Leveraging Time Dependency in Graphs

Arinbjörn Kolbeinsson
Imperial College London
London, UK
ak711
@imperial.ac.uk

Naman Shukla
Deepair Solutions
Dallas, TX
naman
@deepair.io

Akhil Gupta
UIUC
Champaign, IL
akhilg3
@illinois.edu

Lavanya Marla
UIUC
Champaign, IL
lavanyam
@illinois.edu

Background

Most real-world problems are inherently time-variant, and yet temporal information is usually ignored by graph representation learning models [Bronstein et al., 2017, Hamilton et al., 2017]. This is often due to the high complexity required to model time-dependent relationships. However, leveraging this additional information is crucial to capturing many key interactions.

We define the problem of dynamic graph representation learning in continuous time, with dynamic graph structure. This graph is sampled from a continuous distribution $\mathcal{G} \in \mathbb{G}(t)$ with $\mathcal{G}_0 = \mathbb{G}(t = t_0)$ as a boundary condition. Each snapshot $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}^t)$ is a weighted graph with a shared node set $\mathcal{V}$, an edge set $\mathcal{E}^t$, and weighted adjacency matrix $\mathbf{A}^t$ at time t. Unlike some previous work that assumes links can only be added over time, we allow for link removal. Continuous dynamic graph representation learning aims to learn latent representations $\mathbf{e}_v^t \in \mathbb{R}^d$ for each node $v \in \mathcal{V}$ at time t, such that $\mathbf{e}_v^t$ preserves both the local graph structures centred at v and its evolutionary behaviours prior to time t.

Several methods that tackle different variations of this problem have been proposed in the literature. JODIE [Kumar et al., 2019] is a recurrent neural network that predicts instance trajectories for bipartite interactions. CTDNE by Nguyen et al. [2018] has advantages; it works on both directed and undirected graphs. However, it does not generalise to heterogeneous networks and makes use of random-walk methods in its current form, which has known limitations [Ribeiro et al., 2017]. Non random walk based methods [Veličković et al., 2018] have also been extended to incorporate temporal information [Opolka et al., 2019]. Spatial Temporal Graph Convolutional Networks [Yan et al., 2018] utilise graph convolutional networks to join spatial edges and temporal edges in a supervised setting. However, these methods are designed to work with discrete-time sampling and are not always naturally extended into *continuous* time representations. We refer the interested reader to reviews by Holme and Saramäki [2012] and Casteigts et al. [2012] for more comprehensive discussions.

An application in travel

An example use case for continuous graph representations is node embedding in a live travel network. The network consists of transport hubs (node set, $\mathcal{V}$) connected together by routes (edge set, $\mathcal{E}^t$) that have time-dependent properties. Modelling the continuous time aspects of the graph allows us to capture the effects of high-frequency intra-day changes in the network as well as low-frequency changes, assuming sufficient sampling frequency in the data.

A learned representation could be used to augment a recommender system to provide travellers with better personalised suggestions for destinations and travel itineraries. Recommender systems are well suited to be framed as graph learning problems [Monti et al., 2017, Ying et al., 2018], therefore, framing the problem end-to-end as graphs.

Graph Representation Learning Workshop at NeurIPS 2019, Vancouver, Canada.

References

Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017.

Srijan Kumar, Xikun Zhang, and Jure Leskovec. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1269–1278. ACM, 2019.

Giang Hoang Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunyee Koh, and Sungchul Kim. Continuous-time dynamic network embeddings. In *Companion Proceedings of the The Web Conference 2018*, pages 969–976. International World Wide Web Conferences Steering Committee, 2018.

Leonardo FR Ribeiro, Pedro HP Saverese, and Daniel R Figueiredo. struc2vec: Learning node representations from structural identity. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 385–394. ACM, 2017.

Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. *arXiv preprint arXiv:1809.10341*, 2018.

Felix L Opolka, Aaron Solomon, Cătălina Cangea, Petar Veličković, Pietro Liò, and R Devon Hjelm. Spatio-temporal deep graph infomax. *arXiv preprint arXiv:1904.06316*, 2019.

Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Petter Holme and Jari Saramäki. Temporal networks. *Physics reports*, 519(3):97–125, 2012.

Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems*, 27 (5):387–408, 2012.

Federico Monti, Michael Bronstein, and Xavier Bresson. Geometric matrix completion with recurrent multi-graph neural networks. In *Advances in Neural Information Processing Systems*, pages 3697–3707, 2017.

Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 974–983. ACM, 2018.