Task-Focused Dynamic Network Inference from Node Attribute Time Series

dynamic networks, network inference, dynamic attributed network simulation, graph neural network

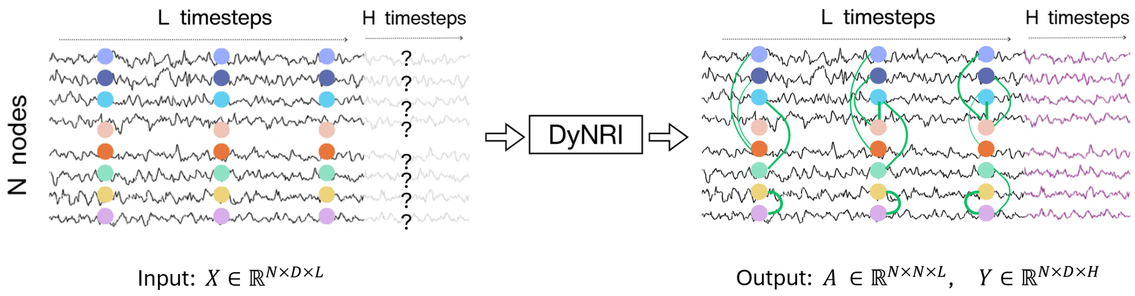

Figure 1: Schematic of the DyNRI framework. Historical node attributes are encoded into dynamic graphs, which guide autoregressive prediction of future attributes.

Extended Abstract

Many systems of scientific and societal importance, ranging from brain connectivity to ecological interactions and financial trust networks, are inherently dynamic. Both the attributes of individual entities and the patterns of interaction between them evolve over time, leading to complex dependencies that are most naturally represented as dynamic networks. Yet in most real-world settings, the underlying network structure is not observed directly, and must instead be inferred from time series of node-level measurements [1, 2]. Network reconstruction has a long history in neuroscience, biology, and economics, but most existing approaches assume static topologies [3–7] or require past observations of edges [8, 9]. These assumptions limit their ability to capture the transient, heterogeneous, and evolving relationships that shape real-world dynamics.

We present **DyNRI**, a framework for *task-focused dynamic network inference* that reconstructs evolving interaction networks directly from node attribute time series. DyNRI builds on Neural Relational Inference (NRI) [10] by extending it to time-varying settings. The model employs an encoder–decoder architecture in which the encoder infers distributions over latent edges at each time step using Gumbel-Softmax sampling [11, 12], and the decoder autoregressively predicts future node attributes conditioned on the inferred dynamic graphs. By aligning network inference with the task of multi-step forecasting, DyNRI learns structures that are not only plausible explanations of past data but also useful for prediction. This distinguishes our approach from purely unsupervised structure discovery and grounds inference in predictive utility.

A major challenge in this domain is evaluation: ground-truth dynamic graphs are rarely available. To address this, we introduce a co-evolutionary simulation [13] framework that generates coupled dynamics of node attributes and connectivity. In our simulator, nodes move in a latent space according to a stochastic policy that balances attraction to the centroid of their local neighborhood with independent goal-directed movement. Edges are formed based on k-nearest neighbors subject to a distance threshold, ensuring that attribute similarity drives connectivity while connectivity, in turn, influences subsequent attribute updates. This bi-directional coupling produces realistic, non-trivial dynamics where attributes and structure mutually influence each other, providing a principled testbed for dynamic inference methods.

Preliminary experiments on synthetic data suggest that DyNRI can recover meaningful dynamic structures without access to ground-truth edges, while maintaining prediction accuracy comparable to or exceeding baselines. Ongoing work is extending these experiments to financial trust networks such as Bitcoin Alpha and OTC [14–16], and to ecological datasets involving animal movement and implicit interactions within groups.

The ethical implications of this work are significant. In contexts involving human or financial systems, reconstructing hidden networks could reveal sensitive or private relationships. If node attributes encode social or demographic biases, inferred networks may amplify these biases. Moreover, misuse of such inferred structures could enable surveillance or manipulation. To mitigate these risks, we validate our methods first on synthetic and ecological case studies, emphasize transparency in modeling assumptions, and commit to reproducibility and open science practices. We see DyNRI not as a tool for covert inference, but as a research framework for understanding the principles of dynamic network reconstruction and evaluating algorithms under controlled conditions.

In conclusion, DyNRI introduces a task-aligned approach to dynamic network inference and a co-evolutionary simulation benchmark for its evaluation. This submission represents work in progress: our goal is to refine the framework, benchmark it against state-of-the-art baselines, and apply it across domains where dynamic networks are central, from neuroscience to ecology to finance.

References

[1] Petter Holme and Jari Saramäki. "Temporal Networks". In: *Physics Reports* 519.3 (2012), pp. 97–125.

[2] Mladen Kolar et al. "Estimating time-varying networks". In: *Annals of Applied Statistics* 4.1 (2010), pp. 94–123.

[3] Bing Yu, Haoteng Yin, and Zhanxing Zhu. "Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting". In: *IJCAI*. 2018. URL: https://www.ijcai.org/proceedings/2018/505.

[4] Yaguang Li et al. "Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting". In: *ICLR*. 2018. URL: https://openreview.net/pdf?id=SJiHXGWAZ.

[5] Shengnan Guo et al. "Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting". In: *AAAI*. 2019. URL: https://ojs.aaai.org/index.php/AAAI/article/view/3881.

[6] Zonghan Wu et al. "Graph WaveNet for Deep Spatial-Temporal Graph Modeling". In: *IJCAI*. 2019. URL: https://www.ijcai.org/proceedings/2019/264.

[7] Lei Zhao et al. "T-GCN: A Temporal Graph Convolutional Network for Traffic Prediction". In: *arXiv:1811.05320* (2018). URL: https://arxiv.org/pdf/1811.05320.

[8] Aldo Pareja et al. “EvolveGCN: Evolving Graph Convolutional Networks for Dynamic Graphs”. In: *AAAI*. 2020, pp. 5363–5370. URL: https://ojs.aaai.org/index.php/AAAI/article/view/5984.

[9] Emanuele Rossi et al. “Temporal Graph Networks for Deep Learning on Dynamic Graphs”. In: *arXiv:2006.10637* (2020). URL: https://arxiv.org/abs/2006.10637.

[10] Thomas Kipf et al. “Neural Relational Inference for Interacting Systems”. In: *ICML*. 2018. URL: https://proceedings.mlr.press/v80/kipf18a/kipf18a.pdf.

[11] Eric Jang, Shixiang Gu, and Ben Poole. “Categorical Reparameterization with Gumbel-Softmax”. In: *arXiv:1611.01144* (2017). URL: https://www.cs.toronto.edu/~bonner/courses/2022s/csc2547/papers/generative/background/categorical-reparameterization%2C-jang%2C-arxiv-2017.pdf.

[12] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. “The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables”. In: *arXiv:1611.00712* (2017). URL: https://arxiv.org/abs/1611.00712.

[13] Thilo Gross and Bernd Blasius. “Adaptive Coevolutionary Networks: A Review”. In: *Journal of the Royal Society Interface* 5.20 (2008), pp. 259–271.

[14] *Bitcoin Alpha trust weighted signed network*. SNAP Dataset page. Accessed 2025-09-05. 2016. URL: https://snap.stanford.edu/data/soc-sign-bitcoinalpha.html.

[15] *Bitcoin OTC trust weighted signed network*. SNAP Dataset page. Accessed 2025-09-05. 2016. URL: https://snap.stanford.edu/data/soc-sign-bitcoin-otc.html.

[16] Srijan Kumar et al. “Edge Weight Prediction in Weighted Signed Networks”. In: *IEEE International Conference on Data Mining (ICDM)*. 2016, pp. 221–230. URL: https://cs.stanford.edu/~srijan/pubs/wsn-icdm16.pdf.