A Survey of Statistical Network Models

Citation

Goldenberg, Anna, Alice X. Zheng, Stephen E. Fienberg, and Edoardo M. Airoldi. A survey of statistical network models. 2009. Machine Learning 2(2): 129-233.

Published Version

doi:10.1561/22000000005

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:4645865>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

A Survey of Statistical Network Models

Anna Goldenberg
University of Toronto

Alice X. Zheng
Microsoft Research

Stephen E. Fienberg
Carnegie Mellon University

Edoardo M. Airoldi
Harvard University

December 2009

Contents

Preface	1
1 Introduction	3
1.1 Overview of Modeling Approaches	4
1.2 What This Survey Does Not Cover	7
2 Motivation and Dataset Examples	9
2.1 Motivations for Network Analysis	9
2.2 Sample Datasets	10
2.2.1 Sampson’s “Monastery” Study	11
2.2.2 The Enron Email Corpus	12
2.2.3 The Protein Interaction Network in Budding Yeast	14
2.2.4 The Add Health Adolescent Relationship and HIV Transmission Study	14
2.2.5 The Framingham “Obesity” Study	16
2.2.6 The NIPS Paper Co-Authorship Dataset	17
3 Static Network Models	21
3.1 Basic Notation and Terminology	21
3.2 The Erdős-Rényi-Gilbert Random Graph Model	22
3.3 The Exchangeable Graph Model	23
3.4 The p_1 Model for Social Networks	27
3.5 p_2 Models for Social Networks and Their Bayesian Relatives	29
3.6 Exponential Random Graph Models	30
3.7 Random Graph Models with Fixed Degree Distribution	32
3.8 Blockmodels, Stochastic Blockmodels and Community Discovery	33
3.9 Latent Space Models	36
3.9.1 Comparison with Stochastic Blockmodels	38
4 Dynamic Models for Longitudinal Data	41
4.1 Random Graphs and the Preferential Attachment Model	41
4.2 Small-World Models	44
4.3 Duplication-Attachment Models	46
4.4 Continuous Time Markov Chain Models	47

4.5	Discrete Time Markov Models	50
4.5.1	Discrete Markov ERGM Model	51
4.5.2	Dynamic Latent Space Model	52
4.5.3	Dynamic Contextual Friendship Model (DCFM)	53
5	Issues in Network Modeling	57
6	Summary	61
	Bibliography	65

Preface

Networks are ubiquitous in science and have become a focal point for discussion in everyday life. Formal statistical models for the analysis of network data have emerged as a major topic of interest in diverse areas of study, and most of these involve a form of graphical representation. Probability models on graphs date back to 1959. Along with empirical studies in social psychology and sociology from the 1960s, these early works generated an active “network community” and a substantial literature in the 1970s. This effort moved into the statistical literature in the late 1970s and 1980s, and the past decade has seen a burgeoning network literature in statistical physics and computer science. The growth of the World Wide Web and the emergence of online “networking communities” such as *Facebook*, *MySpace*, and *LinkedIn*, and a host of more specialized professional network communities has intensified interest in the study of networks and network data.

Our goal in this review is to provide the reader with an entry point to this burgeoning literature. We begin with an overview of the historical development of statistical network modeling and then we introduce a number of examples that have been studied in the network literature. Our subsequent discussion focuses on a number of prominent static and dynamic network models and their interconnections. We emphasize formal model descriptions, and pay special attention to the interpretation of parameters and their estimation. We end with a description of some open problems and challenges for machine learning and statistics.

Chapter 1

Introduction

Many scientific fields involve the study of networks in some form. Networks have been used to analyze interpersonal social relationships, communication networks, academic paper coauthorships and citations, protein interaction patterns, and much more. Popular books on networks and their analysis began to appear a decade ago, [see, e.g., [24](#); [50](#); [318](#); [319](#); [68](#)] and online “networking communities” such as *Facebook*, *MySpace*, and *LinkedIn* are an even more recent phenomenon.

In this work, we survey selective aspects of the literature on statistical modeling and analysis of networks in social sciences, computer science, physics, and biology. Given the volume of books, papers, and conference proceedings published on the subject in these different fields, a single comprehensive survey would be impossible. Our goal is far more modest. We attempt to chart the progress of statistical modeling of network data over the past seventy years and to outline succinctly the major schools of thought and approaches to network modeling and to describe some of their interconnections. We also attempt to identify major statistical gaps in these modeling efforts. From this overview one might then synthesize and deduce promising future research directions. Kolaczyk [[177](#)] provides a complementary statistical overview.

The existing set of statistical network models may be organized along several major axes. For this article, we choose the axis of static vs. dynamic models. Static network models concentrate on explaining the observed set of links based on a single snapshot of the network, whereas dynamic network models are often concerned with the mechanisms that govern changes in the network over time. Most early examples of networks were single static snapshots. Hence static network models have been the main focus of research for many years. However, with the emergence of online networks, more data is available for dynamic analysis, and in recent years there has been growing interest in dynamic modeling.

In the remainder of this chapter we provide a brief historical overview of network modeling approaches. In subsequent chapters we introduce some examples studied in the network literature and give a more detailed comparative description of select modeling approaches.

1.1 Overview of Modeling Approaches

Almost all of the “statistically” oriented literature on the analysis of networks derives from a handful of seminal papers. In social psychology and sociology there is the early work of Simmel and Wolff [268] at the turn of the last century and Moreno [221] in the 1930s as well as the empirical studies of Stanley Milgram [215; 298] in the 1960s; in mathematics/probability there is the Erdős-Rényi paper on random graph models [94]. There are other papers that dealt with these topics contemporaneously or even earlier. But these are the ones that appear to have had lasting impact.

Moreno [221] invented the sociogram — a diagram of points and lines used to represent relations among persons, a precursor to the graph representation for networks. Luce and others developed a mathematical structure to go with Moreno’s sociograms using incidence matrices and graphs (see, e.g., [202; 200; 201; 203; 244; 282; 11]), but the structure they explored was essentially deterministic. Milgram gave the name to what is now referred to as the “Small World” phenomenon — short paths of connections linking most people in social spheres — and his experiments had provocative results: the shortest path between any two people for completed chains has a median length of around 6; however, the majority of chains initiated in his experiments were never completed! (His studies provided the title for the play and movie *Six Degrees of Separation*, ignoring the compleity of his results due to the censoring.) White [321] and Fienberg and Lee [100] gave a formal Markov-chain like model and analysis of the Milgram experimental data, including information on the uncompleted chains. Milgram’s data were gathered in batches of transmission, and thus these models can be thought of as representing early examples of generative descriptions of dynamic network evolution. Recently, Dodds et al. [86] studied a global “replication” variation on the Milgram study in which more than 60,000 e-mail users attempted to reach one of 18 target persons in 13 countries by forwarding messages to acquaintances. Only 384 of 24,163 chains reached their targets but they estimate the median length for completions to be 7, by assuming that attrition occurs at random.

The social science network research community that arose in the 1970s was built upon these earlier efforts, in particular the Erdős-Rényi-Gilbert model. Research on the Erdős-Rényi-Gilbert model (along with works by Katz et al. [166; 168; 167]) engendered the field of random graph theory. In their papers, Erdős and Rényi worked with fixed number of vertices, N , and number of edges, E , and studied the properties of this model as E increases. Gilbert studied a related two-parameter version of the model, with N as the number of vertices and p the fixed probability for choosing edges. Although their descriptions might at first appear to be static in nature, we could think in terms of adding edges sequentially and thus turn the model into a dynamic one. In this alternative binomial version of the Erdős-Rényi-Gilbert model, the key to asymptotic behavior is the value $\lambda = pN$. There is a “phase change” associated with the value of $\lambda = 1$, at which point we shift from seeing many small connected components in the form of trees to the emergence of a single “giant connected component.” Probabilists such as Pittel [243] imported ideas and results from stochastic processes into the random graph literature.

Holland and Leinhardt [149]’s p_1 model extended the Erdős-Rényi-Gilbert model to allow

for differential attraction (popularity) and expansiveness, as well as an additional effect due to reciprocation. The p_1 model was log-linear in form, which allowed for easy computation of maximum likelihood estimates using a contingency table formulation of the model [101; 102]. It also allowed for various generalizations to multidimensional network structures [103] and stochastic blockmodels. This approach to modeling network data quickly evolved into the class of p^* or exponential random graph models (ERGM) originating in the work of Frank and Strauss [110] and Strauss and Ikeda [287]. A trio of papers demonstrating procedures for using ERGMs [316; 241; 254] led to the wide-spread use of ERGMs in a descriptive form for cross sectional network structures or cumulative links for networks—what we refer to here as static models. Full maximum likelihood approaches for ERGMs appeared in the work of Snijders and Handcock and their collaborators, some of which we describe in [chapter 3](#).

Most of the early examples of networks in the social science literature were relatively small (in terms of the number of nodes) and involved the study of the network at a fixed point in time or cumulatively over time. Only a few studies (e.g., Sampson’s 1968 data on novice monks in the monastery [259]) collected, reported, and analyzed network data at multiple points in time so that one could truly study the evolution of the network, i.e., network dynamics. The focus on relatively small networks reflected the state-of-art of computation but it was sufficient to trigger the discussion of how one might assess the fit of a network model. Should one focus on “small sample” properties and exact distributions given some form of minimal sufficient statistic, as one often did in other areas of statistics, or should one look at asymptotic properties, where there is a sequence of networks of increasing size? Even if we have “repeated cross-sections” of the network, if the network is truly evolving in continuous time we need to ask how to ensure that the continuous time parameters are estimable. We return to many of these question in subsequent chapters.

In the late 1990s, physicists began to work on network models and study their properties in a form similar to the macro-level descriptions of statistical physics. Barabási, Newman, and Watts, among others, produced what we can think of as variations on the Erdős-Rényi-Gilbert model which either controlled the growth of the network or allowed for differential probabilities for edge addition and/or deletion. These variations were intended to produce phenomena such as “hubs,” “local clustering,” and “triadic closures.” The resulting models gave us fixed degree distribution limits in the form of power laws — variations on preferential attachment models (“the rich get richer”) that date back to Yule [329] and Simon [269] (see also [218]) — as well as what became known as “small world” models. The small-world phenomenon, which harks back to Milgram’s 1960s studies, usually refers to two distinct properties: (1) small average distance and (2) the “clustering” effect, where two nodes with a common neighbor are more likely to be adjacent. Many of these authors claim that these properties are ubiquitous in realistic networks. To model networks with the small-world phenomenon, it is natural to utilize randomly generated graphs with a power law degree distribution, where the fraction of nodes with degree k is proportional to k^{-a} for some positive exponent a . Many of the most relevant papers are included in an edited collection by Newman et al. [231]. More recently this style of statistical physics models have been used to detect community structure in networks, e.g., see Girvan and Newman [122] and

Backstrom et al. [20], a phenomenon which has its counterpart description in the social science network modeling literature.

The probabilistic literature on random graph models from the 1990s made the link with epidemics and other evolving stochastic phenomena. Picking up on this idea, Watts and Strogatz [320] and others used epidemic models to capture general characteristics of the evolution of these new variations on random networks. Durrett [91] has provided us with a book-length treatment on the topic with a number of interesting variations on the theme. The appeal of stochastic processes as descriptions of dynamic network models comes from being able to exploit the extensive literature already developed, including the existence and the form of stationary distributions and other model features or properties. Chung and Lu [69] provide a complementary treatment of these models and their probabilistic properties.

One of the principal problems with this diverse network literature that we see is that, with some notable exceptions, the statistical tools for estimation and assessing the fit of “statistical physics” or stochastic process models is lacking. Consequently, no attention is paid to the fact that real data may often be biased and noisy. What authors in the network literature have often relied upon is the extraction of key features of the related graphical network representation, e.g., the use of power laws to represent degree distributions or measures of centrality and clustering, without any indication that they are either necessary or sufficient as descriptors for the actual network data. Moreover, these summary quantities can often be highly misleading as the critique by Stouffer et al. [285, 286] of methods used by Barabási [25] and Vázquez et al. [304] suggest. Barabási claimed that the dynamics of a number of human activities are scale-free, i.e., he specifically reported that the probability distribution of time intervals between consecutive e-mails sent by a single user and time delays for e-mail replies follow a power-law with exponent -1 , and he proposed a priority-queuing process as an explanation of the bursty nature of human activity. Stouffer et al. [286] demonstrated that the reported power-law distribution was solely an artifact of the analysis of the empirical data and used Bayes factors to show that the proposed model is not representative of e-mail communication patterns. See a related discussion of the poor fit of power laws in Clauset et al. [74]. There are several works, however, that try to address model fitting and model comparison. For example, the work of Williams and Martinez [323] showed how a simple two-parameter model predicted “key structural properties of the most complex and comprehensive food webs in the primary literature”. Another good example is the work of Middendorf et al. [214] where the authors used network motif counts as input to a discriminative systematic classification for deciding which configuration model the actual observed network came from; they looked at power law, small-world, duplication-mutation and duplication-mutation-complementation and other models (seven in total) and concluded that the duplication-mutation-complementation model described the protein-protein interaction data in *Drosophila melanogaster* species best.

Machine learning approaches emerged in several forms over the past decade with the empirical studies of Faloutsos et al. [97] and Kleinberg [173, 172, 174], who introduced a model for which the underlying graph is a grid—the graphs generated do not have a power law degree distribution, and each vertex has the same expected degree. The strict

requirement that the underlying graph be a cycle or grid renders the model inapplicable to webgraphs or biological networks. Durrett [91] treats variations on this model as well. More recently, a number of authors have looked to combine the stochastic blockmodel ideas from the 1980s with latent space models, model-based clustering [137] or mixed-membership models [9], to provide generative models that scale in reasonable ways to substantial-sized networks. The class of mixed membership models resembles a form of soft clustering [95] and includes the latent Dirichlet allocation model [41] from machine learning as a special case. This class of models offers much promise for the kinds of network dynamical processes we discuss here.

1.2 What This Survey Does Not Cover

This survey focuses primarily on statistical network models and their applications. As a consequence there are a number of topics that we touch upon only briefly or essentially not at all, such as

- *Probability theory associated with random graph models.* The probabilistic literature on random graph models is now truly extensive and the bulk of the theorems and proofs, while interesting in their own right, are largely unconnected with the present exposition. For excellent introductions to this literature, see Chung and Lu [69] and Durrett [91]. For related results on the mathematics of graph theory, see Bollobás [43].
- *Efficient computation on networks.* There is a substantial computer science literature dealing with efficient calculation of quantities associated with network structures, such as shortest paths, network diameter, and other measures of connectivity, centrality, clustering, etc. The edited volume by Brandes and Erlebach [48] contains good overviews of a number of these topics as well as other computational issues associated with the study of graphs.
- *Use of the network as a tool for sampling.* Adaptive sampling strategies modify the sampling probabilities of selection based on observed values in a network structure. This strategy is beneficial when searching for rare or clustered populations. Thompson and Seber [296] and Thompson [293] discuss adaptive sampling in detail. There is also related work on target sampling [294] and respondent-driven sampling [258; 305].
- *Neural networks.* Neural networks originated as simple models for connections in the brain but have more recently been used as a computational tool for pattern recognition (e.g., Bishop [38]), machine learning (e.g., Neal [228]), and models of cognition (e.g., Rogers and McClelland [257]).
- *Networks and economic theory.* A relatively new area of study is the link between network problems, economic theory, and game theory. Some useful entrees to this literature are Even-Dar and Kearns [96], Goyal [131], Kearns et al. [169], and Jackson

[160], whose book contains an excellent semi-technical introduction to network concepts and structures.

- *Relational networks.* This is a very popular area in machine learning. It uses probabilistic graphical models to represent uncertainty in the data. The types of “networks” in this area, such as Bayes nets, dependency diagrams, etc., have a different meaning than the networks we consider in this review. The main difference is that the networks in our work are considered to “be given” or arising directly from properties of the network under study, rather than being representative of the uncertainty of the relationships between nodes and node attributes. There is a multitude of literature on relational networks, e.g., see Friedman et al. [112], Getoor et al. [117], Neville and Jensen [229]; Neville et al. [230], and Getoor and Taskar [116].
- *Bi-partite graphs.* These are graphs that represent measurement on two populations of objects, such as individuals and features. The graphs in this context are seldom the best representation of the data, with exception perhaps of binary measurements or when the true populations have comparable sizes. Recent work on exchangeable Rasch matrices is related to this topic and potentially relevant for network analysis. Lauritzen [186, 187]; Bassetti et al. [29] suggest applications to bipartite graphs.
- *Agent-based modeling.* Building on older ideas such as cellular automata, agent-based modeling attempts to simulate the simultaneous operations of multiple agents, in an effort to re-create and predict the actions of complex phenomena. Because the interest is often on the interaction among the agents, this domain of research has been linked with network ideas. With the recent advances in high-performance computing, simulations of large-scale social systems have become an active area of research, e.g., see [46]. In particular, there is a strong interest in areas that revolve around national security and the military, with studies on the effects of catastrophic events and biological warfare, as well as computational explorations of possible recovery strategies [57; 59]. These works are the contemporary counterparts of more classical work at the interface between artificial intelligence and the social sciences [54; 56; 55].

Chapter 2

Motivation and Dataset Examples

2.1 Motivations for Network Analysis

Why do we analyze networks? The motivation behind network analysis is as diverse as the origin of network problems within differing academic fields. Before we delve into details of the “how” of statistical network modeling, we start with some examples of the “why.” This chapter also includes descriptions of popular datasets for interested readers who may wish to exercise their modeling muscles.

Social scientists are often interested in questions of interpretation such as the meanings of edges in a social network [181]. Do they arise out of friendliness, strategic alliance, obligation, or something else? When the meaning of edges are known, the object is often to characterize the structure of those relations (e.g., whether friendships or strategic alliances are hierarchical or transitive). A large volume of statistically-oriented social science literature is dedicated to modeling the mechanisms and relations of network properties and testing hypotheses about network structure, see, e.g., [280].

Physicists, on the other hand, tend to be interested in understanding parsimonious mechanisms for network formation [28; 235]. For example, a common modeling goal is to explain how a given network comes to have its particular degree distribution or diameter at time t .

Several network analysis concepts have found niches in computational biology. For example, work on protein function classification can be thought of as finding hidden groups in the protein-protein interaction network [7; 8] to gain better understanding of underlying biological processes. Label propagation (node similarity) in networks can be harnessed to help with functional gene annotation [226]. Graph alignment can be used to locate subgraphs that are common among species, thus advancing our understanding of evolution [105]. Motif finding, or more generally the search for subgraph patterns, also has many applications [17]. Combining networks from heterogeneous data sources helps to improve the accuracy of predicted genetic interactions [327]. Heterogeneity of network data sources in biology introduces a lot of noise into the global network structure, especially when networks created for different purposes (such as protein co-regulation and gene co-expression) are combined. [225] addresses network de-noising via degree-based structure priors on graphs. For a review

of biological applications of networks, please see [332].

The task of finding hidden groups is also relevant in analyzing communication networks, e.g., in detecting possible latent terrorist cells [30]. The related task of discovering the “roles” of individual nodes is useful for identity disambiguation [36] and for business organization analysis [207]. These applications often take the machine learning approach of graph partitioning, a topic previously known in social science and statistics literature as blockmodeling [199; 89]. A related question is *functional* clustering, where the goal is not to statistically cluster the network, but to discover members of dynamic communities with similar functions based on existing network connectivity [122; 232; 234; 266].

In the machine learning community, networks are often used to predict missing information, which can be edge related, e.g., predicting missing links in the network [238; 73; 198], or attribute related, e.g., predicting how likely a movie is to be a box office hit [229]. Other applications include locating the crucial missing link in a business or a terrorist network, or calculating the probability that a customer will purchase a new product, given the pattern of purchases of his friends [142]. The latter question can more generally be stated as predicting individual’s preferences given the preferences of her “friends”. This research direction has evolved into an area of its own under the name of *recommender systems*, which has recently received a lot of media attention due to the competition by the largest online movie rental company Netflix. The company has awarded a prize of one million dollars to a team of researchers that were able to predict customer ratings of movies with higher than 10% accuracy than their own in-house system [290].

The concept of information propagation also finds many applications in the network domain, such as virus propagation in computer networks [310], HIV infection networks [222; 163; 164], viral marketing [87] and more generally gossiping [170]. Here some work focuses on finding network configurations optimal for routing, while other research assumes that the network structure is given and focus on suitable models for disease or information spread.

2.2 Sample Datasets

A plethora of data sets are available for network analysis, and more are emerging every year. We provide a quick guided tour of the most popular datasets and applications in each field.

In his ground-breaking paper, Milgram [215] experimented with the construction of interpersonal social networks. His result that the median length of completed chains was approximately 6 led to the pop-culture coining of the phrase “six degrees of separation.” Subjects of subsequent studies ranged from social interactions of monks [259], to hierarchies of elephants [209; 303], to sexual relationships between adults of Colorado [176], to friendships amongst elementary school students [141; 299].

While a lot of biological applications focus on the study of protein-protein interaction networks [114; 115; 184; 248; 328], metabolic networks [158], functional and co-expression gene similarity networks and gene regulatory networks [111; 309], computer science applications revolve around e-mail [207], the internet [97; 63; 151], the web [152; 13], academic paper co-authorship [127] and citation networks [204; 216]. Citation networks have a long history

of modeling in different areas of research starting with the seminal paper of de Solla Price [83] and more recently in physics [190]. With the recent rise of online networks, computer science and social science researchers are also starting to examine blogger networks such as *LiveJournal*, social networks found on *Friendster*, *Facebook*, *Orkut*, and dating networks such as *Match.com*.

Terrorist networks (often simulated) and telecommunication networks have come under similar scrutiny, especially since the events of September 11, 2001 (e.g., see [182; 250; 249; 62]). There has also been work on ecological networks such as foodwebs [323; 16], neuronal networks [188], network epidemiology [306], economic trading networks [123], transportation networks (roads, railways, airplanes; e.g., [113]), resource distribution networks, mobile phone networks [92] and many others.

Several network data repositories are available on public websites and as part of packages. For example, UCINET¹ includes a lot of well known smaller scale datasets such as the Davis Southern Club Women dataset [80], Zachary’s karate club dataset [330], and Sampson’s monk data [259] described below. Pajek² contains a larger set of small and large networks from domains such as biology, linguistics, and food-web. Additional datasets in a variety of domains include power grid networks, US politics, cellular and protein networks and others³. A collection of large and very large directed and undirected networks in the areas of communication, citation, internet and others are available as part of Stanford Network Analysis Package (SNAP)⁴.

We now introduce six examples of networks studied in the literature, describing the data in reasonable detail and including graphs depicting the networks wherever feasible. For each network example we articulate specific questions of interest.

2.2.1 Sampson’s “Monastery” Study

A classic example of a social network is the one derived from the survey administered by Sampson and published in his doctoral dissertation [259]. Figure 2.1 displays the network derived from the “whom do you like” sociometric relations in this dataset. Sampson spent several months in a monastery in New England, where a number of novices were preparing to join a monastic order. Sampson’s original analysis was rooted in direct anthropological observations. He strongly suggested the existence of tight factions among the novices: the loyal opposition (whose members joined the monastery first), the young turks (whose members joined later on), the outcasts (who were not accepted in either of the two main factions), and the waverers (who did not take sides). The events that took place during Sampson’s stay at the monastery supported his observations. For instance, John and Gregory, two members of the young turks, were expelled over religious differences, and other members resigned

¹<http://www.analytictech.com/ucinet/>

²<http://vlado.fmf.uni-lj.si/pub/networks/data/>

³<http://www-personal.umich.edu/~mejn/netdata/>
<http://cdg.columbia.edu/cdg/datasets>

<http://www.nd.edu/~networks/resources.htm>

⁴<http://snap.stanford.edu/data/>

Figure 2.1: Network derived from “whom do you like” sociometric relations collected by Sampson.

shortly after these events. About a year after leaving the monastery, Sampson surveyed all of the novices, and asked them to rank the other novices in terms of four sociometric relations: like/dislike, esteem, personal influence, and alignment with the monastic credo, retrospectively, at four different epochs spanning his stay at the monastery.

The presence of a well defined social structure within the monastery (the factions) that can be inferred from responses to the survey, as well as the social dynamics of subtle ideological conflicts that led to the dissolution of the monastic order, have much intrigued both statisticians and social scientists for the past four decades. Researchers typically consider the faction labels assigned by Sampson to the novices as the anthropological ground truth in their analysis. For example analyses, we refer to [103; 137; 81; 9].

2.2.2 The Enron Email Corpus

The Enron email corpus has been widely studied in recent machine learning network literature. Enron Corporation was an energy and trading company specializing in the marketing of electricity and gas. In 2000 it was the seventh largest company in the United States with reported revenues of over \$100 billion. On December 2, 2001, Enron filed for bankruptcy. The sudden collapse cast suspicions over its management and prompted federal investigations. Thirty-four Enron officials were prosecuted and top Enron executives and associates were subsequently found to be guilty of accounting fraud. During the investigation, the courts subpoenaed extensive email logs from most of Enron’s employees, and the Federal Energy Regulatory Commission (FERC) published the database online.⁵ Subsequently, researchers

⁵<http://www.ferc.gov/industries/electric/indus-act/wec/enron/info-release.asp>

Figure 2.2: E-mail exchange data among 151 Enron executives, using a threshold of a minimum of 5 messages for each link. Source: [153].

in the CALO (Cognitive Assistant that Learns and Organizes) project corrected integrity problems in the dataset.⁶ The original FERC dataset contains 619,446 email messages (about 92% of Enron’s staff emails), and the cleaned-up CALO dataset contains 200,399 messages from 158 users. Another version of the data consists of the contents of the mail folders of the top 151 executives, containing about 225,000 messages covering a period from 1997 to 2004.⁷ Figure 2.2 and Figure 2.3 give network snapshots of the e-mail traffic among these 151 executives with thresholds of 5 and 30 messages, respectively.

Research activity on the Enron dataset range from document classification to social-

⁶<http://www.cs.cmu.edu/~enron/>

⁷<http://www.isi.edu/~adibi/Enron/Enron.htm>

Figure 2.3: E-mail exchange data among 151 Enron executives, using a threshold of a minimum of 30 messages for each link. Source: [153].

network analysis to visualization. A collection of papers working with the Enron corpus were gathered together in a special 2005 issue of *Computational & Mathematical Organization Theory*, see [58].

2.2.3 The Protein Interaction Network in Budding Yeast

The budding yeast is a unicellular organism that has become a de-facto model organism for the study of molecular and cellular biology [47]. There are about 6,000 proteins in the budding yeast, which interact in a number of ways [64]. For instance, proteins bind together to form protein complexes, the physical units that carry out most functions in the cell [184]. In recent years, a large amount of resources has been directed to collect experimental evidence of physical proteins binding, in an effort to infer and catalogue protein complexes and their multifaceted functional roles [e.g. 98; 159; 300; 114; 143]. Currently, there are four main sources of interactions between pairs of proteins that target proteins localized in different cellular compartments with variable degrees of success: (i) literature curated interactions [248], (ii) yeast two-hybrid (Y2H) interaction assays [328], (iii) protein fragment complementation (PCA) interaction assays [291], and (iv) tandem affinity purification (TAP) interaction assays [115; 184]. These collections include a total of about 12,292 protein interactions [162], although the number of such interactions is estimated to be between 18,000 [328] and 30,000 [307]. Figure 2.4 shows a popular image of the interaction network among proteins in the budding yeast, produced as part of an analysis by Barabási and Oltvai [27].

Statistical methods have been developed for analyzing many aspects of this large protein interaction network, including de-noising [32; 8], function prediction [227], and identification of binding motifs [23].

2.2.4 The Add Health Adolescent Relationship and HIV Transmission Study

The National Longitudinal Study of Adolescent Health (Add Health) is a study of adolescents in the United States drawn from a representative sample of middle, junior high, and highschools. The study focused on patterns of friendship, sexual relationships, as well as disease transmissions. To date, four waves of surveys have been collected over the course of fifteen years.

Wave I surveys occurred between 1994 to 1995 and included 90,118 students from 145 schools across the country. Each student completed an in-school questionnaire on his or her family background, school life and activities, friendships, and health status. Administrators from participating schools also completed questionnaires about student demography and school curriculum and services. In addition, 20,745 students were chosen for an in-home interview that included more sensitive topics such as sexual behavior. For 16 selected schools (two large and fourteen small), Add Health attempted to administer the in-home survey to all enrolled students. This saturated sample distinguishes itself from the ego-centric and

Figure 2.4: A popular image of the protein interaction network in *Saccharomyces cerevisiae*, also known as the budding yeast. The figure is reproduced with permission. Source: [27].

snowball samples collected from past studies; it allows for the construction of relationship networks with more accurate global characteristics. The fully observed friendship networks in all the schools are also a valuable resource and an important contribution of this work.

Wave II data collection occurred 18-months after Wave I in 1996 and followed up on the in-home interviews. The dataset covered 14,738 adolescents and 128 school administrators. Based on the data collected from Wave I and II, Bearman et al. [31] constructed the timed sequence of relationship networks amongst students from the two large schools with saturated sampling. The resulting sexual relationship network bears strong resemblance to a spanning tree as opposed to previously hypothesized core or inverse-core structures⁸ (See Figure 2.5.)

Wave III interviews were conducted in 2001 and 2002 with topics including marriage,

⁸A core is a group of inter-connected individuals who sit at the center of the graph and interact with individuals on the periphery. An inverse core is a group of central individuals who are connected to those on the periphery but not to each other.

Figure 2.5: The Add Health sexual relationships network of US highschool adolescents. This figure is reproduced with permission. Source: Bearman et al. [31]

childbearing, and sexually transmitted diseases. Of the original Wave I in-home respondents, 15,170 were interviewed again for Wave III. Of these, 13,184 participants provided oral fluid specimens for HIV testing. Morris et al. [223] studied the prevalence of HIV infections among young adults based on data collected in Wave III.

Wave IV interviews were conducted in 2007 and 2008 with the original Wave I respondents, who are now dispersed across the nation in all 50 states. Of the original respondents, 92.5% were located and 80.3% were interviewed. The interview included a comprehensive survey of the social, emotional, spiritual, and physical aspects of health. Physical measurements, biospecimen, and geographical data were also collected.

For detailed information about the data, as well as access to the public-domain and restricted-access datasets, see <http://www.cpc.unc.edu/projects/addhealth>.

2.2.5 The Framingham “Obesity” Study

One of the most famous and important epidemiological studies was initiated in Framingham, Massachusetts, a suburb of Boston, in 1948 with an originally enrolled cohort of 5209 people. In 1971 investigators initiated an “offspring” cohort study which enrolled most of the children of the original cohort and their spouses. Participants completed a questionnaire and underwent physical examinations (including measurements of height and weight) in three-

year periods beginning 1973, 1981, 1985, 1989, 1992, 1997, 1999. Christakis and Fowler [65] derive body mass index information on a total of 12,067 individuals who appeared in any of the Framingham Heart cohorts (one “close friend” for each cohort member).⁹ There were 38,611 observed family and social ties (edges) to the core 5,124 cohort members.

Through a series of network snapshots and statistical analyses, Christakis and Fowler described the evolution of the “clustering” of obesity in this social network. In particular they claim to have examined whether the data conformed to “small-world,” “scale-free,” and “hierarchical” types of random graph network models. Figure 2.6 depicts data on the largest connected subcomponent (the so-called giant component) for the network in 2000, which consists of 2200 individuals. Other analyses in their paper explore attributions of the individuals via longitudinal logistic-regression models with lagged effects. Subsequently, they have published similar papers focused on the dynamics of smoking behavior over time [66] and on happiness [67], both using the structure of Framingham “offspring” cohort.

This work has come under criticism by others. For example Cohen-Cole and Fletcher note that there are plausible alternative explanations to the network structure based on contextual factors [77], and in a separate paper demonstrate that the same methodology detects “implausible” social network effects for such medical conditions as acne and headaches as well as for physical height [78]. The authors answer to these criticisms can be found in [108]. The question of the magnitude and significance of social network effects is still a subject of an ongoing debate.

2.2.6 The NIPS Paper Co-Authorship Dataset

The NIPS dataset contains information on publications that appeared in the *Neural Information Processing Systems* (NIPS) conference proceedings, volumes 1 through 12, corresponding to years 1987-1999—the pre-electronic submission era. The original collection contained scanned full papers made available by Yann LeCunn. Sam Roweis subsequently processed the data to glean information such as title, authorship information, and word counts per document. In total, there are 2,037 authors and 1,740 papers with an average of 2.29 authors per paper and 1.96 papers per author. The NIPS database is available from Sam Roweis’ website¹⁰ in raw and *MATLAB* formats along with a detailed description and information on its construction.

Various authors have used the NIPS data to analyze author-to-author connectivity in static [126] as well as dynamic settings [264]. Li and McCallum [197] modeled the text of the documents and Sarkar et al. [265] analyzed the two-mode network (author-word-author) in a dynamic context. In Figure 2.7 we reproduce a graphic illustration of the inferred dynamic evolution of the network from [263].

⁹A body-mass index value (weight in kg. divided by the square of the height in meters) of 30 or more was taken to indicate obesity.

¹⁰<http://www.cs.toronto.edu/~roweis/data.html>

Figure 2.6: Obesity network from Framingham offspring cohort data. Each node represents one person in the dataset (a total of 2200 in this picture). Circles with red borders denote women, with blue borders – men. The size of each circle is proportional to the body-mass index. The color inside the circle denotes obesity status - yellow is obese (body-mass index ≥ 30 , green is non-obese. The colors of ties between nodes indicate relationships - purple denotes a friendship or marital tie and orange is a familial tie. This figure is reproduced with permission. Source: [65].

Figure 2.7: NIPS paper co-authorship data. Each point represents an author. Two authors are linked by an edge if they have co-authored at least one paper at NIPS. Left: 1991-1994. Right: 1995-1998. Each graph contains all the links for the selected period. Several well known people in the Machine Learning field are highlighted. The size of the circles around selected individuals depend on their number of collaborations. Colors are meant to facilitate visualization. This figure is reproduced with permission. Source: [263].

Chapter 3

Static Network Models

A number of basic network models are essentially static in nature. The statistical activities associated with them focus on certain local and global network statistics and the extent to which they capture the main elements of actual realized networks. In this chapter, we briefly summarize two lines of research. The first originates in the mathematics community with the Erdős-Rényi-Gilbert model and led to two types of generalizations: (i) the “statistical physics” generalizations that led to power laws for degree distributions—the so-called scale-free graphs, and (ii) the exchangeable graph models that introduce weak dependences among the edges in a controlled fashion, which ultimately lead to a range of more structured connectivity patterns and enable model comparison strategies rooted in information theory. A second line of research originated in the statistics and social sciences communities in response to a need for models of social networks. The p_1 model of Holland and Leinhardt, which in some sense generalizes the Erdős-Rényi-Gilbert model, and the more general descriptive family of exponential random graph models effectively initiate this line of modeling. Some of these models also have a *generative* interpretation that allows us to think about their use in a dynamic, evolutionary setting. We define and discuss popular dynamic interpretations of the data generating process, including the generative interpretation, in [chapter 4](#).

3.1 Basic Notation and Terminology

In theoretical computer science, a graph or network G is often defined in terms of nodes and edges, $G \equiv G(\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is a set of nodes and $\mathcal{E}$ a set of edges, and $N = |\mathcal{N}|$, $E = |\mathcal{E}|$. In the statistical literature, G is often defined in terms of the nodes and the corresponding measurements on pairs of nodes, $G \equiv G(\mathcal{N}, \mathcal{Y})$. $\mathcal{Y}$ is usually represented as a square matrix of size $N \times N$. For instance, $\mathcal{Y}$ may be represented as an adjacency matrix Y with binary elements in a setting where we are only concerned with encoding presence or absence of edges between pairs of nodes. For undirected relations the adjacency matrix is symmetric.

Henceforth we will work with graphs mostly defined in terms of its set of N nodes and its binary adjacency matrix Y containing $\sum_{ij} Y_{ij} = E$ directed edges. Nodes in the network may represent individuals, organizations, or some other kind of unit of study. Edges correspond

to types of links, relationships, or interactions between the units, and they may be directed, as in the Holland-Leinhardt model, or undirected, as in the Erdős-Rényi-Gilbert model.

A note about terminology: in computer science, graphs contain nodes and edges; in social sciences, the corresponding terminology is usually actors and ties. We largely follow the computer science terminology in this review.

3.2 The Erdős-Rényi-Gilbert Random Graph Model

The mathematical biology literature of the 1950s contains a number of papers using what we now know as the network model $G(N, p)$, which for a network of N nodes sets the probability of an edge between each pair of nodes equal to p , independently of the other edges, e.g., see Solomonoff and Rapoport [281] who discuss this model as a description of a neural network. But the formal properties of simple random graph network models are usually traced back to Gilbert [119], who examined $G(N, p)$, and to Erdős and Rényi [93]. The Erdős-Rényi-Gilbert random graph model, $G(N, E)$, describes an undirected graph involving N nodes and a fixed number of edges, E , chosen randomly from the $\binom{N}{2}$ possible edges in the graph; an equivalent interpretation is that all $\binom{N}{E}$ graphs are equally likely.¹ The $G(N, p)$ model has a binomial likelihood where the probability of E edges is

$$\ell(G(N, p) \text{ has } E \text{ edges} \mid p) = p^E (1 - p)^{\binom{N}{2} - E},$$

or, equivalently, in terms of the $N \times N$ binary adjacency matrix Y

$$\ell(Y \mid p) = \prod_{i \neq j} p^{Y_{ij}} (1 - p)^{1 - Y_{ij}}.$$

The likelihood of the $G(N, E)$ model is a hypergeometric distribution and this induces a uniform distribution over the sample space of possible graphs. The $G(N, p)$ model specifies the probability of every edge, p , and controls the expected number of edges, $p \cdot \binom{N}{2}$. The $G(N, E)$ model specifies the number of edges, E , and implies the expected “marginal” probability of every edge, $E / \binom{N}{2}$. The $G(N, p)$ model is more commonly found in modern literature on random graph theory, in part because the independence of edges simplifies analysis [see, e.g., 69; 91].

Erdős and Rényi [94] went on to describe in detail the behavior of $G(N, E)$ as $p = E / \binom{N}{2}$ increases from 0 to 1. In the binomial version the key to asymptotic behavior is the value of $\lambda = pN$. One of the important Erdős-Rényi results is that there is a phase change at $\lambda = 1$, where a giant connected component emerges while the other components remain relatively small and mostly in the form of trees [see 69; 91]. More formally,

- P1. If $\lambda < 1$, then a graph in $G(N, p)$ will have no connected components of size larger than $O(\log N)$, a.s. as $N \rightarrow \infty$.
- P2. If $\lambda = 1$, then a graph in $G(N, p)$ will have a largest component whose size is of $O(N^{2/3})$, a.s. as $n \rightarrow \infty$.

¹Both versions are often referred to as Erdős-Rényi models in the current literature.

- P3. If λ tends to a constant $c > 1$, then a graph in $G(N, p)$ will have a unique “giant” component containing a positive fraction of the nodes, a.s. as $N \rightarrow \infty$. No other component will contain more than $O(\log N)$ nodes, a.s. as $N \rightarrow \infty$.

A summary of a proof using branching processes is given in the appendix of this chapter. Some of the proof concepts will be useful for discussion of exchangeable graph models in [section 3.3](#).

The Erdős-Rényi-Gilbert model has spawned an enormous number of mathematical papers that study and generalize it, e.g., see [\[43\]](#). But few of them are especially relevant for the actual statistical analysis of network data. In essence, the model dictates that every node in a graph has approximately the same number of neighbors. Empirically there are few observed networks with such simple structure, but we still need formal tools for deciding on how poor a fit the model provides for a given observed network, and what kinds of generalized network models appear to be more appropriate. This has led to two separate literatures, one of which has focused on formal statistical properties associated with estimating parameters of network models—the p_1 and exponential random graph models described below—and a second that identifies selected predicted features of models and empirically checks observed networks for those features. The latter is largely associated with papers emanating from statistical physics and computer science, several of which are described in detail in [chapter 4](#).

3.3 The Exchangeable Graph Model

The exchangeable graph model provides the simplest possible extension of the original random graph model by introducing a weak form of dependence among the probability of sampling edges (i.e., exchangeability) that is due to *non-observable* node attributes, in the form of node-specific binary strings. This extension helps focus the analysis, whether empirical or theoretical, on the interplay between connectivity of a graph and its node-specific sources of variability [\[1; 5\]](#).

Consider the following data generating process for an exchangeable graph model, which generates binary observations on pairs of nodes.

1. Sample node-specific K -bit binary strings for each node $n \in \mathcal{N}$

$$\vec{b}_n \sim \text{unif}(\text{vertex set of } K\text{-hypercube}),$$

2. Sample directed edges for all node pairs $n, m \in \mathcal{N} \times \mathcal{N}$

$$Y_{nm} \sim \text{Bern} \left(q(\vec{b}_n, \vec{b}_m) \right),$$

where $\vec{b}_{1:N}$ are K -bit binary strings², and q maps pairs of binary strings into the $[0, 1]$ interval. This generation process induces weakly dependent edges. The edges are conditionally inde-

²Note that the space of K -bit binary strings can be mapped one-to-one to the vertex set of the K -hypercube, i.e., the unit hypercube in K dimensions.

pendent given the binary string representations of the incident nodes. They are *exchangeable* in the sense of De Finetti [82].

From a statistical perspective, the exchangeable graph model we survey here [1; 5] provides perhaps the simplest step-up in complexity from the random graph model [93; 119]. In the data generation process, the bit strings are equally probable but the induced probabilities of observing edges are different. A class of random graphs with such a property has been recently rediscovered and further explored in the mathematics literature, where the class of such graphs is referred to as *inhomogeneous* random graphs [45]. An alternative and arguably more interesting set of specifications can be obtained by imposing dependence among the bits at each node. This can be accomplished by sampling sets of dependent probabilities from a family of distributions on the unit hypercube, $\vec{p}_n \in [0, 1]^K$, and then sampling the bits independently given these dependent probabilities.

1. Sample node-specific K -bit binary strings for each node $n \in \mathcal{N}$

$$\vec{p}_n \sim \text{hypercube}(\vec{\mu}, \sigma, \alpha), \text{ where } \sigma > (K - 1) \cdot \alpha > 0,$$

$$b_{nk} \sim \text{Bern}(p_{nk}), \text{ for } k = 1, \dots, K$$

2. Sample directed edges for all node pairs $n, m \in \mathcal{N} \times \mathcal{N}$

$$Y_{nm} \sim \text{Bern}(q(\vec{b}_n, \vec{b}_m)),$$

In the hypercube distribution³, $\vec{\mu}, \sigma, \alpha$ control the frequency, variability and correlation of the bits within a string, respectively; and q maps binary pairs of strings into the unit interval.

In the exchangeable graph model, the number of bits, K , captures the complexity of the graph. For instance, for $K < N$ the model provides a compression of the graph. For directed graphs the function q is asymmetric in the arguments. The sparsity of the bit strings is controlled by the parameter $\alpha > 0$. A larger value of α leads to larger negative correlation among the bits and thereby a sparser network. In such an exchangeable graph model there are two main sources of variability: (i) the probability of an edge decreases with the number of bits K , as more complexity reduces the chances of an edge, and (ii) the probability of an edge increases with $1/\alpha$, as concentrating density in the corners of the unit K -hypercube improves the chances of an edge. While this model does not quite fit the definition of non-homogeneous models of Bollobás et al. [45], it is tractable enough to allow the analysis of the giant component in (K, α) space, by leveraging the branching process strategy developed by Durrett [91] (see the appendix at the end of the chapter). As in Durrett's analysis, the

³The hypercube distribution can be obtained using a hierarchical construction as follows. Sample $\vec{u} \sim \text{Normal}_k(\vec{\mu}, \Sigma)$, where $u \in \mathbb{R}^k$ and $\Sigma_{ii} = \alpha, \Sigma_{ij} = \beta$ for $i \neq j$. Then define $p_i = (1 + e^{-u_i})^{-1}$ for $i = 1 \dots k$. The resulting density for $\vec{p}$, where $\vec{p} \in [0, 1]^k$ is

$$f_P(\vec{p} \mid \vec{\mu}, \alpha, \beta) = \frac{|2\pi\Sigma|^{-\frac{1}{2}}}{\prod_{j=1}^d p_j(1-p_j)} \exp\left(-\frac{1}{2}(\log(\vec{p}/(1-\vec{p})) - \vec{\mu})' \Sigma^{-1} (\log(\vec{p}/(1-\vec{p})) - \vec{\mu})\right).$$

For more details see [4].

giant component emerges because a number of smaller components must intersect with high probability. In exchangeable graph models however, the giant component has a peculiar structure; connected components are themselves connected to form the giant component as soon as bit strings that match on two bits appear with high probability. Figure 3.1 provides a graphical illustration of this intuition. Nodes that *bridge* two connected components are

Figure 3.1: *Left panel.* An example adjacency matrix that correspond to a fully connected component among 100 nodes. *Right panel.* The clustering coefficient as a function of α on a sequence of graphs with 100 nodes. Here $\sigma = 12$, and $\log(\mu_i) = \frac{1}{K}$ for every $i = 1 \dots K$.

evident in the left panel. Note that there are no nodes that bridge three components, as bit strings that match on three bits is an unlikely event in a graph with 100 nodes.

Given a graph, we can infer the corresponding set of binary strings from data. The likelihood that correspond to an exchangeable graph model is simple to write,

$$\ell(Y|\theta) = \int d \vec{b}_{1:N} \left(\prod_{n,m} \Pr(Y_{n,m}|\vec{b}_n, \vec{b}_m, q) \prod_n \Pr(\vec{b}_n|\theta) \right),$$

where $\theta = (\vec{\mu}, \sigma, \alpha)$ or an appropriate set of parameters. We can apply standard inference techniques [2; 9]. Fitting an exchangeable graph model allows us to assess the complexity of an observed graph, leveraging notions from information theory. For instance, we can use the minimum description length (MDL) principle to decide how many bits we need to explain the observed connectivity patterns with high probability. We can also quantify how much *information* is retained at different bit-lengths, and plot the corresponding information profile for $K < N$ and an entropy histogram for any given value of K .

The exchangeable graph model allows for algorithmic comparison of any set of statistical models that are proposed to summarize an observed graph. As an illustration, consider an observed graph G and two alternative models A and B . Rather than comparing how well models A and B recover the degree distribution of G or other graph statistics, and independently of whether it makes sense to directly compare the two likelihoods of A and B (in fact, these models need not have a likelihood), we can proceed as follows.

1. Given a graph G , fit models $A(\Theta_a)$ and $B(\Theta_b)$ to obtain an estimate of their parameters Θ_a^{Est} and Θ_b^{Est} respectively.
2. Sample M graphs at random from the support of $A(\Theta_a^{Est})$ and $B(\Theta_b^{Est})$.
3. Compute the distributions of summary statistics based on notion from information theory, such as information profile and entropy histogram, corresponding to the $2M$ graphs sampled from A and B .
4. Compare models in terms of the distribution on the statistics above, such as the complexity of the two models' supports and their similarity to the complexity of G .

The exchangeable graph model also allows for evaluation of the distribution of the number of bit strings with I matching bits, for any integer $I < K$. In theory this distribution leads to expectations on the number of nodes that bridge I communities, where the members of each community have only one out of I matching bits. In practice, we may want to specify K in advance so that each bit corresponds to a well defined property. For instance, in applications to biology, nodes may correspond to proteins and the K bits encode presence or absence of specific protein domains. The distribution on the number of I matchings leads to p-values that summarize how unexpected it is to observe binding events among a set of proteins that share a certain combination of domains.

Overall, the exchangeable graph model introduces weak dependences among the edges of a random graph in a controlled fashion, which ultimately lead to a range of more structured connectivity patterns and enable model comparison strategies rooted in notions from information theory. The focus here is not on modeling per se. In fact, the model is kept as simple as possible. Rather, the focus is on modeling as a means to establish a technical link between graph connectivity and node attributes. This technical link is useful to address some of the issues listed in Chapter 5. For more details see [5].

There exist other complex graph models in the network analysis literature that induce exchangeable or partially exchangeable edges. We will discuss latent space models [146; 137] and stochastic blockmodels [236; 7; 9] as examples. These models can all be traced back to an original analysis of multivariate sociometric relations, measurements of relations represented as vectors rather than scalars, that was developed a few decades ago [103]. The difference in these models and the exchangeable graph model lies in the interpretation of the latent variables and in the goal of the analysis. Latent space models interpret the latent variables as latent positions in a social space, and blockmodels interpret the latent membership vectors in terms of functional association or community membership. In the exchangeable graph model, the latent binary strings do not carry semantic meaning, rather they are mathematical artifacts that help to represent a graph and induce an expressive parametric family of distributions [15; 165; 5]. Most importantly, the exchangeable graph model is meant to be a tool to *represent and explore* the space of connectivity patterns in a smooth, principled semi-parametric fashion. In this regard, exchangeable graph models differ substantially from latent space models or stochastic blockmodels.

3.4 The p_1 Model for Social Networks

A conceptually separate thread of research developed in parallel in the statistics and social sciences literature, starting with the introduction of the p_1 model. Consider a directed graph on the set of n nodes. Holland and Leinhardt's p_1 model focuses on dyadic pairings and keeps track of whether node i links to j , j to i , neither, or both. It contains the following parameters:

- θ : a base rate for edge propagation,
- α_i (expansiveness): the effect of an outgoing edge from i ,
- β_j (popularity): the effect of an incoming edge into j ,
- ρ_{ij} (reciprocation/mutuality): the added effect of reciprocated edges.

Let $P(0,0)$ be the probability for the absence of an edge between i and j , $P_{ij}(1,0)$ the probability of i linking to j ("1" indicates the outgoing node of the edge), $P_{ij}(1,1)$ the probability of i linking to j and j linking to i . The p_1 model posits the following probabilities (see [149]):

$$\log P_{ij}(0,0) = \lambda_{ij}, \tag{3.1}$$

$$\log P_{ij}(1,0) = \lambda_{ij} + \alpha_i + \beta_j + \theta, \tag{3.2}$$

$$\log P_{ij}(0,1) = \lambda_{ij} + \alpha_j + \beta_i + \theta, \tag{3.3}$$

$$\log P_{ij}(1,1) = \lambda_{ij} + \alpha_i + \beta_j + \alpha_j + \beta_i + 2\theta + \rho_{ij}. \tag{3.4}$$

In this representation of p_1 , λ_{ij} is a normalizing constant to ensure that the probabilities for each dyad (i,j) add to 1. For our present purposes, assume that the dyad is in one and only one of the four possible states. The reciprocation effect, ρ_{ij} , implies that the odds of observing a mutual dyad, with an edge from node i to node j and one from j to i , is enhanced by a factor of $\exp(\rho_{ij})$ over and above what we would expect if the edges occurred independently of one another.

The problem with this general p_1 representation is that there is a lack of identification of the reciprocation parameters. The following special cases of p_1 are identifiable and of special interest:

1. $\alpha_i = 0$, $\beta_j = 0$, and $\rho_{ij} = 0$. This is basically an Erdős-Rényi-Gilbert model for directed graphs: each directed edge has the same probability of appearance.
2. $\rho_{ij} = 0$, *no reciprocal effect*. This model effectively focuses solely on the degree distributions into and out of nodes.
3. $\rho_{ij} = \rho$, *constant reciprocation*. This was the version of p_1 studied in depth by Holland and Leinhardt using maximum likelihood estimation.

4. $\rho_{ij} = \rho + \rho_i + \rho_j$, *edge-dependent reciprocation*. Fienberg and Wasserman [101, 102] described this model and how to find maximum likelihood estimate for the parameters.

In the constant reciprocation setting, the elevated probability of reciprocal edges does not depend on the dyad, whereas edge-dependent reciprocation dictates multiplicative increases of the reciprocation probability based on node-specific parameters.

The likelihood function for the p_1 model is clearly in exponential family form. For the constant reciprocation version, we have

$$\log Pr_{p_1}(y) \propto y_{++}\theta + \sum_i y_{i+}\alpha_i + \sum_j y_{+j}\beta_j + \sum_{ij} y_{ij}y_{ji}\rho, \quad (3.5)$$

where a “+” denotes summing over the corresponding subscript. The minimal sufficient statistics (MSSs) are y_{i+} , y_{+j} , and $\sum_{ij} y_{ij}y_{ji}$. Then using the usual exponential family theory we know that the likelihood equations are found by setting the MSSs equal to their expectations (cf. [308]). Holland and Leinhardt gave an explicit iterative algorithm for solving these equations with the added constraints that the probabilities for each dyad add to 1.

A major problem with the p_1 and related models, recognized by Holland and Leinhardt, is the lack of standard asymptotics to assist in the development of goodness-of-fit procedures for the model. Since the number of $\{\alpha_i\}$ and $\{\beta_j\}$ increase directly with the number of nodes, we have no consistency results for the maximum likelihood estimates, and no simple way to test for $\rho = 0$, for example. A few ad hoc fixes have been suggested in literature, the most direct of which deals with the problem by setting subsets of the $\{\alpha_i\}$ and $\{\beta_j\}$ equal to one another (see the discussion of blockmodels below) or by considering them as arising from common prior distributions (see, e.g., [311]). Fienberg et al. [104] recently suggested the use of tools from algebraic statistics to find Markov basis generators for the model and the conditional distribution of the data given the MSSs.

Fienberg and Wasserman proposed a slightly different dyad-based data representation for the p_1 model. Conceptually, the dyad considers the two directed measurements together: $\{D_{ij} = (y_{ij}, y_{ji})\}$. In their work, they define

$$x_{ijkl} = \begin{cases} 1 & \text{if } D(y_{ij}, y_{ji}) = (k, l), \\ 0 & \text{otherwise,} \end{cases}$$

where k and l take the values of 1 or 0. This representation converts the dyad $\{D_{ij} = (y_{ij}, y_{ji})\}$ into a 2×2 table with exactly one entry of 1 and the rest 0. Now if we collect the data for the $n(n-1)/2$ dyads together, they form an $n \times n \times 2 \times 2$ incomplete contingency table with “structural” zeros down the diagonal of the $n \times n$ marginal (i.e., no self loops), and “duplicate” data for each dyad above and below the diagonal. In this redundant 4-way table, the model of no second-order interaction corresponds to p_1 with constant reciprocation, and the standard iterative proportional fitting algorithm⁴ can be used to compute the maximum

⁴For details on IPF for contingency tables, see [39; 99]

likelihood estimates. Fienberg et al. [103] show that same type of contingency table representation also works for the correlated p_1 model for multiple relations, and Meyer [213] provides a technical statistical rationale for these contingency table representations.

Holland and Leinhardt analyzed Sampson’s monk dataset (c.f. subsection 2.2.1 and [259]) using the p_1 model. Fienberg et al. [103] analyzed an 8-relation version of the Sampson data (4 positive and 4 negative) using their multiple-relation generalizations of p_1 , but focusing on an aggregation of the 18 monks into the three blocks identified in [322]: a top-esteemed block of 7 monks with an unambivalently positive attitude towards itself, in conflict with a more ambivalent block of 7, and a block of 4 outcasts and waiverers.

3.5 p_2 Models for Social Networks and Their Bayesian Relatives

In the statistical literature, the notion of *fixed effects* typically refers to a set of unknown constant quantities, each of which is used to partly explain the variability of the observations corresponding to a unit of analysis, e.g., an individual or a pair of individuals. This contrasts the notion of *random effects*, which refers to a set of unknown variable quantities that serve a similar purpose and are drawn from the same underlying distribution.

The p_1 model treats expansiveness, $\{\alpha_i\}$, and popularity, $\{\beta_j\}$, as fixed effects associated with unique nodes in the network. Often it makes more sense to think about the ensemble of expansiveness and/or popularity effects as a sample drawn from some underlying distribution, and then estimate the parameters of that distribution. This type of random effects network model has been developed in a series of papers by Snijders and his collaborators and they refer to it as the p_2 network model, e.g., see van Duijn et al. [301]. It is reasonably straightforward to take any of the multivariate variations on p_1 and generate a family of multi-level models with mixtures of fixed and random effects in the spirit of p_2 , e.g., see Zijlstra et al. [333].

Bayesian extensions of frequentist approaches often involve positing a statistical model for fixed effects, thus converting them into random effects. The principal distinction between the p_2 models and Bayesian extensions of p_1 is that, in the latter, the other unknown constant quantities, λ, θ, ρ , may be also converted into random effects. Furthermore, there may be additional levels to the multilevel hierarchy in these models, and there are prior distributions on the parameters at the highest level of the hierarchy (cf. Gill and Swartz [121]; Wang and Wong [311]). It should come as no surprise that authors using the Bayesian approach have worked with Monte Carlo Markov chain (MCMC) methods as have those using versions of p_2 .

MCMC implementations of p_2 models in STOCNET⁵ are well-suited for networks with a relatively large number of nodes, e.g., Zijlstra et al. [333] study network data from 20 Dutch high schools with a total of 1,232 pupils.

⁵STOCNET is a freestanding Software package for the statistical analysis of social networks, available at <http://stat.gamma.rug.nl/stocnet/>.

3.6 Exponential Random Graph Models

Under the assumption that two possible edges are dependent only if they share a common node,⁶ Frank and Strauss [110] proved the following characterization for the probability distribution of undirected Markov graphs:

$$\Pr_{\theta} \{Y = y\} = \exp \left(\sum_{k=1}^{n-1} \theta_k S_k(y) + \tau T(y) + \psi(\theta, \tau) \right) \quad y \in \mathcal{Y}, \quad (3.6)$$

where $\theta := \{\theta_k\}$ and τ are parameters, $\psi(\theta, \tau)$ is the normalizing constant, and the statistics S_k and T are counts of specific structures such as edges, triangles, and k -stars:

$$\begin{aligned} \text{number of edges:} \quad S_1(y) &= \sum_{1 \leq i < j \leq n} y_{ij}, \\ \text{number of } k\text{-stars } (k \geq 2): \quad S_k(y) &= \sum_{1 \leq i \leq n} \binom{y_{i+}}{k}, \\ \text{number of triangles:} \quad T(y) &= \sum_{1 \leq i < j < h \leq n} y_{ij} y_{ih} y_{jh}. \end{aligned}$$

Note that there is a dependence structure to the parameters of this model, with edges being contained in 2-stars, and 2-stars being contained in both triangles and three-stars. Certain variations of this ERGM model that involve directed edges are natural generalizations of the p_1 model. Alternative parameterizations that go beyond Markov graph models have been recently proposed, e.g., see [280; 317; 21].

Frank and Strauss [110] worked mainly with the three parameter model where $\theta_3, \dots, \theta_{n-1} = 0$. They proposed a pseudo-likelihood parameter estimation method [287] that maximizes

$$\ell(\theta) = \sum_{i < j} \log \left(\Pr_{\theta} \{Y_{ij} = y_{ij} \mid Y_{uv} = y_{uv} \text{ for all } u < v, (u, v) \neq (i, j)\} \right).$$

Wasserman and Pattison [316] proposed the current formulation of these *Exponential Random Graph Models* (ERGM), also referred to as p^* models, as a generalization of the Markov graphs of Frank and Strauss. For both directed and undirected graphs, they maintain a similar characterization of the probabilities where the statistics S_k and T are replaced by arbitrary statistics U . This leads to likelihood functions of the form

$$\Pr_{\theta} \{Y = y\} = \exp \left(\theta^{\top} u(y) - \psi(\theta) \right). \quad (3.7)$$

The statistics $u(y)$ are counts of graph structures. Although they are not independent—they count overlapping sets of edges—they are assumed independent in the pseudo-likelihood. Ignoring these correlations is a bad idea; it causes extreme sensitivity of the predicted number of edges to small changes in the value of certain parameters [302]. Park and Newman [240] formally characterized sensitivity issues. Snijders et al. [280] recently proposed a variant of

⁶This is the definition of Markov property for spatial processes on a lattice in [33].

these models where the major problem of double-counting is mitigated but not overcome. Hunter and Handcock [155] estimate likelihood ratios for nearby $\{\theta_i\}$ using a MCMC procedure related to the work of Geyer and Thompson [118]. Their estimation procedure can be used for models based on distributions in the curved exponential family.

Robins et al. [256] describe problems associated with the estimation of parameters in many ERGMs, involving near degeneracies of the likelihood function and thus of methods used to estimate parameters using maximum likelihood. For example, for a certain combination of ERGM statistics, the likelihood function may have multiple, clearly distinct modes, and there are very few network configurations—often radically different from each other—that have non-zero probabilities. This is a topic of current theoretical and empirical investigation rooted in the theory of discrete exponential families [136; 251]. For a discussion of mixing times of MCMC methods for ERGMs and the relevance to convergence and degeneracies, see [35].

There are two carefully constructed packages of routines that are available for analyzing network data using ERGMs: *statnet*⁷ and *SIENA*⁸. These packages focus on the use of MCMC methods for estimating the parameters in ERGMs.

Remark. It is possible to express the current formulation of exponential random graphs using the formalism of undirected graphical models and the Hammersley-Clifford theorem [76; 33]. We can write the likelihood of an arbitrary undirected graph as

$$\Pr(\mathbf{y}|\boldsymbol{\theta}) = \frac{\prod_{c \in \mathcal{C}} \psi(\mathbf{y}_c|\boldsymbol{\theta}_c)}{z}, \quad (3.8)$$

where $\mathbf{y}_c$ denotes the nodes in clique c , $\boldsymbol{\theta}_c$ denotes the corresponding set of parameters, ψ are non-normalized potentials over the cliques, and $z = \sum \prod_{c \in \mathcal{C}} \psi(\mathbf{y}_c|\boldsymbol{\theta}_c)$ is the normalization constant. If the likelihood is in the exponential family, then the log potentials are linear in θ_c and “features” $u(\mathbf{y}_c)$, and we can write:

$$\begin{aligned} \Pr(\mathbf{y}|\boldsymbol{\theta}) &= \exp \left\{ \sum_{c \in \mathcal{C}} \log \psi(\mathbf{y}_c|\boldsymbol{\theta}_c) - \log z \right\} \\ &= \exp \left\{ \sum_{c \in \mathcal{C}} \boldsymbol{\theta}_c^\top u(\mathbf{y}_c) - \log z \right\} \\ &= \exp \left\{ \boldsymbol{\theta}^\top u(\mathbf{y}) - \log z \right\}. \end{aligned}$$

Within the exponential family, the advantage is that computing derivatives and likelihood and deriving the corresponding EM algorithm are feasible, although possibly computationally expensive, by using variational approximation strategies and Monte Carlo methods. A lot of methodology on the subject has been developed in the area of machine learning. There,

⁷A package written for the R statistical environment described at <http://csde.washington.edu/statnet/>. See also the documentation in [138; 157; 224; 129].

⁸Simulation Investigation for Empirical Network Analysis—a freestanding package available at <http://stat.gamma.rug.nl/snijders/siena.html>.

undirected graphs appear primarily in the context of relational learning and imaging. For an in-depth discussion on exact and approximation methods and for references see [247; 308].

3.7 Random Graph Models with Fixed Degree Distribution

The Erdős-Rényi-Gilbert random graph model is fully symmetric and the expected degree (the number of edges associated with a node) is the same for all nodes in the graph, following a binomial distribution. A number of natural extensions of the Erdős-Rényi-Gilbert model result in varying node degrees. For example,

- the preferential attachment model [26] captures the formation of hubs in a graph (see [section 4.1](#));
- the one-parameter “small-world” model [320] interpolates between an ordered finite-dimensional lattice and an Erdős-Rényi-Gilbert random graph in order to produce local clustering and triadic closures (see [section 4.2](#)).

Albert and Barabási [12] describe a number of variants on these themes. Many of the investigators exploring the use of such models often focus on the empirical degree distribution, claiming for example that it follows a power-law in many real world networks (cf. [26; 232; 69; 91]). The papers utilizing these “statistical physics” style models often talk about fixed-degree distributions [e.g., 239], and they either fix the degree-distribution parameters or compute distributions that are conditional on some function of the degree distributions or sequences, such as their expectations (cf. [235; 70]). Software is available to sample from the space of random graphs with a given degree distribution based on Monte Carlo Markov chain methods [42; 138].

There would appear to be a direct link between these ideas and the representation of degree distributions in the family of p_1 models. In the latter, the α_i and β_i parameters represent the out-degree and in-degree for the i th node, and the corresponding sufficient statistics are the empirical values for these. In the statistical literature there is a long tradition of looking at distributions conditional on minimal sufficient statistics, and for network models such a notion was investigated as early as 1975 by Holland and Leinhardt, who looked at the version of p_1 with $\rho = 0$, conditioned on the empirical in-degree and out-degree for all nodes in the network [147]. This allows for the calculation of an exact distribution that is independent of the $\{\alpha_i\}$ and $\{\beta_i\}$ by enumerating all possible adjacency matrices in the reference set with the observed in-degrees and out-degrees. There is the expectation that such an approach could lead to a uniformly most powerful test for $\rho = 0$, but there is no theory to support this expectation as of yet. McDonald, Smith and Forster [211] suggest an iterative approach for such calculations using a Metropolis-Hastings algorithm to generate from the conditional distribution of the triad census given the indegrees, the out-degrees and the number of mutual dyads. In a pair of papers [279; 280], Snijders and colleagues explore such conditioning for maximum likelihood estimation for exponential random graph models,

largely as a mechanism for avoiding the degeneracies and near degeneracies observed when unconditional maximum likelihood is used, cf. [section 3.6](#) and [\[256\]](#). Snijders [\[274\]](#) does something similar for dynamic models for graphs. Roberts [\[252\]](#) suggests an algorithm for the conditional distribution of the p_1 model where $\rho_{ij} = \rho$ given the full set of minimal sufficient statistics, but McDonald et al. [\[211\]](#) offer a counterexample and suggest an alteration of their algorithm to generate the proper exact distribution. Generating such exact distributions is a very tricky matter in discrete exponential families because of the need to utilize appropriate Markov bases, either explicitly as in Diaconis and Sturmfels [\[85\]](#) or implicitly. It is unclear whether the proposals in this literature are in fact reaching all possible tables associated with the distribution.

Blitzstein and Diaconis [\[42\]](#) explore different efficient mechanisms for generating random graphs with fixed degree sequence and explicitly make the link between the “statistical physics” and “sociological” literatures, whereas the earlier papers by Newman [\[232\]](#) and Park and Newman [\[239\]](#) reference exponential random graphs but only approach the notion of fixed degree distributions from a statistical physics perspective, focusing on characteristics of network ensembles rather than maximum likelihood estimation and assessment of goodness-of-fit.

3.8 Blockmodels, Stochastic Blockmodels and Community Discovery

A problem which has been a focus of attention for at least 40 years in the network literature has been the search for an “optimal partition” of the nodes into groups or blocks. In the sociometric literature this was known as blockmodeling. A formalization of networks in terms of non-stochastic blocks goes back at least as far as Lorrain and White [\[199\]](#). Their paper and the discussion of structural equivalence gave rise to innumerable papers in mathematical sociology, (see, e.g., [\[53\]](#)) and algorithmic search strategies for determining blocks (see, e.g., [\[19; 88; 89\]](#)). By embedding these ideas within a framework of random graphs, Holland et al. [\[150\]](#) explained how a special version of p_1 could be used to describe a random graph model with predefined blocks. (See also the related discussion in [\[103\]](#) and [\[311\]](#).)

A true stochastic blockmodel approach, however, involves the *discovery* of the block structure as part of the model search strategy [\[314\]](#), and the first attempts at doing this within the framework of p_1 and its exponential family generalizations was due to Nowicki and Snijders, who focused on technical issues such as non-identifiability in a restricted version of the blockmodel [\[277; 236; 237; 79\]](#). A comprehensive statistical treatment of these models was recently developed for analyzing protein interaction data [\[7; 8\]](#) and then further developed in the context of social network data [\[9\]](#). Handcock et al. [\[137\]](#) approach this stochastic blockmodeling problem through a combination of latent space models and traditional clustering. We describe some of this work in more detail below.

More recently in the statistical physics and computer science literatures the problem has gone under the label of detection of community structure, e.g., see [\[122; 232; 71; 233\]](#);

266; 217]. This literature is now voluminous and seemingly unconnected to the statistical blockmodel work.

The basic idea, in both the model-based and algorithmic approaches as well as the community detection literature, is that nodes that are heavily interconnected should form a block or community. The nodes are reordered to display the blocks down the diagonal of the adjacency matrix representing the network. Moreover, the connections between nodes in different blocks appear in much sparser off-diagonal blocks. In model-based approaches, the partition of the nodes maximizes a statistical criterion linked to the model, e.g., a likelihood function, whereas most algorithmic solutions maximize ad hoc criteria related to the “density” of links within and between blocks.

More formally, a blockmodel is a model of network data that relies on the intuitive notion of *structural equivalence*: two nodes are defined to be structurally equivalent if their connectivity with similar nodes is similar—this is a “soft” definition.⁹ Following up this idea, we can imagine collapsing structurally equivalent nodes together to form a super-node, or a block in the language of blockmodels. Keeping the notion of a block in mind we can now revisit and sharpen the definition of structurally equivalent nodes: given N nodes and K blocks, let $Y_{N \times N}$ be the adjacency matrix of the graph $G(\mathcal{N}, \mathcal{Y})$, then two nodes a and b are structurally equivalent, and thus belong to the same block h , if their connectivity patterns $\mathcal{C}_a$ and $\mathcal{C}_b$ with nodes in other blocks are similar. The equivalence between connectivity patterns of nodes a and b can be formally stated as follows:

$$\mathcal{C}_a \equiv \{ Y(a, i \in h_k) : \forall h_k \neq h \} \approx \mathcal{C}_b,$$

where the index i runs over the nodes other than a, b , the index k runs over the blocks other than h , h_k is the set of nodes in block k , and $\approx$ quantifies similarity according to a suitable distance metric. This definition relies on a pre-specified partitioning of the N nodes into K blocks. A blockmodel is useful, for instance, in the analysis of social relations where blocks may correspond to social factions, as well as in the analysis of protein interactions where blocks may correspond to stable protein complexes.

Collapsing nodes into blocks by leveraging the notion of structural equivalence above is a more general task than clustering. Consider, for example, the green nodes 7–9 in the left panel of [Figure 3.2](#). They are structurally equivalent according to the definition above, as $\mathcal{C}_7 \approx \mathcal{C}_8 \approx \mathcal{C}_9$, although there are no direct connections among the nodes 7–9 themselves. In this sense, nodes 7–9 would not represent a tight cluster according to measures of similarity based on direct connectivity. Blocks that would correspond to clusters can be obtained by pre-specifying an identity blockmodel, $B = \mathbb{I}_K$, in which all off-diagonal blocks equal zero and all diagonal blocks equal one.

At the technical level, we need two sets of parameters in order to instantiate a blockmodel: (i) the blockmodel itself, B , is a $K \times K$ matrix, in which the $B(g, h)$ entry specifies, for instance, the average probability that nodes in block g have connections directed to nodes in block h , and (ii) a mapping between nodes and blocks, $\vec{\pi}_{1:N} = \Pi$, where the node-specific array summarizes some notion of membership. Airolti et al. [9], for instance, specify

⁹The term *stochastic equivalence* is often used in place of structural equivalence, e.g., see [315].

Figure 3.2: Left: An example graph. Right: The corresponding blockmodel, where red nodes have been collapsed into the red block and similarly for the other colors. Note that this problem is not a typical clustering problem, as the green nodes do not share any direct connections; each green node, however, has connections directed to blue nodes, and connections directed from red nodes. In other words, given the partition into monochromatic blocks, the nodes in the green block share patterns of connectivity to nodes in other blocks.

the mapping in terms of mixed membership arrays, in which $\pi_n(h)$ specifies the relative frequency of interactions. Node n participates in $2N-2$ interactions in total and instantiates connectivity patterns that are typical of nodes in its block. These two sets of parameters, B and Π , are two latent sources of variability that compete to explain the observed connectivity. However, the blockmodel B explains global asymmetric block connectivity patterns, while the (mixed) membership mapping Π explains node-specific symmetric connectivity patterns. In this sense, instantiating a blockmodel in terms of B and Π does not introduce any source of non-identifiability beyond the usual multiplicity of parametric configurations that lead to exactly the same likelihood—well characterized in this model by Nowicki and Snijders [236].

As a concrete example, consider the mixed membership stochastic blockmodel (MMB) introduced by [9]; the data generating process for a graph $G = (\mathcal{N}, Y)$ is the following.

1. For each node $p \in \mathcal{N}$:
 - 1.1 Sample mixed membership $\vec{\pi}_p \sim \text{Dirichlet}_K(\vec{\alpha})$.
2. For each pair of nodes $(p, q) \in \mathcal{N} \times \mathcal{N}$:
 - 2.1 Sample membership indicator, $\vec{z}_{p \rightarrow q} \sim \text{mult}_K(\vec{\pi}_p)$.
 - 2.2 Sample membership indicator, $\vec{z}_{p \leftarrow q} \sim \text{mult}_K(\vec{\pi}_q)$.
 - 2.3 Sample interaction, $Y(p, q) \sim \text{Bern}(\vec{z}_{p \rightarrow q}^\top B \vec{z}_{p \leftarrow q})$.

Note that the group membership of each node is *context dependent*. That is, each node may assume different membership when interacting or being interacted with by different peers. Statistically, each node is an admixture of group-specific interactions. The two sets of latent group indicators are denoted by $\{\vec{z}_{p \rightarrow q} : p, q \in \mathcal{N}\} =: Z_{\rightarrow}$ and $\{\vec{z}_{p \leftarrow q} : p, q \in \mathcal{N}\} =: Z_{\leftarrow}$.

Also note that the pairs of group memberships that underlie interactions need not be equal; this fact is useful for characterizing asymmetric interaction networks. Equality may be enforced when modeling symmetric interactions.

Inference in the blockmodel is challenging, as the integrals that need to be solved to compute the likelihood cannot be evaluated analytically. For simplicity, the likelihood is

$$\ell(Y \mid \vec{\alpha}, B) = \int_{\Pi} \int_Z \Pr(Y \mid Z, B) \Pr(Z \mid \Pi) \Pr(\Pi \mid \vec{\alpha}) dZ d\Pi.$$

While the inner integral is easily solvable¹⁰, the outer integral is not. Exact inference is thus not an option. To complicate things, the number of observations scales as the square of the number of nodes, $O(N^2)$. Sampling algorithms such as Monte Carlo Markov chains are typically too slow for real-size problems in the natural, social, and computational sciences. Airolidi et al. [9] suggest a nested variational inference strategy to approximate the posterior distribution on the latent variables, (Π, Z) . (Variational methods scale to large problems without loosing much in terms of accuracy [3; 49; 308].)

Bickel and Chen [37], the most recent contribution to this literature, brings new twists to the model-based approach of community discovery. They use a blockmodel to formalize a given network in terms of its community structure. The main result of this work implies that community detection algorithms based on the modularity score of Newman and Girvan [122] are (asymptotically) biased. It shows that using modularity scores can lead to the discovery of an incorrect community structure even in the favorable case of large graphs, where communities are substantial in size and composed of many individuals. This work also proves that blockmodels and the corresponding likelihood-based algorithms are (asymptotically) unbiased and lead to the discovery of the correct community structure. The proof relies on the exchangeability results developed in the statistics community [15; 165] applied to paired measurements [84].

3.9 Latent Space Models

The intuition at the core of latent space models is that each node $i \in \mathcal{N}$ can be represented as a point z_i in a “low dimensional” space, say $\mathbb{R}^k$. The existence of an edge in the adjacency matrix, $Y(i, j) = 1$, is determined by the distance among the corresponding pair of nodes in the low dimensional space, $d(z_i, z_j)$, and by the values of a number of covariates measured on each node individually. The latent space model was first introduced by Hoff et al. [146] with applications to social network analysis, and has been recently extended in a number of directions to include treatment of transitivity, homophily on node-specific attributes, clustering, and heterogeneity of nodes [144; 137; 183].

¹⁰The inner integral resolves into a series of sums, each one over the support of an individual $\vec{z}$ variable. The support is the same for all such $\vec{z}$ variables, and it is given by the N vertices of the K -dimensional unit hypercube. In other words, the inner integral is a series of sums, each over the same N elements.

The conditional probability model for the adjacency matrix Y is

$$\Pr(Y \mid Z, X, \Theta) = \prod_{i \neq j} \Pr(Y(i, j) \mid Z_i, Z_j, X_{ij}, \Theta),$$

where X are covariates, Θ are parameters, and Z are the positions of nodes in the low dimensional latent space. Each relationship $Y(i, j)$ is sampled from a Bernoulli distribution whose natural parameter depends on Z_i, Z_j, X_{ij} and Θ . In their model, Hoff et al. [146] generated the paired observations $Y(i, j)$ starting from the relevant pair of node representations, (Z_i, Z_j) , through a distance model, pair specific covariates X_{ij} , and parameters $\Theta = (\alpha, \beta)$. The log-odds ratio is then:

$$\log \frac{\Pr(Y(i, j) = 1)}{1 - \Pr(Y(i, j) = 1)} = \alpha + \beta' X_{ij} - |Z_i - Z_j| \equiv \eta_{ij},$$

and the corresponding log likelihood is

$$\log \Pr(Y \mid \eta) = \sum_{n \neq m} \left(\eta_{ij} \cdot Y_{ij} - \log(1 + e^{\eta_{ij}}) \right).$$

One can easily extend the latent space modeling approach to weighted networks. In the general case, paired observations Y may be modeled using a generalized linear model that makes use of $Z_{1:N}$, X_{ij} , and Θ . Following the formalism in [210], a generalized linear model that generates the observed edge weights can be specified in terms of three quantitative elements:

- i. the error model $\Pr(Y_{ij})$, i.e., the model for the observed edge weights with mean $\mu_{ij} = \mathbb{E}[y_{ij}]$;
- ii. the linear model $\eta_{ij} = \eta_{ij}(\beta, Z_i, Z_j)$;
- iii. the link function $g(\mu_{ij}) = \eta_{ij}$, which maps the support of μ_{ij} to that of η_{ij} —typically $\mathbb{R}$.

For example, in the binary graph, the error model is $\Pr(Y_{ij}) = \text{Bern}(\mu_{ij})$, where $\mu_{ij} \in [0, 1]$ for all node pairs $(i, j) \subseteq \mathcal{N}$; the linear model is $\eta_{ij} = \beta + d(Z_i, Z_j)$; the link function is $g(\mu_{ij}) = \log \left(\frac{\mu_{ij}}{1 - \mu_{ij}} \right)$, with its inverse being $\mu_{ij} = \frac{1}{1 + \exp(-\eta_{ij})}$ [146]. In a graph with non-negative, integer edge weights, we can posit $\Pr(Y_{ij}) = \text{Poi}(\mu_{ij})$, where $\mu_{ij} \in \mathbb{R}_+$ for all node pairs $(i, j) \in \mathcal{N}$; the linear model is $\eta_{ij} = \beta + d(Z_i, Z_j)$, the same as in the previous example; the link function is $g(\mu_{ij}) = \log(\mu_{ij})$, and its inverse is $\mu_{ij} = e^{\eta_{ij}}$.

In the general case, the generalized linear model for η_{ij} may also include an explicit distance model d in the latent space $\mathcal{Z}$:

$$\begin{aligned} \eta_{ij} &= \eta_{ij}(\beta, Z_i, Z_j) \\ &= \eta_{ij}(\beta, d(Z_i, Z_j)). \end{aligned}$$

Note that it is possible to re-parametrize $Z_i = \rho_i \omega_i$ to separate the position in a latent reference space, Ω , from its magnitude, ρ_i , a scalar. It is a simple intuition that suggests the use of an explicit distance model in the latent space. In a binary graph, for example, edges are more likely to be generated between pairs of nodes whose representations in the latent space are close. A popular choice of distance measures is Euclidean distance. Estimation can be done via MCMC sampling.

Inference in latent space models has been carried out via Monte Carlo Markov chain in networks with up to several thousand nodes [130]. Scalability issues remain to be addressed before larger networks can be analyzed.

3.9.1 Comparison with Stochastic Blockmodels

The latent space model of Hoff et al. [146] projects nodes onto a latent Euclidean space by inverting the logistic link. While in practice there is often interest in identifying groups of similar nodes, e.g. individuals or proteins, there is no explicit clustering model in the latent space. To identify groups of similar nodes, clustering methods must be used to analyze the set of latent positions inferred by the latent space model. To allow joint inference on latent positions and clusters, Handcock et al. [137] introduce an explicit clustering model in the latent space in the form of a mixture of (spherical) Gaussians.

$$\begin{cases} \Pr(Y \mid Z, X, \Theta) &= \prod_{i \neq j} \Pr(Y(i, j) \mid Z_i, Z_j, X_{ij}, \Theta), \\ Z_i &\sim \sum_k N(\mu_k, \sigma_k^2 \cdot I). \end{cases}$$

This model combines the original latent space model [146] with a finite mixture of Gaussians approach to clustering [297; 205]. It posits that the latent positions $Z_i \in \mathbb{R}^d$ come from a k -dimensional mixture model.

This extension is related to the stochastic blockmodel of [9], which posits a latent membership vector for each node. These vectors can be viewed as cluster assignment probabilities for each node. The observed binary relationships between nodes are mediated by per-pair latent variables, each drawn conditioned on a node’s mixed membership vector. In its general form, the blockmodel allows for multiple relations and covariates. Similarly, the model in [137] is also a hierarchical model, as a Gaussian distribution is placed on the latent positions Z_i . In contrast, however, each node belongs to a single cluster and the corresponding partition governs the observed relationships. There can be variance in the latent position variables, but the idea of belonging to two or more groups cannot be represented. Posterior uncertainty about cluster membership is different from having an explicit distribution that controls mixed membership, which carries with it an additional level of uncertainty. With that said, the latent space in which nodes are projected in [137] is somewhat comparable to the space of cluster proportions in [9]. The former maps nodes to a Euclidean space, while the latter maps nodes to the simplex.

Both models share the same goal: inferring latent structure that explains the variability of the connectivity in an observed network. In the mixed membership model, full MCMC for any but the simplest problems is unreasonably expensive. Airolidi et al. [9] appeal to

variational methods for a computationally efficient approximation to the posterior. These methods can scale to large matrices (e.g., millions of nodes) because of the simplified approximation, but at an unknown cost to accuracy. It would be interesting to explore computational tradeoffs for the latent space cluster model [137] as the sample size grows and when large numbers of covariates are added.

Remark. Blei and Fienberg [40] argue that a stochastic blockmodel and node-specific mixed membership vectors are two sets of parameters that are directly interpretable in terms of notions and concepts relevant to social scientists, and better suited to assist these scientists in extracting substantive knowledge from noisy data, to ultimately inform or support the development of new hypotheses and theories.

Applying the mixed membership stochastic blockmodel (MMB) to Sampsons data demonstrates both similarities and differences [2]. For instance, BIC suggests the existence of three factions among the 18 monks when fitting the MMB, but the groupings differ from those found by the latent space cluster model. One major benefit of applying mixed membership model to the data is the ability to quantitatively identify two out of three of the novices that Sampson labeled as *waverers* in his analysis based on anthropological observations. This could lead to the formation of a social theory of failure in isolated communities, with a possibility to be confirmed with real longitudinal data [2].

In the sociology literature, certain specifications of blockmodels are referred to as *latent class* models, and certain specifications of latent space models are referred to as *latent distance* models. Hoff [145] provides a nice comparison, both theoretical and empirical, of these two types of models with the *eigenmodel*. The eigenmodel is based on a singular value decomposition of the socio-matrix, it can capture more connectivity patterns than the latent class and the latent distance models, for a *given degree* of model complexity, which can be the number of classes, the number of dimensions in the latent space, or the number of eigenvectors. There is a price to pay, however. The eigenmodel is the least amenable to interpretation among the three models, as the inferred patterns that capture connectivity are in terms of eigenvectors. The latent space model can be interpreted in terms of distances. The latent class models can be interpreted in terms of blocks of connectivity, or tight micro-communities; this is the easiest model to interpret.

Appendix: Phase Transition Behavior of the Erdős-Rényi-Gilbert Model

A simple way to analyze the phase transition behavior of Erdős-Rényi-Gilbert models at $\lambda = 1$ is to study the emergence of the giant component as a branching process [91]. Intuitively, consider branching processes that start at every node: for certain values of λ all the branching processes will keep growing with high probability. Their supports, i.e., the sets of nodes involved in each process, will intersect with high probability, leading to the emergence of the

giant component, G , in which each node can be reached from every other node.

The following formal argument comes from lecture notes by Guetz and Constantine [133] based on proofs given by Janson et al. [161]. Pick a node $v \in \mathcal{N}$. If v is connected to all of the nodes in G , then we say that v is *saturated* in G . Now work as follows: pick a node v and place it on the list. Then, identify all its neighbors in G , and add them to the list. Next, take the first unsaturated node on the list and add to the list all of its neighbors which are not already in it. The proof is constructed by considering the distribution of the number of nodes an unsaturated node adds to the list and by using Chernoff bounds to bound the size of the connected component each node belongs to. For details on this proof please see [43]. Bollobás et al. [45] carried out an extensive analysis of the phase transition that mathematically characterizes emergence of the giant component in inhomogeneous random graphs.

Chapter 4

Dynamic Models for Longitudinal Data

In [chapter 3](#) we focused on models for static networks, that consider a cross-section of a real network at a given point in time. However, real networks often contain a dynamic component. In the language of networks, dynamics can be translated into the birth and death of edges and nodes. For example, in a friendship network, new nodes may be introduced at any time and old nodes may drop out due to inactivity; links of friendships and alliances may be even more brittle. Dynamic network modeling has been a neglected sibling of static network modeling, partly due to the added complexity and partly due to a lack of datasets to study. Sampson’s monastery study [\[259\]](#) produced one of the earliest datasets with information on the dynamics in the network of the 18 initiates. The original research, however, focused on the network structure at each given time point, rather than modeling the underlying dynamics explicitly. As online communities gain in popularity, we are beginning to get access to an increasing number of dynamic network datasets of much larger size and longer time span. At the same time, advances in statistical and computational methods for inference and learning have enabled development of richer models. Bearing this in mind, in this chapter we consider three different classes of models. We begin by revisiting the Erdős-Rényi-Gilbert random graph model and its generalizations, viewing them as models for dynamic processes. Then we turn to continuous time Markov process models (CMPPM) and their discrete time cousins (such as a dynamic version of ERGM and other recently proposed models).

4.1 Random Graphs and the Preferential Attachment Model

Many variations on the classical Erdős-Rényi-Gilbert random graph model in [section 3.2](#) are typically considered to be static models, in that they model a single, static snapshot of the network, as opposed to multiple snapshots recorded at different time steps. However, they also contain processes for link addition and modification, which is a dynamic process that may have generated the observed graph, though there is no attempt to fit these dynamic model

properties to observed data. For this reason, we view them as “pseudo-dynamic” models and discuss three examples here: the Erdős-Rényi-Gilbert model, preferential attachment model, and small-world models.

For example, we can view the Erdős-Rényi-Gilbert model $G(N, E)$, itself as a dynamic process used to generate a random graph:

- start from the graph of N unconnected nodes at time 0;
- at each subsequent time step, add a different edge to the network with probability $p = E/\binom{N}{2}$.

By convention, we usually fix the number of nodes at N , although we can extend the process to allow for addition of nodes. This model assumes that edges (and nodes) are not removed once they are added. The degree distribution for $G(N, E)$ is binomial. But as N gets large, Np tends to a constant, so it is approximately Poisson. Durrett [91] provides a rich discussion for situating this dynamic description with the tradition of discrete time random walks and branching processes. In particular, he uses this representation to explore the emergence of the giant component described in [section 3.2](#) (see appendix of [chapter 3](#)).

The Erdős-Rényi-Gilbert model is simple and easy to study but does not address many issues present in real network dynamics. One of the major criticisms [26] of this model centers on the fact that it does not produce a scale-free network, i.e., the resulting node degree distribution does not follow a power law. The network literature is replete with claims that many real networks exhibit the power-law phenomenon, (cf. [12]), and much subsequent research has focused on how various generalizations of the Erdős-Rényi-Gilbert model conform to the power law degree distribution. Molloy and Reed [219] were the first to describe how to construct graphs with a general degree distribution and they went on to describe the emergence of the giant component in that context as well [220].

Barabási and Albert [26] described a dynamic preferential attachment (PA) model specifically designed to generate scale-free networks. At time 0, the model starts out with N_0 unconnected nodes. At each subsequent time step, a new node is added with $m \leq N_0$ edges. The probability that the new node is connected to an existing node is proportional to the degree of the latter. In other words, the new node picks m nodes out of the existing network according to the multinomial distribution

$$p_i = \frac{\delta_i}{\sum_j \delta_j},$$

where δ_i denotes the (undirected) degree of node i . This model, which was described much earlier in the statistical literature by Yule [329] and Simon [269], is intended to describe networks that grow from a small nucleus of nodes and follow a “rich-get-richer” scheme. The assumption is that, for instance, a new web page will more likely link via a URL to a well-known web page as opposed to a little-known one. Mitzenmacher [218] gives a brief history of generative models for power law distributions.

The preferential attachment model of Barabási and Albert results in a network with a power law degree distribution whose exponent is empirically determined to be $\gamma_{BA} =$

2.9 ± 0.1 , whereas the Erdős-Rényi-Gilbert model has a Poisson degree distribution. Many extensions of the model have been proposed that allow for flexible power-law exponents, edge modifications, non-uniform dependence on the node degree distributions, etc. For example, Dorogovtsev and Mendes [90] proposed that creating an edge to node i should be proportional not just to its degree k_i but also to its age, decaying as $(t - t_i)^{-\nu}$, where ν is a tunable parameter. This leads to a power law degree distribution only if $\nu < 1$. Barabási et al. [28] and Durrett [91] provide an account of this and other extensions to the original model of Albert and Barabási. Alternative graph generation mechanisms appear every day—R-MAT [60], ‘winners don’t take all’ [242], ‘forest fire’ [194], ‘butterfly’ [212] and RTG [10], to name a few. The latest, RTG model, proves conformance to 11 empirical laws observed in real networks. The main goal of these random graph models is to describe a process that could generate networks emulating certain known network properties. The generative process could then give an insight into the dynamics that led to the observed network. But these models are often applied to network data are gathered at a few points in time (sometimes only once). Thus the networks are often examined statically.

It has been recently pointed out that, contrary to previous claims, the empirical laws that generative models aim to emulate are not always supported by real data. Visual comparison are not sufficient for determining the goodness of fit of a model. For example, a lot of attention has recently been paid to the degree distribution. Figure 4.1 shows indegree and outdegree distributions for blog and query databases from an unnamed large company. They are plotted on a log-log scale and the downward slopes, if fitted by straight lines, would be visually similar to power law distributions with exponents less than 2. A careful examination of these plots, however, reveals a curvilinear relationship in all cases, which suggests that there is a different generating process than those usually used to justify power laws in empirical network data. Data such as that displayed in Figure 4.1 are often fitted by ordinary least squares or even by eye; often the claim is that a degree distribution is scale-free except for a cutoff at very high or very low degrees, without any adjustment for searching for a cutoff! There have been a number of recent efforts to assess the fit of degree distributions, such as those associated with power laws in log-log plots, with more rigor, e.g., see [74]. As in this example, results from such careful assessments of fit often contradict the assumption of linearity.

Li et al. [196] give a “structural metric” for examining simple connected graphs having identical degree distributions and derive theoretical properties of scale-free graphs. They provide at least one possible way to assess whether a graph corresponding to a network is in fact scale-free. For more informal discussions related to this theoretical work, see [14; 324]. Flaxman et al. [106; 107] describe a class of network models linked to the preferential attachment model that also yield a power-law degree distribution.

Most descriptions of generative models fall short of studying the full parameter space and do not propose procedures for fitting the proposed methods to real data, though there are a few works that suggest maximum likelihood, MCMC and other frameworks for fitting these models to data (for e.g. [34; 75; 214; 323]). One of the notable exceptions is work based on Kronecker graph multiplication. What started as yet another generative procedure [192]

Figure 4.1: Log-log plots of degree distributions for a query data bases and a blog data base from a company database. Left: Blog indegree and outdegree distributions. Right: Query indegree and outdegree distributions. Source: Data from an unnamed large company, stored in iLab, Carnegie Mellon University.

has turned into a well analyzed methodology [195] with an efficient algorithm for model fitting, analysis of the parameter space, and model selection. This work goes further in understanding real network structure and provides a way for principled graph sampling.

4.2 Small-World Models

Watts and Strogatz [320] proposed a small-world model which can be thought of as a “pseudo-dynamic” model in the sense we described in [section 4.1](#). This one-parameter “small-world” model interpolates between an ordered finite-dimensional lattice and an Erdős-Rényi-Gilbert random graph in order to produce local clustering and triadic closures. Bollobás and Chung [44] had previously noted that adding random edges to a ring of N nodes drastically reduces the diameter of the network. The Watts-Strogatz model begins with a ring lattice with N nodes and k edges per node, and randomly rewires each edge with probability p . As p goes from 0 to 1, the construction moves toward an Erdős-Rényi-Gilbert model. They and others who followed, studied the behavior of such small-world networks when $0 < p < 1$. This model is not dynamic although it is often used to describe networks that evolve over time. [Figure 4.2](#) shows a small-world graph for $n = 25$ nodes and 2 rewirings per node.

Kleinberg [174] introduced a variation on the small-world model where random edges are added to a fixed grid. Starting with an underlying finite-dimensional grid, he added shortcut edges, where the probability that two nodes are connected by a long edge depends on the distance between them in the grid. More precisely, the probability that two non-adjacent nodes x and y are connected is proportional to $d(x, y)^{-\alpha}$. With α set to the dimension of the lattice, the greedy routing algorithm can find paths from one node to another in a polylogarithmic number of expected steps.

Figure 4.2: Small-world graph for $N = 25$ nodes and 2 rewirings per node. The red edges form the ring lattice and the blue edges the rewiring. This graph was generated using the Java applet at <http://cs.gmu.edu/~astavrou/smallworld.html>

Several follow-up works have made adjustments to Kleinberg's rewiring procedure in attempt to improve the understanding and efficiency of the navigability of networks. For example, Clauset and Moore [72] suggested to rewire a long distance edge from node x , if while performing a greedy walk over to y , the original topology of the network did not allow to reach y within T_{thresh} steps. The edge was rewired to the place where the search gave up (the node reached after T_{thresh} steps of the walk). They show that through this rewiring procedure the network degree distribution converges to a power law, where $\alpha = \alpha_{\text{rewired}}$. Their work also studied finite size effects and showed that $\alpha_{\text{opt}} \rightarrow d$, as $n \rightarrow \infty$ rather slowly.

Sandberg [260, 261] and Sandberg and Clarke [262] introduced a different rewiring scheme with the end goal to make the network more amenable to statistical analysis. Starting with N nodes on a ring, each with two neighbor links and a long range link, the model of Sandberg [260] randomly rewires a graph in the following steps:

- at each time step $j = 1, 2, 3, \dots$, choose a random starting node x and a target node y and perform greedy routing from x to y ;
- independently and with (small) probability x , update the long-range link of each node on the resulting path to point to y .

This defines a Markov chain on a collection of labeled graphs. Sandberg and Clarke [262] conjecture that when the chain achieves stationarity, the distribution of distances spanned by long-range links is (close to) theoretical optimum for search and the expected length of searches is polylogarithmic. They support the conjecture by a series of simulations. This methodology has been applied to the study of peer-to-peer (P2P) networks.

Durrett [91] discusses links between small-world models and stochastic processes. Typical usage of small-world models include empirical analyses involving aggregate summary statistics (see, e.g., [18; 231]). There are as yet no formal statistical methods for examining the evolution of small-world network models and for assessing their fit to network data measured over time.

4.3 Duplication-Attachment Models

Duplication-Attachment models were originally developed in the computer science theory community to study the world wide web as a directed graph [175; 185]. These models aim at describing properties of a snapshot of the web graph at a specific time, that is, a static directed graph. The data generating process underlying these models, however, is explicitly dynamic. The following example demonstrates some basic assumptions behind the dynamics. Consider a newly added web page A , which provides a new node in the web graph. The creator of web page A will then add *hyper-links* to it, which provide new directed edges in the web graph. In particular, *some* of these hyper-links will point to other web pages regardless of whether their topical content matches the topical content of web page A , but *most* of these hyper-links will point to web pages with a topical content that closely matches the topical content of web page A .

Technically, there are many possible specifications and variants. The basic duplication-attachment model proposed and analyzed by Kumar et al. [185] is as follows. Denote the graph at time t as $G_t = (\mathcal{N}_t, \mathcal{E}_t)$. At each step, say $t + 1$, one new node N is added to G_t . The new node is connected to a *prototype* node m , chosen uniformly at random among those in $\mathcal{N}_t$. Then d out-links are added to node N . The i th out-link is chosen as follows: with probability α the destination node is chosen uniformly at random among those in $\mathcal{N}_t$, and with probability $1 - \alpha$ the destination node is taken to be the i th out-link of the prototype node m . Note that this is possible since the algorithm generates a constant degree graph. Rather than proposing estimation strategies for the two parameters (α, d) of this particular duplication-attachment model, the goal of the analysis of Kumar et al. [185] is on deriving results about topological properties of duplication-attachment graphs, described as functions of the two parameters (α, d) . Recent extensions of this model include a model where fractions of both out-links and in-links of the prototype node m are *copied* by the newly added node N [193]. The goal of the analyses in this line of research, however, remains that of replicating properties of observed graphs, with a few exceptions. In the biological context, duplication-attachment models have appeared to be useful in modeling protein-protein interaction networks. For example, Ratmann et al. [245] proposed a mixture of preferential attachment and duplication divergence with parent-child attachment model to assess evo-

lutionary dynamics of protein interaction networks of *H. pylori* and *P. falciparum*. They proposed a likelihood-free MCMC-based routine to estimate posterior of network summary statistics. A more general review of work in modeling dynamics (evolution) on the basis of protein-protein interaction data is available in [246].

Wiuf et al. [326] have developed a recursive construction of the likelihood for duplication-attachment models, effectively enabling principled statistical data analysis, estimation and inference.

4.4 Continuous Time Markov Chain Models

The use of continuous Markov processes to model dynamic networks was first proposed by Holland and Leinhardt [148] and Wasserman [312] and most recently studied by Snijders and colleagues [275; 276]. As shall become clear in this section, continuous Markov process models (CMPM) are intimately tied to the ERGM models described in [section 3.6](#). Within the CMPM family, network edges are taken to be binary (either absent or present, but not weighted), and the evolution occurs one edge at a time. Model variants arise due to the many possible specifications of edge change probability. Some exceptions to this general approach include the party model of Mayer [206], where multiple edges are allowed to change at the same time, and the work of Koskinen and Snijders [179], which deals with Bayesian parameter inference methods for the case where not all edge modifications are observed.

We begin by providing a quick reminder of continuous Markov processes, borrowing notation from [275]. Define $\{Y(t) \mid t \in \mathcal{T}\}$ to be a stochastic process, where $Y(t)$ has a finite outcome space $\mathcal{Y}$ and $\mathcal{T}$ is a continuous time interval. Suppose that a Markov condition holds: for any possible outcome $\tilde{y} \in \mathcal{Y}$ and any pair of time points $\{t_a < t_b \mid t_a, t_b \in \mathcal{T}\}$,

$$\Pr\{Y(t_b) = \tilde{y} \mid Y(t) = y(t), \forall t : t \leq t_a\} = \Pr\{Y(t_b) = \tilde{y} \mid Y(t_a) = y(t_a)\}. \quad (4.1)$$

In other words, supposing that t_b denotes the future and t_a the present, then conditioning on the past is equivalent to conditioning on the present when it comes to determining the future. If the probability in [Equation 4.1](#) depends only on $t_b - t_a$, then one can prove that $Y(t)$ has a stationary transition distribution, and the transition matrix

$$\Pr(t_b - t_a) := \left[\Pr\{Y(t_b) = \tilde{y} \mid Y(t_a) = y\} \right]_{y, \tilde{y} \in \mathcal{Y}} \quad (4.2)$$

can be written as a matrix exponential

$$\Pr(t) = e^{tQ}, \quad (4.3)$$

where Q is known as the *intensity matrix* with elements $q(y, \tilde{y})$. The elements $q(y, \tilde{y})$ can be thought of as the slope (rate of change) of the probability of state change as a function of time, i.e., $\Pr\{Y(t + \epsilon) = \tilde{y} \mid Y(t) = y\} \approx \epsilon q(y, \tilde{y})$. The diagonal elements $q(y, y)$ are negative and are defined so that the rows of Q sum to zero.

When modeling a social network, the outcome space $\mathcal{Y}$ is taken to be all possible edge configurations of an N -node network, and an individual configuration $\mathbf{y} \in \mathcal{Y}$ is taken to be

a binary vector of length $\binom{N}{2}$. We use the shorthand $q_{ij}(\mathbf{y})$ to denote the propensity for the edge between node i and j to flip into its opposite value under configuration $\mathbf{y}$. The function $q_{ij}(\mathbf{y})$ completely specifies the dynamics of the network model. We now review several variants of CMPM which differ only in their definition of $q_{ij}(\mathbf{y})$.

Independent arc, reciprocity, and popularity models. The *independent arc* model employs the simplest definition of $q_{ij}(\mathbf{y})$:

$$\text{Independent arc model: } q_{ij}(\mathbf{y}) = \lambda_{y_{ij}}, \quad (4.4)$$

i.e., Y_{ij} changes from 0 to 1 at a rate λ_0 , and from 1 to 0 at rate λ_1 . In this model, modification to one edge does not depend on the setting of other edges. The model is simple enough that the transition probabilities $\Pr(t)$ can be derived in closed form (see, e.g., Taylor and Carlin [292] p. 362-364). Maximum likelihood parameter estimation for this model was discussed in [278].

In the *reciprocity* model, the rate of change in y_{ij} depends only on the reciprocal edge y_{ji} :

$$\text{Reciprocity model: } q_{ij}(\mathbf{y}) = \lambda_{y_{ij}} + \mu_{y_{ij}} y_{ji}. \quad (4.5)$$

Thus, if no link currently exists between nodes i and j , then the propensity for adding either directed edge is λ_0 ; if one directed edge exists, then the reciprocal edge is added with propensity $\lambda_0 + \mu_0$. If one directed edge exists, then it is deleted with rate λ_1 . If both edges exist, then the deletion propensity for either is $\lambda_1 + \mu_1$. The transition matrix $\Pr(t)$ can be derived but has a complicated form [189; 272].

Along the same line of development, the *popularity* model and the *expansiveness* model [312; 313] define the change rate for edge y_{ij} to be dependent on y_{+j} , the in-degree of node j , or y_{i+} , the out-degree of node i :

$$\text{Popularity model: } q_{ij}(\mathbf{y}) = \lambda_{y_{ij}} + \pi_{y_{ij}} y_{+j}, \quad (4.6)$$

$$\text{Expansiveness model: } q_{ij}(\mathbf{y}) = \lambda_{y_{ij}} + \pi_{y_{ij}} y_{i+}. \quad (4.7)$$

Edge-oriented dynamics. Snijders [276] outlines two categories of transition dynamics: edge-oriented and node-oriented. In both cases, the intensity matrix is factored into two components: one controls the *opportunity* for change, and the other specifies the propensity of change. More precisely, the continuous time Markov process is now split into two sub-processes; the first operating in the continuous time domain and dictating *when* a change should occur; the second dealing with the probability of the discrete event of individual edge flips. Both edge-oriented and node-oriented dynamics can be interpreted as stochastic optimizations of a potential function $f(\mathbf{y})$ on the network configuration. The difference is that, in the edge-oriented case, f is based on global statistics of the network, whereas in the node-oriented case, f is defined for each node's local neighborhood. Moreover, the choice of *which* edge to flip differs between the two formulations.

Using $y(i, j, z)$ to denote the configuration where the edge e_{ij} has the value $z \in \{0, 1\}$, edge-oriented dynamics can be written in the following general form:

$$q_{ij}(\mathbf{y}) = \rho p_{ij}(\mathbf{y}), \quad (4.8)$$

where

$$p_{ij}(\mathbf{y}) = \frac{\exp(f(y(i, j, 1 - y_{ij})))}{\exp(f(y(i, j, 0))) + \exp(f(y(i, j, 1)))}. \quad (4.9)$$

Thus, in edge-oriented dynamics each edge follows an independent Poisson process, so that the time until the next event has an exponential distribution with parameter ρ . When an event occurs for edge $i \rightarrow j$, the edge flips to its opposite value with probability $p_{ij}(\mathbf{y})$.

The potential function $f(\mathbf{y})$ is usually defined as a linear combination of network statistics:

$$f(\mathbf{y}) = \sum_k \beta_k s_k(\mathbf{y}). \quad (4.10)$$

This should start to look familiar. Indeed the CMPM process with edge-oriented dynamics is equivalent to the Gibbs sampling process for ERGMs (where the next edge to be updated is selected randomly). The statistics $s_k(\mathbf{y})$ for node k take on the usual forms (see [Table 4.1](#)).

Number of directed arcs:	$s_1(\mathbf{y}) = \sum_{ij} y_{ij}$
Number of reciprocated arcs:	$s_2(\mathbf{y}) = \sum_{ij} y_{ij} y_{ji}$
Number of pairs of arcs with the same target:	$s_3(\mathbf{y}) = \sum_{ijk} y_{kj} y_{ji}$
Number of pairs of arcs with the same origin:	$s_4(\mathbf{y}) = \sum_{ijk} y_{ik} y_{ij}$
Number of paths of length two:	$s_5(\mathbf{y}) = \sum_{ijk} y_{ij} y_{jk}$
Number of transitive triplets:	$s_6(\mathbf{y}) = \sum_{ijk} y_{ij} y_{ik} y_{jk}$

Table 4.1: The table of network statistics for a directed social network.

The statistics in [Table 4.1](#) assume directed graphs, however it is easy to come up with the corresponding statistics for undirected graphs. For example, in the undirected case all the edges are “reciprocal” and thus s_1 and s_2 are combined into $s'(\mathbf{y}) = \sum_{i,j \in \mathbb{N}} y_{ij}$.

Due to their close relations to ERGMs, edge-oriented models suffer the same fate of degeneracy. For example, if the parameter β for transitive triplets is not too small, then with high probability the simulated network will be a complete graph. However, compared to static networks, degeneracy in the longitudinal case is not as much a concern, as the complete graph will only emerge at some distant time in the future.

Node-oriented dynamics. Fully node-oriented dynamics [275] defines the intensity matrix as

$$q_{ij}(\mathbf{y}) = \rho_i p_{ij}(\mathbf{y}), \quad (4.11)$$

where

$$p_{ij}(\mathbf{y}) = \frac{\exp(f_i(\mathbf{y}(i, j, 1 - y_{ij})))}{\sum_{h \neq i} \exp(f_i(\mathbf{y}(i, h, 1 - y_{ih})))}. \quad (4.12)$$

Thus the independent Poisson processes for determining edge change *opportunity* are now defined for each node (with intensity ρ_i) as opposed to each edge. Given the opportunity for edge change, each node seeks to optimize its own potential function as defined by

$$f_i(\mathbf{y}) = \sum_k \beta_k s_{ik}(\mathbf{y}). \quad (4.13)$$

The function $f_i(\mathbf{y})$ is similar to the global potential $f(\mathbf{y})$ in Equation 4.10 but only aggregates over the local neighborhood of node i . Node i favors changing the incident edge that would lead to the biggest increase in its potential.

Edge-node mixed dynamics. Snijders [276] also suggested a form of mixed dynamics where the opportunity for change is edge-oriented, but the potential functions are node-oriented:

$$q_{ij}(\mathbf{y}) = \rho \frac{\exp(f_i(\mathbf{y}(i, j, 1 - y_{ij})))}{\sum_{h \neq i} \exp(f_i(\mathbf{y}(i, h, 1 - y_{ih})))}. \quad (4.14)$$

Thus the opportunity to modify each edge $i \rightarrow j$ follows independent Poisson processes with parameter ρ . But given the opportunity for change, the probability of an actual flip depends on node i 's local network configuration.

Remark. Parameter estimation in CPCCM models has until recently been done via method of moments, where the expected values are obtained through MCMC on simulated networks [273]. Koskinen and Snijders [179] proposed a Bayesian inference method that allows for computation of the posterior distribution of the parameters and treats missing values more adequately. For details of the procedure, please refer to Koskinen and Snijders [179].

4.5 Discrete Time Markov Models

In this section, we outline three recent proposals of dynamic network models operating in the discrete time domain (see also [22]). All three models have the Markov property and represent the likelihood as a sequence of factored conditional probabilities

$$\Pr(Y^1, Y^2, \dots, Y^T) = \Pr(Y^T | Y^{T-1}) \Pr(Y^{T-1} | Y^{T-2}) \dots \Pr(Y^2 | Y^1), \quad (4.15)$$

where $\{Y^1, \dots, Y^T\}$ is a sequence of T observed snapshots of the network. Banks and Carley [22] discussed the simplest version of such models. See also [253].

4.5.1 Discrete Markov ERGM Model

Hanneke and Xing [139] proposed a natural extension of the ERGM model in the discrete Markov domain. Unlike the set up in the continuous domain, the potential function in this model involve the statistics of two consecutive configurations of the network:

$$\Pr(\mathbf{y}^t \mid \mathbf{y}^{t-1}) = \frac{1}{Z} \exp\left\{\sum_k \beta_k s_k(\mathbf{y}^t, \mathbf{y}^{t-1})\right\}. \quad (4.16)$$

Table 4.2 lists a few examples of network statistics defined on pairs of network snapshots.

Density of edges:	$s_1(\mathbf{y}^t, \mathbf{y}^{t-1}) = \frac{1}{(n-1)} \sum_{ij} y_{ij}^t$
Stability:	$s_2(\mathbf{y}^t, \mathbf{y}^{t-1}) = \frac{1}{(n-1)} \sum_{ij} [y_{ij}^t y_{ij}^{t-1} + (1 - y_{ij}^t)(1 - y_{ij}^{t-1})]$
Reciprocity:	$s_3(\mathbf{y}^t, \mathbf{y}^{t-1}) = n \sum_{ij} y_{ji}^t y_{ij}^{t-1} / \sum_{ij} y_{ij}^{t-1}$
Transitivity:	$s_4(\mathbf{y}^t, \mathbf{y}^{t-1}) = n \sum_{ijk} y_{ik}^t y_{ij}^{t-1} y_{jk}^{t-1} / \sum_{ijk} y_{ij}^{t-1} y_{jk}^{t-1}$

Table 4.2: The table of network statistics for pairs of network snapshots.

The basic model may be extended to allow for multiple relations, node attributes, and K-th order Markov dependencies of the form

$$\Pr(Y^{K+1}, Y^{K+2}, \dots, Y^T \mid Y^1, \dots, Y^K) = \prod_{t=K+1}^T \Pr(Y^t \mid Y^{t-K}, \dots, Y^{t-1}), \quad (4.17)$$

where

$$\Pr(Y^t \mid Y^{t-K}, \dots, Y^{t-1}) = \frac{1}{Z} \exp\left\{\sum_k \beta_k s_k(Y^t, \dots, Y^{t-K})\right\}. \quad (4.18)$$

The joint distribution of the first K network snapshots may be represented by an ERGM for the first snapshot, and a $(k-1)$ -th order discrete Markov dependency model for Y_k . The paired network statistics may be extended over K network sequences.

Maximum likelihood parameter estimates may be computed via any numerical approximation technique such as the Newton-Raphson method. Computation of the gradient and Hessian requires the mean and covariance of the sequence network statistics, which are exactly computable for a pair of networks, but require Gibbs sampling in the K -sequence case [139]. The likelihood of this model is well behaved if the minimum sufficient statistics involve only dyads, however, similar to its static counterpart, the full dynamic ERGM is prone to likelihood degeneracy.

4.5.2 Dynamic Latent Space Model

Sarkar and Moore [264] extended the static latent space model of Hoff et al. [146] (cf. [section 3.9](#)) in the time domain. Recall that in the static latent space model, the log odds ratio of a link between nodes i and j depends on the distance between their latent positions z_i and z_j . The dynamic latent space model allows the latent positions to change over time in Gaussian-distributed random steps:

$$Z_t \mid Z_{t-1} \sim \mathcal{N}(Z_{t-1}, \sigma^2 I). \quad (4.19)$$

The observation model is a modified version of the original latent space model¹:

$$p_{ij}^L := p^L(y_{ij} = 1) = \frac{1}{1 + \exp(d_{ij} - r_{ij})}, \quad (4.20)$$

where d_{ij} is the Euclidean distance between i and j in latent space, and r_{ij} is a radius of influence defined as $c \times (\max(\delta_i, \delta_j) + 1)$ (δ_i and δ_j being the degrees of node i and j , respectively). The “radius of influence” is based on the assumption that the higher the maximum degree of the two end nodes, the more likely the edge. This may be true in citation networks where prolific authors are more likely to form new co-authorships. The constant 1 is added to ensure that the radius is non-zero, and c is estimated from data by a line-search (a minimization method in one dimension).

The link probability p_{ij} is defined to be a mixture between the modified latent space link probability p_{ij}^L and a noise probability ρ . The idea is that pairs of nodes who are outside of each other’s radius have only a low noise probability of establishing a link, while nodes within each other’s radii follow the probability p_{ij}^L :

$$p_{ij} = \kappa(d_{ij})p_{ij}^L K(d_{ij}) + (1 - \kappa(d_{ij}))\rho. \quad (4.21)$$

The full observation model is then

$$\Pr(Y^t \mid Z^t) = \prod_{i \sim j} p_{ij} \prod_{i \not\sim j} (1 - p_{ij}), \quad (4.22)$$

where $i \sim j$ denotes the presence of an edge from i to j . The latent space positions Z^t are estimated in sequence for $t = 1 \dots T$ by maximizing the likelihood of the observed Y^t :

$$Z^t = \operatorname{argmax}_Z \Pr(Y^t \mid Z) \Pr(Z \mid Z^{t-1}). \quad (4.23)$$

The authors propose conjugate gradient optimization starting from an initial estimate of the latent positions based on a multidimensional scaling (MDS) transform of the observed pairwise distances. To eliminate rotational ambiguity, a Procrustean (rotationally invariant) transform is applied to the MDS transform so that Z^t is aligned with Z^{t-1} .

Applying the model to the NIPS paper co-authorship dataset (cf. [subsection 2.2.6](#)), the authors gave anecdotal evidence of the validity of the changing embeddings of several well

¹Note that in this dynamic version of the latent space model, links are assumed to be undirected.

known machine learning researchers over time. The dynamics of the researchers’ latent positions allowed for an insight into the evolution of the machine learning community.

Sarkar et al. [265] also proposed a richer model based on [124], which improved upon previous work in two ways. One of the differentiating features of this work was the ability to simultaneously embed words and authors into the latent space, which allowed for representation of a two-mode network. The major advantage, however, was the inference method—the authors proposed a Kalman-filter like dynamic procedure, which allowed for estimation of the posterior distributions over the positions of the authors in the latent space. Proposed procedure was applied to a simulated NIPS dataset.

The impact of this line of work is dichotomous: first, it offers an explanation of the network at every time step, and second, it enables an accurate and efficient prediction of the state of the network at a time step in the future. The proposed inference procedures made it possible for network modeling to scale to large dynamic collections of data. The drawback of this approach is the lack of an explicit mechanism that could explain the dynamics behind the real networks.

Another latent model for citation networks was developed in the physics community. Leicht et al. [190] proposed to use latent variables to capture the grouping of papers that have similar citation profiles over time. The network in this case is a directed acyclic graph and the nodes are papers rather than authors. Using as example a set of opinions from the US Supreme Court and their citations between the years of 1789 and 2007, the authors showed how a simple latent model was able to recover, in a completely unsupervised manner, the different eras in US Supreme court opinion references. The parameters of the model, except for the number of latent classes, were estimated using an EM algorithm. Different numbers of latent classes were tested and each revealed something new about the underlying data. The authors also compared the latent method to a clustering based on network *modularity* [233]. Even with the information about time (directionality in the graph) removed, the latent variable model was still able to discover the same split between two groups of opinions that happened around 1937. The network modularity clustering in a way validated the outcome of the latent model.

In a separate experiment, Leicht et al. [190] showed that deterministic approaches such as “hubs and authorities” and eigenvector centrality [171] discovered interesting network properties that were not revealed by the statistical models. The deterministic analyses showed several significant drops in the age of authorities cited, meaning that once in a while, the younger set of opinions became the new authorities and that the process happened in a “decisive” manner, rather than gradually. In this way, deterministic network analysis approaches complement statistical models.

4.5.3 Dynamic Contextual Friendship Model (DCFM)

The dynamic contextual friendship model (DCFM) of Goldenberg and Zheng [128] represents an attempt to capture several aspects of the complexity of the evolution of real social networks over time. In a real-life friendship network, people may meet and interact with each other under different contexts (e.g., school, work projects, social outings, etc.), and the

strength of interpersonal relationships change over time based on these interactions. DCFM offers such a mechanism for network evolution, where edges have weights that indicate the strength of the relationship, and each node is given a distribution over social interaction spheres (contexts). Context is defined to be any activity where people may interact with each other. At each given time step, each node chooses a random context according to the node’s distribution over contexts. Nodes that appear in the same context update the weights of the links between them. The probability of a weight increase (or decrease) depends on whether the pair had a chance to meet (a coin toss in a model) and the “friendliness” parameter of the individuals involved. The possibility of both positive and negative weight updates allows for edge birth and death over time. An extension of the model also allows for addition and deletion of nodes.

The underlying dynamics is captured by a first-order Markov chain model. Letting W^t denote the weighted adjacency matrix at time t , the basic generative process at time t can be formalized as follows:

1. For each node i , sample context $C_i \sim \text{mult}(\theta_i)$, where θ_i denotes the context distribution parameters.
2. For each pair of nodes i and j in the same context, sample meeting variable $M_{ij} \sim \text{Bern}(\nu_i \nu_j)$, where ν_i and ν_j represent the “friendliness” of nodes i and j ;
- 3.

$$W_{ij}^t = \begin{cases} \text{Poi}(\lambda_h(W_{ij}^{t-1} + 1)) & \text{if } M_{ij} = 1, \\ \text{Poi}(\lambda_\ell(W_{ij}^{t-1})) & \text{otherwise,} \end{cases}$$

where λ_h and λ_ℓ are hyperparameters indicating the rates of growth and decay, respectively. The idea is that a meeting should increase the edge weight with high probability, otherwise the weight decays.

The parameters $\theta_i, \nu_i, \lambda_h, \lambda_\ell$ all have conjugate priors and are estimated through Gibbs sampling [331].

The model can generate networks with a number of different properties. For example, Figure 4.3 shows various degree distributions generated by DCFM, while Figure 4.4 demonstrates possible relation dynamics. Pair (47, 45) shows a brief resuming of the relationship, which dissolves again in the next moment. While DCFM is capable of emulating such long-term memory of past relationships, it does so at the cost of added model complexity.

Few datasets contain weighted relationships. The Enron dataset (cf. subsection 2.2.2) contains email exchanges that can be aggregated on a weekly basis to simulate strength of relationships. In the NIPS dataset (cf. subsection 2.2.6), the number of joint publications per year can represent the strength of the coauthorship. In these cases, the DCFM contexts can be taken to be the topics of emails or articles, and the friendliness parameters can be estimated using the method of moments.

One drawback of DCFM is its lack of identifiability; it is impossible to tell without additional knowledge whether an individual formed many friendships because he frequently

Figure 4.3: Log-log plot of the degree distributions of a network with 200 people. ν_i is drawn from Beta(1,3) for the plot on the left, and from Beta(1,8) for the right hand side. Solid lines represent a linear fit and dashed lines quadratic fit to the data. Contexts are drawn every 50-th timesteps.

Figure 4.4: Weight dynamics for 4 different pairs in a DCFM simulated network of 600 people over 600 time steps. Contexts switches occur every 50-th timestep and $b = 3$.

changes contexts and is very friendly or because the contexts themselves tend to be large. Also, weighted network data are hard to come by and thus pseudo-weights often have to be used.

The DCFM model is important in its own right: the life-mimicking, rich generative mechanism is a step towards realistic complex models that ultimately can be used to explain the intricacies of observed data, especially if additional information about contexts and individuals' friendliness is available.

Chapter 5

Issues in Network Modeling

There are a number of major statistical modeling and inferential challenges in the analysis of network data that go well beyond those described in previous sections of this article. These relate to both the quality and the ease of statistical inference and we mention a few of them here:

Network Visualization. With the rise of online social networks and network modeling, we have seen a proliferation of visualization tools, especially those based on variations of constraint-based spring model algorithms, e.g., see the discussion and references in Shneiderman and Aris [267]. The automated algorithms often use node degrees or some form of distance metric between nodes to arrange their placement. For example, *SoNIA*¹ is a popular package for visualizing dynamic or longitudinal network data; it can be used as a platform for the development, testing, and comparison of various static and dynamic layout techniques. However, little is known about how to effectively combine visualization with the kinds of statistical models we review here, especially if one wants to use the visualization as another tool in the analysis of network data.

Computability. Can we do statistical estimation computations and model fitting exactly for large networks, e.g., by full MCMC methods for mixed membership and exponential random graph models, or do we need to resort to approximations such as those involved in the variational approximation employed by [8; 9]?

For ERGM models a newly updated suite of programs and documentation is now available [138; 157; 224; 129]. The SIENA package² developed by Snijders and colleagues contains a complementary suite of programs that are particularly useful for longitudinal network analyses (though Rinaldo et al. [251] speak words of caution). The packages are capable of learning networks of size up to a few thousand nodes.

The truth is that it is unrealistic to expect that really large networks with millions of nodes can be estimated using exact methods. Even variational approximations, which have

¹<http://sonia.stanford.edu/>

²<http://stat.gamma.rug.nl/siena.html>

their own drawbacks such as sensitivity to the starting point, are not realizable for networks on a really large scale. The key to network modeling and parameter estimation is to take into account the sparsity that comes with size. The methods that are good on small or medium-sized but relatively dense networks, might be computationally infeasible or contain invalid assumptions for larger networks. As we gear up to model very large networks, it is important to focus not only on the disadvantages that size brings but also on its advantages.

Asymptotics and Assessing Goodness of Fit. There is no standard large sample asymptotics for networks (e.g., as N goes to infinity) that can be used to assess the goodness-of-fit of models. Thus we may have serious problems with variance estimates for parameters and with confidence or posterior interval estimates. While a few models with a small number of fixed parameters have well-behaved asymptotics, the problems here tend to be the inherent dependence of network data and the growth in the number of parameters to be estimated as N increases. Haberman [134] comments briefly on asymptotics in his discussion of the p_1 model, and notes the similarity to issues for the Rasch model from item response theory. The lack of asymptotics means that we may have problems of consistency of estimators, but it also means that there is no standard basis for model comparison and assessing goodness of fit. Most other authors have addressed these issues either empirically, e.g., Hunter et al. [156], or not at all.

There are two alternative approaches. We can consider assessing fit or comparing models using exact distributions given the minimal sufficient statistics (MSSs). This works for simple models but not obviously for the general class of ERGMs or most dynamic models in the literature. Further, for many of the models, especially those involving latent variables, the MSSs are the data themselves. Alternatively, we could think in terms of some form of cross-validation for model selection and assessment. The problem with cross-validation is the boundary effects associated with subsets of nodes. This is directly related to the problem of sampling in networks.

Bickel and Chen [37] address the problem of asymptotics in the context of blockmodeling or community discovery, and the methods they exploit may be useful in a broader context when the number of parameters to be estimated grows as N increases.

Sampling. Do our data represent the entire network or are they based on only a subnetwork or subgraph? When the data come from a subgraph, even one selected at random, we need to worry about the effects at the boundary³ and the attendant biases they bring to parameter estimates, cf. the negative result in Stumpf et al. for scale-free models in which they show the extent and nature of the bias [289]. Most of the early results on sampling for network data focused on random subgraphs and exploited the traditional statistical theory of design-based sampling, in which the properties of the network are assumed to be fixed, and we evaluate sample quantities by considering their distribution under all possible similarly

³The boundary is the collection of observed nodes which have links to the unobserved nodes. The boundary can potentially include all observed nodes. Only nodes for which the set of known links are certain to be complete are not included in the boundary – the condition that is hard to satisfy in real world networks.

selected subgraphs. For details, see the many papers by Ove Frank [109; 295] and others [125; 135; 258]. Wiuf and Stumpf [325] and Stumpf and Thorne [288] recently adopted a related but different approach focusing on properties such as degree distributions using binomial random sample sizes from “large” graphs. Others such as Leskovec and Faloutsos [191] examine aspects of the question in an empirical but ad hoc fashion. The relevance of sampling for model-based network inference was first addressed by Thompson and Frank [295], and further developed by Handcock and Gile [135], who adapt MCMC algorithms for exponential random graph models to account for sampling designs. To date, these are the only works to seriously explore this important topic. Airoldi and Carley [6] quantify the sensitivity of alternative sampling algorithms to generate graphs that share similar topological properties, as well as the divergence of topological properties of algorithms for sampling popular network models.

We expect the issue of sampling to be of relevance to virtually all of the models and we need to explore their consequences. This will be especially true when we try to update model parameter estimates based on extracts of data in a dynamic fashion.

Missing data. Along with sampling arises a question of the treatment of missing data in statistical networks. Usually, the non-respondents to surveys are excluded from the analysis and the modeling considers only individuals for which all data is available. A few works deal with missing data directly. The empirical impact of nonrespondents in a survey to analysis is considered in [284], the modeling implications and inference for non-respondents in ERGM can be found in [255; 120; 178]. Missing data in longitudinal studies is the subject of [154]. This work makes assumptions about sampling strategies to justify the estimation of missing edges using a Missing at Random assumption. Because this is not in general a correct assumption we have an interesting set of open problems. Kossinets [180] considers three missing data mechanisms: network boundary specification (non-inclusion of actors or affiliations), survey non-response, and censoring by vertex degree (fixed choice design), and examines their effect on a study of a scientific collaboration network. One type of missing data - links or relations - can be treated as a prediction task by treating links between nodes in a given network as probabilistic quantities and using statistical models based on the available data to estimate the likelihood of those edges being there. The problem of prediction is often addressed in the machine learning community and we discuss it next.

Prediction. In our review of the literature on networks across many disciplines we have found limited methodological work focussing on evaluating and comparing the predictive ability of various models, static or dynamic. There are papers on link prediction in the relational network model literature (e.g. [238]). Liben-Nowell and Kleinberg [198] develop approaches to link prediction based on measures for analyzing the “proximity” of nodes in a network, e.g., the WWW. In biological literature, a number of papers examine the problem of predicting missing links in biological networks (e.g. [327] is one of the earlier works). However, these papers focus on how to cleverly combine heterogeneous data in order to discover new links. The evaluation is usually limited to cross-validation on the

known links—information that is incomplete and available only for a few organisms. In the sociological literature on organizations, there is often interest in distinguishing among organizations on the basis of their network structure, so there would clearly be interest in utilizing methodology for prediction based on network structure. Because making predictions of various sorts from dynamic network models fits well within the machine learning paradigm, we expect to see many more papers on the topic in the not too distant future.

Embeddability. Underlying most dynamic network models is a continuous time stochastic process even though the data used to study the models and their implications may come in the form of repeated snapshots at discrete time points (epochs)—a form of time sampling as opposed to node sampling referred to above—or cumulative network links. In such circumstances we need to take special care in how we represent and estimate the continuous-time parameters in the actual data realizations used to fit models. This is known in the statistical literature as the **embeddability** problem and was studied for Markov processes in the 1970s by Singer and Spilerman [270, 271] for social processes, and more recently by Hansen and Scheinkman [140] in the context of econometric models and by others in the computational finance literature. Wasserman [313] and various papers by Snijders and his collaborators illustrate how to address embedding in some simple dynamic models.

Identifiability. Identifiability of model parameters is a technical issue in statistics that refers to the fact that multiple solutions may exist (in the parametric space) that lead to exactly the same likelihood. In this sense, no inference procedure can distinguish between these solutions. For instance, in a mixture model we can permute the assignments of points to mixture components to obtain an equivalent solution. There are a number of papers that describe the issue in various models (e.g., [283; 132]) and from different perspectives (e.g., [51; 52] from the algebraic perspective). A few solutions to address this issue have been proposed recently. Some consider inference on equivalence classes in a blockmodel for network data [236]. Others pre-process the data to identify a reference solution that drives the inference [137].

Combining links with their attributes. In many network data sets, especially those arising in machine learning contexts, there are attributes associated with the network links. For example in e-mail and blog databases, the attributes may be taken to be the contents of the messages or postings. There is an emerging literature focused on cascades of such links but few papers are situated in a full network model setting and few authors attempt to combine the models for links with models for message or posting texts. This is a natural extension to models described here, especially the mixed membership stochastic blockmodels of [section 3.8](#), since the text could naturally be modeled by mixed-membership topic models. McCallum et al. [208] and Chang and Blei [61] suggest different ways to approach this kind of combination model. Dynamic models that combine evolving block and topic structures would be of special interest for such applications.

Chapter 6

Summary

The ubiquity of networks in areas as diverse as the social sciences, biology, computer science, physics, and economics, has spawned extensive literature on the subject. In this review, we discussed in detail a few main trends in the *statistical* network modeling literature, focussing on models that have historically inspired many others as well as a few recent proposals. By charting the evolution of statistical network modeling approaches, we pointed out explicit connections between the discussed models. Figure 6.1 provides a visual diagram of model influence; an arrow pointing from A to B means either that the development of model A influenced the subsequent development of model B, or that B can be viewed as a generalization of A.

The literature on network modeling may be divided along different lines of motivation. Models primarily introduced in the physics literature are motivated by asymptotic properties of networks, whereas the literature stemming from statistics and statistical social science is concerned with the inference step in addition. Thus, the main criticism of the random graph models primarily developed in statistical physics is the lack of the assessment of the fit of the models to the data. The main drawback in the statistical literature is the lack of the comprehensive asymptotic analysis. Though degeneracy found in the limiting case of the earlier versions of the ERGM has been addressed, a more broad analysis is still missing.

In this work we made a distinction between static and dynamic models. Descriptive models such as p_1 , p_2 , and ERGM are clearly static as they infer a set of sufficient statistics from a single snapshot of an existing network. The families of continuous and discrete time Markov models, on the other hand, are clearly dynamic as they seek to model multiple snapshots of an evolving network. The Erdős-Rényi-Gilbert, preferential attachment, and small-world models, while ultimately aim to model a single time point snapshot of a network, are usually described via generative processes, where edges are added one at a time. These models can thus be considered as either static, with respect to what they model, or dynamic, with respect to how they're represented. In this work we refer to them as pseudo-dynamic.

Within the category of static models we discussed two main directions: models that take networks as given (see [section 3.4](#), [section 3.5](#), and [section 3.6](#)) and models that assume and estimate latent structures ([section 3.8](#) and [section 3.9](#)). Latent structure models have to make certain assumptions about the data. Stochastic blockmodels assume structural

Figure 6.1: Network summarizing the relations between models discussed in our review. White nodes denote static models, yellow nodes – “pseudo-dynamic” and green – dynamic models. Arrows indicate inspiration or influence of the model at the source on the model at the target.

equivalence of the nodes, whereas latent space models assume the existence of an embedding of the network in a low dimensional space. These models allow for better understanding of the data in cases where it is believed to contain hidden structure.

We divided the category of dynamic models into continuous time Markov models and discrete time Markov models. CMPM (section 4.4) assumes that the adjacency matrix evolves according to a continuous Markov chain whose intensity matrix can depend on various edge and node dynamics. Discrete time Markov network models deal with a set of network snapshots observed at various time points. Examples of discrete time Markov network models include dynamic extensions of ERGM (subsection 4.5.1) and the latent space model (subsection 4.5.2), the duplication-attachment model, as well as a generative dynamic model for friendship networks (subsection 4.5.3).

Despite the many advances in network modeling over the last decade, there remains a host of unresolved issues. We listed some of the issues in chapter 5. We feel that, from a

statistics or machine learning perspective, the biggest breakthroughs are to be made in the areas of inference and dynamic modeling. Creating a model or perhaps fixing an existing one in such a way that provides realistic generative and inference mechanisms which can identifiably infer parameters of a large real world network would make a great contribution to the statistical network modeling community.

Acknowledgments

This research was partly supported by United States National Institute of General Medical Sciences Center of Excellence grant P50 GM071508, by National Science Foundation grants DBI-0546275, IIS-0513552, by National Institutes of Health grant R01 GM071966 to Princeton University, by National Science Foundation grant DMS-0907009 to Harvard University, and by National Science Foundation grant DMS-0631589 and partial support from U.S. Army Research Office Contract W911NF0910360 to the Department of Statistics, Carnegie Mellon University. Edoardo M. Airolidi was a postdoctoral fellow in the Department of Computer Science and the Lewis-Sigler Institute for Integrative Genomics at Princeton University when a large portion of this work was carried out. We thank three anonymous reviewers for their valuable comments, as well as their helpful additions and corrections to our citation list. We thank Joseph Blitzstein and Pavel Krivitsky for a careful reading and the correction of a number of infelicities. We finally wish to thank László Barabási and Zoltan Oltvai; Peter Bearman, James Moody, and Katherine Stovel; James Fowler and Nicholas Christakis; Purnamrita Sarkar and Andrew Moore for giving permission to re-print figures from their original papers [27; 31; 65; 263].

Bibliography

- [1] E. M. Airoldi. *Bayesian Mixed Membership Models of Complex and Evolving Networks*. PhD thesis, School of Computer Science, Carnegie Mellon University, 2006.
- [2] E. M. Airoldi. Model-based clustering for social networks: Discussion. *Journal of the Royal Statistical Society, Series A*, 170(2):330–331, 2007.
- [3] E. M. Airoldi. Getting started in probabilistic graphical models. *PLoS Computational Biology*, 3(12):e252, 2007.
- [4] E. M. Airoldi. A family of distributions on the unit hypercube. Technical Report 2, Department of Statistics, Harvard University, 2009.
- [5] E. M. Airoldi. The exchangeable graph model. Technical Report 1, Department of Statistics, Harvard University, 2009.
- [6] E. M. Airoldi and K. M. Carley. Sampling algorithms for pure network topologies: A study on the stability and the separability of metric embeddings. *ACM SIGKDD Explorations*, 7(2):13–22, 2005.
- [7] E. M. Airoldi, D. M. Blei, E. P. Xing, and S. E. Fienberg. A latent mixed-membership model for relational data. In *Proceedings of the 3rd International Workshop on Link Discovery: Issues, Approaches and Applications (LinkKDD '05)*, in conjunction with the 11th International ACM SIGKDD Conference, pages 82–89. ACM Press, New York, 2005.
- [8] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership analysis of high-throughput interaction studies: Relational data. <http://arXiv.org/abs/0706.0294>, 2007.
- [9] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership stochastic blockmodels. *Journal of Machine Learning Research*, 9:1981–2014, 2008.
- [10] L. Akoglu and C. Faloutsos. RTG: A recursive realistic graph generator using random typing. In *Data Mining and Knowledge Discovery*, 19(2):194–209, Springer Netherlands, 2009.

- [11] R. D. Alba. A graph-theoretic definition of a sociometric clique. *Journal of Mathematical Sociology*, 3:113–126, 1973.
- [12] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002.
- [13] R. Albert, H. Jeong, and A.-L. Barabási. Diameter of the world wide web. *Nature*, 401:130–131, 1999.
- [14] D. L. Alderson. Catching the ‘network science’ bug: Insight and opportunity for the operations researcher. *Operations Research*, 56(5):1047–1065, 2008.
- [15] D. J. Aldous. Exchangeability and related topics. In *Lecture Notes in Mathematics*, volume 1117, pages 1–198. Springer Berlin / Heidelberg, 1985. (Also in Ecole d’Ete St Flour 1983).
- [16] S. Allesina, D. Alonso, and M. Pascual. A general model for food web structure. *Science*, 320(5876):658–661, 2008.
- [17] U. Alon. Network motifs: Theory and experimental approaches. *Nature Reviews Genetics*, 8:450–461, 2007.
- [18] L. A. N. Amaral, A. Scala, M. Barthélémy, and H. E. Stanley. Classes of small-world networks. *Proceedings of the National Academy of Sciences*, 97(21):11149–11152, 2000.
- [19] P. Arabie, S. A. Boorman, and P. R. Levitt. Constructing blockmodels: How and why. *Journal of Mathematical Psychology*, 17(1):21–63, 1978.
- [20] L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan. Group formation in large social networks: Membership, growth, and evolution. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 44–54. ACM Press, New York, 2006.
- [21] D. Banks and K. M. Carley. Metric inference for social networks. *Journal of Classification*, 11(1):121–149, 1994.
- [22] D. Banks and K. M. Carley. Models for network evolution. *Journal of Mathematical Sociology*, 21:173–196, 1996.
- [23] E. Banks, E. Nabieva, R. Peterson, and M. Singh. NetGrep: Fast network schema searches in interactomes. *Genome Biology*, 9(9):R:138, 2008. <http://genomebiology.com/content/9/9/R138>.
- [24] A.-L. Barabási. *Linked: The New Science of Networks*. Perseus, Cambridge, MA, 2002.
- [25] A.-L. Barabási. The origin of bursts and heavy tails in human dynamics. *Nature*, 435: 207–211, 2005.

- [26] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [27] A.-L. Barabási and Z. Oltvai. Network biology: Understanding the cell’s functional organization. *Nature Reviews Genetics*, 5(2):101–113, 2004.
- [28] A.-L. Barabási, H. Jeong, Z. Neda, E. Ravasz, A. Schubert, and T. Vicsek. Evolution of the social network of scientific collaboration. *Physica A*, 311(3–4):590–614, 2002.
- [29] F. Bassetti, M. Cosentino Lagomarsino, and S. Mandra. Exchangeable random networks. *Internet Mathematics*, 4(4):357–400, 2007.
- [30] J. Baumes, M. Goldberg, M. Magdon-Ismail, and W. A. Wallace. Discovering hidden groups in communication networks. In *Lecture Notes in Computer Science*, volume 3073, pages 378–389. Springer Berlin / Heidelberg, 2004.
- [31] P. S. Bearman, J. Moody, and K. Stovel. Chains of affection: The structure of adolescent romantic and sexual networks. *American Journal of Sociology*, 110(1):44–91, 2004.
- [32] A. Bernard, D. S. Vaughn, and A. J. Hartemink. Reconstructing the topology of protein complexes. In T. Speed and H. Huang, editors, *Research in Computational Molecular Biology 2007 (RECOMB07)*, volume 4453 of *Lecture Notes in Bioinformatics*, pages 32–46. Springer Berlin / Heidelberg, 2007.
- [33] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B*, 36(2):192–236, 1974.
- [34] I. Bezáková, A. Kalai, and R. Santhanam. Graph model selection using maximum likelihood. In *Proceedings of the 23rd International Conference on Machine Learning*, volume 148 of *ACM International Conference Proceeding Series*, pages 105–112. ACM Press, New York, 2006.
- [35] S. Bhamidi, G. Bresler, and A. Sly. Mixing time of exponential random graphs. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 803–812. IEEE Computer Society, Washington, D.C., 2008.
- [36] I. Bhattacharya. *Collective Entity Resolution in Relational Data*. PhD thesis, University of Maryland, 2006.
- [37] P. J. Bickel and A. Chen. A nonparametric view of network models and Newman-Girvan and other modularities. *Proceedings of the National Academy of Sciences*, (to appear), 2009.
- [38] C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.

- [39] Y. M. M. Bishop, S. E. Fienberg, and P. W. Holland. *Discrete Multivariate Analysis: Theory and Practice*. MIT Press, Cambridge, MA, 1975. Reprinted by Springer-Verlag, 2007.
- [40] D. M. Blei and S. E. Fienberg. Model-based clustering for social networks: Discussion. *Journal of the Royal Statistical Society, Series A*, 170(2):332, 2007.
- [41] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [42] J. Blitzstein and P. Diaconis. A sequential importance sampling algorithm for generating random graphs with prescribed degrees. Technical report, Stanford University, 2006.
- [43] B. Bollobás. *Random Graphs*. Cambridge University Press, New York, 2nd edition, 2001.
- [44] B. Bollobás and F. R. K. Chung. The diameter of a cycle plus a random matching. *SIAM Journal on Discrete Mathematics*, 1(3):328–333, 1988.
- [45] B. Bollobás, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. *Random Structures & Algorithms*, 31(1):3–122, 2007.
- [46] E. Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(Suppl. 3):7280–7287, 2002.
- [47] D. Botstein, S. A. Chervitz, and J. M. Cherry. Yeast as a model organism. *Science*, 277(5330):1259–1260, 1997.
- [48] U. Brandes and T. Erlebach, editors. *Network Analysis: Methodological Foundations*, volume 3418 of *Lecture Notes in Computer Science*. Springer Berlin /Heidelberg, 2005.
- [49] M. Braun and J. McAuliffe. Variational inference for large-scale models of discrete choice. <http://arXiv.org/abs/0712.2526>, 2007.
- [50] M. Buchanan. *Nexus: Small Worlds and the Groundbreaking Science of Networks*. W. W. Norton & Company, New York, 2002.
- [51] M.-L. G. Buot and D. S. P. Richards. Counting and locating the solutions of polynomial systems of maximum likelihood equations, I. *Journal of Symbolic Computation*, 41(2): 234–244, 2006.
- [52] M.-L. G. Buot and D. S. P. Richards. Counting and locating the solutions of polynomial systems of maximum likelihood equations, II: The Behrens-Fisher problem. <http://arXiv.org/abs/0709.0957>, 2007.

- [53] R. S. Burt. Models of network structure. *Annual Review of Sociology*, 6:79–141, 1980.
- [54] K. M. Carley. Group stability: A socio-cognitive approach. In E. Lawler, B. Markovsky, C. Ridgeway, and H. Walker, editors, *Advances in Group Processes*, pages 1–44. JAI Press, Greenwich, CT, 1990.
- [55] K. M. Carley. Smart agents and organizations of the future. In L. Lievrouw and S. Livingstone, editors, *The Handbook of New Media*, pages 206–220. Sage, Thousand Oaks, CA, 2002.
- [56] K. M. Carley and A. Newell. The nature of the social agent. *Journal of Mathematical Sociology*, 19(4):221–262, 1994.
- [57] K. M. Carley and J. Reminga. ORA: Organizational Risk Analyzer. <http://www.casos.cs.cmu.edu/projects/ora/>, 2004.
- [58] K. M. Carley and D. Skillicorn. Special issue on analyzing large scale networks: The Enron corpus. *Computational & Mathematical Organization Theory*, 11(3):179–181, Springer Netherlands, 2005.
- [59] K. M. Carley, D. B. Fridsma, E. Casman, A. Yahja, N. Altman, L.-C. Chen, B. Kaminisky, and D. Nave. BioWar: Scalable agent-based model of bioattacks. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 36(2):252–265, 2006.
- [60] D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-MAT: A recursive model for graph mining. In *Proceedings of the 4th SIAM International Conference on Data Mining*, 2004.
- [61] J. Chang and D. M. Blei. Relational topic models for document networks. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS '09)*, 2009.
- [62] H. Chen, E. Reid, J. Sinai, A. Silke, and B. Ganor, editors. *Terrorism Informatics: Knowledge Management and Data Mining for Homeland Security*. Springer-Verlag, New York, 2008.
- [63] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. The origin of power laws in internet topologies revisited. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communication Societies*, 2:608–617, 2002.
- [64] J. M. Cherry, C. Ball, S. Weng, G. Juvik, R. Schmidt, C. Adler, B. Dunn, S. Dwight, L. Riles, R. K. Mortimer, and D. Botstein. Genetic and physical maps of *Saccharomyces cerevisiae*. *Nature*, 387(6632 Suppl.):67–73, 1997.

- [65] N. A. Christakis and J. H. Fowler. The spread of obesity in a large social network over 32 years. *New England Journal of Medicine*, 357(370-379), 2007.
- [66] N. A. Christakis and J. H. Fowler. The collective dynamics of smoking in a large social network. *New England Journal of Medicine*, 358:2249–2258, 2008.
- [67] N. A. Christakis and J. H. Fowler. Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham Heart Study. *British Medical Journal*, 337:a2338, 2008.
- [68] N. A. Christakis and J. H. Fowler. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. Little, Brown and Co., New York, 2009.
- [69] F. Chung and L. Lu. *Complex Graphs and Networks*. American Mathematical Society, Providence, RI, 2006.
- [70] F. Chung, L. Lu, and V. Vu. The spectra of random graphs with given expected degrees. *Proceedings of the National Academy of Sciences*, 100(11):6313–6318, 2003.
- [71] A. Clauset. Finding local community structure in networks. *Physical Review E*, 72(2):026132, 2005.
- [72] A. Clauset and C. Moore. How do networks become navigable? <http://arXiv.org/abs/cond-mat/0309415>, 2003.
- [73] A. Clauset, C. Moore, and M. E. J. Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 453:98–101, 2008.
- [74] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009.
- [75] R. Clegg, R. Landa, U. Harder, and M. Rio. Evaluating and optimising models of network growth. <http://arXiv.org/abs/0904.0785>, 2009.
- [76] P. Clifford. Markov random fields in statistics. In G. R. Grimmett and D. J. A. Welsh, editors, *Disorder in Physical Systems: A Volume in Honour of John M. Hammersley*, pages 19–32. Oxford University Press, 1990.
- [77] E. Cohen-Cole and J. M. Fletcher. Is obesity contagious? Social networks vs. environmental factors in the obesity epidemic. *Journal of Health Economics*, 27:1382–1387, 2008.
- [78] E. Cohen-Cole and J. M. Fletcher. Detecting implausible social network effects in acne, height, and headaches: Longitudinal analysis. *British Medical Journal*, 337:a2533, 2008.

- [79] J. Copic, M. O. Jackson, and A. Kirman. Identifying community structures from network data via maximum likelihood methods. *The B.E. Journal of Theoretical Economics*, 9(1), 2009.
- [80] A. Davis, B. B. Gardner, M. R. Gardner, and J. J. Wallach. *Deep South: A Social Anthropological Study of Caste and Class*. University of Chicago Press, 1941. Reprinted by University of South Carolina Press, 2009.
- [81] G. B. Davis and K. M. Carley. Clearing the FOG: Fuzzy, overlapping groups for social networks. *Social Networks*, 30(3):201–212, 2008.
- [82] B. de Finetti. *Theory of probability, Vol. 1-2*. John Wiley & Sons, New York, 1990. Reprint of the 1974–1975 translation.
- [83] D. J. de Solla Price. Networks of scientific papers: The pattern of bibliographic references indicates the nature of the scientific research front. *Science*, 149(3683):510–515, 1965.
- [84] P. Diaconis and S. Janson. Graph limits and exchangeable random graphs. Technical report, Department of Statistics, Stanford University, 2008.
- [85] P. Diaconis and B. Sturmfels. Algebraic algorithms for sampling from conditional distributions. *Annals of Statistics*, 26(1):363–397, 1998.
- [86] P. S. Dodds, R. Muhamad, and D. J. Watts. An experimental study of search in global social networks. *Science*, 301(5634):827–829, 2003.
- [87] P. Domingos. Mining social networks for viral marketing. *IEEE Intelligent Systems*, 20(1):80–82, 2005.
- [88] P. Doreian, V. Batagelj, and A. Ferligoj. Generalized blockmodeling of two-mode network data. *Social Networks*, 26:29–53, 2004.
- [89] P. Doreian, V. Batagelj, and A. Ferligoj. *Generalized Blockmodeling (Structural Analysis in the Social Sciences)*. Cambridge University Press, 2004.
- [90] S. N. Dorogovtsev and J. F. F. Mendes. Scaling behavior of developing and decaying networks. *Europhysics Letters*, 52(1):33, 2000.
- [91] R. Durrett. *Random Graph Dynamics*. Cambridge University Press, 2006.
- [92] N. Eagle, A. Pentland, and D. Lazer. Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences*, 106(36):15274–15278, 2009.
- [93] P. Erdős and A. Rényi. On Random Graphs, I. *Publicationes Mathematicae*, 6:290–297, 1959.

- [94] P. Erdős and A. Rényi. The evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 5:17–61, 1960.
- [95] E. A. Erosheva, S. E. Fienberg, and J. Lafferty. Mixed-membership models of scientific publications. *Proceedings of the National Academy of Sciences*, 101(Suppl. 1):5220–5227, 2004.
- [96] E. Even-Dar and M. Kearns. A small world threshold for economic network formation. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems (NIPS)*, volume 19, pages 385–392. MIT Press, Cambridge, MA, 2007.
- [97] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM '99)*, pages 251–261. ACM Press, New York, 1999.
- [98] S. Fields and O. Song. A novel genetic system to detect protein-protein interactions. *Nature*, 340(6230):245–246, 1989.
- [99] S. E. Fienberg. *Analysis of Cross-Classified Categorical Data*. MIT Press, Cambridge, MA, 2nd edition, 1980. Reprinted by Springer-Verlag, 2007.
- [100] S. E. Fienberg and S. K. Lee. On small world statistics. *Psychometrika*, 40(2):219–228, 1975.
- [101] S. E. Fienberg and S. S. Wasserman. Categorical data analysis of single sociometric relations. *Sociological Methodology*, pages 156–192, 1981.
- [102] S. E. Fienberg and S. S. Wasserman. An exponential family of probability distributions for directed graphs: Comment. *Journal of the American Statistical Association*, 76(373):54–57, 1981.
- [103] S. E. Fienberg, M. M. Meyer, and S. S. Wasserman. Statistical analysis of multiple sociometric relations. *Journal of the American Statistical Association*, 80:51–67, 1985.
- [104] S. E. Fienberg, S. Petrović, and A. Rinaldo. Algebraic statistics for p_1 random graph models: Markov bases and their uses. In S. Sinharay and N. J. Dorans, editors, *Papers in Honor of Paul W. Holland*. Educational Testing Service, 2009.
- [105] J. Flannick, A. Novak, B. S. Srinivasan, H. H. McAdams, and S. Batzoglou. Græmlin: General and robust alignment of multiple large interaction networks. *Genome Research*, 16(9):1169–1181, 2006.
- [106] A. D. Flaxman, A. M. Frieze, and J. Vera. A geometric preferential attachment model of networks. *Internet Mathematics*, 3(2):187–206, 2006.

- [107] A. D. Flaxman, A. M. Frieze, and J. Vera. A geometric preferential attachment model of networks II. *Internet Mathematics*, 4(1):87–112, 2007.
- [108] J. Fowler and N. Christakis. Estimating peer effects on health in social networks. *Journal of Health Economics*, 27(5):1400–1405, 2008.
- [109] O. Frank. Network sampling and model fitting. In P. J. Carrington, J. Scott, and S. S. Wasserman, editors, *Models and Methods in Social Network Analysis*, pages 31–56. Cambridge University Press, 2005.
- [110] O. Frank and D. Strauss. Markov graphs. *Journal of the American Statistical Association*, 81(395):832–842, 1986.
- [111] N. Friedman. Inferring cellular networks using probabilistic graphical models. *Science*, 303(5659):799–805, 2004.
- [112] N. Friedman, L. Getoor, D. Koller, and A. Pfeffer. Learning probabilistic relational models. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 1300–1309, 1999.
- [113] M. T. Gastner and M. E. J. Newman. Shape and efficiency in spatial distribution networks. *Journal of Statistical Mechanics: Theory and Experiment*, 1:P01015, 2006.
- [114] A.-C. Gavin, M. Bösch, R. Krause, P. Grandi, M. Marzioch, A. Bauer, J. Schultz, J. M. Rick, A.-M. Michon, C.-M. Cruciat, M. Remor, C. Höfert, M. Schelder, M. Brajenovic, H. Ruffner, A. Merino, K. Klein, M. Hudak, D. Dickson, T. Rudi, V. Gnau, A. Bauch, S. Bastuck, B. Huhse, C. Leutwein, M.-A. Heurtier, R. R. Copley, A. Edelmann, E. Querfurth, V. Rybin, G. Drewes, M. Raida, T. Bouwmeester, P. Bork, B. Seraphin, B. Kuster, G. Neubauer, and G. Superti-Furga. Functional organization of the yeast proteome by systematic analysis of protein complexes. *Nature*, 415:141–147, 2002.
- [115] A.-C. Gavin, P. Aloy, P. Grandi, R. Krause, M. Boesche, M. Marzioch, C. Rau, L. J. Jensen, S. Bastuck, B. Dimpelfeld, A. Edelmann, M.-A. Heurtier, V. Hoffman, C. Höfert, K. Klein, M. Hudak, A.-M. Michon, M. Schelder, M. Schirle, M. Remor, T. Rudi, S. Hooper, A. Bauer, T. Bouwmeester, G. Casari, G. Drewes, G. Neubauer, J. M. Rick, B. Kuster, P. Bork, R. B. Russell, and G. Superti-Furga. Proteome survey reveals modularity of the yeast cell machinery. *Nature*, 440(7084):631–636, 2006.
- [116] L. Getoor and B. Taskar, editors. *Introduction to Statistical Relational Learning*. MIT Press, Cambridge, MA, 2007.
- [117] L. Getoor, N. Friedman, D. Koller, and B. Taskar. Learning probabilistic models of link structure. *Journal of Machine Learning Research*, 3:679–707, 2003.

- [118] C. J. Geyer and E. A. Thompson. Constrained Monte Carlo maximum likelihood for dependent data (with discussion). *Journal of the Royal Statistical Society, Series B*, 54:657–699, 1992.
- [119] E. N. Gilbert. Random graphs. *Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.
- [120] K. J. Gile and M. S. Handcock. Model-based assessment of the impact of missing data on inference for networks. CSSS Working paper No. 66, 2006.
- [121] P. S. Gill and T. B. Swartz. Bayesian analysis of directed graphs data with application to social networks. *Applied Statistics*, 53(2):249–260, 2004.
- [122] M. Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002.
- [123] K. S. Gleditsch. Expanded trade and GDP data. *Journal of Conflict Resolution*, 46(5):712–724, 2002.
- [124] A. Globerson, G. Chechik, F. Pereira, and N. Tishby. Euclidean embedding of co-occurrence data. *Journal of Machine Learning Research*, 8:2265–2295, 2007.
- [125] S. Goel and M. J. Salganik. Respondent-driven sampling as Markov chain Monte Carlo. *Statistics in Medicine*, 28(17):2202–2229, 2009.
- [126] A. Goldenberg and A. Moore. Tractable learning of large Bayes net structures from sparse data. In *Proceedings of the 21st International Conference on Machine Learning*, page 44. ACM Press, New York, 2004.
- [127] A. Goldenberg and A. Moore. Bayes net graphs to understand coauthorship networks. In *KDD Workshop on Link Discovery: Issues, Approaches and Applications*, 2005.
- [128] A. Goldenberg and A. Zheng. Exploratory study of a new model for evolving networks. In E. M. Airoldi, D. M. Blei, S. E. Fienberg, A. Goldenberg, E. P. Xing, and A. X. Zheng, editors, *Statistical Network Analysis: Models, Issues and New Directions*, volume 4503 in *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2007.
- [129] S. M. Goodreau, M. S. Handcock, D. R. Hunter, C. T. Butts, and M. Morris. A *statnet* tutorial. *Journal of Statistical Software*, 24(9):1–26, 2008.
- [130] S. M. Goodreau, J. A. Kitts, and M. Morris. Birds of a feather, or friend of a friend? Using exponential random graph models to investigate adolescent social networks. *Demography*, 46(1):103–125, 2009.
- [131] S. Goyal. *Connections: An Introduction to the Economics of Networks*. Princeton University Press, 2007.

- [132] B. Grún and F. Leisch. Dealing with label switching in mixture models under genuine multimodality. *Journal of Multivariate Analysis*, 100(5):851–861, 2008.
- [133] A. Guetz and P. Constantine. Lecture notes for course on Information Networks, 2007. <http://www.stanford.edu/class/msande337/notes/Lec1.pdf>.
- [134] S. J. Haberman. An exponential family of probability distributions for directed graphs: Comment. *Journal of the American Statistical Association*, 76(373):60–61, 1981.
- [135] M. S. Handcock and K. J. Gile. Modeling networks from sampled data. *Annals of Applied Statistics*, 4(1), 2010.
- [136] M. S. Handcock, G. L. Robins, T. A. B. Snijders, J. Moody, and J. Besag. Assessing degeneracy in statistical models of social networks. *Journal of the American Statistical Association*, 76:33–50, 2003.
- [137] M. S. Handcock, A. E. Raftery, and J. Tantrum. Model-based clustering for social networks (with discussion). *Journal of the Royal Statistical Society, Series A*, 170: 301–354, 2007.
- [138] M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, and M. Morris. statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of Statistical Software*, 24(1):12–25, 2008.
- [139] S. Hanneke and E. P. Xing. Discrete temporal models of social networks. In E. M. Airoldi, D. M. Blei, S. E. Fienberg, A. Goldenberg, E. P. Xing, and A. X. Zheng, editors, *Statistical Network Analysis: Models, Issues and New Directions*, volume 4503 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2007.
- [140] L. P. Hansen and J. A. Scheinkman. Back to the future: Generating moment implications for continuous-time Markov processes. *Econometrica*, 63(4):767–804, 1995.
- [141] K. M. Harris, F. Florey, J. Tabor, P. S. Bearman, J. Jones, and R. J. Udry. *The National Longitudinal Study of Adolescent Health: Research Design*. Technical report, Carolina Population Center, University of North Carolina, Chapel Hill, 2003.
- [142] S. Hill, F. Provost, and C. Volinsky. Network-based marketing: Identifying likely adopters via consumer networks. *Statistical Science*, 21(2):256–276, 2006.
- [143] Y. Ho, A. Gruhler, A. Heilbut, G. D. Bader, L. Moore, S.-L. Adams, A. Millar, P. Taylor, K. Bennett, K. Boutilier, L. Yang, C. Wolting, I. Donaldson, S. Schandorff, J. Shewnarane, M. Vo, J. Taggart, M. Goudreault, B. Muskat, C. Alfarano, D. Dewar, Z. Lin, K. Michalickova, A. R. Willems, H. Sassi, P. A. Nielsen, K. J. Rasmussen, J. R. Andersen, L. E. Johansen, L. H. Hansen, H. Jespersen, A. Podtelejnikov, E. Nielsen, J. Crawford, V. Poulsen, B. D. Sørensen, J. Matthiesen, R. C. Hendrickson, F. Gleeson, T. Pawson, M. F. Moran, D. Durocher, M. Mann, C. W. V. Hogue, D. Figeys, and

- M. Tyers. Systematic identification of protein complexes in *Saccharomyces cerevisiae* by mass spectrometry. *Nature*, 415:180–183, 2002.
- [144] P. D. Hoff. Random effects models for network data. In R. Breiger, K. M. Carley, and P. E. Pattison, editors, *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, pages 303–312. The National Academies Press, Washington, D.C., 2003.
 - [145] P. D. Hoff. Modeling homophily and stochastic equivalence in symmetric relational data. In J. C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems (NIPS)*, volume 20, pages 657–664. MIT Press, 2008.
 - [146] P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.
 - [147] P. W. Holland and S. Leinhardt. Local structure in social networks. *Sociological Methodology*, 7:1–45, 1976.
 - [148] P. W. Holland and S. Leinhardt. A dynamic model for social networks. *Journal of Mathematical Sociology*, 5(1):5–20, 1977.
 - [149] P. W. Holland and S. Leinhardt. An exponential family of probability distributions for directed graphs (with discussion). *Journal of the American Statistical Association*, 76(373):33–65, 1981.
 - [150] P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983.
 - [151] P. Holme, J. Karlin, and S. Forrest. An integrated model of traffic, geography and economy in the internet. *ACM SIGCOMM Computer Communication Review*, 38(3):7–15, 2008.
 - [152] B. A. Huberman and L. A. Adamic. Growth dynamics of the world-wide web. *Nature*, 401:131, 1999.
 - [153] S. Huh and S. E. Fienberg. Temporally-evolving mixed membership stochastic blockmodels: Exploring the Enron e-mail database. In *Proceedings of the NIPS Workshop on Analyzing Graphs: Theory & Applications*, Whistler, British Columbia, 2008.
 - [154] M. Huisman and C. Steglich. Treatment of non-response in longitudinal network studies. *Social Networks*, 30(4):297–308, 2008.
 - [155] D. R. Hunter and M. S. Handcock. Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics*, 15(3):565–583, 2006.

- [156] D. R. Hunter, S. M. Goodreau, and M. S. Handcock. Goodness of fit of social network models. *Journal of the American Statistical Association*, 103(481):248–258, 2008.
- [157] D. R. Hunter, M. S. Handcock, C. T. Butts, S. M. Goodreau, and M. Morris. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3), 2008. <http://www.jstatsoft.org/v24/i03/paper>.
- [158] M. Huss and P. Holme. Currency and commodity metabolites: Their identification and relation to the modularity of metabolic networks. *IET Systems Biology*, 1:280–285, 2007.
- [159] T. Ito, K. Tashiro, S. Muta, R. Ozawa, T. Chiba, M. Nishizawa, K. Yamamoto, S. Kuhara, and Y. Sakaki. Toward a protein-protein interaction map of the budding yeast: A comprehensive system to examine two-hybrid interactions in all possible combinations between the yeast proteins. *Proceedings of the National Academy of Sciences*, 97(3):1143–1147, 2000.
- [160] M. O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.
- [161] S. Janson, T. Luczak, and A. Ruciński. *Random Graphs*. John Wiley & Sons, New York, 2000.
- [162] L. J. Jensen and P. Bork. Biochemistry: Not comparable, but complementary. *Science*, 322(5898):56–57, 2008.
- [163] J. H. Jones and M. S. Handcock. Social networks (communication arising): Sexual contacts and epidemic thresholds. *Nature*, 423:605–606, 2003.
- [164] J. H. Jones and M. S. Handcock. An assessment of preferential attachment as a mechanism for human sexual network formation. In *Proceedings of the Royal Society, Series B*, volume 270, number 1520, pages 1123–1128, 2003.
- [165] O. Kallenberg. Probabilistic symmetries and invariance principles. In *Probability and its Applications*. Springer, New York, 2005.
- [166] L. Katz. The distribution of the number of isolates in a social group. *Annals of Mathematical Statistics*, 23(2):271–276, 1952.
- [167] L. Katz and J. H. Powell. Probability distributions of random variables associated with a structure of the sample space of sociometric investigations. *Annals of Mathematical Statistics*, 28(2):442–448, 1957.
- [168] L. Katz and T. R. Wilson. The variance of the number of mutual choices in sociometry. *Psychometrika*, 21(3):299–304, 1956.
- [169] M. Kearns, S. Suri, and N. Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824–827, 2006.

- [170] D. Kempe, J. Kleinberg, and E. Tardos. Influential nodes in a diffusion model for social networks. In *Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 1127–1138. Springer Berlin / Heidelberg, 2005.
- [171] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.
- [172] J. M. Kleinberg. Navigation in a small world—it is easier to find short chains between points in some networks than others. *Nature*, 406:845, 2000.
- [173] J. M. Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, pages 163–170. ACM Press, New York, 2000.
- [174] J. M. Kleinberg. Small-world phenomena and the dynamics of information. In *Advances in Neural Information Processing Systems (NIPS)*, volume 14. MIT Press, Cambridge, MA, 2001.
- [175] J. M. Kleinberg, S. R. Kumar, P. Raghavan, S. Rajagopalan, and A. S. Tomkins. The web as a graph: Measurements, models and methods. In *Computing and Combinatorics*, volume 1627 of *Lecture Notes in Computer Science*, pages 1–17. Springer Berlin / Heidelberg, 1999.
- [176] A. S. Klov Dahl, J. J. Potterat, D. E. Woodhouse, J. B. Muth, S. Q. Muth, and W. W. Darrow. Social networks and infectious disease: The Colorado Springs study. *Social Science & Medicine*, 38(1):79–88, 1994.
- [177] E. D. Kolaczyk. *Statistical Analysis of Network Models*. Springer, New York, 2009.
- [178] J. Koskinen, G. L. Robins, and P. E. Pattison. Analysing exponential random graph (p-star) models with missing data using Bayesian data augmentation. Technical report, Department of Psychology, School of Behavioural Science, University of Melbourne, Australia, 2008.
- [179] J. H. Koskinen and T. A. B. Snijders. Bayesian inference for dynamic social network data. *Journal of Statistical Planning and Inference*, 137(12):3930–3938, 2007.
- [180] G. Kossinets. Effects of missing data in social networks. *Social Networks*, 28(3):247–268, 2006.
- [181] D. Krackhardt. The ties that torture: Simmelian tie analysis in organizations. *Research in the Sociology of Organizations*, 16:183–210, 1999.
- [182] V. E. Krebs. Mapping networks of terrorist cells. *Connections*, 24(3):43–52, 2002.
- [183] P. N. Krivitsky, M. S. Handcock, A. E. Raftery, and P. D. Hoff. Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Social Networks*, 31(3):204–213, 2009.

- [184] N. J. Krogan, G. Cagney, H. Yu, G. Zhong, X. Guo, A. Ignatchenko, J. Li, S. Pu, N. Datta, A. P. Tikuisis, T. Punna, J. M. Peregrín-Alvarez, M. Shales, X. Zhang, M. Davey, M. D. Robinson, A. Paccanaro, J. E. Bray, A. Sheung, B. Beattie, D. P. Richards, V. Canadien, A. Lalev, F. Mena, P. Wong, A. Starostine, M. M. Canete, J. Vlasblom, S. Wu, C. Orsi, S. R. Collins, S. Chandran, R. Haw, J. J. Rilstone, K. Gandi, N. J. Thompson, G. Musso, P. St Onge, S. Ghanny, M. H. Lam, G. Butland, A. M. Altaf-Ul, S. Kanaya, A. Shilatifard, E. O'Shea, J. S. Weissman, C. J. Ingles, T. R. Hughes, J. Parkinson, M. Gerstein, S. J. Wodak, A. Emili, and J. F. Greenblatt. Global landscape of protein complexes in the yeast *Saccharomyces cerevisiae*. *Nature*, 440(7084):637–643, 2006.
- [185] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the web graph. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 57–65, 2000.
- [186] S. L. Lauritzen. Rasch models with exchangeable rows and columns. In J. M. Bernardo et al., *Bayesian Statistics 7*, pages 215–232. Oxford University Press, 2003.
- [187] S. L. Lauritzen. Exchangeable Rasch matrices. *Rendiconti di Matematica, Serie VII*, 28(1):83–95, 2008.
- [188] S. Lee and C. F. Stevens. General design principle for scalable neural circuits in a vertebrate retina. *Proceedings of the National Academy of Sciences*, 104(31):12931–12935, 2007.
- [189] R. T. A. J. Leenders. Models for network dynamics: A Markovian framework. *Journal of Mathematical Sociology*, 20:1–21, 1995.
- [190] E. A. Leicht, G. Clarkson, K. Shedden, and M. Newman. Large-scale structure of time evolving citation networks. *European Physics Journal B*, 59(1):75–83, 2007.
- [191] J. Leskovec and C. Faloutsos. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 631–636. ACM Press, New York, 2006.
- [192] J. Leskovec, D. Chakrabarti, J. Kleinberg, and C. Faloutsos. Realistic, mathematically tractable graph generation and evolution, using Kronecker multiplication. In *Knowledge Discovery in Databases: PKDD 2005*, volume 3721 of *Lecture Notes in Computer Science*, pages 133–145. Springer Berlin / Heidelberg, 2005.
- [193] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: Densification laws, shrinking diameters and possible explanations. In *Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 177–187. ACM Press, New York, 2005.

- [194] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):2, 2007.
- [195] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. <http://arXiv.org/abs/0812.4905v2>, 2009.
- [196] L. Li, D. Alderson, J. C. Doyle, and W. Willinger. Towards a theory of scale-free graphs: Definition, proper ties, and implications. *Internet Mathematics*, 2(4):431–523, 2005.
- [197] W. Li and A. McCallum. Pachinko allocation: DAG-structured mixture models of topic correlations. In *Proceedings of the 23rd International Conference on Machine Learning*, volume 148 of *ACM International Conference Proceeding Series*, pages 577–584. ACM Press, New York, 2006.
- [198] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *Proceedings of the 12th International Conference on Information and Knowledge Management (CIKM '03)*, pages 556–559. ACM Press, New York, 2003.
- [199] F. Lorrain and H. C. White. Structural equivalence of individuals in social networks. *Journal of Mathematical Sociology*, 1:49–80, 1971.
- [200] R. D. Luce. Connectivity and generalized cliques in sociometric group structure. *Psychometrika*, 15(2):169–190, 1950.
- [201] R. D. Luce. Networks satisfying minimality conditions. *American Journal of Mathematics*, 75(4):825–838, 1953.
- [202] R. D. Luce and A. D. Perry. A method of matrix analysis of group structure. *Psychometrika*, 14(2):95–116, 1949.
- [203] R. D. Luce, J. Macy, Jr., and R. Tagiuri. A statistical model for relational analysis. *Psychometrika*, 20(4):319–327, 1955.
- [204] G. S. Mann, D. Mimno, and A. McCallum. Bibliometric impact measures leveraging topic analysis. In *Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 65–74. ACM Press, New York, 2006.
- [205] J.-M. Marin, K. Mengersen, and C. P. Robert. Bayesian modelling and inference on mixtures of distributions. In D. Dey and C. R. Rao, editors, *Handbook of Statistics 25*, pages 15840–15845. Elsevier Sciences, 2005.
- [206] T. F. Mayer. Parties and networks: Stochastic models for relationship networks. *Journal of Mathematical Sociology*, 10:51–103, 1984.

- [207] A. McCallum, A. Corrada-Emmanuel, and X. Wang. Topic and role discovery in social networks. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 786–791, 2005.
- [208] A. McCallum, X. Wang, and N. Mohanty. Joint group and topic discovery from relations and text. In E. M. Airoldi, D. M. Blei, S. E. Fienberg, A. Goldenberg, E. P. Xing, and A. Zheng, editors, *Statistical Network Analysis: Models, Issues and New Directions*, volume 4503 of *Lecture Notes in Computer Science*, pages 28–44. Springer Berlin / Heidelberg, 2007.
- [209] K. McComb, C. Moss, S. M. Durant, L. Baker, and S. Sayialel. Matriarchs as repositories of social knowledge in African elephants. *Science*, 292(5516):491–494, 2001.
- [210] P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman & Hall/CRC, 2nd edition, 1989.
- [211] J. W. McDonald, P. W. F. Smith, and J. J. Forster. Markov chain Monte Carlo exact inference for social networks. *Social Networks*, 29(1):127–136, 2007.
- [212] M. McGlohon, L. Akoglu, and C. Faloutsos. Weighted graphs and disconnected components: Patterns and a generator. In *Proceedings of the 14th International Conference on Knowledge Discovery and Data Mining*, pages 524–532. ACM Press, New York, 2008.
- [213] M. M. Meyer. Transforming contingency tables. *Annals of Statistics*, 10(4):1172–1181, 1982.
- [214] M. Middendorf, E. Ziv, and C. H. Wiggins. Inferring network mechanisms: The *Drosophila melanogaster* protein interaction network. *Proceedings of the National Academy of Sciences*, 102(9):3192–3197, 2005.
- [215] S. Milgram. The small world problem. *Psychology Today*, 1(1):60–67, 1967.
- [216] D. Mimno and A. McCallum. Mining a digital library for influential authors. In *Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 105–106. ACM Press, New York, 2007.
- [217] N. Mishra, R. Schreiber, I. Stanton, and R. E. Tarjan. Finding strongly-knit clusters in social networks. *Internet Mathematics*, 5(1-2):155–174, 2008.
- [218] M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1(2):226–251, 2004.
- [219] M. Molloy and B. Reed. A critical point for random graphs with a given degree sequence. *Random Structures and Algorithms*, 6(2-3):161–180, 1995.

- [220] M. Molloy and B. Reed. The size of the largest component of a random graph on a fixed degree sequence. *Combinatorics, Probability and Computing*, 7:295–306, 1998.
- [221] J. Moreno. *Who Shall Survive?* Nervous and Mental Disease Publishing Company, Washington, D.C., 1934.
- [222] M. Morris and M. Kretzschmar. Concurrent partnerships and transmission dynamics in networks. *Social Networks*, 17(3–4):299–318, 1995.
- [223] M. Morris, M. S. Handcock, W. C. Miller, C. A. Ford, J. L. Schmitz, M. M. Hobbs, M. S. Cohen, K. M. Harris, and J. R. Udry. Prevalence of HIV infection among young adults in the United States: Results from the Add Health Study. *American Journal of Public Health*, 96(6):1091–1097, 2006.
- [224] M. Morris, M. S. Handcock, and D. R. Hunter. Specification of exponential-family random graph models: Terms and computational aspects. *Journal of Statistical Software*, 24(4), 2008. <http://www.jstatsoft.org/v24/i04>.
- [225] Q. Morris, B. Frey, and C. Paige. Denoising and untangling graphs using degree priors. In *Advances in Neural Information Processing Systems (NIPS)*, volume 16. MIT Press, Cambridge, MA, 2003.
- [226] S. Mostafavi, D. Ray, D. Warde-Farley, C. Grouios, and Q. Morris. GeneMANIA: A real-time multiple association network integration algorithm for predicting gene function. *Genome Biology*, 9(Suppl. 1):S4, 2008.
- [227] E. Nabieva, K. Jim, A. Agarwal, B. Chazelle, and M. Singh. Whole-proteome prediction of protein function via graph-theoretic analysis of interaction maps. *Bioinformatics*, 21(Suppl. 1):i302–i310, 2005.
- [228] R. M. Neal. *Bayesian Learning for Neural Networks*, volume 118 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 1996.
- [229] J. Neville and D. Jensen. Collective classification with relational dependency networks. In *Proceedings of the 2nd Multi-Relational Data Mining Workshop, 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003.
- [230] J. Neville, D. Jensen, L. Friedland, and M. Hay. Learning relational probability trees. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 625–630. ACM Press, New York, 2003.
- [231] M. Newman, A.-L. Barabási, and D. J. Watts, editors. *The Structure and Dynamics of Networks*. Princeton University Press, 2006.
- [232] M. E. J. Newman. Detecting community structure in networks. *European Physics Journal B*, 38(2):321–330, 2004.

- [233] M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.
- [234] M. E. J. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, 74(3):036104, 2006.
- [235] M. E. J. Newman, D. J. Watts, and S. H. Strogatz. Random graph models of social networks. *Proceedings of the National Academy of Science*, 99(Suppl. 1):2566–2572, 2002.
- [236] K. Nowicki and T. A. B. Snijders. Estimation and prediction for stochastic blockstructures. *Journal of the American Statistical Association*, 96(455):1077–1087, 2001.
- [237] M. Nunkesser and D. Sawitzki. Blockmodels. In U. Brandes and T. Erlebach, editors, *Network Analysis*, volume 3418 of *Lecture Notes in Computer Science*, pages 253–292. Springer Berlin / Heidelberg, 2005.
- [238] J. O’Madadhain, P. Smyth, and L. Adamic. Learning predictive models for link formation. In *Proceedings of the International Sunbelt Social Network Conference*, 2005.
- [239] J. Park and M. E. J. Newman. Statistical mechanics of networks. *Physical Review E*, 70:066117, 2004.
- [240] J. Park and M. E. J. Newman. Solution of the two-star model of a network. *Physics Reviews E*, 70:066146, 2004.
- [241] P. E. Pattison and S. S. Wasserman. Logit models and logistic regressions for social networks: II. Multivariate relations. *British Journal of Mathematical and Statistical Psychology*, 52(2):169–193, 1999.
- [242] D. M. Pennock, G. W. Flake, S. Lawrence, E. J. Glover, and C. L. Giles. Winners don’t take all: Characterizing the competition for links on the web. *Proceedings of the National Academy of Sciences*, 99(8):5207–5211, 2002.
- [243] B. Pittel. On tree census and the giant component in sparse random graphs. *Random Structures Algorithms*, 1(3):311–342, 1990.
- [244] R. Radner and A. Titter. Communication in networks. Technical Report Ec2098, Cowles Commission, University of Chicago, 1954.
- [245] O. Ratmann, O. Jørgensen, T. Hinkley, M. P. H. Stumpf, S. Richardson, and C. Wiuf. Using likelihood-free inference to compare evolutionary dynamics of the protein networks of *H. pylori* and *P. falciparum*. *PLoS Computational Biology*, 3(11):2266–2278, 2007.
- [246] O. Ratmann, C. Wiuf, and J. W. Pinney. From evidence to inference: Probing the evolution of protein interaction networks. *HFSP Journal*, 3(5):290–306, 2009.

- [247] P. Ravikumar. *Approximate Inference, Structure Learning and Feature Estimation in Markov Random Fields*. PhD thesis, Machine Learning Department, School of Computer Science, Carnegie Mellon University, 2007.
- [248] T. Regulý, A. Breitzkreutz, L. Boucher, B.-J. Breitzkreutz, G. C. Hon, C. L. Myers, A. Parsons, H. Friesen, R. Oughtred, A. Tong, C. Stark, Y. Ho, D. Botstein, B. Andrews, C. Boone, O. G. Troyanskya, T. Ideker, K. Dolinski, N. N. Batada, and M. Tyers. Comprehensive curation and analysis of global interaction networks in *Saccharomyces cerevisiae*. *Journal of Biology*, 5(4):11, 2006.
- [249] E. Reid and H. Chen. Mapping the contemporary terrorism research domain: Researchers, publications, and institutions analysis. In *Intelligence and Security Informatics*, volume 3495 of *Lecture Notes in Computer Science*, pages 322–339. Springer Berlin / Heidelberg, 2005.
- [250] E. Reid, J. Qin, W. Chung, J. Xu, Y. Zhou, R. Schumaker, M. Sageman, and H. Chen. Terrorism knowledge discovery project: A knowledge discovery approach to addressing the threats of terrorism. In *Intelligence and Security Informatics*, volume 3073 of *Lecture Notes in Computer Science*, pages 125–145. Springer Berlin / Heidelberg, 2004.
- [251] A. Rinaldo, S. E. Fienberg, and Y. Zhou. On the geometry of discrete exponential families with application to exponential random graph models. *Electronic Journal of Statistics*, 3:446–484, 2009.
- [252] J. M. Roberts, Jr. Simple methods for simulating sociomatrices with given marginal totals. *Social Networks*, 22(3):273–283, 2000.
- [253] G. L. Robins and P. E. Pattison. Random graph models for temporal processes in social networks. *Journal of Mathematical Sociology*, 25:5–41, 2001.
- [254] G. L. Robins, P. E. Pattison, and S. S. Wasserman. Logit models and logistic regressions for social networks: III. Valued relations. *Psychometrika*, 64(3):371–394, 1999.
- [255] G. L. Robins, P. E. Pattison, and J. Woolcock. Missing data in networks: Exponential random graph (p^*) models for networks with non-respondents. *Social Networks*, 26(3):257–283, 2004.
- [256] G. L. Robins, T. A. B. Snijders, P. Wang, M. S. Handcock, and P. E. Pattison. Recent developments in exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):192–215, 2007.
- [257] T. T. Rogers and J. L. McClelland. *Semantic Cognition: A Parallel Distributed Processing Approach*. MIT Press, Cambridge, MA, 2004.
- [258] M. J. Salganik and D. D. Heckathorn. Sampling and estimation in hidden populations using respondent-driven sampling. *Sociological Methodology*, 34:193–239, 2004.

- [259] F. S. Sampson. *A Novitiate in a Period of Change: An Experimental and Case Study of Social Relationships*. PhD thesis, Cornell University, 1968.
- [260] O. Sandberg. *Searching in a Small World*. PhD thesis, Division of Mathematical Statistics, Department of Mathematical Sciences, Chalmers University of Technology and Göteborg University, Göteborg, Sweden, 2005.
- [261] O. Sandberg. Neighbor selection and hitting probability in small-world graphs. *Annals of Applied Probability*, 18(5):1771–1793, 2008.
- [262] O. Sandberg and I. Clarke. The evolution of navigable small-world networks. <http://arXiv.org/abs/cs/0607025>, 2006.
- [263] P. Sarkar and A. W. Moore. Dynamic social network analysis using latent space models. In *Advances in Neural Information Processing Systems (NIPS)*, volume 18, pages 1145–1152. MIT Press, Cambridge, MA, 2005.
- [264] P. Sarkar and A. W. Moore. Dynamic social network analysis using latent space models. *SIGKDD Explorations: Special Edition on Link Mining*, 7(2):31–40, 2005.
- [265] P. Sarkar, S. M. Siddiqi, and G. J. Gordon. A latent space approach to dynamic embedding of co-occurrence data. In *Proceedings of the 11th International Conference on Artificial Intelligence and Statistics (AI-STATS '07)*, 2007.
- [266] C. R. Shalizi, M. F. Camperi, and K. L. Klinkner. Discovering functional communities in dynamical networks. In E. M. Airoldi, D. M. Blei, S. E. Fienberg, A. Goldenberg, E. P. Xing, and A. Zheng, editors, *Statistical Network Analysis: Models, Issues and New Directions*, volume 4503 of *Lecture Notes in Computer Science*, pages 140–157. Springer Berlin / Heidelberg, 2007.
- [267] B. Shneiderman and A. Aris. Network visualization by semantic substrates. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):733–740, 2006.
- [268] G. Simmel and K. H. Wolff. *The Sociology of Georg Simmel*. The Free Press, New York, 1950.
- [269] H. A. Simon. On a class of skew distribution functions. *Biometrika*, 42(3–4):425–440, 1955.
- [270] B. Singer and S. Spilerman. Social mobility models for heterogenous populations. *Sociological Methodology*, 5:356–401, 1973–1974.
- [271] B. Singer and S. Spilerman. The representation of social processes by Markov models. *The American Journal of Sociology*, 82(1):1–54, 1976.
- [272] T. A. B. Snijders. The transition probabilities of the reciprocity model. *Journal of Mathematical Sociology*, 23(4):241–253, 1999.

- [273] T. A. B. Snijders. The statistical evaluation of social network dynamics. *Sociological Methodology*, 31:361–395, 2001.
- [274] T. A. B. Snijders. Accounting for degree distributions in empirical analysis of network dynamics. In R. L. Breiger, K. M. Carley, and P. E. Pattison, editors, *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, pages 146–161. The National Academies Press, Washington, D.C., 2003.
- [275] T. A. B. Snijders. Models for longitudinal network data. In P. J. Carrington, J. Scott, and S. S. Wasserman, editors, *Models and Methods in Social Network Analysis*, chapter 11. Cambridge University Press, New York, 2005.
- [276] T. A. B. Snijders. Statistical methods for network dynamics. In S. R. Luchini et al., editors, *Proceedings of the XLIII Scientific Meeting, Italian Statistical Society*, pages 281–296, Padova: CLEUP, 2006.
- [277] T. A. B. Snijders and K. Nowicki. Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of Classification*, 14(1):75–100, 1997.
- [278] T. A. B. Snijders and M. A. J. van Duijn. Simulation for statistical inference in dynamic network models. In R. Conte, R. Hegselmann, and P. Terna, editors, *Simulating Social Phenomena*, pages 493–512. Springer, Berlin, 1997.
- [279] T. A. B. Snijders and M. A. J. van Duijn. Conditional maximum likelihood estimation under various specifications of exponential random graph models. In J. Hagberg, editor, *Contributions to Social Network Analysis, Information Theory, and Other Topics in Statistics; A Festschrift in honour of Ove Frank*, pages 117–134. Department of Statistics, University of Stockholm, Stockholm, Sweden, 2002.
- [280] T. A. B. Snijders, P. E. Pattison, G. L. Robins, and M. S. Handcock. New specifications for exponential random graph models. *Sociological Methodology*, 36:99–153, 2006.
- [281] R. Solomonoff and A. Rapoport. Connectivity of random nets. *Bulletin of Mathematical Biology*, 13(2):107–117, 1951.
- [282] S. Spilerman. Structural analysis and the generation of sociograms. *Behavioral Science*, 11:312–318, 1966.
- [283] M. Stephens. Bayesian analysis of mixtures with an unknown number of components—an alternative to reversible jump methods. *Annals of Statistics*, 28(1):40–74, 2000.
- [284] D. Stork and W. Richards. Nonrespondents in communication network studies. *Group & Organization Management*, 17(2):193–209, 1992.
- [285] D. B. Stouffer, R. D. Malmgren, and L. A. N. Amaral. Comment on Barabási, Nature 435, 207 (2005). <http://arXiv.org/abs/physics/0510216>, 2005.

- [286] D. B. Stouffer, R. D. Malmgren, and L. A. N. Amaral. Log-normal statistics in e-mail communication patterns. <http://arXiv.org/abs/physics/0605027>, 2008.
- [287] D. Strauss and M. Ikeda. Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association*, 85(409):204–212, 1990.
- [288] M. P. H. Stumpf and T. Thorne. Multi-model inference of network properties from incomplete data. *Journal of Integrative Bioinformatics*, 3(2):32, 2006. http://journal.imbio.de/index.php?paper_id=32.
- [289] M. P. H. Stumpf, C. Wiuf, and R. M. May. Subnets of scale-free networks are not scale-free: Sampling properties of networks. *Proceedings of the National Academy of Sciences*, 102(12):4221–4224, 2005.
- [290] S. Swasey. Netflix awards \$1 million Netflix prize and announces second \$1 million challenge. Wall Street Journal, September 21, 2009.
- [291] K. Tarassov, V. Messier, C. R. Landry, S. Radinovic, M. M. Serna Molina, I. Shames, Y. Malitskaya, J. Vogel, H. Bussey, and S. W. Michnick. An in vivo map of the yeast protein interactome. *Science*, 320(5882):1465–1470, 2008.
- [292] H. M. Taylor and S. Carlin. *An Introduction to Stochastic Modeling*. Academic Press, New York, 3rd edition, 1998.
- [293] S. K. Thompson. Adaptive web sampling. *Biometrics*, 62(4):1224–1234, 2006.
- [294] S. K. Thompson. Targeted random walk designs. *Survey Methodology*, 32(1):11–24, 2006.
- [295] S. K. Thompson and O. Frank. Model-based estimation with link-tracing sampling designs. *Survey Methodology*, 26(1):87–98, 2000.
- [296] S. K. Thompson and G. A. F. Seber. *Adaptive Sampling*. Wiley, New York, 1996.
- [297] D. M. Titterton, A. F. M. Smith, and U. E. Makov. *Statistical Analysis of Finite Mixture Distributions*. John Wiley & Sons, New York, 1986.
- [298] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32(4):425–443, 1969.
- [299] R. J. Udry. *The National Longitudinal Study of Adolescent Health: (Add health) Waves I and II, 1994–1996; Wave III, 2001–2002*. Technical report, Carolina Population Center, University of North Carolina, Chapel Hill, 2003.
- [300] P. Uetz, L. Giot, G. Cagney, T. A. Mansfield, R. S. Judson, J. R. Knight, D. Lockshon, V. Narayan, M. Srinivasan, P. Pochart, A. Qureshi-Emili, Y. Li, B. Godwin, D. Conover, T. Kalbfleisch, G. Vijayadamodar, M. Yang, M. Johnston, S. Fields, and

- J. M. Rothberg. A comprehensive analysis of protein-protein interactions in *Saccharomyces cerevisiae*. *Nature*, 403(6770):623–627, 2000.
- [301] M. A. J. van Duijn, T. A. B. Snijders, and B. J. H. Zijlstra. p_2 : A random effects model with covariates for directed graphs. *Statistica Neerlandica*, 58(2):234–254, 2004.
- [302] M. A. J. van Duijn, K. J. Gile, and M. S. Handcock. A framework for the comparison of maximum pseudo-likelihood and maximum likelihood estimation of exponential family random graph models. *Social Networks*, 31(1):52–62, 2009.
- [303] E. A. Vance, E. A. Archie, and C. J. Moss. Social networks in African elephants. *Computational & Mathematical Organization Theory*, <http://www.springerlink.com/content/enpk5g428272927m>, 2008. To appear in print, 2009.
- [304] A. Vázquez, J. G. Oliveira, Z. Dezső, K. Goh, I. Kondor, and A.-L. Barabási. Modeling bursts and heavy tails in human dynamics. *Physical Review E*, 73:036127, 2006.
- [305] E. Volz and D. D. Heckathorn. Probability based estimation theory for respondent driven sampling. *Journal of Official Statistics*, 24(1):79–97, 2008.
- [306] E. Volz and L. A. Meyers. Epidemic thresholds in dynamic contact networks. *Journal of the Royal Society Interface*, 6(32):233–241, 2009.
- [307] C. von Mering, R. Krause, B. Snel, M. Cornell, S. G. Oliver, S. Fields, and P. Bork. Comparative assessment of large-scale data sets of protein-protein interactions. *Nature*, 417(6887):399–403, 2002.
- [308] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.
- [309] A. M. Walczak, A. Mugler, and C. H. Wiggins. A stochastic spectral analysis of transcriptional regulatory cascades. *Proceedings of the National Academy of Sciences*, 106(16):6529–6534, 2009.
- [310] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos. Epidemic spreading in real networks: An eigenvalue viewpoint. In *Proceedings of the 22nd International Symposium on Reliable Distributed Systems (SRDS '03)*, pages 25–34, 2003.
- [311] Y. Y. Wang and G. Y. Wong. Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*, 82(397):8–19, 1987.
- [312] S. S. Wasserman. *Stochastic Models for Directed Graphs*. PhD thesis, Department of Statistics, Harvard University, 1977.
- [313] S. S. Wasserman. Analyzing social networks as stochastic processes. *Journal of the American Statistical Association*, 75(370):280–294, 1980.

- [314] S. S. Wasserman and C. Anderson. Stochastic a posteriori blockmodels: Construction and assessment. *Social Networks*, 9(1):1–36, 1987.
- [315] S. S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.
- [316] S. S. Wasserman and P. E. Pattison. Logit models and logistic regression for social networks: I. An introduction to Markov graphs and p^* . *Psychometrika*, 61(3):401–425, 1996.
- [317] S. S. Wasserman, G. L. Robins, and D. Steinley. Statistical models for networks: A brief review of some recent research. In E. M. Airoldi, D. M. Blei, S. E. Fienberg, A. Goldenberg, E. P. Xing, and A. X. Zheng, editors, *Statistical Network Analysis: Models, Issues and New Directions*, volume 4503 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2007.
- [318] D. J. Watts. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press, 1999.
- [319] D. J. Watts. *Six Degrees: The Science of a Connected Age*. W. W. Norton & Company, New York, 2003.
- [320] D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.
- [321] H. C. White. Search parameters for the small world problem. *Social Forces*, 49(2):259–264, 1970.
- [322] H. C. White, S. A. Boorman, and R. L. Breiger. Social structure from multiple networks. I. Blockmodels of roles and positions. *The American Journal of Sociology*, 81(4):730–780, 1976.
- [323] R. J. Williams and N. D. Martinez. Simple rules yield complex food webs. *Nature*, 404(6774):180–183, 2000.
- [324] W. Willinger, D. Alderson, and J. C. Doyle. Mathematics and the internet: A source of enormouse confusion and great potential. *Notices of the American Mathematical Society*, 56(5):586–599, 2009.
- [325] C. Wiuf and M. P. H. Stumpf. Binomial subsampling. *Journal of the Royal Society, Series A*, 462(2068):1181–1195, 2006.
- [326] C. Wiuf, M. Brameier, O. Hagberg, and M. P. H. Stumpf. A likelihood approach to analysis of network data. *Proceedings of the National Academy of Sciences*, 103(20):7566–7570, 2006.

- [327] S. L. Wong, L. V. Zhang, A. H. Y. Tong, Z. Li, D. S. Goldberg, O. D. King, G. Lesage, M. Vidal, B. Andrews, H. Bussey, C. Boone, and F. P. Roth. Combining biological networks to predict genetic interactions. *Proceedings of the National Academy of Sciences*, 101(44):15682–15687, 2004.
- [328] H. Yu, P. Braun, M. A. Yildirim, I. Lemmens, K. Venkatesan, J. Sahalie, T. Hirozane-Kishikawa, F. Gebreab, N. Li, N. Simonis, T. Hao, J. F. Rual, A. Dricot, A. Vazquez, R. R. Murray, C. Simon, L. Tardivo, S. Tam, N. Svrikapa, C. Fan, A. S. de Smet, A. Motyl, M. E. Hudson, J. Park, X. Xin, M. E. Cusick, T. Moore, C. Boone, M. Snyder, F. P. Roth, A.-L. Barabási, J. Tavernier, D. E. Hill, and M. Vidal. High-quality binary protein interaction map of the yeast interactome network. *Science*, 322(5898):104–110, 2008.
- [329] G. U. Yule. A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis, F.R.S. *Philosophical Transactions of the Royal Society of London, Series B, Containing Papers of a Biological Character*, 213:21–87, 1925.
- [330] W. W. Zachary. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*, 33:452–473, 1977.
- [331] A. Zheng and A. Goldenberg. A generative model for dynamic contextual friendship networks. Technical report, Machine Learning Department, Carnegie Mellon University, 2006.
- [332] X. Zhu, M. Gerstein, and M. Snyder. Getting connected: Analysis and principles of biological networks. *Genes Development*, 21(9):1010–1024, 2007.
- [333] B. J. H. Zijlstra, M. A. J. van Duijn, and T. A. B. Snijders. The multilevel p_2 model: A random effects model for the analysis of multiple social networks. *Methodology*, 2(1):42–47, 2006.