SpecAE: Spectral AutoEncoder for Anomaly Detection in Attributed Networks

Yuening Li¹, Xiao Huang¹, Jundong Li^{2,3}, Mengnan Du¹, Na Zou⁴

¹Department of Computer Science and Engineering, Texas A&M University

²Department of Electrical and Computer Engineering, University of Virginia

³Department of Computer Science & School of Data Science, University of Virginia

⁴Department of Industrial & Systems Engineering, Texas A&M University

{yueningl,xhuang,dumengnan,nzou1}@tamu.edu,jl6qk@virginia.edu

ABSTRACT

Anomaly detection in attributed networks (instance-to-instance dependencies and interactions are available) has various applications such as monitoring suspicious accounts in social media and financial fraud in transaction networks. However, it remains a challenging task since the definition of anomaly becomes more complicated and topological structures are heterogeneous with nodal attributes. In this paper, we propose a spectral convolution and deconvolution based framework - SpecAE, to project the attributed network into a tailored space to detect global and community anomalies. SpecAE leverages Laplacian sharpening to amplify the distances between representations of anomalies and the ones of the majority. The learned representations along with reconstruction errors are combined with a density estimation model to perform the detection. Experiments on real-world datasets demonstrate the effectiveness of the proposed SpecAE.

CCS CONCEPTS

• **Computing methodologies** → **Anomaly detection; Neural networks; Mixture models.**

KEYWORDS

Anomaly Detection, Network Embedding, Neural Networks

ACM Reference Format:

Yuening Li¹, Xiao Huang¹, Jundong Li^{2,3}, Mengnan Du¹, Na Zou⁴. 2019. SpecAE: Spectral AutoEncoder for Anomaly Detection in Attributed Networks. In *The 28th ACM International Conference on Information and Knowledge Management (CIKM '19)*, November 3–7, 2019, Beijing, China. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3357384.3358074>

1 INTRODUCTION

Anomaly detection targets at identifying the rare instances that behave significantly different from the majority instances. Conventional detection algorithms are mainly based on the assumption that instances are independent and identically distributed (i.i.d.). But in many practical scenarios, instances are often explicitly or

implicitly connected with each other, resulting in the so-called *attributed networks*. The unique data characteristics of attributed networks bring several challenges to anomaly detection. The definition of anomaly becomes more complicated and obscure. Apart from anomalous nodes whose nodal attributes are rather different from the majority reference nodes from a global perspective, nodes with nodal attributes deviate remarkably from their communities are also considered to be anomalies.

To handle the topological structures and the heterogeneity challenge, in network analysis, a widely-used and effective approach is to embed all information in attributed network into unified low-dimensional node representations. However, it would achieve sub-optimal performance when directly applied to the anomaly detection. To achieve the joint embedding, most existing network embedding algorithms and recent joint embedding based anomaly detection methods mainly rely on the homophily hypothesis, which implies that connected nodes tend to have similar nodal attributes [8, 10]. Based on this assumption, topological relations could be introduced via involving regularization terms to minimize the distance between embedding representations of linked nodes [7, 15]. Another effective way to encode networks is to employ graph convolutional networks (GCN) [4]. It could be considered as a special form of Laplacian smoothing that learns nodes' representations from their neighbors [3, 6]. The homophily hypothesis or smoothing operations are not in line with anomaly detection. They might over-smooth the node representations, and make anomalous nodes less distinguishable from the majority within the community. For example, malicious users would have completely different nodal attributes than their friends. Customers who have written fake reviews might purchase the same products as the normal ones. Thus, existing joint learning models in network analysis could not be directly applied to anomaly detection.

To bridge the gap, we investigate the problem of anomaly detection in attributed networks. Two specific research questions are studied. 1) How to formally define anomalies in attributed network? 2) How to perform joint learning on topological structures and nodal attributes, while being in line with the anomaly detection, such as getting rid of the restriction of homophily and over-smooth problems? Following these research perspectives, our contributions could be summarized as follows.

- We propose a graph convolution and deconvolution based framework, for anomaly detection in attributed networks.
- We develop a tailored model to project the attributed network into a special space for global anomalies and community anomalies. It leverages Laplacian sharpening to amplify the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

CIKM '19, November 3–7, 2019, Beijing, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6976-3/19/11...\$15.00

<https://doi.org/10.1145/3357384.3358074>

Figure 1: A spectral autoencoder (SpecAE) based anomaly detection approach for attributed networks.

distances between embedding representations of anomalies and the ones of other nodes within the communities.

- We conduct evaluations on real-world datasets, demonstrating the effectiveness of SpecAE.

2 PROBLEM STATEMENT

Let $\mathcal{G} = \{A, X\}$ be an input attributed network, with n nodes interconnected by a network. $A \in \mathbb{R}^{n \times n}$ denotes the corresponding adjacency matrix. Each node is associated with a set of m -dimensional nodal attributes. We collect all these nodal attribute as $X \in \mathbb{R}^{n \times m}$.

Definition 1 (Global Anomaly) It refers to a node whose nodal attributes are rare and significantly different from the nodal attributes of majority nodes, from a global perspective.

Definition 2 (Community Anomaly) It is defined as a node with nodal attributes that significantly deviate from the node's neighbors' nodal attributes, based on the topological network structure.

Given the aforementioned definitions, we formally define the anomaly detection in attributed networks problem as follows. Given a set of nodes $\mathcal{V}$ connected by an attribute network $\mathcal{G}$, we target at identifying the global anomalies and community anomalies in $\mathcal{V}$.

3 PROPOSED SPECTRAL AUTOENCODER

We propose a Spectral autoencoder based anomaly detection framework - SpecAE, for attributed networks. The pipeline of SpecAE is illustrated in Fig.1. Given an attributed network $\mathcal{G}$, we leverage a tailored Spectral autoencoder to jointly embed nodal attributes and relations into a carefully-designed space Z . To detect the global anomalies, we apply an autoencoder to all nodal attributes X to learn embedding representations Z_X as well as the reconstruction errors Z_{Xerror} . In such a way, we could globally compare all nodes' attributes. To detect the community anomalies, we design novel graph convolutional encoder and deconvolutional decoder networks to learn nodes' community representations Z_G , based on each node's neighbors. The corresponding reconstruction errors are denoted as Z_{Gerror} . Later on, based on the tailored joint representations Z , we estimate the suspicious level of each node u by measuring its embedding representation $Z_{i,:}$'s energy in the Gaussian Mixture Model.

3.1 Tailored Embedding Spaces for Global Anomalies and Community Anomalies

First, based on the definition of global anomaly, we employ an autoencoder to embed nodal attributes X to learn the first type of tailored representations, i.e., Z_X and Z_{Xerror} . The encoding and

decoding processes are achieved via:

$$Z_X = \sigma(b_e + XW_e), \hat{X} = \sigma(b_d + Z_XW_d), \quad (1)$$

where $\hat{X}$ denotes the reconstructed nodal attributes. Since anomalies are harder to reconstruct than normal nodes, we also include the reconstruction errors $Z_{Xerror} = \text{dis}(X, \hat{X})$, where the operation $\text{dis}(\cdot, \cdot)$ denotes a series of the distance measuring matrices such as the Euclidean distance and the cosine distance. In such a way, global anomalies would tend to have their representations in Z_X being different from the majority, and Z_{Xerror} being large.

Second, to identify community anomalies, we need to jointly consider A and X . Based on their topological dependencies, we develop a novel graph convolutional encoder and graph deconvolutional decoder to learn the second type of tailored representations, i.e., Z_G and Z_{Gerror} . Details are introduced as follows.

3.2 Graph Convolution and Deconvolution

Our goal is to learn nodes' community representations Z_G , which describe the expected nodal attributes according to their neighbors.

A straightforward solution is to apply GCN to embed $\mathcal{G}$. From a formulation perspective [6], GCN could be treated as a special form of *Laplacian Smoothing*. The *Laplacian smoothing* on each channel of the input features can be written as:

$$Y = (I - \gamma \tilde{D}^{-1}(\tilde{D} - \tilde{A}))X, \quad (2)$$

where $0 < \gamma \leq 1$ is a parameter which controls the weighting between the features of the current node and the features of its neighbors, with a normalization trick as:

$$I + D^{-\frac{1}{2}}AD^{-\frac{1}{2}} \rightarrow \tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}, \quad (3)$$

where $\tilde{A} := A + I$ and $\tilde{D} := \sum_j \tilde{A}_{ij}$.

Based on the definition of convolution in graph signals on the spectral domain, to generate a new matrix Y from X by applying the graph convolution as spectral filters, we have:

$$Y = \tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}X, \quad (4)$$

which is a special form of *Laplacian Smoothing* when $\gamma = 1$ after replacing the *normalized graph Laplacian* $L := \tilde{D}^{-1}\tilde{A}$ with the *symmetric normalized graph Laplacian* $L := \tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}$.

However, GCN does not contain nodal attributes reconstruction procedure, which is useful in the anomaly detection task with three reasons. i) Training transformation functions base on one-class observations (only normal instances available) without the reconstruction error may easily lead the objective function converges to a local optimal when all of the node representations collapse into a small area (one point in extreme cases) in the latent space. ii) Repeatedly applying *Laplacian smoothing* might cause the nodal attributes over-mixed with their neighbors and make them indistinguishable, since it will be more difficult to identify each individual instance after the smoothness operation which has weakened their uniqueness of the original attributes. iii) Reconstruction error usually contains useful information as indicators for anomaly detection. For example, in [2], the reconstruction error could be directly used as an anomaly score to rank the anomalous degree of nodes.

Thus, we decide to design the graph decoding (graph deconvolution) from the smoothed features as a complementary inverse

process of graph convolution. Inspired by the digital images field, *sharpening* is an inverse process of blurring/smoothing [9]. Comparing with the smoothing process, which is done in the spatial domain by pixel averaging in neighbors, sharpening is to find the difference by the neighbors, done by spatial differentiation. A Laplacian operator is to restore fine details of an image which has been smoothed or blurred. After Laplacian sharpening, we can reconstruct the nodal attributes from the fusion features caused by the graph convolution process.

If we replace the original GCN function in Eq. (4) as a general *Laplacian smoothing*, we will have:

$$\mathbf{Y} = \mathbf{X} - \alpha(\mathbf{X} - \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X}) = (1 - \alpha)\mathbf{X} + \alpha \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X}. \quad (5)$$

Generalizing the above definition of graph convolution with adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and the nodal attributes $\mathbf{X} \in \mathbb{R}^{n \times m}$ of n instances and m input channels, the propagation rule of the convolution layer can be written as:

$$\text{Conv}(\mathbf{X}, \mathbf{A}) = \sigma \left((1 - \alpha)\mathbf{X} + \alpha \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X} \right) \mathbf{W}_f, \quad (6)$$

where $\mathbf{W}_f$ is the trainable weight matrix in the convolution layer, σ is the activation function, e.g., $\text{ReLU}(\cdot)$. The parameter α is to control the weighting between the features of the current nodal attributes and the features of its neighbors.

Contrary to *Laplacian smoothing*, we compute the new features of nodal attributes through sharpening the features with their neighbors in order to reconstruct the features from the smoothed results. To magnify the difference between the current node and its neighbors, we will have:

$$\mathbf{Y} = \mathbf{X} + \alpha(\mathbf{X} - \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X}) = (1 + \alpha)\mathbf{X} - \alpha \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X}. \quad (7)$$

Given the above definition of deconvolution to an attributed network $\mathcal{G}$, with adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and the nodal attributes $\mathbf{Z} \in \mathbb{R}^{n \times m}$ of n instances and m input channels which after the convolution process, the propagation rule of the deconvolution layer is:

$$\text{Deconv}(\mathbf{Z}, \mathbf{A}) = \sigma \left((1 + \alpha)\mathbf{Z} - \alpha \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{Z} \right) \mathbf{W}_g, \quad (8)$$

where $\mathbf{W}_g$ is the trainable weight matrix in the deconvolution layer. After the sharpening process, we can reconstruct the original attributes from the smoothed features.

Given nodal attributes $\mathbf{X}$ with adjacency matrix $\mathbf{A}$, the compression network computes $\mathbf{Z}$ as:

$$f(\mathbf{Z}_{\mathcal{G}} | \mathbf{X}, \mathbf{A}) = \prod_{i=1}^n f(\mathbf{Z}_{\mathcal{G}_{i,:}} | \mathbf{X}_{i,:}, \mathbf{A}), \quad (9)$$

$$\text{with } f(\mathbf{Z}_{\mathcal{G}_{i,:}} | \mathbf{X}, \mathbf{A}) = \mathcal{N}(\mathbf{Z}_{\mathcal{G}_{i,:}} | \mu_{i,:}, \text{diag}(\sigma_{i,:}^2));$$

$$g(\tilde{\mathbf{X}} | \mathbf{Z}_{\mathcal{G}}, \mathbf{A}) = \prod_{i=1}^n g(\tilde{\mathbf{X}}_{i,:} | \mathbf{Z}_{\mathcal{G}_{i,:}}, \mathbf{A}), \quad (10)$$

$$\text{with } g(\tilde{\mathbf{X}}_{i,:} | \mathbf{Z}_{\mathcal{G}_{i,:}}, \mathbf{A}) = \text{Deconv}(\mathbf{Z}_{\mathcal{G}_{i,:}}, \mathbf{A});$$

$$\mathbf{Z}_{\mathcal{G}_{\text{error}}} = \text{dis}(\mathbf{X}, \tilde{\mathbf{X}}), \quad (11)$$

$$\mathbf{Z} = [\mathbf{Z}_{\mathcal{X}}; \mathbf{Z}_{\mathcal{G}}; \mathbf{Z}_{\mathcal{X}_{\text{error}}}; \mathbf{Z}_{\mathcal{G}_{\text{error}}}], \quad (12)$$

where the $\mu_{i,:} := \text{Conv}_{\mu}(\mathbf{X}, \mathbf{A})$ denote the mean vector in Eq. (7), similarly, $\log \sigma_{i,:} := \text{Conv}_{\sigma}(\mathbf{X}, \mathbf{A})$. The latent variable $\mathbf{Z}$ of each node is shown in Eq. (12), which concatenates from four components. $\mathbf{Z}_{\mathcal{G}}$ is learned by the graph convolutional encoder which

including the neighbor attribution with the topological relations. $\mathbf{Z}_{\mathcal{G}_{\text{error}}}$ is the reconstruction of the nodal attributes $\mathbf{X}$ base on the fusion representation $\mathbf{Z}_{\mathcal{G}}$ and the adjacency matrix $\mathbf{A}$.

3.3 Anomaly Detection via Density Estimation

Given the learned embedding representation $\mathbf{Z}$ of N samples and their soft mixture-component membership prediction $\hat{\mathbf{y}}$ (estimated from a softmax layer based on the low-dimensional representation $\mathbf{Z}$), where $\hat{\mathbf{y}}$ is a K -dimensional vector and K is the number of mixture components in Gaussian Mixture Model (GMM), we can estimate the mixture probability $\hat{\phi}_k$, the mean value $\hat{\mu}_k$, the covariance vector $\hat{\Sigma}_k$ for component k in GMM respectively [16].

The sample energy can be inferred by:

$$E(\mathbf{z}) = -\log \left(\sum_{k=1}^K \hat{\phi}_k \frac{\exp -\frac{1}{2}(\mathbf{z} - \hat{\mu}_k)^T \hat{\Sigma}_k^{-1} (\mathbf{z} - \hat{\mu}_k)}{\sqrt{|2\pi \hat{\Sigma}_k|}} \right), \quad (13)$$

where we estimate the parameters in GMM as follows:

$$\hat{\mathbf{y}} = \text{softmax}(\mathbf{z}), \quad \hat{\phi}_k = \sum_{i=1}^N \frac{\hat{y}_{ik}}{N}, \quad (14)$$

$$\hat{\mu}_k = \frac{\sum_{i=1}^N \hat{y}_{ik} \mathbf{z}_i}{\sum_{i=1}^N \hat{y}_{ik}}, \quad \hat{\Sigma}_k = \frac{\sum_{i=1}^N \hat{y}_{ik} (\mathbf{z}_i - \hat{\mu}_k)(\mathbf{z}_i - \hat{\mu}_k)^T}{\sum_{i=1}^N \hat{y}_{ik}}. \quad (15)$$

3.4 Objective Function

Given an input attributed network $\mathcal{G} = \{\mathbf{A}, \mathbf{X}\}$, the objective function is constructed as follow:

$$\begin{aligned} J(\mathbf{W}) = & \mathbb{E}[\text{dis}(\mathbf{X}, \tilde{\mathbf{X}})] + \mathbb{E}[\text{dis}(\mathbf{X}, \tilde{\mathbf{X}})] + \lambda_1 \mathbb{E}(E(\mathbf{Z})) \\ & + \lambda_2 \sum_{i=1}^K \sum_{j=1}^d \frac{1}{\hat{\Sigma}_{ij}} - \mathcal{KL}[f(\mathbf{Z}_{\mathcal{G}} | \mathbf{X}, \mathbf{A}) || g(\mathbf{Z}_{\mathcal{G}})] \\ & + \mathbb{E}_{f(\mathbf{Z}_{\mathcal{G}} | \mathbf{X}, \mathbf{A})} \log(g(\tilde{\mathbf{X}} | \mathbf{Z}_{\mathcal{G}}, \mathbf{A})). \end{aligned} \quad (16)$$

The objective function includes four components:

- $\mathbb{E}[\text{dis}(\mathbf{X}, \tilde{\mathbf{X}})]$, $\mathbb{E}[\text{dis}(\mathbf{X}, \tilde{\mathbf{X}})]$ are the loss function that characterizes the reconstruction error.
- $E(\mathbf{Z})$ denotes the sample energy of the GMM estimation in Eq. (13). Through minimizing the sample energy, we maximize the likelihood of non-anomalous samples, and predict samples with top-K high energy as anomalies.
- To avoid trivial solutions when the diagonal entries of covariance matrices degenerate to 0, we penalize small values by the third component as a regularizer.
- We optimize the variational lower bound as the last two terms. $\mathcal{KL}[f(\cdot) || g(\cdot)]$ is the Kullback-Leibler divergence between $f(\cdot)$ and $g(\cdot)$. $g(\cdot)$ is a Gaussian prior: $g(\mathbf{Z}_{\mathcal{G}}) = \prod_{i=1}^n g(\mathbf{Z}_{\mathcal{G}_{i,:}}) = \mathcal{N}(\mathbf{Z}_{\mathcal{G}_{i,:}} | \mathbf{0}, \mathbf{I})$.

4 EXPERIMENTS

In this section, we evaluate the performance of our model with experiments on real-world datasets (Cora, Pubmed, and PolBog [11]) with five baseline methods, and ablation analysis with a case study, to verify the effectiveness of the proposed approach.

4.1 Effectiveness evaluation

For both metrics of Accuracy@K (Tab. 1) and ROC-AUC (Fig. 2), SpecAE consistently achieves good performance. In addition, we

Figure 2: ROC/AUC on Cora and Pubmed (from left to right).

K	Cora [14]				Pubmed [14]			
	5	10	15	20	5	10	15	20
LOF [1]	85.60	80.91	76.62	72.01	85.71	81.27	76.89	72.71
OC SVM [13]	85.45	80.76	76.92	72.45	85.54	81.03	76.47	72.27
Deep SVDD [12]	87.89	85.49	81.50	77.25	87.93	84.63	80.96	77.07
Radar [5]	87.70	85.27	82.35	78.58	86.02	81.90	78.03	74.06
GCN [4]	85.23	83.38	76.70	71.05	86.95	83.57	80.11	76.41
SpecAE	94.31	90.81	86.67	82.05	94.98	90.50	85.96	81.52

Table 1: Accuracy@K on two datasets.

Metric	Accuracy	Precision	Recall	F1	AUC
SpecAE-S	54.51	12.54	59.63	20.72	58.45
SpecAE-N	90.51	52.40	52.59	52.50	73.65
SpecAE-nr	82.24	11.07	11.11	11.09	50.53
SpecAE $\alpha = 1$	90.36	51.66	51.85	51.76	73.96
SpecAE $\alpha = 0.7$	91.93	54.98	55.19	55.08	77.15

Table 2: Ablation and hyperparameter analysis.

make the following observations. (1) In general, the proposed SpecAE outperforms all of the baseline methods. Comparing with the methods which only utilize uni-modal information, our proposed approach achieves better performance which validates the importance of applying heterogeneous multi-modal sources. (2) The performance of deep models are superior to the conventional anomaly detection methods in shallow structures. It verifies the effectiveness of extracting features to represent the nodal attributes on attributes networks in deep structures which can extract features in a more effective way. (3) SpecAE shows the importance of enlarging the model capacity by introducing the mutual interactions between nodal attributes with linkage connections and enabling topological relations to contain additional information. As can be observed, Spectral autoencoder can achieve better performance on anomaly detection application than pure GCN based model, which justifies the advantage of our proposed model that overcomes the drawbacks of the GCN, and adopts it into an anomaly detection scenario.

4.2 Ablation Analysis And A Case Study

We conduct ablation studies on Cora to demonstrate the importance of individual components in SpecAE. We perform ablation analysis comparing SpecAE with four alternatives. (i) SpecAE-S: SpecAE without the representations for community anomaly detection, i.e., Z_G and $Z_{G_{error}}$; (ii) SpecAE-N: SpecAE without the representations for global anomaly detection, i.e., Z_X and $Z_{X_{error}}$; (iii) SpecAE-nr: SpecAE without the reconstruction components $Z_{X_{error}}$ and $Z_{G_{error}}$; and (iv) training with an extreme value hyper-parameter $\alpha = 1$ (short-circuit the self features in the graph convolution and deconvolution). Results on are shown in Tab. 2, indicating that the comprehensive method is superior to the variants, since it outperforms most of the ablation test settings where one component

(Anomaly).com	Abnormal Keywords
thismodernworld	siberians, fenno, scandinavia, ultraviolet, primates, colder
directorblue.blogspot	boiler, inflatable, choppers, streamline, rooftop, jettison, sheathed
balloon-juice	recapturing, wooly, buffy, talons, snarled, tails, dismemberment
fringeblog	mosquitos, nerf, redwoods, mammoths, fungi, snail, hawked
nomayo.mu.nu	pancakes, cheese, stewed, pork, brunch

Table 3: Keywords from anomalous nodes on PolBlog.

has been removed or replaced at a time from the full system and performed re-training and re-testing.

We also conduct a case study using the PolBlog [11] dataset. The attributes of each node represent words appeared in each blog based on a given word dictionary. As shown in the Tab. 3, five sampled anomalies (nodes) are listed for more specific explanations. The abnormal keywords refer to the key information which appears in the anomaly bloggers, yet has low frequency and lacks political flavor in a global perspective among all bloggers. To be more specific, some of the posts actually focus on other fields like climates (thismodernworld), food (nomayo.mu.nu), constructions and tools (directorblue.blogspot), huntings and creatures (balloon-juice, fringeblog). Comparing with the most popular political topics, these selected keywords in the abnormal bloggers are less relevance to the political events.

ACKNOWLEDGEMENT

The work is, in part, supported by National Science Foundation (#IIS-1657196, #IIS-1750074, #CNS-1816497).

REFERENCES

- [1] Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. 2000. LOF: identifying density-based local outliers. In *ACM sigmod record*.
- [2] Jinghui Chen, Saket Sathe, Charu Aggarwal, and Deepak Turaga. 2017. Outlier detection with autoencoder ensembles. In *SDM*.
- [3] Xiao Huang, Qingquan Song, Yuening Li, and Xia Hu. 2019. Graph Recurrent Networks with Attributed Random Walks. (2019).
- [4] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. *ICLR* (2017).
- [5] Jundong Li, Harsh Dani, Xia Hu, and Huan Liu. 2017. Radar: Residual analysis for anomaly detection in attributed networks. *IJCAI* (2017).
- [6] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *AAAI*.
- [7] Yuening Li, Ninghao Liu, Jundong Li, Mengnan Du, and Xia Hu. 2019. Deep Structured Cross-Modal Anomaly Detection. *IJCNN* (2019).
- [8] Ninghao Liu, Qiaoyu Tan, Yuening Li, Hongxia Yang, Jingren Zhou, and Xia Hu. 2019. Is a Single Vector Enough? Exploring Node Polysemy for Network Embedding. *arXiv preprint arXiv:1905.10668* (2019).
- [9] Tinghuai Ma, Lu Li, Sai Ji, Xin Wang, Yuan Tian, Abdullah Al-Dhelaan, and Mznah Al-Rodhaan. 2014. Optimized Laplacian image sharpening algorithm based on graphic processing unit. *PHYSICA A* (2014).
- [10] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* (2001).
- [11] Bryan Perozzi, Leman Akoglu, and Emmanuel Müller. 2014. Focused clustering and outlier detection in large attributed graphs. In *SIGKDD*.
- [12] Lukas Ruff, Nico Görnitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Robert Van dermeulen, Alexander Binder, Emmanuel Müller, and Marius Kloft. 2018. Deep one-class classification. In *ICML*.
- [13] Bernhard Schölkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. 2000. Support vector method for novelty detection. In *NeurIPS*.
- [14] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* (2008).
- [15] Wenchao Yu, Wei Cheng, Charu C Aggarwal, Kai Zhang, Haifeng Chen, and Wei Wang. 2018. Netwalk: A flexible deep embedding approach for anomaly detection in dynamic networks. In *SIGKDD*.
- [16] Bo Zong, Qi Song, Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. 2018. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. *ICLR* (2018).