EDITORIAL

Deep Anomaly Analytics: Advancing the Frontier of Anomaly Detection

Feng Xia, *RMIT University, Melbourne, VIC 3000, Australia*

Leman Akoglu, *Carnegie Mellon University, Pittsburgh, PA, 15213-3890, USA*

Charu Aggarwal, *IBM T.J. Watson Research Center, Yorktown, NY, 10598, USA*

Huan Liu, *Arizona State University, Tempe, AZ, 85287-8809, USA*

Deep anomaly analytics is a rapidly evolving field that leverages the power of deep learning to identify anomalies in various datasets. The use of deep anomaly analytics has increased significantly in recent years due to the growing need to detect anomalies in complex data that traditional methods struggle to handle. Deep anomaly analytics has the potential to transform various industries, including, e.g., healthcare, finance, and cybersecurity, by providing valuable insights and helping to diagnose diseases, prevent fraud, and detect cyber threats. However, there are also many challenges associated with deep anomaly analytics. This editorial provides an overview of the field of deep anomaly analytics, and highlights a few key challenges facing this field, i.e., time series anomaly detection, graph anomaly detection, efficiency (of models), and solving real-world problems. Additionally, it serves as an introduction to this special issue that delves further into these topics.

Over recent years, deep learning techniques have shown remarkable success in various areas, such as social computing, virtual assistants, advertising, visual recognition, and natural language processing. One area where deep learning is also proving to be effective is in the detection of anomalies in large and complex datasets. This emerging field of deep anomaly analytics has the potential to revolutionize the way we detect and handle anomalies in various domains, such as health care, cybersecurity, finance, digital media, neuroscience, and manufacturing.

The goal of anomaly detection is to identify unusual events or patterns in data that deviate from normal behavior. Detecting anomalies is a critical task of numerous applications. Traditional methods for anomaly detection rely on, for example, statistical methods, rule-based systems, or clustering techniques. However, these methods may not be effective in handling the complexities of large and heterogeneous datasets.[1,2] Deep anomaly analytics overcomes this limitation by leveraging the power of deep learning to detect anomalies in high-dimensional data.

1541-1672 © 2023 IEEE
Digital Object Identifier 10.1109/MIS.2023.3255590
Date of current version 4 April 2023.

Deep anomaly analytics uses various deep learning techniques such as graph neural networks (GNNs), autoencoders, variational autoencoders (VAEs), deep belief networks, convolutional neural networks, and recurrent neural networks, among others. These techniques enable the detection of complex anomalies that may not be easily detectable using traditional methods. For instance, in health care, deep anomaly analytics can help in the early detection of rare diseases that may be difficult to diagnose using traditional methods.[3]

Despite the significant potential of deep anomaly analytics, many challenges remain. For instance, one of the key challenges is the need for high-quality datasets for training deep learning models. Another challenge is the interpretability and explainability of the results obtained from these models, which are critical in domains such as health care, where the decision-making process needs to be transparent and interpretable. Researchers and practitioners are working to overcome these challenges and develop robust and reliable deep anomaly analytics techniques.

This special issue covers selected recent research results along this line. Nineteen articles were submitted in response to our call for articles, and five articles were accepted after a thorough review process by international experts. Here, we highlight a few key topics that this special issue touches on.

TIME SERIES ANOMALY DETECTION

Time series data are a sequence of observations collected over time. Time series anomaly detection is important because it has numerous applications in diverse fields.[4] For multivariate- or univariate-input and single-dimensional-output time series, the single-dimensional output can normally be transformed into an abnormal score. However, finding an appropriate transformation is challenging. The difficulty lies in identifying abnormal patterns, which can be further compounded by rare or imbalanced anomalous data. While deep learning approaches have shown promise in this area, there is still much research being conducted to address such issues.

One intuitive approach is to use inverse weights in the training loss to alleviate the rarity or imbalance. However, alternative methods have also been explored. For instance, Yang et al.[5] proposed a deep learning solution for PM2.5 anomaly prediction using an extreme value theory-based loss function and an architecture with self-attention. By modeling the data distribution effectively, the extreme value theory provides a better understanding of the underlying pattern of anomalies. Additionally, the self-attention mechanism captures global information from the time series, improving the model's ability to detect subtle anomalies. Given the variety of data and situations that may require different designs, the exploration of appropriate architectures and training losses is likely to remain an ongoing task.

Capturing temporal dynamics from time series data is sometimes insufficient for effective anomaly detection. In the case of unmanned aerial systems, for example, it is crucial to model the co-evolution among multiple features in time series. To tackle this challenge, Ma et al.[6] proposed a novel approach that integrates Bayesian networks and normalizing flows. They developed graphical normalizing flows (GNF), a graph-based autoregressive deep learning model that performs anomaly detection in multivariate time series through density estimation. GNF offers a promising solution to the problem of modeling the complex interactions among multiple features and provides a more accurate and comprehensive approach to anomaly detection in dynamic systems.

GRAPH ANOMALY DETECTION

Graphs are commonly used to represent complex systems, where nodes represent entities, and edges represent relationships between them.[7] Detecting anomalies in such graphs can help identify unusual patterns of behavior, which may be indicative of malicious activity or system failure. For instance, one important application of graph anomaly detection is in social network analysis, where graphs can be used to represent social relationships between individuals. Detecting anomalous behavior in social networks can help identify malicious actors who may be attempting to manipulate the network or spread misinformation. This application can be extended to detect anomalous citations from scholarly big data.[8]

Recent advances in graph anomaly detection[2,9,10] involve the use of deep learning techniques such as GNN and attention mechanisms to learn representations of the graph and detect anomalies. Generative models, such as VAE and generative adversarial networks, can be used to learn the distribution of normal graphs and detect anomalies based on deviations from this distribution.

Graphs can be incredibly large and complex, which makes it challenging to scale anomaly detection algorithms to handle large graphs in real time. In such contexts, efficiency becomes crucial. Ren et al.[11] introduced a contrastive self-supervised learning framework for detecting anomalies on heterogeneous networks. They explored a new perspective for generating positive and negative instances on heterogeneous graphs at the metapath level. This approach preserves rich semantics as a transferable form of knowledge for downstream anomaly detection tasks. The proposed method is highly efficient, conserving computational resources and making it suitable for use in embedded devices.

Despite these advances, it is worth noting that graph anomaly detection faces many challenges, including, for example, scalability, diversity of anomalies, label scarcity, noise and sparsity, dynamics, evaluation metrics, and benchmarks. Addressing these challenges requires ongoing research and the development of new algorithms and techniques that can scale to handle large, complex graphs and can adapt to different types of anomalies in diverse domains.

EFFICIENCY MATTERS

Efficiency is a critical aspect of deep learning that demands constant attention. Deep anomaly analytics typically involves processing large amounts of data from various sources, such as sensors, logs, and user behavior. To process these data in a timely manner, it is important to have efficient algorithms and computing infrastructure that can handle the volume and complexity of the data.[12] In some applications, such as cybersecurity or fraud detection, anomalies need to be detected in real time. This requires algorithms that can process data quickly and efficiently, so that anomalies can be detected as soon as they occur.

Liu et al.[13] studied an important application of anomaly detection in which efficiency is crucial, i.e., hyperspectral imaging (HSI). HSI is a hybrid technology that combines spectroscopy and imaging. Conventional HSI methods require the capture of a large amount of spatial–spectral data within a limited acquisition bandwidth. This presents a challenge for real-time anomaly detection. Improving signal collection and data processing efficiency are two main aspects of enhancing overall efficiency to achieve real-time hyperspectral anomaly detection. However, improving signal collection efficiency using conventional HSI hardware is limited. On the other hand, in terms of data processing efficiency, the optical neural network (ONN) is a novel method that uses optical processors to replace conventional electronic processors as a prospective alternative physical platform for deep learning.[14] In Lingfeng,[13] an ONN deep learning method is proposed to leverage a digital micromirror device to spatially modulate light to enable simultaneous HSI data acquisition and anomaly analytics. It can significantly reduce the sampling time by orders of magnitude without loss of detection accuracy.

SOLVING REAL-WORLD PROBLEMS

Modern society and life depend on the safe operation of many complex safety-critical systems, including, for example, power plants, oil wells, railway control systems, and telecommunication infrastructures. Anomalies in the operation of any of these systems can have potentially catastrophic consequences in the real world. However, employing anomaly detection to facilitate the safe operation of such systems poses several challenges. Deep learning can provide solutions to some of these key problem areas,[1] such as automated feature extraction and preprocessing as well as dealing with high-dimensional and noisy data. Despite this potential, developing effective deep learning architectures for real-world anomaly detection remains challenging and can often lead to suboptimal outcomes.

Magnusson et al.[15] introduced a novel type of recurrent neuron, developed using evolutionary principles,[16] for predicting anomalies in oil wells. They demonstrated that this neuron outperforms state-of-the-art recurrent neurons by a significant margin. The application of evolutionary neuron synthesis is not limited to anomaly detection but offers promising opportunities for solving various other problems. The combination of deep learning and the evolutionary synthesis of new neurons and network architectures could potentially lead to breakthroughs in ambitious pursuits like artificial general intelligence.

In conclusion, deep anomaly analytics has the potential to revolutionize the way we detect and handle anomalies in various domains. This special issue offers a glimpse of the state-of-the-art research in the field. There are still many issues and challenges surrounding deep anomaly analytics that need to be addressed. We anticipate significant advancements in both the research and application of deep anomaly analytics in the years to come.

ACKNOWLEDGMENTS

We would like to thank Prof. Longbing Cao, the former editor-in-chief of *IEEE Intelligent Systems,* for making this possible. We are very grateful to the administrative team for their support in managing this issue. We also thank all authors for their submissions and all reviewers for their diligent work in evaluating these submissions. A few authors helped with the preparation of this editorial. We used ChatGPT to improve the presentation.

REFERENCES

1. G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, “Deep learning for anomaly detection: A review,” *ACM Comput. Surveys*, vol. 54, no. 2, pp. 1–38, 2022, doi: 10.1145/3439950.
2. J. Ren, F. Xia, I. Lee, A. Noori Hoshyar, and C. Aggarwal, “Graph learning for anomaly analytics: Algorithms, applications, and challenges,” *ACM Trans. Intell. Syst. Technol.*, vol. 14, no. 2, pp. 1–29, Feb. 2023, doi: 10.1145/3570906.
3. T. Fernando, H. Gammulle, S. Denman, S. Sridharan, and C. Fookes, “Deep learning for medical anomaly detection – A survey,” *ACM Comput. Surveys*, vol. 54, no. 7, pp. 1–37, 2022, doi: 10.1145/3464423.
4. A. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, “A review on outlier/anomaly detection in time series data,” *ACM Comput. Surveys*, vol. 54, no. 3, pp. 1–33, 2022, doi: 10.1145/3444690.
5. H.-C. Yang, M.-C. Yang, G.-W. Wong, and M. Chang Chen, “Extreme event discovery with self-attention for PM2.5 anomaly prediction,” *IEEE Intell. Syst.*, vol. 99, pp. 1–11, Jan. 2023, doi: 10.1109/MIS.2023.3236561.
6. Y. Ma, M. N. Al Islam, J. Cleland-Huang, and N. V. Chawla, “Detecting anomalies in small unmanned aerial systems via graphical normalizing flows,” *IEEE Intell. Syst.*, early access, Mar. 2023, doi: 10.1109/MIS.2023.3252810.

7. F. Xia et al., “Graph learning: A survey,” *IEEE Trans. Artif. Intell.*, vol. 2, no. 2, pp. 109–127, Apr. 2021, doi: 10.1109/TAI.2021.3076021.
8. J. Liu, F. Xia, X. Feng, J. Ren, and H. Liu, “Deep graph learning for anomalous citation detection,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2543–2557, Jun. 2022, doi: 10.1109/TNNLS.2022.3145092.
9. X. Ma et al., “A comprehensive survey on graph anomaly detection with deep learning,” *IEEE Trans. Knowl. Data Eng.*, early access, Oct. 2021, doi: 10.1109/TKDE.2021.3118815.
10. R. Ma, G. Pang, L. Chen, and A. v. d. Hengel, “Deep graph-level anomaly detection by glocal knowledge distillation,” in *Proc. 15th ACM Int. Conf. Web Search Data Mining (WSDM'22)*, New York, NY, USA, 2022, pp. 704–714, doi: 10.1145/3488560.3498473.
11. J. Ren, M. Hou, Z. Liu, and X. Bai, “EAGLE: Contrastive learning for efficient graph anomaly detection,” *IEEE Intell. Syst.*, early access, 2023, doi: 10.1109/MIS.2022.3229147.
12. G. Menghani, “Efficient deep learning: A survey on making deep learning models smaller, faster, and better,” 2021, *arXiv:2106.08962*.
13. L. Liu, D. Ni, and L. Dai, “Spatial anomaly detection in hyperspectral imaging using optical neural networks,” *IEEE Intell. Syst.*, early access, 2023, doi: 10.1109/MIS.2023.3241431.
14. T. Wang, S.-Y. Ma, L. G. Wright, T. Onodera, B. C. Richard, and P. L. McMahon, “An optical neural network using less than 1 photon per multiplication,” *Nature Commun.*, vol. 13, no. 1, Jan. 2022, Art. no. 123, doi: 10.1038/s41467-021-27774-8.
15. L. Magnusson, R. Olsson, and C. Tran, “Recurrent neural networks for oil well event prediction,” *IEEE Intell. Syst.*, early access, 2023, doi: 10.1109/MIS.2023.3252446.
16. R. Olsson, C. Tran, and L. Magnusson, “Automatic synthesis of neurons for recurrent neural nets,” 2022, *arXiv:2207.03577*.

FENG XIA is a professor at the School of Computing Technologies, RMIT University, Melbourne, VIC 3000, Australia. His research interests include data science, artificial intelligence, graph learning, and systems engineering. Xia received his Ph.D. degree from Zhejiang University, Hangzhou, China. He is a Senior Member of IEEE and ACM, and an ACM Distinguished Speaker. Contact him at f.xia@ieee.org.

LEMAN AKOGLU is the Heinz College Dean’s Associate Professor of Information Systems at Carnegie Mellon University, Pittsburgh, PA, 15213-3890, USA. Her research interests include graph mining, pattern discovery and anomaly detection, with applications to fraud and event detection in diverse real-world domains. Akoglu received her Ph.D. degree from Carnegie Mellon University. Contact her at lakoglu@andrew.cmu.edu.

CHARU AGGARWAL is a Distinguished Research Staff Member at the IBM T. J. Watson Research Center, Yorktown Heights, NY, 10598, USA. His research interests include data streams, privacy, uncertain data, and social network analysis. Aggarwal received his Ph.D. degree in operations research from the Massachusetts Institute of Technology. He is a Fellow of IEEE, ACM, and SIAM. Contact him at charu@us.ibm.com.

HUAN LIU is a Regents Professor and Ira A. Fulton Professor of Computer Science and Engineering at Arizona State University, Tempe, AZ, 85287-8809, USA. He is a Fellow of IEEE, ACM, AAAI, and AAAS. Contact him at huanliu@asu.edu.