
Sequential Edge Clustering in Temporal Multigraphs

Elahe Ghalebi¹, Hamidreza Mahyar², Radu Grosu¹, Graham W Taylor^{3,4}, and Sinead A Williamson⁵

¹Technical University of Vienna, Austria

²Boston University, USA

³University of Guelph, Canada

⁴Vector Institute for Artificial Intelligence, Canada

⁵University of Texas, Austin, USA

Abstract

Interaction graphs, such as those recording emails between individuals or transactions between institutions, tend to be sparse yet structured, and often grow in an unbounded manner. Such behavior can be well-captured by structured, nonparametric edge-exchangeable graphs. However, such exchangeable models necessarily ignore temporal dynamics in the network. We propose a dynamic nonparametric model for interaction graphs that combines properties of edge-exchangeable models with dynamic clustering patterns that tend to reinforce recent behavioral patterns. We show that our method yields improved held-out likelihood over stationary variants, and impressive predictive performance against a range of state-of-the-art dynamic interaction graph models.

1 Introduction

Many forms of social interaction can be represented in terms of a multigraph, where each individual interaction corresponds to an edge in the graph, and repeated interactions may occur between two individuals. For example, we might have multigraphs where the value of an edge corresponds to the number of emails between two individuals, or the number of packets sent between two hosts. Traditional clustering algorithms such as K-means, hierarchical clustering, and stochastic block models have been successfully applied to cluster the nodes of a finite static graph, but only a few have considered edge clustering. Recently, the class of edge exchangeable graphs [1, 2, 3] have been proposed for modeling networks as exchangeable sequences of edges. These models are able to capture many properties of large-scale social networks, such as sparsity, community structure, and power-law degree distribution.

Being explicit models for sequences of edges, edge-exchangeable models are appropriate for networks that grow over time: we can add more edges by expanding the sequence, and their nonparametric nature means that we expect to introduce previously unseen vertices as the network expands. However, their exchangeable nature precludes graphs whose properties change over time. In practice, the dynamics of social interactions tend to vary over time. In particular, in models that aim to capture community dynamics, the popularity of a given community can wax and wane over time.

We propose a new model for multigraphs with clustered edges, that breaks the exchangeability of existing models by preferentially assigning edges to clusters that have been recently active. We show that incorporating dynamics using a mechanism based on the distance-dependent Chinese restaurant process (ddCRP) [4] leads to improved test-set predictive likelihood over exchangeable models. Further, when used in a link prediction task, we show improved performance over both its exchangeable counterpart and a range of state-of-the-art dynamic network models.

We begin in Section 1 and 2 by providing a general introduction to Bayesian models for multigraphs, in both the static and dynamic setting, plus a review of related works. We then introduce our model in Section 3, before empirically evaluating our approach in Section 4.

2 Related Work

Bayesian Models for Multigraphs. Most Bayesian models for graphs and multigraphs fall under the vertex-exchangeable framework, where the distribution over the adjacency matrix is invariant to jointly permuting the row and column indices. In this class, we have models such as the stochastic block model, where vertices are clustered into finitely many communities and a parameter is associated with each community-community pair [5, 6]; the infinite relational model, where the number of communities is unbounded [7]; mixed-membership stochastic block models, where edges are generated according to an admixture model [8]; the latent feature relational model, where parameter are distributed according to a latent feature model based on the Indian buffet process [9]; and Poisson factor analysis, where the parameter are distributed according to a gamma process-based latent factor model [10, 11]. While these models are able to capture interesting community structure, the resulting graphs are dense or empty almost surely [12, 13]. This makes them a poor choice for large real-world networks, which are typically sparse. Further, since they explicitly model zero-edges, they are not well-suited to link prediction tasks.

An alternative way of constructing multigraphs is using an exchangeable sequence of edges [1, 2, 3]. Here, we assume the edges are generated by sequentially sampling pairs of vertices. These pairs are *iid* given some nonparametric prior, such as a Dirichlet process, a normalized generalized gamma process, or a Pitman Yor process. This construction means multigraphs can grow over time by adding new edges, making them well-suited for link prediction tasks. Appropriate choices of priors can yield sparsity and power law degree distribution.

Models for Dynamic Graphs. There has been significant research attention on dynamic (time-evolving) network modelling, ranging from non-Bayesian methods such as dynamic extensions of the exponential random graph model (ERGM) [14], or matrix and tensor factorization-based methods [15], to Bayesian latent variable models [16, 17, 18, 19, 20, 21, 22, 23]. A common approach relies on the extensions of *static* network models to a *dynamic* framework [24]. We focus here on dynamic extensions of Bayesian models of the forms discussed above.

Most dynamic Bayesian networks extend jointly vertex-exchangeable graphs. For example, [25] extends the stochastic blockmodel using an extended Kalman filter (EFK) based algorithm, and the stochastic block transition model [26] relaxes a hidden Markov assumption on the edge-level dynamics, allowing the presence or absence of edges to directly influence future edge probabilities. Several methods have also been used to incorporate temporal dynamics into the mixed membership stochastic blockmodel framework [27, 28, 29] and the latent feature relational model [30, 31, 32]. Lately, several models have extended Poisson factor analysis. The dynamic gamma process Poisson factorization (DGPPF) [33] introduces dependency by incorporating a Markov chain of marginally gamma random variables into the latent representation. The dynamic Poisson gamma model (DPGM) [34] extends a bilinear form of Poisson factor analysis [35] in a similar manner. The dynamic relational gamma process model (DRGPM) [23] also incorporates a temporally dependent thinning process.

3 Sequential Edge Clustering in Temporal Multigraphs

Our model extends the Mixture of Dirichlet Network Distributions (MDND) [3], an edge-exchangeable model that clusters edges to create community-like structure. The MDND is based on a sequence of Dirichlet processes: one controls the global popularity of vertices; one controls a distribution over cluster indicators; and two per cluster control the cluster-specific distributions over “sender” and “recipient” vertices. Any of these distributions could be replaced with dynamic or dependent clustering models to generate a temporally evolving graph. In practice, replacing all of the distributions with dynamic alternatives is likely to lead to overspecification of the dependencies, making inference challenging.

We choose to retain stationary models for H , A_k and B_k , implying that a cluster’s representation stays stable over time, and allow the cluster popularities to vary by making the sequence z_1, z_2, \dots time-varying. A number of methods exist for incorporating temporal dynamics into the Dirichlet process or the related Chinese restaurant process (CRP), *e.g.* [36, 37, 38]. For our purposes, we choose to use the distance-dependent CRP (ddCRP) [4], as it captures the property that we are likely to see clusters that have appeared recently. Recall that the CRP can be described in terms of a restaurant analogy, where customers select tables (clusters) proportional to the number of people already seated at that table, or sit at a new table with probability proportional to a concentration parameter α . The ddCRP modifies this by encouraging customers to sit next to “similar” customers.

In a time-dependent setting, similarity is evaluated based on arrival time using some non-negative, non-increasing decay function f such that $f(\infty) = 0$. Concretely, given t_i, t_j arrival times of i and j ,

$$d_{i,j} = \begin{cases} t_i - t_j & t_i \geq t_j \\ \infty & \text{otherwise} \end{cases}$$

Then the i -th customer picks a customer c_i to sit next to (and therefore a cluster) according to

$$P(c_i = j) \propto \begin{cases} f(d_{i,j}) & i \neq j \\ \alpha & i = j \end{cases} \quad (1)$$

In an interaction network context, this implies that we are likely to see modes of communication that have been popular in recent time periods, over modes of communication that have fallen out of popularity. Another reason to favor the ddCRP is ease of inference: its construction lends itself to an easy-to-implement Gibbs sampler, allowing us to apply our method to larger graphs. By contrast, many other dependent Dirichlet processes have much more complicated inference algorithms, which would limit scalability. While the ddCRP does not yield marginal invariance (i.e. it is not invariant to adding edges at previously observed time points), this is not a concern in our setting, since we are typically able to observe past instances of the full graph, and are interested in predicting future edges. Incorporating the ddCRP into the MDND yields the following generative process:

$$\begin{aligned} H &:= \sum_{i=1}^{\infty} h_i \delta_{\theta_i} \sim \text{DP}(\gamma, \Theta) & P(c_i = j) &\propto \begin{cases} f(d_{i,j}) & i \neq j \\ \alpha & i = j \end{cases} \\ A_k &:= \sum_{i=1}^{\infty} a_i \delta_{\theta_i} \sim \text{DP}(\tau, H) & z_i &= \begin{cases} z_{c_i} & c_i \neq i \\ i & c_i = i \end{cases} \\ B_k &:= \sum_{i=1}^{\infty} b_i \delta_{\theta_i} \sim \text{DP}(\tau, H) & s_n &\sim A_{z_n} \\ & & r_n &\sim B_{z_n} \end{aligned} \quad (2)$$

We perform inference using an MCMC scheme based on that developed in [4].

4 Experiments

Datasets. We evaluate our model on four real-world networks: (1) Face-to-Face Dynamic Contacts (FFDC)¹ [39]; $M = 180$ students with $N = 8,332$ communications for $t = 7$ school days. (2) Social Evolution (SocialEv)² [40]; $M = 68$ nodes and $N = 399$ links. This network has a high clustering coefficient and about 2×10^6 events over $t = 6$ time slots. (3) DBLP [41]; $M = 324$ most connected authors over all time periods which contains $N = 11,154$ edges. (4) Enron³; $N = 8,534$ interactions among $M = 151$ users over 38 months (May 1999- June 2002).

Metrics. We use F1 score, Map@ k and Hit@ k . F1 score is $2 \times (\text{precision} \times \text{recall}) / (\text{precision} + \text{recall})$. Precision is the fraction of edges in the future network present in the true network, Recall is the fraction of edges of the true network present in the future network. MAP@ k is the classical mean average precision measure and Hits@ k is the rate of the top- k ranked edges.

Comparison. In addition to the exchangeable MDND, we consider three state-of-the-art network models, discussed in Section 2: DRGPM [23], DPGM [34], and DGPPF [33]. None of these models are designed for explicit link prediction, yet can be modified to give predictions using the above procedure of selecting the N highest probability edges. These models also have the limitation of assuming a fixed number of vertices.

Results. Figure 1 illustrates the F1 score, Map@ k and Hits@ k for the proposed model, DYNMDND, with all three decay types, EXPONENTIAL, LOGISTIC and CRP compared to DRGPM, DPGM and DGPPF for dynamic link prediction. We use the networks of time slots 1 to T as a training set and predict the network edges of time slot $T + 1$. We report the results on the three datasets: FFDC, DBLP and Enron, using time interval of one day, one year and one month, respectively. For each task,

¹<http://www.sociopatterns.org/datasets/high-school-contact-and-friendship-networks/>

²<http://realitycommons.media.mit.edu/socialrevolution.html>

³<https://www.cs.cmu.edu/~enron/>

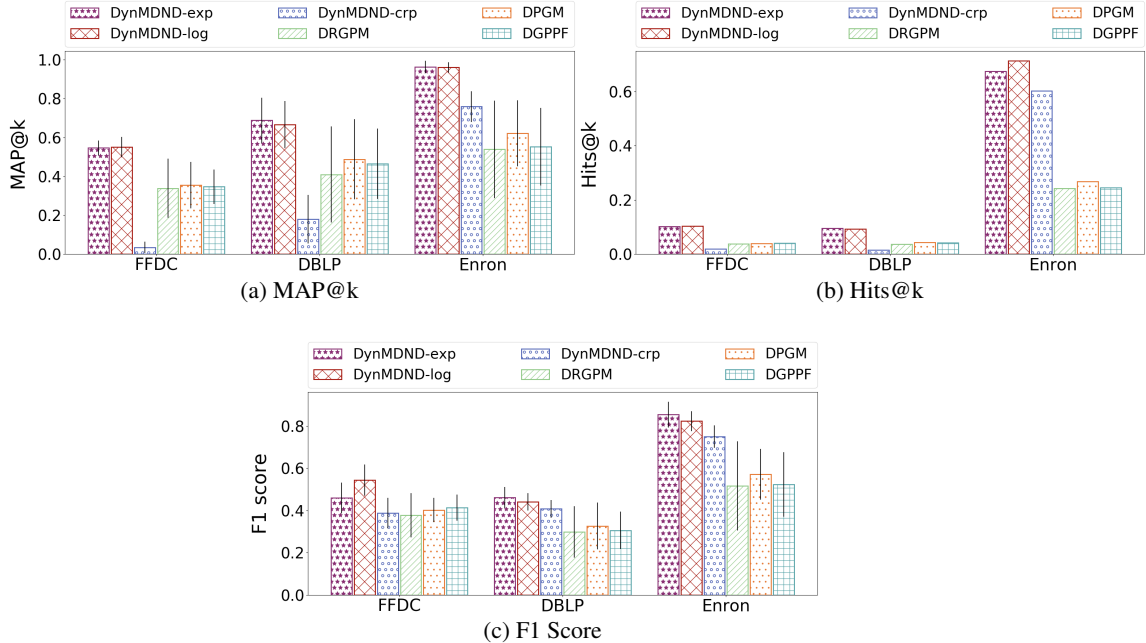


Figure 1: Performance of DYNMDND for dynamic link prediction compared to DRPGM, DPGM and DGPPF on the three datasets.

Table 1: Predictive Log Likelihood for the held-out (test set) data computed by DYNMDND with three different decay functions on four real-world networks (mean \pm standard error).

Dataset	MDND	DYNMDND-LOGISTIC	DYNMDND-EXPONENTIAL
FFDC	-383,094.6 \pm 146	-286,833.9\pm220	-344,683.9 \pm 105
Enron	-1032.94 \pm 147.18	-640.73\pm90.51	-700.88 \pm 22.97
DBLP	-980,928.4 \pm 532.1	-798,568.4 \pm 998.5	-649,521\pm195.0
SocialEv	-173,708.1 \pm 223.6	-23,087.3 \pm 91.0	-19,820.1\pm93.6

we repeat the experiments 10 times and report the mean and standard deviation of each evaluation metric.

We see that DYNMDND significantly outperforms DRPGM, DPGM and DGPPF on all metrics, for the task of dynamic link prediction. We hypothesise that this is due to several reasons. First, DYNMDND is explicitly designed in terms of a predictive distribution over edges, making it well-suited to predicting future edges. Second, DYNMDND is able to increase the number of vertices over time, and is likely better able to capture natural network growth. Conversely, the other methods assume the number of vertices is fixed.

Table 1 shows the predictive log likelihood computed by DYNMDND using two different decays (*i.e.* EXPONENTIAL and LOGISTIC) in comparison with the CRP decay function. The other comparison methods do not incorporate such a log likelihood. At each time slot T , we use 80% of the network data for training the model and the remaining 20% for the test set. It can be seen that considering time dependency into our mixture models results in a better log likelihood for the task of prediction. We use a sample size of 1000. We repeat the experiments 10 times and report the mean and standard deviation of the results for the four real networks.

5 Conclusion

We have presented a new model for interaction networks that can be represented in terms of sequences of links, such as email interaction graphs and collaboration graphs. Using a nonparametric sequence of links makes our model well-suited to predicting future links, and unlike many vertex-based graph models allows for an unbounded number of vertices.

References

- [1] Diana Cai, Trevor Campbell, and Tamara Broderick. Edge-exchangeable graphs and sparsity. In *Advances in Neural Information Processing Systems*, pages 4249–4257, 2016.
- [2] Harry Crane and Walter Dempsey. Edge exchangeable models for interaction networks. *Journal of the American Statistical Association*, 113(523):1311–1326, 2018.
- [3] Sinead A Williamson. Nonparametric network models for link prediction. *The Journal of Machine Learning Research*, 17(1):7102–7121, 2016.
- [4] David M Blei and Peter I Frazier. Distance dependent chinese restaurant processes. *Journal of Machine Learning Research*, 12(Aug):2461–2488, 2011.
- [5] T.A.B. Snijders and T. Nowicki. Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of Classification*, 14(1):75–100, 1997.
- [6] B. Karrer and M.E.J. Newman. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83(1):016107, 2011.
- [7] C. Kemp, J.B. Tenenbaum, T.L. Griffiths, T. Yamada, and N. Ueda. Learning systems of concepts with an infinite relational model. In *National Conference on Artificial Intelligence (AAAI)*, pages 381–388, 2006.
- [8] Edoardo M Airoldi, David M Blei, Stephen E Fienberg, and Eric P Xing. Mixed membership stochastic blockmodels. *Journal of machine learning research*, 9(Sep):1981–2014, 2008.
- [9] Kurt Miller, Michael I Jordan, and Thomas L Griffiths. Nonparametric latent feature models for link prediction. In *Advances in neural information processing systems*, pages 1276–1284, 2009.
- [10] Mingyuan Zhou and Lawrence Carin. Negative binomial process count and mixture modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):307–320, 2013.
- [11] Prem Gopalan, Jake M Hofman, and David M Blei. Scalable recommendation with hierarchical poisson factorization. In *UAI*, pages 326–335, 2015.
- [12] David J Aldous. Representations for partially exchangeable arrays of random variables. *Journal of Multivariate Analysis*, 11(4):581–598, 1981.
- [13] D.N. Hoover. Relations on probability spaces and arrays of random variables. *Preprint. Institute for Advanced Study, Princeton.*, 1979.
- [14] Fan Guo, Steve Hanneke, Wenjie Fu, and Eric P Xing. Recovering temporally rewiring networks: A model-based approach. In *Proceedings of the 24th international conference on Machine learning*, pages 321–328. ACM, 2007.
- [15] Daniel M Dunlavy, Tamara G Kolda, and Evrim Acar. Temporal link prediction using matrix and tensor factorizations. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(2):10, 2011.
- [16] Purnamrita Sarkar, Sajid M. Siddiqi, and Geogrey J. Gordon. A latent space approach to dynamic embedding of co-occurrence data. In *International Conference on Artificial Intelligence and Statistics*, volume 2, pages 420–427, 21–24 Mar 2007.
- [17] Katsuhiko Ishiguro, Tomoharu Iwata, Naonori Ueda, and Joshua B Tenenbaum. Dynamic infinite relational model for time-varying relational data analysis. In *Advances in Neural Information Processing Systems*, pages 919–927, 2010.
- [18] Purnamrita Sarkar, Deepayan Chakrabarti, Michael Jordan, et al. Nonparametric link prediction in large scale dynamic networks. *Electronic Journal of Statistics*, 8(2):2022–2065, 2014.
- [19] Daniele Durante and David B Dunson. Nonparametric bayes dynamic modelling of relational data. *Biometrika*, 101(4):883–898, 2014.
- [20] Aaron Schein, Mingyuan Zhou, David M Blei, and Hanna Wallach. Bayesian poisson tucker decomposition for learning the structure of international relations. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, pages 2810–2819, 2016.
- [21] Konstantina Palla, Francois Caron, and Yee Whye Teh. Bayesian nonparametrics for sparse dynamic networks. *arXiv preprint arXiv:1607.01624*, 2016.

- [22] Yin Cheng Ng and Ricardo Silva. A dynamic edge exchangeable model for sparse temporal networks. *arXiv:1710.04008*, 2017.
- [23] Sikun Yang and Heinz Koepl. Dependent relational gamma process models for longitudinal networks. In *International Conference on Machine Learning*, pages 5547–5556, 2018.
- [24] Elahe Ghalebi, Baharan Mirzasoleiman, Radu Grosu, and Jure Leskovec. Dynamic network model from partial observations. In *Advances in Neural Information Processing Systems*, pages 9862–9872, 2018.
- [25] Kevin S Xu and Alfred O Hero. Dynamic stochastic blockmodels for time-evolving social networks. *IEEE Journal of Selected Topics in Signal Processing*, 8(4):552–562, 2014.
- [26] Kevin Xu. Stochastic block transition models for dynamic networks. In *International Conference on Artificial Intelligence and Statistics*, pages 1079–1087, 2015.
- [27] Wenjie Fu, Le Song, and Eric P Xing. Dynamic mixed membership blockmodel for evolving networks. In *Proceedings of the 26th annual international conference on machine learning*, pages 329–336, 2009.
- [28] Eric P Xing, Wenjie Fu, Le Song, et al. A state-space mixed membership blockmodel for dynamic network tomography. *The Annals of Applied Statistics*, 4(2):535–566, 2010.
- [29] Qirong Ho, Le Song, and Eric Xing. Evolving cluster mixed-membership blockmodel for time-evolving networks. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 342–350, 2011.
- [30] James Foulds, Christopher DuBois, Arthur Asuncion, Carter Butts, and Padhraic Smyth. A dynamic relational infinite feature model for longitudinal social networks. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 287–295, 2011.
- [31] Creighton Heaululani and Zoubin Ghahramani. Dynamic probabilistic models for latent feature propagation in social networks. In *International Conference on Machine Learning*, pages 275–283, 2013.
- [32] Myunghwan Kim and Jure Leskovec. Nonparametric multi-group membership model for dynamic networks. In *Advances in neural information processing systems*, pages 1385–1393, 2013.
- [33] Ayan Acharya, Joydeep Ghosh, and Mingyuan Zhou. Nonparametric bayesian factor analysis for dynamic count matrices. *arXiv preprint arXiv:1512.08996*, 2015.
- [34] Sikun Yang and Heinz Koepl. A poisson gamma probabilistic model for latent node-group memberships in dynamic networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [35] Mingyuan Zhou. Infinite edge partition models for overlapping community detection and link prediction. In *Artificial Intelligence and Statistics*, pages 1135–1143, 2015.
- [36] Steven N MacEachern. Dependent dirichlet processes. *Unpublished manuscript, Department of Statistics, The Ohio State University*, pages 1–40, 2000.
- [37] Dahua Lin, Eric Grimson, and John W Fisher. Construction of dependent dirichlet processes based on poisson processes. In *Advances in neural information processing systems*, pages 1396–1404, 2010.
- [38] Lu Ren, David B. Dunson, and Lawrence Carin. The dynamic hierarchical dirichlet process. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 824–831, New York, NY, USA, 2008. ACM.
- [39] Rossana Mastrandrea, Julie Fournet, and Alain Barrat. Contact patterns in a high school: a comparison between data collected using wearable sensors, contact diaries and friendship surveys. *PloS one*, 10(9):e0136497, 2015.
- [40] Anmol Madan, Manuel Cebrian, Sai Moturu, Katayoun Farrahi, et al. Sensing the " health state" of a community. *IEEE Pervasive Computing*, 11(4):36–45, 2011.
- [41] Sitaram Asur, Srinivasan Parthasarathy, and Duygu Ucar. An event-based framework for characterizing the evolutionary behavior of interaction graphs. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 3(4):16, 2009.