

Prediction and Topological Models in Neuroscience

Bryce Gessell, Matthew Stanley, Benjamin Geib, and Felipe De Brigard

Duke University

Correspondence should be addressed to:

Felipe De Brigard

203A West Duke Building

Duke University

Durham, NC

27708

Abstract: In the last two decades, philosophy of neuroscience has predominantly focused on explanation. Indeed, it has been argued that mechanistic models are the standards of explanatory success in neuroscience over, among other things, topological models. However, explanatory power is only one virtue of a scientific model. Another is its predictive power. Unfortunately, the notion of prediction has received comparatively little attention in the philosophy of neuroscience, in part because predictions seem disconnected from interventions. In contrast, we argue that topological predictions can and do guide interventions in science, both inside and outside of neuroscience. Topological models allow researchers to predict many phenomena, including diseases, treatment outcomes, aging, and cognition, among others. Moreover, we argue that these predictions also offer strategies for useful interventions. Topology-based predictions play this role whether or not they do or can receive a mechanistic interpretation. We conclude by making a case for philosophers to focus on prediction in neuroscience in addition to explanation alone.

Key words: Prediction; Network Science; Topology; Neuroscience

Prediction and Topological Models in Neuroscience

1. Introduction

Contemporary philosophy of neuroscience has in large part been dominated by a focus on *explanation*. This focus follows a more general trend in the philosophy of science, where causal or mechanistic explanations have overtaken law-based explanations as the preferred means for understanding much of the world. Indeed, philosophers have produced compelling arguments for mechanistic models being the standard of explanatory success in the biological sciences (Craver and Darden, 2013). However, when it comes to the enterprise of science, the explanatory power of theoretical models is only one of many virtues (Schindler, 2018). Another virtue of theoretical models, which in recent years has received comparatively less attention in the philosophy of science, is *prediction*, despite it being once heralded as equally relevant as explanation among the goals of science (Hofstadter, 1951; Popper, 1963; Lakatos and Musgrave, 1970; Salmon, 1978). This absence is particularly noticeable in the *philosophy of neuroscience*, as there has been almost no discussion on the predictive power or value of theoretical models in neuroscientific research. When contrasted with the fact that contemporary neuroscience is heavily engaged in generating predictive models (e.g., Yarkoni and Westfall, 2017), the dearth of discussion on prediction in the philosophy of neuroscience is even more remarkable.

Pretheoretically, many people think of prediction as synonymous with prognostication or forecasting, meaning that that which is predicted has not occurred yet. This time-dependent view of prediction contrasts with a knowledge-dependent view, according to which what one predicts may or may not have already occurred, as long as it is not known. In this paper, we adopt this knowledge-dependent or epistemic reading of prediction and side with Barrett and Stanford (2004)

in defining a prediction as “a claim about known matters of fact whose truth or falsity has not already been independently ascertained by some more direct method than that used to make the prediction itself” (586). Moreover, successful predictions in general enhance our epistemic standing, not by way of supplying further explanatory details, but by reducing our uncertainty as to what to expect under certain conditions, and by providing us with strategies to effectively intervene and manipulate phenomena. Of course, successful predictions often lead to improved explanations (Douglas, 2009); however, even without this additional bonus, successful predictions have value in and of themselves.

Perhaps a key reason as to why there is so much emphasis on explanation (as compared to prediction) in the philosophy of science in general, and of neuroscience in particular, is that there is a clear relationship between explanation and intervention. For many scientists and philosophers, the scientific goal of unveiling the real nature of the world is at least as important as that of offering strategies to intervene and control it (e.g., Longino, 2002). Given that mechanistic models provide both descriptions of natural phenomena and approaches to manipulate such phenomena, it is unsurprising that such models are taken as ideal candidates as to how to best pursue research in neuroscience (Craver, 2007). The current chapter, however, puts pressure on this view by highlighting the connection between the *predictive* power of certain theoretical models in neuroscience and their value as strategies for manipulation and intervention (Douglas, 2009). Importantly, the kinds of theoretical models we have in mind are *topological models*, which have recently been the subject of discussion in the philosophy of science, with some arguing that they offer an alternative kind of explanation, different from mere causal or mechanistic explanation (Huneman, 2010; Lange, 2016), and others arguing that they do not (Craver, 2016; Povich and Craver, 2018). We will largely sidestep this discussion, however, as we seek to explore the

predictive rather than explanatory ambitions of topological models in neuroscience (with occasional mention of other disciplines too), and the role they can play in our capacity to intervene, manipulate, and control neural phenomena. To reiterate: our arguments seek neither to support nor to undermine the claim that topological models are explanatory, nor whether or not they are so in virtue of receiving a mechanistic interpretation. We want to argue instead for a different claim, namely that whether or not topological models receive a mechanistic interpretation, they still hold predictive value and can be reliable guides to intervention and manipulation. Moreover, we put forth the more general claim that good predictions ought to be a central goal of neuroscience, whether or not they are afforded by models that have (or even could receive) a complete mechanistic interpretation.

The chapter will proceed as follows. In section 2, we offer a brief discussion on the relationship between prediction and explanation, and we place the role of mechanistic models in general, and in the philosophy of neuroscience in particular, within that dialectic. Next, in section 3, we discuss the nature of topological models and their use in prediction and interventions in a number of different fields before focusing on the use of topological models in network neuroscience for prediction. We also show how these models can be useful for intervention and manipulation even absent a mechanistic understanding of their underpinnings. Finally, in section 4, we draw some general conclusions and questions for future research.

2. Prediction, explanation, and mechanistic models

To fully understand the relationship between intervention (or manipulability), on the one hand, and mechanistic models in neuroscience, on the other, it may be useful to begin with a brief excursus into the history of the debate on the relationship between explanation and prediction in

the philosophy of science (for a recent excellent review see Douglas, 2009). This will allow us to better locate the role of mechanistic models in neuroscience within this dialectic.

2.1. Prediction and explanation: A brief history

Although one can find interesting discussions about the relationship between explanation and prediction in science in the works of Hume (1748), Whewell (1840), and Mill (1843), contemporary scholarship on the subject usually starts with the deductive-nomological (DN) model proposed by Hempel and Oppenheim (1948). According to the DN model, the explanandum (i.e., the statement to be explained) must deductively follow from the explanans: a set of premises that not only should be true but also include boundary conditions and general laws. According to the DN model, in its simplest form, a scientific explanation would have the following structure:

$$\begin{array}{l}
 C_1 \wedge C_2 \wedge C_3 \wedge \dots \wedge C_n \\
 L_1 \wedge L_2 \wedge L_3 \wedge \dots \wedge L_n \quad \left. \vphantom{\begin{array}{l} C_1 \wedge C_2 \wedge C_3 \wedge \dots \wedge C_n \\ L_1 \wedge L_2 \wedge L_3 \wedge \dots \wedge L_n \end{array}} \right\} \textit{Explanans} \\
 \therefore \textit{Explanandum}
 \end{array}$$

Here, each C_i is a true statement of a boundary condition or particular occurrence of an event, and each L_i is a statement of a general law. Thus, a successful explanation of, say, a particular observation of a planet’s location at a particular time would be given by a set of premises involving other relevant observations and empirical conditions, as well as by certain physical laws governing celestial bodies. Importantly, for Hempel and Oppenheim, both explanations and predictions shared the same logical structure, as it is possible for an explanans to state a yet unobserved event.

Predictions, as it were, are explanations of future events; or, alternatively, explanations are predictions of past events (*aka* “postdictions”).

This view, known since as the “symmetry thesis”, was met with substantial backlash in the 1950s and 1960s. It was pointed out, for instance, that while explanations require true propositions, successful predictions need not (Scheffler, 1957). Others argued that the uncertainty that applies to explanations differs from that which applies to predictions (Helmer and Rescher 1959), while still others pointed out that some theoretical models—such as in quantum mechanics (Hanson, 1959) and evolution (Scriven, 1959)—are good at explaining but bad at predicting. By the time Nagel’s *The Structure of Science* (1961) was published, the focus in philosophy of science had almost entirely moved to explanation for, as he remarked, “the distinctive aim of the scientific enterprise is to provide systematic and responsibly supported explanations” (1961, 15). The displacement of prediction—or the “decentering of prediction” as Heather Douglas (2009) calls it—brought explanation to the forefront of philosophical scholarship in the philosophy of science.¹

With prediction relegated to the background, most discussions focused on whether or not the DN model offered a successful analysis of scientific explanation. Philosophers quickly grew dissatisfied with the logical structure of explanations offered by the DN model. Some of the first concerns pertained to the difficulty of distinguishing statements of non-accidental generalizations from those of scientific laws (Hempel, 1965). Soon after, counterexamples to the DN model started to emerge. Some pointed out explanatory asymmetries, as in the example in which the length of a flagpole is deductively derived from the length of its shadow in conjunction with relevant laws about the propagation of light (Bromberger, 1966). Such a derivation, it was argued, conforms to

¹ It is important to note that this decentering may not apply to other related areas of research, such as issues on confirmation and accommodation, both of which are related to the notion of prediction (see, for instance, Eells, 2006). We thank a reviewer for inviting us to note this issue.

the logical structure of the DN model, yet we feel that the explanans (i.e., the length of the flagpole) is not really explained by the explanandum (i.e., the length of the shadow plus laws pertaining to the propagation of light). Other counterexamples pointed at cases of explanatory irrelevance, as with the case of the following derivation (Salmon, 1971):

(P1) All males who take birth control pills regularly fail to get pregnant

(P2) John Jones is a male who has been taking birth control pills regularly

∴ John Jones fails to get pregnant

which seems to conform to the structure of the DN model—i.e., P1 satisfies the criteria of lawfulness, and P2 states particular true observations—and yet does not constitute a successful explanation.

As a consequence, the 1970s and 1980s saw a proliferation of models of scientific explanation, including Salmon's statistical-relevance (SR) model (Salmon, 1971), the causal model (Salmon, 1984), and the unification model (Kitcher, 1989), to name a few. Unsurprisingly, most of the scholarship on scientific explanation during those two decades boiled down to a series of exchanges between counterexamples and defenses of these various models. By the 1990s, no agreed-upon model of explanation was in the offing and, instead, philosophers of science largely moved toward some kind of explanatory pluralism. Arguments turned into discussions as to what sort of explanatory model would be more appropriate for each scientific discipline. This was the intellectual environment in which the mechanistic explanation model was fully articulated (Machamer, Darden and Craver, 2000), and the following years helped to strengthen it as the

paramount model for scientific explanation in the life sciences, including neuroscience (Craver, 2007).

2.2. The mechanistic model of explanation in neuroscience

Although there are several definitions of “mechanism” and “mechanistic explanation” in the philosophy of science (e.g., Machamer et al, 2000; Glennan, 2002; Bechtel and Abrahamsen, 2005), they all seem to agree on what Craver and Tabery (2015) call the “ecumenical” characterization of mechanism, according to which mechanisms consist of four components. First, there is the *phenomenon*, which is understood as the behavior of the system that the mechanism constitutes. Every mechanism, then, is a mechanism of some particular phenomenon—e.g., digestion, long-term potentiation, inattentional blindness—and, depending on the particular phenomenon, a mechanism can produce, underlie, or maintain it. The second component are the *parts* of the mechanism which, in turn, are linked by the third component: *causal relations*. Considerable discussion has ensued regarding the best characterization of causal relations for mechanistic explanations. For our purposes, what matters is that such causal relations are *intervenable*, that is, they can be in principle—even if not in practice—manipulated to make a difference in the phenomenon. Finally, mechanisms are also *organized* in some fashion. In the case of neuroscience, many think of mechanisms as hierarchically organized in different levels (Craver, 2007), but other organizations may be possible too (Craver and Tabery, 2015).

Mechanistic models in neuroscience, then, are useful for the purpose of explanation insofar as they can capture a mechanism. And they can capture a mechanism if they conform to what Craver and Kaplan (2011) call the “3M” mapping requirement:

(3M) A model of a target phenomenon explains that phenomenon when (a) the variables in the model correspond to identifiable components and organizational features of the target mechanism that produces, maintains, or underlies the phenomenon, and (b) the causal relations posited among these variables in the model correspond to the activities or operations among the components of the target mechanism. (Craver and Kaplan, 2011, 272)

It is likely that, as of now, we do not have a single mechanistic model that fully conforms to the 3M requirement and that provides a complete characterization of all the components. At best, we have *schematic* models: abstract or idealized descriptions of a mechanism in which many of the details are omitted and/or that include provisional place-holders for unknown components (Darden, 2002). Moreover, mechanistic schematic models also vary in terms of the degree to which they capture the actual phenomenon. On one extreme, *how-possibly* models describe mechanisms in terms of how the different parts *might* be causally related and organized to produce, maintain, or support a phenomenon. On the other extreme, *how-actually* models depict how they are actually causally related, what all the parts really are, and how the parts are in reality organized to produce, maintain, or support a phenomenon. Unsurprisingly, we likely have very few—if any—how-actually models in neuroscience; these constitute a normative goal that our constantly refined how-possibly mechanistic schematic models seek to reach (Craver and Darden, 2013). Much of the scientific work in contemporary neuroscience consists precisely in discovering the underlying components of a mechanistic model to provide interpretations of the filler terms that can bring a how-possibly model closer to a how-actually one. Consider our current model of long-term potentiation (LTP) in neurons in the dentate gyrus. While this may constitute one of the most thorough mechanistic models in neuroscience, researchers keep discovering new details that help

to make certain assumptions and idealizations more concrete. For instance, while early models postulated that N-methyl-D-aspartate receptors (NMDAR) were necessary to trigger the induction of LTP (Collingridge et al, 1983), more recent discoveries have shown that other receptors, such as metabotropic glutamate (mGluRs) and kainate receptors can do it as well. More recently still it has been shown that even Ca²⁺-permeable α -amino-3-hydroxy-5-methyl-4-isoxazole propionic acid receptor (AMPA) can do the trick, further inviting the revision of the actual components of our mechanistic model of LTP (Park et al., 2014). Despite being one of the most complete mechanistic models in neuroscience, our current model for LTP is not a how-actually model quite yet; at best, it is a how-nearly-actually mechanistic model (Craver and Tabery, 2015).

Nevertheless, mechanistic models seem perfectly appropriate to deliver on what arguably are the two main goals of the scientific enterprise: to uncover the nature of reality, and to enable us to manipulate and control it. Mechanistic models—as opposed to the DN-, the SR-, and some variants of unificationist and mathematical models—are ideally suited to contribute toward the first goal, insofar as they care less about the logical structure of the explanation and more about its ontic commitments, that is, the kinds of actual, real structures that count as legitimately explanatory (Craver, 2014).² The explaining is done by *real* stuff, causally related and organized in various ways in order to produce, sustain or underlie a phenomenon. Mechanistic models not only tell us *why* something happens, but also *what* makes it happen. In turn, mechanistic models contribute to the second goal thanks to their reliance on counterfactual theories of causation, particularly manipulationist views (Woodward, 2003). When the causal relations are thus

² There are some views of mechanistic models that need not have such strong ontic commitments (e.g. Bechtel, 2008) and/or that need not be committed to a manipulationist/counterfactual-dependent account of causation. It is possible that some of the arguments we discuss here do not necessarily apply to these accounts. We don't discuss these accounts in depth, in part because they are not as thoroughly developed in the philosophy of neuroscience. Thanks to a reviewer for inviting us to clarify this point.

understood, the parts of a mechanism that constitute the *relata* can be seen as variables able to make a difference to the phenomenon—i.e., the behavior of the mechanism they are part of. In other words, mechanistic models enable us to tell what would happen to the phenomenon if one were to intervene on a particular variable (i.e., a part) at a certain level of organization. Thus, mechanistic models are ideal to tell us how the phenomenon would behave under counterfactual conditions and, consequently, they seem perfectly suited to offer predictions as well.

Given all these considerations it is hard not to think of mechanistic models as the paradigmatic model for not only scientific explanations, but also scientific *predictions* in neuroscience. In fact, some mechanists seem to suggest as much. They claim that understanding how a phenomenon works via subsuming it under a mechanistic model is perhaps the most reliable way to predict how it will behave in the future, and how it can be manipulated so that we can make it “work for us” (Craver and Kaplan, 2011). A strong reading of this view would imply that models can *only* yield successful predictions if they have strong ontic commitments to the structures they represent, and if they offer, if not a how-actually, at least a how-nearly-actually or a how-plausibly mechanistic schema of the phenomenon.³

In what follows, we argue against this strong reading according to which mechanistic models are paradigmatic models for both explanation *and* prediction, particularly as they apply to neuroscience. Instead, we argue for a weaker view, according to which, even if mechanistic models are ideally fitted for generating explanations in neuroscience, there may be some non-mechanistic models that are well suited to offer not only successful predictions but also strategies to manipulate

³ We see successful predictions as those which accurately model alternative outcomes (and thus support counterfactuals to some degree), or model future states with accuracy significantly above chance. In short, good predictions estimate outcomes above randomness. Note that, on this view, how-actually and how-possibly models can both yield successful predictions; however, how-actually models may not always make predictions that are perfectly accurate, since their use is often limited to certain contexts (consider the difference between Newtonian and relativistic physics, for example).

and control certain phenomena. In particular, we defend this weaker reading in relation to topological models, which have recently been criticized by mechanists who argue that they are not explanatory, or that, if they are, they explain precisely because they ultimately resolve into a mechanistic model (Craver, 2016). The suggestion we put forth in the rest of this essay is that even if topological models only have explanatory value when translated into their mechanistic components, they still have predictive value whether or not they have clear ontic commitments and/or mechanistic interpretations.⁴

3. Network science: prediction and interventions

Network science makes use of the mathematical tools and formalisms from graph theory to empirically investigate real-world networks. In its simplest form, a network can be thought of as a collection of differentiable elements, or *nodes*, and the pairwise relationships between them, or *edges*. Diverse real-world systems can be thought of as networks. For example, protein-protein interaction networks, structural and functional brain networks, infectious disease networks, friendship networks, and air transportation networks have all been modeled as networks for various purposes. Despite the obvious differences in the actual, real-world phenomena, all these networks can be understood as collections of nodes with certain edges between them (Butts, 2009). But of course, what the nodes and edges actually represent in the world will differ across the different kinds of networks (Figure 1).

[Figure 1 about here]

⁴ A clarification: we are *not* saying that Craver is necessarily committed to the strong reading. As far as we know, partisans of mechanisms have said little as to whether or not predictive models also demand the same ontic commitments that explanatory models do. Our view should rather be seen, then, as an admonition to the effect that even if one adopts a mechanistic stance vis-à-vis the way in which neuroscience ought to be pursued, then the strong ontic commitments that have been argued for explanation need not apply to prediction too.

Graph theoretic metrics can then be used to characterize the topological properties of these networks—regardless of how the nodes and edges are defined in practice (Watts and Strogatz 1998; Butts, 2009; Huneman, 2010; Sporns, 2011). A simple example of a topological property is *geodesic distance*: the minimum number of edges required to transverse from one particular node i to another node j in the network. You and your Facebook friend have a geodesic distance of 1, because it only takes one edge to connect you and your friend. But the geodesic distance between you and a friend of that friend who is not also your friend on Facebook would then be 2. Thus, geodesic distance, for instance, can help to calculate the spread of information on your Facebook wall. Relatedly, the *path length* of any node i in a network can be obtained by computing the average shortest number of steps necessary to get from i to each other node in the network (Dijkstra, 1959). Path length offers an indication of how quickly or effectively information can spread throughout a network. Consider, for example, a large hierarchically structured company. The CEO likely has a relatively short path length, and information can be transmitted from the CEO to any employee in relatively few steps, whereas most low-level employees likely have a longer path length, as it takes more steps for them to communicate with members in faraway departments. A more complex graph theoretic metric is *eigenvector centrality*, a measure of the extent to which a node i is connected with other influential nodes in the network (nodes with lots of edges). Nodes with high eigenvector centrality are thought to be highly influential and effective in spreading information throughout a network. On social media (e.g., Twitter), for example, certain celebrities like Justin Bieber tend to have particularly high eigenvector centrality, as they tend to be connected with many other influential celebrities.

Topological and spatial scales can be changed depending on a researcher's interests. To give an example from neuroscience, the hippocampus can be studied as a single structure or unit,

it can be studied as a three-part entity composed of CA1, CA3, and the dentate gyrus, or it can be studied as a more complex structure containing various cell types, layers, and their projections. Although it is often tempting to view phenomena at higher resolutions (e.g., cell types and the properties of those cells) as being the worthiest of serious investigation, it is sometimes not useful or valuable to improve the resolution with which one studies a given brain structure and its relation to cognition—especially when current computational and practical constraints are taken into account. Investigation at a more macroscopic scale often still yields useful and accurate predictions. Because it is unclear which level of granularity is the ground truth, and so unclear how best to demarcate components of the system (i.e., nodes representing functional units in the brain), topological prediction can play a central role. Comparing predictive utility at different levels of granularity can also guide future research and serve useful purposes. It is possible that different “scales of granularity” of network description will yield distinct yet complementary properties for predicting cognitive phenomena, disease states, and disease progression, among other things.

There are many other graph theoretic metrics that capture certain topological properties of networks, such as eigenvector centrality, clustering coefficients, and modularity (Newman, 2010). Critically, one of the central features of network models is that topological properties can be ascertained independently of a system’s physical substrates. That is, the same graph metrics can be computed on any kind of network, no matter what the nodes and edges are representing in the world; networks comprised of differently defined nodes and edges can even possess the same topological properties. To be sure, there are interesting philosophical questions about the nature of such topological properties and their relationship to the actual substrates the network models are based on. We also believe that understanding whether or not the topological properties of network models have any explanatory value above and beyond the mechanisms that underlie the system

they seek to represent, is a worthwhile philosophical question (Huneman, 2010; Craver, 2016, Lange, 2017). That being said, we also believe that the longstanding emphasis on explanation in philosophy of science, as well as the fact that network models have mainly been discussed in reference to alternative explanatory frameworks, have obscured the fact that topological properties in network models have remarkable *predictive* power. Additionally, in some instances, the predictive power of topological properties in network science has enabled us to conduct successful interventions. Let us explore some examples.

3.1. The predictive power of topological properties

We often obtain good predictions when causal information about the components of the system is incorporated into the model. However, in some cases, clear causal information is either unavailable, non-existent, or poorly-defined. Even in such cases, networks can still be characterized topologically, and their topological properties can produce accurate predictions. Studies of co-authorship networks, for example, capture patterns of collaboration in a given field. These networks allow us not only to identify prominent author(s) in a field, but also to successfully predict whether a publication will be well-cited in the future. For example, Sarigol, Pfitzner, Scholtes, Garas, and Schweitzer (2014) analyzed a dataset of over 100,000 publications from the field of computer science, and they investigated how centrality in the co-authorship network differs between authors who have highly cited papers and those who do not. Using a machine learning classifier based only on co-authorship network centrality measures (degree centrality, eigenvector centrality, betweenness centrality, and k-core centrality), they were able to predict whether an article would belong to the 10% most cited articles in five years' time with a precision of 60%, well above chance. Interestingly, in order to not overemphasize one particular dimension of

centrality in networks, they used several complementary measures of network centrality, and this combination of measures was crucial in adequately predicting the publication “success” of the researchers. To compute each centrality metric, however, it was first necessary to define the full set of nodes and edges in the network. By mapping out all connections in the network and computing graph metrics, they quantitatively suggested the existence of a social bias, manifesting itself in terms of visibility and attention, and influencing measurable citation “success” of researchers.

Another example pertains to traffic congestion. Consider the following question: “how can we accurately predict which roads in a city have or will have the highest occurrence of traffic jams?” A network approach might seek to predict whether a road will be congested by examining its topological properties within the larger network. This requires taking into account all other roads in the network (in this case, nodes might represent intersections, and edges might represent the road segments that link the intersections). Note that the kinds of buildings to which the roads provide access is not included in defining nodes and edges, and therefore, this information will not be included in—directly or indirectly—predicting traffic congestion.

Adopting this network approach, Wang et al. (2012) show that by incorporating information about how centrally a road is situated in graph theoretic space, they can accurately predict traffic patterns in San Francisco and Boston. The extent to which a road segment occupies a central place in the city grid is measured in terms of a mathematical property of networks known as *edge betweenness*. Edge betweenness is computed for each individual edge in a graph. To compute the edge betweenness, a search algorithm identifies the shortest possible path between each and every node in the network. It then searches the resulting data structure to determine what proportion of those paths incorporate the road segment in question. That proportion is the edge

betweenness. In this particular case, edges represent road segments defined as stretches of roads between legal intersections, and nodes represent legal intersections. Wang et al. (2012) show that the traffic density on a road segment can be better predicted by modeling both the road's centrality and inherent travel demand (i.e., how often the buildings on the road are frequented) than it can by modeling inherent travel demand alone. That is, incorporating edge betweenness into the predictive model actually provides better predictions above and beyond travel demand alone. Importantly, one can only compute edge betweenness by completely searching the entire topographical structure of the system, even though the measure is computed for an individual road (i.e., edge). With this particular kind of quantitative description, one might then be able predict that in other major metropolitan cities in the United States (e.g., Houston, Phoenix, Chicago, Dallas, etc.), roads with higher edge betweenness will experience more traffic jams, on average.

Another example comes from sexual networks, whereby persons are thought of as nodes and sexual contacts as edges. Long-term and large-scale data collection has led to the production of large-scale sexual networks from Manitoba, Canada, and from Colorado Springs, USA (Woodhouse et al. 1994; Rothenberg et al. 1998; Wylie & Jolly 2001; Jolly & Wylie 2002; Potterat et al. 2002). These kinds of networks highlight the heterogeneities present in sexual networks and show the importance of core groups (i.e., highly and disproportionately interconnected subsets of people with high numbers of contacts) and 'long-distance' connections (linking otherwise distant parts of the network) in disease transmission. Note that it is only possible to uncover these core groups (i.e., network modules) and 'long-distance' connections that interconnect groups by mapping out the full structure of the sexual network. Moreover, edges are defined only by whether two individuals have sex with each other during some time period t . To provide a particularly salient example, Liljeros et al. (2001) showed that sexual networks, like many networks that are

present in the world, have a scale-free degree-distribution (in contrast to, for example, a Gaussian distribution). This property means that the vast majority of individuals in the network have very few sexual contacts, but that there are a few individuals who have had a very large number of sexual contacts. Importantly, the fact the network has a scale-free architecture suggests that some of the individuals with a very large number of partners may bridge relatively isolated communities, i.e. they have long-distance connections in addition to many connections.⁵

On the surface, it may seem as though predictions about human sexual networks are underwritten not by the topological properties of these networks, but instead by our knowledge of the actual causal properties involved. For example, we know a lot about human sexual contact, and can give accurate microbiological explanations of how some STDs pass from one person to another. However, the force of this example is that our predictions about human sexual networks would still be accurate, *even if we had none of this causal and biological knowledge*. Suppose that we were examining sexual networks in an alien species, for example. The topological properties would still be helpful in predicting disease transmission among members of the species, even if we had no knowledge, detailed or otherwise, about the alien biology.

Moreover, this example offers an interesting case in which quantitatively characterizing the topological properties of the network allows us to identify particular individuals in the network who could be targeted for a *subsequent* intervention in response to a sexual-disease outbreak, as

⁵ We say that a scale-free architecture “suggests” this organization of individuals because, while not a mathematical guarantee, it appears likely to be so. In a scale-free network architecture, statically speaking, some of the high-degree nodes will be provincial hubs and some of the high-degree nodes will be connector hubs. Granted, it is not the case that networks must not follow this principle; in some scale-free networks, all the high-degree nodes might be connectors. But this seems statistically unlikely as then distinct modules are unlikely to exist. If the high-degree nodes are “randomly” arranged, then some must be connectors and some must be provincial. In other words, in scale-free networks, the nodes at the far end of the distribution have considerable influence over the other nodes in the network, more so than in other kinds of networks with other kinds of degree distributions. Some of these nodes with very many connections are likely to interconnect many different communities and be essential (in the example from the text) for diseases to propagate throughout the network.

given limited resources, it may not be practical to target all individuals in the network to promote safe-sex practices. But by specifically targeting individuals who have the greatest number of connections and those who tend to disproportionately connect otherwise distant groups or clusters in the network, it may be possible to most efficiently and effectively promote safe-sex practices to reduce the likelihood of disease transmission across the entirety of the network.

Let us summarize the three examples we have seen. The first dealt with a co-authorship networks, and focused on the predictive power of centrality measures. In this example, causal information for entities in the network is either non-existent or poorly defined; nevertheless, centrality measures are still able to help us generate successful predictions. The second example was about traffic congestion. Here, researchers used edge betweenness and other topological measures to describe global properties of the network, thereby making it possible to give accurate predictions about different cities based on the successes of a single network analysis. The third and last example, about sexual networks, illustrated the power of topological models to identify the degree distribution of a particular network—a scale-free distribution in this case, as opposed to, say, a Gaussian distribution. Most importantly, the example of sexual networks shows that causal knowledge—even when it is available and detailed—is not necessary for making predictions in virtue of topological properties.

Taken together, what these and related examples strongly suggest is that the topological properties of network models can be successfully employed to make predictions and to guide interventions on the systems they represent even when no causal or mechanistic information about the system is either known or included in the model. In fact, the examples just discussed are agnostic as to their ontic commitments, as they tend to be independent of the precise physical substrata of the modeled system. Crucially, many of these models offer clear avenues for

intervention. For example, in the case of sexual contact networks, some individuals in the network have more connections than others, and certain individuals are disproportionately responsible for interconnecting relatively segregated communities in the network. By identifying and targeting those individuals, we might be more likely to stop the spread of disease. Of course, it may be possible that some of these predictions illuminate the underlying nature of the phenomenon, and as such may contribute to its explanation. But even if they don't, the topological properties of network models still hold enormous epistemic value by enabling us to make predictions and by offering the possibility of gaining some measure of control over social and natural processes (Douglas 2009, 2013).

3.2. Predictions and interventions in network neuroscience

The aforementioned considerations were confined to network models outside the field of neuroscience. The question now is whether or not we have evidence to the effect that network models and their topological properties can also afford predictions and strategies for manipulation *within* in neuroscience. We believe they can. Some contemporary neuroscientists treat the brain as a large-scale network. Determining what constitutes a node or an edge, however, is tricky, as it depends on the particular level of analysis, the particularities of the research questions, and the idiosyncrasies of the available technologies from which the data is acquired (Stanley et al., 2013). Regarding levels of analysis, brains can be seen as varying along at least three scales. First, there is a spatial scale that ranges from the very micro (e.g., neurons, glia) to the very macro (e.g., gross anatomical regions comprising millions of neurons and even more synapses connecting those neurons). Thus, while networks at the micro-level may include neurons as nodes and synaptic connections as edges, networks at the macro-level may include cytoarchitecturally delimited

portions of brain tissue as nodes and white matter tracts as edges—in the case of structural networks—or voxels as nodes and correlations in signal over time as edges, in the case of functional connectivity networks. Second, brain networks also vary along a topological scale that goes from local (e.g., networks within a brain region) to global (e.g., networks across the whole brain). Finally, brain networks vary along a temporal scale, ranging from very fast (e.g., sub-second neural processes) to the very slow (e.g., life-span or evolutionary changes). Networks whose topological features vary along several scales are known as multi-scale. Therefore, brains can be thought of as multi-scale networks (Betz et al., 2016; De Brigard, 2017).

Different research questions, and their inherent practical limitations, also influence the way in which network models are constructed. We may, for instance, want to construct a brain model in which each individual neuron is represented as a node, with the edges between nodes representing synapses. Unfortunately, while this has been done successfully in the significantly less complex organism *C. Elegans* (Sporns and Kötter, 2004; Towson et al., 2013), it is not currently possible to image, record, or computationally analyze the tens of billions of neurons in the human brain, especially when neurons often have thousands of synapses (Drachman, 2005). Current neuroimaging technology limits functional and structural brain network analyses to nodes above the millimeter scale, meaning that many potentially interacting neurons, synapses, and other structures will be represented as an individual node in human brain networks. The lack of a clear, obvious choice of what should represent a node in a functional brain network has resulted in the analysis of brain networks across a wide range of scales, ranging from 70-node (Wang et al., 2009) to 140,000-node whole brain networks (Eguíluz et al., 2005), using a variety of parcellation schemes dependent on wide-ranging definitional criteria (Stanley et al. 2013). The boundaries representing reasonable functional units (i.e., nodes) in the brain for investigating a particular

phenomenon of interest need not line up with the surfaces of structures or other commonsense loci of demarcation, and the ‘best’ way to define nodes (size, brain-region, etc.) often depends on a researcher’s question. Furthermore, it is possible that these different levels of granularity provide network descriptions that are distinct, yet complementary, when predicting cognitive phenomena. For example, the particular firing patterns of neurons exclusively within the hippocampus support memory encoding and retrieval (Battaglia et al., 2011), and the increased topological centrality of the hippocampus—modeled as a single node in the whole brain network—also supports memory retrieval (Geib et al., 2017a).

Finally, the data from which topological models of brain networks are built also varies as a function of the technology employed to extract them. For instance, functional brain networks have been constructed using functional MRI (fMRI) (Achard & Bullmore 2007, Achard et al. 2006, Eguíluz et al. 2005, Geib et al. 2017a, 2017b, Liu et al. 2008, Salvador et al. 2005a, van den Heuvel et al. 2008), electroencephalography (EEG) (Micheloyannis et al. 2006, Stam et al. 2007), magnetoencephalography (MEG) data (Bassett et al. 2006, Deuker et al. 2009, Stam 2004), and ECoG (Betz et al. 2019). Structural brain graphs have been constructed from diffusion tensor imaging (DTI) or diffusion spectrum imaging (DSI) (Gong et al. 2008, Hagmann et al. 2008), as well as from conventional MRI data (Bassett et al. 2008, He et al. 2007).

Importantly, as in the case of the network models discussed above (section 3.1), recent studies suggest that the topological properties of network models in neuroscience offer extraordinary predictive value. Consider, for instance, research on brain disease. A recent study by Khazaee et al. (2015) combined network analyses of fMRI data with advanced machine learning techniques to investigate brain network differences between patients with Alzheimer’s disease (AD) and healthy, age-matched controls (see also, Khazaee et al., 2017). Alzheimer’s disease is a

progressive neurodegenerative disease that is accompanied by severe decline in cognitive functioning (in memory in particular; Albert et al., 2011). Graph theoretic metrics were obtained from each participant's brain network, and machine learning was used to explore the ability for graph metrics to help in the diagnosis of AD. They applied their method to resting-state fMRI data of twenty patients with AD and twenty age- and gender-matched healthy subjects. The graph measures were computed and then used as the discriminating features in the model. Extracted network-based features were fed to different feature selection algorithms to choose the most significant features. Using a set of graph metrics computed for diverse nodes (brain regions) in the network, the researchers were able to identify patients with AD relative to healthy controls with perfect accuracy (i.e., 100% correctly). So, if a new case were presented to the researchers, they would presumably be able to accurately predict whether that individual had AD based upon a set of graph theoretic metrics obtained from that individual's fMRI data. Results of this study suggest that graph theoretic metrics obtained from functional brain networks can efficiently and effectively assist in the diagnosis of AD. It may be that early diagnosis (before the onset of behavioral symptoms) is also possible by this method, whether or not we have a full mechanistic account explaining what occurs in the brain in AD.

A subsequent study conducted by Hojjati and colleagues (2017) went a step beyond Khazee et al. (2015). Specifically, Hojjati et al. (2017) used similar graph theoretic metrics obtained from brain networks constructed from resting-state fMRI in conjunction with machine learning algorithms to predict which individuals would progress from Mild Cognitive Impairment (MCI) to AD and which individuals would not progress from MCI to AD. MCI is a transitional stage between normal age-related cognitive decline and actual AD. The researchers were able to predict with greater than 90% accuracy which individuals would progress from MCI to AD and which

would not. The ability to accurately predict which individuals are likely to progress to AD offers physicians useful information to better tailor prevention and treatment programs on an individual basis. Additionally, it would be of significant use to family members of when planning for future care. (See delEtoile and Adeli (2017) for a useful recent review of similar research.)

Consider now the case of epilepsy, one of the most common neurological conditions. Epilepsy is characterized by the tendency toward recurrent, unprovoked seizures (Stam, 2014). Treatment for certain severe, drug-resistant cases of epilepsy sometimes involves anterior temporal lobectomy. Using graph theoretic metrics obtained from resting-state brain networks in conjunction with machine learning algorithms, He, Doucet, Pustina, and colleagues (2017) were able to predict surgical outcomes from patients who underwent anterior temporal lobectomy. More specifically, the researchers used graph theoretic measures of centrality during rest *prior* to the lobectomy to predict with high accuracy whether participants would be seizure-free a full year later. This research provides a useful potential biomarker for surgical outcomes, with the potential to usefully guide the decision-making of physicians in future cases by determining which individuals would be most likely to benefit from surgery. As in the examples above, the utility of graph theoretic metrics in predictive surgical outcomes is extremely valuable, whether or not we have a full mechanistic account explaining what occurs in the brain in epilepsy.

Cases of MCI, AD, and epilepsy offer salient examples of the predictive and intervention-guiding value of graph theoretic metrics obtained from brain networks. By “intervention-guiding,” we refer primarily to the possibility of identifying sub-populations which are at greater risk for certain health problems, on which clinicians may focus their treatment (though there are other ways in which brain-based graph theoretic metrics can guide interventions as well). Similar methods have also been used to predict with incredible accuracy which individuals have common

clinical disorders, such as major depressive disorder (Sacchet et al., 2015; Gong & He, 2015) and attention deficit hyperactivity disorder (Colby, Rudie, Brown, et al., 2012). In these cases, graph theoretic metrics obtained from brain networks offer great predictive utility in diagnosis. This, in turn, has the potential to aid in optimal treatment and intervention using other established techniques. Taking all of this together, the findings just reviewed clearly indicate that the topological properties of network models in neuroscience offer extraordinary predictive value and useful information for treatment and intervention, independent of their possible mechanistic interpretations.

4. Conclusion

Science is undoubtedly in the business of offering explanations about natural phenomena. But it is also in the business of offering predictions and strategies to intervene and manipulate reality. Most of the research in contemporary philosophy of science has focused on explanation, and the philosophy of neuroscience has followed suit. The overarching goal of the current paper has been to shed some light on the oft-neglected issue of prediction in neuroscience. We do so through the lens of network models and their topological properties.

While current philosophers of neuroscience disagree as to whether or not network models are truly explanatory, or whether their explanatory power is based in mechanistic schemas (Klein 2012; Muldoon & Bassett, 2016; Craver, 2016), we focused instead on the fact that many network models have predictive value and offer strategies for manipulation and intervention even when no clear causal or mechanistic account of the phenomenon is available. As such, topological models in network neuroscience promise to enhance our epistemic status regarding the brain and its effects by way of informing many vital decisions. The ability to use topological properties from network

models to make predictions also may help us to improve patient outcomes. The progression from pre-MCI to Alzheimer's disease, for example, significantly impacts an individual's life and the lives of loved ones. Predictions offer clear value for addressing these issues, even when a full-fledged mechanistic explanation of the neurological conditions isn't readily available. Cognitive neuroscience may not yet be able to give detailed instructions through mechanistic explanations for manipulating mental states; accurate predictions, on the other hand, do offer a clearer way forward for many applied problems.

References

- Achard, S., & Bullmore, E. (2007). Efficiency and cost of economical brain functional networks. *PLoS computational biology*, 3(2), e17.
- Albert, M. S., DeKosky, S. T., Dickson, D., Dubois, B., Feldman, H. H., Fox, N. C., ... & Snyder, P. J. (2011). The diagnosis of mild cognitive impairment due to Alzheimer's disease: Recommendations from the National Institute on Aging-Alzheimer's Association workgroups on diagnostic guidelines for Alzheimer's disease. *Alzheimer's & dementia*, 7(3), 270-279.
- Barrett, J., and Stanford, P.K. (2006) "Prediction" in *The Philosophy of Science: An Encyclopedia*, ed. Sahotra Sarkar and Jessica Pfeifer. New York: Routledge.
- Bassett, D. S., Bullmore, E., Verchinski, B. A., Mattay, V. S., Weinberger, D. R., & Meyer-Lindenberg, A. (2008). Hierarchical organization of human cortical networks in health and schizophrenia. *Journal of Neuroscience*, 28(37), 9239-9248.
- Bassett, D. S., Meyer-Lindenberg, A., Achard, S., Duke, T., & Bullmore, E. (2006). Adaptive reconfiguration of fractal small-world human brain functional networks. *Proceedings of the National Academy of Sciences*, 103(51), 19518-19523.
- Battaglia, F. P., Benchenane, K., Sirota, A., Pennartz, C. M., & Wiener, S. I. (2011). The hippocampus: hub of brain network communication for memory. *Trends in cognitive sciences*, 15(7), 310-318.
- Bechtel, William & Abrahamsen, Adele (2005). Explanation: a mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences* 36 (2):421-441.
- Betzl, R. F., & Bassett, D. S. (2017). Multi-scale brain networks. *Neuroimage*, 160, 73-83.
- Betzl, R. F., Medaglia, J. D., Kahn, A. E., Soffer, J., Schonhaut, D. R., & Bassett, D. S. (2019). Structural, geometric and genetic factors predict interregional brain connectivity patterns probed by electrocorticography. *Nature biomedical engineering*, 1.

- Bromberger, Sylvain (1966). Questions. *Journal of Philosophy* 63 (20):597-606.
- Butts, Carter T. (2009). Revisiting the foundations of network analysis. *Science* 325 (5939): 414-416.
- Colby, J. B., Rudie, J. D., Brown, J. A., Douglas, P. K., Cohen, M. S., & Shehzad, Z. (2012). Insights into multimodal imaging classification of ADHD. *Frontiers in systems neuroscience*, 6, 59.
- Collingridge, G. L., Kehl, S. J., & McLennan, H. T. (1983). Excitatory amino acids in synaptic transmission in the Schaffer collateral-commissural pathway of the rat hippocampus. *The Journal of physiology*, 334(1), 33-46.
- Craver, C. F., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*. University of Chicago Press.
- Craver, Carl F. & Povich, Mark (2017). The directionality of distinctively mathematical explanations. *Studies in History and Philosophy of Science Part A* 63:31-38.
- Craver, Carl F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press, Clarendon Press.
- Craver, Carl F. (2014). The Ontic Account of Scientific Explanation. In Marie I. Kaiser, Oliver R. Scholz, Daniel Plenge & Andreas Hüttemann (eds.), *Explanation in the Special Sciences: The Case of Biology and History*. Springer Verlag. pp. 27-52.
- Craver, Carl F. (2016). The Explanatory Power of Network Models. *Philosophy of Science* 83 (5):698-709.
- Craver, Carl and Tabery, James. (2015). "Mechanisms in Science", *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), Edward N. Zalta (ed.), forthcoming URL = <https://plato.stanford.edu/archives/sum2019/entries/science-mechanisms/>
- Darden, Lindley (2002). Rethinking Mechanistic Explanation. *Philosophy of Science* 69 (S3):342-353.
- De Brigard, F. (2017). Cognitive systems and the changing brain. *Philosophical Explorations*. 20(2): 224-241.
- delEtoile, J., & Adeli, H. (2017). Graph theory and brain connectivity in Alzheimer's disease. *The Neuroscientist*, 23(6), 616-626
- Deuker, L., Bullmore, E. T., Smith, M., Christensen, S., Nathan, P. J., Rockstroh, B., & Bassett, D. S. (2009). Reproducibility of graph metrics of human brain functional networks. *Neuroimage*, 47(4), 1460-1468.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numer. Math.* 1, 269–271.

- Douglas, Heather & Magnus, P. D. (2013). State of the Field: Why novel prediction matters. *Studies in History and Philosophy of Science Part A* 44 (4):580-589.
- Douglas, Heather E. (2009). Reintroducing prediction to explanation. *Philosophy of Science* 76 (4):444-463.
- Drachman, D. Do we have brain to spare? *Neurology* Jun 2005, 64 (12).
- Eguíluz, V. M., Chialvo, D. R., Cecchi, G. A., Baliki, M., and Apkarian, A. V. (2005). Scale-free brain functional networks. *Phys. Rev. Lett.* 94:018102.
- Geib BR, Stanley ML, Wing EA, Laurienti PJ, Cabeza R. Hippocampal Contributions to the Large-Scale Episodic Memory Network Predict Vivid Visual Memories. *Cereb Cortex*. 2017;27(1):680–693.
- Geib, B. R., Stanley, M. L., Dennis, N. A., Woldorff, M. G., & Cabeza, R. (2017). From hippocampus to whole-brain: The role of integrative processing in episodic memory retrieval. *Human brain mapping*, 38(4), 2242-2259.
- Glennan, Stuart (2002). Rethinking Mechanistic Explanation. *Philosophy of Science* 69 (S3):S342-S353.
- Gong, G., He, Y., Concha, L., Lebel, C., Gross, D. W., Evans, A. C., & Beaulieu, C. (2008). Mapping anatomical connectivity patterns of human cerebral cortex using in vivo diffusion tensor imaging tractography. *Cerebral cortex*, 19(3), 524-536.
- Gong, Q., & He, Y. (2015). Depression, neuroimaging and connectomics: a selective overview. *Biological psychiatry*, 77(3), 223-235.
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., & Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS biology*, 6(7), e159.
- Hanson, N. R. (1959). Copenhagen interpretation of quantum theory. *American Journal of Physics*, 27(1), 1-15.
- He, X., Doucet, G. E., Pustina, D., Sperling, M. R., Sharan, A. D., & Tracy, J. I. (2017). Presurgical thalamic “hubness” predicts surgical outcome in temporal lobe epilepsy. *Neurology*, 88(24), 2285-2293.
- He, Y., Z. J. Chen, and A. C. Evans. 2007. “Small-World Anatomical Networks in the Human Brain Revealed by Cortical Thickness from MRI.” *Cerebral Cortex* 17:2407-2419.
- Helmer, O., & Rescher, N. (1959). On the epistemology of the inexact sciences. *Management science*, 6(1), 25-52.
- Hempel, C. G., & Oppenheim, P. (1948). Studies in the Logic of Explanation. *Philosophy of science*, 15(2), 135-175.
- Hempel, Carl (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. The Free Press.

- Hofstadter, Albert (1951), "Explanation and Necessity", *Philosophy and Phenomenological Research* 11: 339–347.
- Hojjati, S. H., Ebrahimzadeh, A., Khazae, A., Babajani-Feremi, A., & Alzheimer's Disease Neuroimaging Initiative. (2017). Predicting conversion from MCI to AD using resting-state fMRI, graph theoretical approach and SVM. *Journal of neuroscience methods*, 282, 69-80.
- Hume, D. *An Enquiry Concerning Human Understanding*. Glasgow, 1748.
- Huneman, Philippe (2010). Topological explanations and robustness in biological sciences. *Synthese* 177 (2):213-245.
- Jolly, A. M., & Wylie, J. L. (2002). Gonorrhoea and chlamydia core groups and sexual networks in Manitoba. *Sexually transmitted infections*, 78(suppl 1), i145-i151.
- Kaplan, David Michael & Craver, Carl F. (2011). The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective. *Philosophy of Science* 78 (4):601-627.
- Khazae, A., Ebrahimzadeh, A., & Babajani-Feremi, A. (2015). Identifying patients with Alzheimer's disease using resting-state fMRI and graph theory. *Clinical Neurophysiology*, 126(11), 2132-2141.
- Khazae, A., Ebrahimzadeh, A., & Babajani-Feremi, A. (2016). Application of advanced machine learning methods on resting-state fMRI network for identification of mild cognitive impairment and Alzheimer's disease. *Brain imaging and behavior*, 10(3), 799-817.
- Kitcher, Philip (1989). Explanatory unification and the causal structure of the world. In Philip Kitcher & Wesley Salmon (eds.), *Scientific Explanation*. Minneapolis: University of Minnesota Press. pp. 410-505.
- Klein, Colin (2012). Cognitive Ontology and Region- versus Network-Oriented Analyses. *Philosophy of Science* 79 (5):952-960.
- Lakatos, I., & Musgrave, A. (Eds.). (1970). *Criticism and the growth of knowledge: Volume 4: Proceedings of the International Colloquium in the Philosophy of Science*, London, 1965. Cambridge University Press.
- Lange, Marc (2013). What Makes a Scientific Explanation Distinctively Mathematical? *British Journal for the Philosophy of Science* 64 (3):485-511.
- Lange, Marc (2016). *Because Without Cause: Non-Causal Explanations in Science and Mathematics*. Oxford University Press USA.
- Liljeros, F., Edling, C. R., Amaral, L. A. N., Stanley, H. E., & Åberg, Y. (2001). The web of human sexual contacts. *Nature*, 411(6840), 907.
- Longino, Helen (2002). *The Fate of Knowledge*. Princeton University Press.

- Machamer, Peter K. ; Darden, Lindley & Craver, Carl F. (2000). Thinking about mechanisms. *Philosophy of Science* 67 (1):1-25.
- Micheliyannis, S., Pachou, E., Stam, C. J., Vourkas, M., Erimaki, S., & Tsirka, V. (2006). Using graph theoretical analysis of multi channel EEG to evaluate the neural efficiency hypothesis. *Neuroscience letters*, 402(3), 273-277.
- Mill, J. A System of Logic, Ratiocinative and Inductive. London, 1843.
- Monge, Z. A., Geib, B. R., Siciliano, R. E., Packard, L. E., Tallman, C. W., & Madden, D. J. (2017). Functional modular architecture underlying attentional control in aging. *NeuroImage*, 155, 257-270.
- Muldoon, S. F., & Bassett, D. S. (2016). Network and multilayer network approaches to understanding human brain dynamics. *Philosophy of Science*, 83(5), 710-720.
- Nagel, E. (1961). *The Structure of Science* Harcourt, Brace & World, New York.
- Newman, M. (2010). *Networks: an introduction*. Oxford university press.
- Park, P., Volianskis, A., Sanderson, T. M., Bortolotto, Z. A., Jane, D. E., Zhuo, M., ... & Collingridge, G. L. (2014). NMDA receptor-dependent long-term potentiation comprises a family of temporally overlapping forms of synaptic plasticity that are induced by different patterns of stimulation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1633), 20130131.
- Popper, K. (1963). *Conjectures and refutations: The growth of scientific knowledge*.
- Potterat, J. J., Muth, S. Q., Rothenberg, R. B., Zimmerman-Rogers, H., Green, D. L., Taylor, J. E., ... & White, H. A. (2002). Sexual network structure as an indicator of epidemic phase. *Sexually transmitted infections*, 78(suppl 1), i152-i158.
- Rothenberg, R. B., Potterat, J. J., Woodhouse, D. E., Muth, S. Q., Darrow, W. W., & Klovdahl, A. S. (1998). Social network dynamics and HIV transmission. *Aids*, 12(12), 1529-1536.
- S. Achard, R. Salvador, B. Whitcher, J. Suckling, E. Bullmore (2006). A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs *Journal of Neuroscience* 26, 63-72.
- Sacchet, M. D., Prasad, G., Foland-Ross, L. C., Thompson, P. M., & Gotlib, I. H. (2015). Support vector machine classification of major depressive disorder using diffusion-weighted neuroimaging and graph theory. *Frontiers in psychiatry*, 6, 21.
- Salmon, W. C. (1978). *Unfinished business: The problem of induction*. *Philosophical Studies*, 33(1), 1-19.
- Salmon, Wesley (1971). *Statistical Explanation & Statistical Relevance*. University of Pittsburgh Press.

- Salmon, Wesley (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.
- Salvador, R., Suckling, J., Schwarzbauer, C., & Bullmore, E. (2005). Undirected graphs of frequency-dependent functional connectivity in whole brain networks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457), 937-946.
- Sarigöl, E., Pfitzner, R., Scholtes, I., Garas, A., & Schweitzer, F. (2014). Predicting scientific success based on coauthorship networks. *EPJ Data Science*, 3(1), 9.
- Scheffler, I. (1957). Explanation, prediction, and abstraction. *The British Journal for the Philosophy of Science*, 7(28), 293-309.
- Schindler, S. (2018). *Theoretical Virtues in Science: Uncovering Reality Through Theory*. Cambridge University Press.
- Scriven, M. (1959). Explanation and prediction in evolutionary theory. *Science*, 130(3374), 477-482.
- Sporns, O. (2011). The human connectome: a complex network. *Annals of the New York Academy of Sciences*, 1224(1), 109-125.
- Sporns, O., & Kötter, R. (2004). Motifs in brain networks. *PLoS biology*, 2(11), e369.
- Stam, C. J. (2004). Functional connectivity patterns of human magnetoencephalographic recordings: a 'small-world' network?. *Neuroscience letters*, 355(1-2), 25-28.
- Stam, C. J. (2014). Modern network science of neurological disorders. *Nature Reviews Neuroscience*, 15(10), 683.
- Stam, C. J., Nolte, G., & Daffertshofer, A. (2007). Phase lag index: assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources. *Human brain mapping*, 28(11), 1178-1193.
- Stanley, M. L., Moussa, M. N., Paolini, B., Lyday, R. G., Burdette, J. H., & Laurienti, P. J. (2013). Defining nodes in complex brain networks. *Frontiers in computational neuroscience*, 7, 169.
- Towlson, E. K., Vértes, P. E., Ahnert, S. E., Schafer, W. R., & Bullmore, E. T. (2013). The rich club of the *C. elegans* neuronal connectome. *Journal of Neuroscience*, 33(15), 6380-6387.
- Wang, J., Wang, L., Zang, Y., Yang, H., Tang, H., Gong, Q., ... & He, Y. (2009). Parcellation-dependent small-world brain functional networks: A resting-state fMRI study. *Human brain mapping*, 30(5), 1511-1523.
- Wang, P., Hunter, T., Bayen, A. M., Schechtner, K., & González, M. C. (2012). Understanding road usage patterns in urban areas. *Scientific reports*, 2, 1001.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684), 440.

Whewell, W. *The Philosophy of the Inductive Science*. London, 1840.

Woodhouse, D. E., Rothenberg, R. B., Potterat, J. J., Darrow, W. W., Muth, S. Q., Klovdahl, A. S., ... & Reynolds, J. U. (1994). Mapping a social network of heterosexuals at high risk for HIV infection. *Aids*, 8(9), 1331-1336.

Woodward, James (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.

Wylie, J. L., & Jolly, A. (2001). Patterns of chlamydia and gonorrhea infection in sexual networks in Manitoba, Canada. *Sexually transmitted diseases*, 28(1), 14-24.

van den Heuvel, M. P., Stam, C. J., Boersma, M., & Pol, H. H. (2008). Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *Neuroimage*, 43(3), 528-539.

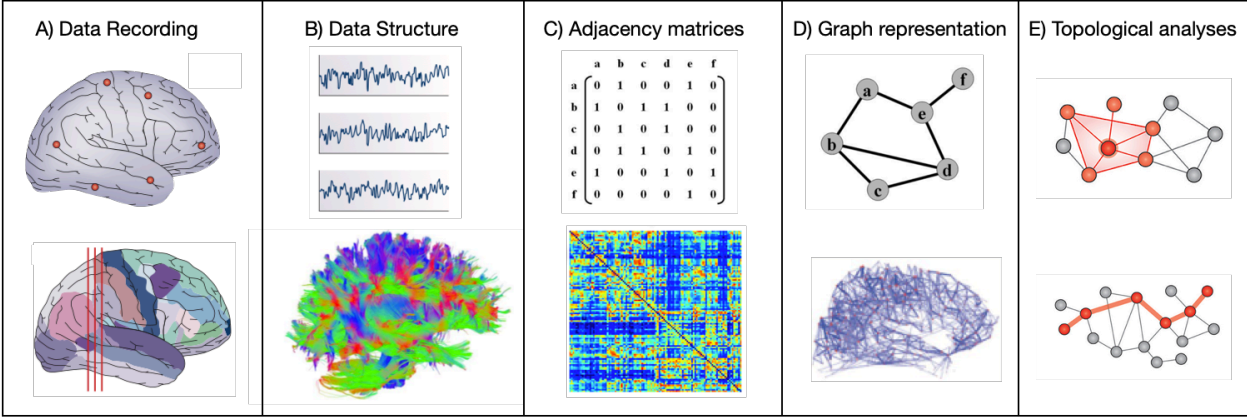


Figure 1: Schematic representation of topological analyses employed in network neuroscience. A) Data acquisition includes several methods, such as functional and structural MRI. B) Depending on the nature of the data, their structure may vary—for example, time series in fMRI or diffusivity measures in diffusion tensor imaging (DTI). C) Data is arranged in adjacency matrices, representing nodes and edges. D) Data can also be represented in graphs, with lines depicting edges connecting nodes. E) Topological analyses are then conducted to identify topological properties (e.g., clustering coefficient, shortest path).