1

Introduction

Dean Lusher, Johan Koskinen, and Garry Robins

Exponential random graph models[1] (ERGMs) are a class of statistical model for social networks. They account for the presence (and absence) of network ties, and so provide a model for network structure. An ERGM models a given network in terms of small local tie-based structures, such as reciprocated ties and triangles. A social network can be thought of as being built of these local patterns of ties, called "network configurations," which correspond to the parameters in the model. Moreover, these configurations can be considered to arise from local social processes, whereby actors in the network form connections in response to other ties in their social environment. ERGMs are a principled statistical approach to modeling social networks. They are theory driven in that their use requires the researcher to consider the complex, intersecting, and, indeed, potentially competing theoretical reasons why the social ties in the observed network have arisen. For instance, does a given network structure occur due to processes of homophily, reciprocity, transitivity, or a combination of these? By including such parameters together in the one model, a researcher can test these effects one against the other, and so infer the social processes that have built the network. Being a statistical model, an ERGM permits inferences about whether, in our network of interest, there are significantly more (or less) reciprocated ties, or triangles (for instance), than we would expect.

ERGMs are fast becoming recognized as one of the central approaches in analyzing social networks. In this short introductory chapter, we describe the intent of this book, how it is structured, how it may be read, the software resources available, and the knowledge that we expect readers to have before following our text.

[1] Throughout this book, we use this now-established convenient term instead of the more correct (but cumbersome) "exponential family random graph models."

1.1 Intent of This Book

In this book, we aim to introduce ERGMs in a way that takes the reader from the basic theoretical assumptions underlying the models, through the technical specifications and statistical detail, to applied examples illustrating how various substantive research questions may be investigated and tested empirically. The goals of this book are twofold: to describe ERGMs and to demonstrate how recent methodological developments allow us to address social network research questions in new and powerful ways. A specific target audience is the growing number of social scientists who are interested in statistical models for social networks and network-based social processes. We aim to provide an intuitive understanding of these models for those readers who may be unfamiliar with ERGMs. We also focus on delivering sufficient technical detail for those with a social network methodological background and who are interested in a deeper understanding of the modeling and estimation. We hope to lay bare the value of a statistical modeling approach in answering core questions about interactive social processes.

1.2 Software and Data

The book includes references to data and software that can be downloaded so that readers can reproduce some of the applications in a "hands-on" fashion. Some parts of this book make explicit use of the PNet suite of programs for ERGMs (Wang, Robins, & Pattison, 2009), and we tend to use PNet terminology for parameters and the like, but the exposition is in no way contingent on the software used (as long as the software does deal with ERGMs, naturally). The estimation algorithms used in PNet and in the R package `statnet` (Handcock et al., 2003; Handcock et al., 2008) are both described in Chapter 12, as is the principle of `Bergm` (Caimo & Friel, 2011). SIENA 3 can be used for fitting ERGMs to cross-sectional data (indeed, the main parts of PNet draw heavily on SIENA) and longitudinal models can be analyzed in RSIENA (or SIENA 4). Models for multiple networks, longitudinal networks, and bipartite networks can be estimated from the PNet suite of programs, XPNet, LPNet, and BPNet, respectively. The autologistic actor attribute models require iPNet. PNet and the example data used in this book can be downloaded to help you work through the chapters in a practical, hands-on fashion (http://www.sna.unimelb.edu.au). The package `statnet` is an ERGM estimation program in the R environment and can be installed from CRAN in standard R fashion. A useful introduction to the `statnet` package is given in a special issue of *Journal of Statistical*

Software (Hunter et al., 2008), and the worked-through examples of Goodreau et al. (2008) are particularly helpful.

1.3 Structure of the Book

This book is divided into four main sections that broadly map onto the rationale and theory, method, and application of ERGMs, with a fourth overview section discussing future directions. We consider that theory, method, and empirical work are fundamentally intertwined and interdependent. Thus, you will note that the first section, despite its flavor of conceptual discussion, still addresses methodological issues and provides empirical examples, whereas theoretical concepts permeate the methods section.

1.3.1 Section I: Rationale

Section I of this book provides an intuitive introduction to ERGMs by connecting to various aspects of social network (and social science) theory, the thinking behind the methods, and the empirical examination of research questions. Chapter 2 provides a very general initial description of ERGMs, including some broad definitions and a discussion of some central elements of the ERGM approach. Chapter 3 introduces some important network ideas about the formation of social ties and social structure more generally, and explains how these relate to an ERGM analysis. Chapter 4 introduces some central methodological details of ERGMs in intuitive terms for readers unfamiliar with the models, with more technical details delayed until Section II on methods. Finally, in this section, Chapter 5 provides an early and simple example of applying an ERGM to network data (an example we explore in greater detail in Chapter 13). In this way, we illustrate the type of inferences and interpretations that can be made about network structure using ERGMs. Our aims in this first section are not only to describe theoretical issues important to ERGMs but also to build intuitions for those coming to this modeling approach for the first time, before they encounter the more technical detail in Section II.

1.3.2 Section II: Methods

Section II of this book presents the basic statistical framework of ERGMs. Throughout the book, the case of a single observation on a (binary, unipartite) network is taken as a standard point of reference, and the technical details of ERGMs are introduced in the context of such a case in Chapter 6. Chapter 7 seeks to give a deeper understanding of ERGMs by way of the notion of the dependence graph. From here, a range of more specific ERGMs are presented, beginning with models that include social selection, dyadic covariate, and geospatial effects in Chapter 8. Chapter 9 covers autologistic actor attribute models (ALAAMs), which are a version of social influence models, for the attributes of the actors. The next chapters describe models for multiple networks and bipartite networks (Chapter 10) as well as longitudinal network data (Chapter 11). The section concludes with descriptions of simulation, estimation, and goodness-of-fit procedures (Chapter 12), together with some illustrations (Chapter 13).

1.3.3 Section III: Applications

Section III of this book demonstrates how ERGMs can be applied empirically to answer social network research questions. In this section, we present an application for each of the different types of ERGMs presented previously. The chapters illustrate to the reader the range of possibilities, issues, and general frame of reference a researcher engages in when using ERGMs. The chapters are presented as separate self-contained case studies, but together they cover a set of network theoretical issues and features of the modeling framework.

In these chapters, we seek to show why and how the models can be used to answer novel and theoretically important questions, formulated in a relational framework. The range of issues covered gives some indication of the breadth of approach that ERGMs offer. For instance, Lusher and Robins (Chapter 14) note how the individual perceptions of attitudes held by others in the network have an independent association on the formation of network ties in two different contexts. They draw on theory from social psychology in their exposition. Lomi and Pallotti (Chapter 15) examine various path closure effects and, ultimately, show the importance of structural equivalence–type patterns in the transfer of patients among hospitals. Zhao and Rank (Chapter 16) examine how multiple types of relations within the one organization relate to each other. Drawing on various leadership theories, Kalish and Luria (Chapter 17) construe leadership as a relational phenomenon and show how leadership networks differ from many other types of social networks. Daraganova and Pattison (Chapter 18) test competing hypotheses about the impact of spatial factors and social ties on unemployment. Igarashi (Chapter 19) demonstrates the importance of homophily in a longitudinal analysis of networks involving different methods of communication. Using bipartite network models, Harrigan and Bond (Chapter 20) show how different forms of capital lead to different types of linking behavior. In the final applications chapter, Quintane

(Chapter 21) examines differences in self-reported ties and observed relational behaviors.

1.3.4 Section IV: Future

There is ongoing methodological development: this book is by no means the last word on ERGMs. To conclude this book, Pattison and Snijders look to future directions for ERGM research and methods.

1.4 How To Read This Book

We suggest that a cover-to-cover reading may not be the best way to approach the book for all readers. Readers with less statistical knowledge might begin with all four chapters in Section I: Rationale, which provides the general basis of ERGMs. Following this, the introductory methods chapter on the fundamentals of ERGMs (Chapter 6) is important, although it could be tackled later if need be. At the very least, the social selection models (Chapter 8) and the illustrations chapter (Chapter 13) in Section II: Methods should be read before moving to relevant chapters of Section III: Applications, and Section IV: Future.

Some readers, however, may want to understand ERGMs from the point of view of the statistical model. Thus, they may want to jump straight to the methods section, especially Chapters 6 and 7, and then work their way through Chapter 12 before returning to the conceptual issues of Section I.

For some readers whose interests are principally directed to fitting models empirically, the chapter on dependency (Chapter 7), which details important conceptual issues with regard to ERGMs, may be skipped at the first or first few readings. For those interested in the theoretical underpinning of the statistical model, however, this chapter is important.

For readers with some familiarity with the ERGM framework and an interest in a specific type of model, it is probably good to have skimmed through Section I first (and perhaps the first chapter of Section II). From there, the reader may go on to the applications chapter of choice, and then work him- or herself back to the relevant chapters in Section II, for clarification of details. You may, for example, have a particular interest in social influence–type models, in which case you might want to start with Chapter 18 and refer back to Chapter 9 in Section II, when necessary.

If you intend to fit your own ERGMs, it is necessary to understand how the models are fitted, whether as introduced in Chapters 4 and 5 of Section I and then Chapter 13, or more in-depth as in Chapter 12 (simulation, estimation, and goodness of fit). Specific applications chapters in Section III may also prove quite useful as a guide. Finally, once you have fitted a model, or several, you may want to delve deeper into the details of Section II.

1.5 Assumed Knowledge of Social Network Analysis

We assume basic familiarity with social network analysis and that the reader understand the general concepts and terminology used in network and graph theory. For this purpose, we refer the reader to one of the many introductory texts, including Wasserman and Faust (1994), Prell (2011), Knoke and Yang (2008), Scott (2000), de Nooy, Mrvar and Batageli (2005), van Duijn and Vermunt (2006), and Hanneman and Riddle (2005), for a more detailed introduction. We assume that the reader is familiar with standard statistical techniques such as regression analysis and logistic regression. Kolaczyk (2009) and Knoke and Yang (2008) both have sections introducing ERGMs. The former provides a good technical and comprehensive treatment of statistical models for networks, which might be useful as further reading for the purposes of introducing alternative models. Reference is occasionally made to more advanced issues in statistics that the interested reader can follow-up on if a more comprehensive picture is desired.

Section I

Rationale

2

What Are Exponential Random Graph Models?

Garry Robins and Dean Lusher

2.1 Exponential Random Graph Models: A Short Definition

Exponential random graph models (ERGMs) are statistical models for network structure, permitting inferences about how network ties are patterned. Put another way, ERGMs are tie-based models for understanding how and why social network ties arise. This focus aligns ERGMs with a principal goal of much empirical social network research, which is to understand a given "observed" network structure (i.e., a network on which a researcher has collected data), and so to obtain insight into the underlying processes that create and sustain the network-based social system.

Much of social network analysis has been concerned with representing the network, a graph G, through various summary measures. From the literature, the reader may be familiar with summary measures $z(G)$ such as the number of edges in G, the number of mutual ties, centrality measures, triad census, and so on. We call these summary measures "network statistics," and in mathematical terms, the ERGM assigns probability to graphs according to these statistics:

$$\mathrm{P}_\theta(G) = c e^{\theta_1 z_1(G) + \theta_2 z_2(G) + \cdots + \theta_p z_p(G)}.$$

The probability of a given network G is given by a sum of network statistics (the zs in this expression) weighted, just as in a regression, by parameters (the θs) inside an exponential (and where c is a normalizing constant). The network statistics are counts of the number of network configurations in the given network G, or some function of those counts. These configurations are small, local subgraphs in the network. In short, the probability of the network depends on how many of those configurations are present, and the parameters inform us of the importance of each configuration.

This expression is explained in much more detail in Section II. However, because the mathematical features are not important for our purposes here, we hope to explain ERGMs in a relatively intuitive way in this introductory section.

To put it as simply as possible, a researcher specifies an ERGM by choosing a set of configurations of theoretical interest. As we will see, there are many sets of plausible configurations that can be used. Then, by applying this particular model to an observed social network, parameters are estimated. This permits inferences about the configurations – the network patterns – in the data, and this in turn allows inferences about the type of social processes that are important in creating and sustaining the network. Thus, ERGMs provide a methodology to investigate network structures and processes empirically.

Note that there is not just one ERGM – there are whole classes of them. The researcher has to choose the specification of an ERGM for the data (just as a researcher has to choose the variables to include in a regression). For an ERGM, the specifications involve choices of configurations that the researcher believes are relevant to the network structure. Although there are some standard ways to do this, the choices are ultimately based on theories about how ties come into being and appear in regular patterns. We discuss some of these theories in greater length in Chapter 3. However, an ERGM itself carries some metatheory about networks, a conceptualization of a social network, and how it is created.

2.2 ERGM Theory

The use of an ERGM is consistent with some basic theoretical assumptions about social networks:

- Social networks are locally emergent.
- Network ties not only self-organize (i.e., there are dependencies between ties), but they are also influenced by actor attributes and other exogenous factors.
- The patterns within networks can be seen as evidence for ongoing structural processes.
- Multiple processes can operate simultaneously.
- Social networks are structured, yet stochastic.

Each idea is discussed more extensively in the following chapters. What is evident is that ERGMs are not "theory free." The preceding list details certain claims consistent with the use of ERGMs that may shape the sort of social network questions that can be asked. Yet many of these claims are not exclusive to ERGMs and overlap with other thinking about social networks. For instance, an ERGM views social tie formation as locally constructed, and much social network theory refers to local processes (e.g., reciprocity, transitivity, homophily – see Chapter 3). On the issue of multiple processes, Monge and Contractor (2003) argued that network research should be multitheoretical – that is, research should examine multiple theoretical perspectives at the same time, and, indeed, ERGMs can do this (see Section 3.1.5).

An explicit and particularly important feature of ERGMs is that network ties depend on one another (i.e., there is network self-organization); thus, the presence of one tie may affect the presence of others (see Sections 3.1.4 and 3.2.1 for general details, and Chapters 6 and 7 for more specific details). Of course, it is well understood that within a social network, individuals are by definition interdependent. In a seminal social networks article, White, Boorman, and Breiger (1976) famously lamented the divergence of theoretical and methodological perspectives on social interaction. On the one hand, they argued, theories are largely concerned with interaction; however, on the other hand, in empirical practice, researchers revert to aggregating individuals by categories. Thus, it is insufficient when we analyze such networks to consider individuals as unrelated "units of analysis" and instead more realistic to consider them as "actors in social relations" (Abbott, 1997, 1152). ERGMs take interdependency one step further by supposing that there is interdependence between network ties. It is in this way that network ties are construed as forming important patterns – the "configurations" parameterized in an ERGM. It is a theoretical and empirical task to delineate the various forms of dependence that are exhibited in actual social structures. We regard this as social network theory at a fundamental level.[1] Dependency among network ties is discussed at many stages of this book, particularly in Chapter 7.

Although there are certain assumptions that one engages when using ERGMs, there is also considerable freedom available to the researcher. As noted previously, there is not just one ERGM, but many. A researcher chooses which model to use by selecting which network structures are important. The utility, as well as the challenge, of ERGMs is the specification of a particular theory (or theories) into social network terms – specifically, in articulating network configurations that, in isolation or combination, reflect a relevant theoretical concept. We hope it is clear

[1] We are referring here to the particular way in which dependence is modeled for cross-sectional data. There are a number of other cross-sectional statistical models that also cater to dependency-related issues such as the random effects *p*2 model (van Duijn, Snijders, & Zijlstra, 2004) and latent social space/variable models (Airoldi et al., 2008; Handcock, Raftery, & Tantrum, 2007; Hoff, Raftery, & Handcock, 2002; Schweinberger & Snijders, 2007). The ERGM approach to dependency, however, is arguably the first and most explicit. See Snijders (2011) for an overview of statistical approaches to network analysis.