Guest Editorial: Big Media Data: Understanding, Search, and Mining

Jingdong Wang, Guo-Jun Qi, Nicu Sebe, and Charu C. Aggarwal

THERE has been an explosive increase in media data, such as images, videos and social media in the Internet, mobile devices, and desktops. Accordingly, the area of big media data has been attracting a lot of research interest from both academia and industry. Big media data offers a promising possibility for in-depth media understanding, as well as exploring the very big scale media data to bridge the well-known semantic gap between high-level semantics and low-level features. It provides richer information, ranging from social relations to context information associated to rich media data of diverse modalities. It also provides us the opportunity to mine reliable and helpful knowledge from big media data for a wide variety of applications.

Big media data is big in terms of various aspects, such as the number of media items, the dimension of the representation, and the number of concepts, and thus entails a lot of research challenges. For example, how do the traditional machine learning algorithms, which have been proven efficient and effective in thousands of data points, scale up to the web-scale big media data with millions and even billions of items? How is the big media data organized in order to enable efficient browsing and retrieval? Besides, it is also important to construct benchmark data to facilitate and validate the newly-developed big-media algorithms.

The first part of this special issue includes three carefully selected examples of current research trends in big media data. The contributions cover three themes: large scale similar image search, image search quality improvement, and semi-supervised multi-label image annotation.

The first paper is "Code Consistent Hashing Based on Information-Theoretic Criterion" by S. Zhang, J. Liang, R. He, and Z. Sun. The authors propose a new learning-to-hash approach, called code consistent hashing, for large scale similar image search. To formulate the problem, they introduce a new code consistency constraint that leverages discriminative information, and utilize the Hadamard code that favors an information-theoretic criterion as the class prototype. Experimental results demonstrate the state-of-the-art performance in both image retrieval and classification applications.

- *J. Wang is with Microsoft Research, Beijing, China. E-mail: jingdw@microsoft.com.*
- *G.J. Qi is with the Department of Computer Science, University of Central Florida, Orlando, FL. E-mail: guojun.qi@ucf.edu.*
- *N. Sebe is with Department of Information Engineering and Computer Science, University of Trento, Italy. E-mail: sebe@disi.unitn.it.*
- *C. C. Aggarwal is with IBM T.J. Watson Research Center, New York, NY. E-mail: charu@us.ibm.com.*

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TBDATA.2015.2504198

In their paper "Exploration of Image Search Results Quality Assessment," X. Tian, Y. Lu, N. Stender, L. Yang, and D. Tao explore the image search results quality assessment problem and propose a RankSVM-based method to compare the quality of a set of ranking result lists for a given query. They design a set of lightweight features that reflect the visual distribution difference between ranking lists with different levels of quality. They demonstrate the capacity of their approach for reranking ability assessment and present several applications, including reranking feature and model selection, merging of image search results, and query suggestion.

F. Wu, Z. Wang, Z. Zhang, Y. Yang, J. Luo, W. Zhu, and Y. Zhuang describe, in their paper "Weakly Semi-supervised Deep Learning for Multi-Label Image Annotation," a weakly semi-supervised deep learning method for multi-label image annotation. They introduce a weakly weighted pairwise ranking loss to handle weakly labeled images, and employ a triplet similarity loss to harness unlabeled images. They train a deep convolutional neural network with the goal of minimizing the two losses with images collected from social networks where images are either only weakly labeled with multiple labels or unlabeled. Experimental results show that their approach significantly improves the performance.

We would like to extend our thanks to the authors and the reviewers for their efforts. We also thank the staff of TBD and especially Kathy Santa Maria and Leigh Ann Testa for their help and guidance throughout the publication process.

Jingdong Wang
Guo-Jun Qi
Nicu Sebe
Charu C. Aggarwal
Guest Editors

Jingdong Wang received the ME and BE degrees in automation from the Department of Automation, Tsinghua University, Beijing, China, in 2001 and 2004, respectively, and the PhD degree in computer science from the Department of Computer Science and Engineering, the Hong Kong University of Science and Technology, Hong Kong. He is a lead researcher in the Visual Computing Group, Microsoft Research Asia. His areas of interest include computer vision, multimedia, and machine learning. At present, he is mainly working on image search including indexing and compact coding for large scale similarity search as well as interactive image search, and visual understanding including image recognition, salient object detection, person re-identification, and image segmentation. He has served as an area chair in ACMMM 2015, an area chair in ICME 2015, a track chair in ICME 2012, a special session chair in ICMR 2014. He has made many contributions to Microsoft products, including Bing image search, Project Oxford, and XiaoIce.

Guo-Jun Qi received the BE degree in automation from the University of Science and Technology of China and the PhD degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign. He is an assistant professor in the Department of Computer Science at the University of Central Florida. His research interests include pattern recognition, machine learning, computer vision and multimedia. He is the senior co-recipient of the best student paper award in IEEE Conference on Data Mining (2014), and the recipient of best paper award in the ACM International Conference on Multimedia (2007). He has served as a program cochair of MMM 2016, an area chair of ACM Multimedia 2015, a senior program committee member of ACM CIKM 2015, and as a reviewer for many academic conferences and journals in the fields of computer vision, pattern recognition, machine learning, and multimedia. He has published more than 50 papers in the venues of computer vision, pattern recognition, machine learning, data mining, and multimedia computing.

Nicu Sebe received the PhD degree from Leiden University, The Netherlands, in 2001. He is a professor in the University of Trento, Italy, and is leading the research in the areas of multimedia information retrieval and human-computer interaction in computer vision applications. He is a senior member of the Association for Computing Machinery and a fellow of the International Association for Pattern Recognition. He was the general cochair of FG 2008 and ACM Multimedia 2013, and the program chair of CIVR 2007 and 2010, and ACM Multimedia 2007 and 2011. He is a cochair of the IEEE Computer Society Task Force on Human-centered Computing and is an associate editor of *IEEE Transactions on Multimedia*, *Computer Vision and Image Understanding*, *Machine Vision and Applications*, *Image and Vision Computing*, *International Journal of Human-Computer Studies*, and *Journal of Multimedia*.

Charu C. Aggarwal received the BS degree from IIT Kanpur in 1993 and the PhD degree from the Massachusetts Institute of Technology in 1996. He is a distinguished research staff member at the IBM T.J. Watson Research Center in Yorktown Heights, New York. He has extensively worked in the field of performance analysis, databases, and data mining. He has published more than 135 papers in refereed conferences and journals, and has been granted more than 50 patents. He has served on the program committees of most major database/data mining conferences, and served as a program vice-chairs of the SIAM Conference on Data Mining, 2007, the IEEE ICDM Conference, 2007, the WWW Conference 2009, and the IEEE ICDM Conference, 2009. He served as an associate editor of the *IEEE Transactions on Knowledge and Data Engineering* from 2004 to 2008. He is an action editor of *Data Mining and Knowledge Discovery*, an associate editor of the *ACM SIGKDD Explorations*, and an associate editor of *Knowledge and Information Systems*. He is a fellow of the ACM (2013) and the IEEE (2010) for "contributions to knowledge discovery and data mining techniques."