Towards Better Evaluation for Dynamic Link Prediction

Farimah Poursafaei*, Shenyang Huang*, Kellin Pelrine, Reihaneh Rabbany

Learning on Dynamic Graphs

Dynamic graphs are used for modeling many real-world networks.

- One important task: dynamic link prediction

  Given a timestamped stream of edges:

  $G = \{(s_1, d_1, t_1), (s_2, d_2, t_2), \dots\}; \qquad 0 \le t_1 \le t_2 \le \dots \le T$

  Objective:

  Predicting the existence of an edge between a pair of nodes in the future.

- Remarkable observation:

  SOTA methods have **near-perfect performance** for dynamic link prediction!

  → Current evaluation settings are not challenging enough.

Historical and Inductive Negative Sampling

Historical Negative Sampling

$e \in E_{train} \cap \overline{E_t}$

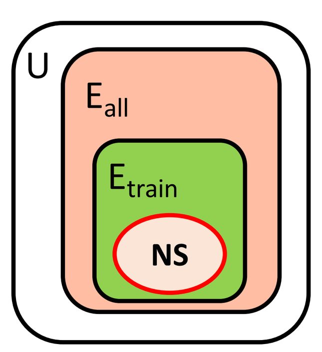

Random Negative Sampling

*The **standard** setting.*

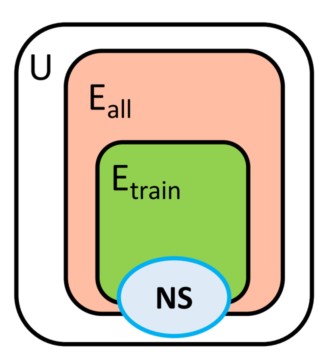

Inductive Negative Sampling

$e \in E_{test} \cap \overline{E_{train}} \cap \overline{E_t}$

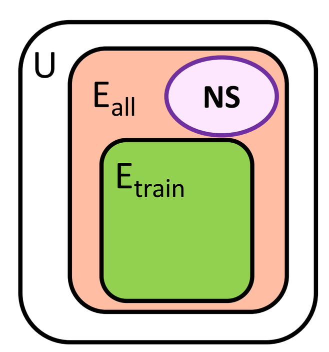

EdgeBank: A Baseline for Dynamic Link Prediction

A **pure memorization-based** approach: A ***bank*** of observed edges

Detects edges with ***frequent recurrence*** patterns but fails for a previously seen negative edge or an unseen positive edge

$EdgeBank_\infty$ memorizes all observed edges

$EdgeBank_{tw}$ only memorizes edges from a time window from the recent past

Summary of Contributions

We help solve the dynamic link prediction challenges:

- **Limited domain diversity.**

  Six new dynamic network datasets:

Domain	Transport	Politics			Economics	Proximity
New Dataset	*Flights*	*Can. Parl.*	*US Legis.*	*UN Trade*	*UN Vote*	*Contact*

  Novel visualization techniques for the investigation of dynamic networks:

  TEA & TET plots

- **Easy negative edges.**

  Novel negative sampling strategies: Historical NS and Inductive NS

- **Lack of a simple baseline.**

  EdgeBank: a baseline for dynamic link prediction

Impact of Negative Sampling

- The ranking of methods changes in the historical and inductive NS.
- EdgeBank shows competitive performance, particularly in the standard setting, and even outperforms in some datasets, e.g. *LastFM.*
- With the alternative negative sampling (i.e. historical and inductive NS):

  A clear gap between the performance of the SOTA models and EdgeBank.

  The models' ranking changes.

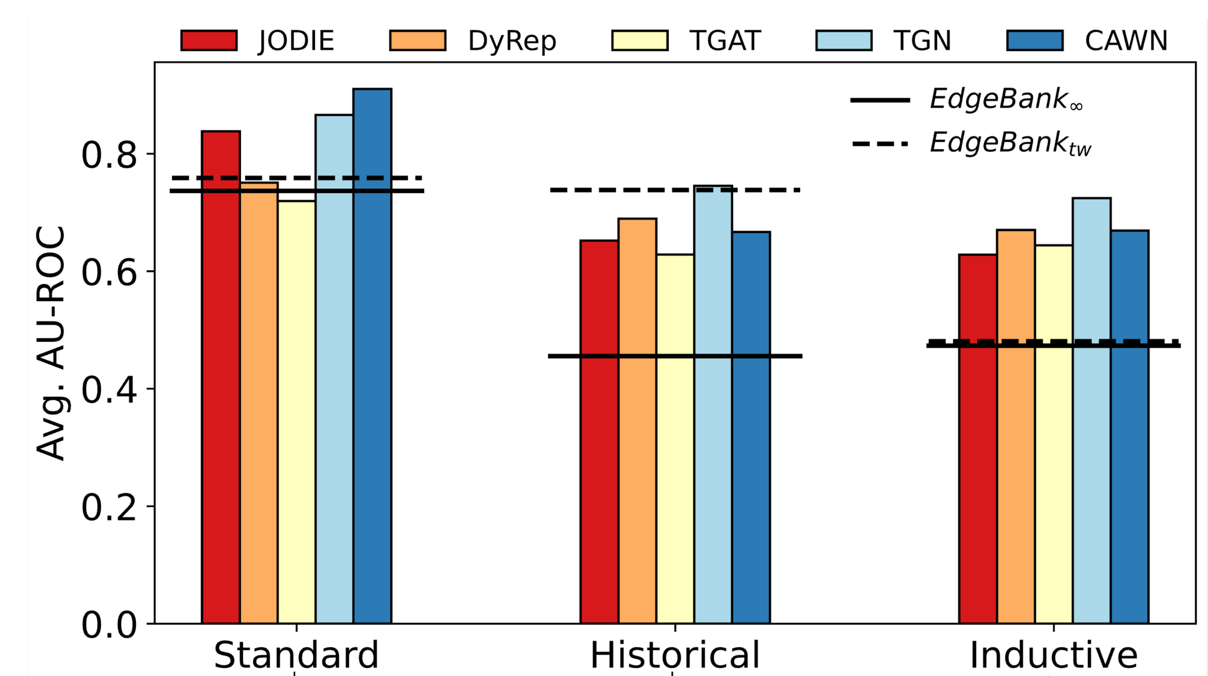

Understanding Dynamic Graph Datasets

TEA (Temporal Edge Appearance) Plots

- High variance in temporal evolutionary patterns.

$$novelty = \frac{1}{T}\sum_{t=1}^{T}\frac{|E^t \setminus E^t_{seen}|}{|E^t|}, where\ E^t = \{(s,d,t_e)|t_e = t\}\ and\ E^t_{seen} = \{(s,d,t_e)|t_e < t\}$$

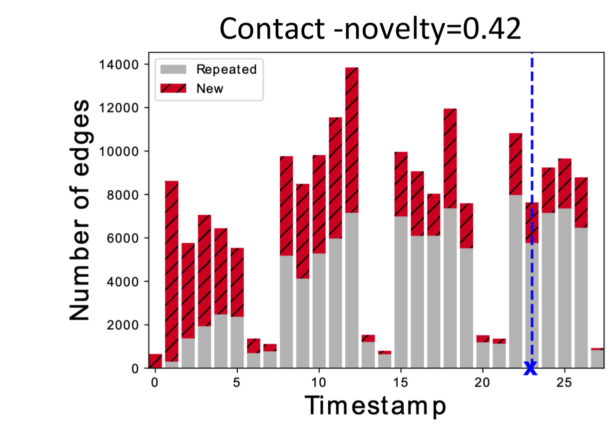

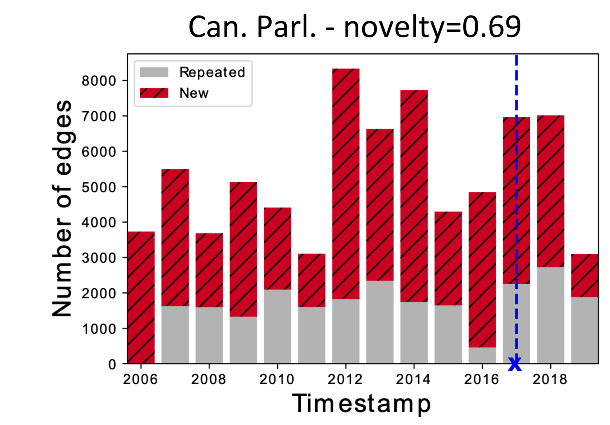

TET (Temporal Edge Traffic) Plots

- Recurrence pattern of edges over time.

$$recurrence = \frac{|E_{train} \cap E_{test}|}{|E\_train|}, \qquad surprise = \frac{|E_{test} \setminus E_{train}|}{|E_{test}|}$$

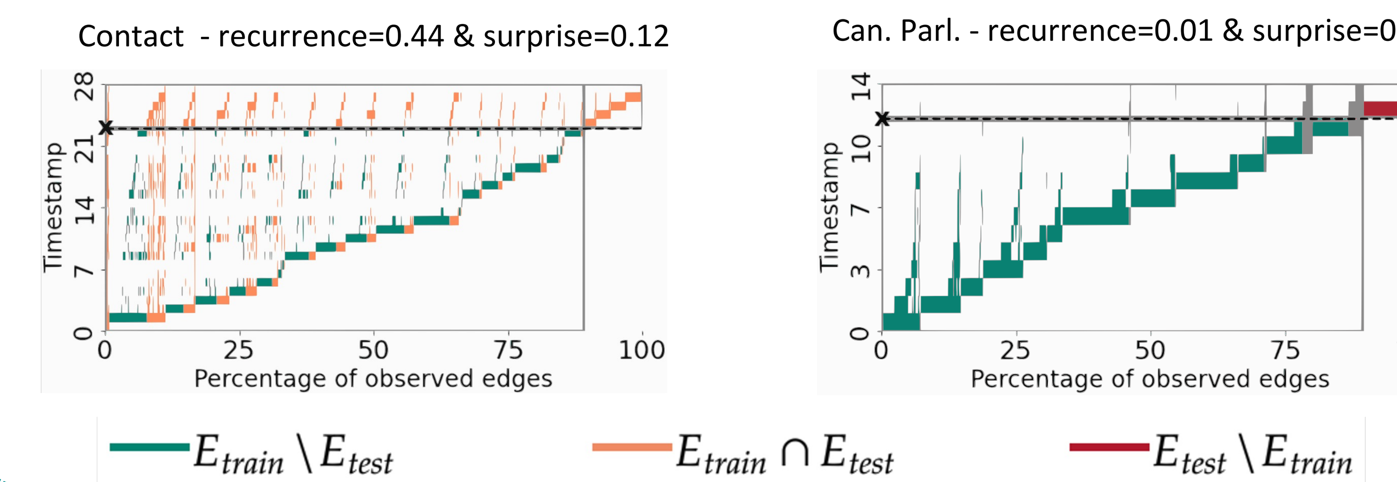

References

Srijan Kumar, et al. (2019). "Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks". In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining.

Rakshit Trivedi, et al. (2019). "DyRep: Learning Representations over Dynamic Graphs". In International Conference on Learning Representations.

Da Xu, et al. (2020). "Inductive Representation Learning on Temporal Graphs". arXiv preprint arXiv:2002.07962.

Emanuele Rossi, et. Al (2020). "Temporal Graph Networks for Deep Learning on Dynamic Graphs". arXiv preprint arXiv:2006.10637.

Yanbang Wang, et al. (2020). "Inductive Representation Learning in Temporal Networks via Causal Anonymous Walks". In International Conference on Learning Representations.

Standard Negative Sampling

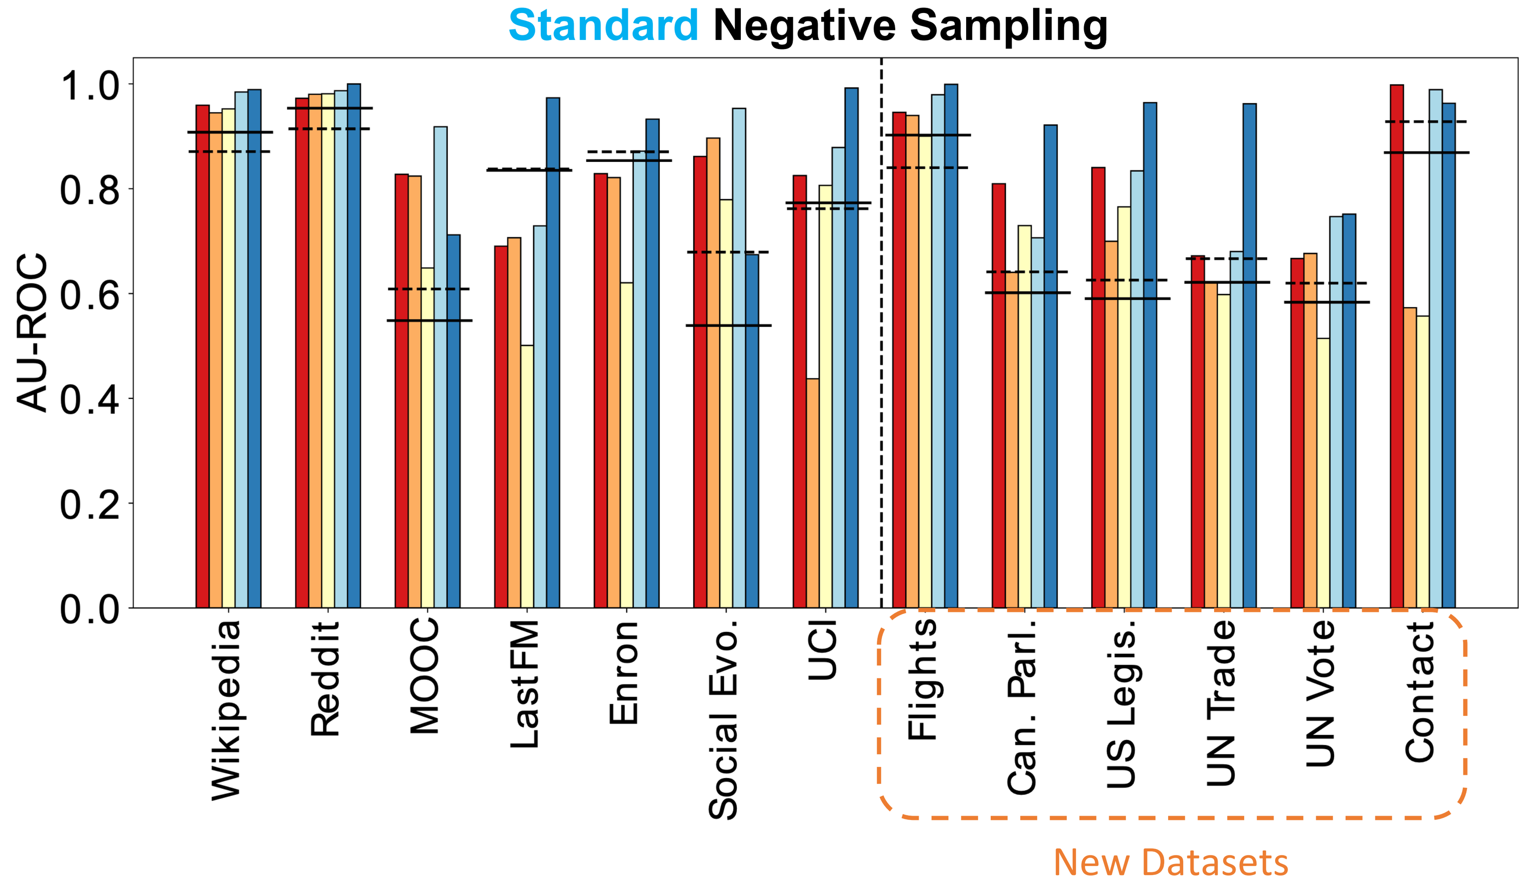

Historical Negative Sampling

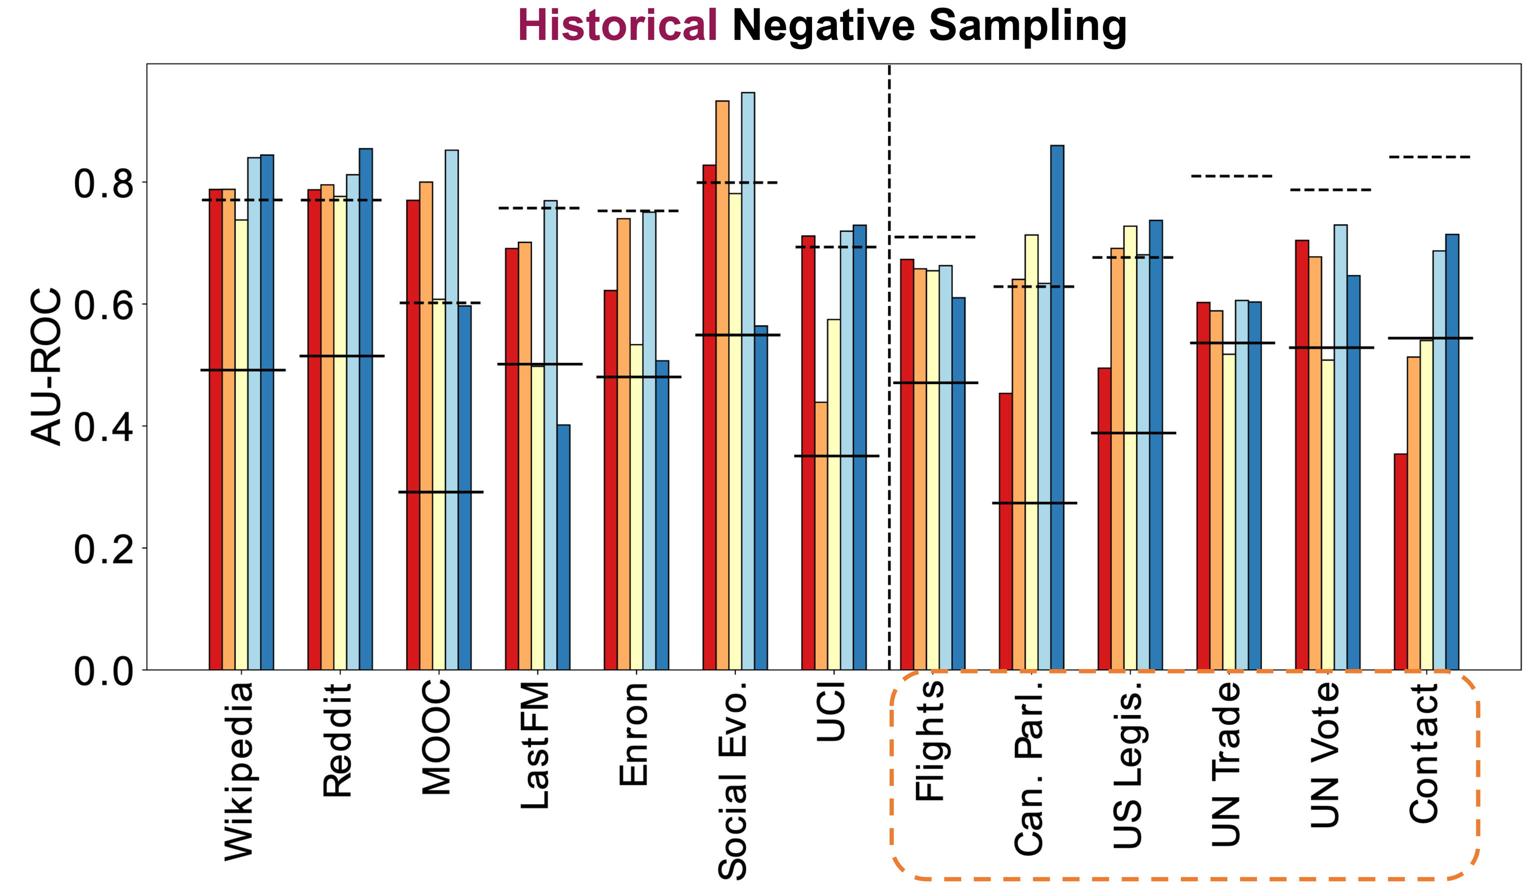

Inductive Negative Sampling