
Changepoint Detection in Highly-Attributed Dynamic Graphs

Emiliano Penalosa^{*12} Nathaniel Stevens³

Abstract

Detecting anomalous behavior in dynamic networks remains a constant challenge. This problem is further exacerbated when the underlying topology of these networks is affected by individual highly-dimensional node attributes. We address this issue by tracking a network’s modularity as a proxy of its community structure. We leverage Graph Neural Networks (GNNs) to estimate each snapshot’s modularity. GNNs can account for both network structure and high-dimensional node attributes, providing a comprehensive approach for estimating network statistics. Our method is validated through simulations that demonstrate its ability to detect changes in highly-attributed networks by analyzing shifts in modularity. Moreover, we find our method is able to detect a real-world event within the #Iran Twitter reply network, where each node has high-dimensional textual attributes. We additionally make our code available on Github¹.

1. Introduction

A network or graph is a collection of interconnected items, called nodes or vertices, found in areas like brain networks (Budzinski et al., 2022), as well as in human-engineered contexts such as computer networks (Mahmood et al., 2015). Network science studies these structures, revealing their benefits in applications including social network mining (Tiukhova et al., 2022) and combatting human trafficking (Rabbany et al., 2018). As the field evolved, subfields like network monitoring emerged, focusing on changes in network behaviour over time. Methods for detecting changes include statistical process monitoring and anomaly detection in time series used in engineering, manufacturing, and fraud

¹Mila - Québec AI Institute ²DIRO, Université de Montréal ³Department of Statistics & Actuarial Science, University of Waterloo. Correspondence to: Emiliano Penalosa <emilianopp550@gmail.com>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

¹<https://github.com/Emilianopp/ChangesHighlyAttributedGraphs>

detection (Tsung et al., 2008; Ramrez & Ramrez, 2018; MacGregor & Kourti, 1995). Network monitoring often combines these methods, such as using statistical control charts or deep learning to detect anomalies in dynamic networks (Malinovskaya & Otto, 2021). Another subfield studies attributed networks, where nodes hold important information. Statistical inference methods are proficient at estimating parameters of the graph generation process but are largely limited by restrictive and unrealistic assumptions (Newman & Clauset, 2016). On the other hand, GNNs have achieved state-of-the-art performance in classification, regression, and generative tasks over attributed and unattributed networks without such assumptions (Dwivedi et al., 2020). Most network monitoring techniques ignore network attributes, focusing only on the network structure itself. Some works consider attributes but use simpler generalized linear models, limiting their capacity (Ebrahimi et al., 2021a; Gahrooei & Paynabar, 2018). This study presents a method for monitoring networks with many attributes by tracking the graph’s community structure. Communities are groups of similar, frequently interacting individuals. We find these groups by maximizing network modularity using a GNN, leveraging structural properties and high-dimensional attributes. When monitoring a time series of graphs, we detect structural changes when modularity changes significantly. Through simulations, we demonstrate that our method can detect a variety of changes in highly attributed networks. Additionally, we showcase our methods’ practicality on a real world Twitter network, where the node attributes are derived from user tweets. We relegate the related work section to App. B, where we note to the best of our knowledge this work is the first to study anomaly detection on *highly-attributed dynamic* graphs, specifically those with textual attributes.

2. Background

2.1. Dynamic Networks

A dynamic network can be described as an ordered sequence of time varying graphs $\mathcal{G}(\mathcal{T}) = \{G_1, \dots, G_T\}$ over some timeframe $\mathcal{T} = \{1, \dots, T\}$. Each of these graphs may be independent, but most often they exhibit temporal dependence. We refer to these instances G_1, G_2, \dots, G_T as network snapshots. Each snapshot contains a set of

nodes, $\mathcal{V}(t) = \{1, 2, \dots, n_t\}$, edges $\mathcal{E}(t) = \{e_{u,v}(t) : u, v \in \mathcal{V}(t), t \in \mathcal{T}\}$, and a set of s -dimensional attributes $\mathcal{X}(t) = \{x_u(t) : u \in \mathcal{V}(t), t \in \mathcal{T}\}$. It is important to note that given our definition of a dynamic network, the set of edges and nodes are expected to change between snapshots. Additionally, we define the neighborhood of first-degree connections of a node u at time t as $\mathcal{N}_u(t) = \{v : e_{u,v}(t) \in \mathcal{E}(\mathcal{T})\}$ at timestamp t . In this work, we encode each network snapshot using an adjacency matrix, denoted \mathbf{A}_t . The matrix can either have binary entries $A_{uv,t} = e_{u,v}(t) \in \{0, 1\}$ or real-valued edge weights $A_{uv,t} = e_{u,v}(t) \in \mathbb{R}$.

2.2. Deep Modularity Networks

Traditionally, graph communities are obtained using self-supervised learning combined with k-means clustering on the derived node representations. However, these methods can be unreliable depending on the graph’s size and sparsity (Xia et al., 2022). To address this, it is effective to directly optimize for a statistic that measures community structure. For example, modularity which measures the difference between the observed number of intra-cluster edges and what would be expected in a randomly generated graph.

$$Q = \frac{1}{2w} \sum_{uv} \left[A_{uv} - \frac{d_u d_v}{2w} \right] \delta_{c_u, c_v} \quad (1)$$

where δ_{c_u, c_v} is the Kronecker delta, w indicates the total number of edges in the graph, $d_u = |\mathcal{N}_u|$ is the degree of node u , and c_u indicates the community assignment for node u . Unfortunately, the computation of the modularity gradient is intractable, making spectral modularity, a convenient alternative (Newman, 2006):

$$Q = \frac{1}{2w} \text{Tr}(C^\top BC) \quad (2)$$

where $C \in \{0, 1\}^{n \times k}$ is a community assignment matrix, k is the number of communities, and $B = A - \frac{dd^\top}{2w}$ with $d \in \mathbb{R}^n$ being the degree vector. Traditionally, one can maximize the optimal modularity by choosing the assignment matrix C to equal the top k eigenvectors of B , but such a derivation is agnostic to node attributes. Recently, Tsitsulin et al. (2023) showed that a GNN can be used to directly optimize spectral modularity and as a byproduct generate cluster assignments which can account for anode ttributes. Deep Modularity Networks (DMoN), achieve this by using a GNN to generate, a modified version of the community affiliation matrix, $C \in \mathbb{R}^{(0,1)^{n \times k}}$, in spectral modularity:

$$C = \text{softmax}(\text{GNN}(\tilde{A}, \mathcal{X})) \quad (3)$$

where the GNN, usually a graph convolutional network (Kipf & Welling, 2017), takes in the degree normalized adjacency matrix, $\tilde{A} = D^{-\frac{1}{2}}(A)D^{-\frac{1}{2}}$. Given this, DMoN

is able to directly optimize spectral modularity, through traditional gradient optimization, yielding a solution that accounts for both graph structure and attributes.

2.3. Methodology

Since our goal is to identify changes in highly attributed *dynamic* graphs, we can optimize community assignments using DMoN and directly track the associated modularity metric. Other than introducing DMoN, Tsitsulin et al. (2023) also introduced the collapse regularizer (CR) that penalizes trivial solutions that arise when optimizing spectral modularity:

$$\text{CR} = \frac{\sqrt{k}}{n} \left\| \text{ColSum}(C) \right\|_2 - 1. \quad (4)$$

For example, this regularization penalizes the case for which all k communities collapse into a single one by taking its maximum value at such a solution. While this regularizer provides stable solutions for cases when the model tries to collapse the clusters onto as single community, it fails to account for the trivial solution of assigning equal weight to all clusters for each node. We instead propose the square root collapse optimizer (SRCO) to address these limitations:

$$\text{SRCO} = \frac{1}{\sqrt{n}} \left\| \text{ColSum}(\sqrt{C}) \right\|_2 - 1. \quad (5)$$

To discourage the model from collapsing onto homogeneous assignments, we take the element-wise square root of the community assignment matrix before applying the column sum. A simple demonstration of the effects of our procedure is shown in App. C. Given this, our full model can be fit end-to-end with the following cost function:

$$\mathcal{L}_{\text{DMoN}} = - \underbrace{\frac{1}{2w} \text{Tr}(C^\top BC)}_{\text{Modularity Loss}} + \underbrace{\frac{1}{\sqrt{n}} \left\| \text{ColSum}(\sqrt{C}) \right\|_2}_{\text{Regularizer}} - 1. \quad (6)$$

We use statistical process monitoring (SPM), a form of anomaly detection to monitor change in modularity. In general, SPM involves modeling a system’s usual behavior and identifying abnormalities by prospectively tracking a key statistic, S_t , in our case Q_t , the estimated modularity at a snapshot t . Monitoring occurs in two phases: Phase I sets control limits based on baseline behavior from sampled statistics, while Phase II monitors new values, flagging anomalies if they fall outside these limits. Control charts plot these limits together with a time series of observations. Different control charts detect various changes, including Shewhart, cumulative sum (CUSUM), and exponentially weighted moving average (EWMA) charts. EWMA charts are favored for detecting gradual and sustained changes by modeling the exponentially weighted moving average, Z_t , of S_t :

$$Z_t = \alpha S_t + (1 - \alpha) Z_{t-1} \quad (7)$$

where $0 < \alpha < 1$ is a smoothing constant, with Z_0 often set as the estimated mean $\hat{\mu} = \bar{S}$. Phase I data determines

Change Type	Change Detection Percentage	Conditional Expected Delay	Average Percentage Over Threshold
Community Splitting	1.0	1.0	1.0
Community Merging	1.0	1.0	1.0
New Community	1.0	1.0	1.0
Attribute Change	1.0	23.38	0.44

Table 1. Performance summaries across $N=100$ replications of the community splitting change

control limits:

$$\hat{\mu} \pm 3\hat{\sigma} \sqrt{\frac{\alpha}{2-\alpha} [1 - (1-\alpha)^{2t}]} \quad (8)$$

with $\hat{\sigma}$ being the estimated standard deviation of S_t . The smoothing parameter α is crucial for sensitivity; values closer to 1 emphasize recent observations, affecting detection rates. An α of 0.2 is commonly used to balance gradual change detection and limit false alarms (Lucas & Saccucci). We note that we directly train our DMoN model on Phase I snapshots and reserve Phase II purely for evaluation. Our full procedure is visualized and further described in App. D

3. Experiments

3.1. Data Generation

We evaluate the performance of the proposed methodology with simulated experiments and synthetic data. Using the data generation technique from Tsitsulin et al. (2022), we create data resembling real-world attributed graphs. This involves two steps: graph construction and attribute generation. First, we construct an initial graph using a degree-corrected stochastic block model (DCSBM) (Karrer & Newman, 2011). The DCSBM models each edge weight via a Poisson distribution, with a mean depending on the nodes' community memberships and their propensity to connect. The probability mass function for the model is:

$$P(A|\theta, \Lambda, \mathbf{c}) = \prod_{u < v} \frac{1}{A_{uv}!} e^{-\lambda_{c_u, c_v} \theta_u \theta_v} (\lambda_{c_u, c_v} \theta_u \theta_v)^{A_{uv}} \times \prod_u \frac{1}{(A_{uv}/2)!} e^{-\lambda_{c_u, c_u} \theta_u^2 / 2} (\lambda_{c_u, c_u} \theta_u^2 / 2)^{A_{uv}/2} \quad (9)$$

Here, θ is the degree propensity vector, with θ_u being the degree propensity for node u , proportional to its expected degree. To ensure identifiability, the added normalization constraint $\sum_u \theta_u \delta_{c_u, r} = 1$ must be applied. Λ is the $k \times k$ community propensity matrix, where an entry $\lambda_{r,s}$ indicates the edge propensity between communities r and s . Here \mathbf{c} is the community assignment vector, with an entry $c_u \in \{1, 2, \dots, n\}$ indicating the community assignment of node u . Node attribute information is generated from cluster-specific multivariate normal distributions, with cluster centers generated using a Gaussian mixture model. For each of the s node attributes, we sample their attribute cluster centers from a k -dimensional multivariate Gaussian distribution with mean $\mu = 0$ and covariance matrix $\Sigma = 3 \times \mathbf{I}_{k \times k}$. Using the sampled cluster centers μ_c , we then sample the node

attributes from an s -dimensional multivariate Gaussian distribution with covariance equal to the identity matrix. The simulation design includes 50 "in control" graphs for Phase I. After injecting a change point, 50 "out of control" graphs are monitored in Phase II. Each graph consists of $n = 1,000$ nodes in one of $k = 4$ equally sized communities. Node propensities, θ_u , are generated using a uniform power law between 2^2 and 2^6 . The intra-community propensity to connect is $\lambda_{s,s} = 18$, and the inter-community propensity is $\lambda_{r,s} = 2$ for $r, s \in \{1, 2, \dots, k\}$. Each node's attributes ($s = 64$) are sampled from their respective multivariate Gaussian distribution with mean μ_c and covariance matrix $\mathbf{I}_{s,s}$. This framework describes a single simulation run, we repeat $N = 100$ times. Performance is evaluated using three metrics: change detection percentage, conditional expected delay, and average percentage of observations exceeding control limits. Change detection percentage is the proportion of runs in which a change was detected in at least one of the 50 out of control networks. Conditional expected delay is the average steps to detect the change, given detection occurs. The average percentage of the 50 out-of-control observations exceeding control limits, averaged over 100 runs, is also calculated. These metrics collectively provide insights into the method's change detection effectiveness and efficiency.

3.2. Large Changes

We evaluate four changes to assess the efficacy of our method: Community splitting, community merging, new community, and attribute change. *Community Splitting*: Given the baseline parameters, we choose one of the four existing communities and randomly split its nodes into two new equally sized communities as well as sample new cluster centers, μ_c , for each of the two partitioned communities. We note this change does not affect the propensities to connect $\lambda_{r,r}$ and $\lambda_{r,s}$ remain unchanged. *Community Merging*: to replicate real-life scenarios, our experimental set-up consists of merging all nodes of two communities into one. Additionally, we average the two community cluster centers to create the merged communities attribute centers. For example, the attribute center for the new community formed by merging communities r and s is $\mu_{c,merged} = \frac{\mu_{c,r} + \mu_{c,s}}{2}$. *New Community*: we model a sudden increase in nodes with unseen attributes by first increasing the number of nodes in the graph by 25%. Since our base scenario has $n = 1,000$ nodes, after the new community is inserted, the graph will

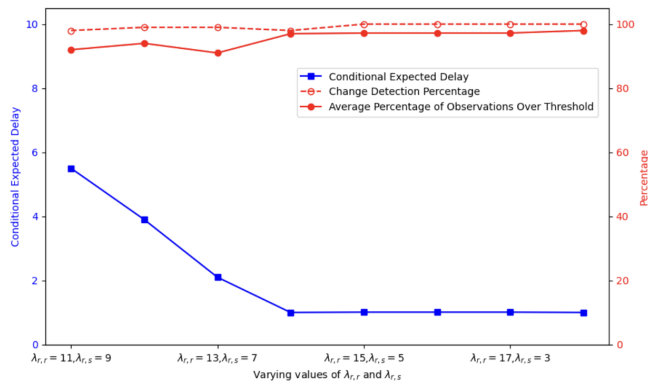


Figure 1. Results of synthetic experiments over varying community propensity, $\lambda_{r,r}$ and $\lambda_{r,s}$, values. Here each point represents $N = 100$ simulations at their respective values of $\lambda_{r,r}$ and $\lambda_{r,s}$

have 1,250 nodes. Thereafter, we repeat the same procedure to generate the attribute information as with the first communities. *Attribute change*: In contrast to sudden and abrupt changes, networks may also undergo gradual shifts in their underlying structure. We model such changes by periodically adding noise to the prior cluster centers μ_c . The process can be described by the following cluster center parameterization $\mu_c^{i+1} = \mu_c^i + \tau$ where τ is a s -dimensional uniformly distributed noise term. That is, $\tau_j \sim U(0, 1)$ for each $\tau_j \in \tau$, and i indicates the current iteration of the noise-adding process. Additionally, we note noise is only added during Phase II. Table 1 shows the evaluation for the reported changes. We observe our method is very effective at detecting these changes, with a change detection percentage of 1 and a very low expected delay. Additionally, we see that when the attributes change slowly, we are still able to detect this change reasonably quickly with a conditionally expected delay of 23.38 and a change detection percentage of 1.

3.3. Structural Change

Another important case is when node attributes are unchanged, but the graph generation process is altered. For this case, we consider imposing a change in the level of interaction between nodes by altering $\lambda_{r,r}$ and $\lambda_{r,s}$. We analyze a grid of changes in which we alter each pair of values sequentially as $\lambda_{s,r} = \lambda_{s,r} + 1$ and $\lambda_{s,s} = \lambda_{s,s} - 1$ and we repeat this process until $\lambda_{s,r} = \lambda_{s,s}$. The resulting grid of λ values are in the range of $\{10, \dots, 18\}$ for $\lambda_{r,s}$ and $\{2, \dots, 10\}$ for $\lambda_{s,s}$. We visualize how the inherent block structure changes in App. E. Figure 1 illustrates the performance of our method across these varying changes, each step illustrates $N = 100$ simulations with those configurations. We observe our method is quite accurate at detecting structural graph changes. Although subtler structural changes, on average, require a few more timestamps, they are still detected fairly quickly.

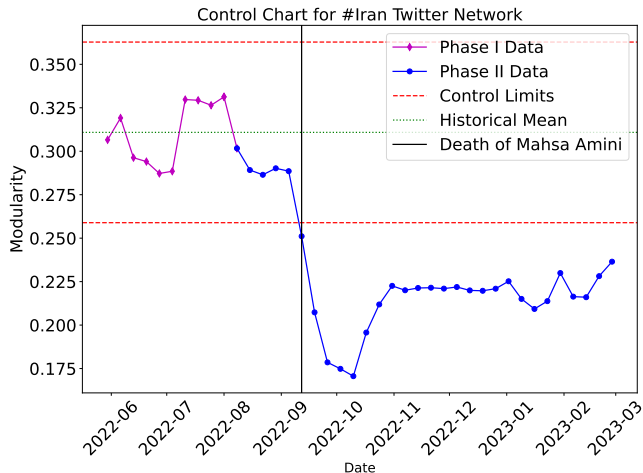


Figure 2. Control Chart for #Iran Twitter Network when $\alpha = 0.2$.

3.4. Twitter Dataset

In addition to simulation studies, we analyze a real-world Twitter dataset. On September 16, 2022, there was massive outrage within the Iranian community due to the death of Mahsa Amini, leading to increased activity in the #Iran Twitter network (Rel, 2023). Here we explore whether our methodology can detect changes in networks during such real-world events. To determine this, we construct a network by using the Twitter API to extract all tweets containing the #Iran hashtag. These tweets were collected over an eleven-month period and aggregated weekly; users are connected if they reply to each other. With these weekly snapshots, we build Phase I data from the first two months and Phase II with the rest. We use the text embeddings of user tweets obtained from a Sentence-Multilingual BERT model (Reimers & Gurevych, 2019; Conneau et al., 2018) as node attributes. Figure 2 shows the control chart for this network. Clearly, the method detects a change point during the week of the incident. This demonstrates the practicality of our method when dealing with a real-world network with high-dimensional textual node attributes. We provide additional training and dataset details in App. G.

4. Conclusion

In this work, we propose a methodology for detecting community changes in highly attributed networks. Additionally, we derive the square root collapse regularizer, which penalizes suboptimal solutions previous regularizers failed to. Through extensive simulation studies, we showcase our method’s ability to detect changes common in real-life networks. Additionally, we showcase the practicality of our method by analyzing a real-world network with textual node attributes. We demonstrate that our method is capable of detecting change associated with real-world events.

References

- On the 1st anniversary of mahsa amini’s death, a un submission sheds light on dire human rights for women girls in iran), Sept 2023. URL <https://ur0.jp/BfJf5>.
- Alotaibi, N. and Rhouma, D. A review on community structures detection in time evolving social networks. *Journal of King Saud University - Computer and Information Sciences*, 34(8):5646–5662, September 2022. doi: 10.1016/j.jksuci.2021.08.016. URL <https://doi.org/10.1016/j.jksuci.2021.08.016>.
- Artico, I. and Wit, E. C. Dynamic latent space relational event model. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 186(3):508–529, April 2023. doi: 10.1093/jrjsssa/qnad042. URL <https://doi.org/10.1093/jrjsssa/qnad042>.
- Atashgahi, Z., Mocanu, D. C., Veldhuis, R., and Pechenizkiy, M. Memory-free online change-point detection: A novel neural network approach, 2022.
- Budzinski, R. C., Nguyen, T. T., Đoà n, J., Mináč, J., Sejnowski, T. J., and Muller, L. E. Geometry unites synchrony, chimeras, and waves in nonlinear oscillator networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 32(3):031104, mar 2022. doi: 10.1063/5.0078791. URL <https://doi.org/10.1063%2F5.0078791>.
- Conneau, A., Lample, G., Rinott, R., Williams, A., Bowman, S. R., Schwenk, H., and Stoyanov, V. Xnli: Evaluating cross-lingual sentence representations, 2018.
- Dwivedi, V. P., Joshi, C. K., Laurent, T., Bengio, Y., and Bresson, X. Benchmarking graph neural networks. *CoRR*, abs/2003.00982, 2020. URL <https://arxiv.org/abs/2003.00982>.
- Ebrahimi, S., Reisi-Gahrooei, M., Paynabar, K., and Mankad, S. Monitoring sparse and attributed networks with online hurdle models. *IISE Transactions*, pp. 1–14, January 2021a. doi: 10.1080/24725854.2020.1861390. URL <https://doi.org/10.1080/24725854.2020.1861390>.
- Ebrahimi, S., Reisi-Gahrooei, M., Paynabar, K., and Mankad, S. Monitoring sparse and attributed networks with online hurdle models. *IISE Transactions*, pp. 1–14, jan 2021b. doi: 10.1080/24725854.2020.1861390. URL <https://doi.org/10.1080%2F24725854.2020.1861390>.
- Gahrooei, M. R. and Paynabar, K. Change detection in a dynamic stream of attributed networks. *Journal of Quality Technology*, 50(4):418–430, October 2018. doi: 10.1080/00224065.2018.1507558. URL <https://doi.org/10.1080/00224065.2018.1507558>.
- Gupta, M., Wadhvani, R., and Rasool, A. Real-time change-point detection: A deep neural network-based adaptive approach for detecting changes in multivariate time series data. *Expert Systems with Applications*, 209:118260, dec 2022. doi: 10.1016/j.eswa.2022.118260. URL <https://doi.org/10.1016%2Fj.eswa.2022.118260>.
- Hashemi, F., Behrouz, A., and Hajidehi, M. R. CS-TGN: Community search via temporal graph neural networks. In *Companion Proceedings of the ACM Web Conference 2023*. ACM, apr 2023. doi: 10.1145/3543873.3587654. URL <https://doi.org/10.1145%2F3543873.3587654>.
- Heard, N. A., Weston, D. J., Platanioti, K., and Hand, D. J. Bayesian anomaly detection methods for social networks. *The Annals of Applied Statistics*, 4(2), June 2010. doi: 10.1214/10-aos329. URL <https://doi.org/10.1214/10-aos329>.
- Huang, S., Hitti, Y., Rabusseau, G., and Rabbany, R. Laplacian change point detection for dynamic graphs. *CoRR*, abs/2007.01229, 2020. URL <https://arxiv.org/abs/2007.01229>.
- Joseph, S. R., Hlomani, H., Letsholo, K., Kaniwa, F., and Sedimo, K. Natural language processing: A review. *International Journal of Research in Engineering and Applied Sciences*, 6(3):207–210, 2016.
- Karrer, B. and Newman, M. E. J. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83(1), jan 2011. doi: 10.1103/physreve.83.016107. URL <https://doi.org/10.1103%2Fphysreve.83.016107>.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks, 2017.
- Lucas, J. M. and Saccucci, S. Exponentially weighted moving average control schemes: Properties and enhancements. 32(1).
- MacGregor, J. and Kourti, T. Statistical process control of multivariate processes. *Control Engineering Practice*, 3(3):403–414, March 1995. doi: 10.1016/0967-0661(95)00014-1. URL [https://doi.org/10.1016/0967-0661\(95\)00014-1](https://doi.org/10.1016/0967-0661(95)00014-1).
- Mahmood, M. A., Seah, W. K., and Welch, I. Reliability in wireless sensor networks: A survey and challenges ahead. *Computer networks*, 79:166–187, 2015.
- Malinovskaya, A. and Otto, P. Online network monitoring. *Statistical Methods & Applications*, 30(5):1337–1364, September 2021. doi: 10.1007/s10260-021-00589-z. URL <https://doi.org/10.1007/s10260-021-00589-z>.

- McCulloh, I., Carley, K., and Webb, M. Social network monitoring of al-qaeda. Technical report, January 2007. URL <https://doi.org/10.21236/ada488339>.
- Motalebi, N., Stevens, N. T., and Steiner, S. H. Hurdle blockmodels for sparse network modeling. *The American Statistician*, 75(4):383–393, feb 2021. doi: 10.1080/00031305.2020.1865199. URL <https://doi.org/10.1080%2F00031305.2020.1865199>.
- Newman, M. E. J. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, 74(3), sep 2006. doi: 10.1103/physreve.74.036104. URL <https://doi.org/10.1103%2Fphysreve.74.036104>.
- Newman, M. E. J. and Clauset, A. Structure and inference in annotated networks. *Nature Communications*, 7(1), June 2016. doi: 10.1038/ncomms11863. URL <https://doi.org/10.1038/ncomms11863>.
- Qi, G.-J., Aggarwal, C. C., and Huang, T. S. Online community detection in social sensing. In *Proceedings of the sixth ACM international conference on Web search and data mining*. ACM, February 2013. doi: 10.1145/2433396.2433474. URL <https://doi.org/10.1145/2433396.2433474>.
- Rabbany, R., Bayani, D., and Dubrawski, A. Active search of connections for case building and combating human trafficking. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & amp Data Mining*. ACM, July 2018. doi: 10.1145/3219819.3220103. URL <https://doi.org/10.1145/3219819.3220103>.
- Ramrez, B. S. and Ramrez, J. G. *Douglas Montgomery’s Introduction to Statistical Quality Control: A JMP Companion*. SAS Institute Inc., USA, 2018. ISBN 1635260221.
- Reimers, N. and Gurevych, I. Sentence-bert: Sentence embeddings using siamese bert-networks, 2019.
- Stevens, N. T. and Wilson, J. D. The past, present, and future of network monitoring: A panel discussion. *Quality Engineering*, 33(4):715–718, October 2021. doi: 10.1080/08982112.2021.1974032. URL <https://doi.org/10.1080/08982112.2021.1974032>.
- Sun, F.-Y., Qu, M., Hoffmann, J., Huang, C.-W., and Tang, J. vgraph: A generative model for joint community detection and node representation learning, 2019.
- Tiukhova, E., Penalzoza, E., Óskarsdóttir, M., Garcia, H., Bahnsen, A. C., Baesens, B., Snoeck, M., and Bravo, C. Influencer detection with dynamic graph neural networks, 2022.
- Tsitsulin, A., Rozemberczki, B., Palowitch, J., and Perozzi, B. Synthetic graph generation to benchmark graph learning, 2022.
- Tsitsulin, A., Palowitch, J., Perozzi, B., and Müller, E. Graph clustering with graph neural networks, 2023.
- Tsung, F., Li, Y., and Jin, M. Statistical process control for multistage manufacturing and service operations: a review and some extensions. *International Journal of Services Operations and Informatics*, 3(2):191, 2008. doi: 10.1504/ijsoi.2008.019333. URL <https://doi.org/10.1504/ijsoi.2008.019333>.
- Voulodimos, A., Doulamis, N., Doulamis, A., and Protopadakis, E. Deep learning for computer vision: A brief review. *Computational Intelligence and Neuroscience*, 2018:1–13, 2018. doi: 10.1155/2018/7068349. URL <https://doi.org/10.1155/2018/7068349>.
- Wang, Y., Chakrabarti, A., Sivakoff, D. J., and Parthasarathy, S. Fast change point detection on dynamic social networks. In *International Joint Conference on Artificial Intelligence*, 2017.
- Wang, Y., Cao, J., Bu, Z., and Leng, M. Temporal dual-attributed network generation oriented community detection model. *IEEE Transactions on Emerging Topics in Computing*, pp. 1–15, 2022. doi: 10.1109/tetc.2022.3223058. URL <https://doi.org/10.1109/tetc.2022.3223058>.
- Wilson, J. D., Stevens, N. T., and Woodall, W. H. Modeling and detecting change in temporal networks via the degree corrected stochastic block model. *Quality and Reliability Engineering International*, 35(5):1363–1378, July 2019. doi: 10.1002/qre.2520. URL <https://doi.org/10.1002/qre.2520>.
- Xia, W., Wang, Q., Gao, Q., Zhang, X., and Gao, X. Self-supervised graph convolutional network for multi-view clustering. *IEEE Transactions on Multimedia*, 24:3182–3192, 2022. URL [10.1109/TMM.2021.3094296](https://doi.org/10.1109/TMM.2021.3094296).
- Xie, J., Chen, M., and Szymanski, B. K. LabelRankT. In *Proceedings of the Workshop on Dynamic Networks Management and Mining*. ACM, June 2013. doi: 10.1145/2489247.2489249. URL <https://doi.org/10.1145/2489247.2489249>.
- Yu, L., Woodall, W. H., and Tsui, K.-L. Detecting node propensity changes in the dynamic degree corrected stochastic block model. *Social Networks*, 54:209–227, July 2018. doi: 10.1016/j.socnet.2018.03.004. URL <https://doi.org/10.1016/j.socnet.2018.03.004>.

Yu, W., Wang, W., Jiao, P., and Li, X. Evolutionary clustering via graph regularized nonnegative matrix factorization for exploring temporal networks. *Knowledge-Based Systems*, 167:1–10, 2019.

Zhao, Z., Li, C., Zhang, X., Chiclana, F., and Viedma, E. H. An incremental method to detect communities in dynamic evolving social networks. *Knowledge-Based Systems*, 163:404–415, January 2019. doi: 10.1016/j.knosys.2018.09.002. URL <https://doi.org/10.1016/j.knosys.2018.09.002>.

A. Acknowledgements

We note the first author of this work is a member of the Latinx community. They conducted all coding, ran all experiments and wrote the majority of this manuscript.

B. Related Work

B.1. Network Monitoring

Network monitoring aims to identify both sporadic and ongoing changes in a dynamic network, and has become increasingly popular within the field of Network Science. One way to approach the problem is by extending process monitoring techniques commonly used in manufacturing and engineering to the network setting ((Stevens & Wilson, 2021)). (McCulloh et al., 2007) were the first to coin the term network monitoring, and apply traditional statistical process monitoring techniques to networks. By surveying network summary statistics, such as betweenness, closeness, and density, they were able to identify changes in communication within the Al-Qaeda terrorist organization. There have been many further extensions of this work, where rather than simply tracking local graph statistics, the graph is assumed to follow a specified generative model whose parameter estimates can then be monitored. Some examples of such generative models are the degree corrected stochastic block model (DCSBM) ((Karrer & Newman, 2011), (Wilson et al., 2019), (Yu et al., 2018)), the hurdle block model ((Motalebi et al., 2021)) or the dynamic latent space model ((Artico & Wit, 2023)) to name a few. Additionally, (Heard et al., 2010) leveraged a Bayesian predictive distribution to define control limits to monitor levels of communication between individuals using either a Poisson or hurdle Poisson conditional distribution.

While impactful, the mentioned methods lack the ability to account for covariate information, which is more often than not crucial to modeling the underlying system dynamics. Fortunately, there are works that take this into account, for example, (Gahrooei & Paynabar, 2018) leverage a logistic regression model, which is able to account for node and/or edge covariates in order to determine the probability of an edge forming between two individuals. This methodology is convenient as they are able to directly monitor the model coefficients and attribute a change directly to a specific covariate. Likewise, (Ebrahimi et al., 2021b) leverage covariates through a hurdle model to explain the propensity of edge weights. Unfortunately, these methods are limited by the linearity of the underlying model, prioritizing interpretability over complexity.

While using statistical process monitoring methods to survey changing networks is quite popular, many other methods exist to identify anomalies in such systems. (Huang et al., 2020) choose to monitor dynamic graphs by observing the behavior of the singular values of the graph Laplacian matrix. Likewise, (Wang et al., 2017) tracks a fixed number of dyads over the observed snapshots. They model the distribution of dyads through a conditionally independent two-state Markov chain over a window of snapshots, detecting a change point if the Kullback–Leibler distance between current and historical snapshots exceeds a threshold.

B.2. Temporal Community Detection

Temporal community detection has become an increasingly important problem over the past decade due to the rise of social-media, online marketing, and the overall increase in people’s online presence. There are many ways to frame the problem of temporal community detection. (Xie et al., 2013) view temporal communities as a static entity, that is, a node’s community label is detected over the proposed series of network snapshots and considered fixed over time. Under this view, they build the LabelRankT algorithm, which is a heuristic algorithm that allows similarly linked nodes to be assigned the same label. In contrast, (Qi et al., 2013) proposes viewing the evolution of network communities in a continuous fashion.

Although useful, these methods are highly sensitive to user assumptions such as number of communities, sensitivity to time decay, and underlying hyperparameter choices ((Alotaibi & Rhouma, 2022)). Many works follow and extend the main ideas of the prior two methods, and similarly experience the same limitations. (Yu et al., 2019), generalize non-negative matrix factorization to build a community matrix that is able to be tracked and modeled through time but is highly sensitive to initial assumptions of the underlying community structure. (Zhao et al., 2019) build a heuristic framework for evolving community detection that does not require an assumed initial number of communities, but its heuristic nature does not allow for a quantifiable way to measure the significance of the change.

Other than the mentioned problems, another common limitation of the aforementioned methods is their lack of covariate modeling. (Wang et al., 2022) consider this by building a generative model that takes into account two types of covariates: *assortative* attributes, those associated with the individual (i.e., age, gender, etc.) and *generative* attributes, those that

influence a link formation (i.e., shared hobbies, occupation, etc). While this work is able to model the evolution of communities over time, it incorporates attributes through a logistic regression model, limiting it by the required assumptions of linear models.

B.3. Deep Learning for Community Based Network Monitoring

Deep learning has become an increasingly popular method when learning high dimensional data and has achieved state-of-the-art performance on text ((Joseph et al., 2016)), image ((Voulodimos et al., 2018)), and graph related tasks ((Dwivedi et al., 2020)). While not in direct application to networks, deep learning has even been used to monitor distribution shifts in dynamic systems, but has mainly been applied to raw time series ((Gupta et al., 2022),(Atashgahi et al., 2022)). Additionally, (Hashemi et al., 2023) developed a DL learning architecture that is able to detect changes to network communities over time and with the ability to adapt to changes in distribution through a meta-learning framework. Unfortunately, their work is rooted in a supervised learning framework that requires ground truth community labels, something that is very rare in real-world datasets.

Unlike the aforementioned methods, there exists DL-based community detection methods that are not restricted by the need for ground truth labels. For example, (Sun et al., 2019) assume edges are generated using a hierarchical procedure with a hidden variable representing node community membership. They then use a variational approach to reconstruct the adjacency matrix, concurrently finding the most likely community memberships. Similarly, (Tsitsulin et al., 2023) propose deep modularity networks (DMoN) which find local graph partitions, synonymous with communities, by directly optimizing the spectral modularity statistic. While these models are able to find high-quality clusters, they operate on a single network and do not generalize to a set of time-varying networks.

C. Regularizers

A simple demonstration of the effects of our procedure are shown in Figure 3 which compares our square root collapse regularizer and the collapse regularizer. For this illustration, we use a simple example with 3 evenly sized communities, $k = 3$, but it generalizes to any $k \geq 2$. When employing the collapse regularizer, assigning a uniform distribution over the community assignments to all the nodes is equivalent to the optimal solution of assigning communities with certainty. On the other hand, the square root collapse regularizer addresses this issue by providing a larger value for the suboptimal solution, which can then be penalized.

D. Training Setup

To be able to utilize a control chart with our method, we define Phase I by first training our model over a dynamic graph with snapshots, G_1, \dots, G_m . This allows our model to learn the stationary behavior of the network. Once converged, we calculate the modularity of the remaining unobserved snapshots, G_{m+1}, \dots, G_T . In Phase II, we track the behavior of the graph using an EWMA control chart, where we choose S_t to be the modularity scores of each successive graph. Through extensive synthetic experimentation, we demonstrate that this methodology is able to robustly detect changes in both graph attributes and graph structural changes. Our methodology is illustrated in Figure 4.

Changepoint Detection in Highly-Attributed Dynamic Graphs

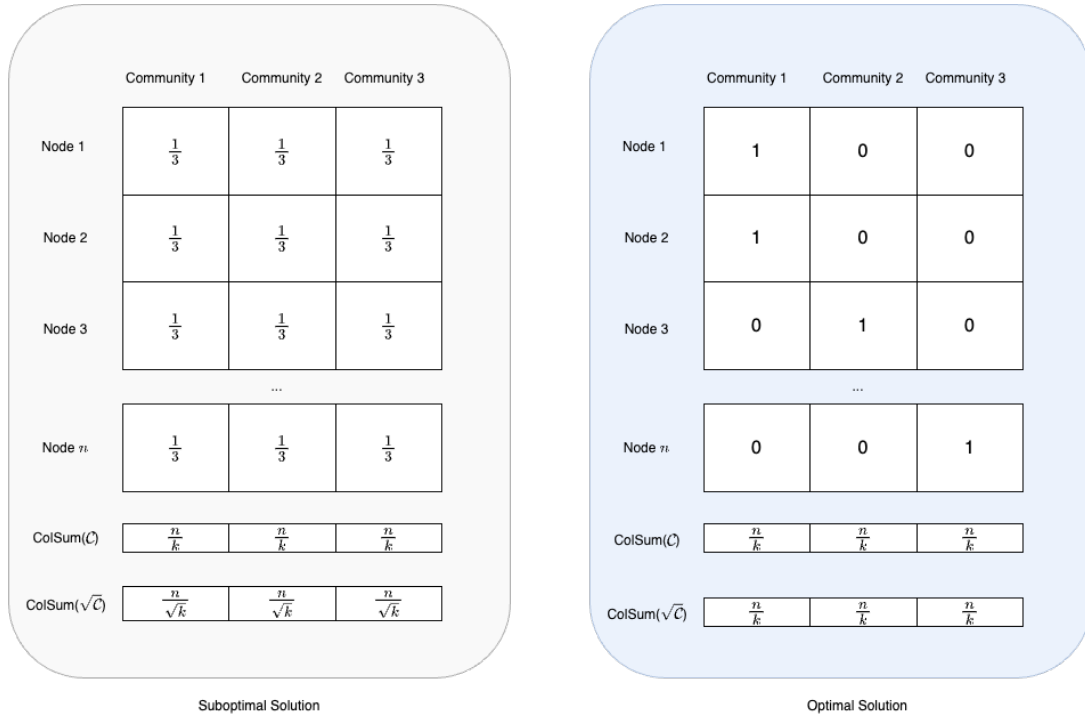


Figure 3. Comparison of square root collapse regularizer and collapse regularizer

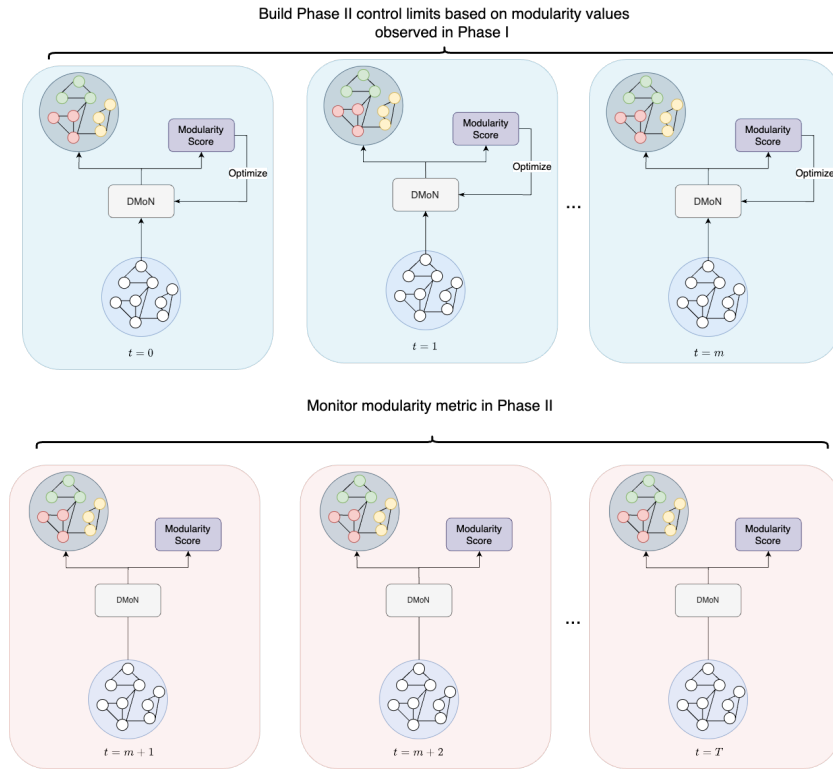


Figure 4. Illustration of model learning during Phase I & monitoring during Phase II

E. Block Visualization for Structural Changes

Figure 5 illustrates how the Λ matrix changes as we alter varying values of $\lambda_{r,r}$ and $\lambda_{r,s}$.

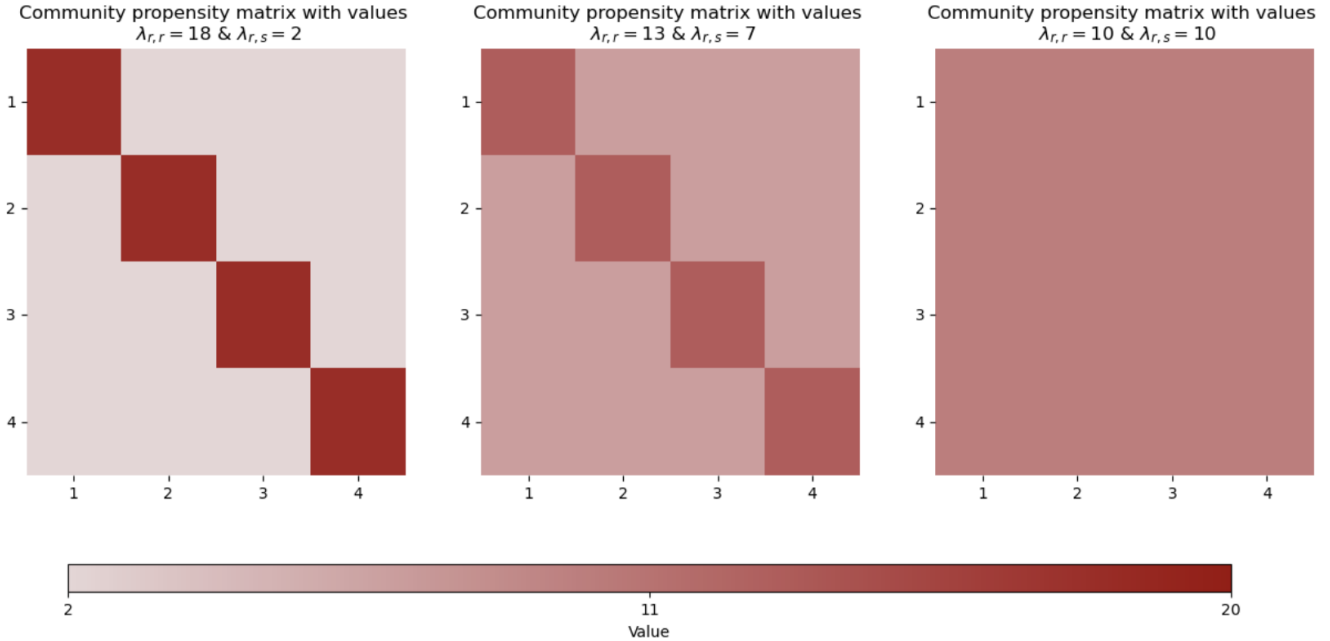


Figure 5. Heatmap representation of varying $\lambda_{r,r}$ and $\lambda_{r,s}$ values

F. No Change

In addition to the explored changes, our method must also be evaluated in scenarios where there is no change to the graph structure or attributes. Unlike the other sections focused on change detection, our objective in this case is different – we do not want to detect changes and hence having a low average percentage over the threshold becomes ideal as it suggests the robustness of our method to type I errors. The performance of our method under these conditions is illustrated in Table 2, where we observe a consistently low average percentage over the threshold, indicating a low false alarm rate. This is slightly larger than the nominal 0.05 value; this small difference is likely due to the monitored modularity scores not precisely following the distributional assumptions made by the EWMA control chart. Figure 6, showcases an example control chart for this scenario. It can be seen there are no anomalous points in this case and the modularity remains stationary throughout both Phase I and II.

α	Average Percentage Over Threshold
0.2	0.07

Table 2. Performance summaries across $N = 100$ replications when no change is injected

G. #Iran Network

G.1. Data

We scrape the #Iran network for a seven-month period, an undirected edge is created when two users reply to each other. We build the text embedding by pooling the tweet embeddings of user tweets and their replies. If a user has multiple tweets in a single snapshot, we create a node embedding by averaging all tweets within that snapshot. We note the number of users may vary between each snapshot with an average of 5254.74 over all the snapshots with a standard deviation of 2,771.93. We collect the data over an eleven-month period from June 3rd 2022, to March 31st 2023. Unfortunately, due to Twitter policies, we cannot make this data public. However, we will open-source all code and simulation data upon acceptance.




Figure 6. EWMA control chart when there is no change injected $\alpha = 0.2$

G.2. Training

We train our DMoN model sequentially on Phase I graphs. We tune the number of communities (k), the learning rate, and dropout by reserving one Phase I snapshot for evaluation. Using the best performing model on the evaluation snapshot, we build the Phase I limits. We use this model throughout Phase II to calculate the modularity of each snapshot.