



## Editorial for the ICMR 2017 special issue

Michael S. Lew<sup>1</sup>

Received: 8 January 2018 / Accepted: 12 January 2018 / Published online: 19 January 2018  
© Springer-Verlag London Ltd., part of Springer Nature 2018

The International Conference on Multimedia Retrieval (ICMR) is the leading ACM multimedia retrieval conference worldwide. In this issue, we present extended versions of the top ICMR papers as selected based upon recommendations from the program chairs (special thanks to the ICMR 2017 program chairs: Jiashi Feng, Rainer Lienhart, Martha Larson and Cees Snoek) and the IJMIR editorial board. The authors were invited to submit extended versions of their work which went through another peer review round. The papers encompass a wide range of important topics in multimedia retrieval which span new areas in content analysis and retrieval to very large scale distributed search to user behavior and social media analysis and are representative of the state-of-the-art.

One of the areas which has seen tremendous growth recently is in the novel design, usage and development of deep neural networks. The paper “Digital Watermarking for Deep Neural Networks” by Yuki Nagai, Yusuke Uchida, Shigeyuki Sakazawa and Shin’ichi Satoh introduces a new problem, that of embedding a digital watermark into the trained neural network. Their findings show that their watermarking approach does not impair the neural network performance and that the watermark is also detectable even after fine-tuning or parameter pruning.

The multimedia retrieval process clearly starts with the user who has a query in mind. The ideal approach for many users would be to allow natural language text queries. The paper “MSRC: Multimodal Spatial Regression with Semantic Context for Phrase Grounding” by Kan Chen, Rama Kovvuri, Jiyang Gao and Ram Nevatia makes significant progress toward answering text queries such as “Big plant on the left” or “Top part of the mountain”. The authors propose a system which applies a spatial regression network (SRN) to predict object locations and a context refinement network (CRN) which encodes context information and uses a novel joint prediction loss to refine the results.

Another user focused research area is in analyzing user behavior and intentions—Why is a user searching for an

item? In the paper “Multimodal Analysis of User Behavior and Browsed Content Under Different Image Search Intents” by Mohammad Soleymani, Michael Riegler and Pal Halvorsen, the authors designed seven different search scenarios involving finding items, re-finding items and entertainment. They trained machine learning systems to predict search intent from the visited images, user interactions and spontaneous responses. Conclusions are given on which features are most effective in classifying search intent.

Multimedia retrieval frequently utilizes the interplay between text and visual information. The paper “Estimating the Information Gap between Textual and Visual Representations” by Christian Henning and Ralph Ewerth presents several contributions which give valuable insight into the possible interrelationships of textual and visual information. Some notable contributions include introducing two measures to describe cross-modal interrelations and also proposing deep learning systems to effectively estimate the measures.

While content analysis is important, another crucial challenge is in adapting the algorithms to very large-scale search. The work “Balancing Search Space Partitions by Sparse Coding for Distributed Redundant Media Indexing and Retrieval” by Andre Mourao and Joao Magalhaes addresses the challenges on a modern distributed search system: load balancing, redundancy on node failure and efficient node usage to answer high-performance queries. The authors propose the balanced KSVD algorithm which distributes data uniformly across codewords. Their findings show that their algorithm does better partition size balancing and that this results in a positive retrieval impact.

Arguably, the most important problem of this era is the detection of misleading content and fake news worldwide because it affects all individuals, societies, countries and governments. Whether the source is a newspaper or a social media message, the editors and readers have a need to know what is credible and what is not. The paper “Detection and Visualization of Misleading Content on Twitter” by Christina Boididou, Symeon Papadopoulos, Markos Zampoglou, Lazaros Apostolidis, Olga Papadopoulou and Yiannis Kompatsiaris presents a system which classifies

✉ Michael S. Lew  
mlew@liacs.nl

<sup>1</sup> Leiden, The Netherlands

Twitter messages as either credible or misleading. In addition to proposing new features, the authors find that the use

of bagging and an agreement-based retraining method results in significant improvements in accuracy.