Dynamic Graph Clustering with Graph Neural Networks:
spectral end-to-end differentiable optimization of temporal modularity

Background

PROBLEM

Node clustering (*community detection*) is a crucial task for many applications [1,2] to obtain coarse-grained graphs. In dynamic graphs with rich, high-dimensional node or edge features, it requires **jointly modeling** structure- and attribute-based signals, while also considering temporal dynamics at micro, macro and mesoscales.

RELATED WORK

A compendium of techniques has been proposed for community detection [2], including spectral, inferential, and greedy heuristic optimization approaches; and more recently, neural networks for graphs, each with their very own (dis)advantages – and often **conflicting definitions** of what constitutes a *community* itself.

LIMITATIONS

Majority of models consider **static communities**; the advantages of GNNs over principled statistical methods is not clear-cut [3]; scarcity of quality real-world data [4,5]; reliance on proxy objectives (e.g., reconstruction error) and memory-based mechanisms (GRUs, LSTM, etc.) increase computational costs and reduce interpretability.

L-MoN 🍋

We propose a **L**ongitudinal **M**odularity **o**ptimization **N**etwork based on **DMoN [6]** for **dynamic** community detection on evolving networks:

INPUT

(continuous-time temporal graph)

HIDDEN LAYER(S)

(any optionally multilayer graph encoder)

OUTPUT

(community assignments)

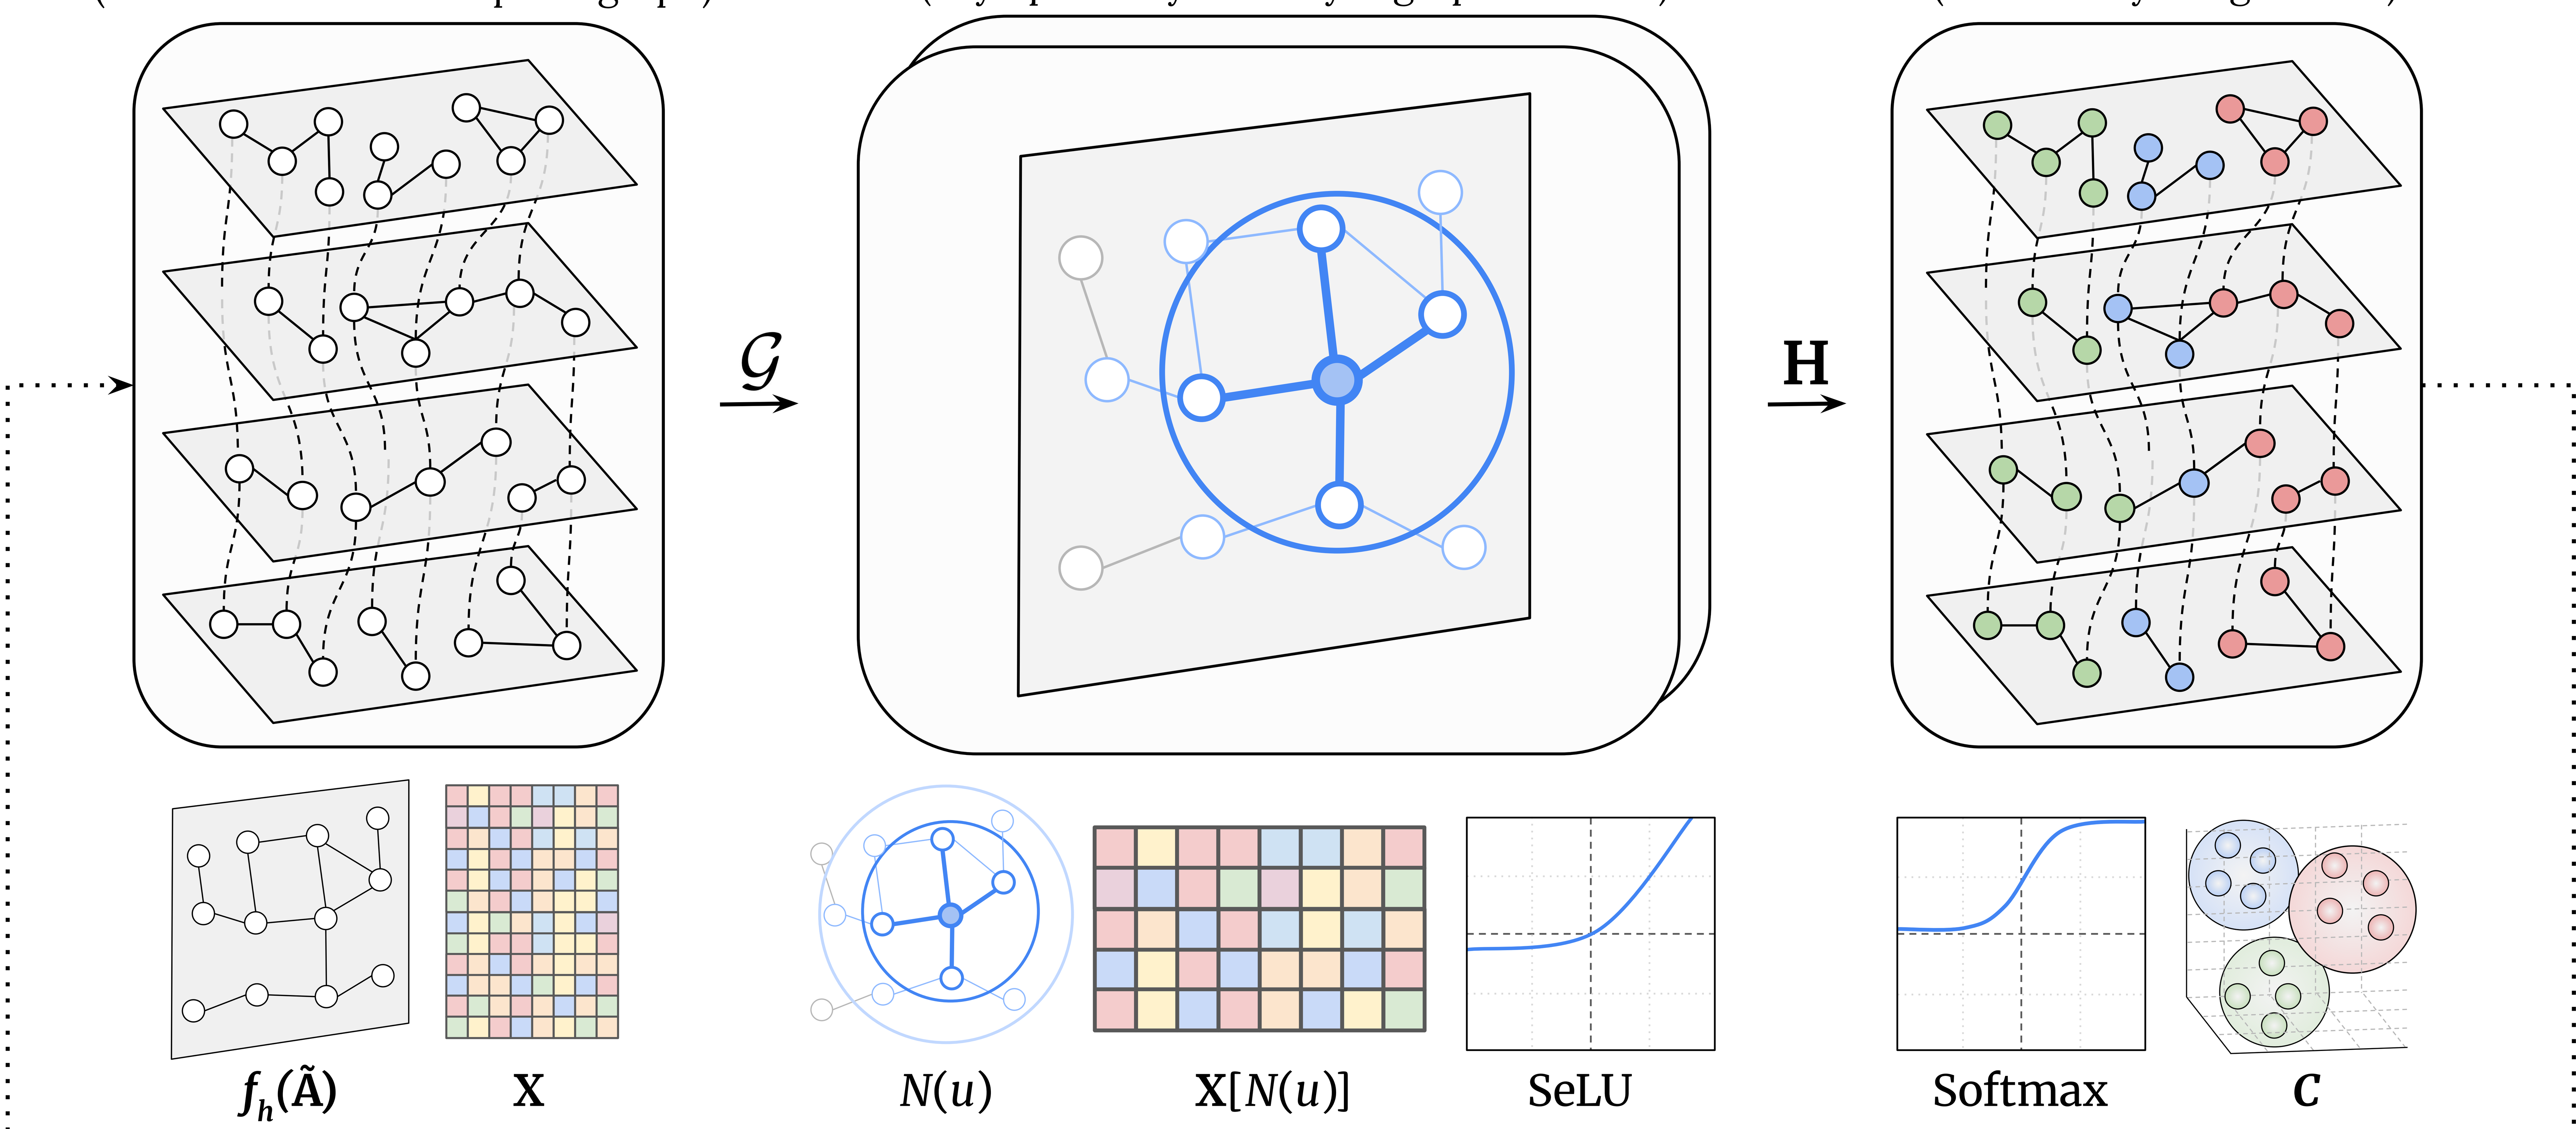

A Hawkes-Laplacian operator f_h models pairwise interaction strength in an **encoder-agnostic** way.

A novel spectral relaxation of **longitudinal modularity** [7] for CTTGs, where a Riemann sum in the loss incentivizes contiguous label assignments.

An unpooling layer obtains clusters from event-level representations maintaining* **best-in-class** complexity $\mathcal{O}(d^2 n + m)$.

$$\mathcal{L}_{\text{LMoN}} = \underbrace{-\mathcal{Q} + \mathcal{S} + \mathcal{R}}_{\substack{\text{Modularity} \\ \text{Sm.+Reg.}}} + \underbrace{\mathcal{H}}_{\substack{\text{Hawkes} \\ \text{NLL}}}$$

(*) Our current and main **limitation**: how can we extend our graph learning pipeline for dynamic feature aggregation preserving its scalability?

[1] Temporal Network Theory. **Holme & Saramäki (Org.)**, Springer, 2012.
[2] Community detection in networks: A user guide. **Fortunato & Hric**, Physics Reports, 2016.
[3] Statistical inference links data and theory in network science. **Peel et al.**, Nature Communications, 2022.
[4] The ground truth about metadata and community detection in networks. **Peel et al.**, Science Advances, 2017.
[5] TADC-SBM: A time-varying, attributed, degree-corrected stochastic block model. **Passos et al.**, ISCC, 2025.
[6] Graph Clustering with Graph Neural Networks. **Tsitsulin et al.**, JMLR, 2023.
[7] Longitudinal modularity, a modularity for link streams. **Brabant et al.**, EPJ Data Science, 2025.

Preprint coming soon! :-)