KDD/MDM 2006: The 7th KDD Multimedia Data Mining Workshop Report

Zhongfei Zhang
Computer Science Dept.
SUNY Binghamton
Binghamton, NY, USA
zhongfei@cs.binghamton.edu

Florent Masseglia
INRIA Sophia Antipolis
France
Florent.Masseglia@sophia.inria.fr

Ramesh Jain
Computer Science Dept.
University of California
Irvine, CA, US
jain@ics.uci.edu

Alberto Del Bimbo
Systems and Informatics
University of Florence
Florence, Italy
delbimbo@dsi.unifi.it

ABSTRACT

This is a report of the 7th ACM KDD/MDM Workshop on Multimedia Data Mining held in Philadelphia, PA, USA, on 20 August 2006. This report outlines the motivations to hold the workshop, the response we have received in terms of paper submissions and participations, as well as the post-workshop events. The central theme of this year's MDM workshop is to merge multimedia and data mining research together to exploit the synergy between the two areas to advance and promote the research and development in multimedia data mining. The workshop was extremely successful. The feedback we have received during and after the workshop was overwhelmingly positive, which indicates the emerging interest and enthusiasm on the topic of multimedia data mining in the related research communities. Due to this fact, we plan to continue this series next year.

Keywords

Multimedia Data Mining, Multimedia, Data Mining.

1. MOTIVATIONS

The 7th Workshop on Multimedia Data Mining was held in Philadelphia, PA, USA, on 20 August 2006 in conjunction with ACM KDD 2006. The previous six workshops were held in conjunction with KDD 2000 (Boston, MA), KDD2001 (San Francisco, CA), KDD 2002 (Edmonton, Canada), KDD 2003 (Washington, DC), KDD 2004 (Seattle, WA), and KDD 2005 (Chicago, IL), respectively. These workshops brought together numerous experts in the following fields: spatial data analysis, digital media, multimedia information retrieval, data mining and knowledge discovery in multimedia database systems, and analysis of data in collaborative virtual environments. Participants were pleased with the workshops and there was consensus about the necessity to continue this annual meeting, where researchers, both from the academia and industry, could exchange and compare both relatively mature and green house theories, methodologies, algorithms and frameworks for multimedia data mining.

This year's workshop focused on the theme of merging multimedia and data mining research together to exploit the synergy between the two areas to promote and advance the research on multimedia data mining. Multimedia and data mining are two very interdisciplinary and multidisciplinary areas with independently and simultaneously rapid developments in recent years. Many successful examples have been seen to address a series of pressing issues, in exploring and exploiting the synergy between both areas. This phenomenon is particularly true in many applications that call for efforts that combine the research from both areas under the theme of Multimedia Data Mining. This workshop aims at dedicating to the research on this theme with the purpose to advance and disseminate the most recent research under this theme.

It is well known that multimedia information is ubiquitous and is often required, if not essential, in many applications. This phenomenon has made multimedia repositories widespread and extremely large. There are tools for managing and searching within these collections, but the need for tools to extract hidden useful knowledge embedded within multimedia collections is becoming pressing and central for many decision-making applications. For example, it is highly desirable for developing the tools needed today for discovering relationships between objects or segments within images, classifying images based on their content, extracting patterns in sound, categorizing speech and music, and recognizing and tracking objects in video streams.

Data mining research today has been advanced far beyond the area of databases. The efforts in this area have been focusing on data collected in fields such as art, design, hypermedia and digital media production, case-based reasoning and computational modeling of creativity, including evolutionary computation, and medical multimedia data. These exotic fields use a variety of data sources and structures, interrelated by the nature of the phenomenon that these structures describe. As a result there is an increasing interest in new techniques and tools that can detect and discover patterns that lead to new knowledge in the problem domain where the data have been collected. There is also an increasing interest in the analysis of multimedia data generated by different distributed applications, such as collaborative virtual environments, virtual communities, and multi-agent systems. The data collected from such environments include a record of the actions in them, a variety of documents that are part of the business process, asynchronous threaded discussions, transcripts from synchronous communications and other data records. These

heterogeneous multimedia data records require sophisticated preprocessing, synchronization and other transformation procedures before even moving to the analysis stage.

On the other hand, researchers in multimedia information systems, in the search of techniques for improving the indexing and retrieval of multimedia information, are looking for new methods for discovering indexing information. Variety of techniques from machine learning, statistics, databases, knowledge acquisition, data visualization, image analysis, high performance computing, and knowledge-based systems, have been used mainly as a research handcraft activity. The development of multimedia databases and their query interfaces recall again the idea of incorporating multimedia data mining methods for dynamic indexing. The emerging international standard for multimedia content description (MPEG-7) promises to foster the collaboration in the field giving a uniform data representation.

The next step for successful KDD applications will certainly involve both fields (Multimedia and Data Mining). Actually, it is well known that real world applications often have data with multiple modalities, from multiple sources and in multiple formats. For example, in homeland security applications, we may need to mine data from an air traveler's credit history, traveling patterns, photo pictures, and videotapes from surveillance cameras in the airport. In the manufacturing domains, business processes can be improved if, for example, part drawings, part descriptions and part flow can be mined in an integrated way instead of separately. In medicine, a disease might be predicted more accurately if the MRI (magnetic resonance imaging) imagery is mined together with other information about the patient's condition. Similarly, in bioinformatics data are available in multiple formats. While mining on structured data or each type of modality of multimedia data such as text data, imagery data, and video data has been broadly addressed, there has not been much effort focusing on integrated approaches to mining information from multiple modalities, multiple sources and multiple formats.

Consequently, it is important to have this workshop to dedicate to this emerging theme as a venue to advance and disseminate the most recent research on this theme.

2. PAPER SOLICITATION AND REVIEW

At the beginning, we were very aggressive to advertise this year's workshop in soliciting paper submissions. In order to further promote the paper submissions to this workshop, we secured a special issue [1] of IEEE Transactions on Multimedia on the same topic of Multimedia Data Mining open for the best papers from the workshop. As a result, we received 22 submissions in total for this workshop relevant to the central theme of the workshop.

In order to secure a high quality review, we formed an international program committee consisting of 25 internationally renowned experts with strong expertise in the related area. Since this was a workshop instead of a well-established conference, during the review process we particularly paid attention to encouraging papers with novel ideas (even without extensive evaluation data). Each paper received as least two anonymous reviews. As a result, we accepted 9 papers as long papers, 7 papers as short papers, and 2 papers as posters. Two papers withdrew for various reasons and two other papers were declined based on negative reviews.

3. PARTICIPATION AND PRESENTATION

The participation into the workshop this year was overwhelming. The workshop room was always almost full during the whole program, and was completely full with many attendees standing in the back during the invited speech. We received in total 68 registrations for this workshop. This is a record-breaking number over the whole workshop series which indicates the emerging interest and enthusiasm on the topic of multimedia data mining among the related research communities.

The program consists of three oral presentation sessions, one poster presentation session, and an invited speech session.

The first oral session was on video mining. This session had three long presentations: *Mining motion patterns using color motion map clustering*, *Mining appearance models directly from compressed video*, and *A novel video annotation method by integrating visual features and frequent patterns*, and two short presentations: *Counting people using video cameras* and *Ontologies for video digital libraries*.

The second oral session was on mining methods. This session had two long presentations: *Adaptive offset subspace self organizing map with an application to handwritten digit recognition* and *A spatial semi-supervised learning method for classification of multi-spectral remote sensing imagery* and three short presentations: *Discovery of latent patterns with hierarchical Bayesian mixed membership models and the issue of model choice*, *Clustering by random projections: application to image segmentation*, and *Visual data mining: integrating machine learning with information visualization*.

The third oral session was on image mining. The session had four long presentations: *Improving image annotations using fuzzy pruning and association rule mining*, *A probabilistic model with multi-modal reinforcement for outlier detection in image database*, *Image classification by a probabilistic model learned from imperfect training data on the Web*, and *Document image retrieval using symbolized bounding box* and two short presentations: *Integrating hyperbolic image browsing towards an integrated multimedia navigator* and *An hybrid approach for semantic information retrieval using multimedia data*.

Due to the tight program schedule, the poster session was held during the lunch hour and consisted of two presentations: *Using closed captions and visual features to classify movies by genre* and *Application of the MACLAW algorithm for bands selection in hyperspectral images*.

4. INVITED SPEECH

Major Amy Magnus, PhD, from US Air Force Office of Scientific Research kindly accepted our invitation to speak in the workshop. She gave the invited presentation at the workshop entitled *Mining for landmarks: introducing new local environments to a global network*. The presentation went extremely well with the whole workshop room full with the audience. This became the highlight of the workshop. We gratefully thank Amy for her time and effort to delivering this wonderful speech.

5. POST-WORKSHOP EVENTS

Per the request from *ACM KDD Explorations* calling for the best paper from each workshop to appear in a special issue of *ACM KDD Explorations*, we went back to look at the reviews received for all the papers. As a result, the paper entitled *Mining motion patterns using color motion map clustering* by Chih Lai, Taras Rafa, and Dwight E. Nelson of the University of St. Thomas, USA, stood out. Congratulations to Chih, Taras, and Dwight!

Due to the successful workshop and the high quality of the papers published in the workshop through the workshop proceedings as well as in the ACM Digital Libraries, two special issues of journals are being organized. In addition to the special issue on multimedia data mining of IEEE Transactions on Multimedia [1], Chabane Djeraba of LIFL, France, is organizing another special issue from the papers in the workshop. Authors of the workshop papers are reminded that they have full choice of contributing to either of the special issue; however, they should not submit their papers to both special issues.

6. CONCLUSIONS

We conclude this report saying that the workshop was a great success and the feedbacks and responses we have received from the workshop attendees were overwhelmingly positive. This fact encourages us to continue this workshop series and we plan to hold the next edition of the workshop next year in conjunction with ACM KDD 2007.

7. WORKSHOP STEERING COMMITTEE

We thank the advice provided by the workshop steering committee:

Fatma Bouali, University of Lille 2, France

Chabane Djeraba, LIFL, France

Latifur Khan, University of Texas at Dallas, USA

Valery A. Petrushin, Accenture Technology Labs, USA

Simeon J. Simoff, University of Technology, Sydney, Australia

8. WORKSHOP PROGRAM COMMITTEE

We thank the effort and time contributed by the workshop program committee members in reviewing all the papers:

K. Slecuk Candan, Arizona State University, USA

Datong Chen, Carnegie Mellon University, USA

Shu-Ching Chen, Florida International University, USA

Alex Duffy, University of Strathclyde, UK

Max J. Egenhofer, University of Maine, USA

William Grosky, University of Michigan,USA

Wynne Hsu , National University of Singapore, Singapore

Manfred Jeusfeld, Tilburg University, Netherlands

Odej Kao, TU Berlin, Germany

Michael Lew, Leiden University, The Netherlands

Bo Long, State University of New York at Binghamton, USA

Dennis McLeod, University of Southern California, USA

Dunja Mladenic, J. Stefan Institute, Slovenia

Andreas Nürnberger, University of Magdeburg, Germany

Vincent Oria, New Jersey Institute of technology, USA

Jia-Yu (Tim) Pan, Carnegie Mellon University, USA

Jian Pei, Simon Fraser University, Canada

Nicu Sebe, University of Amsterdam, The Netherlands

Mei-Ling Shyu, University of Miami, USA

Zhaohui Tang, Microsoft, USA

Guangyou Xu, Tsinghua University, China

Jie Yang, Carnegie Mellon University, USA

Aidong Zhang, State University of New York at Buffalo, USA

Ruofei Zhang, Yahoo!, USA

Zhi-Hua Zhou, Nanjing University, China

9. ACKNOWLEDGMENTS

We finally thank all the authors of this workshop for their contributions and thank all the participants to this workshop for their interest and enthusiasm.

10. REFERENCES

[1]http://www.ieee.org/organizations/society/sp/tmmspecialissues.html

About the authors:

Zhongfei (Mark) Zhang is currently an associate professor of Computer Science at the Computer Science Department at SUNY Binghamton. He received a B.S. in Electronics Engineering (with Honors), an M.S. in Information Sciences, both from Zhejiang University, China, and a PhD in Computer Science from the University of Massachusetts at Amherst. He was on the faculty of Computer Science and Engineering Department, and a research scientist at the Center of Excellence for Document Analysis and Recognition, both at SUNY Buffalo. His research interests include Multimedia Information Indexing and Retrieval, Data Mining and Knowledge Discovery, Computer Vision and Image Understanding, Pattern Recognition, as well as Bioinformatics. His research is sponsored by NSF, AFOSR, AFRL, and NYS, as well as private industries including Microsoft and Kodak. He has served as reviewers/PC members for many conferences and journals, as grant review panelists for governmental and private funding agencies. He has also served as technical consultants for a number of industrial and governmental organizations. He is a recipient of US National Academies/National Research Council Visiting Fellow, Air Force Research Laboratory Faculty Visiting Fellow, Microsoft Research Visiting Researcher, Western New York 2004 Inventor of the Year Individual Category 2nd Place, SUNY Chancellor's Promising Inventor Award, and JSPS International Collaboration Award.

Florent Masseglia did research work in the Data Mining Group at the LIRMM (Montpellier, France) from 1998 to 2002 and received a Ph.D. in computer science from Versailles University, France in 2002. He has been a temporary researcher for the LIRMM from 2002 to 2003. In 2003 he was recruited as a permanent researcher for the INRIA (Sophia Antipolis, France), which is his current position. His research interests include data mining (particularly sequential patterns and mining complex data) and databases.

Ramesh Jain is an educator, researcher, and entrepreneur. Currently he is the Donald Bren Professor in Information & Computer Sciences at University of California, Irvine. Before this he was a Farmer Distinguished Chair at Georgia Institute of Technology. Ramesh is a pioneer in multimedia information systems, image databases, machine vision, and intelligent systems. While professor of computer science and engineering at the University of Michigan, Ann Arbor and the University of California, San Diego, he founded and directed artificial

intelligence and visual information systems labs. Ramesh was also the founding Editor-in-Chief of IEEE Multimedia magazine and serves on the editorial boards of several journals in multimedia, information retrieval and image and vision processing. He has co-authored more than 300 research papers in well-respected journals and conference proceedings. He has co-authored and co-edited several books. He is a Fellow of ACM, IEEE, AAAI, IAPR, and SPIE. He is the Chairman of ACM SIG Multimedia. He recently co-founded Seraja to address needs of emerging EventWeb. Earlier Ramesh co-founded three companies, managed them in initial stages, and then turned them over to professional management. These companies were PRAJA in event-based business activity monitoring (acquired by Tibco); Virage for media management solutions and visual information management (a NASDAQ company acquired by Autonomy); and ImageWare for surface modeling, reverse engineering rapid prototyping, and inspection (acquired by SDRC).

Alberto Del Bimbo is Full Professor of Computer Engineering and the Director of the Master in Multimedia of the University of Florence, Italy. At the present time, he is also Deputy Rector of the University of Florence, in charge of Research and Innovation Transfer. His scientific interests are Pattern Recognition, Image Databases, Human Computer Interaction and Multimedia applications. He has published over 200 publications in some of the most distinguished scientific journals and international conferences, and is the author of the "Visual Information Retrieval" monography on content-based retrieval from image and video databases. Prof. Del Bimbo is Member of IEEE and ACM Fellow of IAPR (International Association for Pattern Recognition). He was the President of the IAPR Italian Chapter, from 1996 to 2000 and Member at Large of the IEEE Publication Board from 1998 to 2000. He was the general Chair of IAPR ICIAP‘97 (International Conference on Image Analysis and Processing) IEEE ICMCS’99 (International Conference on Multimedia Computing and Systems) and in the board of many other primary scientific conferences. He is Associate Editor of Pattern Recognition, Journal of Visual Languages and Computing, Multimedia Tools and Applications, Pattern Analysis and Applications, International Journal of Image and Video Processing and was Associate Editor of IEEE Transactions on Multimedia, and IEEE Transactions on Pattern Analysis and Machine Intelligence. He was the Guest Editor of several special issues on Image databases and image analysis.