Deep neural networks vs. Baselines. Table 1 shows the accuracy of vandal detection using the adjacency matrix $\mathbf{A}$ and spectral coordinates $\mathbf{x}_u$. We can observe that DAE and CNN models outperform the baselines significantly in all settings at the 10^{-8} level with a t-test. In the spectral space, the performance of SVM has a big jump when the percentage of training data increasing from 5% to 10%. In the adjacency matrix space, the performances of SVM do not improve while increasing the size of training data. It indicates that the SVM cannot be well-trained with a small training dataset, especially when the input has high dimensions. The accuracies of k-NN increase steadily while increasing the training data, but the accuracies of k-NN are also much worse than our DAE and CNN.

Spectrum vs. Adjacency matrix. In Table 1, we further observe that using the spectral coordinates as inputs ($\mathbf{x}_u$) achieves significantly better performance than using the adjacency matrix, especially when the percentages of training data are 5% and 10%. Meanwhile, in the spectral space, when the percentage of training dataset is 5%, DAE performs better than CNN since DAE pre-trains the model first. The pre-training phase of DAE encodes the information of nodes into hidden layers, which make the classifier predict the node labels with small training data. On the contrary, when the percentage of training dataset is larger than 5%, CNN outperforms DAE in the spectral space. It indicates CNN has better performance with enough training data. In the adjacency matrix space, the performances of DAE are worse than CNN with various sizes of training data. This is because DAE has much more parameters than CNN when using the adjacency matrix. DAE cannot be well-trained in a high-dimensional space with a small training dataset.

Effect of the dimension of spectral coordinate k. Table 2 compares the deep neural networks with k-NN and SVM on varying the dimension of spectral coordinate k. In this experiment, we use 20% of nodes as training data and the rest of nodes as testing data. When inputs of algorithms are $\mathbf{x}_u$, we can observe that DAE and CNN models outperform classical classifiers with various dimensions of spectral coordinates. We can further discover that DAE and CNN also achieve the most stable accuracy with various dimensions of spectral coordinates. However, the performance of SVM significantly drops while increasing the dimension of the spectral coordinate. This is because both DAE and CNN learn the hidden representations of nodes from their spectral coordinates. The hidden representations of nodes are useful for predicting the labels.

Neighbor inclusion vs. Neighbor exclusion. In our experiment, the inputs of DAE and CNN combine spectral coordinates of nodes and their 1-step neighbors. We further compare the performance of algorithms that adopt the node spectral coordinate with and without combining the 1-step neighbors' spectral coordinates as inputs in Table 2. We can observe that when using the information of the 1-step neighbors' spectral coordinates, the accuracies of all algorithms achieve around 1%-2% improvement. Therefore, combining the information of node neighbors can improve the performance of fraud detection.

Execution time. We also compare the execution time of deep neural networks using spectral coordinates and the adjacency matrix. The DAE and CNN models are trained on a Nvidia Tesla K20 GPU. We observe that when the ratio of training data is 20%, the training time of each epoch for CNN and DAE models with adjacency matrix is 2 seconds. On the contrary, the training time of each epoch for CNN and DAE models with spectral coordinates is less than 1 second. Therefore, using the spectral coordinates with low dimension is also more efficient than using the adjacency matrix.

Table 2: The accuracy of vandal detection with various dimensions of spectral coordinate k when 20% of nodes are used as the training dataset. We further compare algorithms using node spectral coordinate with/without combining neighbor spectral coordinates as inputs to algorithms.

Input	Algorithm	Dimension of spectral coordinate k				
		10	20	30	40	50
$\mathbf{x}_u$	k-NN	79.90%	78.85%	78.19%	78.33%	77.52%
	SVM	81.00%	81.40%	81.15%	80.70%	76.28%
	DAE	81.86%	81.65%	81.92%	81.87%	81.61%
	CNN	**82.24%**	**82.26%**	**82.61%**	**82.42%**	**82.47%**
$\boldsymbol{\alpha}_u$	k-NN	78.00%	77.49%	76.75%	76.55%	76.92%
	SVM	79.79%	79.65%	80.19%	80.31%	77.94%
	DAE	80.47%	81.10%	81.17%	81.20%	**81.40%**
	CNN	**80.52%**	**81.22%**	**81.48%**	**81.44%**	**81.39%**

4 CONCLUSIONS

We have presented a novel framework that applies deep neural networks on the spectral space of a signed graph to identify frauds. In particular, we first conduct graph spectral projections on a signed graph to obtain node spectral coordinates. The node and its s-step neighbors' spectral coordinates are combined together as inputs to the deep autoencoder and convolutional neural network models for fraud detection. The experiment results show that both deep neural networks achieve promising results on fraud detection. Our empirical evaluation further shows that combining the information of node neighbors can improve the effectiveness of deep neural networks on fraud detection.

ACKNOWLEDGMENTS

The authors acknowledge the support from the 973 Program of China (2014CB340404) to Shuhan Yuan and from National Science Foundation to Xintao Wu (1564250), Jun Li (1564348) and Aidong Lu (1564039). This research was conducted while Shuhan Yuan visited University of Arkansas.

REFERENCES

[1] Leman Akoglu, Hanghang Tong, and Danai Koutra. 2015. Graph based anomaly detection and description: a survey. *DMKD* 29, 3 (2015), 626–688.
[2] Pierre Baldi. 2012. Autoencoders , Unsupervised Learning , and Deep Architectures. In *Workshop on Unsupervised and Transfer Learning*. 37–50.
[3] FabrÃŋcio Benevenuto, Gabriel Magno, Tiago Rodrigues, and VirgÃŋlio Almeida. 2010. Detecting spammers on twitter. In *CEAS*.
[4] Y. Bengio, A. Courville, and P. Vincent. 2013. Representation Learning: A Review and New Perspectives. *TPAMI* 35, 8 (2013), 1798–1828.
[5] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. 2014. Uncovering Large Groups of Active Malicious Accounts in Online Social Networks. In *CCS*.
[6] Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, and Shiqiang Yang. 2014. CatchSync: Catching Synchronized Behavior in Large Directed Graphs. In *KDD*.
[7] Srijan Kumar, Francesca Spezzano, and V.S. Subrahmanian. 2015. VEWS: A Wikipedia Vandal Early Warning System. In *KDD*.
[8] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.
[9] X. Ying, X. Wu, and D. Barbará. 2011. Spectrum based fraud detection in social networks. In *ICDE*. 912–923.