

Programme (final):

Day 1 – October 21	
9:00 – 9:15	Workshop opening Fausto Giunchiglia , University of Trento, Italy, Sang "Peter" Chin , Johns Hopkins University, Maryland, US.
Section 1. Social media and events Chair: Francesco De Natale	
9:15 – 10:00	<p>“Event-based Summarization for Media Hyperlinking”</p> <p>Benoit Huet, EURECOM, France</p> <p>Abstract: “The exponential growth of media sharing and demand on the Web comes with a need for effective methods to explore them. Hence, media hyperlinking, which consists in linking together videos based on their content, uncovering the relation between them, is becoming an important functionality for providing users with a way to navigate between video entities and satisfy their information needs. Thanks to such technology, multimedia search can often be replaced by recommendation. A particular usage of hyperlinking is to provide, through a second screen application, extra information or content about the video watched on a main screen (TV). In this talk, we will focus on media hyperlinking from the news: the task at hand consists in locating and identifying relevant media items, and display them on the second screen. The related material is selected based on underlying events that will be detected in the news: events are seen as structuring elements, defined in terms of date, location, intent and attendance. Two approaches for event-based mining of such additional and related information will be presented. Each of them satisfying a different user information need.”</p>
10:00-10:45	<p>“Towards Smart Social Systems”</p> <p>Ramesh Jain, University of California, Irvine, US</p> <p>Abstract: “Availability of enormous volumes of heterogeneous Cyber-Physical-Social (CPS) data streams may allow design and implementation of networks to connect various data sources to detect situations with little latency. In fact, in many cases it may even be possible to predict situations well in advance. This opens up new opportunities in designing smart social systems for specific tasks. Such systems may be very useful for many important problems at local as well as regional and even global level. We believe that such systems offer many novel challenges to researchers in multimedia, particularly in social and cross-modal media systems. We will present our ideas and early approach towards building smart social systems.”</p>
11:00-11:30	<i>coffee break and networking</i>

11:30-12:15	<p>“Insights from Big Data: Interaction, Design, and Innovation”</p> <p>Alejandro Jaimes, Yahoo! Research-Barcelona, Spai</p> <p>Abstract: “In recent years, our ability to process large amounts of data has increased significantly, creating many opportunities for innovation. Having large quantities of data, however, does not necessarily turn into actionable insights that make a difference for users in consumer applications. In this talk I will give a quick overview of some ways in which “big data” can be used in industry, with a particular focus on Human-Centered approaches to innovation. In particular, I will discuss how the combination of qualitative and quantitative methods can be of benefit, giving examples around social media and giving an overview of some of the areas of research I am currently focusing on at Yahoo!. Within this context, I will outline a blueprint for a research framework as it applies to innovation, and discuss specific technical approaches within that framework. I will argue on the importance of taking a human-centered view and highlight what I consider the most fundamental problems in computer science today from that perspective.”</p>
12:15-13:00	<p>“Understanding Events and Message Popularity in Media-rich Social Networks”</p> <p>Lexing Xie, Australian National University, Australia</p> <p>Abstract: “Multimedia is growing to take up more than 50% of the internet traffic. Understanding these content and their social traces presents new research challenges and opportunities at the intersection of rich-media content understanding and mining the social web. Several recent work in my group focuses on analyzing real-world event traces in social media, including: use hyperlink patterns to diffusion flow about news events, track large-scale video remix on Youtube, analyzing rich-media microblogs with cross-media topic model, and predicting user preference with fine-grained social interactions. I will share current results in mapping the macro-structure of the event web, and predicting message popularity from social and content features.”</p>
13:00-14:30	<i>lunch break</i>

Section 2. Event indexing and summarization
Chair Pil Ho Kim

14:30-15:15	<p>"Classifying Images and Videos by Learning from Web Data" Jiebo Luo, University of Rochester, US</p> <p>Abstract: "Everyday, increasingly rich and massive social multimedia data are being posted to the web. Such image and video data are generally accompanied by rich and valuable contextual information (e.g., tags, categories, and captions). Given any textual query (e.g., picnic), keywords (also called tags) based search can be readily used to collect a large number of relevant Flickr images or YouTube videos for classifying new images and videos. In the first part of our talk, we will introduce a visual event recognition framework for consumer videos by leveraging a large amount of loosely labeled web videos (e.g., from YouTube). At its core, we develop a new domain adaptation method, referred to as Adaptive Multiple Kernel Learning (A-MKL), in order to 1) fuse the information from multiple pyramid levels and features (i.e., space-time features and static SIFT features) and 2) cope with the considerable variation in feature distributions between videos from two domains (i.e., web video domain and consumer video domain). Extensive experiments demonstrate the effectiveness of our proposed framework that requires only a small number of labeled consumer videos by leveraging web data. In the second part of our talk, we will describe a new approach to learn a robust classifier for text-based image retrieval (TBIR) using relevant training web images (e.g. from Flickr), in which we explicitly handle noise in the loose labels of training images. Specifically, we first partition the relevant training web images and the randomly selected irrelevant training web images into clusters. By treating each cluster as a "bag" and the images in each bag as "instances", we formulate this task as a multi-instance learning problem with constrained positive bags, where each positive bag contains at least a portion of positive instances. We present a new algorithm called MIL-CPB to effectively exploit such constraints on positive bags and predict the labels of test instances (images). Comprehensive experiments on two challenging real-world web image data sets demonstrate the effectiveness of our approach. Finally, we will discuss several future directions on how to effectively and efficiently exploit the freely available web data for visual recognition with minimal human supervision."</p>
15:15-16:00	<p>"Event Mining in Social Multimedia". Symeon Papadopoulos, CERTH-ITI, Greece</p> <p>Abstract: "The presentation will discuss different approaches for social event detection on large collections of user-contributed multimedia content. Social events are defined as real-world events that are planned and attended by people and that are represented by media content captured by people attending them. Two main event detection settings will be presented: (a) a discovery scenario, where events of all types are of interest, and (b) an event detection scenario, where specific types (or classes) of events are sought. Approaches and insights will be presented for both settings, e.g. for the discovery of events in large media collections, as well as for the detection of events of given types. Supervised learning and clustering constitute the main components of these approaches. Several case studies and evaluation results will be presented using Flickr as the source of social media content. Notably, insights will be presented from the participation of the presenter in the two Social Event Detection contests (in the context of MediaEval '11 and '12), and an outline will be provided of pertinent research challenges and future work in this area."</p>

16:00-16:45	<p>“Event Duality: Exploitation of Personal and Social Dimensions for Photo Indexing”</p> <p>Ivan Tankoyeu, University of Trento, Italy</p> <p>Abstract: “Within the scope of the talk I will introduce the distinction between social and personal events. Following this strategy I will describe separate techniques for mining events from multimedia streams. These techniques solely exploit spatiotemporal metadata for detection personal and social events. Further analysis of detected events allows us to unveil interesting and useful information about an owner of a photo collection.”</p>
-------------	---

Day 2 – October 22	
9:00 – 9:15	Workshop opening
Section 3. Event semantics and modeling Chair: Symeon Papadopoulos	
9:15 – 10:00	<p>“Events in Multimedia: Theory, Model, and Application” Ansgar Scherp, University of Mannheim, Germany</p> <p>Abstract: “We introduce the notion of events and objects. While events are said to occur or happen (i.e., they extend over time), objects are said to exist (and unfold in space). Events and their objects allow for representing human experiences and can be related in manifold ways. Implementing this theory of events in a formal model for representing events and event relations enables for a better interoperability of distributed multimedia event-based systems. Our formal model provides comprehensive support to represent time and space, objects and persons, as well as mereological, causal, and correlative relationships between events. In addition, the model provides extensible means for event composition, modeling event causality and event correlation, and representing different interpretations of the same event. Selected features of the model will be presented and discussed. Finally, a mobile event-based application will be presented that implements an instance of an event model for social media data. The application allows for exploring events such as concerts, weekly markets, opening hours, etc. and at the same time explore places such as sights, restaurants, organizations, and persons extracted from different sources. The mobile client retrieves the data through a proxy server, which applies an incremental matching algorithm that integrates complementing information as well as eliminates duplicates from the social media sources.”</p>
10:00-10:45	<p>“Semantics and modeling of events and contexts” Opher Etzion, IBM Research Lab Haifa, Israel</p> <p>Abstract: “People are event-driven creatures, a lot of our daily behaviour is reaction to event we observe or infer; in contrast computerized applications are mainly follow the request-response paradigm (the computer responds to explicit request by human). The availability of real-time data based on the Internet of Things and mobile devices, and the pressure to increase business velocity and provide real-time analytics decisions and actions, are the roots of a paradigm shift towards event-driven computing. In this talk we will concentrate around two main concepts: situation and context. Situation is a (possibly derived) event that requires a reaction, while context is a (possibly multi-dimensional) condition that provides semantic partitions over the flowing events. The first part of the talk drills down to the modelling aspects of deriving situations from events using event patterns, and discuss the evolution of modelling schemes; the second part of the talk discusses the different dimensions of context: temporal, spatial, segmentation and states, and shows examples of hybrid event-state oriented contexts.”</p>
11:00-11:30	<i>coffee break and networking</i>

11:30-12:15	<p>“Discovering Event Media Semantics using Games with a Hidden Purpose” Francesco De Natale, University of Trento, Italy</p> <p>Abstract: “Automatic tools that allow discovering the semantics of a media object from its content show intrinsic limitations, due to the fact that current image content description and recognition approaches still suffer of a rather limited accuracy. The possibility of outsourcing part of these tasks to user crowds, exploiting the power of human computation, has been explored by various researchers, either for directly handling the problem or to produce a ground-truth for further elaboration by means of machine learning approaches. Although a well-designed crowdsourcing mechanism can provide good results, it is not time effective and requires investments for every job launched. In this talk we will introduce a different approach to achieve human cooperation in complex media analysis tasks, with the introduction of specifically designed games with a hidden purpose. In detail, we will show how one can produce a competitive game in which the evident goal and reward is simply entertain, playing and possibly winning matches against other players and gaining reputation, while the hidden purpose is to produce new knowledge on the media objects handled within this contests. A couple of examples will be presented, one conceived to detect event-related salient areas in event media, and the other designed to propagate the annotation across images with related contents. Tests will be presented for both games to demonstrate the viability of these approaches in solving complex tasks.”</p>
12:15-13:00	<p>“Five Recommendations for Recognizing Video Events by Concept Vocabularies” Cees G.M. Snoek, University of Amsterdam, Netherlands</p> <p>Abstract: “Representing videos using vocabularies composed of concept detectors appears promising for generic event recognition. While many have recently shown the benefits of concept vocabularies for recognition, studying the characteristics of a universal concept vocabulary suited for representing events is ignored. In this talk, we present the findings of a study on how to create an effective vocabulary for arbitrary-event recognition in web video. We consider five research questions related to the number, the type, the specificity, the quality and the normalization of the detectors in concept vocabularies. From the analysis we derive a set of five recommendations for recognizing video events by concept vocabularies, which provide guidelines for future work.”</p>
13:00-14:30	<i>lunch break</i>
14:30-16:00	<p>"Future trends in events for media" Panel session</p> <ul style="list-style-type: none"> • coordinator: Bogdan Ionescu • events and media: Nicu Sebe • event summarization: Benoit Huet • social media and events: Lexing Xie • event modeling: Ansgar Scherp • event semantics: Opher Etzion
16:00-16:15	<p>Closing remarks Fausto Giunchiglia, University of Trento, Italy Sang "Peter" Chin, Johns Hopkins University, Maryland, US</p>