
Learning to Generate Networks

James Atwood

School of Computer Science
University of Massachusetts Amherst
Amherst, MA
jatwood@cs.umass.edu

Don Towsley

School of Computer Science
University of Massachusetts Amherst
Amherst, MA
towsley@cs.umass.edu

Krista Gile

Department of Mathematics and Statistics
University of Massachusetts Amherst
Amherst, MA
gile@math.umass.edu

David Jensen

School of Computer Science
University of Massachusetts Amherst
Amherst, MA
jensen@cs.umass.edu

Abstract

We investigate the problem of learning to generate complex networks from data. Specifically, we consider whether deep belief networks, dependency networks, and members of the exponential random graph family can learn to generate networks whose complex behavior is consistent with a set of input examples. We find that the deep model is able to capture the complex behavior of small networks, but that no model is able capture this behavior for networks with more than a handful of nodes.

1 Introduction

Researchers in network science have long studied the behavior of synthetic models of network structure whose intent is to explain how an observed network behavior arises. For instance, Price’s model demonstrates that the heavy-tailed degree distribution observed in citation networks can be thought of as a consequence of a preferential attachment process [1]. Generally speaking, a researcher defines a generative process (such as preferential attachment), generates multiple networks via that process, and then determines whether the generated networks “look like” real networks.

This work is concerned with reversing the process described above. Can a generative process be learned from data? Can a single model encode many different generative processes? That is, rather than accounting for a specific complex behavior through a predefined process, the model should be able to represent as broad a set of networks as possible with minimal input from a user.

The goal of this work is to find a flexible model that is able to learn a generative network process from data. We consider whether models such as deep belief networks, dependency networks, and members of the exponential random graph family can learn to generate networks whose complex behavior is consistent with a set of input examples. We find that deep belief networks are the only model to perform well at any network size, and we find that all models perform poorly when generating moderate to large networks. The good performance of the deep model for small network sizes suggests that it is learning latent representations that encode limited complex network behavior.

The remainder of the paper is structured as follows. Section 2 presents the task and a set of candidate models. Section 3 contains an evaluation method for the experiments described in Section 4. Results and conclusions are discussed in Sections 5 and 6, respectively.

2 Problem Formulation

Consider a fixed set of nodes V and an adjacency matrix \mathbf{Y} , where each element Y_{ij} takes on value 1 if there is an edge from node i to node j and 0 otherwise. Given a set of observed network examples, the task is to learn a joint distribution over the $|V|^2$ binary random variables that comprise \mathbf{Y} in such a way that complex network structure is maintained. For instance, if the input consists of scale-free networks that were generated from a preferential attachment process, the output networks should be scale-free as well. Note that this is a substantial departure from most work in network generation, which focuses on either understanding a predefined process that is not learned from data, or estimating potential generative processes from a single network instance rather than many.

A naive method would be to place no constraints on the dependencies between edges; that is, to model $P(\mathbf{Y} = \mathbf{y})$ as a fully-connected Markov random field (MRF). Learning and exact inference in this model are computationally intractable, however. We could mitigate this intractability by learning the dependence relationships from data; that is, learning the structure of the MRF. However, learning the structure of undirected discrete graphical models is itself intractable [2, Chapter 26.8].

An alternative approach is to learn *approximate* models; that is, models that provide tractable inference and learning of an approximation to the joint distribution. We present an application of two such models: dependency networks and deep belief networks.

2.1 Dependency Networks

A dependency network [3] is a form of probabilistic graphical model. A dependency network can be thought of as a collection of regressions or classifications among variables that can be combined via Gibbs sampling to define a joint distribution. This representation facilitates straightforward and efficient algorithms for structure learning and inference.

A dependency network is constructed over a set of variables D (in this case, the elements of the adjacency matrix) by independently fitting a sparse full-conditional distribution to each variable [2, Chapter 26.2.2]. Any available method for fitting these sparse conditional distributions can be selected; Heckerman et al. use probabilistic decision trees and support vector machines, for instance [3]. The graphical component of the dependency network is a bidirected unweighted network in which the elements of D are nodes, and there is an edge from $d_1 \in D$ to $d_2 \in D$ if d_2 's conditional model depends on d_1 .

We choose L1-penalized logistic regression for learning the sparse conditional models due to its appropriateness for binary data, conceptual and computational simplicity, and readily available implementations [4]. With this choice of model, the dependency network can be represented as a collection of conditional distributions

$$P(Y_{ij} = 1 | \mathbf{Y} \setminus \{Y_{ij}\}) = \frac{\exp(\beta_k \cdot \mathbf{Y} \setminus \{Y_{ij}\})}{1 + \exp(\beta_k \cdot \mathbf{Y} \setminus \{Y_{ij}\})} \quad (1)$$

2.2 Restricted Boltzmann Machines

A Boltzmann machine is a variety of Markov Random Field with a set of hidden (latent) nodes \mathbf{h} and a set of visible (observable) nodes \mathbf{x} [5], [2, Chapter 27.7]. In this work, the visible variables are the elements of the adjacency matrix. Exact inference in a Boltzmann machine is intractable and approximate inference can be slow [2, Chapter 27.7]. The speed of inference can be improved by only allowing edges between elements of \mathbf{h} and \mathbf{x} ; such a model is known as a restricted Boltzmann machine (RBM). We consider binary RBMs which constrain the hidden and visible variables to be binary-valued [6]. One can interpret the hidden variables in a binary RBM as ‘gates’ which govern first-order correlations between visible variables.

More formally, a binary restricted Boltzmann machine is defined as an undirected fully-connected bipartite graph with ‘visible’ nodes \mathbf{x} and ‘hidden’ nodes \mathbf{h} , and a $|\mathbf{x}| \times |\mathbf{h}|$ real matrix of parameters W . This model defines a joint distribution

$$P(\mathbf{x}, \mathbf{h}) = Z^{-1} \exp(\mathbf{x}^T W \mathbf{h}) \quad (2)$$

configuration	p1	markov	higher-order
# edges	X	X	X
# reciprocated edges	X	X	X
two in stars		X	X
two out stars		X	X
two mixed stars		X	X
transitive triads		X	X
geometric sum of in-degree			X
geometric sum of out-degree			X
two-paths			X
alternating k-triangles			X
alternating k-paths			X

Table 1: ERGM model specification.

where $Z = \sum_{\mathbf{x}', \mathbf{h}'} \exp(\mathbf{x}'^T W \mathbf{h}')$ is the normalization constant. The probability of the visible variables \mathbf{x} is found by marginalizing over the hidden variables \mathbf{h}

$$P(\mathbf{x}) = Z^{-1} \sum_{\mathbf{h}'} \exp(\mathbf{x}^T W \mathbf{h}') \quad (3)$$

There are many available methods for learning the parameters (W) of restricted Boltzmann machines. We select the stochastic maximum likelihood method described by Marlin et al. due to its ‘consistently better performance in terms of density estimation’ [7].

2.3 Deep Belief Networks

Deep belief networks (DBNs) can be thought of as stacked RBMs. We can think of the higher levels of the deep belief network as gates which mitigate higher-order correlations between visible variables. This allows deep belief networks to represent very complex dependence structures. However, exact learning and inference for such a model is highly intractable.

Fortunately, Hinton, Osindero, and Teh introduced a greedy approximate algorithm for tractably learning the parameters of this structure [8]. Each layer of the hierarchical structure is learned greedily; that is, as its own restricted Boltzmann machine. Data at higher levels are inferred using the parameters learned at lower levels. This approximate method has been very successful in practice, providing state-of-the-art performance for the classification of handwritten digits, for instance [8]. The algorithm is theoretically justified using the idea of complementary priors; for more details, please see ‘A fast learning algorithm for deep belief nets’ Appendix A [8].

2.4 Exponential Random Graph Models

An exponential random graph model (ERGM) models a network as a collection of edges whose probabilities are given as a normalized exponential function of configurations. Configurations can be constructed from any collection of dyads; some examples include a single edge, a reciprocated pair of edges, or three edges forming a cycle. Configurations must be specified by the experimenters. The form of the model is

$$P(\mathbf{Y} = \mathbf{y}) = \left(\frac{1}{\kappa}\right) \exp \left\{ \sum_A \eta_A g_A(\mathbf{y}) \right\} \quad (4)$$

where the summation is over all configurations, η_A is the parameter corresponding to configuration A , $g_A(\mathbf{y}) = \prod_{y_{ij} \in A} y_{ij}$ is the network statistic corresponding to A , and κ is the normalization constant. The parameter η_A can be interpreted as an index of the level of configuration A . For a more complete treatment, please see ‘An introduction to exponential random graph (p^*) models for social networks’ by Robins et al [9].

The ERGM framework accommodates a family of models of varying complexity. The simplest model assumes that all edges are independent. This is known as the Bernoulli graph model [9] or

Directed ER										
Network Size	deep		dependency		p1		markov		higher-order	
	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$
4	0.2810	0.0520	0.0000	0.0000	0.0000	0.3290	0.4690	0.0260	0.4430	0.0020
6	0.0140	0.0550	0.0000	0.0000	0.5880	0.0000	0.0020	0.0010	0.0040	0.5920
8	0.6390	0.9350	0.0020	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
10	0.5310	0.7860	0.0000	0.3330	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
20	0.0000	0.0000	0.0650	0.7640	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Configuration Model										
Network Size	deep		dependency		p1		markov		higher-order	
	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$
4	0.0630	0.0400	1.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0010	0.8510
6	0.0000	0.0000	0.0000	0.0000	0.0000	0.0020	0.4170	0.0350	0.4190	0.0410
8	0.0000	0.0130	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
10	0.0000	0.2250	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
20	0.0000	0.0000	0.0000	0.0010	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Krapivsky's Model										
Network Size	deep		dependency		p1		markov		higher-order	
	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$
4	0.5480	0.0000	0.0000	1.0000	0.0000	0.0000	0.1760	0.0000	0.1710	0.0000
6	0.1210	0.0120	0.0000	0.0060	0.0000	0.0300	0.0000	0.0000	0.0000	0.0000
8	0.7760	0.2760	0.3990	0.0670	0.7450	0.0000	0.0000	0.0300	0.0000	0.4590
10	0.3870	0.0150	0.0000	0.0000	0.0070	0.5290	0.1830	0.8160	0.1770	0.3800
20	0.0010	0.4290	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Twitter										
Network Size	deep		dependency		p1		markov		higher-order	
	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$	$\bar{\mu}_d$	$\bar{\mu}_c$
4	0.0000	0.0770	0.0000	0.0000	0.0000	0.0000	0.0170	0.2480	0.0190	0.1230
6	0.6030	0.6650	0.0590	0.0000	0.0000	0.0060	0.0000	0.5550	0.0000	0.0000
8	0.5950	0.3090	0.0000	0.0010	0.0000	0.0020	0.0000	0.0000	0.0000	0.0000
10	0.1220	0.2910	0.0000	0.0000	0.9160	0.0750	0.1350	0.3200	0.4690	0.2610
20	0.0000	0.0000	0.0000	0.0020	0.0000	0.0000	0.0000	0.0000	0.0010	0.0000

Table 2: The application of the permutation test describe in Section 3.1. Each column lists the p -value of the permutation test for a particular generative model and network statistic when applied to networks of different sizes. **Bold** values indicate that we can not reject the null hypothesis that the networks are indistinguishable at significance $\alpha = 0.05$; i.e. that the model succeeded in reproducing the ground-truth networks with respect to this statistic. $\bar{\mu}_d$ is the mean in-degree and $\bar{\mu}_c$ the mean clustering coefficient. Note that the deep model is most consistently able to reproduce the data.

the Erdos-Renyi (ER) model [10]. A slightly more complex formulation is found in the p_1 model, which assumes that $dyads$ are independent of each other. The p_2 model assumes a looser dyadic independence, where dyads are independent given node-specific attributes [11, 9].

Markov random graphs extend these dependence notions. In the model of Frank and Strauss, an edge Y_{ij} depends on any other edge involving i or j , conditioned on the rest of the graph [12]. This implies a set of possible configurations involving two or three nodes; please see Figure 1 of [9] for more details.

Finally, there is the higher-order model of Snijders and collaborators [13, 14], which includes a higher order of dependencies than those that can be represented by Markov random graphs. In contrast to other methods, the higher-order specification includes statistics derived from the entire network, such as a geometrically-weighted sum of the distribution of in-degree.

In this work, we consider the p_1 , Markov random, and new specification graph models. We omit p_2 because this work is not concerned with nodal attributes. Table 1 specifies the configurations used in each model. It should be noted that we did not include the ‘‘fixed effects’’ component of the p_1 model.

It should also be noted that members of the exponential random graph model are not typically used to generate networks. In fact, some ERGMs have been observed to exhibit ‘near degeneracy’. A graph distribution $P(\mathbf{Y} = \mathbf{y})$ is termed ‘near-degenerate’ if it puts most of its probability mass on only a few (possibly one or two) distinct graphs [15, 16, 17]. The implied graphs are usually highly unrealistic; for instance, fully-connected meshes or graphs with no edges at all. This property has been observed across members of the ERGM family, from the relatively uncomplicated Markov graph model of Frank and Strauss [14] to the higher order model of Snijders and collaborators [17, 13], although the higher-order models do a better job of avoiding it.

ERGMs are also typically fit to a single network rather than a set. We adapt ERGMs to the multiple input setting by assuming that the input networks are independent and identically distributed and maximizing their joint likelihood.

Network Size	$\bar{\mu}_d$	$\bar{\mu}_c$
4	0.5480	0.0000
6	0.6030	0.6650
8	0.5950	0.3090
10	0.1220	0.2910
20	0.0000	0.0000
100	0.0000	0.0000
200	0.0000	0.3160
400	0.0000	0.0000
600	0.0000	0.0000

Table 3: The permutation test applied to larger networks generated by the deep belief network. Note that the performance of the deep model degrades sharply as the networks increase in size.

3 Assessing the Quality of Generated Networks

The task of this work is to generate samples that ‘look like’ data. Ideally, we could formulate a testable hypothesis that one set of networks looks like another without appealing to the models that generated them. In the section below we describe a method for testing a minimum condition for effective network generation.

3.1 Permutation Test

For each network, we estimate a mean from a set of values corresponding to the nodes in the network. We then test whether the distribution of those estimated mean values from a set of generated networks is significantly different from those of the original networks from which the generator was learned. If this condition is satisfied, we still may not be learning good generative processes, but failing this condition implies that we are learning poor generators.

More formally, let Ω be the sample space defined by the set of all networks with $|V|$ nodes and let $f : \Omega \rightarrow \mathbb{R}^d$ be a function which maps a network to a real vector; i.e., f defines a vector of random variables. Let N be some observed network, let X_N be the result of the application of f to N , and μ_{X_N} be the mean of X . For instance, if f returns the degree distribution of a network, then μ_{X_N} is the mean degree of N .

Consider two sets of networks \mathcal{N}_1 and \mathcal{N}_2 ; let $\mu_{\mathcal{X}_i}$ be the mean of the mean values obtained by applying f to network set \mathcal{N}_i . For instance, if f returns the degree distribution of a network, then $\mu_{\mathcal{X}_1}$ is the mean of all of the mean degrees of \mathcal{N}_1 .

We wish to assess the hypothesis that sample means $\mu_{\mathcal{X}_1}$ and $\mu_{\mathcal{X}_2}$ are drawn from the same distribution. We assess this by computing an observed test statistic $t_{obs} = |\mu_{\mathcal{X}_1} - \mu_{\mathcal{X}_2}|$ and a set of n ‘permuted’ test statistics $t_{perm} = \left\{ |\mu_{\mathcal{X}_1^k} - \mu_{\mathcal{X}_2^k}| : k = 1, 2, \dots, n \right\}$ where $\mu_{\mathcal{X}_1^k}$ and $\mu_{\mathcal{X}_2^k}$ are the sample means obtained by repartitioning \mathcal{N}_1 and \mathcal{N}_2 according to some label permutation k .

Let p be the proportion of permuted means with a larger absolute difference than the true absolute difference. If $p \geq \alpha$, we conclude that we can not reject the null hypothesis that the means are drawn from the same distribution at some significance level α [18], implying that \mathcal{N}_1 and \mathcal{N}_2 are not distinguishable with respect to f .

This test tells us whether the generative model is able to learn the first moment of a specified complex network behavior (such as the degree distribution). As previously mentioned, if this condition is satisfied, we still may not be learning good generative processes (because the generated networks may not properly represent the higher moments of the complex behavior of the training networks), but failure does imply that we are learning poor generators.

4 Experiments

In order to empirically evaluate the ability of each model to reproduce sets of networks, we train each model on a set of ground-truth networks, generate a new set of networks from that model, and assess whether or not the model reproduced the ground truth using the techniques described

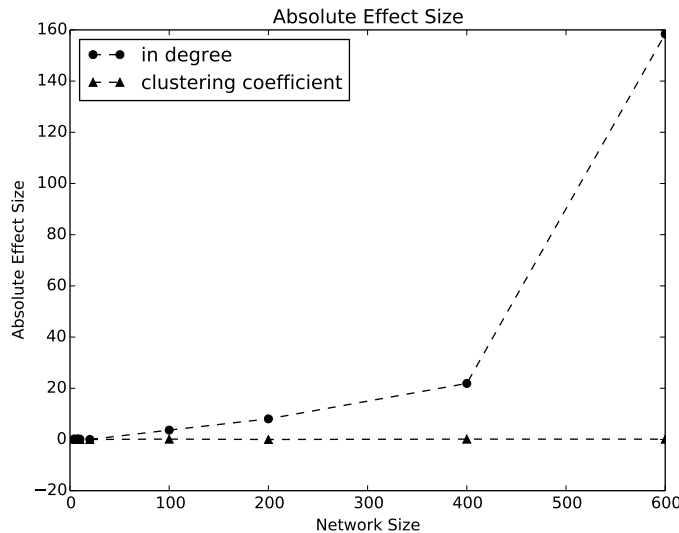


Figure 1: The effect size as a function of network size for larger networks generated by the deep belief network.

in Section 3. These methods all require that we specify a set of network statistics; we select the in-degree and clustering coefficient statistics, which are considered to be fundamental properties of networks [19]. Other network statistics will be investigated in future work.

4.1 Synthetic Data

We compare the relative ability of the candidate ERGMs and the new models to reproduce network data generated from existing models of network structure by including a variety of generators with different underlying dependence assumptions. The ground-truth models are the (directed) Erdos-Renyi (ER) model, the configuration model (which produces graphs with a given degree sequence), and the Krapivsky model (which produces graphs with dependent marginally heavy-tailed distributions over in- and out-degree via preferential attachment) [20].

4.2 Real Data

We also assess the ability of each model to reproduce a set of real networks. Specifically, we train each model on the Stanford Network Analysis Project (SNAP) Twitter ego network dataset [21]. This dataset contains 973 one-hop ego networks sampled from the online social network Twitter. Each network is defined by an ‘ego’ individual, the individuals that they follow, and the ‘follows’ edges amongst the nodes so defined. The networks contain between 6 and 248 nodes and 10 and 18,143 edges, respectively. This dataset is well-suited to evaluating the models at hand because it provides networks that are of the same scale which can reasonably be considered to be independent and identically distributed.

We limit the size of networks considered by sampling from each ego network. A set of ego networks with n nodes is constructed by, for all ego nets with at least n nodes, creating a network with the ego node and $n - 1$ randomly chosen nodes and the edges between them.

5 Results

We provide results for the directed ER, configuration, and Krapivsky models Table 2. In all cases, we assessed whether the models reproduced the data with respect to mean in-degree and clustering coefficient. 200 networks \mathcal{N}_1 were generated from each ground-truth model and 30 samples \mathcal{N}_2 were generated from each candidate model.

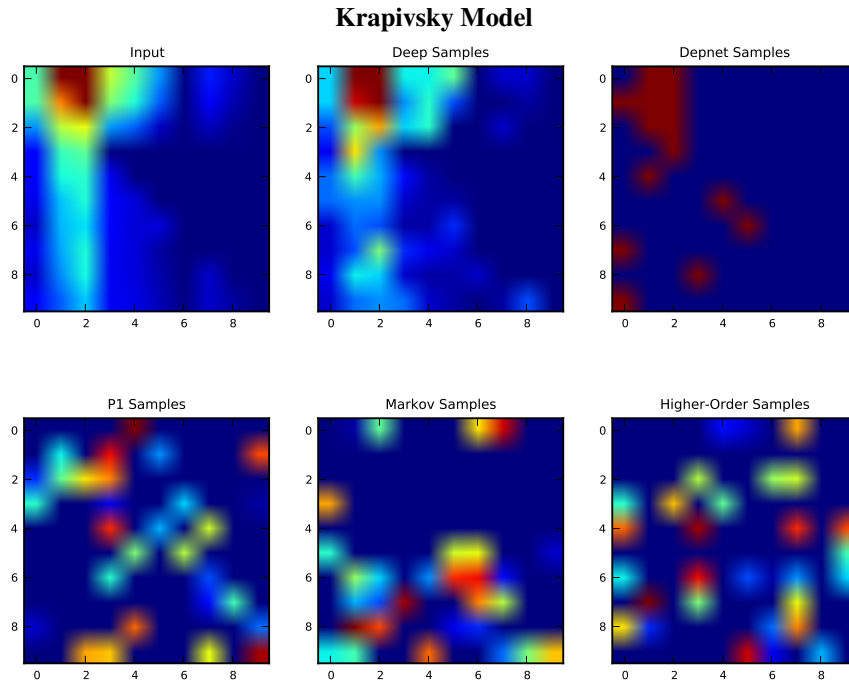


Figure 2: A visual assessment of network reproduction for Krapivsky networks with 10 nodes. Note that the deep model produces an average adjacency matrix which most closely resembles the input.

Table 2 shows the results of applying the permutation test described in Section 3.1 to the synthetic ground truth \mathcal{N}_1 and model-generated networks \mathcal{N}_2 of various sizes. We find that the deep model consistently provides the best performance; \mathcal{N}_1 and \mathcal{N}_2 are most often found to not be distinguishable with respect to the in-degree or clustering coefficient. The deep model outperforms the p1, Markov, and higher-order exponential random graph models as well as the dependency network model.

We were able to apply the deep learner to larger networks than the other models due to its shorter wall-clock runtime; results are shown in Table 3. Unfortunately, we found that the deep model does not maintain its performance for larger networks, and is almost completely unable to reproduce complex behavior for networks with more than a handful of nodes.

Note that the decreasing p-values seen in Table 3 could be a natural consequence of increasing statistical power rather than an indication that the deep model has failed to learn complex behavior. To assess this, we plotted the absolute effect size as a function of network size, as shown in Figure 1. The sharp increase in the effect size for the in-degree distribution suggests that the deep model is not learning to generate larger networks with correct degree distributions. On the other hand, the flat effect size for clustering coefficient suggests that the deep model is learning to generate larger networks with reasonable distributions of clustering coefficients. This may explain why the deep model is able to pass the clustering coefficient permutation test for networks with 200 nodes.

Figure 2 shows a visual assessment of each model’s reproduction of networks with 10 nodes generated from the Krapivsky models. The heatmaps display the average adjacency matrix of the ground-truth set and the adjacency matrices sampled from each model. Strikingly, the deep model’s reproduction is very similar in character to the input; conversely, the average adjacency matrices produced by the exponential random graph models look almost nothing like the input. This is at least partially due to an important difference between the model classes: the deep model encodes *particular edges*, whereas the ERGMs we consider, working only with a distribution over configurations, do not. The deep model will accordingly provide much better results when the task entails learning the structure of labeled edges.

6 Conclusion

We consider whether deep belief networks, dependency networks, and members of the exponential random graph family can learn to generate networks whose complex behavior is consistent with a set of input examples. We find that deep belief networks are the only model to perform well at any network size, but that all models perform poorly when generating moderate to large networks.

References

- [1] Derek de Solla Price. A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, 27(5):292–306, 1976.
- [2] Kevin P Murphy. *Machine learning: a probabilistic perspective*. Cambridge, MA, 2012.
- [3] David Heckerman, David Chickering, C Meek, Robert Rounthwaite, and Carl Kadie. Dependency Networks for Inference, Collaborative Filtering, and Data Visualization. *Journal of Machine Learning Research*, 1, 2000.
- [4] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [5] David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.
- [6] Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- [7] Ben Marlin, Kevin Swersky, Bo Chen, and Nando de Freitas. Inductive Principles for Restricted Boltzmann Machine Learning. *Journal of Machine Learning Research*, pages 1–24, April 2010.
- [8] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, pages 1–16, 2006.
- [9] Garry Robins, Pip Pattison, Yuval Kalish, and Dean Lusher. An introduction to exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):173–191, May 2007.
- [10] P Erdos and A R WI. On random graphs I. *Publ Math Debrecen*, 1959.
- [11] E Lazega and M Van Duijn. Position in formal structure, personal characteristics and choices of advisors in a law firm: A logistic regression model for dyadic network data. *Social Networks*, 1997.
- [12] Ove Frank and David Strauss. Markov Graphs. *Journal of the American Statistical Association*, pages 1–12, September 1986.
- [13] Tom AB Snijders, Philippa E Pattison, Garry L Robins, and Mark S Handcock. New specifications for exponential random graph models. *Sociological methodology*, 36(1):99–153, 2006.
- [14] Garry Robins, Tom Snijders, Peng Wang, Mark Handcock, and Philippa Pattison. Recent developments in exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):192–215, May 2007.
- [15] Mark S Handcock. Statistical models for social networks: Inference and degeneracy. *Dynamic social network modeling and analysis*, 126:229–252, 2003.
- [16] Mark S Handcock, Garry Robins, Tom AB Snijders, Jim Moody, and Julian Besag. Assessing degeneracy in statistical models of social networks. Technical report, Working paper, 2003.
- [17] Michael Schweinberger. Instability, Sensitivity, and Degeneracy of Discrete Exponential Families. *Journal of the American Statistical Association*, 106(496):1361–1370, December 2011.
- [18] Meyer Dwass. Modified randomization tests for nonparametric hypotheses. *The Annals of Mathematical Statistics*, 28(1):181–187, 03 1957.
- [19] M Newman. *Networks: An Introduction*. 2010.
- [20] P Krapivsky, G Rodgers, and S Redner. Degree Distributions of Growing Networks. *Physical Review Letters*, 86(23):5401–5404, June 2001.

[21] J J McAuley and J Leskovec. Learning to Discover Social Circles in Ego Networks. *NIPS*, 2012.