How network models contribute to science*

Charles Rathkopf†

May 4, 2023

Abstract

Network models represent complex empirical systems by means of graphs, composed of nothing more than nodes and edges, which themselves lack internal structure. Graphs can be constructed from empirical data, on the basis of simple rules that do not require much theoretical insight into the target system. Moreover, network models do not compress the data they represent. In typical network models, the mapping from data to graph is invertible. For these reasons, network modeling can seem more like a trendy format for data summary than the powerful modeling framework it is sometimes claimed to be. This chapter shows that, despite the apparent simplicity of the graph construction process, network modeling is indeed an inferentially powerful modeling framework that enables novel forms of discovery, prediction, and explanation. Thereafter, the chapter explores the fact that network properties seem to crop up repeatedly, across a wide variety of empirical domains. How surprising is this fact? Does it occur because the relevant empirical domains are intrinsically network-like, or for more pragmatic reasons to do with the way we are disposed to reason about them?[1]

*This is a draft for a chapter that will appear in the forthcoming *Routledge Handbook on the Philosophy of Scientific Modeling*, edited by Tarja Kuuttila, Natalia Carrillo and Rami Koskinen.

†Institute for Brain and Behavior, Jülich Research Center

[1]Many thanks to Kareem Khalifa and Daniel Kostić for their insightful and detailed comments on an earlier draft of this chapter.

Contents

1 Introduction

Network models are surprisingly easy to construct. There are at least two reasons for this. First, the construction process typically requires rather little theoretical guidance. Network models represent empirical objects as graph nodes, and relations between objects as graph edges. The objects in question are usually discrete and, at least when viewed within the relevant scientific context, easily individuated. Examples include people, corporations, power plants, academic publications, tree species, and protein types. Moreover, the rules that govern the mapping between the graph and the data are straightforward. On the mathematical side, the nodes and edges are simple mathematical objects, mostly devoid of internal structure. On the empirical side, data sets typically include only a few types of object. Often, there is just one. Data sets also typically include only a few types of relation between objects. Again, there is often just one.

Another reason that network models are easy to construct is that the construction of the model does not involve any attempt to capture patterns hidden in the data, and therefore does not involve data compression. Typically, once the data set has been cleaned, every object and relation gets represented. In this respect, network models are radically different from the compact, closed-form equations that have historically been viewed as the standard-bearer of scientific representation. These two observations about the construction of network models might lead one to think that such models must be superficial. They may look more like a trendy format for data summary than an innovative modeling strategy, capable of supporting profound scientific insight.

This last thought is mistaken, and it is the burden of this chapter to show why. The central thesis is that network models are capable of supporting profound insight into a surprisingly diverse range of phenomena. This view is supported with three case studies, selected to show that network models play an indispensable role in prediction, discovery, and explanation. Before getting into the case studies, the chapter provides a brief overview of the history that led to modern network modeling, and introduces a few of the most common mathematical concepts. The chapter concludes with a discussion of the fact that network models are applicable to an enormously diverse range of empirical phenomena. This point has been emphasized by advocates of network modeling, and has been used, at least occasionally, to support grandiose claims about the role of network models within the larger scientific enterprise. The account

developed here is comparatively tempered, but not dismissive. It suggests that the trans-domain applicability of network models may sometimes offer us new and currently under-appreciated opportunities for scientific unification.

2 The emergence of network science

Modern network modeling emerged from the confluence of two historical research traditions, one in pure mathematics, and the other in social science. On the mathematical side, the paper "On the Evolution of Random Graphs," by Paul Erdős and Alfred Rényi, introduced modern techniques for studying large graphs analytically (Erdős et al., 1960). In that paper, Erdős and Rényi imagine a large set of nodes, along with all of the possible graphs that can be constructed from that set, where a graph is simply a configuration of edges that connect the nodes. They prove that the set of all possible graphs with n nodes has several interesting properties. For example, they prove that, as n tends to infinity, the size of the largest connected subgraph follows a Poisson distribution. On the social science side, Mark Granovetter's paper "The Strength of Weak Ties," (1973) showed how quantitative properties of social graphs could provide sociological insight. It showed, on the basis of both empirical data and hypothetical reasoning, that weak social ties play an outsized role in generating macroscopic sociological phenomena. The crux of his reasoning is that, unlike friends, mere acquaintances move in social circles different from one's own. As a result, acquaintances provide links to social groups that are both valuable and otherwise inaccessible. When it comes to finding a job, for example, acquaintances tend to be more advantageous than friends.

Modern network modeling can be viewed as a synthesis of the two research traditions that emerged, respectively, from these two papers. To see this, it helps to note some of the most salient differences between the two traditions. First, the sociological tradition used networks as a means of representing empirical data, while the mathematical tradition did not. Second, the sociological tradition focused primarily on networks with complex, non-random structure, while the mathematical tradition focused primarily on either random or lattice-like networks, both of which are more susceptible to mathematical analysis than complex graphs. Third, the mathematical tradition focused on networks that were large or infinite, while the sociological tradition, especially early on, focused on networks that were quite small (Granovetter's paper, for example, was

based on data from just 54 people.)

Modern network analysis blends these two traditions together. It studies graphs that are large, but based on empirical data, and therefore finite. Moreover, most empirical networks are neither perfectly ordered nor perfectly random, and are, therefore, difficult to study using purely analytical techniques. To understand how large and complex networks behave under different parameter settings, computer simulations are required. Today, a large part of what is sometimes called network science involves the discovery of algorithms that can compute interesting properties of large complex graphs. Because many of these properties are probabilistic, one typically needs to study a whole ensemble of graphs, which is computationally demanding. It is therefore no accident that modern network modeling emerged only after the rise of cheap computing power.

The first papers to undertake this synthesis, and which, in so doing, launched the modern era of network modeling, appeared between 1998 and 2000. The two most frequently cited are Watts and Strogatz (1998), in which the small-world model was introduced, and Barabási and Albert (1999) in which the so-called *BA preferential attachment model* was introduced. In 2005, a crucial and underappreciated historical landmark in the development of network modeling was the release of an open-source Python library called *NetworkX*. *NetworkX* made it possible to convert lists and matrices into networks, compute common network properties, and visualize networks graphically (Hagberg et al., 2008). Once that software was released, scientists in many other fields began to use network analysis on their data, which in turn drove the development of new software for network analysis.

3 Common graph-theoretic concepts

Network models are based on the mathematics of graphs. A graph consists of a set of nodes and a set of edges, where an edge is just a two-element set of nodes (Trudeau, 1976). Graphs are typically visualized as points and lines on a plane, but for the purposes of computing, a graph is represented as a type of matrix. The most common type of matrix for representing a graph is an adjacency matrix. In an adjacency matrix, both axes are defined by the set of nodes. If a direct connection exists between two nodes, their intersection is marked with a 1, indicating the presence of an edge; otherwise, it is marked

with a 0.

$$A = \begin{matrix} 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{matrix}$$

The same information can also be represented as an incidence matrix, which is a matrix defined by the nodes on one axis and the edges on the other. Incidence matrices are used less often than adjacency matrices but are favored for the representation of networks that are both large and sparse, because, in those cases, incidence matrices can be represented more compactly than adjacency matrices. (Where n is the number of nodes, and m is the number of edges, a sparse graph is one in which $n >> m$. An adjacency matrix has dimensions $n \times n$, which makes it larger than the corresponding incidence matrix with dimensions $n \times m$.)

The introduction sketched an intuitively appealing inference from the claim that networks are easy to construct to claim that they are superficial, or inferentially weak. One way to resist this cynical inference is to emphasize the distinction between constructing a network model, and using it productively, once constructed. To use a network model to make inferences about the target system, you have to (i) choose the appropriate network properties to measure, and (ii) interpret the theoretical significance of those measurements. While the second of these two steps certainly does demand domain-specific empirical knowledge, it is less clear what sort of knowledge is required for the first step. A natural assumption would be that you need domain-specific knowledge of the target system in order to know which properties of the associated network representation are worth measuring. However, virtually all popular accounts of network science defend (or at least assert) the idea that, regardless of which empirical domain you are working in, the same set of network properties end up being important. This claim is both fascinating and puzzling, and it will be discussed it more detail below. Here, I only want to mention it as justification for the suggestion that one can understand a surprisingly large swath of network science modeling ideas on the basis of a rather small number of graph-theoretical

concepts.

The following short list of network science concepts captures some of the most basic and most frequently used concepts. The characterizations are not rigorous definitions, but provide enough information to render the subsequent discussion accessible.

1. Node degree: the number of nodes with which a given node is connected.
2. Path length: the length of the shortest path that can be traversed between two nodes.
3. Clustering coefficient: a measure of how likely it is that there is an edge between nodes A and C, given that there is an edge between A and B, and another between B and C.
4. Small-worldness: the ratio of clustering coefficient to average path length.
5. Scale-free network: a network whose node degree distribution follows a power law.
6. Random graph: a graph in which the edges between nodes is determined by some random selection process.
7. Regular graph: a lattice-like graph in which every node has the same degree.

Although each of these properties is a property of a graph, not all of them can be found in books on graph theory per se. For an overview of network properties as they pertain to network modeling, readers are advised to consult one of the many textbook treatments of network modeling ideas. Newman (2005) does a particularly good job of balancing ease of exposition with mathematical rigor.

4 Reasoning with networks

Each of the concepts listed above is exploited by the reasoning in the case studies below. Before turning to those, one source of potential confusion must be addressed. The reasoning in each case study draws not only on a graph and a graph-to-data mapping, but also on additional modeling apparatus. When additional modeling apparatus is required, a defender of the view that network models are superficial might say that the case studies described here fail to

support the primary non-superficiality thesis, because the models in question are not pure network models. Their success, therefore, may have little to do with networks per se. Although one can find examples of pure network models in the relevant sense (such as models that account for why author citation networks have the distributions they do), I do not discuss them here, since such pure network models have played a relatively minor role in the advancement of scientific knowledge over the past 25 years. More dramatic progress has been made by combining graph theoretical representation with other kinds of modeling apparatus (such as, for example, a system of differential equations.) The focus here on hybrid models, as we might call them, might prompt an objection of the following sort. One could only gather support for the thesis defended here (that network modeling supports profound forms of scientific inference) if one could first work out, with respect to any given inference, how to distinguish cleanly between the insight contributed by the network model, and the insight contributed by the other modeling apparatus. This demand for a crisp criterion of model individuation is asking for too much. The claim that network modeling has made a substantive and distinctive contribution to science need not rely on specific criteria for counting network models. Adequate support for the claim can be provided simply by identifying network properties that are practically indispensable for novel forms of scientific inference. If at least some of those inferences can be described as profound, the central thesis defended in this chapter follows logically. By formulating the central thesis in terms of properties rather than models, it can be reconciled with a wider variety of views about the nature of scientific models. In particular, a given collection of network properties can be viewed either as constitutive of a pure network model, which, as a contingent matter, was used in conjunction with another model, or they can be viewed merely as a subset of elements within a larger, more multifaceted model. Both views are compatible with the central thesis of this chapter.

4.1 Discovery

One area in which network modeling has been used to make new discoveries is in molecular biology. In a landmark paper, Spirin and Mirny (2003) undertook a network-based analysis of an existing, open-source database of protein-protein interactions in yeast, which were themselves detected by well-established experimental methods. Protein-protein interactions are biochemical interactions

between proteins that allow them to function together as part of a molecular machine that accomplishes some cell function. Most mesoscale cell functions are carried out by a large family of proteins, not all of which engage in direct biochemical reactions with one another. In addition, many of the protein-protein interactions involved in any given mesoscale cell function simply have yet to be probed experimentally. For both reasons, there will often be proteins that play an important role in a given cell function, but which are not yet known to do so. Spirin and Mirny used network modeling to facilitate a new process for protein discovery that is radically more efficient than what was previously possible.

The network they constructed consisted of 3,992 nodes, each representing a protein type, and 6,500 edges, each representing a known protein-protein interaction. Their primary goal was to locate biologically significant clusters within this network. This is trickier than it sounds. Even the problem of identifying the single largest cluster in a graph is NP-hard, so developing efficient search algorithms is a non-trivial mathematical problem. Spirin and Mirny designed an algorithm to find the maximum of the function:

$$Q(m, n) = 2m/(n(n-1))$$

where m is the number of interactions between n nodes. Q characterizes the density of the cluster. The algorithm uses a Monte Carlo procedure that starts with a set of nodes, selected at random, and then replaces members of that set, re-computing Q for each new set until it converges. They then selected all clusters with a Q value high enough to make it statistically significant. (Statistical significance is evidence that supports a rejection of the null hypothesis, which is itself typically formulated as the claim that the observation in question appeared by chance. In this setting, the operational meaning of "appeared by chance" is that it appears in a graph which is itself a member of an ensemble of graphs that was generated by a random graph construction procedure.)

With that done, Spirin and Mirny worked out which cell function the cluster of proteins contributes to. In each complex, at least some proteins were already known, and their functions were annotated in the open-source database. They hypothesized that the other proteins in the cluster would contribute to the same function; an inference strategy known as guilt-by-association. This led to a suite of predictions about the functional role of proteins that were in the cluster, but not yet known to be involved in the cell function associated with that cluster.

The Spirin and Mirny paper counts as a significant contribution to scientific

discovery for two reasons. First, their predictions radically reduced the space of proteins to be tested experimentally, and thereby made it easier to choose experiments that were likely to have valuable results. The second reason is that, since it was first published in 2003, their predictions have been largely confirmed by experiment (Omranian et al., 2022). Moreover, the methods they developed for identifying protein complexes have been widely reused by other labs which have themselves made valuable discoveries with it.

Spirin and Mirny's work shows that network properties play a role in scientific discovery. There are good reasons to believe, furthermore, that the role they play is practically ineliminable. The space of possible protein-protein interactions is enormous. One could not practically perform the sort of experimental screening (such as two-hybrid screening) required to detect each possible interaction. In the absence of that brute force approach, one needs to make predictions about which proteins are likely to interact with each other from some set of known protein-protein interactions. Predictions of this sort can be divided into two classes: those that rely on theoretical knowledge of the proteins involved, and those that do not. If you go with the former class, you may get some predictive traction, but your predictions will be painstaking and slow. In the latter class, you make many predictions at scale. If you want to generate predictions at scale, it is necessary to represent the full suite of structural relations (interaction vs. no-interaction) that characterize the pre-theoretical domain. To construct an uncompressed data representation of a suite of objects and structural relations is to construct a network. Therefore, if you want to make a large class of predictions about biologically significant protein clusters, network representation is practically ineliminable.

4.2 Prediction

One can hardly write about the use of network models in 2022 and fail to discuss their use in modeling the Covid-19 pandemic. One of the puzzling facts about the early phase of the Covid-19 pandemic was that in many countries, after an initial wave in which the infection rate grew exponentially, it continued to grow linearly, even though, according to standard epidemiological models, the probability of sustained linear growth is effectively zero. The models in question, commonly known as susceptible-immune-recovered (SIR) models, predict either exponential growth or exponential decay whenever the reproduction number R deviates even slightly from 1. The reproduction number can be defined as the

expected number of secondary infections an infected person will cause, and it is rarely precisely equal to one.

The fact that steady linear growth was observed over long periods of time suggests, therefore, that the SIR models were missing something important about Covid-19 dynamics. Of course, epidemics are intrinsically difficult processes to predict. They are stochastic, they are influenced by many social and biological variables, and, at least in early stages, they are non-linear. Consequently, one cannot expect high precision predictions. Even when holding all parameter values constant, infection curves can look quite different from one run of the simulation to the next, and the total number of infected people can vary by a factor of two. Still, by casting the structure of the population as a network, it becomes possible to represent the critical degree of a population explicitly and use it to improve predictive traction.

Following the work of Pastor-Satoris and Vaspigani (2001), Thurner et al. (2020) simulated the SIR model in a network environment. Nodes represent people, and edges between people represent physical proximity sufficient for viral transmission. Each person is represented as having a particular number of contacts per day to whom they could theoretically transmit the virus. That number is called the node degree and it varies from person to person, following a distribution that is designed to mimic the contact structure of real human populations. The network was based on empirical data, but Thurner et al. constructed their network by algorithmic means. Algorithmic construction allowed them to vary the parameters of the network systematically, but also shouldered them with the burden of having to run model fitting tests to check how well the algorithmically generated model fits the empirical data. Of particular relevance are the facts that (i) during lockdown, the average degree D drops to near family size, and (ii) there will be some interactions between families, but these interactions will not take the form of giant hubs.

Intuitively, the higher the average degree, the more easily a virus will spread. In their simulations, Thurner et al. observed a qualitative shift in disease dynamics when D drops below a critical threshold. Above the threshold, the epidemic grows exponentially; below it, growth remains linear. This qualitative shift roughly captures the effect of lockdown policies during the summer of 2020. Once lockdown measures were in place, the contact structure of the population dropped to a value only slightly higher than the average family size. (The exact value of D at which this qualitative shift occurs is not a universal property of epidemics. It depends on the transmission rate of the virus, among

other factors.) Using this model, Thurner et al. managed to outperform extant predictions of the Covid-19 infection curves in both Austria and the USA. The choice of these two countries was significant because they differ so dramatically. Austria adopted strict lockdown polices early in the pandemic, while the USA introduced weaker lockdown policies later on. Despite these and many other differences between the two countries, Thurner et al. achieved this predictive improvement by choosing a value of D to fit empirical estimates of contact structure before and after lockdown measures were in place.

Crucially, this work deserves to be counted as a case of prediction, rather than model fitting, because the representative infection curves were captured without having to fit any other model parameters, all of which were chosen at the outset on the basis of measurement.

4.3 Explanation

At the neural level, an epileptic seizure is an episode of synchronized hyperexcitatory spiking activity. One of the puzzling things about epilepsy is that it is often caused by injuries that induce a substantial loss of neural connectivity. Intuitively, connectivity should facilitate hyperexcitability, since, when connectivity is high, there are more paths between the input activation and those neurons farthest from the input layer. In light of this intuition, a natural question is: why does a loss of connectivity lead to hyperexcitability? Answering this question is crucial to understanding why epileptic seizures occur.

The core of the answer to this question is that hyperexcitability is not due to the loss of connectivity itself, but to the new pattern of connectivity that emerges in the wake of that loss. In other words, the answer depends on a change in the topological structure of the network. This idea was first suggested by Percha et al. (2005). Although that paper suggested the correct topological answer to our question, it was based on a small simulation of 144 neurons. As a result, it was unclear whether the simulation could justifiably be interpreted as a guide to post-injury epilepsy in humans. Dyhrfjeld-Johnson et al. (2007) published a radically more extensive model that confirmed and expanded the initial results.

The Dyhrfjeld-Johnson paper focused on the dentate gyrus, a part of the temporal cortex known both to be involved in the generation of seizures, and to be unusually sensitive to injury. Dyhrfjeld-Johnson et al. built a nearly full-scale model of the dentate gyrus of the rat brain, with 50,000 neurons and over one billion connections. On top of each node in the graph, they built a

compartmental model neuron, which captures the spiking behavior of neurons as a response to electrical input.

The dependent variable in this study is the degree of hyperexcitability in the network, which is defined as a function of (i) the proportion of the neurons in the network that get activated after a particular input, (ii) the length of the interval between initial activation and the activation of the last neuron to be activated, and (iii) the duration of the whole network activation, once achieved. The primary independent variable is the degree of small-worldness of the network topology.

How does getting hit on the head lead to an increase in the small-worldness of your dentate gyrus? The answer to this question is incomplete, but interesting. Some cell types are more susceptible to injury than others. Hilar cells are both particularly susceptible to injury and highly connected. So, when these cells die out after injury, connectivity drops drastically. Soon afterward, granule cells (GCs) begin to form new connections at a greater rate than usual. Moreover, they form excitatory recurrent connections to other GCs, which, in the healthy brain is very rare (about 0.05 percent of all possible GC-GC pairs share a synaptic connection.) Some of these GCs connect at rates that are extreme, in comparison with the expected level of connectivity, and thereby become network hubs. These facts are supported directly by physiological observation, but are not themselves well-understood. So let us set aside the question of why GCs sprout new recurrent connections after the injury. Instead, we want to focus on what the new topology is like, and how that topology influences hyperexcitability in the spiking neuron model.

Two topological characteristics stand out. First, although the healthy dentate gyrus is already estimated to have the small-world property to some degree, the post-injury dentate gyrus has a very high degree of small-worldness, with low average path length, and, nevertheless, high local connectivity. In this case, the path length dropped to an extremely low value: on average, two neurons are connected by less than three edges, even though there are 50,000 neurons, and despite the fact that connectivity in the post-injury model is only 4.7 percent of what it was previously. The topology leads to hyperexcitability gradually, until reaching a threshold. Until that threshold is reached, you get monotonic increases in hyperexcitability with small-worldness. Since other properties of the network are held fixed, only the topological property can explain the hyperexcitability. Robustness analysis reveals that the effect is stable. Under a parameter sweep, the link between small-worldness and hyperexcitability re-

mains largely unchanged.

The Dyhrfjeld-Johnson model answers the following why-question: why do injuries to the temporal lobe increase susceptibility to epileptic seizures, even though they trigger substantial loss of connectivity? The answer is that (i) injury causes an increase in the small-worldness of the topology of the dentate gyrus, (ii) the increase in small-worldness promotes hyperexcitability, even in the absence of other physiological changes (iii) hyperexcitability is a state of increased susceptibility to epileptic seizure. For a full defense of the view that this case study deserves to be counted as an explanation, rather than as a mere description, it would be necessary to lay out one or more philosophical theories of explanation. There is no room for that project here, but one can find accounts of explanation congenial to the view in Rathkopf (2018), Kostic (2018), and Kostic and Khalifa (2021).

5 Network science is not superficial

In the introduction, two characteristics of the process of constructing a graph from empirical data were described. The first was that, once you have an appropriately structured data set, constructing a graph from the data does not require additional theoretical knowledge of the empirical domain. The second was that constructing a network model does not involve data compression. When you construct a network model, all the data gets recapitulated in graph-theoretic form. These two characteristics can give the misleading impression that network modeling is a superficial enterprise, in the sense that network models are likely to facilitate only rather shallow empirical inferences. The case studies above were selected to illustrate that this impression is incorrect. Network models are practically indispensable for certain kinds of scientific inference, some of which are profound. Here I will attempt to make the case for this claim more systematically.

Let us start with the subsidiary claim that network models are practically indispensable. The term "practically indispensable" refers to a weak form of necessity: it is not logically impossible to draw the conclusions at issue by means of some other modeling strategy, or by means of some other representational apparatus. Rather, the claim is that, given the contingent constraints involved in real scientific practice, a network model of some kind is the only viable option. In the first case study, Spirin and Mirny used their network model to draw a

host of conclusions about the functional contributions of various proteins. While it is logically possible that someone might have reached the same conclusions experimentally, the number of experiments required would be in the tens of millions. In the second case study, Thurner et al. used a network representation of the contact structure of a population under lockdown to improve predictions about the Covid-19 infection curve. In that case, the only alternative to network representation is to invoke the so-called mean-field assumption, which says that the probability of anyone coming into contact with anyone else is the same. As Thurner et al. argue, however, the mean-field assumption breaks down under conditions of lockdown. In the third case study, Dyhrfjeld-Johnson et al. showed that the propensity for epilepsy in patients with head trauma is explained by the degree of small-worldness of the dentate gyrus. You cannot invoke the small-world property in an explanation without some form of network representation. In this case, therefore, the network representation is not only practically necessary, but also logically necessary.

Furthermore, it does seem that the inferences enabled by network representations are at least sometimes profound. In each of the case studies, the inference enabled by appeal to network representation delivered a non-obvious answer to a substantive and important scientific question: (i) What is the function of that protein? (ii) why does Covid-19 last so long? (iii) Why do people sometimes get epilepsy after head injuries? If a scientific inference delivers a non-obvious and substantive answer to questions like these, the inference itself may be regarded as profound.

One might still suspect that there is another sense in which network modeling is superficial. Network modeling may appear, from an epistemological point of view, rather like a free lunch: it is substantial, but nevertheless undemanding. That is, although the incorporation of network representation into a given scientific enterprise may substantially enrich that enterprise, it does not demand any expertise. People from each discipline may find it useful to use network modeling, in much the same way they may find it useful to use elementary arithmetic, but the relevant techniques are so straightforward as to make claims of expertise in network science overblown. However, as Elek and Babarkzy (2022) argue, network modeling has become a field in its own right, and one in which it is possible to gain expertise. This expertise is visible in our three case studies. Spirin and Mirny needed to know how to build an algorithm that identifies clusters satisfying quantitative criteria. The design of such algorithms demands considerable computational expertise, even though such expertise is

not domain-specific empirical knowledge. In both the second and third case studies, the authors needed to integrate a static network model with a dynamical model, vary the topological properties of the network model systematically, and record the effects of that intervention on the dynamical model. To achieve that end, both studies went far beyond the simple task of constructing a graph from empirical data. Both studies involved the algorithmic construction of a graph, along with a statistical analysis to check how well the algorithmically-constructed graph fit the empirical data. In summary, we can identify at least four kinds of expert knowledge involved in network modeling.

1. Constructing graphs by algorithm and varying their properties systematically.
2. Devising algorithms to find empirically significant substructures in empirically constructed graphs.
3. Using model fitting statistics to assess how well an algorithmically constructed graph fits an empirically constructed graph.
4. Integrating other representational apparatus into the network model.

Because network modeling is practically necessary for generating at least some profound scientific inferences, and because it is, increasingly, a field in which computational expertise can be accumulated, network modeling cannot reasonably be regarded as a superficial science.

6 Trans-domain applicability

Perhaps the most philosophically interesting thing about network modeling is the fact that the same network properties seem to crop up in many otherwise unrelated empirical domains. Network modelers often suggest that this fact has far-reaching implications. One of the leading voices in network modeling, Albert Barabási, says:

> A key discovery of network science is that the architecture of networks emerging in various domains of science, nature, and technology are similar to each other, a consequence of being governed by the same organizing principles. Consequently we can use a common set of mathematical tools to explore these system ((Barabási, 2016, p.8)).

Is the similarity among observed network architectures really a consequence of the fact that networks across domains are governed by the same organizing principles? One of the most discussed principles in the network science literature is that all or most empirical networks have a scale-free distribution over node degree. That is, where k is the degree of a single node, the distribution over values of k is given by: $P(k) = k^{(-\alpha)}$, where the critical exponent α takes a value between 2 and 3. There is a robust and quite technical debate about how broadly this power law relation actually applies to measured networks (Voitalov et al., 2019; Broido and Clauset, 2019; Zhang et al., 2015; Newman, 2005). The debate proceeds by gathering node-distribution estimates from many different data sets, fitting those estimates to a power law, and summarizing the results in a large table. The curious thing about this debate is that little effort has been given toward determining the appropriate reference class. That is, if we think of the power law relation above as a first-order property of empirical systems, then there should be a way of expressing the generalization in terms of the logical schema: $\forall x(Fx \rightarrow Gx)$ where G refers to the property of satisfying the power law distribution. The reference class problem is simply that we should be able to say which property plays the role of F. As far as I can tell, there seems to be no hard conceptual boundaries on the class of systems that can be modeled as networks. If that impression is correct, the property F may not exist. In that case, the frequency with which scale-free node degree distributions appear in nature is simply undefined, and the debate about the relative frequency of scale-free networks is conceptually muddled. Perhaps all that is needed to rectify the situation is a more careful exposition of the goal of gathering such data sets. Rather than framing the work as an attempt to capture the frequency of network properties in nature, it can be framed as an attempt to examine the many varieties of scale-free networks in nature, and then, to work out whether they have other interesting properties in common.

Thus far, I have talked about trans-domain applicability of networks or network models. If we want to think more rigorously about the unit of scientific representation that has the capacity to be applied across empirical domains, these shorthand expressions are likely to be misleading. The term "network" is ambiguous between an object in nature and our representation of it. Arguably, the term "network model" refers to something that is not applicable across domains at all. This will depend on philosophical theories about what models are. As an example, however, consider the view a network model consists of two parts: (i) a graph, and (ii) a mapping between the graph and the target

system. If you conceive of models that way, then their individuation conditions are too fine-grained for them to be applicable across empirical domains. Instead, as suggested in Humphreys (2004), we should invoke the broader notion of a computational template. Computational templates are syntactic objects that come with an intended interpretation, but which are flexible enough to be applied to new phenomena. Their primary attraction is the fact that they serve as a point around which computational expertise can be gathered. As Richard Feynman likes to repeat in his famous lectures on physics, the same equations have the same solutions Feynman et al. (2010). Science is therefore easier if you can re-use equations whose solutions, whether computational or analytic, have been worked out by others.

In the case of network modeling, we should say that the unit of scientific representation that has the capacity for trans-domain applicability is the network template. Although this idea is new, there is a rich and growing literature on trans-domain modeling and computational templates which begins with Humphreys 2004 and includes many more recent articles. Of particular relevance to the discussion of networks is Knuuttila and Loettgers (2016).

7 Generating conditions for networks

If we accept that the claims about network properties being realized in vastly different systems are both true and non-trivial, we will then naturally want to ask why this is the case. This sort of question might have many answers, but one important answer will point to the generating conditions for networks of that kind. If we can show that models of network generation can be described in terms of abstract conditions, and if we can get enough clarity about how to interpret those conditions in empirical terms, then perhaps we can show that different systems that appear radically distinct when described in terms of their more superficial empirical features, are actually quite similar with respect to the fact that they both satisfy these abstract conditions. For example, it may be the case that citation networks of certain kinds can be explained by means of a preferential attachment model. It also may be the case that the distribution of city sizes can be explained by a similar preferential attachment model. We would then have a kind of model-based unification of the two processes that is grounded in an etiological mechanism, but which is nevertheless formulated at a level of abstraction that seems to leave implementation details far behind.

Preferential attachment, as formalized in the well-known BA model mentioned at the beginning, is only one of several models of network construction. Others include the Initial Attractiveness Model, Internal Links Model, Node Deletion Model, Accelerated Growth Model, and the Aging Model (for details see Barabási (2016)). All of these models of network generation are domain-general. So, although they could perhaps furnish an etiological explanation in terms of causes, they would presumably differ from mechanistic explanations, which are domain-specific. An interesting avenue of further research is to articulate the kind of explanation we have when we show that the process by which some empirical system was generated is well-described by one of these network generation models above. If we could gather a long list of such systems, we could group them in terms of their abstract generating model, and thereby broaden the unification base.

To conclude, network science shows that what you can learn about some phenomenon depends on the way you represent it. Most of the prominent examples of network science are cases in which we learn new facts about target systems that are composed of elements that are already familiar. Nevertheless, by choosing to represent these systems as networks, there is much to be learned that would have been inaccessible without the network representation.

References

Barabási, A.-L. (2016). *Network Science*. Cambridge University Press.

Barabási, A.-L., Albert, R., and Jeong, H. (1999). Mean-field theory for scale-free random networks. *Physica A: Statistical Mechanics and its Applications*, 272(1-2):173–187.

Broido, A. D. and Clauset, A. (2019). Scale-free networks are rare. *Nature communications*, 10(1):1–10.

Dyhrfjeld-Johnsen, J., Santhakumar, V., Morgan, R. J., Huerta, R., Tsimring, L., and Soltesz, I. (2007). Topological determinants of epileptogenesis in large-scale structural and functional models of the dentate gyrus derived from experimental data. *Journal of neurophysiology*, 97(2):1566–1587.

Elek, G. and Babarczy, E. (2022). Taming vagueness: The philosophy of network science. *Synthese*, 200(2):1–31.

Erdős, P., Rényi, A., et al. (1960). On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60.

Feynman, R. P., Leighton, R. B., and Sands, M. (2010). *The Feynman lectures on physics; New millennium ed.* Basic Books, New York, NY. Originally published 1963-1965.

Granovetter, M. S. (1973). The strength of weak ties. *American journal of sociology*, 78(6):1360–1380.

Hagberg, A. A., Schult, D. A., and Swart, P. J. (2008). Exploring network structure, dynamics, and function using networkx. In Varoquaux, G., Vaught, T., and Millman, J., editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA.

Humphreys, P. (2004). *Extending ourselves: Computational science, empiricism, and scientific method*. Oxford University Press.

Knuuttila, T. and Loettgers, A. (2016). Model templates within and between disciplines: from magnets to gases–and socio-economic systems. *European journal for philosophy of science*, 6(3):377–400.

Kostić, D. (2018). Mechanistic and topological explanations: an introduction. *Synthese*, 195(1):1–10.

Kostić, D. and Khalifa, K. (2021). The directionality of topological explanations. *Synthese*, (5-6):14143–14165.

Newman, M. E. (2005). Power laws, pareto distributions and zipf's law. *Contemporary physics*, 46(5):323–351.

Omranian, S., Nikoloski, Z., and Grimm, D. G. (2022). Computational identification of protein complexes from network interactions: Present state, challenges, and the way forward. *Computational and Structural Biotechnology Journal*.

Pastor-Satorras, R. and Vespignani, A. (2001). Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200.

Percha, B., Dzakpasu, R., Żochowski, M., and Parent, J. (2005). Transition from local to global phase synchrony in small world neural network and its possible implications for epilepsy. *Physical Review E*, 72(3):031909.

Rathkopf, C. (2018). Network representation and complex systems. *Synthese*, (1).

Spirin, V. and Mirny, L. A. (2003). Protein complexes and functional modules in molecular networks. *Proceedings of the national Academy of sciences*, 100(21):12123–12128.

Thurner, S., Klimek, P., and Hanel, R. (2020). A network-based explanation of why most covid-19 infection curves are linear. *Proceedings of the National Academy of Sciences*, 117(37):22684–22689.

Trudeau, R. (1976). *Introduction to Graph Theory*. Dover Publishers.

Voitalov, I., van der Hoorn, P., van der Hofstad, R., and Krioukov, D. (2019). Scale-free networks well done. *Phys. Rev. Research*, 1:033034.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442.

Zhang, L., Small, M., and Judd, K. (2015). Exactly scale-free scale-free networks. *Physica A: Statistical Mechanics and its Applications*, 433:182–197.