

GeoMM 2014: The Third ACM Multimedia Workshop on Geotagging and Its Applications in Multimedia

Liangliang Cao
IBM Research
Yorktown Heights, NY, USA
liangliang.cao@us.ibm.com

Gerald Friedland
ICSI
Berkeley, CA, USA
fractor@icsi.berkeley.edu

Lexing Xie
Australian National University
Canberra, Australia
lexing.xie@anu.edu.au

ABSTRACT

The Third ACM Workshop on Geotagging and Its Applications in Multimedia (GeoMM'14), continues the workshops in 2012 and 2013, with the goal of building a forum for the presentation and synthesis of vision and insight from leading experts and practitioners on the developing directions of geotagging research related to multimedia. Following the success in previous years, the GeoMM workshop serves as a venue for the premier research in geotagging and multimedia, and continues to attract submissions from a diverse set of researchers, who address newly arising problems within this emerging field. Five regular papers are presented in this workshop, covering a number of novel applications and new methodologies. An invited paper is also presented to introduce the related MediaEval 2014 Placing task, which consists of 5 million geotagged photos and 25,000 geotagged videos. We believe this workshop will benefit more and more research works in the broad research field.

Categories and Subject Descriptors

I.5.4 [Applications]: Computer Vision

General Terms

Overview, Experiments

Keywords

Geo-tagged Multimedia; Location based Services; Overview

1. INTRODUCTION

In recent years, the rise of wide-spread use of GPS sensors and the increasing availability of open geographical databases has motivated a large volume of work on geotagging. The increase in the use of geotagging and improvements in geolocation support systems open up a new dimension for the description, organization and manipulation of multimedia data. This new dimension radically expands the usefulness

of multimedia data both for daily users of the Internet and social networking sites as well as for experts in particular application scenarios.

The goal of this workshop is to bring together cutting-edge research in geotagging as well as novel applications related to geotagging. It will help facilitate in-depth discussions, share existing tools, and ultimately enhance the research efforts in this area. The workshop will serve as a forum for the presentation and synthesis of vision and insight from leading experts and practitioners on the developing directions of geotagging research related to multimedia.

Topics of interests include, but not limited to

- Collection and estimation of geotags
- Location estimation: inferring geolocation from visual and audio cues
- Mining unstructured web text (e.g., news, blogs) for location
- Knowledge discovery from open geographical databases (e.g., Google earth, GeoNames) and web photos
- Mining geotags from mobile devices and sensor network
- Applications of geotagging
- Photo and video annotation
- Business applications (e.g., geomarketing, targeted advertising etc.)
- Social networking
- Privacy and security issues
- Visualization of geotagged information
- 3D construction of world model
- Visualization of tourism, travel routes, and other spatial-temporal patterns

The call for papers attracted a number of submissions from both industry and academia institutions. We finally select five papers, which cover rich aspects of geommm, including profiling of business locations, capture-time classification, recommending photo-taking locations, urban computing, and inferring home locations. We are excited that a number of novel applications have been proposed and hope these work will make a big impact in the future.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

MM'14, November 3–7, 2014, Orlando, Florida, USA.

ACM 978-1-4503-3063-3/14/11.

<http://dx.doi.org/10.1145/2647868.2647873>

In addition, we also invite a paper from the organizers of MediaEval 2014 Placing Task. The Placing task 2014 is a multimodal location estimation task with record-breaking 5 million geotagged photos and 25,000 geotagged videos. We would like to embrace the opportunity and challenge with this new event, and hope the datasets would be useful for many researchers in this field.

Finally, we are proud to host a number of keynote speakers who have done influential work in this field. Their talks will provide cutting edge overviews of the state of the art in this field.

We would like to send our warmest greetings and invite everyone who is interested in geommm to attend our workshop. After organizing this series of workshop for three years, we have witness the surge of this emerging area and we believe we will see more success in the future.

2. WORKSHOP ORGANIZERS

Liangliang Cao is a Research Staff Member in IBM T. J. Watson Research Center, and also an adjunct assistant professor at Columbia University. His research lies in the intersection of computer vision, multimedia and big data analytics. His work has won three prestigious visual recognition competitions, include ImageCLEF Medical Image Classification (2012, 2013), and ImageNet Large Scale Visual Recognition Challenge (2010). His contribution has been awarded the IBM Outstanding Accomplishment (2012), the Best Paper Award in the First International Workshop on Big Data Mining (2012), IBM Watson Emerging Leader in Multimedia and Signal Processing (2010), Facebook Fellowship Finalist (2010), and UIUC Computational Science and Engineering Fellowship (2009-2010). He has authored more than 40 papers in top conferences and journals, including ICCV, CVPR, ECCV, NIPS, ACM Multimedia, WWW, TPAMI, and PIIEEE. He is a guest editor of ACM Transactions on Multimedia Computing, Communications (TOMCCAP) and Applications and also Computer Vision and Image Understanding (CVIU) Journal. He is a general chair of New York Area Multimedia and Vision Meeting in Greater in 2012 and 2013. Dr. Cao is an area chair of ACM Multimedia 2012 and IEEE WACV 2014. He fulfills review duties for more than 15 journals and various conferences.

Gerald Friedland is the Director of Audio and Multimedia Research at the International Computer Science Institute, a private non-profit lab affiliated with the University of California, Berkeley. His leads a group of currently 18 researchers that focus on acoustic techniques for video retrieval, multimodal location estimation, and the privacy implications of multimedia research. Dr. Friedland has published about 200 peer-reviewed articles in conferences, journals, and books and has authored a textbook on Multimedia Computing together with Dr. Ramesh Jain. More importantly, he has participated in MediaEval Placing since the beginning, has authored the inaugural Brave New Topic paper on Multimodal Location Estimation and has co-authored the book "Multimodal Location Estimation of Images and Videos" together with Jaeyoung Choi. He is the recipient of several research and industry recognitions, among them the European Academic Software Award and the Multimedia Entrepreneur Award by the German Federal Department of Economics. Most recently, he lead the team that won the ACM Multimedia Grand Challenge in 2009. Gerald Friedland is an associate editor for the ACM Transactions on

Multimedia Computing, Communications, and Applications (TOMCCAP), IEEE Multimedia Magazine and regularly reviews for IEEE Transactions on Acoustics, Speech, and Language Processing, IEEE Transaction on Multimedia, IEEE Multimedia, Springer's Machine Vision and Application, and other journals. He has been involved in the organization of ACM Multimedia since 2009 and was program co-chair at IEEE ICME 2012. He is a senior member of IEEE and ACM and is also chairing the Special Interest Group on Audio and Speech Processing for Multimedia (SIGASP) which is part of the IEEE TC on Multimedia.

Lexing Xie is Senior Lecturer in the Research School of Computer Science at the Australian National University. She was research staff member at IBM T.J. Watson Research Center in New York from 2005 to 2010, and adjunct assistant professor at Columbia University 2007-2009. She received B.S. from Tsinghua University, Beijing, China, and M.S. and Ph.D. degrees from Columbia University, all in Electrical Engineering. Her research interests are in applied machine learning, multimedia, social media. Her recent projects are on multimedia content analysis, social media tracking and recommendation, visual semantics, large-scale image and video search. Lexing's research has received five best student paper and best paper awards between 2002 and 2011, and a Grand Challenge Multimodal Prize at ACM Multimedia 2012. She was the 2005 IBM Research Josef Raviv Memorial Postdoc fellow in Computer Science and Engineering. She currently serves an associate editor of IEEE Transactions on Multimedia, and ACM Transactions on Multimedia. Her service roles include the program and organizing committees of major multimedia, machine learning, web and social media conferences.

Acknowledgment

We would like to thank all authors who submitted contributions and the program committee members for their excellent work to review the papers. We are also very grateful for the invited keynote speakers for their presentations in the workshop.