

Editorial for the ICMR 2016 special issue

Susanne Boll¹ · Winston Hsu² · Jiebo Luo³

Received: 27 January 2017 / Accepted: 30 January 2017 / Published online: 25 February 2017
© Springer-Verlag London 2017

In 2016, the *ACM International Conference on Multimedia Retrieval—ICMR’16*—continued a 17-year tradition of being the premier forum for presenting research results and experience reports on multimedia retrieval research and systems. This annual international conference, started in 2011, inherits the leadership of two former ACM CIVR (from 2002) and ACM MIR (from 2000) series in showcasing the state-of-the-art in the analysis, storage, comparison, retrieval, and presentation of image, video, audio, and text working together on diverse platforms. Of particular interest are developments that investigate multimedia retrieval in practice, including challenges related to scale, effectiveness, efficiency, multi-modal integration, user interfaces, and multimedia retrieval deployment across a range of industries.

Based on the engaging presentations at the conference and the very supporting reviewer comments, we specifically recommend three papers for the special issue of *IJMIR*: “Multilingual Visual Sentiment Concept Clustering and Analysis,” by Pappas et al., “Investigating Country-specific Music Preferences and Music Recommendation Algorithms with the LFM-1b Dataset,” by Markus Schedl, and “Learning Hierarchical Video Representation for Action Recognition,” by Li et al. These papers, representing advanced and diverse themes in the multimedia retrieval community, highlight the timely works for several active aspects ranging from deep learning methods for fine-grained video presentations, rich and multi-facet open dataset for music recommendation

and user profiling, to cross-culture/cross-lingual semantics discovery. The papers presented in this special issue are substantial extensions of the articles presented at ICMR 2016 and underwent a peer-review with international experts before they were finally selected for this special issue.

The paper by Pappas et al., “Multilingual Visual Sentiment Concept Clustering and Analysis,” makes use of the current technology on visual concept detection to match visual sentiments in various languages. The authors carefully designed the analysis method and the annotation of the dataset by crowdsourcing. In addition, they showed an example of the analysis on portraits, with interesting results and insights.

Markus Schedl, in “Investigating Country-specific Music Preferences and Music Recommendation Algorithms with the LFM-1b Dataset,” presents a new dataset with 1 billion music listening events from 120,000 different Last.FM users. LFM-1b has two unique features: (1) its substantial size and (2) a wide range of additional user descriptors that reflect their music taste and consumption behavior. The detailed user- and listening-specific information are important for a variety of research tasks.

In “Learning Hierarchical Video Representation for Action Recognition,” Li et al. attempt to address the problem that different actions may have different granularities, i.e., different lengths in time, by analyzing videos in multiple granularities. The proposed hierarchical framework fuses multiple scores for different granularities by training a classifier on multi-granular distributions.

We would like to thank the authors for their contributions, and the reviewers for their constructive comments that helped strengthen the papers. We are also grateful for the assistance by the editorial office of *IJMIR*.

✉ Susanne Boll
susanne.boll@informatik.uni-oldenburg.de

¹ University of Oldenburg, Oldenburg, Germany

² National Taiwan University, Taipei, Taiwan

³ University of Rochester, Rochester, NY, USA