Edges between the two layers can represent which restaurants are visited by which individuals, allowing one to examine phenomena such as relationships between friendship groups and restaurant-patronage patterns.

Temporal Networks. In a network, nodes and edges (and edge weights) often change over time. For instance, in the social network of a town, people move in and out (changing

nodes), so the relationships between people change (i.e., time-dependent edges) over time. It is often convenient to represent a temporal network using a multilayer network, with each layer encoding the network at a specific time or aggregated over a specific time period. Research on multilayer representations of temporal networks is related to analysis of temporal networks more generally (for reviews, see Holme & Saramäki, 2012 and Holme, 2015), and investigating a temporal network may be useful to researchers who are interested in relating individual cognition to dynamically changing social environments.

Figure 4.

(a) Multiplex Network

(b) Multilayer Network

Examples of multilayer networks. **(a)** A multiplex network is a type of multilayer network in which each layer has a different type of edge, and interlayer edges can exist only between corresponding nodes in different layers. In this example, each layer represents a different type of social network. The nodes in this example are the same individuals in each layer, and the edges in different layers encode different types of social relationships. We do not show any interlayer edges. In the first layer, edges represent friendships between individuals, whereas edges represent professional relationships between individuals in the second layer and recreational relationships between individuals in the third layer. **(b)** In this example of a multilayer network, the first layer represents the same friendship network that we showed in panel **(a)**. The second layer represents a restaurant network, where nodes encode restaurants and intralayer edges encode culinary collaborations between restaurants. Interlayer edges represent restaurant patronage of a restaurant by an individual, with an edge indicating that an individual has visited a restaurant. This type of multilayer network can help one understand

possible relationships between friendship and restaurant-patronage patterns. In this example, friends tend to eat at the same restaurants.

Learning from other fields. As we discussed in this section, there is great potential for using multilayer networks to advance the study of complex human behavior and social systems. It seems especially promising for social neuroscientists who are interested in studying individual cognition in the context of individuals' broader social contexts. Multilayer networks can provide integrated representations of the diversity of networks that surround an individual, enabling researchers to draw insight and test how different layers of a network influence not only each another but also processes that occur on them. Although the analysis of multilayer networks is a relatively novel methodology in network science, it has been used to enrich the study of diverse topics, including transportation systems (Gallotti & Barthelemy, 2015), coauthorship networks (Berlingerio, Coscia, Giannotti, Monreale, & Pedreschi, 2013), ecological networks (Pilosof, Porter, Pascual, & Kéfi, 2017), brain networks (Vaiana & Muldoon, 2018), and animal social networks (Barrett, Henzi, & Lusseau, 2012). Researchers who study human behavior can learn and draw inspiration from such prior work. For example, see Finn et al. (2019) for a detailed discussion of the use of multilayer network analysis to study animal behavior and Aleta & Moreno (2019) and Kivela et al. (2014) for broader reviews on multilayer networks. We recommend this literature to researchers who are interested in implementing multilayer network analysis to study human behavior.

Methods to Obtain Networks

In this section, we discuss some of the most common methods for obtaining networks.

Self-report surveys and questionnaires. A very common approach for obtaining social networks is through self-report surveys and questionnaires. Using a name generator, one asks participants to list people with whom they are connected in a social network (e.g., "With whom are you friends?"). In the same survey, one can generate multilayer networks by asking a

selection of questions (e.g., “With whom are you friends?” and “To whom do you turn for advice?”). Name generators can be either fixed choice (e.g., “Name the top-7 friends with whom you are closest.”) or free choice, without imposing limits on the number of people who a person can list. When it is possible for a researcher to obtain all of the names of individuals in a network prior to data collection, it is also possible to use roster-based methods. In a roster-based approach, one gives participants a list of all of the individuals in a network and asks them to characterize their relationship with each individual (e.g., indicating whether they are friends with each person or not, the strength of their friendship, and so on). Roster-based approaches have fewer recall issues than other approaches, and it is preferable to use them when possible. As with all self-reported data, these methods entail potential concerns about bias and inaccuracy because of desirability concerns of participants and question-order effects (Pustejovsky & Spillane, 2009). However, this potential disadvantage of self-report surveys can become an object of interest in itself. For instance, a researcher who is interested in understanding how people understand and represent their own social networks, even if they are not accurate, can use the framework of cognitive social structures (Krackhardt, 1987).

Direct observation. Another method to obtain networks is through direct observation. This is a frequently-used option that is available for researchers who study animal social networks, as they use it for observing and recording animal behavior (Noonan et al., 2014; Sallet et al., 2011), although many recent studies of animal social networks have employed technology such as radio frequency identification (RFID) data (Bonter & Bridge, 2011; Firth et al., 2017; Krause et al., 2013). In humans, direct observation can be labor-intensive, and it is typically feasible only for small groups. For instance, a researcher may observe the classroom behavior of children to construct a friendship network (Gest, Farmer, Carolina, & Cairns, 2003; Santos, Daniel, Fernandes, & Vaughn, 2015).

Archival and third-party records. It is also possible to reconstruct social networks using archival or third-party records. If we are interested in understanding intermarriage of royal

families in Europe during the 1500s, we can look at historical marriage records to reconstruct such a network. For instance, historical data were used to characterize and analyze the social network of political elite families in 13th Century Florence, and they were able to identify network characteristics that contributed to the rise of the powerful Medici family (Padget & Ansell, 1993). Today, we can leverage technological advances to obtain data such as e-mail, phone, and geographic-location records to reconstruct not only networks that encode the existence of communication ties, but also the frequency and patterns of communication. This method has been used in studying communication within organizations (Campbell, Maglio, Cozzi, & Dom, 2003), face-to-face contact in academic conferences and museums (Isella et al., 2011), features of social structures inferred from mobile-phone data (Eagle & Pentland, 2006), and other examples. We expect that these methods will also be useful for social neuroscientists.

Advantages of these methods include that they do not rely on self-reported data, are relatively low-effort (with the caveat that such data may be hard to access), and can provide a wealth of different types of data (and an abundance of data of each type). However, researchers should be mindful when interpreting the social significance of a tie in these networks. For instance, an e-mail exchange in an organization may encode only formal ties between coworkers and fail to capture less formal ties, which can also affect the phenomena that a researcher is hoping to capture. For example, an employee may exchange frequent e-mails with their supervisor and none at all with a coworker who sits in the cubicle next to them, although they may have a closer relationship with their coworker. Consequently, measuring the distance between these people in a network that one constructs using exclusively e-mail data is unlikely to provide a complete picture of these individuals' social relationships. Therefore, researchers should be mindful of these considerations while drawing inferences from calculations that use such networks. Researchers should also be mindful of the privacy concerns that may arise from accessing potentially sensitive personal information of participants, particularly when considering posting the data online (which ordinarily is desirable, as it helps promote open

science initiatives). It is possible to reconstruct even fully anonymized data, especially when there is a lot of data for each participant, to identify individuals (Herschel & Miori, 2017).

The rise of online social networking websites, such as Facebook and Twitter, has also afforded researchers the opportunity to “scrape” them (and otherwise acquire data from them) and study online social networks (Lewis, Kaufman, Gonzalez, Wimmer, & Christakis, 2008), although the policies of the companies that own the networks may entail some limitations. Additionally, when studying large online social networks, it is also necessary to pay close attention to the characteristics both of the network at large and of smaller local networks of interest, as these they may influence salient network measures (see, e.g., Jeub et al., 2015; Ugander, Backstrom, Marlow, & Kleinberg, 2012). Furthermore, social networks obtained from online websites are often 1-ego networks (encoding information about an individual ego and their friends), which have limitations, as we discussed previously in our introduction to egocentric approaches. (See the section “Sociocentric Networks versus Egocentric Networks”.) One also needs to be careful when interpreting the social significance of ties in online social networks. For instance, a large degree centrality on Facebook or a large in-degree centrality on Twitter may be an indication that an individual frequently uses the platform, rather than being related to the kinds of individual differences in socio-behavioral tendencies that may be of more interest to social neuroscientists. In other words, a person with a small degree centrality (i.e., few “friends”) on Facebook may actually have a large degree centrality in their offline life, which can be problematic if one uses degree centrality from Facebook data alone as a measure to relate to a neural or behavioral measure. Moreover, in data from social networking websites, there may be additional uncertainty in the ability of researchers to detect an effect, because they are characterizing information about an individual’s network based on a small slice of their social world (Ugander et al., 2012). Although this issue is particularly salient for nuances of analyzing online-social-network data, researchers need to be careful more generally to ensure that they are obtaining sufficient relevant information about an individual’s social world

whenever they attempt to relate individual differences in network centrality values (or other differences in the network characteristics of individuals) to different socio-behavioral tendencies. For instance, similar issues can arise if we use individual differences in centrality measures (e.g., degree centrality) based on a bounded social group (e.g., a high school), while failing to capture sufficiently many relevant aspects of individuals' social worlds. For example, in an analogous offline situation to the aforementioned online one, an individual may have small degree centrality in their high school but have many friends outside of the high school who are not captured if one calculates degree centrality based only on a high-school network. Therefore, when researchers are interested in interpreting a difference in social network position² as an individual difference measure (i.e., as a trait), it is advantageous to construct network data that captures people's full social worlds. When this is not possible (as is often the case), it is desirable to ask participants about their relationships outside of the social network that one is obtaining.

Tutorial: Example Social Network

Now that we have discussed some key concepts in network analysis that are particularly relevant for people who are interested in studying human social networks, we present a tutorial using a sample network. In this artificial network, we are interested in characterizing the network of a first-year dorm (with $n = 50$ students). Suppose that we obtained this data by asking participants to go through the list of everyone in the network and to identify whether or not they are friends with each individual (i.e., using a roster-based approach). This gives directed edges, because some friendships may not be reciprocated. If we are interested in understanding how individuals cognitively represent different members of the network or how individual differences in network measures are correlated with differences in cognition during a task, we can also try to obtain brain data from all or some of the members of the network. (We do not cover this idea in the tutorial.) The tutorial uses an artificial network with 50 nodes, which we label with people's names to facilitate exposition. (In real research situations, one should anonymize the nodes

prior to analysis.) We use the IGRAPH package in R (Csardi & Nepusz, 2006) to visualize the data and calculate various network measures—such as degree centrality, eigenvector centrality, and betweenness centrality—and to illustrate community detection. Our tutorial includes detailed comments on the practical application of the concepts that we have discussed in this paper. We also present a separate tutorial to illustrate visualization of multilayer networks using the PYMNET library in Python (Kivelä, 2017). Both tutorials can be accessed at https://github.com/elisabaek/social_network_analysis_tutorial. We hope that the two tutorials will be helpful for researchers who are interested in incorporating network measures in studies of individual cognition.

Future Directions

In the present paper, we have given an introductory overview of basic network ideas and concepts that we hope will provide a helpful starting point for social neuroscientists who are new to network analysis. Although the incorporation of network-analysis tools in social neuroscience is in its nascent stages, recent work that uses such tools has produced fascinating insights into how features of an individual's social world are represented in their brain. There are many open questions in the area, so it is a particularly exciting time to do research in it. In this, we highlight areas of future growth. We discuss both how social neuroscientists can integrate common network methods (see our earlier discussions of them) in new lines of inquiry and how to productively incorporate new developments and tools in network science and mathematics into future work in social neuroscience.

Open questions that leverage existing network tools. We begin by highlighting some of the many open questions in social neuroscience that can benefit from applying tools from network analysis. Although we will of course not be exhaustive, we hope to highlight the broad range of exciting research opportunities for social neuroscientists who are interested in using network analysis.

Information about different types of relationships. Several of the findings that we discussed highlight how the brain has mechanisms in place to track and spontaneously retrieve information about different aspects of social networks that are characterized by friendship, such as the extent to which individual members are popular (Zerubavel et al., 2015), socially valuable (Morelli et al., 2018), well-connected to well-connected others (Parkinson et al., 2017), and serve as brokers (Parkinson et al., 2017). These studies barely scratch the surface of the many different types of information about the social world that our brains may be tracking. People's lives consist not only of different types of social groups (e.g., friendship, professional, and family), but also of different types of information about the same social groups that may be important for successful social navigation. For instance, in the same group of friends, individuals may turn to different people when seeking emotional support versus when seeking career advice. Indeed, recent findings suggest that centralities in a social network can have different implications, depending on how one characterizes the relationships. For example, a recent paper (Morelli et al., 2018) related in-degree centrality in two different social networks—one with edges that encode trust and the other with edges that encode shared fun—in the same college first-year dorms to two different personality traits. People with better well-being were located more centrally in the fun network, and people with higher empathy were located more centrally in the trust network. Such findings suggest that where an individual is located in different social networks (i.e., with different edges) of the same social group may be associated with different behavioral outcomes.

Although this was not tested explicitly by Morelli et al. (2018), one possibility is that perceivers are also able to track the extent to which other individuals are located centrally in the two networks with the different types of edges (i.e., trust and fun), as this information may be important for guiding behavior in different contexts. For example, when seeking empathic support, it seems advantageous to seek individuals who are located centrally in the trust network. However, when looking to have fun, it may be appropriate to seek individuals who are

located centrally in the fun network. Future studies that explore how individual brains encode and retrieve information about social networks with different types of connections (i.e., edges) about the same social group may be particularly fruitful. For example, does prompting different motivations and situations lead individuals to switch between the types of social network from which they draw information, and how does the brain support this ability? One way to examine such questions may be through obtaining information about multiple types of relationships that individuals have with one another and then comparing how the brain retrieves information about the structure of these different relationships based on the motivational context. One can encode and analyze these relationships using a multiplex network, with each layer representing different types of relationships.

Characterizing different types of relationships in a social group can also help increase understanding of not only who is popular, but also those to whom others turn for support or empathy. Given that individuals who are more likely to seek social support to help regulate their emotions (i.e., interpersonal emotion regulation) tend to have better well-being and more supportive relationships (Williams, Morelli, & Zaki, 2018), one fruitful future direction may be to use centrality measures to identify supportive individuals (e.g., Morelli et al., 2018) and test how people's cognitive and affective processes may be affected by their social distance to these individuals or by the amount of time that they spend with these individuals (e.g., by incorporating weighted edges).

Individual differences in network features. A small body of research has also begun to explore associations between individual differences in network positions and differences in brain activity. Popular individuals (specifically, individuals with large in-degree centrality in a network in which edges represent being liked by others) tend to have greater neural sensitivity in the brain's valuation system in tracking the popularity of others in a network (Zerubavel et al., 2015), and people who are high in brokerage (in this case, the proportion of someone's Facebook friends who are friends with one another on Facebook) tend to exhibit greater activity in the

brain's mentalizing system when considering and incorporating social recommendations to make their own recommendations of consumer products to others (O'Donnell et al., 2017). It has also been illustrated that social status in non-human primates covaries with structural and functional differences in brain regions that are associated with social cognition (Noonan et al., 2014). In combination, these findings suggest that an individual's social-network position is associated with neural activity and behavioral responses to various everyday situations. There are many open questions, as only a few studies have related individual differences in social-network position to neural responses and even fewer have done so in the context of social decision-making (see, e.g., O'Donnell et al., 2017). Future studies that explore how individual differences in social-network position relate to neural responses during different social tasks (e.g., social influence, emotion regulation, interpersonal communication, and so on) may be particularly fruitful. Findings from such studies have the potential to advance understanding of how to distinguish influential individuals based on how they use their brains and in their responses to various social situations.

Causal relationships. With a few exceptions, the majority of the research that integrates neuroscience with network analysis has been cross-sectional or correlational (c.f. Sallet et al., 2011). Accordingly, many open questions remain for exciting developments for inferring causal directions of the various correlative findings that we have discussed in this paper. It remains unclear, for instance, whether (1) differences in neural responses cause differences in social-network characteristics or (2) differences in neural responses result from differences in an individual's social-network position. Experimental findings from nonhuman primates offer some clues, as it has been demonstrated that social-network characteristics (e.g., network size) causally affect the structure and functional responses in regions of the macaque brain that are associated with social cognition (Sallet et al., 2011). Although experimental manipulation of social networks in humans is admittedly more challenging to implement because of both practical and ethical concerns, longitudinal studies can also elucidate some of the ambiguity

about causality. Longitudinal studies that span key neural and social developmental periods, such as adolescence or older adulthood, may be particularly fruitful for providing insight into questions about the causal directions of effects.

Despite the challenging nature of experimental manipulation of social networks in humans, there are a few possible approaches for researchers to pursue. One possibility is by recruiting participants to join either offline or online interest-based communities and then randomly assigning participants to different social networks that one controls experimentally to vary in network characteristics of interest. For example, perhaps one wants the network to have a specific degree distribution, such as many people with small degrees and few people with large degrees. Such methods have been used previously to test how social-network characteristics influence the spread of behavior in online social networks (e.g., how similarity of contacts influence adoption of health behavior; Centola, 2010, 2011), but to our knowledge they have not yet been used with neuroimaging tools. Future studies that use similar experimental methods while also obtaining neural responses before and after individuals' experiences in a social network may further elucidate the causal directions of such observations. However, it remains unclear whether, and to what extent, an individual's cognitive and affective processes are influenced by artificially constructed social networks. Nevertheless, if successful, future studies that employ such approaches may provide valuable insights into causal relationships between social and neural phenomena.

Potential of incorporating new methods of network analysis. We now provide a brief overview of a few new methods in network analysis and related subjects that may be insightful for developing richer characterizations of social-network structures. We keep our descriptions brief because of the introductory nature of this paper, and we point interested readers to additional resources on the various topics.

Multilayer networks. As we discussed in previous sections, multilayer and temporal networks afford rich opportunities to examine how individual brains interact over time with the

social world in which they live (Aleta & Moreno, 2019; Kivela et al., 2014). For instance, multilayer network analysis will be useful for longitudinal studies to help understand how characteristics of a social network change over time, so such analysis may be able to inform the causal relationships of some of the previously observed findings that link brain activity and social-network characteristics. There is potential to use multilayer networks to examine interactions between brain networks and social networks over time to help predict behavior. It is also possible to analyze cognitive social structures using multilayer networks (Kivela et al., 2014). Tools from network science (including multilayer network analysis) have been used to analyze functional and anatomical networks in the brain (Bassett et al., 2011; Fair et al., 2008; Hutchison et al., 2013; Vaiana & Muldoon, 2018; van den Heuvel & Sporns, 2013), as well as to link these brain networks with social-network structures (Schmälzle et al., 2017) and with cognition and behavior (Bassett & Mattar, 2017; Mattar, Thompson-Schill, & Bassett, 2018). Recently, researchers have begun to recognize and highlight the potential advantages of using multilayer network analysis to represent such complex relationships; these efforts have the potential to advance understanding of processes of interest to social neuroscientists (Falk & Bassett, 2017). One potential fruitful application is to investigate how health behaviors change over time (Christakis & Fowler, 2007). For instance, one can use multilayer and temporal networks to study how to predict health-behavior change (e.g., quitting smoking) from changes in an individual's social network (e.g., joining a support group to stop smoking) through changes in functional networks in the brain (e.g., how regions in the brain's valuation system respond to smoking cues). Such a research question can contribute to broader understanding of how people's social environments impact neural processing and behavior.

Beyond pairwise connectivity in networks. We anticipate that it will also be fruitful to examine relationships between nodes beyond the usual pairwise connections. The simplest way to do this is with hypergraphs, which allow edges (which are called “hyperedges” in this context) to connect more than two nodes, which may be useful for representing relationships that involve

more than two people (Newman, 2018). One example is a coauthorship network, where a single hyperedge connects all of the coauthors of a manuscript, instead of connecting them through multiple pairwise edges. Another example is a network of college roommates, where it may be desirable to use a single hyperedge to connect all occupants of one room, which may be shared by more than two people. Moreover, it is possible that some pairs of roommates may also be connected directly in a pairwise fashion, so using hypergraphs gives a sensible way to simultaneously include both pairwise connections and other connections in a network structure. A more complicated, but likely very useful, approach to study relationships among arbitrarily many actors in a social network is to use “simplicial complexes” (Ghrist, 2014), an idea from algebraic topology that many researchers have leveraged for topological data analysis (Otter, Porter, Tillmann, Grindrod, & Harrington, 2017; Topaz, 2016). One can use tools from topological data analysis to systematically examine a diversity of structural features of social networks, such as by algorithmically finding topological “holes” (e.g., gaps) in coauthorship network. Perhaps such holes may help uncover barriers to academic collaboration, and it seems plausible to try to relate such topological holes to Burt’s notion of “structural holes” in social networks (Burt, 1992). The most popular approach from topological data analysis is “persistent homology”, which allows one to track many types of topological holes over multiple scales in a network. We anticipate that persistent homology and other tools from topological data analysis will be used increasingly in the study of social networks. See Topaz (2016) for a brief popular introduction and Otter et al. (2017) for a more mathematical introduction and a tutorial of available software.

Community-level characteristics. As we discussed briefly, one can examine densely connected communities of nodes in a network. There are numerous algorithms to study community structure; some of them involve assigning nodes to single communities, and others allow overlapping communities (Fortunato & Hric, 2016; Porter et al., 2009). One potential future direction that involves community structure and other large-scale network structures is to

simultaneously relate individuals' brain data to features of the local structures of their networks, characteristics of their intermediate-scale (i.e., "mesoscale") structures (such as communities), and global network characteristics of the network. As we have discussed, there exist numerous algorithms for identifying communities in a network (Fortunato & Hric, 2016; Porter et al., 2009). There are also methods for characterizing other types of intermediate-scale structures. One example is "core-periphery structure", in which one attempts to detect one or more cores of densely connected nodes, along with sparsely connected peripheral nodes (Csermely et al., 2013; Rombach et al., 2017). Another example is "role structure", in which one attempts to detect similar role structures of nodes (e.g., perhaps the ego networks of graduate students, postdoctoral scholars, and professors have different structural characteristics), regardless of the density of connections (Rossi & Ahmed, 2015). Future research that integrates tools for detection of communities and other mesoscale structures in networks may be fruitful for elucidating how the features of such large-scale structures impact individuals' cognitive processes and behavior.

Conclusions and Outlook

Recent research in social neuroscience that relates characteristics of people's social networks to their individual cognition offers fascinating insights into how the brain represents and may be influenced by its social context. Tools from network analysis provide rich opportunities for social neuroscientists who are interested in studying (1) how people navigate and interact with their complex social environments and (2) the mental architecture that supports these processes. Researchers can leverage existing tools and measures to study new questions, and they also have the opportunity to incorporate developing methods in network analysis for new lines of inquiry. The use of network analysis in social neuroscience is in its emerging stages, so this is a particularly exciting time, with wonderful opportunities to contribute to shaping the direction of the field.

References

- Aleta, A., & Moreno, Y. (2019). Multilayer networks in a nutshell. *Annual Review of Condensed Matter Physics*, 10(1), 45–62. <https://doi.org/10.1146/annurev-conmatphys-031218-013259>
- Ballinger, G. A., Cross, R., & Holtom, B. C. (2016). The right friends in the right places: Understanding network structure as a predictor of voluntary turnover. *Journal of Applied Psychology*, 101(4), 535–548. <https://doi.org/10.1037/apl0000061>
- Banerjee, A., Chandrasekhar, A., Duflo, E., & Jackson, M. (2014). Gossip: Identifying central individuals in a social network. In *SSRN Electronic Journal*. <https://doi.org/10.3386/w20422>
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2013). The diffusion of microfinance. *Science*, 341(6144). <https://doi.org/10.1126/science.1236498>
- Barrett, L., Henzi, S. P., & Lusseau, D. (2012). Taking sociality seriously: The structure of multi-dimensional social networks as a source of information for individuals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1599), 2108–2118. <https://doi.org/10.1098/rstb.2012.0113>
- Bassett, D. S., & Mattar, M. G. (2017). A network neuroscience of human learning: Potential to inform quantitative theories of brain and behavior. *Trends in Cognitive Sciences*, 21(4), 250–264. <https://doi.org/10.1016/j.tics.2017.01.010>
- Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., & Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences of the United States of America*, 108(18), 7641–7646. <https://doi.org/10.1073/pnas.1018985108>
- Bassett, D. S., Zurn, P., & Gold, J. I. (2018). On the nature and use of models in network neuroscience. *Nature Reviews Neuroscience*, 19(9), 566–578. <https://doi.org/10.1038/s41583-018-0038-8>
- Berlingerio, M., Coscia, M., Giannotti, F., Monreale, A., & Pedreschi, D. (2013). Multidimensional networks: Foundations of structural analysis. *World Wide Web*, 16(5–6),

567–593. <https://doi.org/10.1007/s11280-012-0190-4>

Bonacich, P. (1972). Factoring and weighting approaches to status scores and clique identification. *The Journal of Mathematical Sociology*, 2(1), 113–120.

<https://doi.org/10.1080/0022250X.1972.9989806>

Bonter, D. N., & Bridge, E. S. (2011). Applications of radio frequency identification (RFID) in ornithological research: A review. *Journal of Field Ornithology*, 82(1), 1–10.

<https://doi.org/10.1111/j.1557-9263.2010.00302.x>

Brent, L. J. N. (2015). Friends of friends: Are indirect connections in social networks important to animal behaviour? *Animal Behaviour*, 103, 211–222.

<https://doi.org/10.1016/j.anbehav.2015.01.020>

Brent, L. J. N., Ruiz-Lambides, A., & Platt, M. L. (2017). Family network size and survival across the lifespan of female macaques. *Proceedings of the Royal Society B: Biological Sciences*, 284(1854). <https://doi.org/10.1098/rspb.2017.0515>

Bringmann, L. F., Elmer, T., Epskamp, S., Krause, R. W., Schoch, D., Wichers, M., ... Snippe, E. (2019). What do centrality measures measure in psychological networks? *Journal of Abnormal Psychology*, 136(1), 23–42. <https://doi.org/10.1037/abn0000446>

Burt, R. S. (1992). *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press.

Burt, R. S. (2004). Structural Holes and Good Ideas 1. *American Journal of Sociology*, 110(2), 349–399. <https://doi.org/10.1086/421787>

Campbell, C. S., Maglio, P. P., Cozzi, A., & Dom, B. (2003). Expertise identification using email communications. *International Conference on Information and Knowledge Management, Proceedings*, 528–531.

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>
- Cherkassky, B. V., Goldberg, A. V., & Radzik, T. (1996). Shortest paths algorithms: Theory and experimental evaluation. *Mathematical Programming, Series B*, 73(2), 129–174.
<https://doi.org/10.1007/BF02592101>
- Christakis, N. A., & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *The New England Journal of Medicine*, 357(4), 370–379.
<https://doi.org/10.1056/NEJMsa066082>
- Christley, R. M., Pinchbeck, G. L., Bowers, R. G., Clancy, D., French, N. P., Bennett, R., & Turner, J. (2005). Infection in social networks: Using network analysis to identify high-risk individuals. *American Journal of Epidemiology*, 162(10), 1024–1031.
<https://doi.org/10.1093/aje/kwi308>
- Crossley, N., Bellotti, E., Edwards, G., Everett, M. G., Koskinen, J., & Tranmer, M. (2015). *Social network analysis for ego-nets*. London, United Kingdom: SAGE Publications.
- Csardi, G., & Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal, Complex Sy*, 1695. Retrieved from <http://igraph.org>
- Csermely, P., London, A., Wu, L. Y., & Uzzi, B. (2013). Structure and dynamics of core/periphery networks. *Journal of Complex Networks*, 1(2), 93–123.
<https://doi.org/10.1093/comnet/cnt016>
- Davis, A., Gardner, B. B., & Gardner, M. R. (1941). *Deep south: A social anthropological study of caste and class*. Chicago, IL: University of Chicago Press.
- Dodds, P. S., Muhamad, R., & Watts, D. J. (2003). An experimental study of search in global social networks. *Science*, 301(5634), 827–829. <https://doi.org/10.1126/science.1081058>
- Dunbar, R. I. M. (2008). Why humans aren't just Great Apes. *Issues in Ethnology and Anthropology*, 3(3), 15–33.
- Eagle, N., & Pentland, A. (Sandy). (2006). Reality mining: Sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4), 255–268. <https://doi.org/10.1007/s00779-005->

0046-3

Ellwardt, L., Labianca, G. J., & Wittek, R. (2012). Who are the objects of positive and negative gossip at work?. A social network perspective on workplace gossip. *Social Networks*, 34(2), 193–205. <https://doi.org/10.1016/j.socnet.2011.11.003>

Everett, M. G., & Valente, T. W. (2016). Bridging, brokerage and betweenness. *Social Networks*, 44(2010), 202–208. <https://doi.org/10.1016/j.socnet.2015.09.001>

Fair, D. A., Cohen, A. L., Dosenbach, N. U. F., Church, J. A., Miezin, F. M., Barch, D. M., ... Schlaggar, B. L. (2008). The maturing architecture of the brain's default network. *Proceedings of the National Academy of Sciences*, 105(10), 4028–4032. <https://doi.org/10.1073/pnas.0800376105>

Falk, E. B., & Bassett, D. S. (2017). Brain and social networks: Fundamental building blocks of human experience. *Trends in Cognitive Sciences*, 21(9), 674–690. <https://doi.org/10.1016/j.tics.2017.06.009>

Finn, K. R., Silk, M. J., Porter, M. A., & Pinter-Wollman, N. (2019). The use of multilayer network analysis in animal behaviour. *Animal Behaviour*, 149, 7–22. <https://doi.org/10.1016/j.anbehav.2018.12.016>

Firth, J. A., Voelkl, B., Crates, R. A., Aplin, L. M., Biro, D., Croft, D. P., & Sheldon, B. C. (2017). Wild birds respond to flockmate loss by increasing their social network associations to others. *Proceedings of the Royal Society B: Biological Sciences*, 284(1854), 20170299. <https://doi.org/10.1098/rspb.2017.0299>

Fortunato, S., & Hric, D. (2016). Community detection in networks: A user guide. *Physics Reports*, 659, 1–44. <https://doi.org/10.1016/j.physrep.2016.09.002>

Fowler, J. H., & Christakis, N. A. (2008). Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ*, 338(7685), 23–26. <https://doi.org/10.1136/bmj.a2338>

Gallotti, R., & Barthélemy, M. (2015). Anatomy and efficiency of urban multimodal mobility.

Scientific Reports, 4(1), 6911. <https://doi.org/10.1038/srep06911>

Gest, S. D., Farmer, T. W., Carolina, N., & Cairns, B. D. (2003). Identifying children's peer social networks in school classrooms: Links between peer reports and observed interactions.

Social Development, 12(4), 513–529. <https://doi.org/10.1111/1467-9507.00246>

Ghrist, R. (2014). *Elementary applied topology*. CreateSpace Independent Publishing Platform.

Gleich, D. F. (2015). PageRank beyond the web. *SIAM Review*, 57(3), 321–363.

<https://doi.org/10.1137/140976649>

Grund, T. U. (2012). Network structure and team performance: The case of English Premier League soccer teams. *Social Networks*. <https://doi.org/10.1016/j.socnet.2012.08.004>

Hampton, W. H., Unger, A., Von Der Heide, R. J., & Olson, I. R. (2016). Neural connections foster social connections: A diffusion-weighted imaging study of social networks. *Social Cognitive and Affective Neuroscience*, 11(5), 721–727. <https://doi.org/10.1093/scan/nsv153>

Heide, R. Von Der, Vyas, G., & Olson, I. R. (2014). The social network-network: Size is predicted by brain structure and function in the amygdala and paralimbic regions. *Social Cognitive and Affective Neuroscience*, 9(12), 1962–1972.

<https://doi.org/10.1093/scan/nsu009>

Herschel, R., & Miori, V. M. (2017). Ethics & Big Data. *Technology in Society*, 49, 31–36.

<https://doi.org/10.1016/j.techsoc.2017.03.003>

Holme, P. (2015). Modern temporal network theory: A colloquium. *The European Physical Journal B*, 88(9), 234. <https://doi.org/10.1140/epjb/e2015-60657-4>

Holme, P., & Saramäki, J. (2012). Temporal networks. *Physics Reports*, 519(3), 97–125.

<https://doi.org/10.1016/j.physrep.2012.03.001>

House, J. S., Landis, K. R., & Umberson, D. (1988). Social relationships and health. *Science*, 241(4865), 540–545.

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., ... Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations.

- NeuroImage*, 80(3), 360–378. <https://doi.org/10.1016/j.neuroimage.2013.05.079>
- Isella, L., Stehlé, J., Barrat, A., Cattuto, C., Pinton, J. F., & Van den Broeck, W. (2011). What's in a crowd? Analysis of face-to-face behavioral networks. *Journal of Theoretical Biology*, 271(1), 166–180. <https://doi.org/10.1016/j.jtbi.2010.11.033>
- Jeub, L. G. S., Balachandran, P., Porter, M. A., Mucha, P. J., & Mahoney, M. W. (2015). Think locally, act locally: Detection of small, medium-sized, and large communities in large networks. *Physical Review E*, 91, 012821. <https://doi.org/10.1103/PhysRevE.91.012821>
- Kaluza, P., Kölzsch, A., Gastner, M. T., & Blasius, B. (2010). The complex network of global cargo ship movements. *Journal of the Royal Society Interface*, 7(48), 1093–1103. <https://doi.org/10.1098/rsif.2009.0495>
- Kivelä, M. (2017). Multilayer Networks Library for Python (Pymnet 0.1). Retrieved from <http://www.mkivela.com/pymnet/>
- Kivela, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., & Porter, M. A. (2014). Multilayer networks. *Journal of Complex Networks*, 2(3), 203–271. <https://doi.org/10.1093/comnet/cnu016>
- Kleinberg, J. M. (2000). The small-world phenomenon: An algorithmic perspective. *Proceedings of the Annual ACM Symposium on Theory of Computing*, 163–170. <https://doi.org/10.1145/335305.335325>
- Kleinberg, J. M. (2011). Navigation in a small world. *The Structure and Dynamics of Networks*, 9781400841356(August), 551.
- Kolaczyk, E. D. (2009). *Statistical Analysis of Network Data: Methods and Models*. New York, NY: Springer Science+Business Media, LLC.
- Krackhardt, D. (1987). Cognitive social structures. *Social Networks*, 9, 109–134.
- Krackhardt, D. (1990). Assessing the political landscape: Structure, cognition, and power in organizations. *Administrative Science Quarterly*, 35(2), 342–369.
- Krause, J., Krause, S., Arlinghaus, R., Psorakis, I., Roberts, S., & Rutz, C. (2013). Reality

- mining of animal social systems. *Trends in Ecology & Evolution*, 28(9), 541–551.
<https://doi.org/10.1016/j.tree.2013.06.002>
- Lewis, K., Kaufman, J., Gonzalez, M., Wimmer, A., & Christakis, N. A. (2008). Tastes, ties, and time: A new social network dataset using Facebook.com. *Social Networks*, 30(4), 330–342.
<https://doi.org/10.1016/j.socnet.2008.07.002>
- Marsden, P. V. (2002). Egocentric and sociocentric measures of network centrality. *Social Networks*, 24(4), 407–422. [https://doi.org/10.1016/S0378-8733\(02\)00016-3](https://doi.org/10.1016/S0378-8733(02)00016-3)
- Masuda, N., Porter, M. A., & Lambiotte, R. (2017). Random walks and diffusion on networks. *Physics Reports*, 716–717, 1–58. <https://doi.org/10.1016/j.physrep.2017.07.007>
- Mattar, M. G., Thompson-Schill, S. L., & Bassett, D. S. (2018). The network architecture of value learning. *Network Neuroscience*, 2(2), 128–149.
https://doi.org/10.1162/netn_a_00021
- Milgram, S. (1967). The small-world problem. *Psychology Today*, 1(1), 61–67.
- Milgram, S. (1969). An experimental study of the small world problem. *Sociometry*, 32(4), 425–443. Retrieved from <http://www.jstor.org/stable/2786545>
- Morelli, S. A., Leong, Y. C., Carlson, R. W., Kullar, M., & Zaki, J. (2018). Neural detection of socially valued community members. *Proceedings of the National Academy of Sciences*, 201712811. <https://doi.org/10.1073/pnas.1712811115>
- Newman, M. (2006). Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, 74(3). <https://doi.org/10.1103/PhysRevE.74.036104>
- Newman, M. (2018). *Networks* (2nd ed.). New York, NY: Oxford University Press.
- Noonan, M. P., Sallet, J., Mars, R. B., Neubert, F. X., O'Reilly, J. X., Andersson, J. L., ... Rushworth, M. F. S. (2014). A neural circuit covarying with social hierarchy in macaques. *PLoS Biology*, 12(9). <https://doi.org/10.1371/journal.pbio.1001940>
- O'Donnell, M. B., Bayer, J. B., Cascio, C. N., & Falk, E. B. (2017). Neural bases of recommendations differ according to social network structure. *Social Cognitive and*

- Affective Neuroscience*, 12(1), 61–69. <https://doi.org/10.1093/scan/nsw158>
- Otter, N., Porter, M. A., Tillmann, U., Grindrod, P., & Harrington, H. A. (2017). A roadmap for the computation of persistent homology. *EPJ Data Science*, 6(1).
<https://doi.org/10.1140/epjds/s13688-017-0109-5>
- Padget, J. F., & Ansell, C. K. (1993). Robust action and the rise of the Medici, 1400-1434. *American Journal of Sociology*, 98(6), 1259–1319. <https://doi.org/10.1126/science.1165821>
- Parkinson, C., Kleinbaum, A. M., & Wheatley, T. (2017). Spontaneous neural encoding of social network position. *Nature Human Behaviour*, 1(5), 0072. <https://doi.org/10.1038/s41562-017-0072>
- Parkinson, C., Kleinbaum, A. M., & Wheatley, T. (2018). Similar neural responses predict friendship. *Nature Communications*, 9(1), 332. <https://doi.org/10.1038/s41467-017-02722-7>
- Peixoto, T. P. (2017). *Bayesian stochastic blockmodeling*. Retrieved from <http://arxiv.org/abs/1705.10225>
- Pilosof, S., Porter, M. A., Pascual, M., & Kéfi, S. (2017). The multilayer nature of ecological networks. *Nature Ecology and Evolution*. <https://doi.org/10.1038/s41559-017-0101>
- Porter, M. A., Onnela, J.-P., & Mucha, P. J. (2009). Communities in networks introduction: Networks and communities. *Notices of the American Mathematical Society*, 56(9), 19. Retrieved from <http://www.ams.org/notices/200909/rtx090901082p.pdf>
- Pustejovsky, J. E., & Spillane, J. P. (2009). Question-order effects in social network name generators. *Social Networks*, 31(4), 221–229. <https://doi.org/10.1016/j.socnet.2009.06.001>
- Read, J. M., Eames, K. T. D., & Edmunds, W. J. (2008). Dynamic social networks and the implications for the spread of infectious disease. *Journal of the Royal Society Interface*, 5(26), 1001–1007. <https://doi.org/10.1098/rsif.2008.0013>
- Rombach, M. P., Porter, M. A., Fowler, J. H., & Mucha, P. J. (2017). Core-periphery structure in networks. *SIAM Review*, 59(3), 619–646. <https://doi.org/10.1137/120881683>
- Rossi, R. A., & Ahmed, N. K. (2015). Role discovery in networks. *IEEE Transactions on*

Knowledge and Data Engineering, 27(4), 1112–1131.

<https://doi.org/10.1109/TKDE.2014.2349913>

Sallet, J., Mars, R. B., Noonan, M. P., Andersson, J. L., O'Reilly, J. X., Jbabdi, S., ...

Rushworth, M. F. S. (2011). Social network size affects neural circuits in macaques.

Science, 334(6056), 697–700. <https://doi.org/10.1126/science.1210027>

Salmivalli, C., Lagerspetz, K., Björkqvist, K., Österman, K., & Kaukiainen, A. (1996). Bullying as a group process: Participant roles and their relations to social status within the group.

Aggressive Behavior, 22(1), 1–15.

Santos, A. J., Daniel, R., Fernandes, C., & Vaughn, B. E. (2015). Affiliative subgroups in preschool classrooms: Integrating constructs and methods from social ethology and sociometric traditions. *PLoS ONE*, 10(7), e0130932.

<https://doi.org/10.1371/journal.pone.0130932>

Schmälzle, R., O'Donnell, M. B., Garcia, J. O., Cascio, C. N., Bayer, J., Bassett, D. S., ... Falk, E. B. (2017). Brain connectivity dynamics during social interaction reflect social network structure. *Proceedings of the National Academy of Sciences*, 114(20), 5153–5158.

<https://doi.org/10.1073/pnas.1616130114>

Seeman, T. E. (1996). Social ties and health: The benefits of social integration. *Annals of Epidemiology*, 6(5), 442–451.

Topaz, C. (2016). Topological data analysis. *Dynamical Systems Magazine*.

Uchino, B. N. (2006). Social support and health: A review of physiological processes potentially underlying links to disease outcomes. *Journal of Behavioral Medicine*, 29(4), 377–387.

<https://doi.org/10.1007/s10865-006-9056-5>

Ugander, J., Backstrom, L., Marlow, C., & Kleinberg, J. M. (2012). Structural diversity in social contagion. *Proceedings of the National Academy of Sciences of the United States of America*, 109(16), 5962–5966. <https://doi.org/10.1073/pnas.1116502109>

Ugander, J., Karrer, B., Backstrom, L., & Marlow, C. (2011). *The anatomy of the Facebook*

- social graph*. 1–17. Retrieved from <http://arxiv.org/abs/1111.4503>
- Vaiana, M., & Muldoon, S. F. (2018). Multilayer brain networks. *Journal of Nonlinear Science*, 1–23. <https://doi.org/10.1007/s00332-017-9436-8>
- van den Heuvel, M. P., & Sporns, O. (2013). Network hubs in the human brain. *Trends in Cognitive Sciences*, 17(12), 683–696. <https://doi.org/10.1016/j.tics.2013.09.012>
- Wang, H., Hernandez, J. M., & Van Mieghem, P. (2008). Betweenness centrality in a weighted network. *Physical Review E*, 77(4), 1–10. <https://doi.org/10.1103/PhysRevE.77.046105>
- Wasserman, S., & Faust, K. (1994). *Structural analysis in the social sciences. Social network analysis: Methods and applications*. New York, NY: Cambridge University Press.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, 393(June), 440–442.
- Waugh, A. S., Pei, L., Fowler, J. H., Mucha, P. J., & Porter, M. A. (2009). *Party polarization in Congress: A social networks approach*. Retrieved from <https://arxiv.org/abs/0907.3509>
- Weaverdyck, M. E., & Parkinson, C. (2018). The neural representation of social networks. *Current Opinion in Psychology*, 24, 58–66. <https://doi.org/10.1016/J.COPSYC.2018.05.009>
- White, D. R., & Borgatti, S. P. (1994). Betweenness centrality measures for directed graphs. *Social Networks*, 16(4), 335–346. [https://doi.org/10.1016/0378-8733\(94\)90015-9](https://doi.org/10.1016/0378-8733(94)90015-9)
- Williams, W. C., Morelli, S. A., & Zaki, J. (2018). Interpersonal emotion regulation: Implications for affiliation, perceived support, relationships, and well-being. *Journal of Personality and Social Psychology*, 115(2), 224–254. <https://doi.org/10.1037/pspi0000132>
- Wuchty, S., & Uzzi, B. (2011). Human communication dynamics in digital footsteps: A study of the agreement between self-reported ties and email networks. *PLoS ONE*, 6(11). <https://doi.org/10.1371/journal.pone.0026972>
- Zaheer, A., & Bell, G. G. (2005). Benefiting from network position: Firm capabilities, structural holes, and performance. *Strategic Management Journal*, 26(9), 809–825. <https://doi.org/10.1002/smj.482>

Zerubavel, N., Bearman, P. S., Weber, J., & Ochsner, K. N. (2015). Neural mechanisms tracking popularity in real-world social networks. *Proceedings of the National Academy of Sciences*, 112(49), 15072–15077. <https://doi.org/10.1073/pnas.1511477112>