Back to the Graphs: A Collection of Datasets and Quality Criteria for Temporal Networks Layout and Visualization

Velitchko Filipov* TU Wien Vienna, Austria

Alessio Arleo* TU Wien Vienna, Austria

Tatiana von Landesberger† University of Cologne Cologne, Germany

Daniel Archambault‡ Newcastle University Newcastle, UK

ABSTRACT

Dynamic network visualization is a rapidly evolving field and computing a node-link layout for these graphs is one of the most studied problems in this discipline. Benchmark datasets and quality criteria are lacking in order to enable comparative evaluations of the layouts produced by different dynamic network visualization techniques. This paper proposes a collection of both discrete time and event-based dynamic quality criteria for evaluating dynamic graph layouts. Furthermore, we present and discuss datasets, generated from real data, which we include in a benchmark set for future researchers to use in their evaluations.

Index Terms: Human-centered computing—Visualization—Visualization techniques—Graphs;

1 INTRODUCTION

Node-link representations of graphs on a plane are the most common way to visualize graphs [16]. Typically, these graphs $G = (V, E)$ encode the vertices $v \in V$ as circles linked together by edges $e \in E$, depicted as lines. Dynamic network visualization develops approaches to depict the temporal dynamics of nodes and edges. Broadly speaking, time can be represented in two ways [2, 6]: mapping time to space (i.e., small multiples) or time to time (i.e., animation).

As the dynamic network visualization discipline evolves, different time abstractions are introduced. In graph drawing, time has been modeled as snapshots of the graph at equally spaced, discrete moments in time (the *timeslices*) [6]. Recent techniques model the temporal dynamics of nodes and edges (e.g., addition/deletion) as *events* with time coordinates as real numbers and duration, hence the name *temporal* (or *event-based*) networks [10]. Modeling the graph dynamics in this fashion better approximates real-world phenomena, as data from microblogging messages, network communications, and contact tracing do not naturally fall neatly into timeslices. This new perspective motivates this presentation of quality criteria and benchmark datasets for designing and evaluating visualizations, an otherwise largely under-investigated topic.

Our Contribution. We compile a list of quality criteria for event-based layouts; second, we survey and include real-world both timesliced and temporal graphs as benchmark datasets for further research in the field. Our objective is to provide the community with a starting set of resources to support transparent, reproducible, and comparative evaluations of the techniques developed in this design space. We also advocate for more research in this direction given the advantages of temporal networks in dynamic network visualization [4, 17].

*e-mail: {velitchko.filipov—alessio.arleo@tuwien.ac.at}

†e-mail: tlandesb@uni-koeln.de

‡e-mail: daniel.archambault@newcastle.ac.uk

2 RELATED WORK

Research on static graph quality metrics focuses on understanding which factors influence layout readability, e.g., through user studies (e.g., Purchase et al. [16]). From this work, it became clear that it was not possible to optimize all quality criteria simultaneously, thus requiring tradeoffs. Brandes et al. [7] conducted an experimental study exploring the tradeoff between quality and stability in the context of timesliced graph drawing. Beck et al. [5] identified a set of aesthetic criteria for network animations. Among these, it also mentions the user mental map preservation, which has been investigated in several studies [3]. As all related work focuses on timesliced dynamic graphs and considering recent developments in temporal network visualization, this highlights gaps in literature and motivates research on benchmark datasets and quality criteria.

3 TEMPORAL NETWORKS: QUALITY CRITERIA & DATASETS

Quality Criteria. We present quality criteria for drawing temporal networks that do not depend on the temporal encoding but are rather indicators of the readability of the embedding in a $2D+t$ space-time cube. These methods have been used in prior work [4, 17].

○ *Node Movement* tracks the average distance traveled by the nodes during the evolution of the dynamic graph [17]. This metric is strongly related to layout *stability* which is required to support the cognitive map [3] of the viewer. This metric has been used in both timesliced [7] and event-based [4, 17] evaluations. Lower movement reflects stable graph drawings but some networks require, by their nature, higher movement – for example when there is a fundamental change in a node's local neighborhood in a short period of time.

○ *Crowding* counts the number of times a node passes close to another node during the evolution, meaning that node identities could be confused. It has been used in previous studies to evaluate drawings of temporal networks [4, 17]. Low crowding means that nodes are easily distinguishable and high crowding indicates that it could be hard to follow the node through the evolution of the dynamic network. High crowding may also result from low movement drawings with the nodes very close to each other.

○ *Stress* measures the difference between the Euclidean and graph theoretic distance in the drawing and has been extensively used to evaluate dynamic graph drawings [4, 7, 17]. Low stress drawings place nodes at Euclidean distances proportional to their distance in terms of the number of hops on the graph. In this context, we identify two methods for computing stress: (i) *Timesliced* Stress, computed and averaged on a per timeslice basis, and (ii) *Continuous* Stress, calculated and averaged in continuous time, using the exact node and edge appearances. In literature [4, 17], they are called *StressOn* and *StressOff* respectively.

Timesliced & Temporal Datasets. We present a collection of timesliced and temporal networks that can be used for evaluating visualizations (see Table 1). The datasets have a variety of event distributions, matching the characteristics of several real-world scenarios (e.g., alternating peaks and valleys, periods of stability, high volatility). The datasets and code are openly available on GitHub [1].

○ *Timesliced Graphs.* Van Debunt: encodes the relationship between 32 freshman at seven different time points [18]. Timeslices and edges

Dataset Name	Van Debunt	Newcomb	InfoVis	Rugby	Pride	MOOC	RM	MSG	RAMP
Reference	[18]	[14]	[11]	[17]	[9]	[12]	[8]	[15]	[13]
Dataset Type	Timesliced	Timesliced	Timesliced	Event-based	Event-based	Event-based	Event-based	Event-based	Event-based
# of Events	0.1k	0.6k	2.8k	3.1k	4.0k	15k	28k	15k	0.8k
Trends	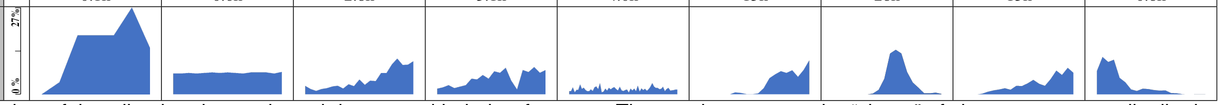								

Table 1: Overview of timesliced and event-based datasets with their references. The trends represent the "shape" of changes, or event distribution (i.e., % of total number of events over time) for each graph dataset. Each graph has a different temporal extent but the same Y-axis.

are selected as in the paper by Brandes and Mader [7]. Newcomb: contains the sociometric preference of 17 members of a fraternity in the University of Michigan in the fall of 1956 [14]. Timeslices are selected as in previous work [7]. InfoVis: is a co-authorship network for papers published in the InfoVis conference from 1995 to 2015 [11]. Authors of a paper are connected in a clique at the time of publication. This is not a cumulative network as authors can appear, disappear and appear again. The dataset has 21 timeslices (one per year). Although timesliced, previous studies have shown that event-based drawings of this data [4, 17] outperform timesliced based methods due to the drastic changes between timeslices.
○ *Temporal Networks.* Rugby is derived from over 3000 tweets involving teams in the 'Guinness Pro12' rugby competition. The tweets were posted between 1 September 2014 and 23 October 2015. Each edge has a time of publication with a precision down to the second. Pride: lists the dialogues between characters in the novel 'Pride and Prejudice' in order [9]. The book has 61 chapters with 4000 interactions between characters. When the algorithm required timeslices as input, we divided this data into 61 of them (one for each chapter). The graph can naturally be considered a temporal network by simply considering the character interactions in order. MOOC: represents the actions taken by users on a popular massive open online class platform [12]. The nodes represent users and course activities (targets), and the edges represent the actions by users. We pick and elaborate the first 15 thousand events. RM: is the data coming from The Reality Mining study [8]. It followed 94 participants using mobile phones pre-installed with several pieces of software that recorded their actions. We only consider voice calls and consider the first 28 thousand events. MSG: is comprised of private messages sent on an online social network at the University of California, Irvine [15]. Each message is represented as a temporal edge. We consider the first 15 thousand events. RAMP: consists of the simulated spread of COVID-19. It was developed by the Scottish COVID-19 Response Consortium contact tracing model [13]. Nodes are infected individuals and edges occur on the day when the illness was transmitted. We consider the first 800 events.

4 Conclusion

This paper presents a collection of datasets and quality criteria for supporting further research on the visualization and layout of temporal networks. Event-based layout showed higher quality drawings on temporal graphs, due to all temporal information being used, and not (partially) lost in discretization. We believe that these resources can benefit both researchers and practitioners pursuing this research direction, by enabling them to develop, compare, and evaluate techniques. Future work can extend this collection of temporal network metrics and datasets by incorporating more graphs and expanding the quality criteria to cover additional aspects of network visualization, such as user interaction, engagement, and perception.

Acknowledgements

For the purpose of open access, the author has applied a Creative Commons Attribution (CC-BY) license to any Author Accepted Manuscript version arising from this submission. This work is supported by the Austrian Science Fund (FWF) grant *ArtVis* [P35767].

References

[1] Multidynnos github repository. `https://github.com/EngAAlex/MultiDynNos/`. Online, accessed 2023-08-03.
[2] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of Time-Oriented Data*. Second edition ed., 2023. doi: Forthcoming
[3] D. Archambault and H. C. Purchase. Can animation support the visualisation of dynamic graphs? *Information Sciences*, 330:495–509, 2016. doi: 10.1016/j.ins.2015.04.017
[4] A. Arleo, S. Miksch, and D. Archambault. Event-based Dynamic Graph Drawing without the Agonizing Pain. *Computer Graphics Forum*, 41(6):226–244, 2022. doi: 10.1111/cgf.14615
[5] F. Beck, M. Burch, and S. Diehl. Towards an Aesthetic Dimensions Framework for Dynamic Graph Visualisations . In *2009 13th International Conference Information Visualisation*, pp. 592–597, 2009. doi: 10.1109/IV.2009.42
[6] F. Beck, M. Burch, S. Diehl, and D. Weiskopf. A Taxonomy and Survey of Dynamic Graph Visualization. *Computer Graphics Forum*, 36(1):133–159, 2017. doi: 10.1111/cgf.12791
[7] U. Brandes and M. Mader. A quantitative comparison of stress-minimization approaches for offline dynamic graph drawing . In *International Symposium on Graph Drawing*, pp. 99–110. Springer, 2011.
[8] N. Eagle and A. (Sandy) Pentland. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4):255–268, 2006. doi: 10.1007/s00779-005-0046-3
[9] S. Grayson, K. Wade, G. Meaney, and D. Greene. The Sense and Sensibility of Different Sliding Windows in Constructing Co-occurrence Networks from Literature . In *Computational History and Data-Driven Humanities*, pp. 65–77, 2016.
[10] P. Holme and J. Saramäki. Temporal networks. *Physics Reports*, 519(3):97–125, 2012.
[11] P. Isenberg, F. Heimerl, S. Koch, T. Isenberg, P. Xu, C. Stolper, M. Sedlmair, J. Chen, T. Möller, and J. Stasko. vispubdata.org: A Metadata Collection about IEEE Visualization (VIS) Publications . *IEEE Transactions on Visualization and Computer Graphics*, 23, 2017. doi: 10.1109/TVCG.2016.2615308
[12] S. Kumar, X. Zhang, and J. Leskovec. Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, p. 1269–1278, 2019. doi: 10.1145/3292500.3330895
[13] S. Mohr, L. Matthews, S. Brett, V. Mano, J. Nonweiler, R. Takahashi, and E. Townsend. Contact-tracing-model. *See https://github.com/ScottishCovidResponse/Contact-Tracing-Model*, 2020.
[14] T. M. Newcomb. *The acquaintance process as a prototype of human interaction.*, pp. 259–261. 1961. doi: 10.1037/13156-015
[15] P. Panzarasa, T. Opsahl, and K. M. Carley. Patterns and dynamics of users' behavior and interaction: Network analysis of an online community . *Journal of the American Society for Information Science and Technology*, 60(5):911–932, 2009. doi: 10.1002/asi.21015
[16] H. C. Purchase. Metrics for Graph Drawing Aesthetics. *Journal of Visual Languages & Computing*, 13(5):501–516, 2002. doi: 10.1006/jvlc.2002.0232
[17] P. Simonetto, D. Archambault, and S. Kobourov. Event-Based Dynamic Graph Visualisation. *IEEE Transactions on Visualization and Computer Graphics*, 26(7):2373–2386, 2020.
[18] G. G. Van De Bunt, M. A. Van Duijn, and T. A. Snijders. Friendship Networks Through Time: An Actor-Oriented Dynamic Statistical Network Model . *Computational & Mathematical Organization Theory*, 5(2):167–192, 1999. doi: 10.1023/A:1009683123448