CIKM 2021
1-5 NOVEMBER

ANEMONE: Graph Anomaly Detection with Multi-Scale Contrastive Learning

Ming Jin
Monash University
ming.jin@monash.edu

Yixin Liu
Monash University
yixin.liu@monash.edu

Yu Zheng
La Trobe University
yu.zheng@latrobe.edu.au

Lianhua Chi
La Trobe University
l.chi@latrobe.edu.au

Yuan-Fang Li
Monash University
yuanfang.li@monash.edu

Shirui Pan
Monash University
shirui.pan@monash.edu

Presenter : Ming Jin

MONASH University

LA TROBE UNIVERSITY

Graph Anomaly Detection

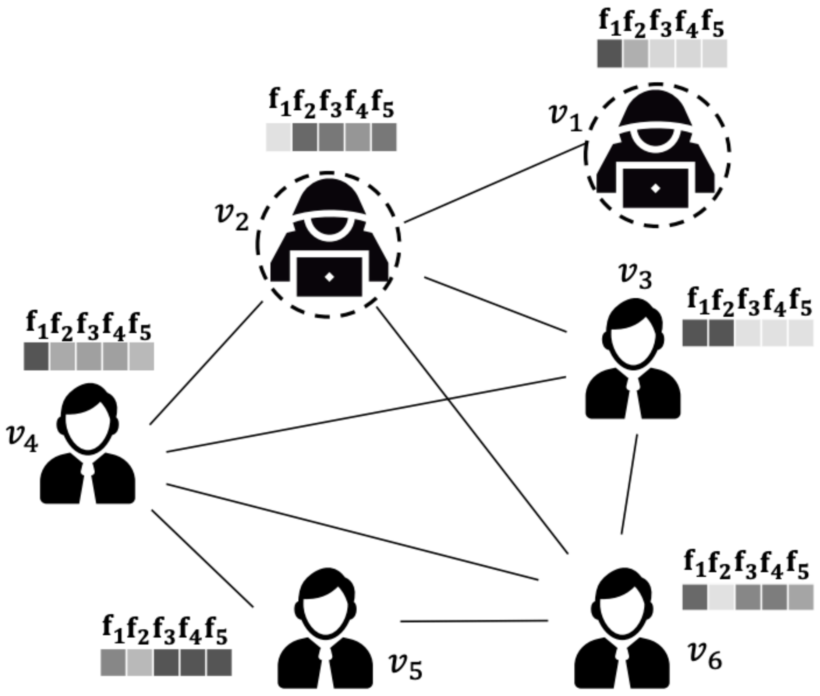

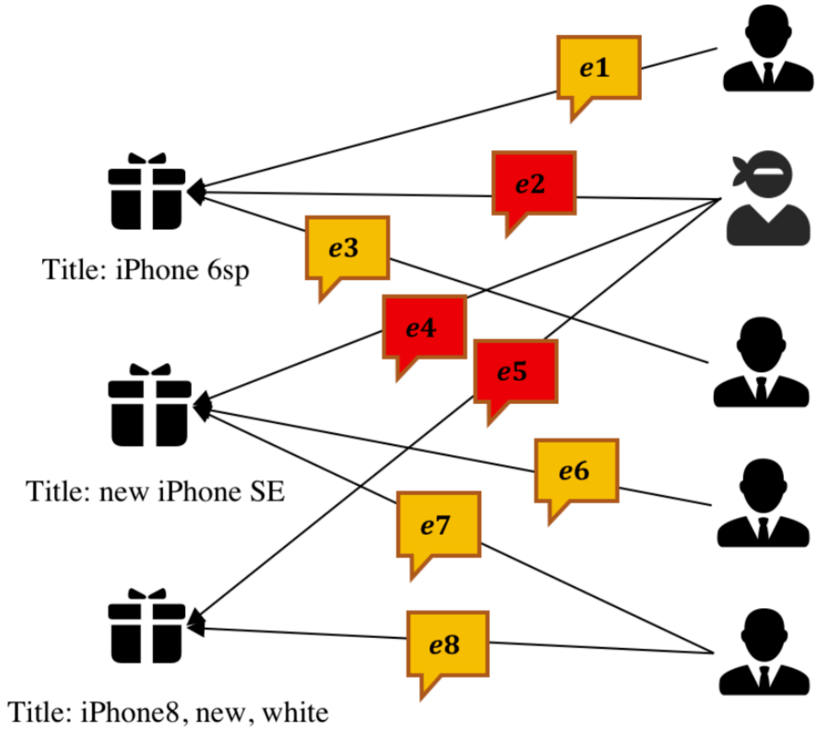

Graph Self-Supervised Learning

(Semi-)Supervised Graph Learning

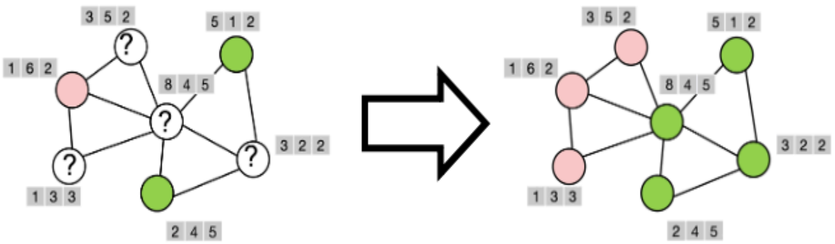

Input: A partially labeled attributed graph

Output: Inferring the labels of unlabeled nodes

Graph Self-Supervised Learning

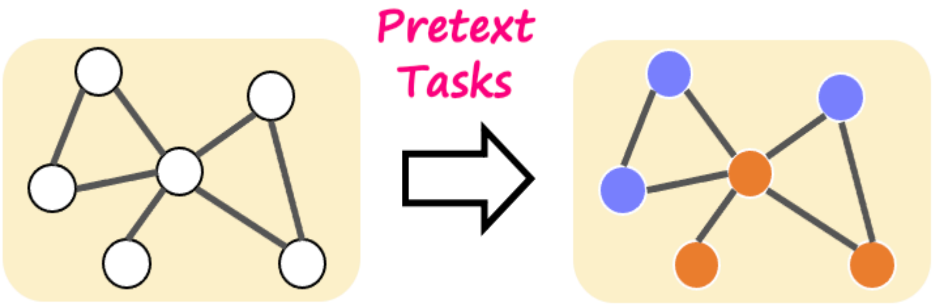

Input: An unlabeled attributed graph

Output of downstream task: Inferring the labels of unlabeled nodes

Contrastive Graph Anomaly Detection

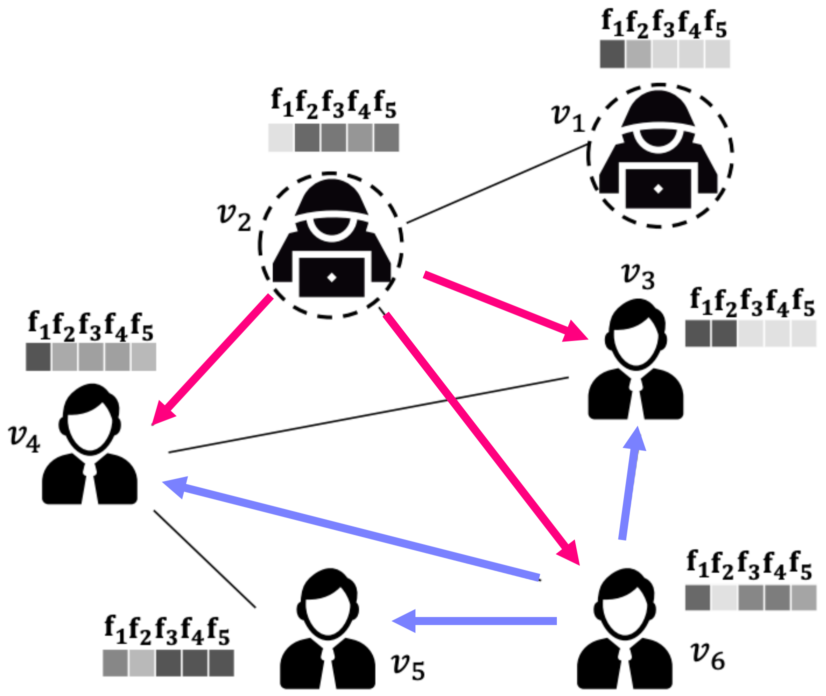

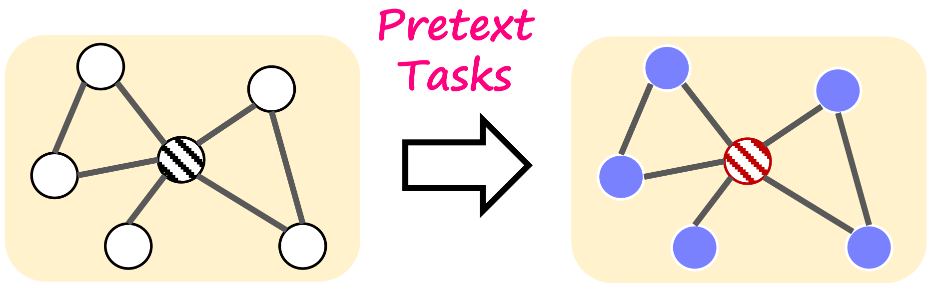

Input: An unlabeled attributed graph

Output of downstream task: Inferring node anomaly scores

The **mismatch** between a node and its surrounding contextual information reflects its abnormality

ANEMONE

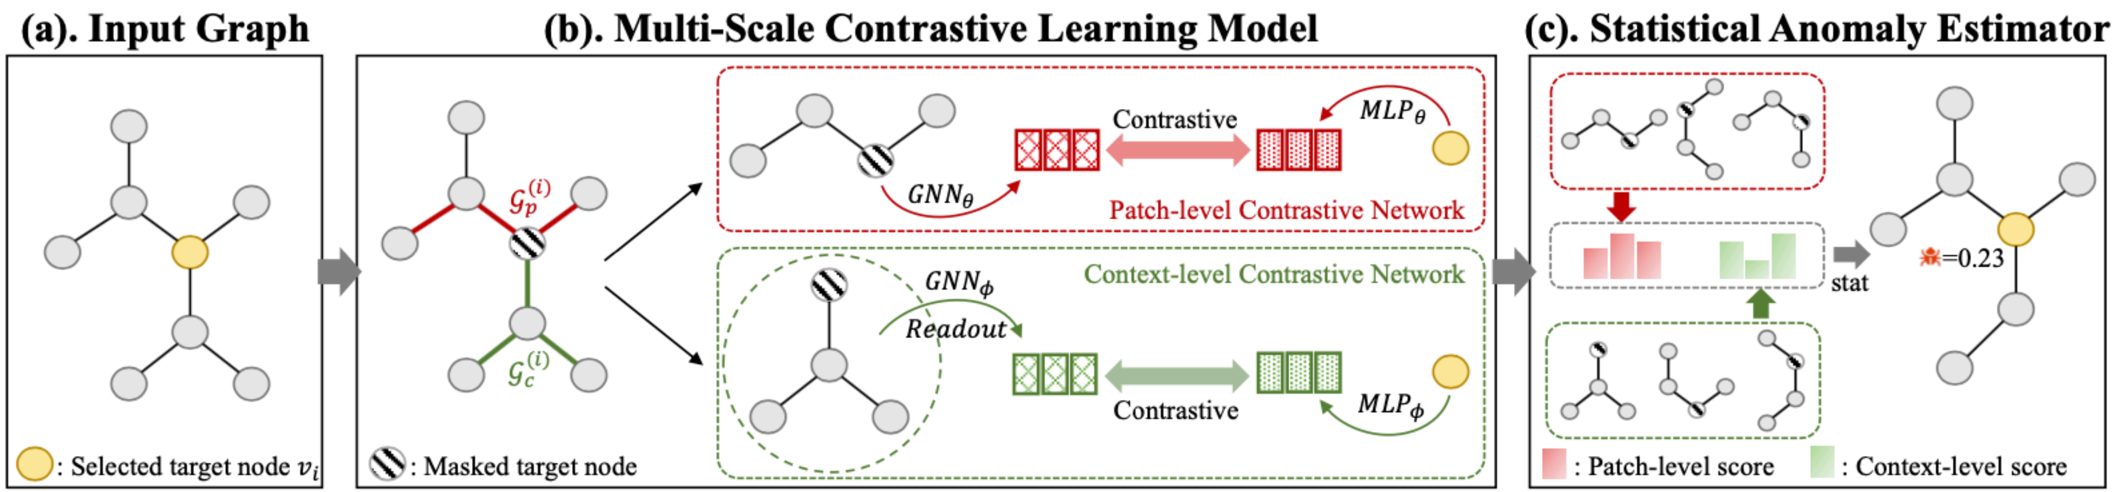

- Given a target node, **two contrastive pretext tasks** are created to predict the anomaly score of this node
- A **patch-level** task contrasts the embedding of a masked target node with the mapping of its raw information
- A **context-level** task contrasts a target node with the contextual embedding obtained from its surrounding neighbors
- Finally, the abnormality of a node is **statistically estimated** by referring two contrastive scores

Patch-Level Contrastiveness

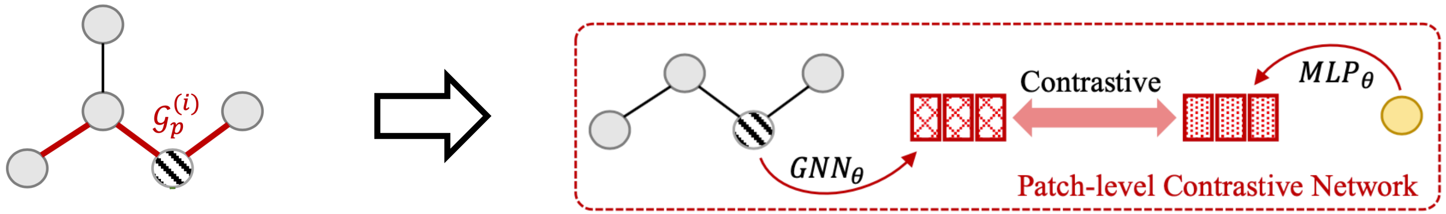

- Firstly, the masked target node embedding is obtained via a GCN parameterized by θ:

$$\mathbf{H}_p^{(i)} = GNN_\theta\left(\mathcal{G}_p^{(i)}\right) = GCN\left(\mathbf{A}_p^{(i)}, \mathbf{X}_p^{(i)}; \mathbf{\Theta}\right) = \sigma\left(\widetilde{\mathbf{D}_p^{(i)}}^{-\frac{1}{2}} \widetilde{\mathbf{A}_p^{(i)}} \widetilde{\mathbf{D}_p^{(i)}}^{-\frac{1}{2}} \mathbf{X}_p^{(i)} \mathbf{\Theta}\right),$$

- Then, the target node representation is calculated via a MLP:

$$\mathbf{z}_p^{(i)} = MLP_\theta\left(\mathbf{x}^{(i)}\right) = \sigma\left(\mathbf{x}^{(i)} \mathbf{\Theta}\right)$$

- Finally, we maximize their agreement based on the assumption that most nodes in a graph is **NOT** anomalies

$$\mathcal{L}_p = -\frac{1}{2n} \sum_{i=1}^{n} \left(log\left(s_p^{(i)}\right) + log\left(1 - \tilde{s}_p^{(i)}\right)\right)$$

$$s_p^{(i)} = Bilinear\left(\mathbf{h}_p^{(i)}, \mathbf{z}_p^{(i)}\right) = \sigma\left(\mathbf{h}_p^{(i)} \mathbf{W}_p \mathbf{z}_p^{(i)\top}\right)$$

$$\tilde{s}_p^{(i)} = Bilinear\left(\mathbf{h}_p^{(j)}, \mathbf{z}_p^{(i)}\right) = \sigma\left(\mathbf{h}_p^{(j)} \mathbf{W}_p \mathbf{z}_p^{(i)\top}\right)$$

Context-Level Contrastiveness

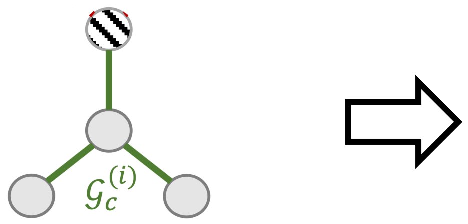

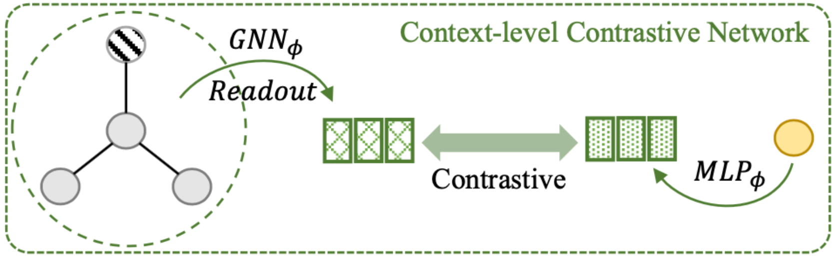

- Firstly, the contextual embedding of a target node is obtained via a GCN parameterized by ϕ:

$$\mathbf{H}_c^{(i)} = GNN_\phi\left(\mathcal{G}_c^{(i)}\right) = \sigma\left(\widetilde{\mathbf{D}_c^{(i)}}^{-\frac{1}{2}}\widetilde{\mathbf{A}_c^{(i)}}\widetilde{\mathbf{D}_c^{(i)}}^{-\frac{1}{2}}\mathbf{X}_c^{(i)}\mathbf{\Phi}\right) \qquad \mathbf{h}_c^{(i)} = readout\left(\mathbf{H}_c^{(i)}\right) = \frac{1}{K}\sum_{j=1}^{K}\mathbf{H}_c^{(i)}[j,:]$$

- Then, the target node representation is calculated in the same way but with another MLP

- Finally, we maximize their mutual information with another estimator:

$$\mathcal{L}_c = -\frac{1}{2n}\sum_{i=1}^{n}\left(log\left(s_c^{(i)}\right) + log\left(1 - \tilde{s}_c^{(i)}\right)\right)$$

Thus, our overall objective is minimizing this contrastive loss: $\mathcal{L} = \alpha\mathcal{L}_c + (1-\alpha)\mathcal{L}_p$

Statistical Anomaly Estimator

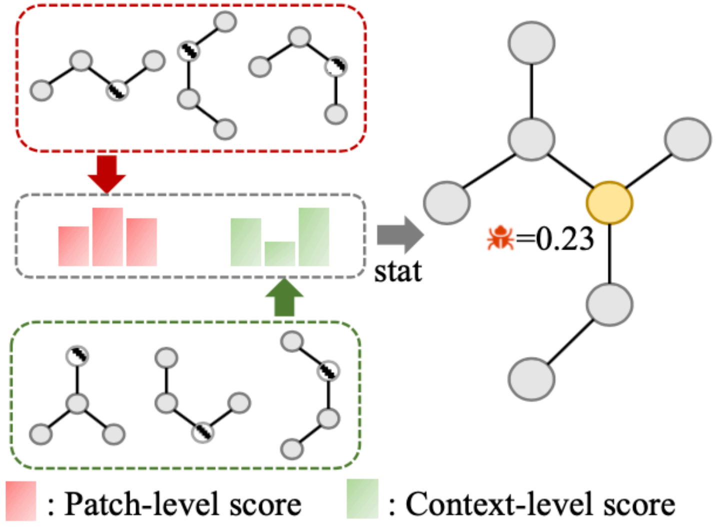

- For a target node v_i, we generate R ego-nets for patch- and context-level contrastive learning

$$[\underbrace{s_{p,1}^{(i)}, \cdots, s_{p,R}^{(i)}, s_{c,1}^{(i)}, \cdots, s_{c,R}^{(i)}}_{\text{Positive scores}}, \underbrace{\tilde{s}_{p,1}^{(i)}, \cdots, \tilde{s}_{p,R}^{(i)}, \tilde{s}_{c,1}^{(i)}, \cdots, \tilde{s}_{c,R}^{(i)}}_{\text{Negative scores}}]$$

- We denote the base anomaly score as follows:

$$b_{view,j}^{(i)} = \tilde{s}_{view,j}^{(i)} - s_{view,j}^{(i)}$$

where the subscript "view" represents "p" or "c" and $j \in [1, \cdots, R]$

- The final anomaly score is calculated via:

$$y^{(i)} = \alpha y_c^{(i)} + (1-\alpha) y_p^{(i)} \qquad y_{view}^{(i)} = \bar{b}_{view}^{(i)} + \sqrt{\sum_{j=1}^{R} \left(b_{view,j}^{(i)} - \bar{b}_{view}^{(i)}\right)^2 / R} \qquad \bar{b}_{view}^{(i)} = \sum_{j=1}^{R} b_{view,j}^{(i)} / R$$

Experiments

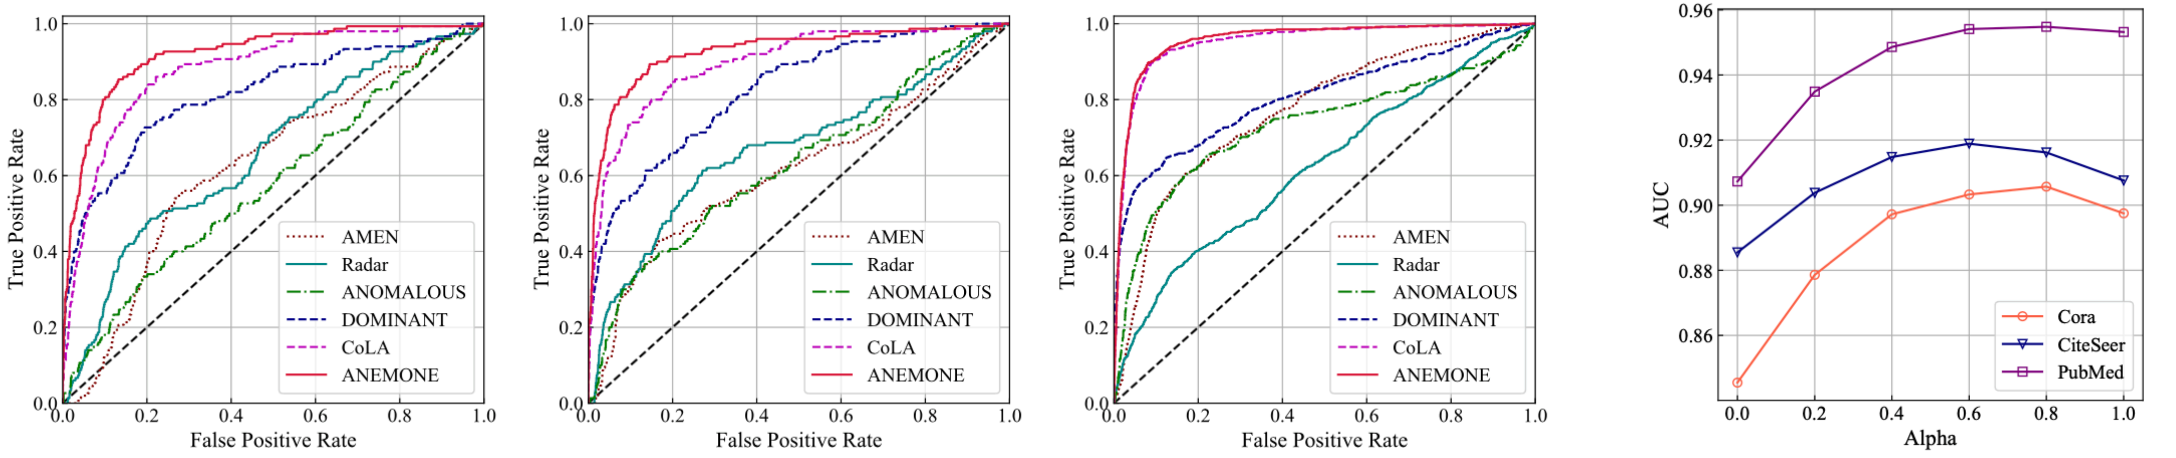

(a) ROC curve of Cora. (b) ROC curve of CiteSeer. (c) ROC curve of PubMed. (d) Trade-off parameter α w.r.t. AUC values.

Table 1: Basic statistics of the three datasets.

Datasets	# Nodes	# Edges	# Attributes	# Anomalies
Cora	2,708	5,429	1,433	150
CiteSeer	3,327	4,732	3,703	150
PubMed	19,717	44,338	500	600

Table 2: AUC of ANEMONE, its competitors and variants.

Methods	Cora	CiteSeer	PubMed
AMEN	0.6266	0.6154	0.7713
Radar	0.6587	0.6709	0.6233
ANOMALOUS	0.5770	0.6307	0.7316
DOMINANT	0.8155	0.8251	0.8081
CoLA	0.8779	0.8968	0.9512
CoLA_{stat}	0.8869	0.9047	0.9532
ANEMONE_{mean}	0.8963	0.9066	0.9524
ANEMONE_{std}	0.5402	0.7077	0.7440
ANEMONE	**0.9057**	**0.9189**	**0.9548**

References

- Li, A., Qin, Z., Liu, R., Yang, Y., & Li, D. (2019, November). Spam review detection with graph convolutional networks. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (pp. 2703-2711).
- Ding, K., Li, J., Bhanushali, R., & Liu, H. (2019, May). Deep anomaly detection on attributed networks. In Proceedings of the 2019 SIAM International Conference on Data Mining (pp. 594-602). Society for Industrial and Applied Mathematics.
- Tang, J. (2021). *Graph Neural Networks and Self-supervised Learning*. Jie Tang's Homepage. https://keg.cs.tsinghua.edu.cn/jietang/publications/GNN-Pre-train.pdf
- Efros, A. (2021). *Self-Supervision for Learning from the Bottom Up*. ICLR 2021 Invited Talks. https://iclr.cc/virtual/2021/invited-talk/3720