

Social media mining and search

Jinhui Tang · Meng Wang · Xian-Sheng Hua ·
Tat-Seng Chua

Published online: 24 May 2011
© Springer Science+Business Media, LLC 2011

Recent years have witnessed the proliferation of social media and the success of many social websites, including Flickr, Youtube, Facebook, Twitter, etc., which drastically increased the volume of community-shared media resources, including images and videos. These websites allow users not only to create and share media data but also to rate and annotate them. Thus lots of meta-data, such as user-provided tags, comments, geo-tags, capture time and EXIF information, associated to multimedia resources, are available in the social media websites. On the one hand, the rapid increase of social media data makes many related applications challenging, such as categorization, recommendation and search. On the other hand, the rich information clues associated with the data also offer us opportunities to attack many well-recognized difficulties encountered in multimedia analysis and understanding, e.g., insufficiency of labeled data for semantic learning.

The multimedia research community has widely recognized the importance of learning effective models for understanding, organization, and access but has failed to make rapid progress due to the insufficiency of labeled data, which typically comes from users in an interactive labor-intensive manual process. In order to reduce this manual effort, many semi-supervised learning or active learning approaches have been proposed. Nevertheless, there is still a need to manually annotate a large set of images or videos to bootstrap and steer the training. The rich information clues associated with the multimedia data in the social media websites offer a way out. If we can learn the models for semantic concepts effectively from user-shared data by using their associated meta-text as training labels, or if we can infer the semantic concepts of the multimedia data directly from the data in the

J. Tang (✉)
Nanjing University of Science and Technology, Nanjing, China
e-mail: jinhuitang@mail.njust.edu.cn

M. Wang · T.-S. Chua
National University of Singapore, Singapore, Singapore

X.-S. Hua
Microsoft Bing Multimedia Search, Redmond, WA, USA

Internet, manual efforts in multimedia annotation can be reduced. Consequently, semantic-based multimedia retrieval can benefit much from the community-contributed resource. There is, however, a problem in using the associated meta-information as training labels: they are often very noisy. Thus how to remove the noise in the training labels or how to handle the noise in the learning process are urgent research topics.

At the semantic level, the key information we need for social imedia data is the so-called 5W's and 1H, i.e., who, where, when, what, why and how. As why and how tend to be abstract, difficult and may not be relelant to most social media applications, automatic tools for analyzing and mining the social media should help to understand the first 4W's, namely, who, where, when and what. The aforementioned challenges are mostly related to the "what" problem. For the "who" issue, the key techniques are face detection and recognition, which have been extensively studied in computer vision research. Besides the "what" problem, "where" and "when" issues are also hot topics in multimedia research community, that is, to mining the EXIF and geotag information for multimedia applications.

Besides modeling media items (e.g. an image or a video), the social media sites are providing with incredible resources to model users, through the aggregation of their traces on social media sites (e.g. the images they upload, the tags they use, the people whose content they comment on). So in addition to model multimedia data only, how to model people's behaviors or events is also important.

Pivotal to many tasks in relation to social media mining and research is the availability of sufficiently large dataset and its corresponding ground truth. Currently available datasets for multimedia research are either too small, too specific, or without ground truth. While it is relatively easy to crawl and store a huge amount of data, the creation of ground-truth necessary to systematically train, test, and evaluate the performance of various algorithms and systems is a major problem. For this reason, more and more research groups are individually putting efforts into the creation of such corpus in order to carry out research on social media dataset, such as the MIRFlickr and NUS-WIDE.

Recently, more and more research efforts have been dedicated to the aforementioned challenges and opportunities. Therefore, we edit this special issue named Social Media Mining and Search. The goals of this special focus on threefold: (1) introduce novel research in learning from resources in the Internet; (2) survey on the progress of this area in the past years; (3) discuss new applications based on the newly learned models.

We had 18 submissions came from an open call for paper. With the assistance of dedicated referees, nine papers have been selected out of them after two rounds of rigorous reviews. These papers cover widely subtopics of social media mining and search.

1 Brief introduction to the papers in this issue

In the first paper "social multimedia: highlighting opportunities for search and mining of multimedia data in social media applications", Namman gives a survey on the social media applications, and describes two specific applications that are built on mining of multimedia data by the Yahoo! Research Berkeley team. These applications could help illustrate some of the new opportunities embodied in social multimedia. The author generalizes the approach for the two applications, to suggest a general approach for social multimedia analysis and applications.

The second paper “leveraging community metadata for multimodal image ranking” by Richter et al. focuses on the goal of selecting relevant images given a query term, i.e. finding images showing content that most people associate with the query term. More specifically they aim to solve this image search problem on a large-scale community database such as Flickr, where images are often associated with different types of user generated metadata, e.g. tags, date & time, and location.

In the third paper “social media filtering based on collaborative tagging in semantic space”, Kim et al. proposes a semantic collaborative filtering method to enhance recommendation quality derived from user-generated tags. In addition, we explore several advantages of semantic tagging for ambiguity, synonymy, and semantic interoperability, which are notable challenges in information filtering. The proposed approach first determines semantically similar users using social tagging and subsequently discovers semantically relevant items for each user.

The fourth paper “social image annotation via cross-domain subspace learning” by Si et al. envisions the cross-domain discriminative subspace learning and provides an effective solution to cross-domain subspace learning. In particular, they propose the cross-domain discriminative locally linear embedding (CDLLE), which connects the training and the testing samples by minimizing the quadratic distance between the distribution of the training samples and that of the testing samples. They also apply this subspace learning method for social image annotation.

In the next paper “using manual and automated annotations to search images by semantic similarity”, Magalhães and Rürger utilize the manual and automatic annotations to search images by semantic similarity. More specifically, they study the accuracy of manual annotations versus automatic annotations, the influence of manual annotations with different accuracies as a result of incorrect annotations, and revisit the influence of the keyword space dimensionality, to examine different aspects of search images by semantic similarity.

The sixth paper “inferring photographic location using geotagged web images” and the seven paper “geotag propagation in social networks based on user trust model” are both on the geo-tag information mining. In the sixth paper, Joshi et al. leverage user tags along with image content to infer the geolocation of images. Their model builds upon the fact that the visual content and user tags of pictures can together provide significant hints about their geolocations. Using a collection of over a million geotagged pictures, they build location probability maps for commonly used image tags over the entire globe. These maps reflect the collective picture-taking and tagging behaviors of thousands of users from all over the world. In the seventh paper, Ivanov et al. present a system for efficient geotag propagation based on a combination of object duplicate detection and user trust modeling. The geotags are propagated by training a graph based object model for each of the landmarks on a small tagged image set and finding its duplicates within a large untagged image set. Based on the established correspondences between these two image sets and the reliability of the user, tags are propagated from the tagged to the untagged images. The user trust modeling reduces the risk of propagating wrong tags caused by spamming or faulty annotation.

The eighth paper “mining diversity on social media networks” by Liu et al. gives a comprehensive study on the concept of *diversity*, which characterizes how diverse a given node connects with its peers, in the social network. They lay out two criteria that capture the semantic meaning of diversity, and then propose a compliant definition which is simple enough to embed the idea. Based on the approach, not only a user’s sociality and interest diversity but also a social media’s user diversity can be measured. An efficient top-k diversity ranking algorithm is also developed for computation on dynamic networks.

In the last paper “inferring competitive role patterns in reality TV show through nonverbal analysis”, a new facet of social media, namely that depicting social interaction, is introduced by Raducanu and Gatica-Perez. They address this problem from the perspective of nonverbal behavior-based analysis of competitive meetings. Their analysis is centered on two tasks regarding a person’s role in a meeting: predicting the person with the highest status, and predicting the fired candidates.

2 Further challenges and directions

Social media, such as photos, videos, tags, comments and user relations have proved a valuable resource for multimedia information mining, search and further developing valuable applications. We summarize the research topics in social media mining and search, as well as the related applications, in Fig. 1. Current a lot of the effort was devoted to research on annotation, tagging, classification, ranking, retrieval, tag recommendation, geo-tag based applications and social media information network, while the papers in this special issue are all fell into these topics. In this promising research area, there are several open issues worth further study:

- When utilizing the social media as training data to mine new knowledge or learn model, how to remove the noise in the tags and comments or how to handle the noise in the learning process is an urgent issue to tackle;
- Social media sites provide incredible resources to model users, through the aggregation of their traces on social media sites. Thus besides modeling the multimedia data itself, how to model people’s behaviors or events is valuable for social media mining and understanding;
- Different tags have different representativeness to certain images, while most of the existing works deal them equally. Thus how to measure the tag representativeness or relevant to the given image and how to rank tags for each image is also important;
- The scalability of the techniques for social media mining and search in billion-scale collections, especially when the content of interest (object, landmark, etc.) only occurs in a small portion of the images or videos; and
- How to unify the individual efforts to create a unified web-scale repository for experimental evaluation of newly proposed techniques.

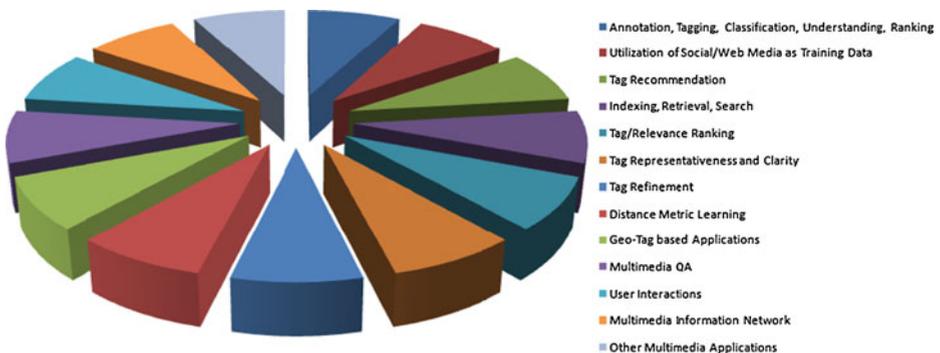


Fig. 1 The research topics in social media mining and search, as well as the related applications



Jinhui Tang is a Professor of School of Computer Science and Technology, Nanjing University of Science and Technology. He received his B.E. and Ph.D. degrees in July 2003 and July 2008 respectively, both from the University of Science and Technology of China (USTC). From July 2008 to Dec. 2010, he worked as a research fellow in School of Computing, National University of Singapore. During that period, he visited School of Information and Computer Science, UC Irvine, as a visiting research scientist, from Jan. 2010 to Apr. 2010. His current research interests include large-scale multimedia search, social media analysis and computer vision. He has over 60 journal and conference papers in these areas.

Dr. Tang served as a guest editor of *IEEE Transactions on Multimedia*, *ACM Transactions on Intelligent Systems and Technology*, *ACM/Springer Multimedia Systems Journal*, and *Springer Journal of Multimedia Tools and Applications*, a co-chair of the ACM the First International Workshop on Web-Scale Multimedia Corpus 2009, a Technical Committee Member for over 30 international conferences, such as ACM Multimedia and SIGIR, and a reviewer for over 30 prestigious international journals. Dr. Tang is a recipient of the 2008 President Scholarship of Chinese Academy of Science, and a co-recipient of the Best Paper Award in ACM Multimedia 2007. He is a member of ACM, IEEE, and IEEE Computer Society.



Meng Wang is a research staff member in National University of Singapore. He received the B.E. degree and Ph.D. degree in the Special Class for the Gifted Young and signal and information processing from the University of Science and Technology of China (USTC), Hefei, China, respectively. He previously worked as an associate researcher at Microsoft Research Asia, and then a core member in a startup in Bay area. His current research interests include multimedia content analysis, search, mining, recommendation, and large-scale computing. He has authored 6 book chapters and over 70 journal and conference papers in these areas, including TMM, TCSVT, TOMCCAP, ACM MM, WWW, SIGIR, ICDM, etc.

Dr. Wang serves as an associate editor of Information Sciences, an editorial board member of International Journal of Multimedia Intelligence and Security, a guest editor of ACM/Springer Multimedia Systems journal for the special issue on “Interactive Multimedia Computing”, Elsevier Journal of Visual Communication and Image Representation for the special issue on “Large-Scale Image and Video Search: Challenges, Technologies, and Trends”, and Springer Multimedia Tools and Applications for the special issue on “Social Media Mining and Search”, a Technical Committee Member and reviewer for various international conferences and journals. He received the best paper awards continuously in the 17th and 18th ACM International Conference on Multimedia and the best paper award in the 16th International Multimedia Modeling Conference. He is a member of IEEE and ACM.



Xian-Sheng Hua is now a Principle Research and Development Lead for Bing Multimedia Search with Microsoft. He is responsible for driving a team to design and deliver thought-leading media understanding and indexing features. Before joining Bing in 2011, Dr. Hua was a Lead Researcher with Microsoft Research Asia. During this time, his research interests are in the areas of multimedia search, advertising, understanding, and mining, as well as pattern recognition and machine learning. He has authored or co-authored more than 180 publications in these areas and has more than 60 filed patents or pending applications.

Dr. Hua received the B.S. and Ph.D. degrees from Peking University, Beijing, China, in 1996 and 2001, respectively, both in applied mathematics. He serves as an Associate Editor of IEEE Transactions on Multimedia, Associate Editor of ACM Transactions on Intelligent Systems and Technology, Editorial Board Member of Advances in Multimedia and Multimedia Tools and Applications. He also has successfully served as vice program chair of VCIP 2005, senior TPC members or area chairs of ACM Multimedia, ACM KDD, and SIGIR. He will be a program co-chair of ACM Multimedia 2012.

Dr. Hua won the Best Paper Award and Best Demonstration Award in ACM Multimedia 2007, Best Poster Award in 2008 IEEE International Workshop on Multimedia Signal Processing, Best Student Paper Award in ACM Conference on Information and Knowledge Management 2009, and Best Paper Award in International Conference on MultiMedia Modeling 2010. He also won 2008 MIT Technology Review TR35 Young Innovator Award for his outstanding contributions to video search, and named as one of the “Business Elites of People under 40 to Watch” by Global Entrepreneur. He is a member of IEEE and a senior member of ACM.



Tat-Seng Chua is the KITHC Chair Professor at the School of Computing, National University of Singapore (NUS). He was the Acting and Founding Dean of the School of Computing during 1998–2000. He joined NUS in 1983, and spent three years as a research staff member at the Institute of Systems Science (now I2R) in the late 1980s. Dr Chua's main research interest is in multimedia information retrieval, in particular, on the analysis, retrieval and question-answering (QA) of text and image/video information. He is currently working on several multi-million-dollar projects: interactive media search, local contextual search, and real-time live media search. His group participates regularly in TREC-QA and TRECVID video retrieval evaluations.

Dr Chua has organized and served as program committee member of numerous international conferences in the areas of computer graphics, multimedia and text processing. He is the conference co-chair of ACM Multimedia 2005, CIVR 2005, and ACM SIGIR 2008. He will serve as Technical PC Co-Chair of SIGIR 2011. He serves in the editorial boards of: ACM Transactions of Information Systems (ACM), Foundation and Trends in Information Retrieval (NOW), The Visual Computer (Springer), and Multimedia Tools and Applications (Springer). He is the member of steering committee of ACM ICMR, Computer Graphics International, and Multimedia Modeling conference series; and as member of International Review Panels of two large-scale research projects in Europe. He received the Best Paper Award in ACM Multimedia 2010.