6-2022

Deep learning for anomaly detection

Guansong PANG
Singapore Management University, gspang@smu.edu.sg

Charu AGGARWAL

Chunhua SHEN

Nicu SEBE

Editorial
Deep Learning for Anomaly Detection

ANOMALY detection aims at identifying data points which are rare or significantly different from the majority of data points. Many techniques are explored to build highly efficient and effective anomaly detection systems, but they are confronted with many difficulties when dealing with complex data, such as failing to capture intricate feature interactions or extract good feature representations. Deep-learning techniques have shown very promising performance in tackling different types of complex data in a broad range of tasks/problems, including anomaly detection. To address this new trend, we organized this Special Issue on Deep Learning for Anomaly Detection to cover the latest advancements of developing deep-learning techniques specially designed for anomaly detection. This editorial note provides an overview of the paper submissions to the Special Issue, and briefly introduces each of the accepted articles.

I. Overview of the Special Issue

Due to the significance to many critical domains like cybersecurity, fintech, healthcare, public security, and AI safety, anomaly detection has been one of the most active research areas in various communities, such as data mining, machine learning, and computer vision [1], [2]. Many techniques are explored to build highly efficient and effective anomaly detection systems, but they are confronted with many difficulties when dealing with complex data, such as failing to capture intricate feature interactions or extract good feature representations. Deep-learning techniques, including different types of deep neural network architectures such as recurrent neural networks, convolutional neural networks, generative adversarial networks, and graph neural networks, and a broad range of regularization and training techniques for learning expressive representations, have shown very promising performance in tackling different types of complex data in a broad range of tasks/problems [3]. Nevertheless, its development in the area of anomaly detection is relatively slow due to some unique characteristics of anomalies, such as rare and heterogeneous distributions [2]. The goal of this Special Issue is to discuss the latest advancements in developing deep-learning techniques specially designed for anomaly detection.

In total, we received 77 valid submissions from a wide range of 22 countries, including China (30), India (7), the USA (6), the U.K. (5), Australia (5), and Brazil (3). An overview of the country of submitting author is presented in Fig. 1. After a rigorous review procedure, 22 submissions were accepted for publication in this Special Issue. The majority of the accepted manuscripts focus on devising novel deep anomaly detection methods. In addition, there are two comparative studies performing a large-scale evaluation of existing state-of-the-art deep anomaly detectors on different types of data. We also have four studies that present interesting real-world applications of anomaly detection. In the following, we provide a brief introduction to each of these accepted manuscripts.

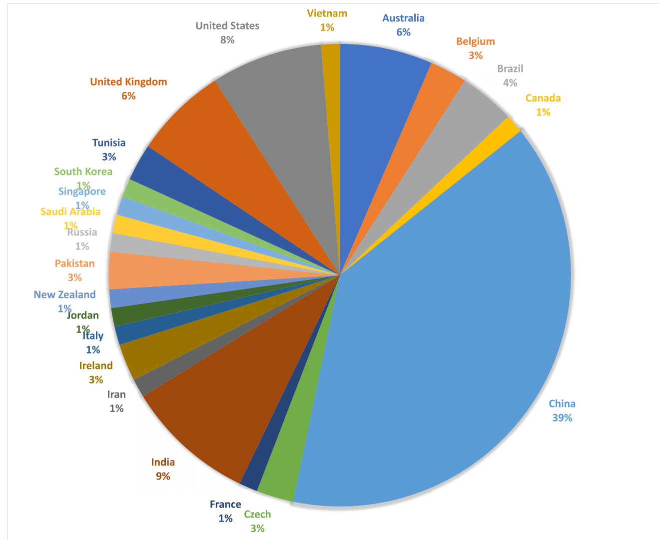

Fig. 1. Distribution of countries of submitting author. We received 77 submissions from 22 countries, with China, India, the USA, the U.K., and Australia as the top five sources of submission country.

II. Novel Anomaly Detection Methods

Anomaly detection is a general problem in different types of data, such as video data, image data, graph data, temporal data, and tabular data. Below we briefly introduce the accepted articles from a data perspective.

A. On Visual Data

In [A1], Nguyen *et al.* tackle video anomaly detection with normal videos available during training. The work extends the model that includes one shared encoder to learn features and two decoders to enforce texture coherence and successive motion coherence by reconstruction error minimization and image translation prediction, respectively. The method shows competitive performance on diverse benchmarks from pedestrian and traffic video surveillance scenes. The work is an extended version of its conference version published in ICCV 2019 [4].

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TNNLS.2022.3162123.

Digital Object Identifier 10.1109/TNNLS.2022.3162123

In [A2], Wang *et al.* introduce a future frame prediction-based method for video anomaly detection. The model uses multiple path network structure and non-local blocks to learn the feature representations, and introduces a loss function that helps increase the robustness to irrelevant background information. By doing so, the model does not require generative adversarial models (as needed in previous work, such as the seminal work [5] in this line) to effetively predict future frames and detect abnormal events in videos. In addition, multiresolution modeling of normal patterns is found to be useful in detecting anomalies, which is also confirmed in [A3] in which Massoli *et al.* design a multi-layer one-class classifier on feature representations learned in different layers. Similar observations are also found in several other recent studies [7], [8].

In [A4], Yang *et al.* explore the incorporation of a memory module into a generative adversarial network (GAN)-based anomaly detection models. The resulting model can map normal samples into a latent space where they reside in a convex hull of the memory units, facilitating better discrimination when performing data reconstruction from this latent space. Learning an auxiliary memory module has been emerging as a popular approach to enable autoencoder- or GAN-based anomaly detection in recent years [9]–[11].

In [A5], Zhou *et al.* propose a structure-texture-based data reconstruction method for anomaly detection. The method learns internal structures underlying normal images and uses the learned structure representations to reconstruct input images, alleviating the overfitting to anomaly images in traditional autoencoder-based detectors that can often reconstruct anomalous images comparably well to normal images. The work extends its ECCV conference version [12] by the introduction of a memory module to learn the structure-texture correspondence and the incorporation of extra structure information into the model.

In [A6], Macêdo *et al.* study the problem of out-of-distribution (OOD) detection where the target is to detect OOD samples while at the same time accurately classifying in-distribution samples into fine-grained classes. The work introduces a novel loss function based on exclusively distances between samples and class prototypes, showing effective performance improvement when plugging into the current OOD detection methods to replace the popular softmax loss.

In [A7], Li *et al.* explore the case of automated anomaly detection, aiming at automatically searching for an optimal neural network-based anomaly detector out of a large number of possible detection model candidates. They introduce a reinforcement learning-based search method that uses curiosity-based internal rewards, in addition to the use of accuracy on a labeled validation set as external rewards, to bind the model search with anomaly detection. The method is evaluated on six image datasets with injected OOD samples from external data sources as anomaly samples, and compared with four baseline OOD detectors and anomaly detectors.

B. On Graph Data

In [A8], Liu *et al.* introduce a contrastive self-supervised anomaly detection method for attributed graph data. The self-supervised pretext task is to predict the relationship between a target node and a small subgraph, in which we may generate a positive subgraph by sampling the subgraph around the target node, or a negative subgraph that is randomly sampled from the other part of the graph. This subgraph can be seen as a context for the target node. Thus, if the target node does not fit well with the subgraph, it indicates anomalousness of the node. In doing so, the pretext task and the anomaly detection task are unified into one framework. As only subgraph inputs are required, the method can also reduce time and space complexities compared to many existing graph neural network methods that work on the full graph.

In [A9], Zhao *et al.* introduce a new graph anomaly detection method, which combines pattern mining methods and graph neural network-based methods to learn global and local relevant features into node feature representations for detecting abnormal nodes in a graph. The model is evaluated on two real-world graph anomaly detection datasets, including one user-posts-hashtag graph dataset and one Bitcoin trading user dataset, and compared with 14 competing methods.

In [A10], Ding *et al.* explore the problem of cross-domain graph anomaly detection, in which we aim to adapt knowledge learned from a fully labeled source graph data to an unlabeled target graph data. They use a shared-weight graph neural network encoder to embed nodes of source and target graphs into the same feature space, and then use three loss functions, including anomaly classification loss, domain discriminator loss, and node attribute-based data reconstruction loss, to train the model.

C. On Temporal Data

In [A11], Deng *et al.* aim to detect abnormal events in spatial-temporal data derived from urban traffic data. The work formulates the problem into a graph mining problem and introduces a graph convolutional adversarial network to tackle the problem. The use of graph networks enables the modeling of correlations between spatially adjacent data points. The model can also incorporate previously occurred trend features and some other pre-defined auxiliary features. The model is evaluated on two real-world traffic anomaly detection datasets and shows good detection recall rates.

In [A12], Zhou *et al.* introduce a class prototype learning-based method for novelty detection while at the same time updating classification models with newly detected novel classes. It dynamically learns a class center for each class for know class classification and uses the classification uncertainty and novel class prediction probability to detect novel classes. The effectiveness of the method is justified on five real-world datasets with ten relevant methods as competing methods.

D. On Tabular Data

In [A13], Xie *et al.* introduce a deep-generative model for semi-supervised anomaly detection on multidimensional data. The model makes use of a small number of labeled anomaly samples to train the generative model using a classification-based objective rather than a data reconstruction objective in existing work. It concludes that if there are labeled anomaly examples available during training, we should utilize them to achieve substantially improved detection accuracy. Similar observations/conclusions are discussed in a number of previous studies (e.g., [13]–[16]), though the approach to leverage the labeled anomaly data is different.

In [A14], Zhou *et al.* explore how to fully leverage a limited number of labeled anomaly samples to train anomaly-informed detection models. The work introduces an encoder that combines autoencoder-based hidden representation, residual vectors between original data and reconstructed data, and a reconstruction error to form a feature representation of a given data point. It then introduces two loss functions that enforce a large margin between the reconstruction errors (deviation-based anomaly scores) of unlabeled data and the labeled anomaly examples. The resulting model is evaluated on eight tabular datasets and compared with both "supervised" and unsupervised detection models.

In [A15], Zhao *et al.* address the semi-supervised anomaly detection problem. The work first uses a local density measure to identify a set of prototypical instances, dubbed "landmarks" and then introduces a new loss function that minimizes the distance between unlabeled samples while guaranteeing a large margin between the unlabeled samples and anomaly examples. The effectiveness of the method is evaluated on 12 real-world tabular datasets and compared with 11 semi-supervised/unsupervised competing methods.

In [A16], Li *et al.* introduce a new deep anomaly detection method by imposing several constraints on top of the popular deep support vector data description-based method [17] to learn more reasonable one-class center description. The key idea is to incorporate the importance of each data instance into the one-class center generation and to dynamically update the one-class center. The model is evaluated on 16 tabular datasets, and compared with nine competing methods. The model is also evaluated on two image datasets.

III. Large-Scale Empirical Evaluation

There have been a large number of shallow and deep anomaly detection methods developed over the years. However, a systematic empirical comparison of these methods is missing in the literature. In this Special Issue, we have two studies that are specifically dedicated to fill this gap [A17], [A18]. More specifically, in [A17], Škvára *et al.* present a comprehensive evaluation of both shallow and deep anomaly detection methods (with a focus on generative models) on a large collection of tabular and image datasets. The study evaluates how the difference in the dataset type, the hyperparameter selection strategy, or the computational resource is associated with the performance of each detector. The code and results of the study are publicly available for download.

In [A18], Garg *et al.* perform a large-scale evaluation of deep unsupervised and semi-supervised anomaly detection methods on multivariate time series data. The work uses 11 deep methods from three general categories of approach, including generic normality feature learning, anomaly-measure dependent feature learning, and end-to-end anomaly scoring. Seven real-world time-series anomaly detection datasets from several domains, such as water resource management, manufacturing process, spacecraft, and server monitoring, are used.

IV. Applications of Anomaly Detection

One main reason that anomaly detection is of great interest to the community is due to its application potential in broad domains. Some interesting and relevant applications we accepted to this special issue include malware detection, rumor detection, and detection of abnormal citation behaviors. Particularly, in [A19], Sun *et al.* study the problem of malware detection and introduce a two-stage method for both efficient and accurate detection, in which a traditional classifier is first used to detect malware with high confidence while a deep classifier is used to handle malware candidates that the traditional classifier are not confident with.

In [A20], Li *et al.* explore the problem of rumor detection and introduce a hierarchical heterogeneous graph-based method to jointly optimize stance and rumor detection, resulting in largely improved performance in both stance detection and rumor detection.

In [A21], Liu *et al.* tackle the problem of identifying abnormal citations in an academic paper citation network, facilitating the identification of citation manipulation and inflation behaviors. A graph learning method is introduced to achieve a joint learning of node features and edge features in the citation network and the relevant citation context.

In [A22], Cheng *et al.* address the problem of blade icing detection, in which we are interested in automatically assessing whether there is a blade icing issue based on data of wind turbines installed at high-latitude places. Class imbalance is an inherent issue in blade icing data, as turbines mostly operate under normal conditions. The work introduces a prototypical-based neural network to address the class-imbalance problem. The proposed model is evaluated on one real-world dataset that contains about 1,000 hours of wind turbine data, accompanied by 26 features defined by domain experts. The model is compared with unsupervised anomaly detectors and imbalanced classifiers to show its effectiveness.

Acknowledgment

The authors would like to extend sincere appreciation to the former and current Editors-in-Chief, Prof. Haibo He and Prof. Yongduan Song, the journal administrative team, authors, and reviewers for their support of this Special Issue.

Guansong Pang, *Guest Editor*
School of Computing and Information Systems
Singapore Management University
Singapore 178902

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

Chunhua Shen, *Guest Editor*
State Key Laboratory of CAD&CG
Zhejiang University
Hangzhou 310027, China

NICU SEBE, *Guest Editor*
Department of Information Engineering and Computer Science
University of Trento
38122 Trento, Italy

APPENDIX: RELATED ARTICLES

[A1] T.-N. Nguyen, S. Roy, and J. Meunier, "SmithNet: Strictness on motion-texture coherence for anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2287–2300, Jun. 2022.

[A2] X. Wang *et al.*, "Robust unsupervised video anomaly detection by multipath frame prediction," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2301–2312, Jun. 2022.

[A3] F. V. Massoli, F. Falchi, A. Kantarci, C. S. Akti, H. K. Ekenel, and G. Amato, "MOCCA: Multilayer one-class classification for anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2313–2323, Jun. 2022.

[A4] Z. Yang, T. Zhang, I. S. Bozchalooi, and E. Darve, "Memory-augmented generative adversarial networks for anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2324–2334, Jun. 2022.

[A5] K. Zhou *et al.*, "Memorizing structure-texture correspondence for image anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2335–2349, Jun. 2022.

[A6] D. Macêdo, T. I. Ren, C. Zanchettin, A. L. Oliveira, and T. Ludermir, "Entropic out-of-distribution detection: Seamless detection of unknown examples," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2350–2364, Jun. 2022.

[A7] Y. Li *et al.*, "Automated anomaly detection via curiosity-guided search and self-imitation learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2365–2377, Jun. 2022.

[A8] Y. Liu, Z. Li, S. Pan, C. Gong, C. Zhou, and G. Karypis, "Anomaly detection on attributed networks via contrastive self-supervised learning," *IEEE Trans. neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2378–2392, Jun. 2022.

[A9] T. Zhao, T. Jiang, N. Shah, and M. Jiang, "A synergistic approach for graph anomaly detection with pattern mining and feature learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2393–2405, Jun. 2022.

[A10] K. Ding, K. Shu, X. Shan, J. Li, and H. Liu, "Cross-domain graph anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2406–2415, Jun. 2022.

[A11] L. Deng, D. Lian, Z. Huang, and E. Chen, "Graph convolutional adversarial networks for spatiotemporal anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2416–2428, Jun. 2022.

[A12] D.-W. Zhou, Y. Yang, and D.-C. Zhan, "Learning to classify with incremental new class," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2429–2443, Jun. 2022.

[A13] Q. Xie, P. Zhang, B. Yu, and J. Choi, "Semisupervised training of deep generative models for high-dimensional anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2444–2453, Jun. 2022.

[A14] Y. Zhou, X. Song, Y. Zhang, F. Liu, C. Zhu, and L. Liu, "Feature encoding with autoencoders for weakly supervised anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2454–2465, Jun. 2022.

[A15] J. Zhao, F. Deng, Y. Li, and J. Chen, "Attract-repel encoder: Learning anomaly representation away from landmarks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2466–2479, Jun. 2022.

[A16] D. Li, Q. Tao, J. Liu, and H. Wang, "Center-aware adversarial autoencoder for anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2480–2493, Jun. 2022.

[A17] V. Skvara, J. Franca, M. Zorek, T. Pevny, and V. Smidl, "Comparison of anomaly detectors: Context matters," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2494–2507, Jun. 2022.

[A18] A. Garg, W. Zhang, J. Samaran, R. Savitha, and C.-S. Foo, "An evaluation of anomaly detection and diagnosis in multivariate time series," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2508–2517, Jun. 2022.

[A19] R. Sun *et al.*, "Learning fast and slow: Propedeutica for real-time malware detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2518–2529, Jun. 2022.

[A20] C. Li *et al.*, "Joint stance and rumor detection in hierarchical heterogeneous graph," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2530–2542, Jun. 2022.

[A21] 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.

[A22] X. Cheng, F. Shi, X. Liu, M. Zhao, and S. Chen, "A novel deep class-imbalanced semisupervised model for wind turbine blade icing detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2558–2570, Jun. 2022.

REFERENCES

[1] C. C. Aggarwal, "An introduction to outlier analysis," in *Outlier Analysis*. Springer, 2017, pp. 1–34.

[2] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, "Deep learning for anomaly detection: A review," *ACM Comput. Surv.*, vol. 54, no. 2, pp. 1–38, Mar. 2021.

[3] S. Pouyanfar *et al.*, "A survey on deep learning: Algorithms, techniques, and applications," *ACM Comput. Surv.*, vol. 51, no. 5, pp. 1–36, Sep. 2019.

[4] T.-N. Nguyen and J. Meunier, "Anomaly detection in video sequence with appearance-motion correspondence," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Dec. 2019, pp. 1273–1283.

[5] W. Liu, W. Luo, D. Lian, and S. Gao, "Future frame prediction for anomaly detection—A new baseline," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6536–6545.

[6] M. Salehi, N. Sadjadi, S. Baselizadeh, M. H. Rohban, and H. R. Rabiee, "Multiresolution knowledge distillation for anomaly detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 14902–14912.

[7] Y. Chen, Y. Tian, G. Pang, and G. Carneiro, "Deep one-class classification via interpolated Gaussian descriptor," 2021, *arXiv:2101.10043*.

[8] D. Gong *et al.*, "Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1705–1714.

[9] H. Park, J. Noh, and B. Ham, "Learning memory-guided normality for anomaly detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Dec. 2020, pp. 14372–14381.

[10] J. Hou, Y. Zhang, Q. Zhong, D. Xie, S. Pu, and H. Zhou, "Divide-and-assemble: Learning block-wise memory for unsupervised anomaly detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Sep. 2021, pp. 8791–8800.

[11] K. Zhou *et al.*, "Encoding structure-texture relation with P-net for anomaly detection in retinal images," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2020, pp. 360–377.

[12] G. Pang, L. Cao, L. Chen, and H. Liu, "Learning representations of ultrahigh-dimensional data for random distance-based outlier detection," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 2041–2050.

[13] G. Pang, C. Shen, and A. van den Hengel, "Deep anomaly detection with deviation networks," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2019, pp. 353–362.

[14] L. Ruff *et al.*, "Deep semi-supervised anomaly detection," 2019, *arXiv:1906.02694*.

[15] G. Pang, C. Shen, H. Jin, and A. van den Hengel, "Deep weakly-supervised anomaly detection," 2019, *arXiv:1910.13601*.

[16] G. Pang, A. van den Hengel, C. Shen, and L. Cao, "Toward deep supervised anomaly detection: Reinforcement learning from partially labeled anomaly data," in *Proc. 27th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, Aug. 2021, pp. 1298–1308.

[17] L. Ruff *et al.*, "Deep one-class classification," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 4393–4402.

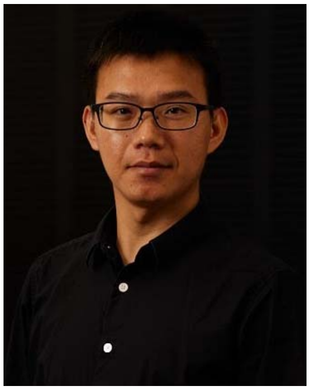

Guansong Pang received the Ph.D. degree from the University of Technology Sydney, Ultimo, NSW, Australia, in 2019.

He is a tenure-track Assistant Professor of computer science with the School of Computing and Information Systems, Singapore Management University (SMU), Singapore. Before joining SMU, he was a Research Fellow with the Australian Institute for Machine Learning (AIML), The University of Adelaide, Adelaide, SA, Australia. His research investigates novel data mining and machine learning techniques and their applications, with a particular focus on abnormality and rarity learning.

Charu Aggarwal received the B.Tech. degree in computer science from IIT Kanpur, Kanpur, India, in 1993, and the Ph.D. degree in operations research (focus: mathematical optimization) from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1996.

He is a Distinguished Research Staff Member (DRSM) at the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. He has worked extensively in the field of data mining, with particular interests in data streams, privacy, uncertain data, and social network analysis.

Dr. Aggarwal is a fellow of the ACM in 2013 and the SIAM in 2015 for contributions to knowledge discovery and data mining algorithms.

Chunhua Shen was a Professor at The University of Adelaide, Adelaide, SA, Australia. He is a Professor of computer science at Zhejiang University, Hangzhou, China.

Nicu Sebe is a Professor of computer science at the Department of Information Engineering and Computer Science and the Leader of the Multimedia and Human Understanding Group (MHUG), University of Trento, Trento, Italy. His main research interests relate to the investigation and implementation of new techniques in the fields of computer vision and multimedia.

Dr. Sebe is a fellow of IAPR and a Senior Member of ACM. He was the General Chair of ACM Multimedia 2013 and ICMR 2017 and the Program Chair of ACM Multimedia 2011 and 2007, ECCV 2016, ICCV 2017, and ICPR 2020.