

# Special Section on Social Media as Sensors

**Abstract**—Social media have transformed the Web into an interactive sharing platform where users upload data and media, comment on, and share this content within their social circles. Each content item is associated with an abundance of metadata and related information such as location, tags, comments, favorites and mood indicators, access logs, and so on. At the same time, all this information is implicitly or explicitly interconnected based on various properties such as social links among users, groups, communities, and sharing patterns. These properties transform social media into data sources of an extremely dynamic nature that reflect topics of interests, events, and the evolution of community opinion and focus. Social media processing offers a unique opportunity to structure and extract information and to benefit multiple areas ranging from computer vision to psychology and marketing. The objective of this special section is to provide an overview of the current research in emerging topics related to applications where social media can act as sensors of real-life phenomena and case studies that reveal valuable insights usually not possible with existing, limited, controlled, and laboratory-based datasets.

**Index Terms**—Social media, social media analytics, social media sensors, social sensing, crowdsourcing.

## I. INTRODUCTION

SOCIAL media sharing websites, such as Twitter, Flickr, and Instagram, host billions of text messages, images, and videos, which have been annotated and shared among friends, or published in groups that cover a specific topic of interest. The fact that users and devices annotate the content in the form of tags, ratings, preferences, social connections, location, time, and so on, and that these activities are constantly updated, gives such social media data sources an extremely dynamic nature that reflects topics of interests, events, and the evolution of community opinion and focus.

However, the current technologies have mostly focused on producing large volumes of such media and on the piecemeal consumption and processing of any of such tweets, images, or songs. It is the utilization of aggregated collections of such media and the combination of their visual, textual, temporal, location, social, and affective context, that provides the ingredients for a more thorough understanding of the various events, patterns, and situations embedded or emerging in such data. Therefore, there is a need for scalable and multi-modal approaches that are able to handle the massive amount of available data and transform Web 2.0 to a sensor of topics, events, and contents and ultimately facilitate planning, prediction, and action, for example, in the form of recommendations.

Such approaches can be applied to various areas including computer vision, machine learning, data mining, and multimedia fields, where social media sites can be used as a source of huge amounts of weakly labeled data for solving large-scale problems. In personalization and retrieval applications, social-based indexing and user profiles can incorporate information about the structure and activity of the users social network directly into the personalization, search, and ranking processes. They can also be applied to address problems in

fields such as tourism and culture, social sciences, politics, economics, and marketing. In tourism and culture, uploaded media can reveal off-the-beaten-path points of interest and events, otherwise difficult to discover through usual Web sources. Social media can have important contributions to scientific fields such as computational sociology and psychology for the development and testing of theories of complex social processes through bottom-up modeling of social interactions. Furthermore, it is possible to explore and predict trends by analyzing the sharing patterns of uploaded and downloaded social multimedia. For example, in politics and news, Web data may contain the most up-to-date information that can reflect the popularity of elections candidates and their support among populations.

Social media analysis approaches open a brand-new arena for the research community with a potentially big impact on multimedia research. This Special Section targets a mixed audience of readers from several communities, i.e., multimedia analysis, data mining, machine learning, information extraction, social networks, large-scale techniques, computational psychology, and sociology. This Special Section attempts to present a representative sample of ongoing research focusing on social media analysis approaches that exploit explicit and implicit features and connections in order to extract useful information for a variety of applications ranging from tweet analysis for news to estimating the correlation between users' mood, social activity, and personality. The Special Section Call for Papers received a strong response from the community. A total of 28 manuscripts were submitted for consideration. Of these submissions, 8 papers were accepted following a rigorous review process coordinated by the guest editors. We are grateful to all the authors for their contributions and to the reviewers for their effort to ensure the highest possible quality in all the accepted papers.

In *Learning Crowdsourced User Preferences for Visual Summarization of Image Collections*, Rudinac *et al.* demonstrate the value of crowdsourcing in understanding how observers summarize image sets. The authors apply this knowledge as part of an automatic image selection method for visual summary generation using Flickr images, which integrates content, context, and sentiment from associated image comments.

In *Tracking Large-Scale Video Remix in Real-World Events*, Xie *et al.* investigate the phenomenon of visual memes, i.e., video segments that are often repurposed on sites like YouTube and that emerge as a trend when world-events unfold. The authors discuss methods to extract them and to address the problem of estimating meme popularity, using real-world events to illustrate their approach.

Also related to online video, in *Towards Cross-Domain Learning for Social Video Popularity Prediction*, Roy *et al.* investigate cross-domain relations between videos in YouTube and messages in Twitter from the perspective of popularity. The authors present a transfer learning technique that uses knowledge extracted from tweets to classify YouTube videos as bursty.

Aiello *et al.* in *Sensing Trending Topics in Twitter* compare six topic detection methods on three Twitter datasets related to

major real-world events, which differ in their time scale and topic churn rate. Performance not only depends on the detection method used but also on the specific nature of the data stream. They show that standard natural language processing techniques can perform well for social streams on very focused topics, but novel techniques designed to mine the temporal distribution of concepts are needed to handle more heterogeneous streams containing multiple stories evolving in parallel.

In *Travel Recommendation by Mining People Attributes and Travel Group Types from Community-Contributed Photos*, Chen *et al.* address the problem of generating personalized travel recommendations by exploiting large-scale photo sites like Flickr. Their recommendation approach combines individual demographic attributes extracted from user-generated geo-tagged images along with social link information.

In *Inferring Contexts from Facebook Interactions: A Social Publicity Scenario*, Servia-Rodriguez *et al.* propose to infer user social contexts from textual evidence of their interactions on Facebook using NLP and data mining techniques. The knowledge of these social contexts allows to build an application to benefit both companies and potential customers such that the former can promote their products among the most likely buyers and the latter only receive ads of interest.

In *Quantitative Study of Music Listening Behavior in a Social and Affective Context*, Yang *et al.* use data from LiveJournal and Last.fm sites in order to study the correlation between user mood, the affective content of music, and individuals' personality. A rich set of audio and text features (from both blogs and song lyrics) are extracted to establish the tripartite relationship between music emotion, user mood, and music listening behavior, taking also into account the personal traits of the bloggers based on a text-based Personality Recognizer. Results show that social functions of music can be well explored from a real-life, in-situ data set.

In *Connectivity, Online Social Capital and Mood: A Bayesian Nonparametric Analysis*, Phung *et al.* investigate the formulation of online social capital, and establish the link between online social capital defined from social connectivity and user mood. Through quantitative methods and Bayesian analysis, the authors show that social media can be a barometer of mood and thus play an important part in monitoring well-being.

This Special Section demonstrates the broad diversity and evolving nature of social media processing. Existing problems such as visual summary generation find new solutions through crowdsourcing and the availability of social networking applications, while new problems emerge including visual meme and YouTube popularity estimation through social media processing. In many papers, interesting multi-modal approaches are presented exploiting cross-domain relationships, which take into account a variety of features including content, context, and sentiment and employ, for example, transfer learning techniques. Therefore, specific issues arise including accessing large-scale real-world datasets from social networks, as well as designing and implementing efficient, cross-domain, and cross-modal processing approaches. Social data also impose the need for new studies since in many cases, performance depends on the specific nature of the social data stream and not only on the processing method applied. Regarding applications, social media open a new path towards addressing challenging problems, since in many cases, they offer access to real-life experiences, which are much more relevant than in controlled or laboratory settings. Socially enriched personalization and

recommendation approaches include information from social interactions, and mood estimation related applications can only be developed thanks to the availability and richness of social media and networking datasets. We see this as an important step in multimedia since social media availability allows the development of sophisticated multi-modal, cross-media approaches. Multimedia analysis should not be limited to audio-visual and textual features, and should also involve social interactions. At the same time, extracted metadata can be much richer to encompass mood related annotations and new aspects such as popularity. We hope the readers find the Special Section papers as exciting as we did.

Ioannis Kompatsiaris  
Information Technologies Institute, CERTH  
Thessaloniki, Greece  
ikom@iti.gr

Daniel Gatica-Perez  
Idiap Research Institute and EPFL  
Lausanne, Switzerland  
gatica@idiap.ch

Xing Xie  
Microsoft Research Asia  
Beijing, China  
Xing.Xie@microsoft.com

Jiebo Luo  
University of Rochester  
Rochester, NY USA  
jluo@cs.rochester.edu

**Ioannis (Yiannis) Kompatsiaris** (SM'11) is a Senior Researcher (Researcher B) with the CERTH-ITI in Greece. His research interests include semantic multimedia analysis, indexing and retrieval, social media analysis, knowledge structures, reasoning, and personalization for multimedia applications. He is the co-author of more than 200 papers in refereed journals and international conferences and of 7 patents. He has been the co-organizer of various international conferences and workshops and has served as a regular reviewer for a number of journals.

**Daniel Gatica-Perez** is the Head of the Social Computing Group at Idiap Research Institute and the Swiss Federal Institute of Technology (EPFL) in Lausanne, Switzerland. His research interests span social and mobile computing and computational social science. He has served as an Associate Editor of the IEEE TRANSACTIONS ON MULTIMEDIA, *Image and Vision Computing*, and the *Journal of Ambient Intelligence and Smart Environments*.

**Xing Xie** (SM'09) is a lead researcher in Microsoft Research Asia, Beijing, China. His research interests include spatial data mining, location based services, social networks, and ubiquitous computing. During the past years, he has published over 140 referred journal and conference papers. He has more than 50 patents filed or granted. He is a member of the Joint Steering Committee of the UbiComp and Pervasive Conference Series.

**Jiebo Luo** (F'09) joined the University of Rochester, Rochester, NY, USA, in Fall 2011 after fifteen years at Kodak Research Laboratories. His research spans computer vision, machine learning, social media data mining, and ubiquitous computing. He has published extensively in these fields with over 200 papers and 70 US patents.

Dr. Luo is a Fellow of the SPIE, and IAPR.