COLUMN: From the Editors

Reflecting on Two Decades of Services Computing

Patrick C. K. Hung
University of Ontario Institute of Technology

M. Brian Blake
Drexel University

As young researchers more than 15 years ago, Patrick Hung and Brian Blake participated in the 2003 IEEE International Conference on Web Services. This year in July, Hung and Blake attended the most recent iteration of the conference, the 2018 IEEE World Congress on Services in San Francisco. Although the topics have varied over the decades, the 2018 Congress covered aspects of innovative services research and current and emerging applications in services computing. The Congress now contains seven different conferences in the areas of big data, cloud computing, edge computing, cognitive computing, the Internet of Things, web services, and services computing. Over lunch, Hung and Blake reflected on papers from this year's Congress and considered current trends in the community.

In 2003, *services computing* represented a relatively new research area that encompassed a new class of paradigms and technologies including web services, service-oriented architecture (SOA), business process integration and management, utility/grid computing, and autonomic computing. At the time, the IEEE Computer Society had officially launched the Technical Steering Committee for Services Computing (TCSVC). The discipline of services computing covers the science and technology of bridging the gap between business services and information technology services. This year's IEEE World Congress on Services had interesting contributions in four research categories: big data analytics/cognitive computing, mobile edge computing, machine learning, and robotic computing. Here, we discuss some of the papers from these areas and their contributions.

BIG DATA ANALYTICS

Big data is a term that has been gaining considerable attention in recent years. It describes a large amount of organized or unorganized data that is analyzed to make informed decisions or

Published by the IEEE Computer Society
1089-7801/18/$33.00 ©2018 IEEE

evaluations. The data can be taken from a variety of sources including browsing history, geolocation, social media, purchase history and medical records. There are three main characteristics associated with big data:

1. *volume* is used to describe the vast amounts of data that is utilized by big data;
2. *variety* is used to describe the many different types of data sources used as part of a big data analytics system; and
3. *velocity* is used to describe the speed at which data is generated.[1]

Here, we share a few readings from the Congress on the latest findings.

Referring to data management and quality evaluation, Ikbal et al.[2] presented an across-the-board quality management framework. It includes a roadmap for data scientists that considers the assessment of quality as early as possible and end-to-end integration across the following areas:

- implementation of continuous quality improvement and enforcement mechanisms in quality management;
- specification of data quality metrics that should cope with the data's dynamic nature and its unconventional characteristics;
- development of new quality dimensions with specific measurement attributes for unstructured and schema-less data;
- enforcement of quality requirements, generation of quality reports and feedback to support assessment activities;
- development of automated real-time dashboards for data quality monitoring;
- application of higher degrees of statistical proof in different data quality evaluation processes including sampling, regression, correlation, and matching;
- development of effective quality outcome predictions; and
- evaluation of the quality of representative sets of data samples and generation of quality models to apply to the whole data.

It's widely accepted that quality is the most important foundation to support big data analytics.

For big data analytics, researchers establish approaches to mine and discover unknown patterns and insights from huge volumes of raw data.[3] Big data analytics has become very popular in the areas of marketing and customer-relationship management. Many industries have adopted the use of big data analytics and are experiencing fantastic results. For example, the healthcare, retail, insurance, and telecommunications industries have all displayed the endless possibilities of implementing big data into their operations.[1]

Khalajzadeh et al.[3] studied data analytics software tools for domain experts who are not computing specialists. The tools have the following functions:

- to cover data preprocessing operations such as cleaning, wrangling, anomaly detection, and so on;
- to incorporate a variety of algorithms for each stage of data processing, modeling, and evaluation processes; and
- to cover software development life cycle (SDLC) stages, including business problem descriptions, requirements, design, implementation, testing, and deployment.

A research topic related to big data analytics is sentiment analysis. Sentiment analysis techniques determine the overall sentiment orientation for topics discussed in the text as positive, negative, or neutral, while emotion detection from text identifies the categories of emotions the text expresses.[4] Analyses of text using emotional categories have proven valuable with the growth of social media tools, such as Twitter, for communication and collaboration. Traditional sentiment analysis only visualizes an aggregation of opinions expressed in the content, while neglecting the presence of the creators of the content and the impact of their varying levels of participation. Hemmings-Jarrett et al.[5] addressed this gap and concluded that differences in the level of user participation potentially impact the samples extracted for sentiment analysis and interpretation in their study.

MOBILE EDGE COMPUTING

Mobile edge computing is a network architecture concept that enables the cloud computing service environment at the edge of the cellular network by running applications and performing related processing tasks closer to the cellular customer. Mobile edge computing is designed to decrease latency and network congestion for mobile users. Zhang et al.[6] presented a quality of experience (QoE) aware control plane for adaptive streaming service over mobile edge computing infrastructures with the following features:

- a timeslot system with a look-ahead window for calculating the cost of edge node switch and video quality adaption (to balance network load and reduce latency);
- conducting service adaption via a set of cooperative action components running on client devices, edge nodes, and center nodes (to ensure a smooth viewing experience); and
- constructing a flexible QoE model and extending the scope and meaning of user-perceived experience.

Edge servers are usually deployed at the edge of the network so that computation is performed at the proximity of data source. This has two advantages: on downstream data, edge servers play a role of cloud service provider, making computing resources close to end users so that the latency of service request can be very low; and on upstream data, it helps to improve the network transmission on the core network. Li and Wang[7] studied the problem of energy-aware edge server placement as a multi-objective optimization problem and found a more effective placement scheme with low energy consumption.

MACHINE LEARNING

Intelligence in computing is essential to achieve service excellence for the ever more complicated requirements of the rapidly evolving global environment, as well as to discover useful patterns among the vast amount of data. This involves knowledge from various disciplines such as computer science, industrial and systems engineering, management sciences, operation research, marketing, contracts, and negotiations. It also involves cultural transformation and integration methods based on beliefs, assumptions, principles, and values among organizations and humans. For example, machine learning has been used in recent years for processing and analyzing service-oriented architecture, providing insights to businesses and policymakers for making intelligent decisions. More recently, deep learning technology promises to further revolutionize such processing, leading to better and more accurate results. Deep learning employs software tools from advanced analytics disciplines such as data mining, predictive analytics, text mining, and machine learning based on a set of algorithms that attempt to model high-level abstractions in data by using multiple processing layers with complex structures or nonlinear transformations. As such, the processing and analysis of deep learning applications present methodological and technological challenges together with opportunities. For example, Ishtiaq et al.[8] presented a semi-supervised clustering-based diagnosis recommendation model in healthcare via machine learning techniques based on an unstructured textual dataset.

Referring to intelligent transport systems, Abbas et al.[9] presented a short-term road traffic density prediction based on long short-term memory (LSTM) neural networks. The model is trained by using traffic data collected by the Motorway Control System in Stockholm, that monitors highways and collects flow and speed data per lane every minute from radar sensors based on partitioning the road network into stretches and junctions with one or more LSTM neural networks. On the other side, Duan et al.[10] discussed a neural network-based method to simulate the cognitive process of how human beings read Earth science articles and identify implicitly cited dataset entities from the articles.

Deep learning employs software tools from advanced analytics disciplines such as data mining, predictive analytics, text mining, and machine learning based on a set of algorithms that attempt to model high-level abstractions in data by using multiple processing layers with complex structures or nonlinear transformations. However, the processing and analysis of deep learning applications present methodological and technological challenges. Further, deep learning applications are advantaged by a rise in sensing technologies as witnessed by both the number and rich diversity of sensors ranging from smartphones, personal computers, and health tracking appliances to the Internet of Things (IoT) technologies. These technologies are designed to give contextual, semantic data to entities in a ubiquitous environment that could apply intelligence to decision making. Recently, deep learning technologies have been applied to service-oriented computing. For example, Cai et al.[11] presented an example of applying advanced deep learning techniques on a large-scale, geo-tagged, and image-based dataset measuring urban tree cover using Google Street View (GSV) images to efficiently estimate important urban metrics, particularly in deep convolutional neural networks.

Referring to the brain-computer interfaces, Bellman et al.[12] described an experiment to determine if modern machine learning techniques could be used to accurately detect and classify unaware and aware facial recognition. The experiment consisted of participants viewing a variety of images. Over a period of three phases across two days, participants were first trained on a number of images that they were to implicitly learn for unaware recognitions on the following day. On the second day of the experiment, participants were shown these implicitly learned images, among others including a single memorized face for aware recognition, and then the electroencephalogram signals were recorded for later analysis.[12]

ROBOTIC COMPUTING

Robotic computing is a branch of artificial intelligence (AI) technologies and their synergistic interactions that enable and are enabled by robots. James Kuffner at Google coined the term “cloud robotics” to describe a new approach to robotics that takes advantage of the Internet as a resource for massively parallel computation and real-time sharing of vast data resources. For example, Li et al.[13] investigated the task assignment and scheduling in collaborative cloud robotic systems (CCRS), in which robotic agents can work cooperatively, not only by sharing their processing resources with each other but also by supporting cloud services, making them more intelligent, efficient, and knowledgeable.

CCRS is a technical solution to fulfill complex tasks, such as multi-robot Simultaneous Localization And Mapping (SLAM). However, the challenges stem from not only the computation complexity of large-scale map merging but also the inefficiency of enabling parallel computing in this process, which is indispensable to make available the frontier of computing technology, such as cloud infrastructure. For example, Zheng et al.[14] presented a scalable real-time multi-robot visual SLAM framework based on the cloud robotic paradigm that can distribute the SLAM process to multiple computing hosts in a cluster, which enables map building in parallel. Further, Silva Filho et al.[15] described a robotic platform developed within Baker Hughes, a GE company (BHGE) and GE Global Research Centers (GE-GRC), discussing its use in an industrial inspection case study for remote methane inspection in oilfields.

CONCLUSION

Services computing continues to evolve. We were delighted that the area continues to thrive both in research and in industry. Services computing is integrated into the physical world more today than ever before. The future of this area will continue to be connected to big data in addition to practical applications in artificial intelligence and human-centered computing.

We certainly hope to be having this conversation again in another decade but can only imagine how the area will evolve. This month’s issue has been crafted from a group of papers that overlap human-centered interaction with the Internet and Web Media. We hope you enjoy the papers selected for this issue.

REFERENCES

1. *Big Data Applications and Use Cases*, Patrick C.K. Hung, ed., Springer International Series on Applications and Trends in Computer Science, Springer, 2016.
2. I. Taleb, M.A. Serhani, and R. Dssouli, "Big Data Quality: A Survey," *Proceedings of the 2018 IEEE International Congress on Big Data*, 2018, pp. 166–173.
3. H. Khalajzadeh et al., "A Survey of Current End-user Data Analytics Tool Support," *Proceedings of the 2018 IEEE International Congress on Big Data*, 2018, pp. 41–48.
4. A. Agarwal et al., "Sentiment analysis of Twitter data," *Proceedings of the 2011 Workshop on Languages in Social Media* (LSM 11), 2011, pp. 30–38.
5. K. Hemmings-Jarrett, J. Jarrett, and M.B. Blake, "Sentiment Analysis of Twitter Samples that Differentiates Impact of User Participation Levels," *Proceedings of the 2018 IEEE International Conference on Cognitive Computing*, 2018, pp. 65–72.
6. L. Zhang, S. Wang, and R.N. Chang, "QCSS: A QoE-aware Control Plane for Adaptive Streaming Service over Mobile Edge Computing Infrastructures," *Proceedings of the 2018 IEEE International Conference on Web Services*, 2018, pp. 139–146.
7. Y. Li and S. Wang, "An energy-aware Edge Server Placement Algorithm in Mobile Edge Computing," *Proceedings of the 2018 IEEE International Conference on Edge Computing*, 2018, pp. 66–73.
8. I. Ahmed et al., "Diagnosis Recommendation Using Machine Learning Scientific Workflows," *Proceedings of the 2018 IEEE International Congress on Big Data*, 2018, pp. 82–90.
9. Z. Abbas et al., "Short-Term Traffic Prediction Using Long Short-Term Memory Neural Networks," *Proceedings of the 2018 IEEE International Congress on Big Data*, 2018, pp. 57–65.
10. X. Duan et al., "A Neural Network-Powered Cognitive Method of Identifying Semantic Entities in Earth Science Papers," *Proceedings of the 2018 IEEE International Conference on Cognitive Computing*, 2018, pp. 9–16.
11. B.Y. Cai et al., "Treepedia 2.0: Applying Deep Learning for Large-scale Quantification of Urban Tree Cover," *Proceedings of the 2018 IEEE International Congress on Big Data*, 2018, pp. 49–56.
12. C. Bellman, M.V. Martin, and S. MacDonald, "On the Potential of Data Extraction by Detecting Unaware Facial Recognition with Brain-Computer Interfaces," *Proceedings of 2018 IEEE International Conference on Cognitive Computing*, 2018, pp. 99–105.
13. S. Li et al., "Latency-Aware Task Assignment and Scheduling in Collaborative Cloud Robotic Systems," *Proceedings of the 2018 IEEE 11th International Conference on Cloud Computing*, 2018, pp. 66–72.
14. P. Zhang et al., "Cloud-based Framework for Scalable and Real-time Multi-robot SLAM," *Proceedings of the 2018 IEEE International Conference on Web Services*, 2018, pp. 147–154.
15. R.S. Filho et al., "Semi-Autonomous Industrial Robotic Inspection: Remote Methane Detection in Oilfield," *Proceedings of the 2018 IEEE International Conference on Edge Computing*, 2018, pp. 17–24.

ABOUT THE AUTHORS

Patrick C. K. Hung is a professor at the University of Ontario Institute of Technology. His research interests include services computing, toy computing, robotic computing, security and privacy. Hung received a PhD in computer science from the Hong Kong University of Science and Technology. Contact him at patrick.hung@uoit.ca.

M. Brian Blake is the provost of Drexel University and the editor in chief of *IEEE Internet Computing*. His research interests include service-oriented computing, workflow, and collaboration. Blake received a PhD in software engineering from George Mason University. Contact him at mb3545@drexel.edu.