

A UNIQUE GRAPH SIMILARITY METRIC FOR ANOMALY DETECTION

Nathanael J. K. Brown
Christopher R. Frazier
Katherine R. Cauthen

Sandia National Laboratories
Post Office Box 5800
Albuquerque, NM 87185, USA

Linda K. Nozick

School of Civil and Environmental Engineering
Cornell University
220 Hollister Hall
Ithaca, NY 14853, USA

ABSTRACT

The ability to detect changes between similar graphs can be viewed as an anomaly detection problem. This paper presents a graph similarity metric that is sensitive enough to detect subtle changes in link strength or graph structure without early saturation when dramatic changes must be detected. The algorithm achieves this level of fidelity by combining multiple graph analysis algorithms in an efficient manner.

1 INTRODUCTION

Graphs are used to represent a variety of network phenomena, such as social interactions in humans (Michalski et al. 2012) and animals (Hobson et al. 2013), web searches (Papadimitriou et al. 2010), traffic (Ramdlani et al. 2017), semantic analysis (Drieger 2013), brain function (Mheich et al. 2015), and molecular physics (Hernandez et al. 2016). Often these networks are observed to evolve over time, and therefore it is of interest to detect changes in the graph. In order to detect differences in the graph at different points in time, it is necessary to quantify the amount of similarity between two graphs. Several graph similarity metrics have been proposed in the literature, each of which leverages different network properties.

There has been a great deal of work regarding graph similarity and how to measure it, but relatively little attention has been paid to methods for continuously monitoring these changes for anomalous behavior. In dynamic networks one would expect a certain amount of natural shifting in nodes and edges over time, which may be characterized as noise. However, a substantial change in the network topology would indicate anomalous activity. By quantifying and tracking the similarity between instances of the graph over time, anomaly detection methods could be applied to understand when a network's baseline level of shifting has been punctuated by an event of interest. Our contribution is twofold: 1) We present a new graph similarity metric which *simultaneously* and *directly* incorporates: node existence, edge existence, edge weight, and higher-order structure and 2) we propose using control charts of the similarity metric to detect anomalies in a network over time.

The literature examining and quantifying graph similarity is extensive. Graph isomorphism, in which two graphs with the same number of nodes are connected in an identical way, is one way of determining whether two graphs are similar (Chartrand 1985). Graph isomorphism is a combinatorial optimization problem that is difficult to solve, and many procedures have been proposed for doing so (Corneil and Gotlieb 1970; Pelillo 1999; Ullman 1976). In addition to algorithms that determine whether two graphs are isomorphic there is also a similarity measure based upon data fusion of isomorphic and nonisomorphic subgraphs (Ruiz et al. 2007). Research on common graphs seeks to identify the maximum or minimum common isomorphic subgraph or supergraph, respectively (Fernandez and Valiente 2001; Bunke et al. 2000).

While graph isomorphism determines whether two graphs are identical, a graph edit distance approach allows one to quantify the similarity between graphs at a higher resolution. Graph edit distance is the smallest cost needed to transform one graph to another via edit operations (i.e., addition and deletion of nodes and edges) (Bunke 1999; Messmer and Bunke 1998; Gao et al. 2010; Papadimitriou et al. 2010). However, calculating graph edit distance is computationally expensive, and there is no effective, general method since calculating cost is application-dependent. Cao et al. (2013) propose using the Levenshtein distance (i.e., string edit distance) in combination with a canonical labeling system as an alternative. Related to the idea of graph edit distance, Michalski et al. (2012) characterize the similarity between graphs as a graph differential tuple, consisting of the set of added and removed nodes and edges as well as the set of modified edge weights. They propose several metrics combining information from these three forms of graph edits, and they propose a separate metric based solely upon the difference in edge weights. However, the authors do not combine the graph edits with changes in edge weights for a single, comprehensive metric.

Other researchers have leveraged aggregate measures of graph topology using a statistical framework (Albert and Barabasi 2002; Dill et al. 2002). By defining and quantifying network structural properties such as betweenness and degree distribution, similar graphs can be compared and their structural differences can be described in an interpretable way. Michalski et al. (2014), for example, employ a forgetting mechanism in their calculation of such structural measures to characterize structure in an evolving network. Spectral graph theory has emerged as another option for assessing graph similarity. These methods calculate the distance between the eigen decomposition of the two graphs' matrices (Lagunas et al. 2018; Wilson and Zhu 2008). Koutra et al. (2013) proposed the DeltaCon similarity metric which compares pair-wise node affinities of two graphs, although the metric is heavily influenced by which characterization of node affinity is utilized.

Several attempts have been made to combine approaches. Mheich et al. (2015) propose an algorithm that incorporates graph edit distance, string edit distance, and physical location to compare graphs of brain connectivity. Bopche and Mehtre (2017) fuse graph edit distance and maximum common subgraph into a graph distance metric in an anomaly detection setting for detecting changes on the attack surface of dynamic computer networks. Papadimitriou et al. (2010) also consider the problem of anomaly detection on consecutive graphs in time. They propose sequence and signature similarity measures, adapted from document and vector similarity methods, to detect anomalies on web graphs.

2 FORMULATION

The similarity metric is composed both of simple network metric calculations as well as a small optimization for determining the node matches between clusters.

2.1 Similarity Calculation

Assume we have two similar networks A and B which each have nodes with unique numeric IDs. Nodes are considered to match when their numeric IDs match. Two links are considered to match when they connect the same two nodes with the same IDs. Note that our implementation assumes bidirectional links for simplicity, but the calculation is still valid when used with directed networks. The similarity between the two networks is calculated as the weighted sum of the four components defined below and can range in value from zero to one.

Link Strength Similarity provides a measure of the link strength similarity between networks and is the sum of the absolute strength differences normalized by the sum of the strengths across both networks. If there is a link in one network that does not exist in the other, then the “missing” link has a strength value of zero. If all link strengths match, this component has value zero. If there are no matching links, then this component has value one.

$$L_S = \frac{\sum_{i=1}^N |S_i^A - S_i^B|}{\sum_{i=1}^N [S_i^A + S_i^B]} \quad (1)$$

- S_i^A – strength of link i from network A
- S_i^B – strength of link i from network B
- N – maximum number of links between both networks

Matching Link Ratio is the ratio of the number of matching links between networks to the total number of unique links across both networks. Two links are a match if they connect two nodes with the same IDs. If all links match, this component has value one. If no links match, this component has value zero.

$$L_M = \frac{\text{Number of matching links between networks}}{\text{Total number of unique links in both networks}} \quad (2)$$

Matching Node Ratio is the ratio of the number of matching nodes (nodes with equal-valued IDs) between networks to the total number of unique nodes in both networks. If all nodes match, this component has value one. If none of the nodes match, this component has value zero.

$$N_M = \frac{\text{Number of matching nodes between networks}}{\text{Total number of unique nodes in both networks}} \quad (3)$$

Matching Cluster Ratio is the ratio of the number of matching node-cluster assignments to the total number of unique nodes across both networks. It is assumed that there are an equal number of clusters in each network. If all nodes have matching cluster assignments, this component has value one. If none of the nodes have matching cluster assignments, this component has value zero.

$$N_C = \frac{\text{Number of matching node cluster assignment between networks}}{\text{Total number of unique nodes in both networks}} \quad (4)$$

Total Similarity, S , is the weighted sum of these components:

$$S = 0.25 * (1 - L_S) + 0.25 * L_M + 0.25 * N_M + 0.25 * N_C \quad (5)$$

Note that the components could be weighted differently based on the application so long as the similarity still has a range from 0 to 1. For example, if the nodes in both network are guaranteed to match (resulting in a constant value of 1 for N_M), the coefficient of N_M could be changed to zero and the coefficients for all other components could be changed to 1/3.

2.2 Node Cluster Mapping Optimization

In order to calculate the matching cluster ratio component of the similarity, the best cluster mapping between the two networks must be determined. To find the best cluster mapping for a large numbers of clusters, it becomes too expensive to examine all $C!$ permutations, where C is the number of clusters in each network (assumed to be the same). The goal is to find the mapping with the greatest number of overlapping nodes across all cluster-cluster assignments between networks A and B . To find this mapping, a population-based metaheuristic is employed. A selection of initial solutions is generated by randomly

assigning clusters in network A to clusters in network B and calculating the node overlap. The set of random solutions is augmented with a set of greedy solutions which attempt to assign the highest overlapping clusters to each other in a systematic fashion. The overlap counts for each cluster in network A to the highest overlapping cluster in network B can be summed to give the theoretical maximum overlap count for the best mapping solution. A population-based local search technique is used to gradually improve the collection of initial solutions generated.

2.3 Node Clustering Optimization

To compare the clustering of nodes in each network, we must first determine the node-cluster assignments. The standard Louvain algorithm (Blondel et. al. 2008) determines the optimal clustering for a collection of linked nodes by maximizing the modularity. The final number of clusters is determined by the algorithm and cannot be a fixed value. Our modified version of the algorithm allows the number of clusters to have a fixed upper bound and is a population-based metaheuristic using modularity as the objective value. The optimization is broken down into two main pieces: initialization followed by solution tuning via local search. The initialization procedure essentially follows the first phase of the Louvain Method by separating all nodes into separate communities then finding the best community for each node by calculating the relative gain in modularity. If the number of remaining communities is still greater than the maximum allowed number of clusters, the communities are optimally combined until there are no more than the allowed number remaining. The resulting solution is augmented by creating a collection of solutions with random node-cluster assignments. All solutions are then tuned by executing a local search of moving single nodes or fully-connected groups of nodes to new communities in an effort to maximize the modularity. The solution with the highest modularity is then retained as the final solution.

2.4 Experiments

In order to determine the efficacy of the algorithm for detecting changes to a network, a simple network consisting of 25 nodes in five distinct clusters is constructed as shown in figure 1.

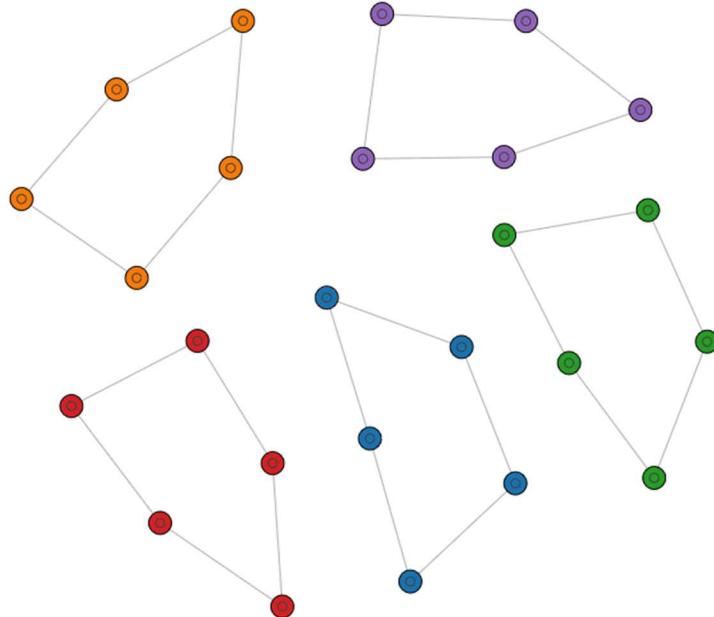


Figure 1: Simple 25-node, 5-cluster network

The nodes are minimally connected within each cluster so that overall network modularity is 0.8 with no cross-cluster links. Two experiments are conducted: one to demonstrate the efficacy of the similarity metric when comparing a static network to a randomly modified version of the same network and another which studies the effects on similarity of a gradually changing social network.

2.4.1 Random Network Modification Experiment

In this experiment, a series of modifications is applied to a second, initially identical network (identified as “network 2”) to slowly modify it in specific ways to determine the resultant similarity value. Over a period of 275 iterations, the following modifications are applied to alter the network:

- At each iteration, a random increment or decrement operation is selected at random (50% chance of each).
 - If the action is to decrement, the link strength is decremented by a random amount between one and 15. If the strength goes below zero, the link is removed.
 - If the action is to increment, two nodes are selected at random. If a link currently exists between them, the strength is incremented by a random amount between one and 15. If a link does not exist, a new one is created with strength between one and 15.
- Every 15 iterations a node is dropped from the network, meaning that all connected links are removed. This missing node is not included in any of the node-matching or clustering ratio calculations for the similarity metric so this action will eventually force network 2 to be completely different from network 1 (similarity is zero since there are no common nodes or links).

Since each random change may result in either an increase or a decrease in similarity, the sampled changes are sorted by decreasing similarity value in the figure 2 plots below. These plots demonstrate that the metric can span its full range depending on the changes made to the network.

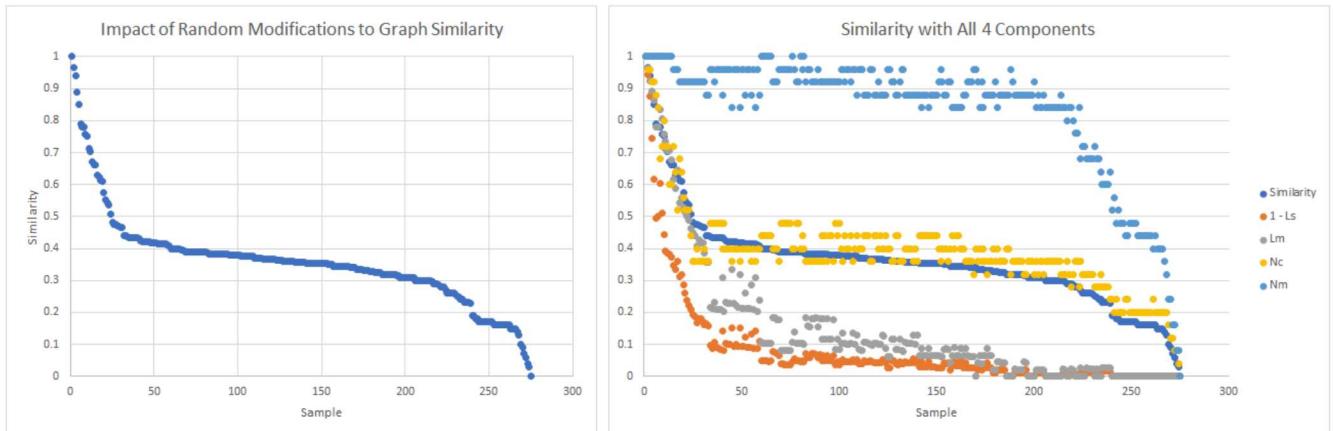


Figure 2: Impact of Random Network Modifications to Similarity Metric and Associated Components

During the experiment, intermediate networks at similarity values close to 0.75 and 0.50 are captured and examined to see how they differ from the original network. Table 1 summarizes the differences in the node clustering.

The structural clustering difference can also be seen in figure 3 where the interior of each node has the original cluster color and the exterior of each node has the new clustering assignment. The width of each link is proportional to its strength and the size of each node is proportional to the strength of the links connected to it.

Table 1: Impact of Random Network Modifications to Network 2 Node Clustering

Cluster	Seed Network Nodes	0.76 Network Nodes	0.51 Network Nodes
c1	1-5	10, 23	2, 6
c2	6-10	6-9	5, 8, 9, 16
c3	11-15	1-4, 11-15	1, 12, 13
c4	16-20	5, 16-20	4, 7, 10, 11, 14, 15, 17-20, 22, 23
c5	21-25	21, 22, 24, 25	21, 25

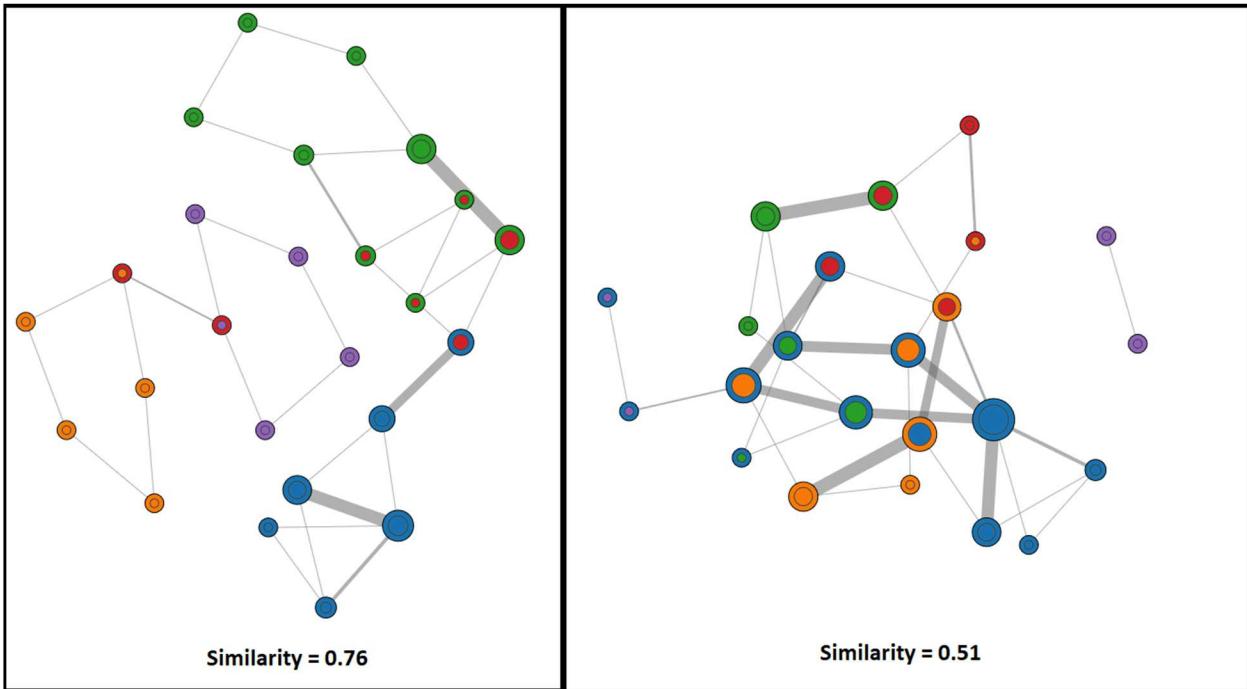


Figure 3: Network 2 at similarity values 0.76 and 0.51

2.4.2 Gradual Change of a Social Network

For this experiment, the same node clusters shown in figure 1 are preserved as the underlying “ground truth” clustering and each node is assumed to be a person in a social network. From this seed network, all possible in-cluster and cross-cluster links are defined. A series of networks are then generated where the link structure is maintained but the level of weekly communication is modeled by a Poisson process with mean two. The sparsest matrix has 60% in-cluster communication with no cross-cluster communication. The most dense network has 100% of the in-cluster links with 40% of the cross-cluster links. Note that all links (in- and cross-cluster) are a progressively larger superset of the 60-00 network links as would be (somewhat) expected in a gradually changing network (i.e., new relationships are formed but old relationships are not lost). For example, the 70-30 network contains 70% of the possible in-cluster links and 30% of the possible cross-cluster links and holds all of the same links as the 60-20 network in addition

to several new links. The 70-30 in-cluster links include all 25 of the 60-20 network in-cluster links as well as 5 new links and also include all 50 of the 60-20 cross-cluster links with 25 additional links.

Table 2 summarizes the number of links for each percentage combination based on a total of 300 possible links where 50 are in-cluster and 250 are cross-cluster. 25 different networks are generated with link percentages ranging from 60-0 (in-cluster/cross-cluster percentage) to 100-40.

Table 2: Number of links for each permutation of in-cluster and cross-cluster links

In-Cluster %	In-Cluster Link Count	Cross-Cluster %	Cross-Cluster Link Count
100	50	40	100
90	45	30	75
80	40	20	50
70	35	10	25
60	30	0	0

In order to do a comparison to an “average”, we chose the 70-10 permutation as a representative network and set the strength of each link equal to the average of 100 draws from the Poisson process described above. For each of the 25 network permutations, a series of 100 random instances is generated where the link structure is preserved but the link strength is generated from the same Poisson process. If the strength of a link is zero for a draw, it is considered to be nonexistent with respect to the similarity calculation. Similarly, if a node has no associated non-zero links it is also considered to be nonexistent with respect to the metric calculation. The similarity metric is calculated for each randomly generated instance to create a virtual time series of length 2500 across all instances, grouped by network permutation.

Since the overall goal is to create a control chart which detects when the network has significantly changed from the reference point, we performed a distribution fit on the 70-10 network random draws. We compared the fits of three possible distributions: truncated normal, beta, and a mixture of normals. The best fit was selected using a voting procedure based on three model fit criteria: AIC (Akaike Information Criterion), BIC (Bayesian Information Criterion), and the Kolmogorov-Smirnov test. For the 70-10 similarity dataset, the best fit is a mixture of normal with means 0.79372 and 0.85084; standard deviations 0.02444 and 0.016695; and lambdas (weights) of 0.84438 and 0.15562. Figure 4 is the smoothed histogram of similarity values (black dotted line) overlaid with the various distribution fits.

Once the underlying distribution is selected, the control chart limits can be established using the techniques described in Montgomery (2009), namely:

- Pick a “running length”, the average number of samples from the in-control distribution before a sample is observed that appears to be out-of-control. In our case we selected 75.
- Calculate the probability limits for the out-of-control limits:
 - $\alpha = 1/(\text{running length}) = 1/75$
 - probability limits = $(\alpha/2, 1-\alpha/2) = (1/150, 149/150)$
- Using a numerical approximation to the inverse mixture of normals distribution, find the similarity values corresponding to the probability limits and set as the control chart limits.

Figure 5 illustrates the application of these control limits to the virtual time series. The similarity values from the 70-10 series is nearly fully encapsulated by the control limits whereas even the small network change to 70-20 results in the bulk of the samples being outside of the control limits. Also of note is that the XX-0 networks go down in similarity because they do not have any of the cross-cluster links present in the 70-10 network structure.

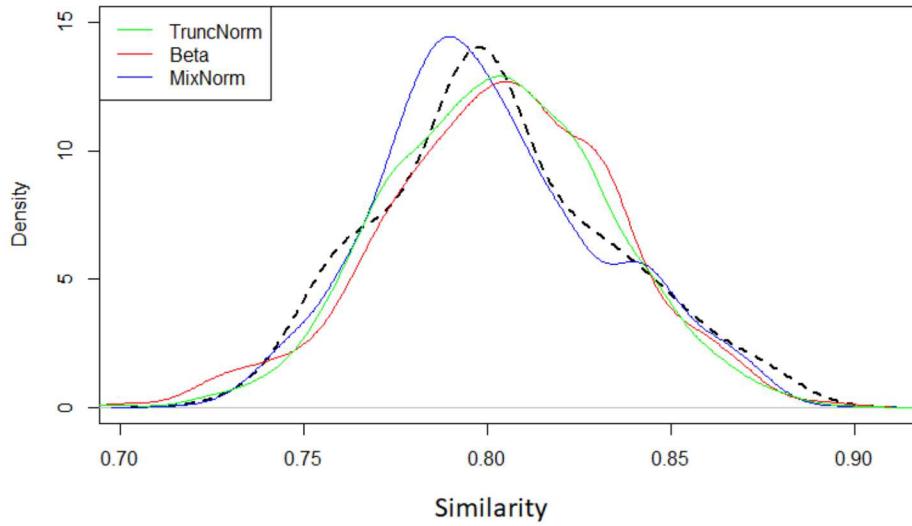


Figure 4: Selecting the best distribution fit for the 70-10 network similarity values (black dotted line).

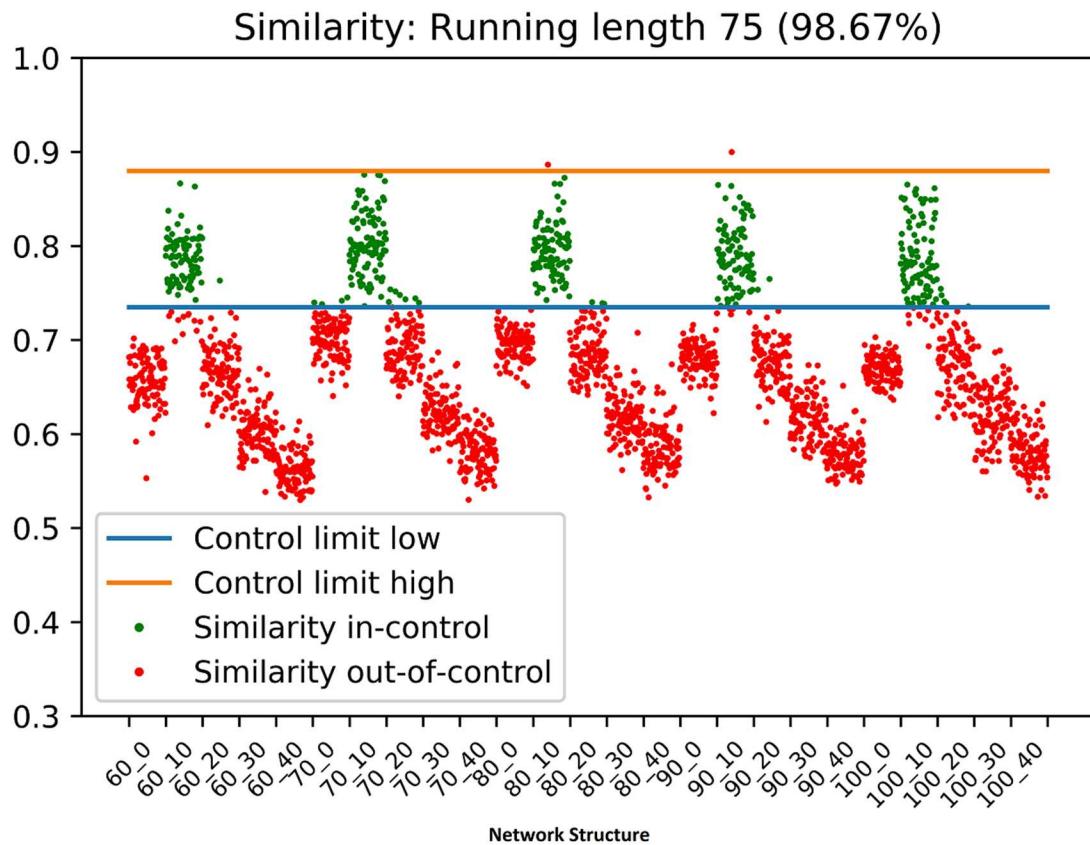


Figure 5: Application of control chart limits to the full similarity of the simulated network data.

When comparing our similarity metric to that proposed by Michalski et al. (2012) the main difference is our inclusion of the matching cluster ratio. This added information is structural and not readily apparent by doing basic comparisons between links and nodes. To illustrate the value of this addition, we pulled out the matching cluster ratio contribution to the similarity metric and renormalized so that the new metric has value ranging from zero to one (this is equivalent to setting the weights on all components in equation 5 to be 1/3 except for N_m which is given a weight of zero). We found that a beta distribution with shape 1 parameter 188.37 and shape 2 parameter 31.14 had the best fit, and used this to create control chart limits for this modified similarity metric. The results of applying the control chart are shown in figure 6.

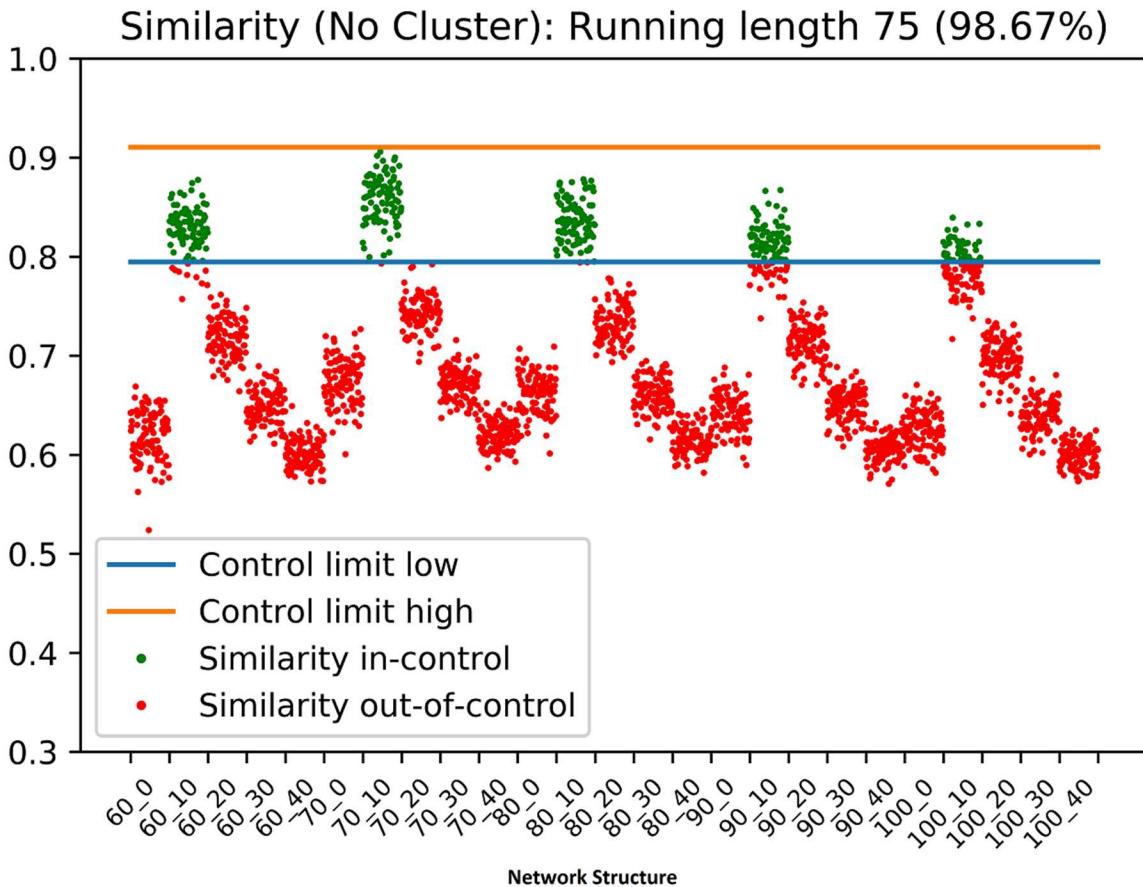


Figure 6: Application of control chart limits to modified similarity metric with no clustering component.

A few items to note when comparing full similarity metric (left plot) to similarity with no clustering information (right plot):

- The XX_0 networks in figure 6 have a much greater deviation from the XX_10 network similarity values as compared to figure 5 even though their clustering is quite similar. The histogram for 90_0 in figure 7 illustrates how that network clustering structure is actually quite similar to the 70_10 with a typical cluster overlap of 80%.
- The 90_10 and 100_10 networks in figure 6 have a much larger percentage of out-of-control samples than their counterparts in figure 5, which shows the effect of eliminating the consideration

of clustering from the similarity metric. Though these networks exhibit similar clustering to the 70_10 network, not including clustering causes them to appear out-of-control more often than one would intuitively expect.

- For the 60_40 network, it is expected that the similarity should be fairly low, as shown in figure 6, given that the cluster overlap is typically in the 40 – 45% range as shown in the figure 7 histogram. The fact that the similarity is higher in figure 5 is also expected, since it does not include the cluster overlap metric, which demonstrates why it's important to include that component.

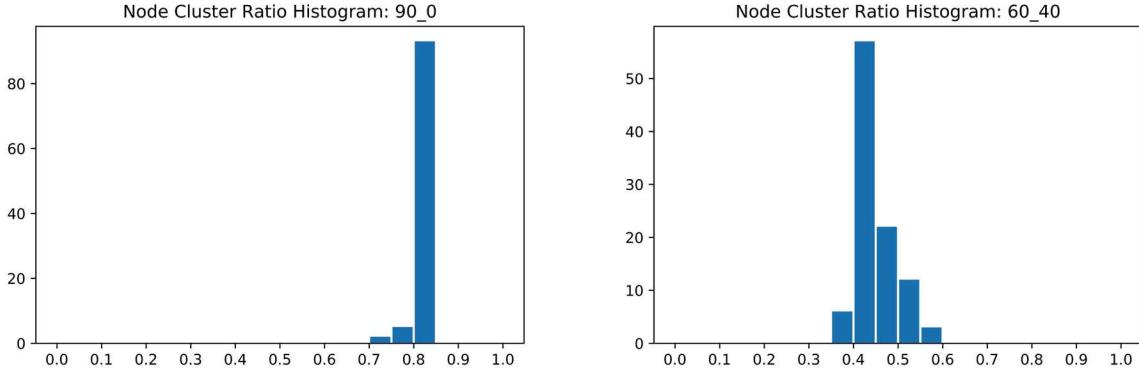


Figure 7: Histograms of the node cluster ratio component for the 90-0 and 60-40 network structures.

CONCLUSION

Having a capability that can detect significant changes within a network without being overly reactive to smaller changes is a difficult balance to strike, especially with a single metric. Our graph similarity metric appears to have achieved this goal by aggregating several graph analysis techniques in a unique manner to produce a value that is easy to interpret and decompose if necessary. By using our metric in conjunction with a control chart we demonstrated that graphs which differ in small ways will be recognized as being similar but that significant (though not hugely dramatic) changes will still be detected. This temporal anomaly detection could be useful in determining if organized crime networks have changed in some meaningful way. Though our case study focused on social network analysis, the metric could be equally valuable to areas such as chemical graph theory and semantic analysis where the clustering of similar nodes has as much or more meaning than the simple presence or absence of particular nodes and arcs.

One extension to this model would be to explore how the component weights could be optimally adjusted for specific applications. For example, if changes in the clustering structure of the network are more important than the presence or absence of weak links, the matching cluster ratio could have a higher relative weighting with respect to the other components.

ACKNOWLEDGMENTS

Special thanks to Drew Levin of Sandia National Laboratories for acting as a sounding board when designing the experiments and with his help creating the network graphics. This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. SAND NO. 2019-xxxx J.

REFERENCES

Albert, R. and A.L. Barabasi. 2002. "Statistical Mechanics of Complex Networks". *Reviews of Modern Physics* 74(1):47-97.

Blondel, V. D., J.-L. Guillaume, R. Lambiotte, E. Lefebvre. 2008. "Fast Unfolding of Communities in Large Networks". In *Journal of Statistical Mechanics: Theory and Experiment* (10), P10008.

Bopche, G.S. and B.M. Mehtre. 2017. "Graph Similarity Metrics for Assessing Temporal Changes in Attack Surface of Dynamic Networks". *Computers and Security* 64(C):16-43.

Bunke, H. 1999. "Error Correcting Graph Matching: On the Influence of the Underlying Cost Function". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21(9):917-922.

Bunke, H., X. Jiang, and A. Kandel. 2000. "On the Minimum Common Supergraph of Two Graphs". *Computing* 65(1):13-25.

Cao, B., Y. Li, and J. Yin. 2013. "Measuring Similarity Between Graphs Based on the Levenshtein Distance". *Applied Mathematics and Information Sciences* 7(1L):169-175.

Chartrand, G. 1985. *Introductory Graph Theory*. New York: Dover.

Corneil, D.G. and C.C. Gotlieb. 1970. "An Efficient Algorithm for Graph Isomorphism". *Journal of the Association for Computing Machinery* 17(1):51-64.

Dill, S., R. Kumar, K.S. McCurley, S. Rajagopalan, D. Sivakumar, and A. Tomkins. 2002. "Self-Similarity in the Web". *ACM Transactions on Internet Technology* 2(3):205-223.

Drieger, P. 2013. "Semantic Network Analysis as a Method for Visual Text Analytics". *Procedia – Social and Behavioral Sciences* 79:4-17.

Fernandez, M.L. and G. Valiente. 2001. "A Graph Distance Metric Combining Maximum Common Subgraph and Minimum Common Supergraph". *Pattern Recognition Letters* 22(6-7):753-758.

Gao, X., B. Xiao, D. Tao, and X. Li. 2010. "A Survey of Graph Edit Distance". *Pattern Analysis and Applications* 13:113-129.

Hernandez, M., A. Zaribafyan, M. Aramon, and M. Naghibi. 2016. "Quantum Approaches to Graph Similarity". 1QB Information Technologies.

Hobson, E. A., M.L. Avery, and T.F. Wright. 2013. "An Analytical Framework for Quantifying and Testing Patterns of Temporal Dynamics in Social Networks". *Animal Behavior* 85:83-96.

Koutra, D., J.T. Vogelstein, and C. Faloutsos. 2013. "DeltaCon: A Principled Mass-Graph Similarity Function". *SIAM International Conference on Data Mining*, May, Austin, Texas.

Lagunas, E., A.G. Marques, S. Chatzinotas, and B. Ottersten. 2018. "Graph Similarity Based on Graph Fourier Distances". *26th European Signal Processing Conference*, September, Rome, Italy.

Messmer, B.T. and H. Bunke. 1998. "A New Algorithm for Error-Tolerant Subgraph Isomorphism Detection". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(5):493-504.

Mheich, A., M. Hassan, V. Gripon, M. Khalil, C. Berrou, O. Dufor, and F. Wendling. 2015. "A Novel Algorithm for Measuring Graph Similarity: Application to Brain Networks". *7th Annual International IEEE EMBS Conference on Neural Engineering*, April 22-24, Montpellier, France, 1068-1071.

Michalski, R., P. Brodka, P. Kazienko, and K. Juszczyszyn. 2012. "Quantifying Social Network Dynamics". In *Proceedings of the 4th International Conference on Computational Aspects of Social Networks*, November 21-23, Sao Carlos, Brazil, 69-74.

Michalski, R., T. Kajdanowicz, P. Brodka, and P. Kazienko. 2014. "Seed Selection for Spread of Influence in Social Networks: Temporal vs. Static Approach". *New Generation Computing* 32:213-235.

Montgomery, D.C. 2009. *Introduction to Statistical Quality Control*. 6th ed. Hoboken, New Jersey: John Wiley and Sons, Inc.

Papadimitriou, P., A. Dasdan, and H. Garcia-Molina. 2010. "Web Graph Similarity for Anomaly Detection". *Journal of Internet Server Applications* 1:19-30.

Pelillo, M. 1999. "Replicator Equations, Maximal Cliques, and Graph Isomorphism". *Neural Computation* 11(8):1933-1955.

Ramdlani, A., G.A. Putri Saptawati, Y. Asnar. 2017. "Graph Analysis on ATCS Data in Road Network for Congestion Detection". *International Conference on Data and Software Engineering*, November 1-7, Palembang, Indonesia.

Ruiz, I.L., M. U. Cuadrado, and M.A. Gomez-Nieto. 2007. "New Graph Similarity Measurements Based on Isomorphic and Nonisomorphic Data Fusion and their Use in the Prediction of the Pharmacological Behavior of Drugs". *International Journal of Pharmacological and Pharmaceutical Sciences* 1(2):47-51.

Ullman, J.R. 1976. "An Algorithm for Subgraph Isomorphism". *Journal of the Association for Computing Machinery* 23(1):31-42.

Wilson, R.C. and P. Zhu. 2008. "A Study of Graph Spectra for Comparing Graphs and Trees". *Pattern Recognition* 41:2833-2841.

Brown, Frazier, Cauthen, and Nozick

AUTHOR BIOGRAPHIES

NATHANAEL BROWN is a Principal Member of the Technical Staff at Sandia National Laboratories in Albuquerque, NM. He has over 25 years of experience in the design and development of advanced software algorithms ranging from handwriting recognition to multi-objective optimization using heuristic techniques. He has worked extensively in the area of investment planning optimization as applied to infrastructure reinforcement and resiliency against multiple threats. His email address is njbrown@sandia.gov.

CHRISTOPHER FRAZIER is a Senior Member of the Technical Staff at Sandia National Laboratories in Albuquerque, NM. He has extensive experience in statistical modeling, data analysis, and software engineering. His email address is crfrazi@sandia.gov.

KATHERINE CAUTHEN is a Member of the Technical Staff at Sandia National Laboratories in Albuquerque, NM. She has applied her background in statistics and data science to anomaly detection and modeling and simulation, primarily in the domains of public health and cyber security. Her email address is kcauthe@sandia.gov.

LINDA NOZICK is Professor and Director of Civil and Environmental Engineering at Cornell University. She has been the recipient of several awards including a CAREER award from the National Science Foundation and a Presidential Early Career Award for Scientists and Engineers from President Clinton for "the development of innovative solutions to problems associated with the transportation of hazardous waste." She has authored over 60 peer-reviewed publications, many focused on transportation, the movement of hazardous materials and the modeling of critical infrastructure systems. Her email address is lkn3@cornell.edu.