Best papers in multimedia information retrieval

Michael S. Lew

Published online: 20 February 2013
© Springer-Verlag London 2013

An important area in multimedia information retrieval is event recognition in video retrieval. This area is of particular interest because it directly addresses the grand challenge of bridging the semantic gap between the user and the machine. We open this issue with the survey paper “High-level event recognition in unconstrained videos” by Yu-Gang Jiang, Subhabrata Bhattacharya, Shih-Fu Chang and Mubarak Shah, which presents a comprehensive review, breakdown and analysis of the current and critical technologies in state-of-the-art high-level event recognition systems.

On behalf of the editorial board, it is an honor to present extended versions of the four best conference papers from the ACM multimedia retrieval research community. The papers were selected based on the reviews of the program committee by the ACM ICMR program chairs, Alan Smeaton, Alex Hauptmann, and Chong-Wah Ngo and the ACM MIR program chairs, Nuria Ramirez and Apostol Natsev.

Due to the prevalence of 3D objects in architecture, CAD applications, computer games, and numerous internet applications, 3D object retrieval has grown quickly in societal importance. The paper “3D object retrieval using salient views” by Indriyati Atmosukarto and Linda G. Shapiro presents a novel approach based on silhouettes which is both competitive in accuracy with the top research community methods and also significantly higher performance.

A growing trend in the research community is to use very large external knowledge sources to improve the search process. The paper “Exploiting semantics on external resources to gather visual examples for video retrieval” by David Vallet, Iván Cantador and Joemon M. Jose empirically evaluates the feasibility of external knowledge sources such as Flickr, DBpedia and Google in the context of example-based video search.

While the search aspect of a retrieval system is of clear importance, an equally important element is summarizing the multimedia information. In the paper “Beyond Audio and Video Retrieval: Towards Multimedia Summarization” by Florian Metze, Duo Ding, Ehsan Younessian and Alexander Hauptmann, the authors present a novel multi-modal approach that automatically summarizes multimedia content by generating a representative paragraph of natural language.

Understanding the user and the user’s relationship to the information in the search process is a major challenge. The paper “Content Analysis meets Viewers: Linking Concept Detection with Demographics on YouTube” by Adrian Ulges, Damian Borth and Markus Koch, investigates the mapping between multimedia content and viewer demographics. It also introduces the notion of a demographic signal which may be useful to disambiguate concept detection in cases of similar concepts.

I would like to thank the authors and the program chairs from the ACM multimedia retrieval conferences for their contributions and collaboration. I also want to thank the multimedia retrieval community for their strong support of this journal.

M. S. Lew (✉)
Leiden Institute of Advanced Computer Science, Leiden University,
Niels Bohrweg 1, 2333 CA Leiden, The Netherlands
e-mail: mlew@liacs.nl

Michael S. Lew
Editor-in-Chief
International Journal of Multimedia Information Retrieval