

Fast Change Point Detection on Dynamic Social Networks

Yu Wang* Aniket Chakrabarti* David Sivakoff# Srinivasan Parthasarathy*

* Department of Computer Science and Engineering

Department of Statistics

The Ohio State University, Columbus, Ohio, USA

wang.5205@osu.edu, srini@cse.ohio-state.edu

Abstract

A number of real world problems in many domains (e.g. sociology, biology, political science and communication networks) can be modeled as dynamic networks with nodes representing entities of interest and edges representing interactions among the entities at different points in time. A common representation for such models is the snapshot model - where a network is defined at logical time-stamps. An important problem under this model is change point detection. In this work we devise an effective and efficient three-step-approach for detecting change points in dynamic networks under the snapshot model. Our algorithm achieves up to 9X speedup over the state-of-the-art while improving quality on both synthetic and real world networks.

1 Introduction

Dynamic network analysis is increasingly used in complex application domains ranging from social networks (Facebook network evolution [Leskovec and Sosič, 2016]) to biological networks (protein-protein interaction [Shih and Parthasarathy, 2012]), from political science (United Nations General Assembly voting network [Voeten, 2012]) to communication networks ([Klimt and Yang, 2004; Asur *et al.*, 2007]). Such dynamic networks are often represented using the snapshot model. Under this model, every network snapshot (represented by a graph) is defined at a logical timestamp. Two questions of fundamental importance are – (i) how does a network evolve? (ii) when does a network change significantly so as to raise suspicion that something fundamentally different is happening?

Various generative models [Peixoto and Rosvall, 2015; Zhang *et al.*, 2016] have been proposed to address question (i) - to explain the evolution of a network. They study network evolution under certain generative models [Erdős and Rényi, 1960; Karrer and Newman, 2011]. In reality, the generative model itself might change, as addressed in question (ii) above. Existing work [Akoglu *et al.*, 2014; Ranshous *et al.*, 2015] use complex methods to detect such changes. A drawback of these methods is that they are time-consuming, and hence not scalable (in terms of both network size and number of snapshots). We seek to find an efficient

and effective solution which can scale up both with network size and with the number of snapshots.

In this paper, we present a simple and efficient algorithm based on likelihood maximization to detect change points in dynamic networks under the snapshot model. We demonstrate the utility of our algorithm on both synthetic and real world networks drawn from political science (congressional voting, UN voting), and show that it outperforms two recent approaches (DeltaCon[Koutra *et al.*, 2016], and LetoChange[Peel and Clauset, 2015]) along the axes of both quality and efficiency. Our work has the following contributions:

1. Our approach is general purpose – it can accommodate various snapshot generative models (see Table 1).
2. We model network evolution as a first order Markov process and consequently our algorithm accounts for the temporal dependency while computing the dissimilarity between snapshots.
3. Our algorithm is effective and efficient on both synthetic and real world networks and has constant memory overhead that can be tuned by a user controlled parameter.

2 Related Work

Ranshous *et al.* [Ranshous *et al.*, 2015], and Akoglu *et al.* [Akoglu *et al.*, 2014] recently survey network anomaly detection methods which share related goals. Closest to our work are those ideas that fall under the theme of “Event and Change Detection”, defined as: given a network sequence, a dissimilarity scoring function, and a threshold, a change is flagged if the dissimilarity of two consecutive snapshots is above the threshold[Ranshous *et al.*, 2015]. A key differentiator is that the work surveyed there is somewhat ad-hoc while we take a more principled approach and assume there is a latent generation model governing the network dynamics and seek to detect the change in the latent space. Moreover, we consider the temporal dependency across the snapshots while none of the works surveyed explicitly account for temporal dependency or the bias introduced therein. We also note that our work is orthogonal to methods that seek to identify the temporal scale resolution for networks[Caceres and Berger-Wolf, 2013].

DeltaCon [Koutra *et al.*, 2016] uses a graph similarity-based [Berlingerio *et al.*, 2012] approach to detect change

points in dynamic networks. It derives the features of a snapshot based on sociological theories. And the feature similarity of each consecutive snapshot pair is calculated. That work is model agnostic (has no assumption on the generation model of networks), and is the state-of-the-art in terms of efficiency. We compare our algorithm against this approach.

Moreno and Neville [Moreno and Neville, 2013], Bridges et al. [Bridges et al., 2015] and Peel and Clauset [Peel and Clauset, 2015] develop network hypothesis testing based approaches. The advantage is that one can get a p-value of the test, which quantifies the confidence of the conclusion. However, these approaches have two shortcomings: firstly, they need to assume a specific generation model of the networks (mKPGM, GBTER and GHRG respectively); secondly, they are extremely slow, mostly due to the bootstrapping for p-value calculation. La Fond et al.’s work [La Fond et al., 2014] can also generate a p-value. It is tested against DeltaCon without reporting running time and efficiency concern is also mentioned in the paper. These algorithms will not work in our setting where the detection is done real time under bounded memory constraints. We compare our model agnostic algorithm against [Peel and Clauset, 2015].

Finally, Li et al. [Li et al., 2016] propose an online algorithm, and consider temporal dependency. The problem they study is different from ours in that they study information diffusion on network with fixed structure and use continuous time. A recent work by Zhang et al. [Zhang et al., 2016] also studies the dynamic network in a Markov chain setting. They focus on community detection while we focus on change point detection.

3 Problem Formulation

This paper studies how to detect the times at which the fundamental evolution mechanism of a dynamic network changes. We assume that there is some *unknown* underlying model that governs the generative process. Our change point detection algorithm is agnostic to this model. We assume that the observed network snapshots are samples that depend on some generative model and the previous snapshot. Networks have fluctuation across snapshots even when the generative model stays unchanged. Only when the generative model changes do we consider it a fundamental change. We represent the evolutionary process as a Markov Network (Figure 1).

In Figure 1, M_t is the network generation model at time t . It is a triad $M_t = \langle \text{Type}_t, \Theta_t, \alpha_t \rangle$, where α_t is the continuity parameter at time t , Type_t specifies the model, while Θ_t represents the model parameters (Table 1 consists of some generative models we experiment on). G_t is the network (graph) observable at time t . We assume the number of vertices in G_t is fixed to be N for all snapshots (the union of all nodes is used when there is node addition/deletion, as in [Peel and Clauset, 2015]), so each G_t has $2^{\binom{N}{2}}$ possible configurations, and T is the total number of snapshots we observe. As per Figure 1 the configuration of the network at time t , G_t , depends on the generation model at time t , M_t (unobserved) and the network configuration at time $t - 1$, G_{t-1} (observed). Hence the networks in the observed sequence are samples from a conditional distribution (samples are not independent). The continuity rate parameter α_t controls the




Figure 1: Representation of the underlying generative process. Our inference is agnostic to M_t s.

fraction of edges and non-edges that are retained from the previous snapshot, G_{t-1} . The network at time t is assumed to be generated in the following way: for each dyad, keep the connection status from time $t - 1$ with probability $1 - \alpha_t$, and with probability α_t , resample the connection according to the generation model at time t . Consequently, the smaller α_t is, the more overlap between two consecutive snapshots there is. Note that two consecutive network configurations may differ substantially if $\alpha_t > 0$, even though the underlying generation model may be the same. Moreover, the changes of the generation model are assumed to be rare across the time span ($M_t \neq M_{t+1}$ is a rare event).

Problem Definition Our goal is to efficiently find a set $S \subset \{2, \dots, T\}$ such that $t \in S \iff M_t \neq M_{t-1}$, that is, to efficiently find all the time points at which the network generation model is different from the previous time point.

4 Methodology

Given the graphical formulation of the problem, exact inference is impossible since we do not know the underlying generative model, and our observations are stochastic. However, even without prior knowledge of the generative model, we can still design an approximate inference technique based on MCMC sampling theory.

The framework is straightforward, as mentioned in Section 2: we first extract a “feature vector” from each snapshot, then quantify the dissimilarity between consecutive snapshots, and flag out a change point when the dissimilarity score is above a threshold. We use the joint edge probability as the “feature vector” (Section 4.1), exploit Kolmogorov-Smirnov statistic, Kullback-Leibler divergence and Euclidean distance for dissimilarity measure (Section 4.2), and use a permutation test like approach to determine the threshold (Section 4.3).

4.1 Edge Probability Estimation

In this subsection, we discuss how to (approximately) estimate the joint distribution of the dyads¹. We track the presence or absence of a small fixed number of dyads throughout the entire observed sequence of network snapshots. We break down the observation sequence into fixed-length windows, and for each window we infer the joint distribution of the dyads in our sample. We model each dyad to be a conditionally independent *two-state Markov chain* (Figure 2, we use α instead of α_t in this section for brevity) given the sequence of generative models $(M_t)_{t \geq 0}$ (this conditional independence assumption is satisfied for the choices of models in Table 1). Note that even for generative processes that may

¹we refer node pairs, which may or not be linked, as dyads

Table 1: Edge probability between a dyad in each model

Model	Edge probability	Explanation
Erdős-Rényi (ER)	$p(\langle n_i, n_j \rangle M) = p$	p : edge probability
Chung-Lu (CL)	$p(\langle n_i, n_j \rangle M) = \beta \bar{w}_i \bar{w}_j / \sum_i \bar{w}_i$	\bar{w}_i : weight of node i ([Pfeiffer III et al., 2012]); β : edge density
Stochastic Block Model (SBM)	$p(\langle n_i, n_j \rangle M) = p(c_i, c_j)$	c_i : community assignment of node n_i $p(r, s)$: probability of edges between communities r and s
SBM-CL	$p(\langle n_i, n_j \rangle M) \propto p(c_i, c_j) \bar{w}_i \bar{w}_j$	notation as above
BTER	$p(\langle n_i, n_j \rangle M) = p_{ER} I[c_i = c_j] + p_{CL} I[c_i \neq c_j]$	Intra-community edge probability follows ER, inter-community CL; [Seshadhri et al., 2012] $I[\cdot]$ is the indicator function




Figure 2: Two state Markov Chain

result in greater dependence among dyads (such as the configuration model), in many cases such dependence will be local, and if the number of dyads sampled is small, then these dyads will be spread out enough to be considered independent. Moreover, the conditional independence assumption significantly improves computational efficiency ([Hunter et al., 2012]). The marginal probabilities of these dyads can then be estimated using the observed samples within each time window.

We formalize the estimation procedure below. Given a network sequence G_1, \dots, G_T , we group the networks into sliding windows. We define W_t to be a subsequence of s consecutive observed networks ending at network G_t , so $W_t \equiv (G_{t-s+1}, G_{t-s+2}, \dots, G_t)$. We use equal sized sliding windows with a step size η , and we obtain a window sequence $W_s, W_{s+\eta}, \dots, W_{s+(i-1)\eta}, \dots, W_T$. Non-overlapping window setting uses $\eta = s$. In each window i , we can estimate the joint edge distribution (for the selected dyads) $P(e_1, e_2, \dots, e_k | M_{s+(i-1)\eta})$, where $e_j = 1$ indicates an undirected edge between the j -th dyad, and $k \ll \binom{N}{2}$ is the number of dyads tracked. For each of the models in Table 1, the joint distribution can be factorized into $P(e_1, e_2, \dots, e_k | M_{s+(i-1)\eta}) = \prod_j P(e_j | M_{s+(i-1)\eta})$. (conditional independence, see method description above)

We can view a dyad across time as a two state Markov chain, and the chain length is the window size. We call a dyad across time a chain in the following text for brevity. Let $p_j \equiv P(e_j = 1 | M_{s+(i-1)\eta})$, $q_j \equiv 1 - p_j$, and suppose we are interested in k chains.

1) Maximum Likelihood Estimator (MLE)

The joint probability of the chains is (Figure 2)

$$P(N_{**}^{(j=1, \dots, k)} | \alpha, \vec{p}) = c_k \prod_{j=1}^k (\alpha p_j)^{N_{01}^{(j)}} (\alpha q_j)^{N_{10}^{(j)}} (1 - \alpha p_j)^{N_{00}^{(j)}} (1 - \alpha q_j)^{N_{11}^{(j)}} \quad (1)$$

where $N_{01}^{(j)}$ is the number of transitions from 0 to 1 for a chain (non-edge to edge for the dyad j within the window), c_k stands for the combinatorial coefficients independent of α, \vec{p} . and $\sum_{p, q \in \{0, 1\}} N_{pq}^{(j)} = s - 1$ for all j s. And hence the log-likelihood (omitting the coefficient c_k) is:

$$L(\alpha, \vec{p} | N_{**}^{(j)}) = \sum_{j=1}^k [N_{01}^{(j)} \ln(\alpha p_j) + N_{10}^{(j)} \ln(\alpha q_j) + N_{00}^{(j)} \ln(1 - \alpha p_j) + N_{11}^{(j)} \ln(1 - \alpha q_j)] \quad (2)$$

MLE for a single chain First consider there is only one chain. Solving the zero-derivative Equations (4) and (5), leads to estimators of α, p . And the estimators indeed lead to a negative definite Hessian, and therefore is the MLE. Hence we have

$$\hat{\alpha}_{MLE} = \frac{N_{01}N_{1*} + N_{10}N_{0*}}{N_{0*}N_{1*}} \quad \hat{p}_{MLE} = \frac{N_{01}N_{1*}}{N_{01}N_{1*} + N_{10}N_{0*}} \quad (3)$$

MLE for multiple chains The MLE for multiple chains essentially involves solving a high degree polynomial, which in general does not have a closed form solution.

$$\frac{\partial L}{\partial p_j} = \frac{N_{01}^{(j)}}{p_j} - \frac{N_{10}^{(j)}}{q_j} - \frac{\alpha N_{00}^{(j)}}{1 - \alpha p_j} + \frac{\alpha N_{11}^{(j)}}{1 - \alpha q_j} = 0 \quad (4)$$

$$\frac{\partial L}{\partial \alpha} = \sum \frac{N - N_{00}^{(j)} - N_{11}^{(j)}}{\alpha} - \frac{p_j N_{00}^{(j)}}{1 - \alpha p_j} - \frac{q_j N_{11}^{(j)}}{1 - \alpha q_j} = 0$$

$$\xrightarrow{\alpha \neq 0} \sum \frac{N_{00}^{(j)}}{1 - \alpha p_j} + \frac{N_{11}^{(j)}}{1 - \alpha q_j} = kN \quad (5)$$

where $1 - \alpha$ is the continuity rate. If $\alpha = 0$ then all snapshots are identical, which is uninteresting, so we have $\alpha \neq 0$ in Equation (5).

A common way to solve such maximization problems is to employ numerical methods such as gradient descent. The drawback of such an approach is that it can be computationally expensive with hundreds of dyads and windows. Therefore, we settle for an approximation of the MLE that empirically approximates numerical values well². Intuitively, the estimator for α should depend on all the chains, but chains that spend more time in both states 0 and 1 provide more information about α than chains that spend most time in one state (the latter may be due to small α or to a value of p_j far from 1/2). Since we can easily compute the MLE for α for a single chain, we estimate α with a weighted average of the MLEs from the individual chains, with chains that spend more time in both states being weighted more heavily. We then estimate each p_j by the MLE for the j th chain, since the

²At significant level 0.05, two sample t -test shows the approximated values equal to the numerical values.

chains are conditionally independent given α . This results in the following estimators.

$$\hat{\alpha}_{\text{approx}} = \sum_j w_j \hat{\alpha}_{j\text{MLE}} = \sum_j w_j \frac{N_{01}^{(j)} N_{1*}^{(j)} + N_{10}^{(j)} N_{0*}^{(j)}}{N_{0*}^{(j)} N_{1*}^{(j)}} \quad (6)$$

$$\hat{p}_{j\text{approx}} = \hat{p}_{j\text{MLE}} = \frac{N_{01}^{(j)} N_{1*}^{(j)}}{N_{01}^{(j)} N_{1*}^{(j)} + N_{10}^{(j)} N_{0*}^{(j)}}$$

where $w_j = \frac{[N_{0*}^{(j)} N_{1*}^{(j)}]^p}{\sum_j [N_{0*}^{(j)} N_{1*}^{(j)}]^p}$, and $N_{0*}^{(j)} = N_{00}^{(j)} + N_{01}^{(j)}$

Empirically we find the exponent $p = \infty$ works best, which means $\hat{\alpha}_{\text{approx}}$ is a simple average of the $\alpha_{j\text{MLE}}$ corresponding to the chains with maximal value of $N_{0*}^{(j)} N_{1*}^{(j)}$. The continuity rate describes the temporal dependency among networks, and can help us determine a proper window size.

Drawbacks of MLE Though MLEs are consistent in general, there is no guarantee of unbiasedness for these particular estimators with limited samples. Moreover, they have three random quantities (N_{01} , N_{10} , N_{1*} , N_{0*} in (6) have 3 degrees of freedom for fixed s) and hence require more samples to estimate, making it prohibitive in practice.

2) Simplified Estimator

To overcome the drawbacks of MLE, we propose a simple estimator for the edge probability which is consistent and unbiased, has only one random quantity and therefore requires fewer samples. The simple estimator essentially estimates the edge frequency in each window. If we know changes happen rarely, and the process stays in equilibrium in most of the time, we can show the following estimator to be consistent and unbiased in equilibrium:

$$\hat{p}_{j\text{eq}} \equiv \frac{N_{1*}^{(j)} + e_j^{[s+(i-1)\eta]}}{s} = \frac{\# \text{ of 1s in the chain}}{s} \equiv \frac{N_1^{(j)}}{s} \quad (7)$$

which is the proportion of snapshots in which the dyad j being an edge within the window.

Proposition 1 In equilibrium, $\hat{p}_{j\text{eq}}$ is consistent as chain length (window size) increases.

Proof. $\pi_0^{(j)} \equiv P(\text{non-edge of chain } j \text{ in equilibrium})$, $s \equiv N_0^{(j)} + N_1^{(j)}$ (fixed), $N_1^{(j)} \equiv \# \text{ of 1s in the chain}$

By ergodic theorem [Givens and Hoeting, 2012], $\pi_1^{(j)} \xrightarrow{\text{almost surely}} \lim_{s \rightarrow \infty} N_1^{(j)} / s = \lim_{s \rightarrow \infty} \hat{p}_{j\text{eq}}$, where the first equation means almost sure convergence, and implies convergence in probability (estimator being consistent). \square

Proposition 2 In equilibrium, $\hat{p}_{j\text{eq}}$ is unbiased.

Proof. $p_{01}^{(j)} \equiv p_j \alpha$, $p_{10}^{(j)} \equiv (1 - p_j) \alpha$. (Figure 2)

$$\begin{aligned} \pi_0^{(j)} p_{01}^{(j)} + \pi_1^{(j)} p_{11}^{(j)} &= \pi_1^{(j)} \implies \pi_0^{(j)} p_{01}^{(j)} = \pi_1^{(j)} p_{10}^{(j)} \\ &\implies \pi_0^{(j)} \alpha p_j = \pi_1^{(j)} \alpha (1 - p_j) \\ &\implies p_j = \pi_1^{(j)} = E N_1^{(j)} / (N_1^{(j)} + N_0^{(j)}) = E N_1^{(j)} / s \\ &\implies E \hat{p}_{j\text{eq}} = E N_1^{(j)} / s = p_j \implies \hat{p}_{j\text{eq}} \text{ is unbiased} \end{aligned}$$

\square




Figure 3: Comparison between MLE (red, top) and Simplified Estimator (green, bottom) for change point detection on the same sequence as Figure 4 and Table 2. Distance measure KL is used and other measures have consistent results; horizontal bars are corresponding thresholds. MLE has large fluctuation, and increasing window size reduces fluctuation; MLE misses true changes at ③, ④ and ⑤, while edge frequency estimator has perfect recall and precision (see Section 5 for detail). (re-scaled and shifted for visualization)

The above propositions imply that the larger the window size the better the estimation, and that in equilibrium, the temporal dependency (continuity rate) has no impact on estimating the onset probability of a Markov chain, and hence no impact on estimating the edge probability of a snapshot.

Although MLE is close to the true value when the chain is long enough (100 or longer), we do not use so large a window size in practice (20 usually, no larger than 50). Experiments (Figure 3) show that the simplified estimator is much better than MLE for change point detection in practice.

4.2 Distance Measure

Now, we need to compare the probability distributions of edges across consecutive windows. Kolmogorov-Smirnov (KS) statistic and Kullback-Leibler (KL) divergence are two common measures for comparing distribution. Their calculations require the enumeration of the whole state space and hence exponential to the number of variables for joint distributions. Although KS statistic is designed for univariate distribution, we can map the joint distribution, which has multivariate binary variables, to one dimension by decoding the binary vectors as an integer and use KS statistic. We bootstrap from empirical distributions of two consecutive windows respectively and use two sample KS test to quantify the difference of two distributions. We can use divide-and-conquer to alleviate the exponential complexity: partition the dyads into g groups, compute KL/KS dissimilarity within each small group, and record the median among all the g groups as the final dissimilarity.

Both of the above measures have good quality in terms of change point detection (Figure 4), but KS statistic is extremely slow (see extended version [Wang *et al.*, 2017]), mostly due to the large sample bootstrap from each window. Euclidean distance, though lacks probabilistic interpretation, has linear complexity and has reasonable quality in practice.

4.3 Threshold Determination

Suppose we have w windows, then we compare $w - 1$ pairs of distributions and get $w - 1$ difference/distance scores. How do we choose a threshold to determine at which window the

Table 2: Model Change Explanation for an SBM-CL Experiment in Figure 4

Order	Window Index	Type of Change
①	15	The weight sequence of 1/3 of the nodes is re-generated
②	30	The weight sequence of 2/3 of the nodes is re-generated
③	60	Half of the communities change their (inter- and intra-community) connection rate, overall density retained
④	75	All of the communities change their (inter- and intra-community) connection rate, overall density retained
⑤	90	Half of the communities change their (inter- and intra-community) connection rate, overall density changed
⑥	105	All of the communities change their (inter- and intra-community) connection rate, overall density changed
⑦	135	Community assignments of all the nodes are changed

network changes? We use a permutation test [Pitman, 1937] based approach to determine the threshold. For a desired significance level α_s , we bootstrap from the $w - 1$ distance scores, and use the upper $100\alpha_s\%$ quantile as the threshold.

4.4 Complexity Analysis

The algorithm is linear to the number of windows and constant to the network size for moderately large network. Only a small fraction of dyads in the network is sampled and tracked. The sampling of the dyads is only performed once at the beginning, and hence irrelevant to the number of snapshots. For each snapshot, selecting a specific set of dyads has linear cost to the number of edges. Each window is only scanned once and therefore the time cost is linear in the number of windows. Moreover, since the number of windows is linear to the total number of snapshots even in the worst case (windows are overlapping, and window step is one), the algorithm is linear to the number of snapshots. Therefore, the time complexity is $O(MT)$, where M is the averaged number of edges in each snapshot, and T is the number of snapshots.

The memory cost is low, and can be viewed as constant: for each snapshot, only the information of the tracked dyads is stored; information of dyads within the same window is aggregated; dyads information in the old window is overwritten once it is compared against the new window. And the space complexity is $O(c)$, where c is a prescribed sample size. Theoretically the sample size should be proportional to the network size for good estimation. Our experiments show that a fixed sample size (to track 250 out of 1.2G dyads $\approx \binom{50k}{2}$) works well on a moderately large network.


5 Experiments and Results

We did thorough evaluation of our edge probability estimation based change point detection algorithm (called EdgeMonitoring for simplicity) on synthetic and real world datasets. For the synthetic datasets, the generative process is known, and we can compute the ground truth in the form of likelihood, which is naturally a baseline choice. We also use the state-of-the-art DeltaCon [Koutra *et al.*, 2016] and LetoChange [Peel and Clauset, 2015] as two baselines.


5.1 Synthetic Data

Data generation³ We generate a sequence of networks from a fixed generation model. The snapshots are not independent, each snapshot depends on the preceding one through the continuity parameter $\alpha(\alpha_t \equiv \alpha)$. For each snapshot, each edge

³Generated using SNAP[Leskovec and Sosič, 2016]



(a) Likelihood (ground truth)



(b) Algorithms comparison. Curves are dissimilarity scores and horizontal bars are thresholds, and they two have corresponding colors and line shapes. (re-scaled and shifted for visualization)

Figure 4: SBM, ground truth changes explained in Table 2. $1 - \alpha = 0.51$ and window size = 20. DeltaCon (Fig b-top) and EM-KL (Fig b-bottom) have the smallest variance, but DeltaCon has two false negatives at ④ and ⑤.

is selected independently with probability α , and if selected, the edge is again sampled from the generative model (Table 1). We introduce the change points by changing the generative model in the middle of the sequence of snapshots. Note that this change may be simply a change of parameter values for a given model (Eg. ER0.4 to ER0.6), or a change in the model type (Eg. SBM to ER), as well. Since our algorithm makes no assumptions about model specifics, we are able to detect both kinds of changes. We only inject parameter change in the synthetic experiment since the latter change is easily detectable. Sample changes are displayed in Table 2. The likelihood of the snapshot sequence is also provided.

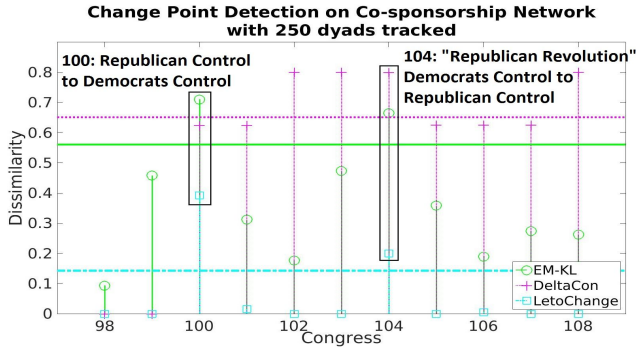


Figure 5: Change point detection on US Senate co-sponsorship network. Change points at the 100th and the 104th Congresses (boxed) correspond to partisan domination shifts. Both EM-KL (green) and LetoChange (cyan) have perfect recall and precision, while DeltaCon (pink) has 3 false positives and 1 false negative.

We ran experiments with network sizes ranging from 1k to 50k, window size to be from 10 to 100 and continuity rate $1 - \alpha$ to be 0.51 and 0.9. We generated a total of 5000 snapshots and sampled 250 edges uniformly at random to track. Both overlapping window ($s=2\eta$) and non-overlapping window have similar results, yet the latter is faster simply due to fewer windows. Hence we display non-overlapping window results only. For KL and KS, edges are grouped into 25 equal-sized groups. We use upper 5% quantile as the threshold.

Results Figure 4 shows the qualitative comparison and Table 3 reports the efficiency. Figure 4a shows the likelihood of the network drops dramatically after the generative model changes, and recovers to new equilibrium afterwards. Our EdgeMonitoring (EM-Eu, EM-KL) approach can successfully identify all change points with 5X speed up over DeltaCon. The changes are explained in Table 2. We can see that EM-KL has the best performance: little fluctuation and perfect precision and recall. DeltaCon, though has smallest fluctuation, misses two change points. Both EM-KS and EM-Eu have large fluctuation. The quality of EM-KS heavily relies on the joint probability estimation, and we do see smaller fluctuation and higher recall for larger window size. EM-Eu in general has large fluctuation. EM-KL has the best overall performance, in terms of both quality and time efficiency. We believe grouping together with median selection contribute to its superiority.

5.2 Real World Data

Senate cosponsorship network([Fowler, 2006]) We construct a co-sponsorship network from bills (co-)sponsored in US Senate during the 93rd-108th Congress. An edge is formed between two congresspersons if they cosponsored the same bill. Each bill corresponds to a snapshot, and forms a clique of co-sponsors. A window is set to include all bills in a single Congress (Biennially).

We randomly selected 250 dyads and tracked their fluctuations across the Congresses. We start from the 97th Congress since full amendments data is available only from 97th session onwards. Figure 5 compares EdgeMonitoring+KL, DeltaCon and LetoChange. All methods were able to detect the most significant change point at the 104th Congress. Fowler [Fowler, 2006] points out that there was a “Repub-

Table 3: Time efficiency comparison (5k snapshots)

Model	Network Size	Window Size	EM Time ¹	DC Time (speedup)	LC Time
CL	1k	20	18s	91s (5X)	DNF
SBM-CL	1k	10	27s	125s (5X)	DNF
SBM-CL	1k	50	9s	43s (5X)	DNF
SBM-CL	5k	20	54s	309s (6X)	DNF
SBM-CL	10k	20	232s	1920s (8X)	DNF
SBM-CL	50k	20	26m	240m (9X)	DNF
BTER ²	1k	20	3s	12s (4X)	6h
Figure 4	1k	20	21s	103s (5X)	DNF
Figure 5	100	biennial	4s	16s (4X)	13h
[Voeten, 2012]	≈200	annual	10s	93s (9X)	DNF
Enron	150	weekly	1s	1s (1X)	60h

¹ EM for EdgeMonitoring (running time includes both KL and Euclidean), DC for DeltaCon, LC for LetoChange. EM and DC are implemented in MATLAB while LC in Python. All run on a commercial desktop with 60hrs as time limit. Each running time averaged over 5 runs.

² BTER dataset has 800 snapshots

lican Revolution” in the 104th Congress which “*caused a dramatic change in the partisan and seniority compositions.*” The author also points out the significance of the 100th (highest clustering coefficient, significant collaboration) and 104th Congress (lowest clustering coefficient, low point in collaboration) as inflection points in the Senate political process. Both our EdgeMonitoring approach and LetoChange classify these two Congresses as change points, but the latter takes much more time. DeltaCon picks up on one (104th) and not the other (100th). This provides evidence that our algorithm is able to capture the changes in network evolution effectively while being significantly faster than the state-of-the-art.

6 Conclusion

We develop a change point detection algorithm for dynamic networks that is efficient and accurate. Our approach relies on sampling and comparing the estimated joint edge (dyad) distribution. We first develop a maximum likelihood estimator, and analyze its drawbacks for small window sizes (the typical case). We then develop a consistent and unbiased estimator that overcomes the drawbacks of the MLE, resulting in significant quality improvement over the MLE. We conduct a thorough evaluation of our change point detection algorithm against two state-of-the-art approaches on both synthetic and real-world data. Our results indicate that our method is up to 9X faster than the state-of-the-art while achieving better quality. In the future we plan to extend our work to track higher order structures of the network such as 3-profiles [Elenberg *et al.*, 2015] or 4-profiles and see how they evolve over time.

Acknowledgments

This work is supported in part by National Science Foundation (NSF) grants DMS-1418265, IIS-1550302 and IIS-1629548. Opinions and findings herein are of the author(s) and not necessarily the views of the NSF.

References

- [Akoglu *et al.*, 2014] Leman Akoglu, Hanghang Tong, and Danai Koutra. Graph-based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery (DAMI)*, 28(4), 2014.
- [Asur *et al.*, 2007] Sitaram Asur, Srinivasan Parthasarathy, and Duygu Ucar. An event-based framework for characterizing the evolutionary behavior of interaction graphs. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '07*. ACM, 2007.
- [Berlingerio *et al.*, 2012] Michele Berlingerio, Danai Koutra, Tina Eliassi-Rad, and Christos Faloutsos. Net-simile: a scalable approach to size-independent network similarity. *arXiv preprint arXiv:1209.2684*, 2012.
- [Bridges *et al.*, 2015] Robert A Bridges, John P Collins, Erik M Ferragut, Jason A Laska, and Blair D Sullivan. Multi-level anomaly detection on time-varying graph data. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*, pages 579–583. ACM, 2015.
- [Caceres and Berger-Wolf, 2013] Rajmonda Sulo Caceres and Tanya Berger-Wolf. Temporal scale of dynamic networks. In *Temporal Networks*, pages 65–94. Springer, 2013.
- [Elenberg *et al.*, 2015] Ethan R Elenberg, Karthikeyan Shanmugam, Michael Borokhovich, and Alexandros G Dimakis. Beyond triangles: A distributed framework for estimating 3-profiles of large graphs. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 229–238. ACM, 2015.
- [Erdős and Rényi, 1960] Paul Erdős and A Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 5:17–61, 1960.
- [Fowler, 2006] James H Fowler. Legislative cosponsorship networks in the US House and Senate. *Social Networks*, 28(4):454–465, 2006.
- [Givens and Hoeting, 2012] Geof H Givens and Jennifer A Hoeting. *Computational statistics*, volume 710. John Wiley & Sons, 2012.
- [Hunter *et al.*, 2012] David R Hunter, Pavel N Krivitsky, and Michael Schweinberger. Computational statistical methods for social network models. *Journal of Computational and Graphical Statistics*, 21(4):856–882, 2012.
- [Karrer and Newman, 2011] Brian Karrer and Mark EJ Newman. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83(1):016107, 2011.
- [Klimt and Yang, 2004] Bryan Klimt and Yiming Yang. The enron corpus: A new dataset for email classification research. In *Machine learning: ECML 2004*, pages 217–226. Springer, 2004.
- [Koutra *et al.*, 2016] Danai Koutra, Neil Shah, Joshua T Vogelstein, Brian Gallagher, and Christos Faloutsos. Deltacon: Principled Massive-Graph Similarity Function with Attribution. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 10(3):28, 2016.
- [La Fond *et al.*, 2014] Timothy La Fond, Jennifer Neville, and Brian Gallagher. Anomaly detection in networks with changing trends, 2014.
- [Leskovec and Sosič, 2016] Jure Leskovec and Rok Sosič. Snap: A general-purpose network analysis and graph-mining library. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(1):1, 2016.
- [Li *et al.*, 2016] Shuang Li, Yao Xie, Mehrdad Farajtabar, and Le Song. Detecting weak changes in dynamic events over networks. *arXiv preprint arXiv:1603.08981*, 2016.
- [Moreno and Neville, 2013] Sebastian Moreno and Jennifer Neville. Network hypothesis testing using mixed kronecker product graph models. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pages 1163–1168. IEEE, 2013.
- [Peel and Clauset, 2015] Leto Peel and Aaron Clauset. Detecting change points in the large-scale structure of evolving networks. 2015.
- [Peixoto and Rosvall, 2015] Tiago P Peixoto and Martin Rosvall. Modeling sequences and temporal networks with dynamic community structures. *arXiv preprint arXiv:1509.04740*, 2015.
- [Pfeiffer III *et al.*, 2012] Joseph J Pfeiffer III, Timothy La Fond, Sebastian Moreno, and Jennifer Neville. Fast generation of large scale social networks with clustering. *arXiv preprint arXiv:1202.4805*, 2012.
- [Pitman, 1937] Edwin JG Pitman. Significance tests which may be applied to samples from any populations. *Supplement to the Journal of the Royal Statistical Society*, 4(1):119–130, 1937.
- [Ranshous *et al.*, 2015] Stephen Ranshous, Shitian Shen, Danai Koutra, Steve Harenberg, Christos Faloutsos, and Nagiza F Samatova. Anomaly detection in dynamic networks: a survey. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7(3):223–247, 2015.
- [Seshadhri *et al.*, 2012] C Seshadhri, Tamara G Kolda, and Ali Pinar. Community structure and scale-free collections of erdős-rényi graphs. *Physical Review E*, 85(5):056109, 2012.
- [Shih and Parthasarathy, 2012] Yu-Keng Shih and Srinivasan Parthasarathy. Identifying functional modules in interaction networks through overlapping markov clustering. *Bioinformatics*, 28(18):i473–i479, 2012.
- [Voeten, 2012] Erik Voeten. Data and analyses of voting in the UN general assembly. *Available at SSRN 2111149*, 2012.
- [Wang *et al.*, 2017] Yu Wang, Aniket Chakrabarti, David Sivakoff, and Srinivasan Parthasarathy. Fast Change Point Detection on Dynamic Social Networks. *arXiv preprint arXiv:1705.07325*, 2017.
- [Zhang *et al.*, 2016] Xiao Zhang, Cristopher Moore, and MEJ Newman. Random graph models for dynamic networks. *arXiv preprint arXiv:1607.07570*, 2016.