Tackling Fake News Detection by Continually Improving Social Context Representations using Graph Neural Networks

BY NIKHIL MEHTA ET. AL. (ACL 2022)

PRESENTED BY:

GROUP 53

Members:

Presenter 1: Edwin Wang
Presenter 2: Shashank Rangarajan
Chia-Yu Tung
Rishabh Ghosh
Michael Guastalla

Motivation & Problem Description

Motivation:

- The popularity of social media has given rise to the spread of Fake News, hence, detecting it has become urgent and challenging

Problem Description:

- **Task:** Given sources, articles they publish, and social media users that interact with them, what is fake news?

- **Main Question:** Given social media context (in-terms of a graph) can we reason about fake news?

- **Main idea - *Social Homophily*:** the tendency of individuals to form social ties with others who share their views and preferences

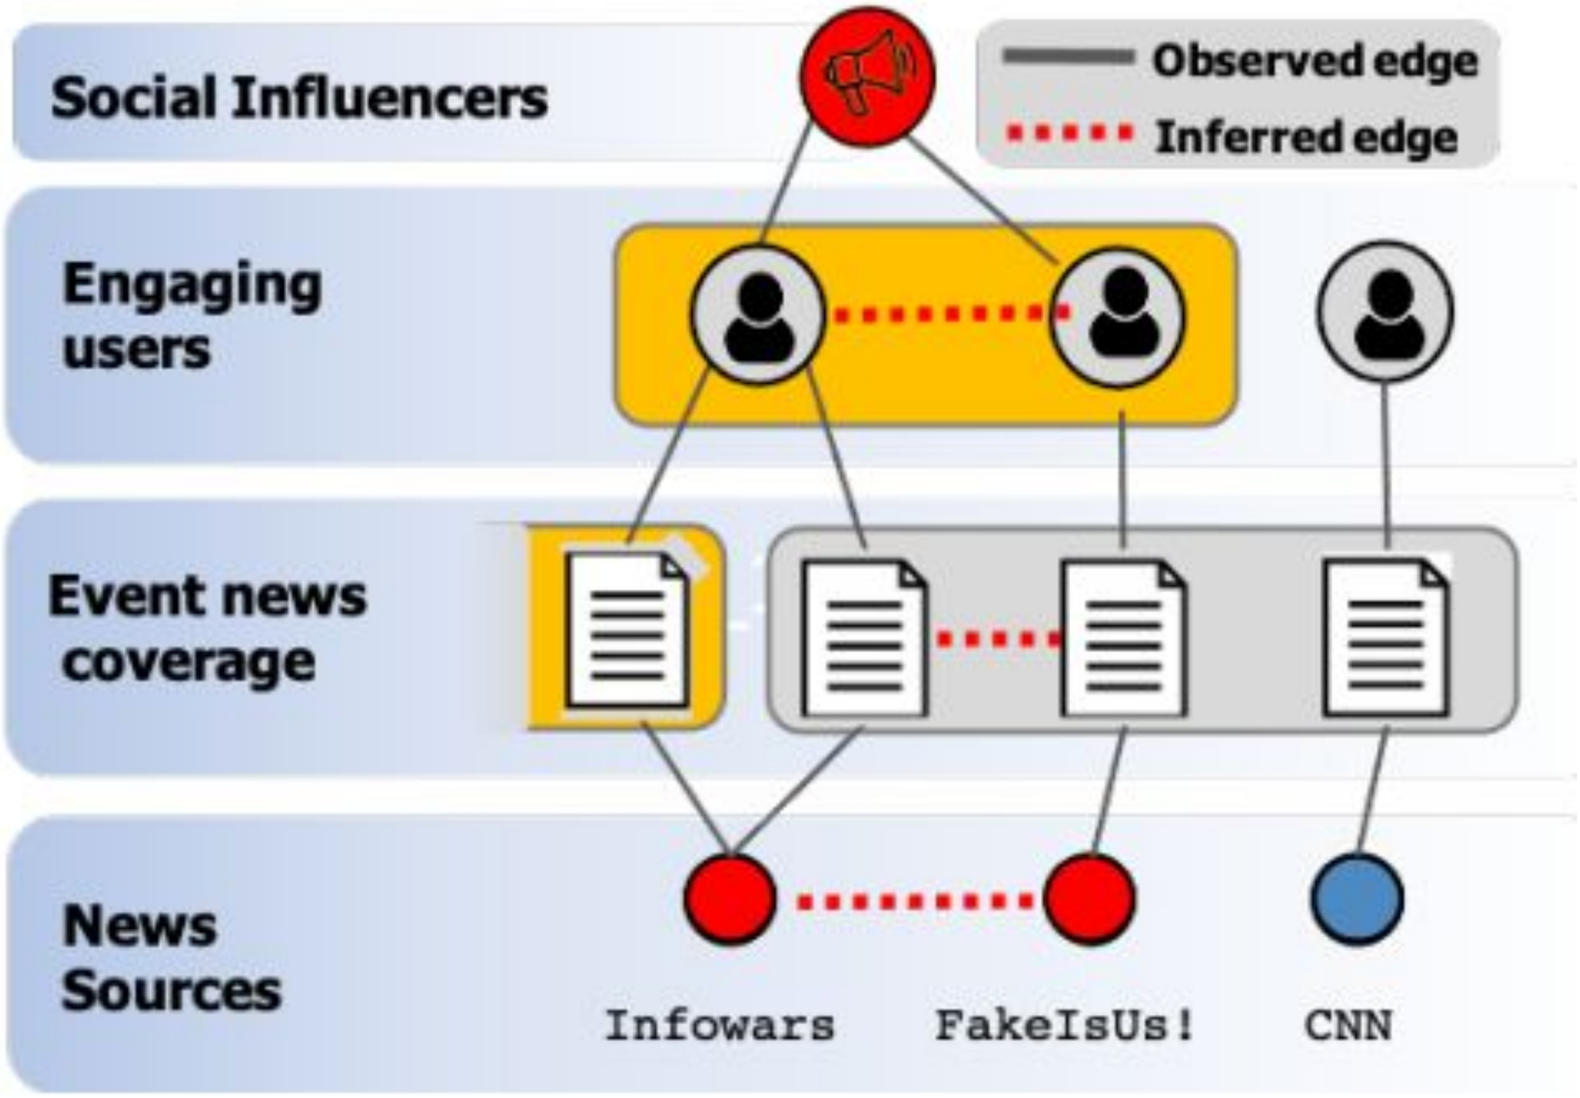
Social Influencers
Observed edge
Inferred edge
Engaging users
Event news coverage
News Sources
Infowars
FakeIsUs!
CNN

Related Works

- **Fake News Detection:**
    - Earlier works: Treat it as a classification problem, where they assess the factuality level from representations of news and their social context
        - *Disadvantage:* They do not capture the interactions between the users and sources that share fake news on social media
    - Recent works: Use Graph Neural Networks (GNNs) which counter the above disadvantage.
        - *Disadvantage:* They do not tend to uncover the hidden relationship between social media entities like - sources, articles, users, etc.

- **Iterative Graph Learning:**
    - These works learn to augment graphs such as by using end-to-end neural models optimized for the final task
        - *Disadvantage:* end-to-end learning can be prone to be task specific (edges may be created solely for achieving higher classification accuracy), rather than learning a high quality social media representation.

Proposed Solution

1. **Graph Creation Using Social Context**

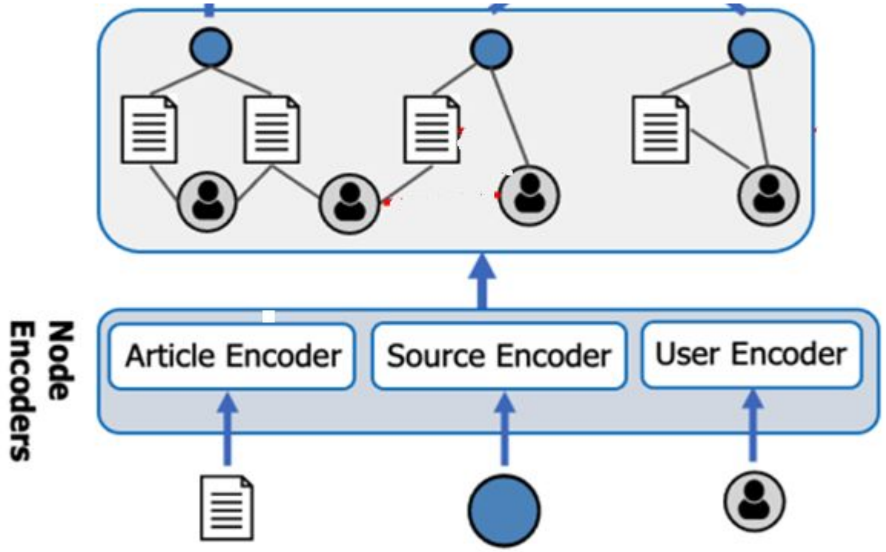

Nodes:

- **S (Sources):** Twitter and YouTube profiles embeddings
- **A (Articles):** Sentences embeddings (SBERT, RoBERTa)
- **U (Users):** Twitter profile embeddings of the Twitter users that interacted with articles and sources

Edges:

- **e(S, U)** - User U follows source S (upto 5000)
- **e(S, A)** - Article A produced from Source S (upto 300)
- **e(A, U)** - User U tweets Article A (within 3 months)
- **e(U, U)** - User follow each other

$$h_i^{l+1} = ReLU\left(\sum_{r\in R}\sum_{u\in U_r(v_i)} \frac{1}{z_{i,r}} W_r^l h_u^l\right)$$

2. Graph Embedding Using RGCN (Relational Graph Convolutional Network)

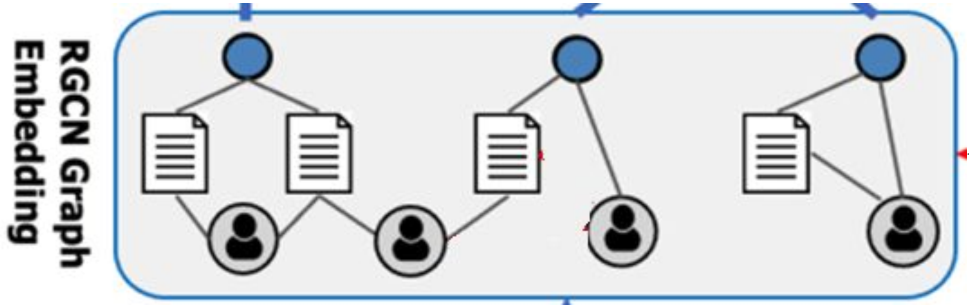

h - hidden representation
z - normalization factor
W - trainable parameter
i - index of node
U - neighbouring node representation
R - relationship
l - layer number

3. Inference Operators

Two broad types:

1. **Recommendation Style:**
    a. User-Source
    b. User-User
    c. User-Article
2. **New Edge Types:**
    a. Source-Source
    b. Article-Articles
    c. Influencers

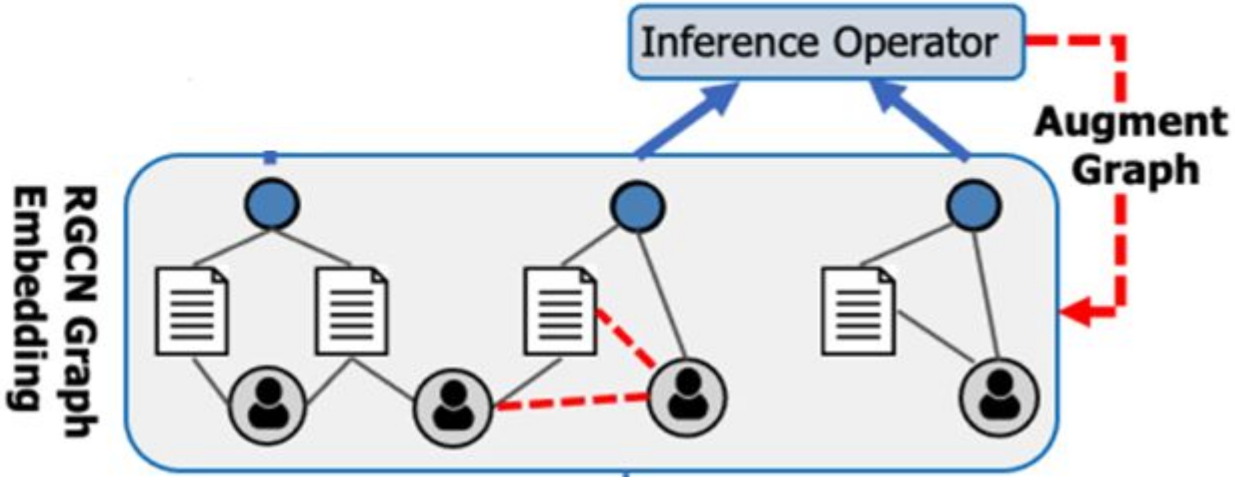

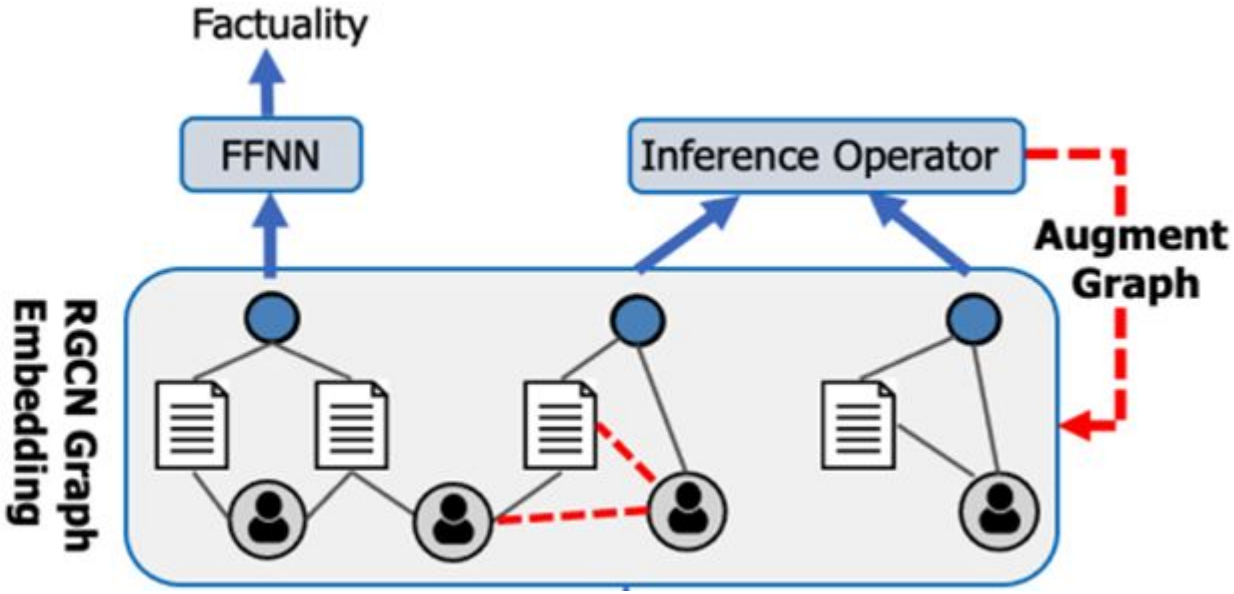

4. Joint Inference & Representation

Iteratively repeat Steps 2, 3 until convergence (determined by dev set performance for the task)

Algorithm 1 *Joint Representation Learning and Reasoning for Characterizing Information Sources*

1: **Input:** $G_0 = S, A, U$

2: **Output:** L_s (labels of sources), R-GCN parameters ϕ

3: **Initialization:** $\phi_0 \leftarrow \mathrm{L}_{final}$ over G_0

`Initial graph embedding uses Node Classification (NC)`

4: **while** not converged **do**

5: **Infer:** $G_i = G_{i-1} \cup \texttt{InferOperators}(G_{i-1}, \phi_{i-1})$

`Use inference operators to augment the graph`

6: **Learn:** $\phi_i \leftarrow L_{nc}$ over G_i

`Retrain R-GCN over new graph`

7: **end while**

8: **Final Training:** $\phi_{final} \leftarrow \mathrm{L}_{nc}$ over G_{final}

`Reset parameters and train final graph using NC`

9: **return** $L_s \leftarrow \phi_{final}$ over G_{final}

`Predict unknown sources, using the final R-GCN`

Results

Datasets Used:

- Media Bias/Fact Check dataset (Baly et al., 2018, 2020b)
  - The public dataset consists of 859 sources, each labeled on a 3-point factuality scale: low, mixed, and high
  - Final graph consists of 69,978 users, 93,191 articles, 164,034 nodes, and 7,196,808 edges
  - Used for news source classification

- Fake news article detection dataset (Nguyen et al., 2020)
  - For each article, the dataset provides its source and a list of engaged users.
  - the followers for each user were also collected, leading to a graph with 48,895 users, 442 sources, and 1,050 articles
  - Used for news article classification

Results

Model	Performance		
	Acc	Macro F1	# Edges
M1 : Majority class	52.43	22.93	-
M2 : (Baly et al., 2018)	66.45	61.08	-
M3 :(Baly et al., 2020b)	71.52	**67.25**	-
M4 : Replication of (Baly et al., 2020b)	69.38	63.63	-
M5 : Node classification (NC)	68.90	63.72	-
M6 : InfOp Best Model	**72.55**	66.89	32K

Table 1: Results on (Baly et al., 2020b). Our best model (M6) achieves a 3.17% acc improvement compared to our re-implementation (which uses the same data -> M4 vs M6), and the state-of-the-art on this dataset (which used different data - outside of train/test source labels - that was not released). Further, applying the inference operators (InfOp with 32K added edges) improves acc by 3.65% compared to Node Classification

Model	Split	Performance	
		AUC	# New Edges
FANG	90%	75.18	-
SVM	90%	75.89	-
NC	90%	83.48	-
InfOp	90%	**85.89**	10,000
FANG	70%	72.32	-
SVM	70%	59.18	-
NC	70%	73.15	-
InfOp	70%	**77.76**	10,000
FANG	50%	71.66	-
InfOp	50%	**73.88**	10,000
FANG	30%	70.36	-
InfOp	30%	**72.63**	10,000
FANG	10%	66.83	-
InfOp	10%	**67.51**	10,000

Table 2: On (Nguyen et al., 2020), we achieve the SOTA on all data splits (% of data used for training). Our model beats the strong FANG model (Nguyen et al., 2020), SVM, and the Node Classification (NC).

Summary & Conclusion

Summary:

- The work proposes a new approach for tackling fake news detection that involves continually improving social context representations using an iterative representation learning and inference framework.

- The framework involves learning an initial graph embedding and applying different inference operators to reveal hidden relationships in the graph, capturing more knowledge about the social dynamics that allow fake news to propagate.

Conclusion:

- The learnt hidden embeddings prove to be useful for various downstream applications - including user segmentation/ clustering, etc.

- The proposed approach is suitable for the social media space as the interactions between various entities continue to grow, and the approach can easily accommodate such interactions

References

- Nikhil Mehta, Maria Leonor Pacheco, and Dan Goldwasser. 2022. Tackling Fake News Detection by Continually Improving Social Context Representations using Graph Neural Networks. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1363–1380, Dublin, Ireland. Association for Computational Linguistics.

- Ramy Baly, Georgi Karadzhov, Jisun An, Haewoon Kwak, Yoan Dinkov, Ahmed Ali, James Glass, and Preslav Nakov. 2020b. What was written vs. who read it: News media profiling using text analysis and social media context. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL '20.

- Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2020. Fang: Leveraging social context for fake news detection using graph representation. In Proceedings of the 29th ACM International Conference on Information & Knowledge Management, pages 1165–1174.

Thank You !