


DETECTING INFORMATION CHANNELS IN CONGRESSIONAL TRADING VIA TEMPORAL GRAPH LEARNING



Benjamin Pham Roodman*, Eugene Sy*,
Yu-Shiang Huang, J. Xavier Atero Vázquez,
Che Lin, Chuan-Ju Wang



INTRODUCTION

Publicly disclosed stock trades by American congresspeople are commonly interpreted as price signals. **We assembled a complex dataset of public records and a custom Temporal Graph Network (TGN) architecture to extract insights from trades and the robust corporate-political contexts in which they occur.** Because standard models typically discard interactions during the unresolved gap between a trade and its performance outcome, GAP-TGN is designed to maintain a fresh information state by processing interactions throughout the label resolution gap.

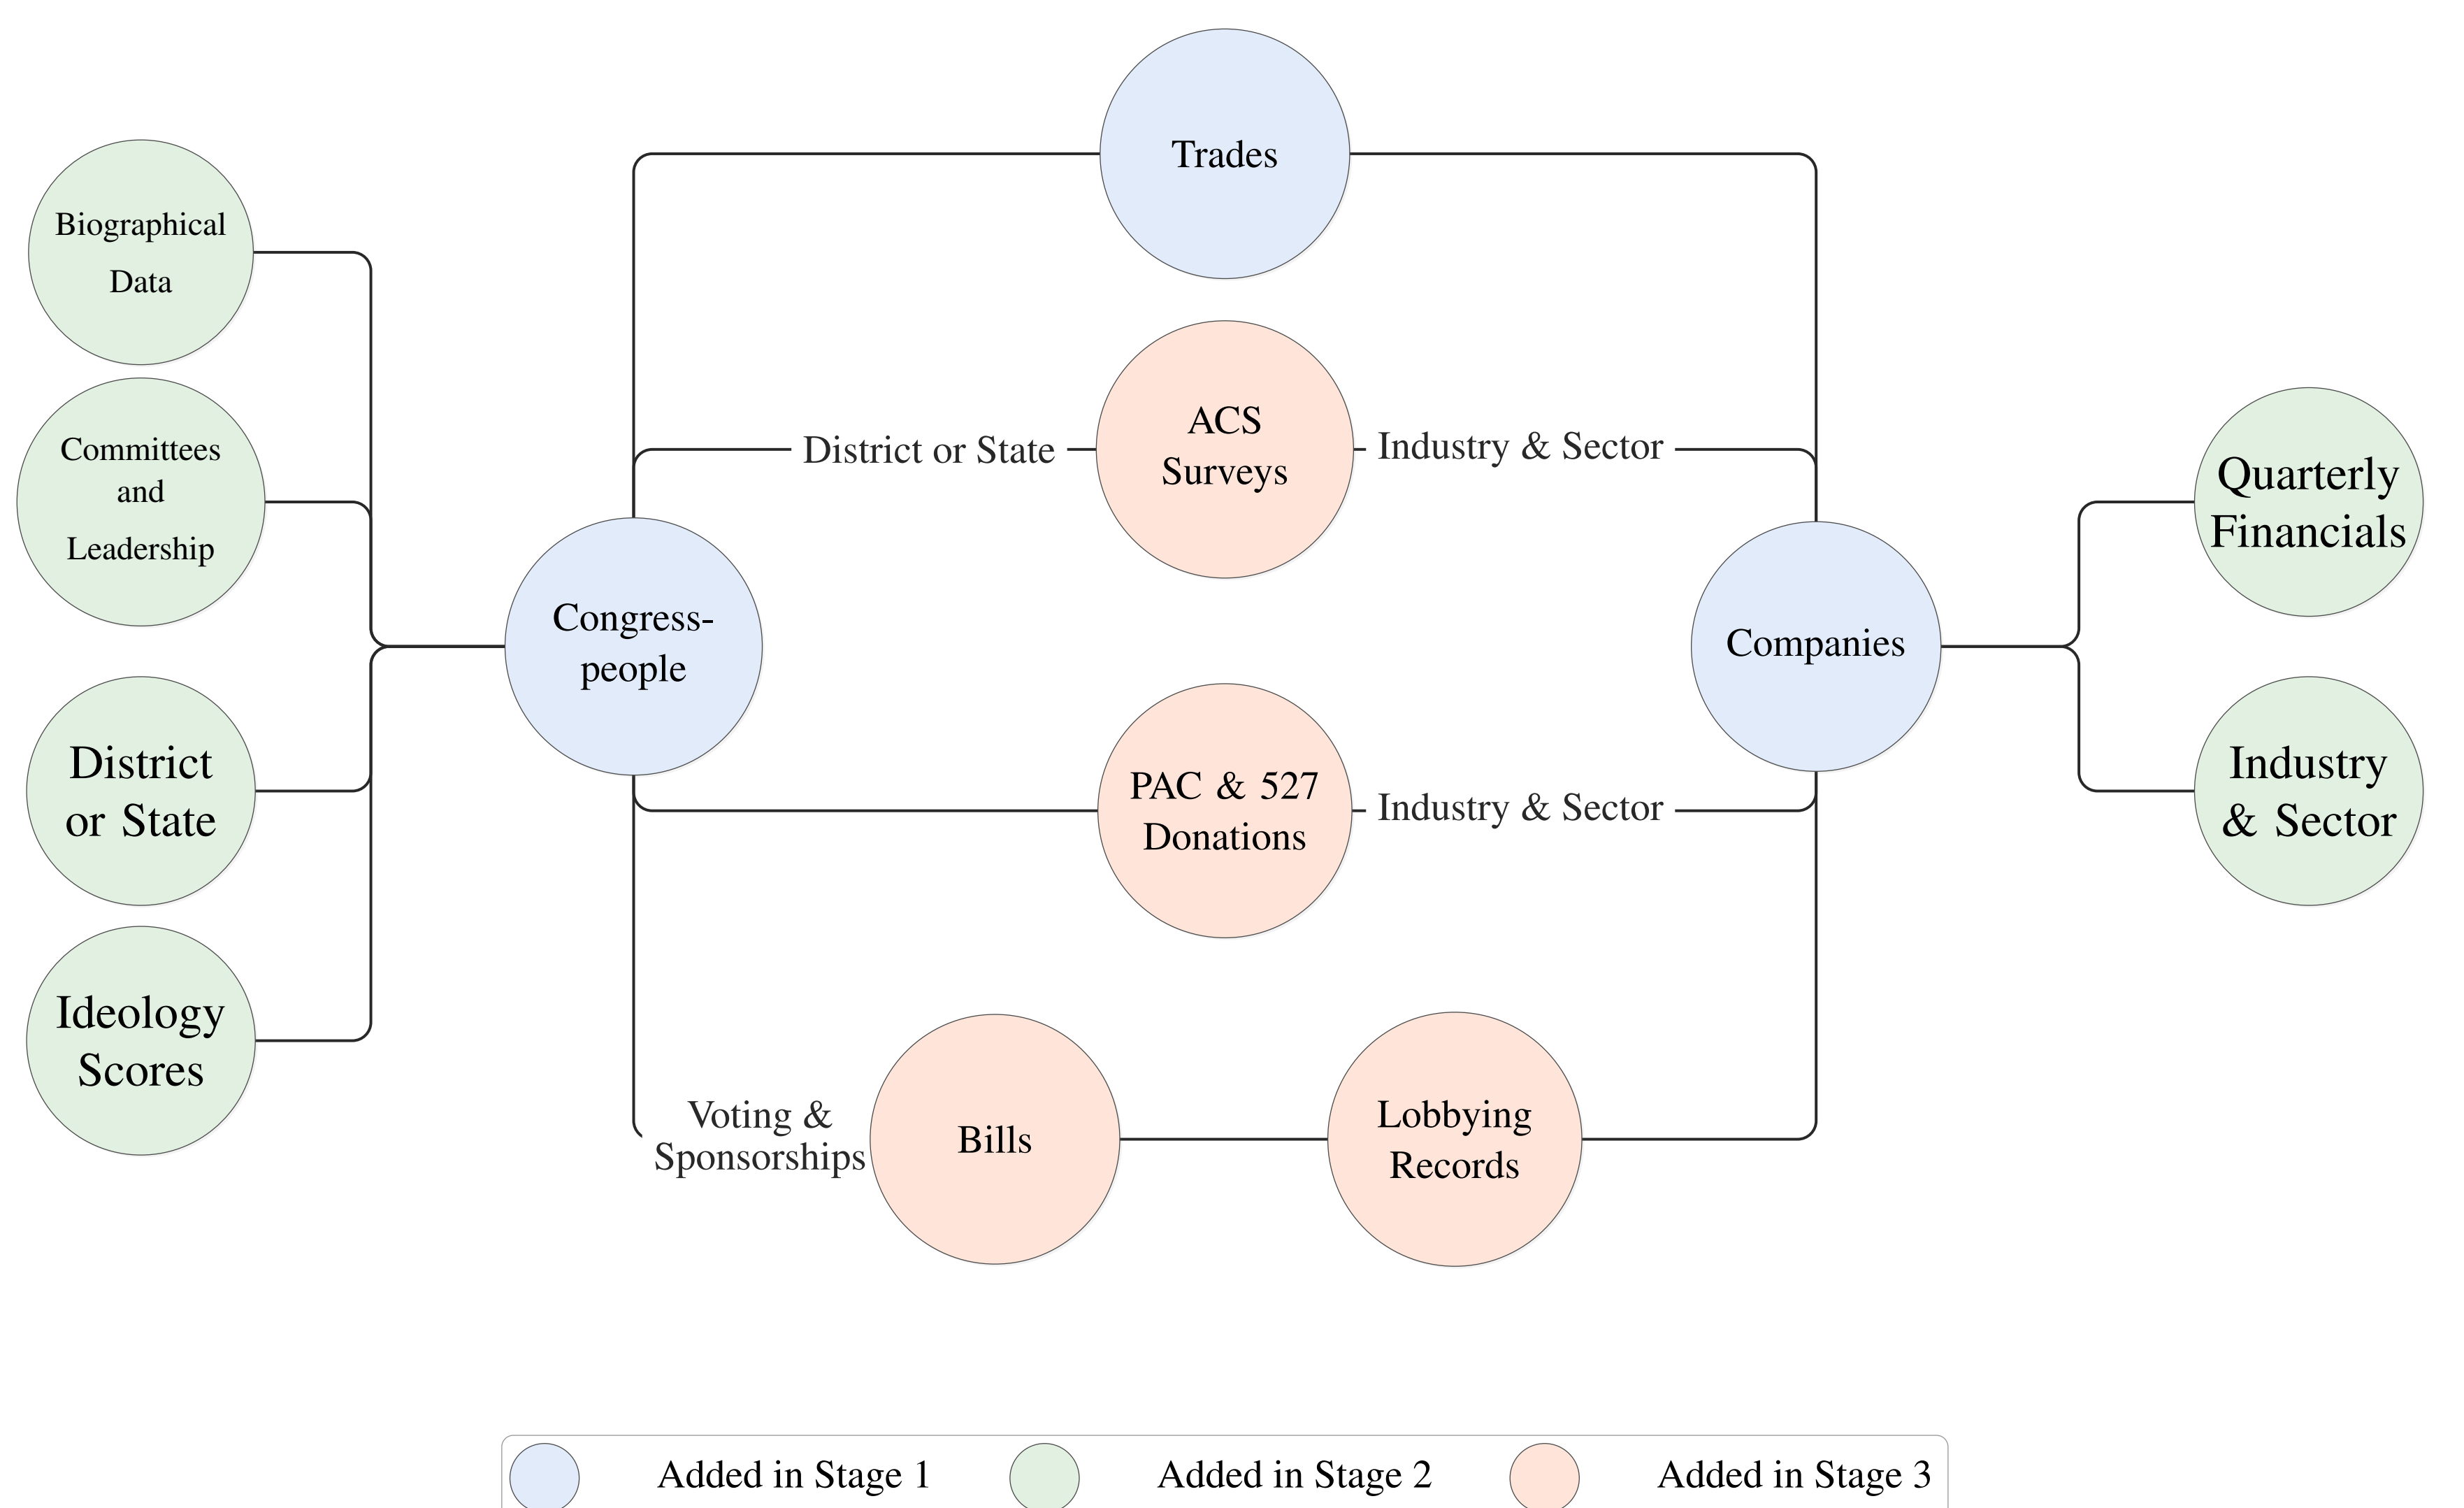
DATA CONSTRUCTION

The multimodal dataset:

- **Aggregates nine public and non-profit record sets** concerning behavior of American congresspeople and the companies they trade,
- **Maps politicians to companies** by channels of influence, information, and trade histories, as edges in a temporal graph structure,
- **Embeds detailed information** of congresspeople and companies as node attributes,
- **Fuses unstructured economic time series data** with the structured graph data,
- **Avoids data leakage** by indexing records by disclosure dates over reference dates.

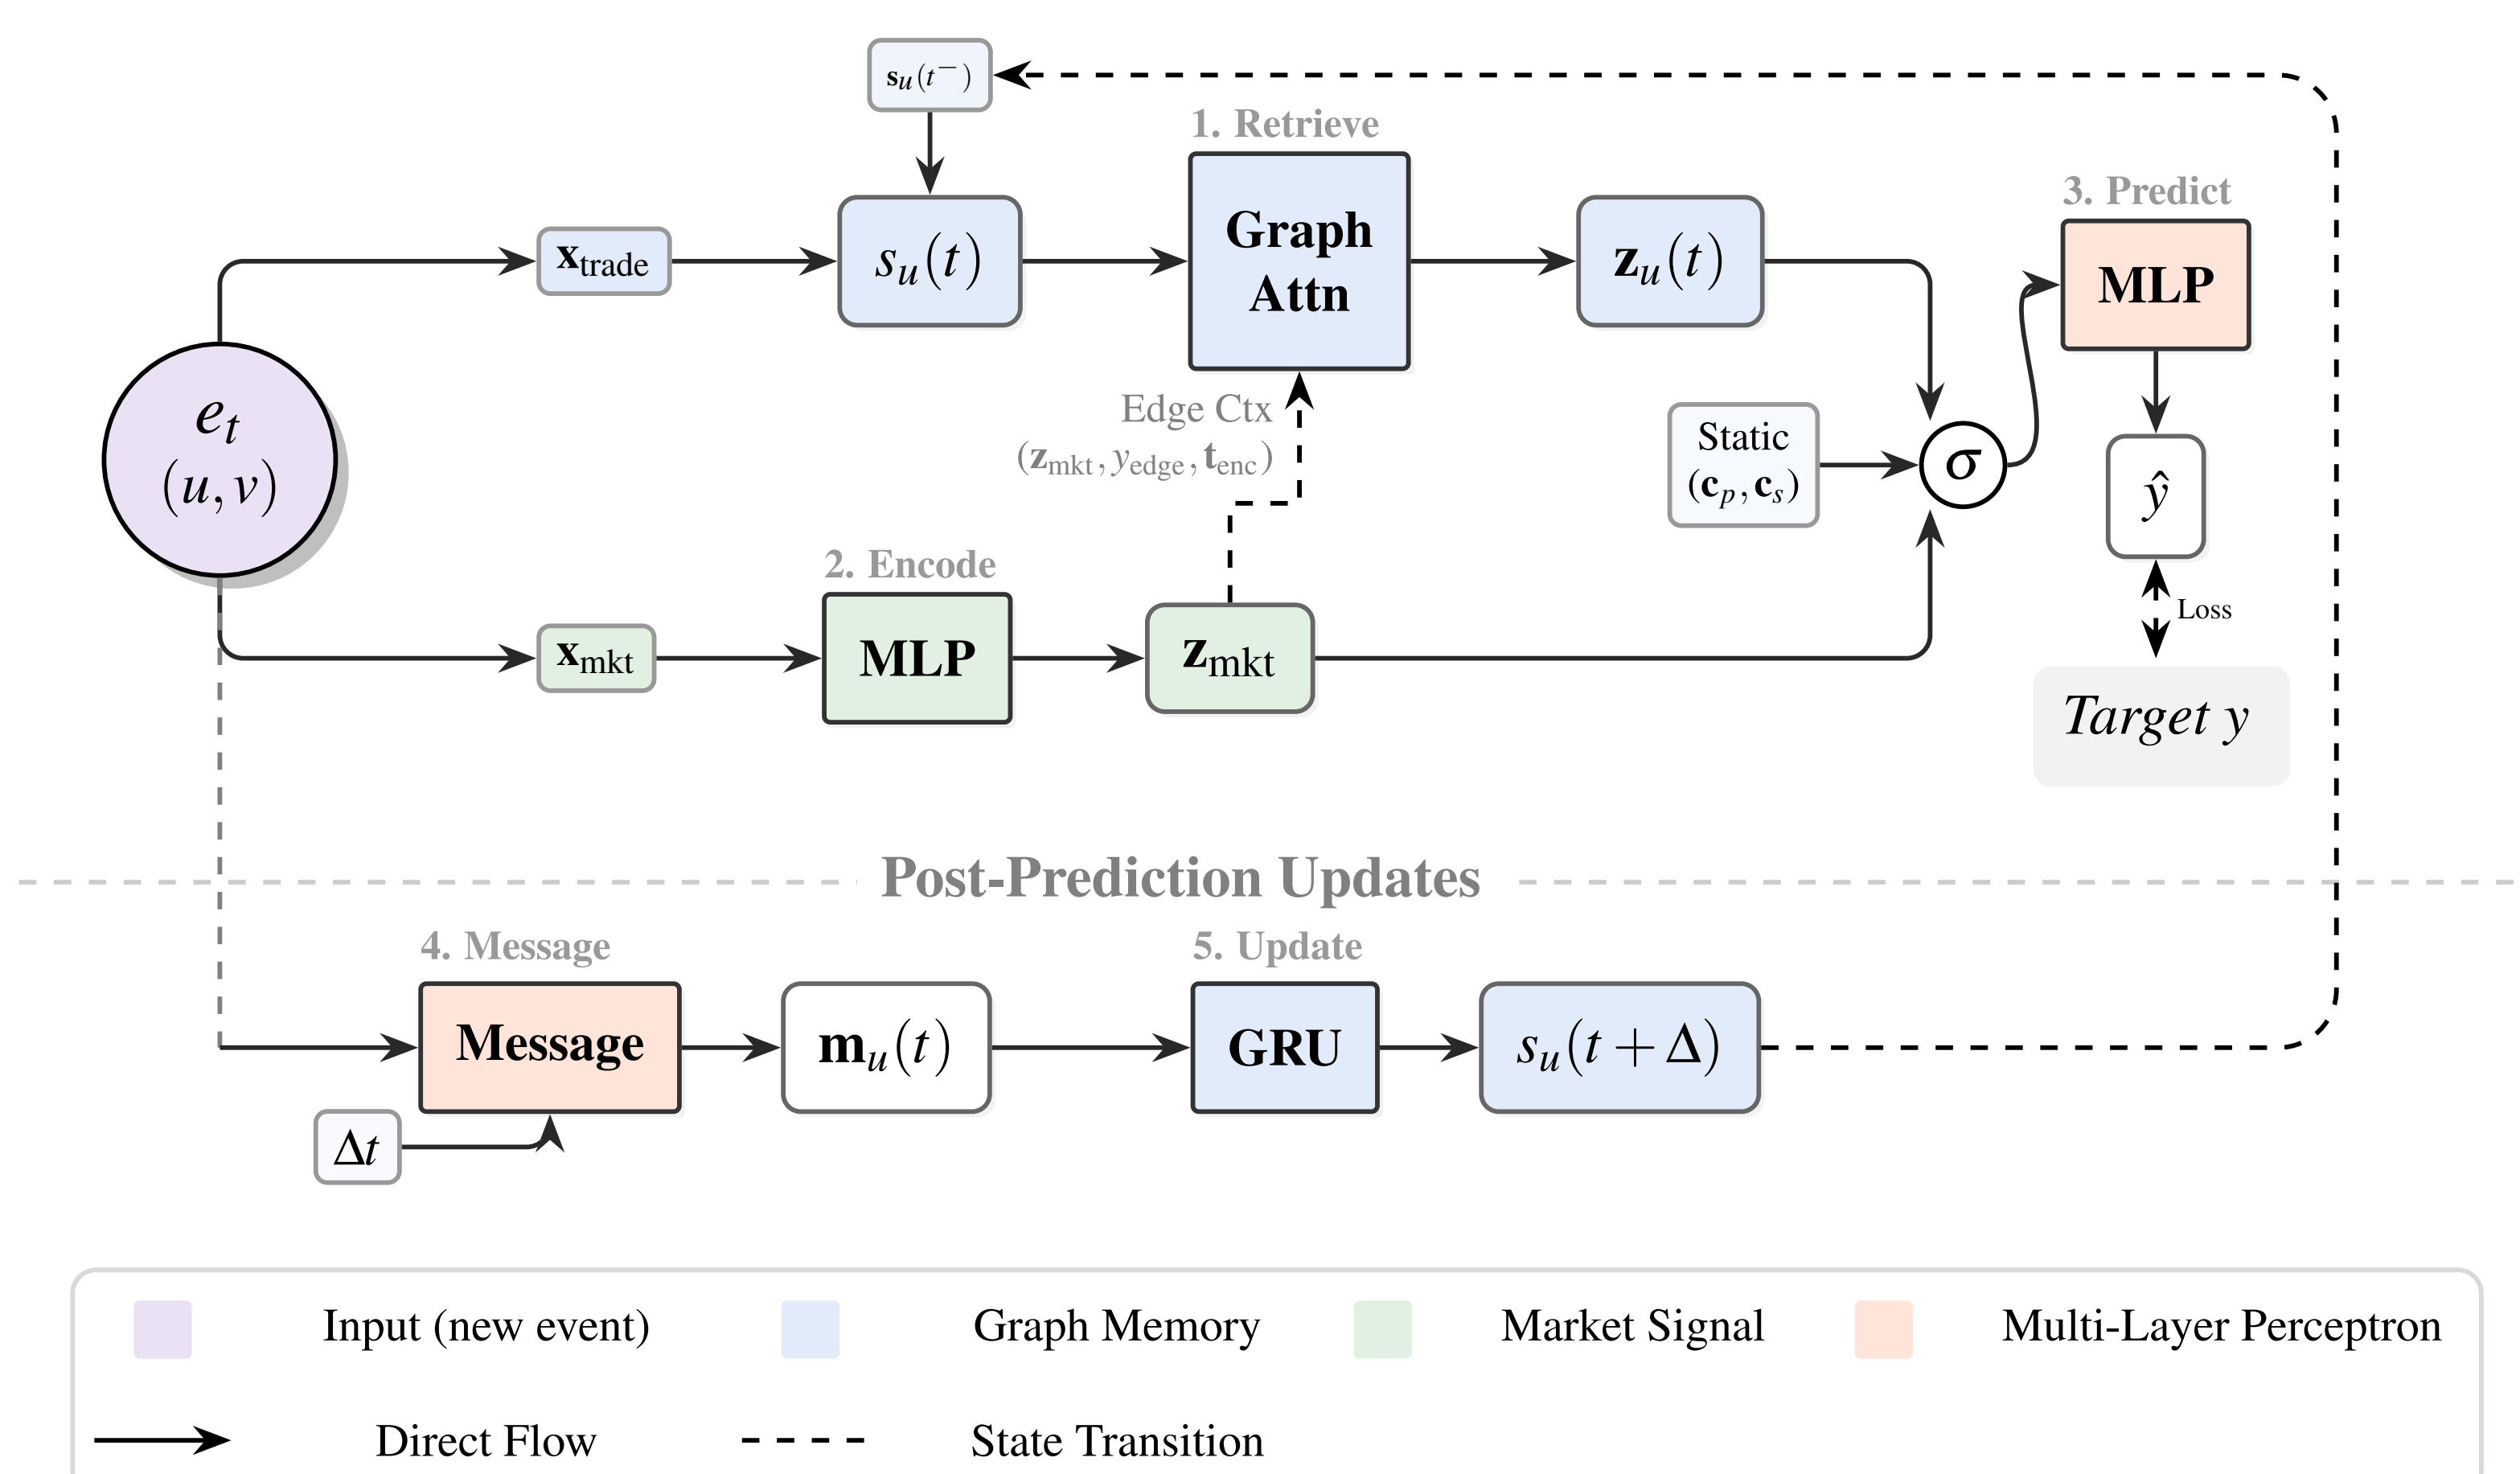
Table	Entries
Congressional Trades	97,431
Congressional Metadata	45,525
Roll Call Votes	26,262,296
Congressional Ideology Scores	32,740
Congressional District Employment	1,544,968
Congressional Lobbying Reports	1,690,362
PAC Campaign Donations	423,350,920
527 Org Campaign Transactions	8,714,809
Corporate Financial Indicators	20,309,853

DATA SCHEMA



We integrated political and corporate records into a temporal graph mapping congresspeople to companies, in three experimental stages.

GAP-TGN ARCHITECTURE



GAP-TGN uses graph attention and gated fusion to maintain state during the label resolution gap.

METHODOLOGY AND EXPERIMENTS

GAP-TGN handles training data when graph edges representing trades have appeared but are not yet labeled, as trade outcomes are not immediately known.

- **Detection of informed trades is cast as an edge classification task.**
- GAP-TGN assigns intermediate forecasted values between 0 (trade resolves as unprofitable) and 1 (trade resolves as profitable).
- Information propagates via a graph attention mechanism at each update. Edges represent trades, or other links between congresspeople and companies.
- A gated multi-modal fusion layer merges graph structural insights against temporal market signals to predict profitable trades.

GAP-TGN beats benchmarks on AUROC, when labeling trades by predicted profitability after several months.

Prediction Horizon (Months)	AUROC		F1-Score	
	18M	24M	18M	24M
Logistic Reg.	0.486	0.465	0.300	0.195
MLP	0.498	0.496	0.402	0.378
XGBoost	0.525	0.524	0.334	0.291
GAP-TGN	0.490	0.518	0.438	0.440



ICLR
International Conference On
Learning Representations

