

Overview

The 4th ACM International Conference on Multimedia Retrieval (ICMR) took place on April 1st – 4th 2014 at the University of Glasgow. The event was co-organised by staff at the University of Glasgow and Glasgow Caledonian University. The focus and aim of this conference was to facilitate the exchange of leading-edge multimedia retrieval ideas among researchers, practitioners and other potential users of multimedia retrieval systems. The first day (April 1st) of the conference was the workshop and tutorial day. In total there were 2 tutorials and 3 workshops, which are a record number for ACM ICMR, a short abstract for each workshop is provided later in this document. The cost to attend the workshops and tutorials was £150 per person, to assist with this the SICSA HCI theme provided £1000 to provide free registration for SICSA institutions. In total 16 people from 5 SICSA institutions availed of the free places (University of Glasgow 3, Glasgow Caledonian University 4, University of Strathclyde 1, University of Edinburgh 1, and Robert Gordon University 7). Each workshop and tutorial was attended by at least 1 SICSA attendee (we allowed attendees to move between events so conceivably people could attend more than one) the breakdown of registered events is as follows:

- Tutorial: Music Information Retrieval - Theory and Applications 5
- Tutorial: Visual recognition in large collections 2
- Workshop: Social Multimedia and Storytelling 6
- Workshop: Social Events in Web Multimedia 2
- Workshop: Environmental Multimedia Retrieval (EMR) 1

As well as attending on the day each attendee was also given a USB stick with the full proceedings for the entire conference. A brief abstract outlining each workshop and tutorial follows.

Workshops

Environmental Multimedia Retrieval (EMR) (<http://mklab2.itι.gr/emr2014/>)

The rapid advancements of digital technologies, as well as the progress and wide availability of digital cameras and sensors have resulted in a great increase of multimedia data production worldwide. This is also the case for multimedia data that describe the state of the environment, which include huge amounts of data streams from model systems, dedicated stations and amateur sensors, as well as visual environmental information, such as heatmaps and forest satellite images. In parallel, the success of citizen sciences and social networking tools has fostered the emergence of large and structured communities of nature observers (e.g. e-bird, xeno-canto, Tela Botanica, etc.), who started to produce outstanding collections of biodiversity multimedia records. Citizens have become increasingly aware of the important role that environmental data (e.g. weather forecast, air quality, life species distributions) play on health issues (e.g. allergies), as well as to a variety of other human activities (e.g. agriculture, trip planning). In addition, such data are very important for environmental issues and phenomena, such as the greenhouse effect, the global warming and the climate change.

Social Events in Web Multimedia (<http://mklab2.itι.gr/sewm14/>)

The modeling, detection and processing of events is an area that has started to receive considerable attention by the multimedia community. This Workshop aims to attract and present 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. By

social events, we refer to events that are planned by people, are attended by people, and are represented by multimedia and in general social media content that is captured by people.

Social Multimedia and Storytelling

[\(https://sites.google.com/site/socialmultimedystorytelling/\)](https://sites.google.com/site/socialmultimedystorytelling/)

The increasing availability of social multimedia content around real-world objects and activities holds great potential for deriving richer representations of the depicted places, experiences and events. This is not only due to the abundance of diverse multimedia content, but also due to the availability of a large variety of contextual information, such as location and textual metadata, and online interactions (e.g. comments/likes on shared videos). Therefore, leveraging such content and its surrounding context holds the potential for better capturing and representing the real world and for building novel engaging applications. However, the nature of social multimedia raises a number of research challenges that call for new multimedia retrieval approaches. To this end, the workshop seeks research contributions that address different stages of the lifecycle of social multimedia content, from the moment an event is captured in media, through its online sharing, collection, and processing to its remixing, repurposing, retrieval, and presentation. Of particular importance are submissions that cross-cut different research topics of the workshop in creative ways, and that propose new engaging applications with strong multimedia and social networking elements.

Tutorials

Visual recognition in large collections (Hervé Jégou, INRIA, France)

The tutorial covers state-of-the-art methods and systems for visual recognition in large collections of images, considering different trade-offs with respect to efficiency and search quality. The tutorial first focuses on three important ingredients of such systems: 1) The image description itself, including the bag-of-words model, Fisher vectors, and more generally match kernels; 2) Algorithms for efficient similarity search, in particular approximate nearest neighbour search, compressed-domain search and Min-hash; 3) Complementary techniques, aiming in particular at exploiting the spatial information on a large scale and query expansion. Then, the tutorial gives a brief overview of the problem of automatic discovery of visual patterns or related images in unstructured image collections. The tutorial finally describes a recent approach for precise search in very large collections of videos with limited resources.

Music Information Retrieval - Theory and Applications (George Tzanetakis, University of Victoria, Canada)

The goal of this tutorial is to provide a thorough 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 can be used for research in this area. Specific emphasis will be given on how MIR techniques relate to other fields of current multimedia research. MIR is an inherently interdisciplinary area touching on several research areas such as digital signal processing, machine learning, perception, visualization, human computer interaction, content-based retrieval and digital libraries. Music has several unique characteristics that differentiate it from other areas of multimedia research. The different problems and techniques proposed to solve them that will be described in the first part of the tutorial will be followed with concrete practical examples of applying these techniques using existing software tools and datasets.