

---

## CONFERENCE REPORT

---

# ICMR 2014: 4<sup>th</sup> ACM International Conference on Multimedia Retrieval

**Martin Halvey**

University of Strathclyde, UK  
[martin.halvey@strath.ac.uk](mailto:martin.halvey@strath.ac.uk)

**Philip McParlane**

University of Glasgow, UK  
[p.mcparlane.1@research.gla.ac.uk](mailto:p.mcparlane.1@research.gla.ac.uk)

**Joemon M. Jose**

University of Glasgow, UK  
[jj@dcs.gla.ac.uk](mailto:jj@dcs.gla.ac.uk)

**Keith van Rijsbergen**

University of Glasgow, UK  
[keith@dcs.gla.ac.uk](mailto:keith@dcs.gla.ac.uk)

**Stefan Rueger**

The Open University, UK  
[s.rueger@open.ac.uk](mailto:s.rueger@open.ac.uk)

**R. Manmatha**

University of Massachusetts,  
Amherst, USA  
*emails*  
[manmatha@cs.umass.edu](mailto:manmatha@cs.umass.edu)

**Mohan Kankanhalli**

National University of Singapore,  
Singapore  
[mohan@comp.nus.edu.sg](mailto:mohan@comp.nus.edu.sg)

## 1 Introduction

ICMR was initially started as a workshop on challenges in image retrieval (in Newcastle in 1998) and later transformed into the Conference on Image and Video Retrieval (CIVR) series. In 2011 the CIVR and the ACM Workshop on Multimedia Information Retrieval were combined into a single conference that now forms the ICMR series. The 4<sup>th</sup> ACM International Conference on Multimedia Retrieval took place in Glasgow, Scotland, from 1–4 April 2014. This was the largest edition of ICMR to date with approximately 170 attendees from 25 different countries. ICMR is one of the premier scientific conference for multimedia retrieval held worldwide, with the stated mission “to illuminate the state of the art in multimedia retrieval by bringing together researchers and practitioners in the field of multimedia retrieval.” According to the Chinese Computing Federation Conference Ranking (2013), ACM ICMR is the number one multimedia retrieval conference worldwide and the number four conference in the category of multimedia and graphics. Although ICMR is about multimedia retrieval, in a wider sense, it is also about automated multimedia understanding. Much of the work in that area involves the analysis of media on a pixel, voxel, and wavelet level, but it also involves innovative retrieval, visualisation and interaction paradigms utilising the nature of the multimedia —be it video, images, speech, or more abstract (sensor) data. The conference aims to promote intellectual exchanges and interactions among scientists, engineers, students, and multimedia researchers in academia as well as industry through various events, including a keynote talk, oral, special and poster sessions focused on research challenges and solutions, technical and industrial demonstrations of prototypes, tutorials, research, and an industrial panel. In the remainder of this report we will summarise the events that took place at the 4<sup>th</sup> ACM ICMR conference.

## 2 Venue and Social Program

ICMR took place in Glasgow Scotland. The main conference was situated at the University of

---

---

Glasgow ([www.gla.ac.uk](http://www.gla.ac.uk)) The University of Glasgow was founded in 1451 and is the fourth-oldest university in the English-speaking world. The workshops and tutorials took place in the School of Computing Science, with the remainder of the events situated in the Western Infirmary Lecture Theatre in the main university campus.

As well as conference activities at the university a number of social events were held in other parts of the city. On Tuesday April 1<sup>st</sup> an optional private tour of the Auchentoshan whisky distillery (<http://www.auchentoshan.com/>) was organised. This tour allowed attendees to taste one of Scotland's most famous export and also learn more about the distillation process. A welcome reception took place in Glasgow City Chambers in George Square in Glasgow City Centre on Wednesday April 2<sup>nd</sup>. This building is the headquarters of Glasgow City Council and is also a fine example of Victorian architecture. The city council provided the reception as well as a welcome from the bailie on behalf of the Lord Provost (Mayor) of Glasgow. The conference dinner took place on Thursday April 3<sup>rd</sup> Glasgow Science Centre (<http://www.glasgowsciencecentre.org/>), located beside the River Clyde. The attendees were greeted by a bagpiper before a drinks reception where the attendees had an opportunity to try out the museum exhibits. A dinner followed highlighting the best of Scottish cuisine, including haggis. During the conference dinner the best paper awards were presented. Photographs from the conference and social events are available online - <https://www.flickr.com/photos/122120613@N04/>




Figure 1: Welcome reception at Glasgow City Chambers



Figure 2: Conference dinner at Glasgow Science Centre

### 3 Tutorials and Workshops

The ICMR 2014 workshops and tutorials took place on Tuesday April 1<sup>st</sup> 2014. There were 2 tutorials and 3 workshops, this was the largest number of these types of events at ICMR to date. Over 100 attendees attended the workshops and tutorials.

---

---

The tutorial “Visual Recognition in Large Collections” was delivered by Hervé Jégou, INRIA, France. The tutorial focused on visual recognition in large collections of images, paying particular attention to image descriptions, algorithms for efficient similarity search, complementary techniques and considering different trade-offs in search. The tutorial “Music Information Retrieval – Theory and Applications” was delivered by George Tzanetakis, University of Victoria, Canada. This tutorial provided the theoretical overview of the state-of-the-art in music information retrieval followed by a practical hands-on demonstration of several existing tools and resources that are used for research in this music information retrieval.

The workshop on Environmental Multimedia Retrieval looked at the need for the development of advanced techniques for analysing, interpreting and aggregating environmental data provided in multimedia formats. The Social Events in Web Multimedia workshop covered the topic of modeling, detection and processing of events; in particular it solicited papers on the latest developments and results on the discovery of social events from web multimedia content, and on techniques for the detection and retrieval of media items that are related to such social events. The workshop proceedings are available online - [http://mklab2.itl.gr/sewm14/wp-content/uploads/2014/03/SEWM\\_2014\\_Proceedings.pdf](http://mklab2.itl.gr/sewm14/wp-content/uploads/2014/03/SEWM_2014_Proceedings.pdf) The Social Multimedia and Storytelling workshop covered a number of aspects of multimedia retrieval, including content analysis and understanding, search and retrieval, HCI, image/video summarisation and visualisation etc. around the area of social multimedia and storytelling. The proceedings of the workshop are available online - <http://ceur-ws.org/Vol-1198/>

#### **4 Main Technical Program**

The main technical program of the conference consisted of a keynote, full research papers, short research papers, demonstrations and special technical sessions. The conference was opened with a keynote by Prof. Susanne Boll. Prof. Boll is Professor of Media Informatics and Multimedia Systems in the Department of Computing Science at the University of Oldenburg, in Germany. She also serves on the executive board of the OFFIS Institute for Information Technology, in Oldenburg. The keynote covered user experience and multimedia content.

In terms of the main content of the conference in recent years multimedia retrieval research has matured and become a distinct field in its own right and the papers of the main track covered many of the core topics of the field. For example ICMR 2014 had papers on automated image annotation that made use of innovative cross-media models, tackled large classes of images, and utilised social activities, all while reporting better image annotation performance than ever before. Events in videos have also become a focus of attention, and related research explores detecting events in unconstrained videos and using detected events to support semantic concept discovery. The complexity of these tasks calls for complex mathematical modeling. ICMR 2014 included papers focusing on dedicated methods from machine learning and statistics via graph theory to data mining. ICMR has become the premier place to present, or discover, new models that make progress for automated multimedia analysis with a view to improve multimedia search engines. At the other end of the spectrum from theory and practice is work that exploits the user and capitalises on intuitive interfaces that allow the exploration of multimedia collections—whether they contain videos, celebrities, 3D models, citizen science data, or social media TV snippets. Much of this work was shown in short paper and demos at ICMR 2014. The full range of topics is illustrated in the Scotland shaped tag cloud in Figure 3.

---



**Figure 3: Main topics from ICMR 2014 papers, keywords are taken from the abstracts of accepted papers**

In terms of numbers, we received 111 valid submissions for the full regular research papers of which 21 were accepted for oral presentation (18.9% acceptance rate) and 21 were accepted for poster presentation (giving an overall acceptance of 42 papers with 37.8% acceptance rate). For short papers, we received 56 submissions, of which 27 were accepted for poster presentation (48.2% acceptance rate). We also had two special sessions Harvesting and Analyzing Live Social Media Data (5 papers accepted out of 10 submissions) and User-centric Video Search and Hyperlinking (3 papers accepted out of 9 submissions). The submissions to the special sessions were organised separately by the special organisers and are not included in the other submission figures above. Further, we had 13 technical demonstrations of systems showcasing key research contributions. The demonstrations and poster presentations were split across two dedicated sessions, these sessions were also accompanied by a wine and cheese reception to help make the event more social (see Figure 4).



**Figure 4: Poster and demonstration session from the last day of the conference**

From the main technical program 4 papers were selected to be considered for the best paper award and these were presented in a best paper session. The nominated papers were:

- Subhabrata Bhattacharya, Felix Yu and Shih-Fu Chang Minimally Needed Evidence for Complex Event Recognition in Unconstrained Videos
- Sean Moran and Victor Lavrenko Sparse Kernel Learning for Image Annotation
- Lu Jiang, Wei Tong, Deyu Meng and Alex Hauptmann. Towards Efficient Learning of Optimal Spatial Bag-of-Words Representations
- Winn Voravuthikunchai, Bruno Cremilleux and Frederic Jurie. Image re-ranking based on statistics of frequent patterns

The best paper award was awarded to “Minimally Needed Evidence for Complex Event Recognition in Unconstrained Videos” by Subhabrata Bhattacharya, Felix Yu and Shih-Fu Chang from Columbia University USA. The best student paper award was awarded to “Sparse Kernel Learning for Image Annotation” by Sean Moran and Victor Lavrenko from the University of Edinburgh, Scotland.

## **5 Industry Day**

On Friday, 4<sup>th</sup> April, the industry day opened with a panel drawn for industry and academia

---

---

discussing the past, present, and future of multimedia retrieval. The panel members were Alberto Del Bimbo (University of Firenze, Italy), Nicu Sebe (University of Trento, Italy), Hervé Jégou (INRIA, France), Alejandro Jaimes (Yahoo! Research) and Rogerio Ferris (IBM). Following the panel some industry representatives also gave talks on recent advancement in product, tools, and services. Chordify's Bas de Haas showed how music analysis and retrieval can help to play the chords of your favourite music song. Rogerio Ferris of IBM talked about the latest results in multimedia surveillance applications; Alexandr Krainov of Yandex, Russia talked about Yandex video retrieval techniques; Matthew Cooper of FXPAL discussed the over emphasis of quantitative evaluation and interesting augmented reality; and Pieter van der Linden of Technicolor, the Netherlands, talked about the latest advances made through the Quaero program.

## **6 Sponsorship and Student Support**

ICMR 2014 was generously supported by a number of organisations. FXPAL, Springer and AXES provided sponsorship for the main conference. The Scottish Informatics and Computer Science Alliance (SICSA) provided fees for 10 students from Scottish universities to attend ICMR, these fees were used to cover the costs of student volunteers from 4 Scottish universities. The SICSA HCI theme sponsored the workshop and tutorial day which allowed attendees from SICSA institutions to attend those events for free. ACM through SIGMM and ACM SIGIR provided sponsorship which was used to provide 10 travel grants for students travelling from overseas. 36 student travel grant applications were received, preference was given to students who had work accepted to ICMR 2014 and considerations were also taken to ensure that the grants went to as diverse a group of students as possible. The welcome reception was provided free of charge by Glasgow City Council.

## **7 Concluding Remarks**

ICMR 2014 was highly successful and the largest ICMR event so far, with 170 attendees from 25 countries, 42 full papers, 27 short papers, 2 special sessions, 2 tutorials and 3 workshops. The growth of the conference demonstrates the growth of the community and the interest in multimedia information retrieval. As well as the academic sessions in the conference there were extensive break out and social sessions to foster discussion amongst the attendees. We are extremely grateful to all of the attendees and every single person who was involved in the organization of the conference as without them it would not have been a successful event. ICMR 2015 will take place in Shanghai, China and promises to be every bit as interesting and exciting as ICMR 2014.

---