

Home
Important Dates
Info for Authors
People
Conference Program
Arts Program
Workshops
Tutorials
Registration
Hotel, Travel, Local
Sponsors



SPECIAL INTEREST GROUP  
on MULTIMEDIA

# ACM Multimedia 2006

October 23-27, 2006  
Fess Parker's DoubleTree Resort Hotel  
Santa Barbara, CA

[\\* Flickr photos of ACM Multimedia 2006 \\*](#)

## ACM MM 2006 Awards

**Best full technical paper:** Tiling Slideshow

Authors: Jun-Cheng Chen, Wei-Ta Chu, Jin-Hau Kuo, Chung-Yi Wenb, Ja-Ling Wu  
National Taiwan University

**Best short technical paper:** Fourth Frame Forums: Interactive Comics for Collaborative Learning

Authors: Andrew Gordon  
University of Southern California

**Best Art Contribution:** Archeology of Multimedia

Authors: Fabrizio Nunnari, Vincenzo Lombardo, Andrea Valle, Francesco Giordana, Andrea Arghinetti  
University of Torino

**Best Demonstration:** A Real-Time Multi-modal Biofeedback System for Stroke Patient Rehabilitation

Authors: Yin-peng Chen, Wei-wei Xu, Richard Isaac Wallis, Hari Sundaram, Thanassis Rikakis, Todd Ingalls,  
Loren Olson, Ji-ping He  
Arizona State University

**Open Source Competition Winner:** CLAM: C++ Library for Audio and Music

Authors: Xavier Amatriain, Pan Arumi, David Garcia  
University of California, Santa Barbara and Pompeu Fabra University, Barcelona

**Best Video Program:** Globe4D, Time-traveling with an Interactive Four Dimensional Globe

Authors: Rick Companje, Nico van Dijk, Hanco Hogenbirk, Danica Mast  
Globe4D.com

## News and Updates

- ACM Multimedia 2006 had over 400 attendees and was a smashing success. We hope to see you all in Augsburg, Germany (September 23-29, 2007) for [ACM Multimedia 2007!](#)
- There were a couple problems with the conference DVD - their DVD duplicator was defective. If any registered conference or workshop participants would like a replacement DVD, please email Bob at [typedept3@sheridanprinting.com](mailto:typedept3@sheridanprinting.com) with your name and address, and we will mail one to you. Here are the online fixes:
  - The VSSN'06 Workshop proceedings may not be readable. They are now available online [here](#).
  - One of the video program files did not copy properly. It is available [here](#).
- The report from the SIGMM 2003 Retreat can be found [here](#).

## General Information

ACM Multimedia 2006 invites your participation in the premier annual multimedia conference, covering all

aspects of multimedia computing: from underlying technologies to applications, theory to practice, and servers to networks to devices.

The technical program will consist of plenary sessions and talks with topics of interest in:

- (a) Multimedia content analysis, processing, and retrieval;
- (b) Multimedia networking and systems support;
- (c) Multimedia tools, end-systems, and applications; and
- (d) Foundational sciences of multimedia.

## Keynote Speakers

- [Ken Goldberg](#), UC Berkeley
- [Bradley Horowitz](#), Yahoo!

---

## The Conference Includes:

**State-of-the-Art Tutorials** by leading experts will precede the technical program. The full- and half-day offerings will span a wide variety of topics.

**Awards** will be given to to the best paper and the best student paper. In order to encourage sharing of implementations, this year will also initiate awards for best demo, best art program paper, and the best-contributed open-source software.

**Short papers** will be presented in poster format and are an opportunity for researchers to present new work and ideas in an interactive setting.

**Panels** will consist of discussions on timely and controversial topics.

**Brave New Topics** track is a special sessions track containing papers, which establish foundational sciences and extend the boundaries of multimedia research.

**Technical Demos** will include leading edge work in every area of multimedia technology and its application. An award will be given to the best demo.

**Interactive Art** Program will include long and short papers describing interactive multimedia art works, tools, applications, and technical approaches for creative uses of multimedia content and technology. It will also include an art exhibition.

**Video Demonstrations** allow researchers and artists to demonstrate their tool, system or application without having to bring the equipment for a "live" demo.

**Workshops** on topics of great current interest to members of the multimedia research community will precede the technical program.

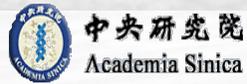
**The Doctoral Symposium** is a venue for doctoral students to present their research and receive feedback from members of the multimedia research field.

**The Open-source Software Competition** is a recent addition to the ACM Multimedia program and 2006 will be our third year in running the competition.

---

Santa Barbara, California

Sponsors and Benefactors



# ACM Multimedia 2006

October 23-27, 2006

Fess Parker's DoubleTree Resort Hotel  
Santa Barbara, CA

\* [Flickr photos of ACM Multimedia 2006](#) \*

## ACM MM 2006 Awards

**Best full technical paper:** Tiling Slideshow

Authors: Jun-Cheng Chen, Wei-Ta Chu, Jin-Hau Kuo, Chung-Yi Wenb, Ja-Ling Wu  
National Taiwan University

**Best short technical paper:** Fourth Frame Forums: Interactive Comics for Collaborative Learning

Authors: Andrew Gordon  
University of Southern California

**Best Art Contribution:** Archeology of Multimedia

Authors: Fabrizio Nunnari, Vincenzo Lombardo, Andrea Valle, Francesco Giordana, Andrea Arghinenti  
University of Torino

**Best Demonstration:** A Real-Time Multi-modal Biofeedback System for Stroke Patient Rehabilitation

Authors: Yin-peng Chen, Wei-wei Xu, Richard Isaac Wallis, Hari Sundaram, Thanassis Rikakis, Todd Ingalls, Loren Olson, Ji-ping He  
Arizona State University

**Open Source Competition Winner:** CLAM: C++ Library for Audio and Music

Authors: Xavier Amatriain, Pan Arumi, David Garcia  
University of California, Santa Barbara and Pompeu Fabra University, Barcelona

**Best Video Program:** Globe4D, Time-traveling with an Interactive Four Dimensional Globe

Authors: Rick Companje, Nico van Dijk, Hanco Hogenbirk, Danica Mast  
Globe4D.com

## News and Updates

- ACM Multimedia 2006 had over 400 attendees and was a smashing success. We hope to see you all in Augsburg, Germany (September 23-29, 2007) for [ACM Multimedia 2007!](#)
- There were a couple problems with the conference DVD - their DVD duplicator was defective. If any registered conference or workshop participants would like a replacement DVD, please email Bob at [typedept3@sheridanprinting.com](mailto:typedept3@sheridanprinting.com) with your name and address, and we will mail one to you. Here are the online fixes:
  - The VSSN'06 Workshop proceedings may not be readable. They are now available online [here](#).
  - One of the video program files did not copy properly. It is available [here](#).
- The report from the SIGMM 2003 Retreat can be found [here](#).

## General Information

ACM Multimedia 2006 invites your participation in the premier annual multimedia conference, covering all aspects of multimedia computing: from underlying technologies to applications, theory to practice, and servers to networks to devices.

The technical program will consist of plenary sessions and talks with topics of interest in:

- (a) Multimedia content analysis, processing, and retrieval;
- (b) Multimedia networking and systems support;
- (c) Multimedia tools, end-systems, and applications; and
- (d) Foundational sciences of multimedia.

## Keynote Speakers

- [Ken Goldberg](#), UC Berkeley
- [Bradley Horowitz](#), Yahoo!

---

## The Conference Includes:

**State-of-the-Art Tutorials** by leading experts will precede the technical program. The full- and half-day offerings will span a wide variety of topics.

**Awards** will be given to to the best paper and the best student paper. In order to encourage sharing of implementations, this year will also initiate awards for best demo, best art program paper, and the best-contributed open-source software.

**Short papers** will be presented in poster format and are an opportunity for researchers to present new work and ideas in an interactive setting.

**Panels** will consist of discussions on timely and controversial topics.

**Brave New Topics** track is a special sessions track containing papers, which establish foundational sciences and extend the boundaries of multimedia research.

**Technical Demos** will include leading edge work in every area of multimedia technology and its application. An award will be given to the best demo.

**Interactive Art** Program will include long and short papers describing interactive multimedia art works, tools, applications, and technical approaches for creative uses of multimedia content and technology. It will also include an art exhibition.

**Video Demonstrations** allow researchers and artists to demonstrate their tool, system or application without having to bring the equipment for a "live" demo.

**Workshops** on topics of great current interest to members of the multimedia research community will precede the technical program.

**The Doctoral Symposium** is a venue for doctoral students to present their research and receive feedback from members of the multimedia research field.

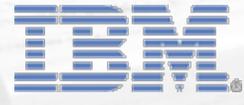
**The Open-source Software Competition** is a recent addition to the ACM Multimedia program and 2006 will be our third year in running the competition.

---



Santa Barbara, California

Sponsors and Benefactors



## Important Dates

<b>Feb. 15, 2006</b>	Early Bird Workshop Proposal Deadline
<b>Mar. 01, 2006</b>	Early Bird Workshop Decision
<b>Apr. 10, 2006</b>	Research Paper Abstract Submission Deadline
<b>Apr. 17, 2006</b>	Research Paper/Panel/Workshop Submission Deadline
<b>Apr. 24, 2006</b>	Tutorial Submission Deadline
<b>May. 15, 2006</b>	Brave Session Proposal Submission Deadline
<b>Jun. 01, 2006</b>	Interactive Art Program Long Paper/Exhibition Submission Deadline
<b>Jun. 01, 2006</b>	Short Paper/Video Program/Interactive Art Program Short Paper/Open Source/Doctor Program/Demo Proposal Submission Deadline
<b>Jun. 02, 2006</b>	Open Source Submission Deadline
<b>Jun. 10, 2006</b>	Video Program Submission Deadline
<b>Jun. 21, 2006</b>	PC Meeting, Microsoft Redmond
<b>Jul. 01, 2006</b>	Notification of acceptance (full papers)
<b>Jul. 10, 2006</b>	Notification of acceptance (short papers)
<b>Aug. 01, 2006</b>	Camera-ready papers due - conference
<b>Aug. 14, 2006</b>	<b>Camera-ready papers due - workshops (strict deadline - no extensions!)</b>
<b>Sept. 23, 2006</b>	<b>Advanced Registration Deadline</b>
<b>Sept. 29, 2006</b>	<b>Hotel cutoff date for guaranteed reservations</b>
<b>Oct. 23, 2006</b>	<b>MM06 Tutorials</b>
<b>Oct. 24-26, 2006</b>	<b>Main conference</b>
<b>Oct 26-27, 2006</b>	<b>MM06 Workshops</b>
<b>Oct 22-28, 2006</b>	<b>MM06 Interactive Arts Program</b>

## Instructions for Presenters

**Posters:** Easels will be provided to attach your printed poster to, as well as staples, tape, or an appropriate method of attachment. The **easels** are large enough to hold the ISO AO size (33" width by 46" height, or 89cm by 118cm). Posters should be set up before the poster session begins. An author should be at the poster at the specified times (see the conference schedule).

You want your poster to be visually appealing and interesting, but most of all, **make sure that your poster is easy to read from a distance**. Fonts for text should be at a minimum **24 point**, without exception. If your text does not fit on the poster at this size, then this is a good indication that you need to reduce the amount of text, not the font size. Use an easy to read font (e.g., Times, Palatino or other serif font) for the text. A good font size for titles is 72 point.

**Oral presentations:** See [here](#)

**Demo presentations:** Please arrive at least 30 minutes prior to the beginning of your demo session to set up your demo. Each demo station will be provided with a table (3x4 ft ?), and power. There will be wireless Internet connection in the room. There is (NO) provision to display posters at the demo session.

This is NO provision for a slideshow or talk about your demo. The audience will walk around the room and will interact with you on a one-on-one basis. You are expected to be available at your demo station for the duration of the demo session.

**Video presentations:** Authors of the video demonstrations are asked to be present for the presentation of their videos. Each video will be followed by a short question and answer session from the audience. The authors should be prepared to answer questions regarding their videos.

---

## Instructions for preparing camera-ready papers

[Click here for detailed camera-ready paper instructions.](#)

Camera-ready papers are due **August 1st** for the following

- Full papers, short papers, demo papers, panel descriptions, Art full papers, Art short papers, Art exhibition papers, Foundational Sciences, video papers, doctoral symposium papers, open source contest

Camera-ready papers are due **August 14th** for the following:

- All workshop papers

Please note that the page deadlines are firm - there is no opportunity to purchase extra pages.

**[Note: [Click here for information on student conference participation grants](#)]**

# ACM MM 2006 People

## Organizing Committee

### Program Committee

### Contacts

## Organizing Committee

### General Co-Chairs

- [Klara Nahrstedt](#) (UIUC)
- [Matthew Turk](#) (UCSB)

### Program Co-Chairs

- [Yong Rui](#) (Microsoft Research)
- [Wolfgang Klas](#) (Universität Wien)
- [Ketan Mayer-Patel](#) (UNC)

### Local Arrangement

- [Xavier Amatriain](#) (Lead, UCSB)
- [Zoran Dimitrijevic](#) (Google)
- [Ankur Jain](#) (UCSB)

### Short Paper Co-Chairs

- [Brian Bailey](#) (UIUC)
- [Belle Tseng](#) (NEC USA)
- [Nalini Venkatasubramanian](#) (UCI)

### Tutorial Co-Chairs

- [Dick Bulterman](#) (CWI)
- [Radu Marculescu](#) (CMU)

### Workshop Co-Chairs

- [Rainer Lienhart](#) (U. Augsburg, Germany)
- [Alan Hanjalic](#) (Delft University of Technology, The Netherlands)

### Panel Co-Chairs

- [Wei-Ying Ma](#) (Microsoft Research, Asia)
- [Lawrence A. Rowe](#) (UC, Berkeley)

### New Foundational/Application Topics

- [Jonathan Foote](#) (FX Palo Alto)
- [John Smith](#) (IBM)

### Video Program

- [Wuchi Feng](#) (PSU)

### Demonstration Co-Chairs

- [Baochun Li](#) (U. Toronto)
- [Raju Rangaswami](#) (FIU)

### Doctorial Symposium

- [Reza Rejaie](#) (U. Oregon)

### Publicity Co-Chairs

- [Wolfgang Effelsberg](#) (U. Mannheim, Germany)
- [Mark Liao](#) (Academia Sinica, Taiwan)
- [Qi Tian](#) (UT San Antonio)

### Proceedings Chair

- [Roger Zimmermann](#) (USC)

### Registration

- [Kingshy Goh](#) (Proximex)
- [Yi Wu](#) (Intel Research)

### Treasurer

- [Alex Kouznetsov](#) (UCSB)

### Travel Grant

- [Carsten Griwodz](#) (U. Oslo)

### Web Chair

- [Gang Wu](#) (UCSB)

### SigMM Chair

- [Ramesh Jain](#) (UCI)

## Open Source Competition

- [Scott Brandt](#) (UCSC)

## Interactive Art Program

- Alejandro Jaimes (Fuji Xerox, Japan)
- [George Legrady](#) (UCSB)
- [Lonce Wyse](#) (Institute for Infocomm Research & NUS, Singapore)

## Program Committee

## Content Track

- [Yong Rui](#) (Microsoft Research, USA, Program Co-Chair)
- Kiyoharu Aizawa (University of Tokyo, Japan)
- Lalitha Agnihotri (Philips Research, USA)
- Shih-Fu Chang (Columbia University, USA)
- Tat-Seng Chua (National U Singapore, Singapore)
- Alberto Del Bimbo (Univ of Florence, Italy)
- Chitra Dorai (IBM T.J. Watson, USA)
- Abed El Saddik (U Ottawa, Canada)
- Stephan Fischer (SAP Research, Germany)
- Daniel Gatica-Perez (IDIAP, Switzerland)
- Yihong Gong (NEC Labs, USA)
- William Grosky (Univ of Michigan, USA)
- Amarnath Gupta (UCSD, USA)
- Alan Hanjalic (Delft University, Netherlands)
- Alexander Hauptmann (CMU, USA)
- Kien Hua (University of Central Florida, USA)
- Benoit Huet (Eurecom, France)
- Mohan S Kankanhalli (National U Singapore, Singapore)
- John Kender (Columbia University, USA)
- Jiebo Luo (Kodak, USA)
- Michael Lew (Liacs, Netherlands)
- Rainer Lienhart (Univ of Augsburg, Germany)
- Mark Liao (Academia Sinica, Taiwan)
- Chih-Jen Lin (National Taiwan U)
- Weiying Ma (Microsoft Research Asia, China)
- Nasir Memon (Polytechnic University, USA)
- Bernard Merialdo (Eurecom, France)
- Apostol Natsev (IBM T J Watson, USA)
- Shin'ichi Satoh (NII, Japan)
- Nicu Sebe (U. of Amsterdam, The Netherlands)
- Alan Smeaton (Dublin City University, Ireland)
- John Smith (IBM Research, USA)
- Hari Sundaram (Arizona State University, USA)
- Svetha Venkatesh (Curtin University, Australia)
- Paul Viola (Microsoft Research, USA)
- James Wang (Penn State U, USA)

- Lynn Wilcox (FXPAL, USA)
- Marcel Worring (University of Amsterdam, Netherlands)
- Changsheng Xu (I2R, Singapore)
- Cha Zhang (Microsoft Research, USA)

## Systems Track

- **Ketan Mayer-Patel** (UNC, USA, Program Co-Chair)
- Sid Ahuja (Bell Labs, USA)
- Kevin Almeroth (UCSB, USA)
- Nina Bhatti (HP, USA)
- Surendar Chandra (Univ of Notre Dame, USA)
- Alberto Del Bimbo (Univ of Florence, Italy)
- Mark Claypool (Worcester Polytechnic Institute, USA)
- Pedro Cuenca (U. de Catilla-La Mancha, Spain)
- Zoran Dimitrijevic (Google, USA)
- Carole Dulong (Intel, USA)
- Wu-Chi Feng (Portland State U, USA)
- David Gotz (IBM Research, USA)
- Carsten Griwodz (U. of Oslo, Norway)
- Baochun Li (U. of Toronto, Canada)
- Jie Liu (Microsoft Research, USA)
- Laurent Mathy (Lancaster U., UK)
- Andreas Mauthe (Lancaster U., UK)
- Wei Tsang Ooi (National U. of Singapore, Singapore)
- B. Prabhakaran (U. of Texas at Dallas, USA)
- Reza Rejaie (U. of Oregon, USA)
- Prashant Shenoy (U. of Massachusetts, USA)
- Shervin Shirmohammadi (U. of Ottawa, Canada)
- Michael Vernick (Avaya Labs Research, USA)
- Dongyan Xu (Purdue U., USA)
- Michael Zink (U. of Massachusetts, USA)

## Applications Track

- **Wolfgang Klas** (Universit t Wien)
- Ramazan Aygun (University of Alabama in Huntsville, USA)
- Noboru Babaguchi (Osaka University, Japan)
- Susanne Boll (University of Oldenburg, Germany)
- Dick Bulterman (CWI, The Netherlands)
- Stavros Christodoulakis (Technical University of Crete, Greece)
- Jana Dittman (Otto-von-Guericke-University, Germany)
- Borko Furht (Florida Atlantic Univ., USA)
- Jim Gemmell (Microsoft Research, USA)
- Forouzan Golshani (Wright State University, USA)
- Thomas Haenselmann (University of Mannheim, Germany)
- Lynda Hardman (CWI, The Netherlands)
- Brigitte Kerhervé (Université du Québec à Montréal, Canada)
- Ross King (ARC Research Studio Digital Memory, Austria)
- Kiyoshi Kogure (Advanced Telecommunications Research Institute International, Japan)
- Stefan Leitich (University of Vienna, Austria)
- Peiya Liu (Siemens Corporate Research Inc., USA)
- Max Muehlhaeuser (Technical University Darmstadt ,Germany)
- Frank Nack (V2\_ Institute for the Unstable Media, The Netherlands)
- Chandrasekhar Narayanaswami (IBM T. J. Watson, USA)
- Silvia Pfeiffer (CSIRO Australia ,Australia)
- Gopal Pingali (IBM TJ Watson Research Center, USA)
- Thomas Plagemann (University of Oslo, Norway)
- Christoph Rensing (Technical University of Darmstadt, Germany)
- Yuqing Song (University of Michigan at Dearborn, USA)

- Savitha Srinivasan (IBM Almaden Research Center, USA)
- VS Subrahmanian (University of Maryland, USA)
- Ye Wang (National University of Singapore, Singapore)
- Gerd Utz Westermann (University of California at Irvine, USA)
- Maia Zaharieva (University of Vienna, Austria)

## Contacts

### General Conference Information

For any information regarding the conference (dates, submissions, etc.) email the [General Co-Chairs](#).

### Website Information

For any questions or comments regarding this website email [Gang Wu and Matthew Turk](#)

# Conference Program

## General Schedule

<b>Monday, Oct. 23</b>	Tutorials, arts program
<b>Tuesday, Oct. 24</b>	Main conference, arts program, evening reception
<b>Wednesday, Oct. 25</b>	Main conference, arts program, conference banquet dinner
<b>Thursday, Oct. 26</b>	Main conference, workshop, arts program
<b>Friday, Oct. 27</b>	Workshops, arts program

## Overview

Monday	am	Half-day tutorials	Full-day tutorials	Interactive Arts Exhibition (at UCSB)
	pm	Half-day tutorials		
Tuesday	am	Keynote speaker	Best Paper Session	
	pm	Content, Applications, Systems, Arts	Poster Session	
	pm	Conference reception		
Wednesday	am	Content, Applications, Systems, Arts	Demos	
	pm	Content, Applications	Posters, Demos, Panel Brave New Topics Doctoral Symposium	
	pm	Conference banquet		
Thursday	am	Keynote speaker	Content, systems sessions Open source and video demos	
	pm	Content, Foundations discussion	MIR Workshop	
Friday	am	Workshops		
	pm			

[Click here for a detailed conference schedule](#)

## Conference Events

- **Full Papers**
- **Short Papers**
- **Workshops**
- **Tutorials**
- **Demos**
- **Doctoral Symposium**
- **Video Program**
- **Open Source Contest**
- **Panel**
- **Foundational Sciences**
- **Interactive Arts Program Papers and Exhibition**

**Keynote Speakers**



**Invited Talk:** Networked Robots  
**Invited Speaker:** Ken Goldberg, UC Berkeley

Tuesday, October 24th, 8:45am



**Invited Talk:** Title TBD  
**Invited Speaker:** Bradley Horowitz, Yahoo!

Thursday, October 26th, 9:00am

# Workshops

ACM Multimedia 2006 plans to include seven day-long workshops on topics in new and emerging areas of interest to members of the multimedia research community. The workshops are scheduled to be on 26 - 27 October 2006, after the main conference.

## Scheduled Workshops:

### 1. Workshop on Multimedia Information Retrieval (MIR 2006)

Extending beyond the borders of culture, art, and science, the search for digital information is one of the major challenges of our time. MIR 2006 is a peer-reviewed meeting for scientific researchers and users to discuss important challenges in multimedia retrieval. Following the success of the five previous MIR workshops held in conjunction with the ACM Multimedia Conferences, the purpose of the 8th ACM SIGMM International Workshop on Multimedia Information Retrieval (MIR 2006) is to bring together researchers, developers, and practitioners from academia and industry. We are soliciting original papers that address a wide range of issues in multimedia information retrieval.

#### Important Dates

**11 July** MIR'2006 Submission Deadline  
(extended)

#### Organizers

- **James Z. Wang** (PSU, USA)
- **Nozha Boujema** (IMEDIA Research Group, France)

---

### 2. Workshop on Capture, Archival and Retrieval of Personal Experiences (CARPE 2006)

Personal storage of all one's media throughout a lifetime has been desired and discussed since at least 1945, when Vannevar Bush published *As We May Think*, positing the "Memex" device "in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility." His vision was astonishingly broad for the time, including full-text search, annotations, hyperlinks, virtually unlimited storage and even stereo cameras mounted on eyeglasses. Storage, sensor, and computing technology have progressed today to the point of making Memex feasible and even affordable. Indeed, we can now look beyond Memex at new possibilities. In particular, while media capture has typically been sparse throughout a lifetime, we can now consider continuous archival and retrieval of all media relating to personal experiences.

Following the success of the two previous CARPE workshops held in conjunction with the ACM Multimedia Conferences, and the Pervasive 04 workshop on Memory and Sharing of Experience, this one-day workshop aims to foster deeper and wider discussion on issues related to capture, archival and retrieval of personal experiences. We invite regular and position papers as well as demonstrations on relevant topics, including, sensors, wearable, data storage and management, content analysis, user interfaces, applications, etc.

#### Important Dates

**14 July** CARPE'2006 Submission Deadline  
(extended)

#### Organizers:

- **Kenji Mase** (Nagoya U., Japan)

---

### 3. Workshop on Multimedia Content Protection and Security (MCPS 2006)

#### Important Dates

**14 July** MCPS'2006 Submission Deadline  
(extended)

#### Organizers:

- **Edward J. Delp** (Purdue Univ.)
  - **Dulce B. Ponceleon** (IBM Almaden Research Center)
  - **Ginger Myles** (IBM Almaden Research Center)
- 

### 4. Workshop on Human-Centered Multimedia (HCM 2006)

Human-centered computing (HCC) lies at the crossroads of multiple disciplines and research areas that are concerned both with understanding human beings and with the design of computational methods. Researchers and designers of human-centered computing methods and systems include engineers in multimedia information processing and systems, scholars in psychology, cognitive science, sociology, and graphic designers. The research in this area deals with understanding humans, both as individuals and in social groups, by focusing on the ways that human beings adopt, adapt, and organize their lives around computational technologies. Human-centered design of computational tools attempts to address problems that traditional human computer interaction does not generally address. For example, the design of algorithms and systems, and not just of interaction, are some non-traditional design problems that are of interest to HCC researchers. The focus of the workshop is the multimedia aspect of the HCC paradigm.

#### Important Dates

**10 July** HCM'2006 Submission Deadline  
(extended)

#### Organizers:

- **Daniel Gatica-Perez** (IDIAP, Switzerland)
  - **Alejandro Jaimes** (Fuji Xerox, Japan)
  - **Nicu Sebe** (Univ. of Amsterdam, The Netherlands)
- 

### 5. Workshop on Video Surveillance and Sensor Networks (VSSN 2006)

#### Important Dates

**5 July** VSSN'2006 Submission Deadline  
(extended)

#### Organizers:

- **J. K. Aggarwal** (University of Texas, USA)
  - **Rita Cucchiara** (University of Modena and Reggio Emilia, Italy )
  - **Andrea Prati** (University of Modena and Reggio Emilia, Italy)
- 

### 6. Workshop on Audio and Music Computing for Multimedia (AMCMM 2006)

Multimedia is usually defined as the combination of text, graphics, full-motion video, and sound into an integrated application. The final frontier, sound, which includes general-purpose audio, music, and speech, is an integral, but often neglected, component of the field. Traditionally, the audio and the image research communities have developed independently, and almost in perfect isolation one with another. Our goal is to set a new trend of active contribution by the audio and music community through the initiating of a workshop on **Audio and Music Computing for Multimedia (AMC MM)**, specifically targeted to researchers in this community. The workshop will also feature special sessions and keynotes that will be of relevance and interest to the general multimedia community, so as to foster a better understanding of the unique challenges and problems of interest in the audio realm. The proposed AMCMM workshop will be held in the last day of the conference in parallel with other successful workshops

### Important Dates

**5 July**      **AMCMM'2006 Submission Deadline**

### Organizers:

- **Xavier Amatriain** (UCSB, USA)
- **Elaine Chew** (USC, USA)
- **Jonathan Foote** (Rotorbrain Industries)

# Tutorials

## Scheduled Tutorials:

Title (click on title for more information)	Presenters	Length
<a href="#">Peer-to-Peer Multimedia Applications</a>	Jin Li	Full-day
<a href="#">Interactive Digital Television and Multimedia Systems</a>	Pablo Cesar, Konstantinos Chorianopoulos	Half-day
<a href="#">Flexible Modelling and Performance Debugging of Real-Time Embedded Multimedia Systems</a>	Samarjit Chakraborty	Half-day
<a href="#">Computer Audition: An introduction and research survey</a>	Shlomo Dubnov	Half-day
<a href="#">Data Mining and Information Retrieval in Time Series/Multimedia Databases</a>	Eamonn Keogh	Half-day
<a href="#">Recent developments in video compression standards and their impact on embedded platforms: from scalable to multi-view video coding</a>	Iole Moccagatta	Half-day
<a href="#">Multimedia Content Protection</a>	Dulce Ponceleon, Nelly Fazio	Half-day
<a href="#">Semantic Indexing and Retrieval of Video</a>	Marcel Worring, Cees Snoek	Half-day

## Call For Tutorials

Proposals are solicited for tutorials of either a half day (3 hours plus breaks) or full day (6 hours plus breaks). Submissions should be made to the Tutorials Chair and should include a cover sheet and an extended abstract. The cover sheet should specify:

1. The length of the tutorial
2. The intended audience (introductory, intermediate, advanced);
3. Complete contact information for the contact person and other presenters
4. A brief biography ( max. 2 paragraphs) for each presenter.

The extended abstract should be 4 to 6 pages, and should include an outline of the tutorial, along with descriptions of the course objectives and course materials.

## Important Dates

**24 April Tutorials Submission Deadline**

## Submission Instructions

**The submission deadline for tutorial proposals is extended to April 24th 2006.** Proposals in the form of a regular Call for Papers plus some background information are more than welcome. This background information should include:

- Names and CVs of the organizers
- Motivation for having a tutorial on the proposed topic
- Explanations why the organizers think their proposed tutorial would be successful and of a high impact.

Please submit all this information to the co-chairs.

## Contacts

For any questions regarding tutorials please email the tutorial co-chairs:

- **Dick Bulterman** (CWI)
- **Radu Marculescu** (CMU)



# Registration

**On-line registration is now closed. On-site registration will be available at the conference.**

Electronic registration is now available. Please review the registration information below before proceeding. [Click here to register on-line](#). Alternatively, a [PDF registration form](#) is provided for those who do not wish to pay on-line. Mailing instructions can be found on the form. Note that advanced registration closes on Sept 23.

**Note:** [Click here for information on student conference participation grants](#).

## Important Dates

**23 Sept 2006** Advance Conference Registration Deadline is 11:59pm EST, 23 September 2006.

**Each main conference accepted paper must have at least one author registered by the advance registration deadline, for the paper to be included in the final conference program and the proceedings.**

For workshop accepted papers, the same registration policy applies with respect to workshop registration.

## Registration Fees

Advance registrations are encouraged. On-site registration will be available at the conference site.

There are several registration options, everything from a single half-day tutorial to an "all-inclusive" package deal. [Click here for the registration fee schedule](#).

## Notes

- All conference fees are in US Dollars (USD).
- Main conference registration includes:
  - + Attendance to the following sessions: keynote addresses, full papers, short paper presentation, technical demos, panels, video program, panel, doctoral symposium, and arts exhibition
  - + Main conference proceedings (on DVD)
  - + Attendance to social programs: receptions, conference banquet, coffee breaks and lunches (Note: student registration does not include the conference banquet; **it does include the reception and the conference lunches.**) [Updated 9/18/06]
- Extra banquet tickets (Wed. evening) may be purchased for \$75 each.
- Extra reception tickets (Tues. evening) may be purchased for \$25.
- Extra lunch tickets may be purchased for \$35 each.
- Extra DVD proceedings (while available) may be purchased for \$25.

**Payment:** We accept payment through Master Card, Visa Card, Amex Card and Check. Checks should be drawn on US Dollars and from a US bank. Checks should be made payable to "ACM/MM06". Please indicate your registration reference number on the check. The check and a copy of the registration invoice should be mailed to: **ACM, Attn: Brooke Hardy, 2 Penn Plaza, Ste. 701, New York, NY 10121-0701, USA**. Note that checks are not accepted for on-site registration during the conference, only credit cards will be accepted.

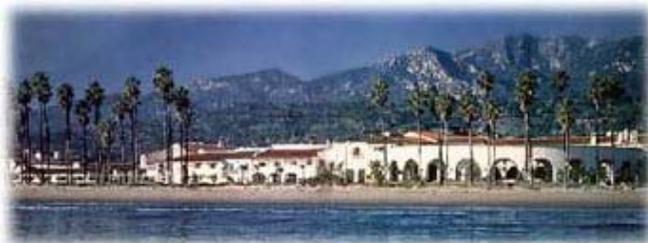
- **Registration Refund Policy:** Full refunds (minus a \$50 handling fee) may be granted before 28 September, 2006. After 28 September, 2006, no refunds will be made. Refund requests must be made in writing (by mail or fax) to the registration chairs. Please fax requests to **Attn: Yi Wu, 408-653-9677**, or mail in the request to: **Attn: Yi Wu, 2200 Mission College Blvd, SC12-303, Santa Clara, CA, 95054-1537, USA**.
- Registration receipts will be issued during the conference at the registration desk.
- Student registrants need to show proof of student status (e.g., Student ID or a letter from advisor or a letter from the department) when picking up the conference package. Also, ACM members need to include their ACM membership number on the registration form.
- **Visa Letters:** International registrants should be particularly aware and careful about visa requirements, and should plan travel well in advance. All visa inquiries must be handled by ACM HQ. Please send your request for a letter in support of a visa application to Stephanie Smith (smith\_s@hq.acm.org), and include your name, mailing address, and fax number, as well as the name of the conference you are attending and the registration reference number. (Authors of papers/posters should also include the title). Please note that ACM does not issue formal letters of invitation to any of its conferences.

The sites [www.unitedstatesvisas.gov](http://www.unitedstatesvisas.gov) and [travel.state.gov](http://travel.state.gov) have information about obtaining a visa for those traveling to the United States. Both sites have links to websites for U.S. embassies and consulates worldwide. The embassy and consulate websites have very helpful information about procedures, timelines, communities served, required documentation, and fees.

## Contact

For queries regarding registration, please contact registration chairs [Kingshy Goh](#) and [Yi Wu](#).

## Accommodations, Travel, and Local Information



The main conference hotel is **Fess Parker's DoubleTree Resort Hotel**, right across from the Pacific Ocean in beautiful Santa Barbara, California. The conference room rate is \$169 (+ tax). Parking is free. Internet access in the guest rooms is included. There will also be wireless internet access in the meeting rooms.



### Fess Parker's Doubletree Resort Hotel

633 East Cabrillo Boulevard  
Santa Barbara, CA 93103  
(805)-564-4333



### Maps and Directions

From Route 101 North or Route 101 South

Click [here](#) to make reservations at the Fess Parker's DoubleTree Resort Hotel. Reservations must be made online by **September 21, 2006**.

Click [here](#) for further information about the Doubletree (cancellation policy, special needs, extra guests, taxes, check-in time, etc.).

### Other Hotels (Please mention ACM Multimedia for preferred rate):

- [The Hotel Oceana Santa Barbara](#)

To make a reservation call 800-965-4577.

Book by 21 September to get a preferred rate of \$145 (+ tax).

CheckIn: 4pm, CheckOut: 12 noon.

- [Santa Barbara Inn](#)

Book by 21 September to get a preferred rate of \$169.00 - \$189.00 (+ tax).

Complimentary breakfast, Free Wireless Internet.

CheckIn: 4pm, CheckOut: 12 noon.

- [Best Western Beachside](#)

Preferred rate: \$120 (+ tax).

Call 800-5286556 to make a reservation.

- [Santa Barbara Ramada Limited](#)

Book by calling 800-654-1965 or 805-964-3511.

Preferred rates:

Sunday - Thursday (Mountainside room - King size Bed - \$99.00 / Gardenside room - 1 King or 2 Queen size beds - \$119.00)

Friday & Saturday (Mountainside room - King size Bed - \$129.00 / Gardenside room - 1 King or 2 Queen size beds - \$149.00)

Complimentary breakfast, Internet Connection,

## Local Information

### Traveling to Santa Barbara



Santa Barbara is about 90 miles north of Los Angeles. Conference attendees can fly directly into [Santa Barbara Airport](#) (SBA) and take a shuttle or taxi to the hotel, or fly into [Los Angeles International Airport](#) (LAX) and drive up to Santa Barbara - which can take from 1.5 to 2.5 hours, depending on traffic.

The Santa Barbara Airport has direct service to and from Los Angeles, San Jose, San Francisco, Denver, Phoenix and Las Vegas. The airport is right next to UCSB and 15-20 minutes from downtown Santa Barbara (and the Doubletree Resort Hotel).

### Weather and What to Wear

[Santa Barbara weather](#) in late October is typically quite nice: daily highs of around 68-70 degrees F (20-21 degrees C) and nighttime lows around 50 degrees F (10 degrees C). This is not "swimming in the ocean" weather (the water is freezing anyway!), but it's very nice "walking on the beach" weather. Some of the conference events will be outside (lunches, the reception, the banquet), so be sure to bring a jacket or sweater, as it does get cool in the evenings.

### Places to eat or purchase food near the conference hotel

[Here](#) is a list of nearby places (with map), and [here](#) is an online map showing their locations.

### Things To Do in Santa Barbara



We would like to thank [santabarbara.com](#) for providing us a very useful brochure about many interesting things in Santa Barbara. Please following [this link](#) to find out 101 free things to do in Santa Barbara.

Other Santa Barbara links:

- <http://www.newspress.com/insidesb>
- <http://www.santabarbaraca.com/>
- <http://www.sbchamber.org/>
- <http://www.lonelyplanet.com/worldguide/destinations/north-america/usa/santa-barbara/>
- <http://www.santabarbaradirect.com/default.htm>
- <http://www.santabarbaraca.gov/home.htm>
- <http://www.totalsantabarbara.com/beach.shtml>

## Sight Seeing in Santa Barbara

Bamboo from the Asia Tsunami



Santa Barbara's Beaches



Sunsets in Santa Barbara



The University of California, Santa Barbara's Lagoon



## Sponsor Opportunities and Benefits

The conference organizers warmly invite your support of ACM Multimedia 2006. Corporate support publicizes your organization's interest in and commitment to the multimedia field. An event with international visibility, ACM MM 2006 is the perfect place to inform leaders and students in the field, as well as the interested general public, about your company's activities and products. Regarding the details on how to become an ACM MM 2006 corporate supporter, please read the [linked document](#) in the pdf format.

ACM Multimedia 2006 gratefully acknowledges the support of the following benefactors and sponsors:

### Benefactors:



### Sponsors:



中央研究院  
Academia Sinica



PHILIPS



### Student Travel Grant:



# ACM SIGMM Retreat Report on Future Directions in Multimedia Research

(Final Report March 4, 2004)

*Lawrence A. Rowe*

Computer Science Division – EECS  
University of California  
Berkeley, CA 94720-1776

*Ramesh Jain*

School of Elec. & Comp. Eng.  
Georgia Institute of Technology  
Atlanta, GA 30332-0250

## **Abstract**

The ACM Multimedia Special Interest Group was created ten years ago. Since that time, researchers have solved a number of important problems related to media processing, multimedia databases, and distributed multimedia applications. A strategic retreat was organized as part of ACM Multimedia 2003 to assess the current state of multimedia research and suggest directions for future research. This report presents the recommendations developed during the retreat. The major observation is that research in the past decade has significantly advanced hardware and software support for distributed multimedia applications and that future research should focus on identifying and delivering applications that impact users in the real-world.

The retreat suggested the community focus on solving three grand challenges: 1) make authoring complex multimedia titles as easy as using a word processor or drawing program, 2) make interactions with remote people and environments nearly the same as interactions with local people and environments, and 3) make capturing, storing, finding, and using digital media an everyday occurrence in our computing environment. The focus of multimedia researchers should be on applications that incorporate correlated media, fuse data from different sources, and use context to improve application performance.

## **1. Introduction**

The word *multimedia* has many definitions, but most people incorporate the idea of combining different media into one application such as an educational title that uses text, still images, and sound to explain or demonstrate a concept. Another example is a distributed collaboration that allows people at different locations to work together on a document or jointly operate a remote instrument (e.g., a telescope) using streaming audio, video, and application-specific data.

The first computer-based use of multimedia occurred in the early 1960's when text and images were combined in a document. Soon thereafter applications incorporated continuous media. Audio, video, and animations are examples of continuous media. These media are combined, in other words synchronized, using a time-line to specify when they should be played.

The multimedia research community is inherently multi-disciplinary in that it includes people with interests in operating systems, networking, signal processing, graphics, databases, human-computer interfaces, cognitive science, and various application communities (e.g., education, architecture, art, etc.). Consequently, several ACM Special Interest Groups joined

together to organize the First International Conference on Multimedia in 1993 (MM'93). The conference was co-located with the SIGGRAPH Conference being held that year in Anaheim, California. MM'93 was very successful and led to the formation of SIG Multimedia (SIGMM) in early 1994 to serve the multimedia research community.

The annual Multimedia Conference was separated from SIGGRAPH in 1994 to encourage more interaction amongst participants. This annual conference is the premier conference for the publication of multimedia research as measured by attendance and the selectivity of the papers published. Most years more than 300 papers are submitted to the program committee of which 15-18% are accepted for publication. The conference has grown over the years to include formal papers, posters, demonstrations, videos, and a dissertation review program.

Discussions at MM'01 in Ottawa, Canada suggested it was time for senior members of the research community to meet and discuss the current state and future direction for multimedia research. Some members of the community believed that multimedia research, as represented by publications at the annual conference, was addressing narrow topics with limited impact rather than addressing major problems that will have wider impact on technology for real and emerging applications. At the same time, people inside and outside the community questioned why the Multimedia Conference has not grown into a major event similar to the SIGGRAPH conference. The belief is that multimedia is such a hot topic that the conference should attract several thousand people rather than the 200-300 people that typically attend.

Professors Lawrence A. Rowe and Ramesh Jain, the past and current SIGMM Chairs, respectively, organized the retreat with advice from the SIGMM Executive Committee. A two-day retreat was held in conjunction with MM'03 in Berkeley, California. The Executive Committee selected the retreat attendees. Twenty-six researchers, shown in Table 1, participated in the retreat. The goal was to include both academic and industrial researchers from a variety of areas as well as young and old members of the community. Each participant was invited to write a short position paper briefly responding to questions about past research successes, future research directions, and the current state of SIGMM and the annual conference. These position papers were distributed to attendees before the retreat and are being published jointly with this report [SIGMM 2003].

The first day of the retreat was dedicated to discussions about future directions for multimedia research, and the second day focused on organizational issues. This report covers the research recommendations developed during the retreat. These recommendations have been modified somewhat after a public presentation and discussion at MM'03. The organizational issues report will be published separately on the SIGMM Website (<http://www.acm.org/sigmm>).

The remainder of this report is organized as follows. Section 2 presents background on multimedia research over the past decade. Section 3 presents unifying themes, which underlie the field. Section 4 presents three *Grand Challenges* identified as the problems that multimedia researchers should be trying to solve and funding agencies should be supporting. Finally, section 5 discusses topics mentioned at the retreat and in public and private discussions since the initial findings were presented.

**Table 1: SIGMM Retreat Participants**

Sid Ahuja (Lucent)	Wolfgang Klas (U Vienna)
Brian Bailey (UIUC)	Joseph Konstan (U Minn) *
Dick Bulterman (CWI)	Dwight Makaroff (U Saskatchewan) +
Shih-Fu Chang (Columbia)	Ketan Mayer-Patel (U North Carolina)
Tat-Seng Chua (Singapore)	Klara Narhstedt (UIUC) *
Marc Davis (UC Berkeley)	Arturo Pizano (Siemens SCR)
Nevenka Dimitrova (Philips Research)	Thomas Plagemann (U Oslo)
Wolfgang Effelsberg (TU Mannheim)	Lawrence A. Rowe (UCB) *
Jim Gemmell (Microsoft Research)	Henning Schulzrinne (Columbia)
Forouzan Golshani (Arizona State U)	Ralf Steinmetz (TU Darmstadt) *
Nicolas Georganas (U Ottawa) *	Michael Vernick (Avaya)
Ramesh Jain (GaTech) *	Harrick Vin (U Texas)
Martin Kienzle (IBM Research)	Lynn Wilcox (FX PAL)

\* Member SIGMM Executive Committee

+ SIGMM Information Director

## 2. Multimedia Research Background

Multimedia research through the middle 1990's focused on the development of infrastructure to support the capture, storage, transmission, and presentation of multimedia data. Researchers and product developers worked on I/O devices, scheduling algorithms, media representations, compression algorithms, media file servers, streaming and real-time network protocols, multimedia databases, and tools for authoring multimedia titles. Driving applications included CD-ROM playback, non-linear audio/video editing, videoconferencing, multimedia content analysis and search, lecture webcasting, video-on-demand (VOD), and video games. While many companies focused on stand-alone multimedia applications (e.g., CD-ROM playback), the research community recognized early on that the most important and difficult applications involved distributed multimedia, sometimes called "networked multimedia," and multimedia database applications. Examples are VOD, videoconferencing, and algorithms to analyze and search music and image databases.

Research on compression algorithms, which began in the 1950's, has led to the development of standards for low bandwidth audio and video coding that support video conferencing applications and wireless telephony. Low-latency coding is important for these applications because human communication requires bounded end-to-end delay. Compression standards were developed in the 1980's and early 1990's for low-bandwidth, high-quality audio coding to support telephony and for high-quality video coding to support home entertainment applications (e.g., satellite receivers and personal video recorders) and transmission of broadcast programming (e.g., delivering live news and sporting events from anywhere in the world). This research on coding has yielded many algorithms that operate at different points in the space, time, bandwidth, and computational complexity space. While research will continue on further improvements in coding algorithms, many researchers believe dramatic improvements in coding will require significant breakthroughs.

Computer network research has been an active area of multimedia research since the 1980's. New protocols were developed with standard wire formats for packet audio and video that enabled continuous media players to recover when packets are lost. Significant changes to the standard Internet model were explored to support bounded delay protocols and resource management to insure that time-critical streaming media packets are delivered before less time-critical packets (e.g., FTP). Multicast protocols were designed, implemented, and deployed to support broadcast and small group collaboration applications. Today researchers are developing new protocols for wireless networks. Considerable progress has been made on systems and protocols for media streaming, but resource management, scalable multicast protocols, and wireless networking continue to be a challenge.

The conversion to digital media, whether still images taken by a digital camera, an mp3 song downloaded from a music archive, or an mpeg video captured by a desktop video camera or cellphone, and the development of large media databases, which were enabled by the dramatic increase in storage capacity over the past two decades, has led to research on algorithms to automate the analysis, indexing, summarization, and searching of this content. Internet search engines that operate on text data have proven extremely valuable. Next generation search engines will incorporate other media. While some limited successes have been achieved in multimedia analysis and search, digital asset management that solves real-world problems continues to be a challenge.

Many researchers and companies have developed tools for authoring multimedia content. Content examples are video games, web-based hypermedia (i.e., media data with links between components), and CD-/DVD-ROM titles. Non-linear audio and video editors are notable successes. However, creating multimedia content and using it in everyday applications (e.g., email, documents, web titles, presentations, etc.) is still not possible for most users. For example, many colleges and universities regularly webcast lectures of various sorts (e.g., classes, seminars, conferences, etc.). Using this material in an assignment or creating a study guide that includes links to selected clips with annotations is difficult. Better tools are also needed for professional content authors. Specifically, current tools poorly serve artistic content and multi-player game authors.

While early multimedia systems required special-purpose hardware to decode and play continuous media, regardless of whether it was streamed across a network or read from a local storage system, the continuing improvement in semiconductor technology, the introduction of special-purpose instruction sets (e.g., Intel MMX), and the addition of special-purpose processors on graphics adapters has made multimedia playback and media processing a software application available on all modern PC's. Software media processing coupled with the deployment of broadband networking suggests that distributed multimedia applications will become increasingly important over the next decade.

In summary, research over the past several decades has focused on the "nuts & bolts" infrastructure required by multimedia applications. These applications are inherently *real-time*, which means events or processes must respond within a bounded time to an event (e.g., an I/O interrupt), and *isochronous*, which means processing must occur at regular time intervals (e.g., decode and display a video frame 24 times per second). Two fundamental principals were developed: *statistical guarantees* and *adaptation*. Because continuous media has a presentation time-line, a video frame or audio block that is not available at the scheduled playout time is worthless. Early researchers studied resource allocation algorithms that could guarantee on-time availability. However, guaranteed service requires the reservation of too

many resources to prepare for an infrequent worst-case scenario. The development of statistical guarantees allows improved utilization of resources while providing the user a high-quality experience. This high-quality experience is possible because applications can adapt to lost data and limited resources. For example, a decoder can fill-in data that is lost by using redundant information sent in previous packets or by constructing plausible values for missing data. The term *quality-of-service* (QoS) refers to the allocation of resources to provide a specified level of service. QoS management and the development of algorithms and technologies to produce the highest user-perceived quality experience is an important contribution of multimedia research.

### 3. Unifying Themes

The multimedia field, as mentioned above, is inherently multi-disciplinary. Few researchers identify multimedia as their primary research area. More often, researchers identify themselves as being in signal processing, computer systems, databases, user interfaces, graphics, vision, or computer networking. This list of areas ignores the content side of multimedia, whether it be artistic, entertainment, or educational, which must also be considered part of the multimedia research community. One goal of the retreat was to identify the unifying or overarching themes that unite the multimedia field. These themes help to inform us about the nature of multimedia research.

Many important unifying themes were identified during discussions at the retreat. These themes can be organized into three areas. First, a multimedia system or application is composed of more than one media that are correlated. The media can be discrete (e.g., an image or text document) or time-based (e.g., weather samples collected by a sensor network or a video).<sup>1</sup> Different media are correlated but not necessarily time-based or co-located. For example, an artist might put together a still image and a video to evoke a particular response in the viewer. Or, musicians at different locations playing together are using multiple streams of time-based media (audio) created at different geographic locations. Someone listening to the performance hears one sound, most likely from a stereo or multiple channel surround sound system. Notice that this example is multiple streams of the same media type. A virtual clock correlates the different streams. The representation of time and synchronization of time-based events is a fundamental concept for multimedia applications.

The second unifying theme is integration and adaptation. Any distributed application with user interactions must deal with end-to-end performance and user perception. Multimedia applications are cross-layer (e.g., network protocols, software abstractions, etc.) and multi-level (e.g., high-level abstractions down to low-level representations). For example, streaming media requires application-level framing, that is, the application is best at deciding how to pack media data into network packets because only the application knows enough about the data to make framing decisions that will enable recovery when packets are lost. A simple example of multi-level media is the use of different sized images in an application (e.g., thumbnail images in a summary view and large images in a detailed view). Similar ideas have been applied to other media (e.g., video summarization, audio skimming, etc.) and to different applications (e.g., hierarchical coding). Distributed multimedia applications should provide transparent delivery of dynamic content. Content must adapt to the user's

---

<sup>1</sup> While analog data can be continuous, digital data is always discrete. Digital data can be sampled frequently to capture an acceptable representation of the continuous data.

environment. For example, content displayed on a PDA might look and behave differently than content displayed on a large projection screen in a classroom or theatre.

Media integration means that information is conveyed by the relationship between media as well as by the media itself. A simple example is a video composed of a sequence of audio blocks synchronized with a sequence of still images. The user requires both sequences to understand the video. Either media by itself is insufficient. Much of the current research in analysis, compression, and organization considers these sequences separately. Media must be considered separately and jointly to address emerging problems.

A second facet of integration and adaptation is ubiquitous interaction with multiple media. One retreat participant cited the following example. A user should be able to enter a room and interact with various devices and sensors in that space. For example, the user's laptop computer or PDA should sense or query the environment to locate cameras, microphones, printers, presentation projectors and the applications available to manage and use them. It should be easy to access, display, annotate, and modify the media. Contrast this situation with the reality today. A user must explicitly configure his or her computer to tell it what devices and sensors exist and how to use them. Unfortunately, retrieving data from a large collection for display to remote participants or capturing a reference to a subpart of this media along with an annotation currently requires detailed knowledge about media representations, networking, and other components of the system infrastructure. The focus should be on "ease of use" to solve a problem, not on system configuration and operation.

A third facet of integration and adaptation is the emphasis on using multiple media and context to improve application performance. Early research on multimedia content analysis, summarization, and search focused on one media type (e.g., still image or music archive query) and limited context. Researchers are now exploring systems that use information derived from correlated media and context. For example, executing a query to find information about the election of a state governor might involve restricting the search to TV news programs and identifying segments in which a person is shown in the video stream who uses the words "election" and "governor" in the audio stream. Using the type of program (e.g., a chemistry lecture, baseball game, etc.) is an example of using context to guide the search and improve the results. Research on parsing news programs and producing synchronized text transcripts has been very successful. The challenge now is to extend this research to less well-structured environments where transcripts are not provided and the speaker uses different vocabularies. Lecture webcasts or discussion seminars are examples of less well-structured content in an education setting.

The third unifying theme is that multimedia applications are multi-modal and interactive. The conventional interface to a desktop or laptop computer, that is, a two-dimensional windows, mouse, and keyboard interface, is being replaced with new interface modalities (e.g., pen, voice, gesture, touch, smell, etc.) and multiple devices (e.g., PDA's, active badges, tablet computers, projectors with embedded computers connected directly to the network, etc.) and smart spaces. Most applications that use these interfaces will be interactive and require coordination between different ways of specifying an operation (e.g., pushing a button on a computer, gesturing with your hand or speaking a command might specify "move to the next slide" in a presentation). These applications are also inherently multimedia because they incorporate traditional continuous media (e.g., audio and video). Human-computer interactions as well as communication among humans using computer-based applications (e.g., Voice-over-IP, video conferencing, immersive environments, etc.) are

important themes for research on multimedia applications. Several people noted that human-computer interactions in the future would be more like human-to-human communication.

## 4. Grand Challenges

The primary goal of the retreat was to identify a short list of “grand challenges” that multimedia researchers should solve. The idea was to raise awareness of the importance of research that involves “multiple media” and impacts real users as opposed to narrow results that can be published in a journal.

Three grand challenges were identified. The first challenge is to *make authoring complex multimedia titles as easy as using a word processor or drawing program*. Content authoring is expensive and difficult. Most groups that produce hypermedia content use teams of experts supervised by producers and directors. Specialized tools are used for different media (e.g., a word processor for text, a non-linear editor for audio and video, an image editing tool for still images, a 3D modeling system for animations, etc.). These content elements are then combined to produce the title. Combining the different media with rules for synchronization and user interaction requires a programmer and yet another collection of tools. Producing the title, that is, coding the material and physically publishing it (e.g., pressing a DVD or uploading the material to one or more servers) is time-consuming and complex. Moreover, different versions of the title are typically produced for different environments (e.g., TV set-top box, game console, desktop computer, PDA, etc.), which is itself a challenge. Few people have the experience required to use these tools and produce multiple versions of a title.

The multimedia research community should develop the algorithms, heuristics, and tools that will allow average users to produce compelling multimedia content. Users need tools to support creation of different types of content. Some examples are:

1. A teacher needs tools to prepare educational material that includes video demonstrations to show an object and simulations and animations to illustrate dynamic behaviors. Good educational material allows students to explore the underlying principles and objects by modifying the input parameters to a simulation and examining related objects.
2. A travel agent needs tools to prepare material showing places and experiences potential customers might want to visit. This material might include pictures, videos, sounds, and other immersive experiences. It might include live interactions with people at a remote location. It might also include links to material authored by someone who has taken the trip. This title might be composed of slide shows, videos, trip summaries, and links to detailed information about places visited, artifacts seen, places to stay, and methods of transportation.
3. A family member needs tools to prepare material that documents a significant life event such as a birthday, wedding, or birth of a child. For example, the title created for a wedding might include a time-line that relates different events (e.g., proposal, engagement, wedding ceremony, and celebration), formal and informal pictures of the participants and the event, audio and video captured at various stages in the event, and detailed information about the clothes worn by the bride and groom and the locations of the ceremony and reception.

Some excellent tools exist either for particular media (e.g., Photoshop for images, Dreamweaver for websites, Premiere for audio/video, etc.) and particular applications (e.g., PowerPoint for presentations, iMovie for home movies, FrameMaker and Word for documents, etc.). The problem is that the tools are not integrated, do not encourage content re-use, run on different platforms, and are targeted at different user communities. For example, Photoshop is an excellent tool for graphic design experts, and iMovie is an excellent tool for less sophisticated tech-oriented end-users. The problem is that expert-user tools require too much learning and end-user tools are typically too restrictive. Authoring tools and systems are needed that can incorporate editors for different media, perhaps different versions of the editor depending on user experience or application requirements. These tools must work together seamlessly with content acquired from different sources (e.g., uploaded from a capture device, downloaded from an archive, created during the authoring process, derived from other content, etc.). And lastly, the tools must incorporate features to support production of different versions of the title and on-going enhancement and bug fixing.

Tools can make a significant difference. PowerPoint is a good example. Before the development of this end-user tool, expensive custom-designed commercial systems were used to produce 35mm slides. These systems provided expert-user interfaces that ran on specialized computer systems. The introduction of PowerPoint and similar tools in the 1980's made the creation of formal presentation material available to nearly every PC user. The grand challenge is the development of the systems and tools required to support widespread authoring of multimedia content.

This grand challenge may sound like an engineering problem, that is, build a wonderful software package. Our intent is not an engineering solution, rather the challenge to the research community is to develop new user-interface paradigms, software abstractions, media processing algorithms, display presentations and operations for editing media, and media databases that will significantly reduce the effort required to produce high-quality multimedia titles. In other words, make all data types first-class citizens in our applications. To begin with, it might make sense to carefully observe and measure how titles are created today. Research on multimedia authoring will likely include better understanding of media aesthetics, storytelling, and how people relate to a multimedia experience.

The second grand challenge is to *make interactions with remote people and environments nearly the same as interactions with local people and environments*. This grand challenge incorporates two problems: distributed collaboration and interactive, immersive three-dimensional environments. Videophones and videoconferencing have been around for a long time. RCA conducted early experiments with video telephone booths in the 1930's, and A.T.&T. demonstrated a videophone at the 1960 World's Fair in New York. While notable successes have been achieved, for example, small group videoconferencing using H.323 systems (e.g., Polycom), web-based on-line meeting services (e.g., WebEx), person-to-person video chats (e.g., NetMeeting), and webcasting audio or video programming, the telephone is still the dominant medium for remote collaboration. It is still too complicated and expensive to incorporate n-way collaboration using a variety of media into our day-to-day lives. Many problems can be identified including: 1) the difficulty of setting up and operating the equipment, 2) the cost of bandwidth required for high-quality n-way communication is too expensive, 3) the poor support for flexible and scalable multicast services, 4) service limitations (e.g., H.323 does not allow multiple people to view each other and carry on

parallel conversations at the same time nor does it scale to large groups with different meeting interaction styles such as lectures, town meetings, etc.), and 5) collaboration tools, that is tools to encourage working together on a document, viewing results produced by remote equipment (e.g., a telescope or CAT scanner) or meeting to review project schedules, are inadequate. But, the grand challenge is more than solving the distributed collaboration problem. The promise of interactions with remote people, places, and virtual environments, may dramatically change the way we live. New sensors (e.g., touch, smell, taste, motion, etc.) and output devices (e.g., large immersive displays and personal displays integrated with eye glasses) offer the opportunity for more intimate and sensitive interaction with a remote environment. And, continued development of semiconductor technology will bring real-time three-dimensional virtual environments to every computing and communication platform. As one participant said, “interacting with a remote environment should be better than being there.” The grand challenge is to understand the opportunities these new hardware technologies offer and to develop user interfaces and interaction paradigms that allow seamless communication and interactions with remote and virtual environments.

Many research problems are incorporated into this grand challenge including exploring the use of multiple streams of data, whether it be images, sounds, or sensor readings, and developing interaction hardware and software that allow humans to use this data. For example, users interacting with educational or entertainment programs, whether it be live sporting events, lecture webcasts, or broadcast programs want a variety of services that allow them to locate interesting or important events, view program summaries (e.g., a lecture outline or baseball game summary) with links to allow them to watch detailed segments of the program, skim through stored programs rapidly, record material for viewing at a different time (i.e., time shifting), and translations that allow viewing a program on a different platform (e.g., a TV or cellphone). Moreover, it should be easy to create derived works from the content. These services and many more will be required in the future for interactions with remote people and environments.

The third grand challenge is to *make capturing, storing, finding, and using digital media an everyday occurrence in our computing environment*. The widespread adoption of digital cameras and the emergence of cellphones with built-in video cameras are adding to the information glut. Of course, increases in storage capacity and reductions in cost make it possible to store massive amounts of this data. The challenge is to make it useful. Past research has addressed multimedia database models, algorithms to analyze media data, and algorithms to search for relevant or interesting data (e.g., query a music archive by humming a tune or find pictures with similar color palettes). However, significant challenges remain before the multimedia storage and search problems are solved. For example, the following problems cannot be solved today.

1. Search an archive of radio broadcasts to find an interview with a particular individual and a picture archive to find a photo of the person visiting a particular city. Text-to-speech requires context to disambiguate the words being spoken (e.g., technical terms interspersed in a news broadcast are often misunderstood) and identifying where a particular photo was taken might require extensive image analysis or automatic capture of metadata when the photo was taken (e.g., geographic location of the camera at the time the picture was captured). The problem is complicated by the fact that the data in the broadcast archive is not fused with the photo archive.

2. Find lectures by a particular person published on the web. This problem might be solved by looking at the text associated with a streaming media file published on a web page. However, it may be difficult to identify the text associated with a video clip if the web page is generated dynamically. Problems arise too because most commercial webcasting systems use proprietary media coding, storage representations, and network packet formats.
3. Who is that person across the room? The idea is to point your cellphone camera at the person and have it tell you the name of the person. Solving this problem takes context and data fusion as well as connecting to a shared database and a processing server. The obvious solution is to do face matching on the person using the captured image. But, this approach might return too many possible matches or take too much time. What the system should do is use the context of the situation (e.g., a holiday party for a company or a workshop at a conference) to restrict the candidate matches to people who might actually be at the event.
4. Make the billions of hours of home video currently stored in shoeboxes useful. People shoot video but there are no good tools to organize and store it in a form so a user can say, “show me the shot in which Jay ordered Lexi to get the ball.” The solution to this problem may require developing semi-automatic analysis tools coupled with powerful tagging and indexing to organize data so it is easily accessible using unified indexes.

Again, these problems might look like engineering problems because they are drawn from applications that people might actually want to use. However, the grand challenge is to work on the fundamental algorithms (e.g., query planning, parallel search, media-specific search and restriction, combining partial results, unified indexing, and tagging multimedia data) so the problem can be solved sufficiently well that a system could be built and deployed that people will use.

Finally, underlying this last grand challenge is the problem of digital rights management. While this topic is not directly related to multimedia, it will have a dramatic impact on the development and use of content. Discussions at the retreat identified the need for access and propagation rights particularly for fair-use and educational-use rights, the need to track the source of a media asset, and the need for an economic model to pay content owners and creators.

## 5. Discussion

During the retreat several topics were discussed that merit comment. First, people asked whether research on text and images (e.g., a web browser) or research on analyzing or querying a single media (e.g., an image archive) is acceptable multimedia research today. The answer is “no” to text and image research unless it contributed something new to our understanding of combining these media. Ten years ago such research was definitely valid. The sense at the retreat is that we should raise the bar. Music query is a more difficult question. On the one hand, it is a single media, which suggests “no.” Truly innovative research on a single media will always be of interest. But, the contribution must be greater than “here is my nifty algorithm to query still images using color histograms and frequency domain filtering.” Repeatable experiments using published benchmarks are required for the field to progress. Several times during the retreat two ideas were emphasized: 1) compare

new algorithms with previously published algorithms using published benchmarks, and 2) use software developed by other researchers. For example, while some people raised issues about specific details, the development of the video TREC benchmark is viewed as an important contribution to the research community [NIST 2003].

Using software developed by others has been a continuous refrain in the multimedia research community. Everyone recognizes the need to avoid re-implementing known algorithms. The problem is that funding to support a widely adopted common toolkit is too expensive for traditional funding sources. One participant remarked on the investment in Unix by A.T.&T. or Java by Sun. Building an acceptable software toolkit for distributed multimedia applications will likely cost \$20M or more with on going support costs. It is difficult to fund these types of projects without a strong financial incentive for the funding organization. Meanwhile, many researchers build incomplete toolkits and targeted platforms and they spend too much time trying to use a limited commercial toolkit to test one idea. The community collectively spends many times the money required to build an open source portable toolkit re-implementing common components. The lack of software toolkits also causes researchers to forego possible experiments because the effort required to do them is too high.

Several interesting comments were made during the authoring tools discussion. The trade-off between tools for expert versus novice users is well known in the human-computer interface community. Experts want tools that provide more features and finer-grain control over the presentation and behavior of the content being created. However, more features and finer-grain control typically leads to more complex user-interfaces (e.g., they have more options and operations), which makes the authoring tool more difficult for less experienced people to learn and use. On the other hand, a tool designed for inexperienced people might be missing the features required for an expert. The Holy Grail is a system that adapts to user experience, but that remains a challenge for the research community to achieve. The notions of incorporating critics or agents to watch user behavior and either automatically change the object or suggest a change to the user to improve it is an appealing idea. But, more research is needed on these types of interactive environments before they can be made to work.

A common metric used in multimedia applications is QoS. This metric optimizes an internal parameter related to the application (e.g., network delay, CPU cycles dedicated to decompression, etc.). Discussions during the retreat supported the observation that *Quality of Experience* (QoE) is more important than QoS because it relates the user-perceived experience directly rather than the implied impact of QoS. QoE is related to QoS, but it might be a complex function of several parameters, including human perception, rather than easily quantified engineering parameters. The multimedia research community should focus on QoE as the primary metric to be optimized. Research that incorporates the user is more difficult because human behavior is so variable. Nevertheless, the goal of nearly all multimedia applications is to solve a problem for a user, so user perception must be incorporated in an evaluation metric for the algorithm or application.

The last three topics that received considerable attention and generated interesting discussions during the retreat relate to new media, user interfaces, and system configurations. Looking at new media (e.g., haptic, smell, and other sensors) will encourage researchers to think about new ways that humans can interact remotely with people or equipment or with the computer itself. Multimedia research during the past decade has focused on audio and video media. It is time to explore other media.

The windows, mouse, and keyboard human computer interface using a two-dimensional output display has been the standard UI for the past twenty-five years. Many new interface devices and metaphors have been explored during this time including three-dimensional output displays and pen, speech, gesture, and tangible input interfaces. An important theme is multi-modal interfaces that allow the user to interact with the system using several media (e.g., pen and speech). The user should use different devices for an operation (e.g., gesture or mouse) depending on the situation. The multimedia research community should participate in these research activities, even though they are primarily UI problems, because the multimedia community understands the underlying time-oriented media and multimedia processing.

Lastly, many participants discussed the ubiquitous computing metaphor for human-computer interaction. The idea is that many sensors and smart devices with embedded computers will be present in our environment either carried by the user or permanently located in the space. Applications should be written to exploit this collection of devices. They should adapt to the availability of equipment and processing to solve a user's problem. Distributed multimedia is inherent in this new world.

The past decade has seen significant progress in multimedia research. Now is the time to raise expectations for the future. The focus should be on incorporating new media and devices and exploiting multiple media to create applications that solve an important problem and produce high quality user experiences.

## **6. References**

- [NIST 2003] Digital Video Retrieval at NIST, <http://www-nlpir.nist.gov/projects/trecvid>, 2003.
- [SIGMM 2003] ACM SIG Multimedia Strategic Retreat Participant Position Papers, <http://www.acm.org/sigmm/>, December 2003.

## Keynote Speakers

### Sensitivity Analysis: Unexpected Outcomes in Art and Engineering



[Ken Goldberg](#)

Professor of IEOR and EECS at UC Berkeley. His research addresses robot manipulation, geometric algorithms for automation, and networked robots. More information and online projects are linked from <http://goldberg.berkeley.edu>.

#### Abstract:

Contemporary art and engineering research are both at their best when things don't turn out as planned. I'll present selected examples based on artworks developed with students and other collaborators involving robots and networks over the past 20 years. These projects set out to investigate intersections of technology and nature, such as the Telegarden, a robot installation that allowed online participants to remotely tend a living garden; Ballet Mori, a classical dance performed to sounds triggered by live seismic data; and Demonstrate, where an ultra high-resolution video camera raised eyebrows at the 40th anniversary of the Free Speech Movement. Every project led to unexpected twists and complications.

I'll also argue that the languages of contemporary art and engineering research are complex, dynamic, and often frustratingly impenetrable to outsiders. In art, a blue disk can be a cliché, or, in the right place at the right time, profound. In engineering, analogous contexts determine the beauty of a coordinate frame or mathematical equation. In both spheres, aesthetic interpretation is based on knowledge of prior art and contemporary dialogues. Being so similar, it is not surprising that unexpected forces arise when these two spheres are brought together.

#### Bio:

Ken Goldberg is an artist and professor of engineering at UC Berkeley. Goldberg's art installations such as the Telegarden have been exhibited at the Whitney Biennial, Venice Biennale, Pompidou Center (Paris), Walker Art Center, Ars Electronica (Linz Austria), ZKM (Karlsruhe), ICC Biennale (Tokyo), Kwangju Biennale (Seoul), Artists Space, and The Kitchen (New York). He has held visiting positions at San Francisco Art Institute, MIT Media Lab, and Pasadena Art Center. The Tribe, a short film he co-wrote, was selected for the Sundance and Tribeca Film Festivals. Ballet Mori, a multi-media project he developed to commemorate the 1906 Earthquake, was performed by the SF Ballet at the San Francisco Opera House.

Goldberg is an IEEE Fellow and Vice President of Technical Activities for the Robotics and Automation Society. His PhD is in Computer Science from Carnegie Mellon University and he has published over 150 research papers on robotics, automation, and geometric algorithms. He is editor of several books, including *The Robot in the Garden: Telerobotics and Telepistemology in the Age of the Internet* (MIT Press, 2000). Goldberg is Founding Director of Berkeley's Art, Technology, and Culture Colloquium.

Goldberg was awarded the National Science Foundation Young Investigator Award in 1994, the NSF Presidential Faculty Fellowship in 1995, the Joseph Engelberger Robotics Award in 2000, the IEEE Major Educational Innovation Award in 2001.

#### Implicit Participation

**Bradley Horowitz**



VP, Yahoo! Product Strategy Group

**Abstract:**

A new generation of simple, affordable and easy to use tools has led to what has been called the democratization of publishing . Anyone with a camera has become a photographer , with a keyboard an author , a microphone a podcaster , etc. As the means for production and distribution of content have become readily accessible, the most valuable inelastic commodity has become attention. In this talk we will describe both how Yahoo has been working to lower barriers to participation and turn every consumer into a creator , but also how we are using data mining techniques to help identify and leverage high-value content. Yahoo! is creating systems that engender mass participation but also that allow the cream to rise by ensuring that every user of the system creates value in his or her wake. Examples from widely available real-world products (such as Flickr, Yahoo! Answers, del.icio.us, upcoming.org, etc.) will be used as illustrations.

**Bio:**

Bradley Horowitz, vice president of Yahoo!'s product strategy group, leads Yahoo!'s efforts in building innovative products and technologies across the company. Horowitz is driving innovation and leveraging Yahoo!'s platform to deliver compelling Yahoo! products and services to a community of 500 million users. In addition, he is responsible for the company's initiative to open up its platform which includes overseeing the Yahoo! Developer Network (Y!DN). Previously, he managed a portfolio of products for Yahoo! including media search, desktop search and the Yahoo! Toolbar.

Prior to joining Yahoo!, Horowitz served as both the chief technical officer and the vice president of engineering for the Virage division of Autonomy, where he was responsible for the technical delivery of five major product lines. Prior to Autonomy, he founded Virage, the company widely recognized as the market creator and leader for advanced media indexing and analysis. Horowitz helped grow the company from "a garage startup" through its NASDAQ IPO.

Horowitz was a PhD candidate at the MIT Media Lab. While at the Media Lab, he worked on a number of topics related to computer vision, graphics and image processing, which resulted in a patented new technique for the recovery of structure, motion and camera parameters from video sequences.

Horowitz holds an MS in Media Science from MIT and a BS in Computer Science from the University of Michigan.

## Typical Easel for Posters

Not all of the easels will be the same, but below is a picture of one type we'll be using. It fits two of the AO sized posters, but we hope to use only one poster per side.



## Oral Presentations

The oral presentations at ACM MM 2006 will be scheduled to start on the half-hour (at 8:30am, 9:00am, 9:30am, etc.) so that the audience may freely move between parallel sessions and see complete talks. Session chairs should enforce this, and thus make sure talks do not go overtime or start too soon. Since there will be 30 minutes for everything - the talk, questions and answers, and changing to the next speaker - the talk should run **22-24 minutes** (no longer) to leave adequate time for questions and changeover. Session chairs will give speakers 5-minute and 1-minute notices, and they will feel free to stop the talk at the 24 minute mark.

For all oral presentations, a video projector (for PC or Mac laptop) and audio will be provided. Bring your own laptop, and have it ready to go when you step up to give your presentation (that is not the time to boot the computer, look for the file, etc.!).

Presenters, please find the session chair and introduce yourself to him or her before the session begins. You may want to tell the session chair how to pronounce your name, and ask how he/she will give the 5-minute and 1-minute warnings.

Session chairs, please be at the session early to check on A/V and to meet the speakers. Find out who is giving the talk for each paper, and how to pronounce the speaker's name. Make sure they understand about the 5-minute and 1-minute warnings, and how you will deliver these.

## Presentation Guidelines: Do's and Don'ts

This information is reprinted in part from the ACM CHI and IUI conference series.

### 1. Checking Content Appropriateness

DON'T give a presentation that will be comprehensible and interesting only to people who work in the same area as you. Please be aware that MM is a multidisciplinary conference, with researchers and practitioners in attendance.

DO ensure that even people who have little familiarity with your subarea of MM can understand at least the main points:

- what questions you addressed,
- why they're important,
- what methods you used (not necessarily the details),
- what your main results were, and
- why they are interesting.

In fact, even the experts in your area don't need to understand more than these points; for the rest, they can read the paper.

DON'T try to squeeze in so much material as to leave hardly any time for questions.

DO aim to be finished in your allotted time:

Your audience will love and admire you for it, and you will be rewarded with a relatively deep discussion. (If the discussion flags, the Session Chair or the program chairs will ask a question to get it going again.) Note that any speaker who exceeds 15 or 20 minutes will be interrupted mercilessly by the Session Chair, and time for questions will be reduced accordingly.

DON'T subject your audience to an "ordeal by bulleted list." Bulleted lists especially those with large amounts of text should be used only in exceptional cases. They are generally boring, abstract, unconvincing, and hard to read while the speaker is talking.

DO present a series of "exhibits": images, videos, system demos, diagrams, graphs, or tables. You can explain and elaborate on these exhibits while people are looking at them, but in general you don't need to write what you say on the slides. Anyone who wants to see the points you made in black and white can read your paper. Carefully preparing an exhibit can take at least 10 times as long as dashing off a bulleted list, but your audience and your research, case study, or discussion deserve nothing less.

DON'T use text smaller than a 28-point font. Your audience will not be able to read your slides otherwise.

DO use text sparingly: Keep your points in short, concise, outline form. This will inform the viewer about the topic, and will also help you remember your key points for discussion. There is no real need to write in full sentences, as this will unnecessarily clutter your slides. If you have that much text on the screen, break part of it out to another slide.

## **2. Polishing the Details**

DON'T put material on a slide that only the people in the front rows can read.

DO pay special attention to types of material that often turn out to be illegible: screen shots and complex graphics. If an exhibit like this can't be shown legibly as a whole, find a way to zoom in on individual parts of it as they are discussed.

DON'T clutter each slide with distracting logos and superfluous information such as the title of the talk or the name and date of the conference.

DO present only material that helps you to convey your points effectively. If you must include your institution's logo on each slide, make sure that it is not the most conspicuous and interesting element on any slide.

## **3. Giving the Presentation**

DON'T risk fumbling desperately with the laptop at the beginning of your talk.

DO arrive 20 minutes before your session to test the compatibility of your laptop with the projector. If you bring your presentation on a CD or memory stick to present on someone else's laptop, do everything possible to maximize its portability, and test the presentation at the earliest opportunity, leaving plenty of time to fix any problems (e.g., replacing missing fonts).

DON'T talk in such a way that only a fraction of the listeners can understand you.

DO keep in mind the people in the back row who are not especially experienced in listening to English-language presentations. Native speakers of English need to avoid speaking too fast or colloquially; non-native speakers should enunciate especially clearly so that any foreign accent does not impair comprehension.

DO use your microphone, even if there are not many attendees in your session. Session rooms are still enormous, and you will be on a stage. Remember that the use of a microphone does not in itself guarantee that people in the back can hear you easily: Speak up in a lively manner!

DON'T ignore your Session Chair's time warnings.

DO pay attention to the Session Chair's countdown cards. You will receive warnings at five minutes prior, one minute prior, and when time is up. If you do not stop when time is called, your Session Chair will come to the stage to start the Q&A session.

DON'T rush to cover your remaining content if you are running out of time.

DO rehearse your presentation before attending MM, and if you find yourself with a lot more content to cover at the 5-minute mark, resist the temptation to speak faster to finish. Your audience will not remember that much material!

#### **4. Answering Questions**

DON'T end your presentation with a slide that contains only an uninformative text like "Any questions?"

DO conclude with a slide summarizing your main contributions, leaving it on the screen except when a question requires you to switch to another slide. (This is one of those rare cases where a bulleted list may be appropriate.) This slide will help people to think of important questions to ask - and also to be impressed by your achievements.

DON'T use a question from the audience as a springboard to leap into the five minutes of your talk that you had to leave out because of the time limit.

DO answer each question directly and concisely, without digressing into related topics. Give others a chance to ask their questions as well.

## Detailed instructions for preparing camera-ready papers

[Note: Each accepted paper must have at least one author **registered** for the conference **at the regular full rate (not the student rate)** by the advance registration date (September 23) for the paper to remain part of the ACM MM conference and proceedings. If **all** of the paper's authors are students, then a student registration suffices for that paper.]

The website for submitting the final version of the accepted papers is now open. Please go to the relevant links below to access the **copyright form** and **submission web site**:

The instructions are organized as follows:

- [Overview of deadlines and maximum number of pages](#)
- [Camera ready copy and the copyright form](#)
- [Formatting instructions and uploading camera ready copy](#)
- [The ACM computing classification](#)
- [The Shepherd process](#) (full papers only)
- [Printed proceedings and ACM digital library](#)
- [Contact the proceedings coordinator](#), Lisa M. Tolles at Sheridan Printing
- [Contact the MM05 Proceedings Chair](#), Roger Zimmermann

### Overview of the page limits and deadlines

**Note: These deadlines (Aug 1 and Aug 14) must be met - no extensions!**

Paper type	Max. number of pages (no exceptions)	Deadline
Full papers	10	August 1, 2006
Short papers	4	August 1, 2006
Demo papers	2	August 1, 2006
Panel descriptions	2	August 1, 2006
Art full papers	10	August 1, 2006
Art short papers	4	August 1, 2006
Art exhibit papers	2	August 1, 2006
Foundational sciences	10	August 1, 2006
Video papers	2	August 1, 2006
Doctoral symposium papers	3	August 1, 2006
Open source contest	4	August 1, 2006
MIR Workshop	10	August 14, 2006
VSSN Workshop	10	August 14, 2006
MSC Workshop	10	August 14, 2006
CARPE Workshop	10	August 14, 2006
HCM Workshop	10	August 14, 2006
MCPS Workshop	10	August 14, 2006
AMCMM Workshop	10	August 14, 2006

### Overview of camera ready and copyright form

Authors of accepted full papers/short papers/demos/videos/art exhibit/doctoral symposium for the ACM MM 2006 conference will receive an email from Lisa M. Tolles from Sheridan Printing, the publisher that will produce the conference proceedings. The email contains all the details for submitting the camera ready copy and copyright form.

### Formatting instructions and uploading camera ready copy

The camera ready copy of each submission must strictly adhere to the Sheridan Printing Instructions for the

ACM MM 2006 Proceedings.

We kindly ask you to read carefully through the [instructions from Sheridan](#) and start preparing the camera ready copy as early as possible to avoid any delays due to formatting problems, paper length, or lacking a copyright form. For an appealing print of your contribution in the conference proceedings please do NOT change the given paper style, e.g., by increasing text width and height or reducing line spacing in order to get more space for your paper.

The Sheridan web page also contains information on the files and their required naming for the submission as well as the URL where to upload the files of the final version.

If you have any questions left regarding the preparation of the camera ready copy, please contact the proceedings coordinator Lisa M. Tolles at Sheridan printing. For questions regarding the shepherding process and further general questions, please contact the proceedings chair Roger Zimmermann.

### **The ACM computing classification**

Please note that each paper must provide a classification of it's content and general terms according to the [ACM computing classification](#). If you are not familiar with this classification please make sure that you read carefully through the guide "How to classify works using the ACM's computing classification system" to select the most suitable categories and general terms for your paper and add some additional keywords to it. Do NOT omit classification, general terms or keywords in your camera ready copy, e.g., in order to save space!

### **The shepherding process (for Full Papers only)**

As was started in ACM 2003, the inclusion of the reviewers' comments in the final version of the paper is guided by a "Shepherd" - a member of the program committee - to help the authors in the process of reflecting the reviewers' comments in the final version. The shepherd will ensure that all review comments are addressed properly in the final version. This process contributes to the high quality of the conference and also honors the reviewers' effort in writing reviews. Your shepherd should already be in touch with you; please negotiate a time schedule with your shepherd.

The proceedings chair must get a confirmation from the shepherds that all review comments are appropriately considered in your final camera-ready version. Please do not hesitate to ask your shepherd if you have any questions. In order to facilitate the task of the shepherd, we ask you to establish a document that lists all review comments and discusses how you have taken them into account in your final paper.

Email this document to the proceedings chair Roger Zimmermann before August 1, 2006. Please include [ACM MM06] in the subject of the e-mail. This information, acknowledged by your shepherds as complete, is important for being included in the conference proceedings.

### **Printed proceedings and ACM digital library**

Printed conference proceedings will be produced by Sheridan Printing; the proceedings will also be entered in the ACM Digital Library.

### **Contact the publisher**

Lisa Tolles

Email: [Lisatolles@sheridanprinting.com](mailto:Lisatolles@sheridanprinting.com)

Phone +1-908-213-8988

Fax +1-908-454-2554

Address

Sheridan Printing Co.

Attn: Lisa M. Tolles

1425 Third Ave.

Alpha, NJ 08865 (USA)

### **Contact proceedings chair**

In case you have any questions left after having carefully read [the instructions provided by Sheridan](#)

**Printing**, please [contact the proceedings chair](#):

[Through August 3rd, contact **both** of the following:]

Roger Zimmermann, University of Southern California

Email: [rzimmerm@imsc.usc.edu](mailto:rzimmerm@imsc.usc.edu)

Phone +1-213-740-7654

Fax +1-213-740-5807

Min Qin, University of Southern California

Email: [mqin@usc.edu](mailto:mqin@usc.edu)

Phone +1-213-740-4177

Fax +1-213-740-5807

# Student Conference Participation Grant

The ACM Multimedia conference is the main annual conference of the ACM SIGMM. It provides a great opportunity for researchers to present their work, get acquainted with the work of others, and discuss with colleagues both junior and senior. To promote this form of cooperation, the SIG and its sponsors make conference participation grants available for several students.

A grant is meant to support personal attendance of one student at the main conference by paying the **advance registration fee for students**. The grant will not cover fees for workshops or tutorials. Individual grants may cover additional expenses subject to conditions of individual sponsors. Students will have to make their own arrangements for travel and accommodation.

All students can apply, but regional and institutional diversity will be considered in the selection process.

To apply for a grant, a **faculty member** of the student's host institution must send an email to [griff@simula.no](mailto:griff@simula.no) with the subject "ACM MM Student Conference Participation Grant."

The application must be sent in ASCII in the email body no later than **August 14, 2006**. Applications arriving after the deadline will not be considered. The notification date is ~~August 26, 2006~~ **August 30, 2006** (coming soon!). Decisions made are final.

Please include the following information in an application:

1. the student's full name and email address
2. the name and nationality of the institution
3. the student's academic status and expected year of graduation
4. whether the student is (co-)author of a paper that was accepted for the main conference or if the student contributes to any other part of the conference
5. if the student is a (co-)author, the paper title, author list, the student's contribution to the paper, and whether the student will present the paper
6. the reason for applying for the grant, including the expected benefits that the students will have from attending the conference and including the funding status of the student
7. a brief statement by the student about his or her current research interests and research accomplishments to date

We look forward to receiving your applications!

# Detailed Conference Schedule

[Further information is available on the [tutorials](#), [workshops](#), and the [Interactive Arts Program](#).]

Note: The "Brave New Topics 2" session was rescheduled - it is now on Wednesday 2:00-3:30pm (and now called "Brave New Topics 1").

## Monday, October 23, 2006

**Noon-6:00pm: Interactive Arts Program Exhibition (at UCSB)**

**Full-day Tutorial (8:30am-5:30pm)**

Peer-to-Peer Multimedia Applications  
*Jin Li*

**Half-day Tutorials: Morning (8:30am-noon)**

Interactive Digital Television and Multimedia Systems  
*Pablo Cesar and Konstantinos Chorianopoulos*

Flexible Modeling and Performance Debugging of Real-Time Embedded Multimedia Systems  
*Samarjit Chakraborty*

Computer Audition: An introduction and research survey  
*Shlomo Dubnov*

Recent developments in video compression standards and their impact on embedded platforms: from scalable to multi-view video coding  
*Iole Moccagatta*

**Half-day Tutorials: Afternoon (2:00pm-5:30pm)**

Data Mining and Information Retrieval in Time Series/Multimedia Databases  
*Eamonn Keogh*

Multimedia Content Protection  
*Dulce Ponceleon, Julian Cerruti*

Semantic Indexing and Retrieval of Video  
*Marcel Worring, Cees Snoek*

## Tuesday, October 24, 2006

**8:30-10:00am Conference opening remarks and Keynote presentation**  
**Session Chair:** Matthew Turk, UC Santa Barbara

Sensitivity Analysis: Unexpected Outcomes in Art and Engineering  
*Ken Goldberg, UC Berkeley*

**10:00-10:30am Coffee break** (sponsored by FXPAL)

**10:30am-noon Best Papers Session**

**Session Chair:** Wolfgang Klas, Univ. of Vienna

An Innovative Three-Dimensional User Interface for Exploring Music Collections Enriched with Meta-Information from the Web

*Peter Knees, Markus Schedl, Tim Pohle, Gerhard Widmer*

Tiling Slideshow

*Jun-Cheng Chen, Wei-Ta Chu, Jin-Hau Kuo, Chung-Yi Weng, Ja-Ling Wu*

Video Search Reranking via Information Bottleneck Principle

*Winston Hsu, Lyndon Kennedy, Shih-Fu Chang*

**Noon-6:00pm: Interactive Arts Program Exhibition (at UCSB)**

Shuttle vans available approximately hourly between hotel and UCSB

**Noon-2:00pm Conference Lunch** (sponsored by Microsoft Research)

Plaza Del Sol (outdoors plaza at the hotel)

**Noon-4:00pm Short Papers Poster Session 1**

Authors at posters 1:00-2:00pm and 3:30-4:00pm

**Session Chair:** Brian Bailey, University of Illinois at Urbana-Champaign

Maximum Unfolded Embedding: Formulation, Solution and Application for Image Clustering

*Huan Wang*

An Efficient Approach to Generic Multimedia Adaptation

*Joseph Thomas-Kerr*

Atomic Topical Segments Detection for Instructional Videos

*Ying Li*

Building Concept Ontology for Medical Video Annotation

*Jianping Fan*

Real-time Automatic 3D Scene Generation from Natural Language Voice and Text Descriptions

*Lee Seversky*

Training Combination Strategy of Multi-stream Fused Hidden Markov Model for Audio-visual Affect Recognition

*Zhihong Zeng*

Fourth Frame Forums: Interactive Comics for Collaborative Learning

*Andrew Gordon*

Improving the Experience of Controlling Avatars in Camera-Based Games Using Physical Input

*Na Li*

Effective and Efficient Object-based Image Retrieval Using Visual Phrases

*Qing-Fang Zheng Zheng*

Music Emotion Classification: A Fuzzy Approach

*Yi-Hsuan Yang*

Shape from Regularities for Interactive 3D Reconstruction of Piecewise Planar Objects from Single Images

*Zhenguo Li*

To Construct Optimal Training Set For Video Annotation

*Jinhui Tang*

Low Complexity Controllable Scrambler/Descrambler for H.264/AVC in Compressed Domain

*Ho-Jae Lee*

Automatic Function selection for Large Scale Salient Object Detection

*Yuli Gao*

Tracking Users through a Projection Screen

*Ismo Rakkolainen*

3D Object Retrieval with Easy 3D Query Generation and Graph Based Relevance Feedback

*Liangliang Cao*

Does Ontology Help in Image Retrieval - A Comparison between Keyword, Text Ontology and Multi-Modality Ontology Approaches

*Huan Wang*

Automatic Document Orientation Detection and Categorization through Document Vectorization

*Shijian Lu*

Distortion-Aware Video Communication with Pipeline Forwarding

*Mario Baldi*

3D Model Retrieval Based on Volumetric Extended Gaussian Image and Hierarchical Self Organizing Map

*Jiqi Zhang*

On the Significance of Cluster-Temporal Browsing for Generic Video Retrieval - a Statistical Analysis

*Mika Rautiainen*

Text Segmentation Based on Stroke Filter

*Cheolkon Jung*

Video Inpainting and Implant via Diversified Temporal Continuations

*Timothy Shih*

Synchronization of Multiple Video Recordings Based on Still Camera Flashes

*Prarthana Shrestha*

GVU-PROCAMS: Enabling Novel Projected Interfaces

*Jay Summet*

An Unsupervised Method for Clustering Images Based on their Salient Regions of Interest

*Gustavo Borba*

Quaternion Image Watermarking using the Spatio-Chromatic Fourier Coefficients Analysis

*Tsz Kin Tsui*

Supporting Guaranteed Continuous Media Streaming in Mobile Ad-hoc Networks with Link Availability Prediction

*Min Qin*

Transductive Inference using Multiple Experts for Brushwork Annotation in Paintings Domain  
*Yelizaveta Marchenko*

Hybrid Watermarking for Improving Detector Performance  
*Won-gyum Kim*

Design and Implementation of a Multi-stream CableCARD with a High-Speed DVB-Common Descrambler  
*Joon-Young Jung*

News Video Search with Fuzzy Event Clustering using High-level Features  
*Shi-Yong Neo*

Scalable Relevance Feedback Using Click-Through Data for Web Image Retrieval  
*En Cheng*

**Noon-4:00pm Arts Short Papers Poster Session 1**

Authors at posters 1:00-2:00pm and 3:30-4:00pm

**Session Chair: Marko Peljhan**

Cyborglogging with Camera Phones: Steps Toward Equiveillance  
*Steve Mann, Raymond Lo, James Fung*

The Andantephone: Teaching Music by Walking on Patio Stones with Sensors Activated by Sequential Foot Steps  
*Steve Mann*

Interactive Immaterial Screen for Performing Arts  
*Ismo Rakkolainen*

Interactive Composition, Performance and Music Generation  
*Paolo Bottoni, Stefano Faralli, Anna Labella, Mario Pierro, Claudio Scozzafava*

Flow An Interactive AJAX-based Internet Information Requesting System  
*Yu-Chuan Tseng, Chia-Hsiang Lee*

Talk2Me: The Art of Augmenting Conversations  
*Ann Morrison, Peta Mitchell, Ralf Muhlberger*

**1:30-2:00pm Arts Exhibit Overview**

**2:00-3:30pm Content 1: Multi-Modal Analysis**

**Session Chair:** Marcel Worring, Univ. of Amsterdam

Segmentation, Categorization, and Identification of Commercial Clips from TV Streams Using Multi-modal Analysis  
*Ling-Yu Duan, Jinqiao Wang, Yantao Zheng, Jesse Jin, Hanqing Lu, Changsheng Xu*

Multimodel Fusion using Learned Text Concepts for Image Categorization  
*Qiang Zhu, Mei-Chen Yeh, Kwang-Ting Cheng*

Live Sports Event Detection Based on Broadcast Video and Web-casting Text  
*Changsheng Xu, Jinjun Wang, Yiqun Li, Kongwah Wan, Ling-Yu Duan*

**2:00-3:30pm Applications 1: Media Presentation**

**Session Chair:** Gopal Pingali, IBM Research

Multimedia Thumbnails for Documents  
*Berna Erol, Kathrin Berkner, Siddharth Joshi*

Video Retargeting: Automating Pan and Scan  
*Feng Liu, Michael Gleicher*

Progressive Cut  
*Chao Wang, Qiong Yang, Mo Chen, Zhongfu Ye, Xiaoou Tang*

**2:00-3:30pm Arts Session 1: Installations and Media Archaeology**

**Session Chair:** Alejandro Jaimes

Buzz: Telling Compelling Stories  
*Sara Owsley, Kristian Hammond, David Shamma, Sanjay Sood*

Archeology of Multimedia  
*Fabrizio Nunnari, Vincenzo Lombardo, Andrea Valle, Francesco Giordana, Andrea Arghinenti*

Instant Archaeologies: Digital Lenses to Probe and to Perforate the Urban Fabric  
*Petra Gemeinboeck, Atau Tanaka, Andy Dong*

**3:30-4:00pm Coffee break** (sponsored by FXPAL)

**4:00-6:00pm Content 2: Machine Learning in Multimedia**

**Session Chair:** Mohan Kankanhalli, Univ. of Singapore

Content-based Image Retrieval using More Expressive Salient Point Representations  
*Hui Zhang, Rouhollah Rahmani, Sharath Cholleti, Sally Goldman*

Learning Image Manifolds by Semantic Subspace Projection  
*Jie Yu, Qi Tian*

Learning from Facial Aging Patterns for Automatic Age Estimation  
*Xin Geng, Zhi-Hua Zhou, Yu Zhang, Gang Li, Honghua Dai*

Efficient Top-k Hyperplane Query Processing for Multimedia Information Retrieval  
*Navneet Panda, Edward Chang*

**4:00-6:00pm Systems 1 Streaming**

**Session Chair:** Surendar Chandra, Notre Dame University

Tavarua: Video Streaming with WWAN Striping  
*Asfandyar Qureshi, Jennifer Carlisle, John Guttag*

Scalable Streaming for Heterogeneous Clients  
*Liqi Lishi, Phillipa Sessini, Anirban Mahanti, Zongpeng Li, Derek Eager*

Towards Scalable Delivery of Video Streaming to Heterogeneous Receivers  
*Bashar Qudah, Nabil Sarhan*

Scalable and Adaptive Streaming for Non-Linear Media  
*David Gotz*

**4:00-5:30pm Applications 2: Searching Media I**

**Session Chair:** Jim Gemmell, Microsoft Research

EnjoyPhoto - A Vertical Image Search Engine for Enjoying High-Quality Photos  
*Lei Zhang, Le Chen, Feng Jing, Kefeng Deng, Wei-Ying Ma*

IGroup: Web Image Search Results Clustering  
*Feng Jing, Changhu Wang, Yuhuan Yao, Kefeng Deng, Lei Zhang, Wei-Ying Ma*

Extreme Video Retrieval: Joint Maximization of Human and Computer Performance  
*Alexander Hauptmann, Wei-Hao Lin, Rong Yan, Jun Yang, Mingyu Chen*

**7:00-9:00pm Conference Reception** (sponsored by the University of California Institute for Research in the Arts)

**Location: California Nanosystems Institute at UCSB**

Bus transportation will be provided

## Wednesday, October 25, 2006

**8:30-10:00am Applications 3: Entertainment & Home Environments**

**Session Chair:** Dick Bulterman, CWI

Presence and Portrayal: Video for Casual Home Dialogues  
*David Chatting, Josie Galpin, Judith Donath*

Edge Indexing in a Grid for Highly Dynamic Virtual Environments  
*Beomjoo Seo, Roger Zimmermann*

Exploring Composite Acoustic Features for Efficient Music Similarity Query  
*Jialie Shen, Bin Cui, Gao Cong, Heng Tao Shen, Cui Yu*

**8:30-10:00am Content 3: Semantic Concepts**

**Session Chair:** Alan Smeaton, Dublin City Univ.

The Challenge Problem for Automated Detection of 101 Semantic Concepts in Multimedia  
*Cees Snoek, Marcel Worring, Jan van Gemert, Jan-Mark Geusebroek, Arnold Smeulders*

Player Action Recognition in Broadcast Tennis Video with Applications to Semantic Analysis of Sports Game  
*Guangyu Zhu, Changsheng Xu, Qingming Huang, Wen Gao, Liyuan Xing*

Learning Concepts from Large Scale Imbalanced Data Sets Using Support Cluster Machines  
*Jinhui Yuan, Jianmin Li, Bo Zhang*

**8:30-10:00am Arts Session 2: Interactive Spaces and Performance**

**Session Chair:** George Legrady, UC Santa Barbara

Choreographic Buttons: Promoting Social Interaction through Human Movement and Clear Affordances  
*Andrew Webb, Andruid Kerne, Eunyee Koh, Pranesh Joshi, YoungJoo Park, Ross Graeber*

Motion Swarms: Video Interaction for Art in Complex Environments  
*Quoc Nguyen, Jeffrey Boyd, Christian Jacob, Gerald Hushlak*

Movement-based Interactive Dance Performance  
*Jodi James, Todd Ingalls, Gang Qian, Daniel Whiteley, Loren Olson, Siew Wong, Thanassis Rikakis*

**8:30am-noon Demo Session 1**

**Session Chair:** Baochun Li, Univ. of Toronto

Eye/gaze Tracking in Web, Image and Video Documents

*Chabane Djeraba, Stanislas Lew, Dan Simovici*

Geelix.com: Sharing Gaming Experiences

*Ole-Ivar Holthe, Leif Arne Rønningen*

Fotofiti: Web Service for Photo Management

*Benjamin Lee, Wen-Yen Chen, Edward Y. Chang*

Multimedia Thumbnails for Documents: Implementation and Demonstration

*Berna Erol, Kathrin Berkner, Siddharth Joshi*

Blazingly Fast Image Copyright Enforcement

*Herwig Lejsek, Fridrik H. Smundsson, Björn Jönsson, Laurent Amsaleg*

MAGICAL Demonstration: System for Metadata Automated Generation for Instructional Content

*Chitra Dorai, Robert Farrell, Amy Katriel, Galina Kofman, Ying Li, Youngja Park*

NN<sup>k</sup> Networks and Automated Annotation for Browsing Large Image Collections from the World Wide Web

*Daniel Heesch, Alexei Yavlinsky, Stefan Rieger*

Enabling Secure Distribution of Digital Media to SD-Cards

*Dulce Ponceleon*

IGroup: A Web Image Search Engine with Semantic Clustering of Search Results

*Feng Jing, Changhu Wang, Yuhuan Yao, Kefeng Deng, Lei Zhang, Wei-Ying Ma*

Event-Centric Multimedia Data Management for Reconnaissance Mission Analysis and Reporting

*Utz Westermann, Srikanth Agaram, Bo Gong, Ramesh Jain*

A Real-Time, Multimodal Biofeedback System for Stroke Patient Rehabilitation

*Yinpeng Chen, Weiwei Xu, Richard Isaac Wallis, Hari Sundaram, Thanassis Rikakis, Todd Ingalls, Loren Olson, Jiping He*

3WNews: Who, Where, and When in News Video

*Jun Yang, Alexander G. Hauptmann*

3D TV Using MPEG-2 and H.264 View Coding and Autostereoscopic Displays

*Lakis Christodoulou, Liam M. Mayron, Hari Kalva, Oge Marques, Borko Furht*

PartyPeer: a P2P based Massively Multiplayer Online Game

*Leslie S. Liu, Roger Zimmermann, Baoxuan Xiao, Jon Christen*

**10:00-10:30am Coffee break** (sponsored by IBM)

**10:30am-noon Arts Session 3: Tools for Creativity and Art Analysis**

**Session Chair:** Lonce Wyse

Handheld Electronic Camera Flash Lamp as a Tangible User-Interface for Creating Expressive Visual Art Works.

*Steve Mann, Corey Manders*

Hydraulophone Design Considerations: Absence, Displacement, and Velocity-Sensitive Music Keyboard in which Each Key is a Water Jet

*Steve Mann, Ryan Janzen, Mark Post*

Semi-supervised Annotation of Brushwork in Paintings Domain using Serial Combinations of Multiple Experts

*Yelizaveta Marchenko, Tat-Seng Chua, Ramesh Jain*

### **10:30am-noon Systems 2: Distributed Systems**

**Session Chair:** David Gotz, IBM

Fundamental Scaling Laws of Peer-to-Peer Live Multimedia Streaming

*Tara Small, Ben Liang, Baochun Li*

DANS: Decentralized, Autonomous, and Network-wide Service Delivery and Multimedia Workflow Processing

*G. Kwon, Kasim Candan*

Design and Implementation of an Adaptive Distributed On-line Video Correlation System

*Xiaohui Gu, Zhen Wen, Ching-Yung Lin, Philip Yu*

### **10:30am-noon Applications 4: Searching Media II**

**Session Chair:** Qi Tian, UT San Antonio

Efficient Benchmarking of Content-based Image Retrieval via Resampling

*Jialie Shen, Shepherd John*

Virtual Observers in a Mobile Surveillance System

*Stewart Greenhill, Svetha Venkatesh*

Scalability of Local Descriptors: A Comparative Study

*Herwig Lejsek, Fridrik Asmundsson, Björn Jönsson, Laurent Amsaleg*

### **Noon-6:00pm: Interactive Arts Program Exhibition (at UCSB)**

Shuttle vans available approximately hourly between hotel and UCSB

### **Noon-2:00pm Conference Lunch (sponsored by Google) and SIGMM Meeting**

SIGMM meeting participants: Reagan Room

Others: Plaza Del Sol

### **Noon-4:00pm Short Papers Poster Session 2**

Authors at posters 1:00-2:00pm and 3:30-4:00pm

**Session Chair:** Belle Tseng, NEC Labs

VirtualTour: An Online Travel Assistant Based on High Quality Images

*Feng Jing*

Audio Similarity Measure by Graph Modeling and Matching

*Yuxin Peng*

Image Annotation by Large-Scale Content-based Image Retrieval

*Xirong Li*

Mixed-Initiative Multimedia for Mobile Devices: Unified Design of User-Driven & System-Prompted Modalities

*Jeannie Lee*

PhotoArcs: Ludic Tools for Sharing Photographs

*Morgan Ames*

Inserting 3D Projected Virtual Content into Broadcast Tennis Video

*Xinguo Yu*

Manifold-Ranking Based Video Concept Detection on Large Database and Feature Pool

*Xun Yuan*

Towards Content-Based Relevance Ranking for Video Search

*Wei Lai*

Animation Movies Trailer Computation

*Patrick Lambert*

A Trustworthy End-to-end Key Management Scheme for Digital Rights Management

*Yeonjeong Jeong*

Detecting Irregularity in Videos Using Kernel Estimation and K-D Trees

*Yun Li*

User-friendly H.264/AVC for Real-Time Browsing

*Pengpeng Ni*

Image Annotation Refinement using Random Walk with Restarts

*Changhu Wang*

An Architecture for Viewer-Side Enrichment of TV Content

*Dick Bulterman*

Interactive Mosaic Generation for Video Navigation

*Kihwan Kim*

Syllabic Level Automatic Synchronization of Music Signals and Text Lyrics

*Denny Iskandar*

Automatic Detection of Player's Identity in Soccer Videos Using Faces and Text Cues

*Marco Bertini*

The Differential Structure of Sub Pixels Interpolated from Integer Pixels Using N-Tab FIR Filters for High Definition H.264 Video Encoding

*Jinwuk Seok*

Leveraging Community Annotations for Image Adaptation to Small Presentation Formats

*Patrick Schmitz*

Interactive Audio-Visual Video Browsing

*Wolfgang Hrst*

Automatic Annotation and Semantic Retrieval of Video Sequences using Multimedia Ontologies

*Marco Bertini*

Mapping Learning in Eigenspace for Harmonious Caricature Generation

*Junfa Liu*

RITZ: a Real-time Tool for Interactive spatilZation

*Joshua Nixdorf*

Remote Rendering and Streaming of Progressive Panoramas for Remote Virtual Walkthrough on Mobile Devices

*Azzedine Boukerche*

The Influence of Cross-Validation on Video Classification Performance

*Jan van Gemert*

A Scalable Service for Photo Annotation, Sharing, and Search

*Benjamin Lee*

Semantic Image Retrieval Based on Probabilistic Latent Semantic Analysis

*Amin Shah-Hosseini*

Diversifying the Image Retrieval Results

*Kai Song*

Browsing Personal Media Archives with Spatial Context Using Panoramas

*Brett Adams*

Model Generation for Video-based Object Recognition

*Humera Noor Yusuf*

Measuring Movement Expertise in Surgical Tasks

*Kanav Kahol*

Collaborative Dancing in Teleimmersive

*Zhenyu Yang*

Reading the Fine Print: The Effect of Text Legibility on Perceived Video Quality in Mobile TV

*Hendrik Knoche*

**Noon-4:00pm Arts Short Papers Poster Session 2**

Authors at posters 1:00-2:00pm and 3:30-4:00pm

**Session Chair: ???**

User Authorship and Creativity within Interactivity

*Karl Willis*

Variations 10b: A Digital Realization of Cage's Variations II

*Nicholas Knouf*

Tabletop Community: Artwork for Visualization of Social Interactions Using a Bipartite Network

*Noriyuki Fujimura, Satoshi Fujiyoshi, Tom Hope, Takuichi Nishimura*

Taking Sides: Dynamic Text and Hip-Hop Performance

*Jason Lewis, Yannick Assogba*

The Computational Extraction Of Spatio-Temporal Formal Structures in the Interactive Dance Work

22

*Hari Sundaram, Thanassis Rikakis, Jodi James, Vidyarani Dyaberi*

**2:00-3:30pm Panel Session: Multimedia and Web 2.0 - Hype, Challenge, Synergy**

**Session Chair:** Susanne Boll, Univ. of Oldenburg

Panelists:

Edward Chang, Google R&D

Marc Davis, Yahoo

Patrick Schmitz, UC Berkeley

**2:00-3:30pm Applications 5: Multimedia Applications Potpourri**

**Session Chair:** Jan Pieper, IBM

Electronic Clipping System with Invisible Barcodes  
*Koichi Kamijo, Noboru Kamijo, Masaharu Sakamoto*

The Design of a Real-Time, Multimodal Biofeedback System for Stroke Patient Rehabilitation  
*Hari Sundaram, Yinpeng Chen, Thanassis Rikakis*

Video Object Segmentation by Motion-based Sequential Feature Clustering  
*Mei Han, Wei Xu, Yihong Gong*

**2:00-3:30pm Brave New Topics Session 1: Multimedia Signal Processing and Systems in Healthcare and Life Sciences**

**Session Chair:** Nevenka Dimitrova, Philips Corporation

Personalized Concept-Based Multimedia Health Records: Opportunities and Challenges  
*Shahram Ebadollahi, Shih-Fu Chang, Tanveer Syeda-Mahmood, Anni Coden, Arnon Amir, Michael A. Tanenblatt*

Multimedia Signal Processing for Behavioral Quantification in Neuroscience  
*Peter Andrews, Sigal Saar, Haibin Wang, Dan Valente, Jihene Serkhane, Ofer Tchernichovski<sup>2</sup>, Ilan Golani<sup>3</sup>, Partha Mitra<sup>1</sup>*

Analysis and Visualization of DNA Spectrograms: Open Possibilities for the Bioinformatics Research  
*Nevenka Dimitrova, Yee Him Cheung, Michael Zhang*

**2:00-5:30pm Demo Session 2**

**Session Chair:** Raju Rangaswami, Florida International Univ.

Large-Scale News Video Retrieval via Visualization  
*Hangzai Luo, Jianping Fan, Yuli Gao, Shin ichi Satoh, William Ribarsky*

The MediaMill Large-lexicon Concept Suggestion Engine  
*Marcel Worring, Cees G.M. Snoek, Bouke Huurnink, Jan van Gemert, Dennis Koelma, Ork de Rooij*

MOM: Multimedia Ontology Manager. A Framework for Automatic Annotation and Semantic Retrieval of Video Sequences  
*Marco Bertini, Alberto Del Bimbo, Carlo Torniai, Rita Cucchiara, Costantino Grana*

Metadata Production Framework and Metadata Editor  
*Masanori Sano, Yoshihiko Kawai, Hideki Sumiyoshi, Nobuyuki Yagi*

Mobile Camera Supported Document Redirection  
*Qiong Liu, Paul McEvoy*

PEANO: Pictorial Enriched ANnotation of Video  
*C. Grana, R. Vezzani, D. Bulgarelli, G. Gualdi, R. Cucchiara, M. Bertini, C. Torniai, A. Del Bimbo*

ZooMICSS: A Zoomable Map Image Collection Sensemaking System (The Katrina Rita Context)  
*Ross Graeber, Andruid Kerne, M. Kathryn Henderson*

International Remix: Video Editing for the Web  
*Patrick Schmitz, Peter Shafton, Ryan Shaw, Samantha Tripodi, Brian Williams, Jeannie Yang*

Speakr: Auditory Skimming and Scrolling

*S H Srinivasan*

3dB: A System for Geometric Tagging  
*Subhajit Sanyal, S H Srinivasan*

Fotowiki - Distributed Map Enhancement Service  
*Wen-Yen Chen, Ben N Lee, Edward Y Chang*

GLS: Simulator for Online Multi-player Games  
*Wladimir Palant, Carsten Griwodz, P I Halvorsen*

Interfaces for Interactive Audio-Visual Media Browsing  
*Wolfgang H rst, Tobias Lauer, Robert Kashuba*

A System for 3D Projected Virtual Content Insertion into Broadcast Tennis Video  
*Xin Yan, Xinguo Yu, Tran Thi Phuong Chi*

Searching and Browsing Large Scale Image Database Using Keywords and Ontology  
*Yuli Gao, Hangzai Luo, Jianping Fan*

Vocal Telekinesis; Physical Control of Inanimate Objects with Minimal Paralinguistic Voice Input  
*Sama'a Al Hashimi, Gordon Davies*

**3:30-4:00pm Coffee break** (sponsored by IBM)

**4:00-6:00pm Content 4: Event and Copy Detection**

**Session Chair:** James Wang, Penn State Univ.

Visual Attention Detection in Video Sequences Using Spatiotemporal Cues  
*Yun Zhai, Mubarak Shah*

Towards Optimal Audio "Keywords" Detection for Audio Content Analysis and Discovery  
*Lie Lu, Alan Hanjalic*

Robust Voting Algorithm Based on Labels of Behavior for Video Copy Detection  
*Julien Law-To, Olivier Buisson, Val ririe Gouet-Brunet, Nozha Boujemaa*

Fast Tracking of Near-Duplicate Keyframes in Broadcast Domain with Transitivity Propagation  
*Chong-Wah Ngo, Wanlei Zhao, Yu-Gang Jiang*

**4:00-6:00pm Brave New Topics Session 2: Human-Centered Multimedia**

**Session Chair:** Nicu Sebe, University of Amsterdam

Human-Centered Computing: A Multimedia Persepctive  
*Alejandro Jaimes, Nicu Sebe, Daniel Gatica-Perez*

Human Computing for Interactive Digital Media  
*Alex Pentland, Jonathan Gips, Wen Dong, Will Stoltzman*

Human-Centered Design Meets Cognitive Load Theory: Designing Interfaces that Help People Think  
*Sharon Oviatt*

Multimedia Production and Human-Centered Computing  
*Marc Davis*

**4:00-5:30pm Doctoral Symposium**

**Session chair:** Yi Wu, Intel Corp.

A Multi-Stream Adaptation Framework for Tele-immersion  
*Zhenyu Yang*

Large-scale Video Retrieval via Semantic Classification  
*Hangzai Luo*

Rate-Accuracy Tradeoff in Automated, Distributed Video Surveillance Systems  
*Pavel Korshunov*

**7:00-9:00pm Conference Banquet**

## Thursday, October 26, 2006

**8:30-10:00am Keynote presentation**

**Session Chair:** Klara Nahrstedt, UIUC

Implicit Participation  
*Bradley Horowitz, Yahoo!*

**10:00-10:30am Coffee break** (sponsored by Yahoo! Research)

**10:30am-noon Content 5: Image Annotation**

**Session Chair:** Lynn Wilcox, FXPAL

SmartLabel: A Object Labeling Tool using Iterated Harmonic Energy Minimization  
*Wen Wu, Jie Yang*

Automatic Image Annotation by Incorporating Feature Hierarchy and Boosting to Scale-up SVM Classifier Training  
*Jianping Fan, Ramesh Jain*

Real-time Computerized Annotation of Pictures  
*Jia Li, James Wang*

**10:30am-noon Systems 3: Assorted Topics**

**Session Chair:** Balakrishnan Prabhakaran, UT Dallas

Event on Demand with MPEG-21 Video Adaptation System  
*Min Xu, Jiaming Li, Liang-Tien Chia, Jesse Jin, Yiqun Hu, Bu Sung Lee, Deepu Rajan*

Very Low Complexity MPEG-2 to H.264 Transcoding Using Machine Learning  
*Gerardo Fernandez-Escribano, Hari Kalva, Pedro Cuenca, Luis Orozco-Barbosa*

Modelling Dependency in Multimedia Streams  
*Alexander Eichhorn*

**10:30am-noon Open Source and Video Program**

**Session Chair:** Scott Brandt, UC Santa Cruz, and Wu-Chi Feng, Portland State Univ.

[Open Source Winner]

CLAM: C++ Library for Audio and Music

*Xavier Amatriain, Pau Arumi, David Garcia*

Audiovisual Slideshow: Present Your Journey by Photos

*Jun-Cheng Chen, Wei-Ta Chu, Jin-Hau Kuo, Chung-Yi Weng, Ja-Ling Wu1*

Automatic Scaling and Cropping of Videos for Devices with Limited Screen Resolution

*Stephan Kopf, Fleming Lampi, Thomas King, Wolfgang Effelsberg*

Globe4D, Time-Traveling with an Interactive Four-Dimensional Globe

*Rick Comanje, Nico van Dijk, Hanco Hogenbirk, Danica Mast*

Virtual Videography

*Rachel Heck, Michael Wallick, Michael Gleicher*

Video Color Adaptation for Mobile Devices

*Stephan Kopf, Thomas King, Fleming Lampi, Wolfgang Effelsberg*

Video Inpainting and Implant via Diversified Temporal Continuations

*Timothy K. Shih, Nick C. Tang, Wei-Sung Yeh, Ta-Jen Chen*

### **Noon-6:00pm: Interactive Arts Program Exhibition (at UCSB)**

Shuttle vans available approximately hourly between hotel and UCSB

### **Noon-2:00pm Lunch (on your own)**

### **2:00-3:30pm Content 6: Multimedia Exploration**

**Session Chair:** Apostol Natsev, IBM TJ Watson

Automatic Video Annotation by Semi-supervised Learning with Kernel Density Estimation

*Meng Wang, Xian-Sheng Hua, Yan Song, Xun Yuan, Li Shipeng, Hong-Jiang Zhang*

Toward Bridging the Annotation-Retrieval Gap in Image Search

*Ritendra Datta, Weina Ge, Jia Li, James Wang*

Extraction of social context and application to personal multimedia exploration

*Brett Adams, Dinh Phung, Svetha Venkatesh*

### **2:00-5:30pm MIR Workshop**

**3:30-4:00pm Coffee break** (sponsored by Yahoo! Research)

### **4:00-5:30pm Discussion: Foundations and Directions of Multimedia Research**

**Session Chair:** Matthew Turk, UCSB, and Klara Nahrstedt, UIUC

## **Friday, October 27, 2006**

**8:30am-5:30pm Workshop on Multimedia Information Retrieval (MIR 2006)**

**8:30am-5:30pm Workshop on Capture, Archival and Retrieval of Personal Experiences (CARPE 2006)**

**8:30am-5:30pm Workshop on Multimedia Content Protection and Security (MCPS 2006)**

**8:30am-5:30pm Workshop on Human-Centered Multimedia (HCM 2006)**

**8:30am-5:30pm Workshop on Video Surveillance and Sensor Networks (VSSN 2006)**

**8:30am-5:30pm Workshop on Audio and Music Computing for Multimedia (AMCMM 2006)**

**Noon-6:00pm: Interactive Arts Program Exhibition (at UCSB)**

# Full Papers

## Call for Full Papers

The 14th Annual ACM International Conference on Multimedia (MM 2006) will be held in Santa Barbara, California, October 22-27, 2006. ACM Multimedia is the premier technical multimedia conference attended by an international community of researchers from both academia and industry.

MM 2006 invites your participation in the premier annual multimedia conference, covering all aspects of multimedia computing: from underlying technologies to applications, theoretical foundations to experimental systems, and servers to networks to devices. MM 2006 seeks high-quality, original papers that will set the standard in the field and stimulate the trends for years to come. We especially encourage introduction of novel media such as haptic, olfactory, multiple sensors, etc.

**Note:** For papers for the interactive art program (full/short/exhibition) please visit the "[Interactive Arts Program](#)" website.

## Technical Program

The technical program will consist of plenary sessions and talks with topics of interest in:

- **Multimedia analysis, processing, and retrieval**, including multimedia semantics, aesthetics, modeling, assimilation/fusion, audio/video/multi-modal processing, multi-sensor processing, multimedia content description and indexing, multimedia digital rights management (protection and attribution), content-based retrieval with emphasis on multiple and novel media.
- **Multimedia tools, end-systems, and applications**, including new UI metaphors, usable distributed collaboration, authoring, multi-modal interaction and integration, multimedia in e-learning, entertainment, personal media, assisted living, and virtual environments.
- **Multimedia networking and systems**, including context-aware multimedia communications, Internet telephony, peer-to-peer streaming, audio/video streaming, multimedia content distribution, wireless multimedia, adaptive support for scalable media, Internet protocols, multimedia servers, operating systems, middleware and QoS.

The above list is not exhaustive. We encourage submissions in new and emerging areas. We particularly encourage submissions of paper addressing **foundational sciences** of multimedia.

## Important Dates

**1 July** Notification of acceptance

**1 August** Camera-ready papers

## Submission Systems

**NEW!** The online submission systems for conference papers have already closed. However, the submissions for some workshops are still open (see [Workshops](#) page for further information).

## Submission Instructions

Prepare a paper (no more than 10 pages) using the ACM template for the conference -- Portable Document Format (PDF) or PostScript (version 2 or later), formatted in two-column conference style. Submissions should present original reports of substantive new work. Papers should properly place the work within the field, cite related work, and clearly indicate the innovative aspects of the work and its contribution to the field. We will not accept any paper that, at the time of submission, is under review for or has already been published or accepted for publication in a journal or another conference. The names of the violated authors will be shared with other major conferences.

Please see the ACM proceedings template available at <http://www.acm.org/sigs/pubs/proceed/template.html>.

## Double-Blind Reviewing

All research papers submitted to MM 2006 will undergo a "double-blind" reviewing process: the program committee members and referees who review the paper will not know the identity of the authors. *The double-blind reviewing is applicable for only research papers and is not applicable to other submissions.*

For the purpose of double blind review, please remove names of authors and affiliations from the paper heading, and all references to (your own) papers or systems that may reveal your identity.

To ensure anonymity of authorship, authors must prepare their manuscript as follows:

1. Authors' names and affiliations must not appear on the title page or elsewhere in the paper.
2. Funding source(s) must not be acknowledged on the title page or elsewhere in the paper.
3. In place of names of authors and affiliations in the heading, please replace it by the paper id as "Paper xxx". For some references to your own papers, you may want to leave the reference id but remove the details of references by stating: "reference removed for the purpose of anonymous review."
4. Omit all personal acknowledgements. Research group members or other colleagues or collaborators must not be acknowledged anywhere in the paper. There should also be no acknowledgement section in the paper.
5. It is strongly suggested that the submitted file is named with the assigned submission number. For example, if assigned paper number is 352, then name your submitted file **352.pdf**.
6. Source file naming must also be done with care. For example, if your name is Jane Smith and you submit a PDF file generated from a **.dvi** file called **Jane-Smith.dvi**, one can infer your authorship by looking into the PDF file.

**Common sense** and careful writing can go a long way toward preserving anonymity without diminishing the quality or impact of a paper. The goal is to preserve anonymity while still allowing the reader to fully grasp the context (related past work, including your own) of the submitted paper. In past years this goal has been achieved successfully by hundreds of papers. If you need specific guidance, please contact the Program Chairs.

*It is the responsibility of authors to do their very best to preserve anonymity. Papers that do not follow the guidelines here, or otherwise potentially reveal the identity of the authors, are subject to immediate rejection.*

## Contacts

For any questions regarding full papers, please email the co-chairs:

- [Yong Rui](#) - Content
- [Ketan Mayer-Patel](#) - Systems
- [Wolfgang Klas](#) - Applications

# Short Papers

## Call for Short Papers

Short papers should present interesting recent results or novel thought-provoking ideas that are not quite ready for a regular full-length paper. We are looking for interesting and bold papers in all areas of multimedia technology and its application. Submissions are encouraged in areas including multimedia databases, content analysis, media processing, compression, multimedia storage, networking, multimedia and hypermedia authoring, multimedia user interfaces and innovative applications. Submissions will be peer-reviewed to ensure quality. *Accepted short papers will be presented in an interactive poster format at the conference and appear in the ACM DL along with the accepted long papers.*

## Important Dates

<b>1 June</b>	<b>Short Paper Submission Deadline (Firm Deadline)</b>
<b>10 July</b>	<b>Author notification</b>

## Submission Systems

**NEW!** The online submission system for short papers is available now. Please follow the link below for submission. After login, the authors will need to select to which track (systems, content, or applications) the papers belong.

- [Short Papers](#)

## Submission Instructions

Prepare a four page paper using the ACM template for the conference. The paper should be in Portable Document Format (PDF) and formatted in two-column conference style. Short papers are in the same format as the long papers. Please see the ACM proceedings template available at <http://www.acm.org/sigs/pubs/proceed/template.html>. **NO abstract need to be submitted prior to the submission of the actual paper. Moreover, this year the short papers will NOT undergo a blind review process. For uniformity, please include the names and affiliations of all authors in the submissions.**

## Technical Program Committee

- [Content Track](#)
- [Applications Track](#)
- [Systems Track](#)

## Contacts

For any questions regarding short papers, please email the co-chairs:

- [Brian Bailey](#) (UIUC)
- [Belle Tseng](#) (NEC, USA)
- [Nalini Venkatasubramanian](#) (UCI)

# Demo Program

## Call For Demos

This year's Technical Demonstrations will showcase leading edge work and work in progress in every area of multimedia, including:

- Multimedia analysis, processing, and retrieval
- Multimedia networking and system support
- Multimedia tools, end-systems, and applications

Submissions are particularly encouraged in the areas of haptics, smell, sensors, novel interface design, immersive technologies, wireless multimedia applications, multimedia-based security applications, digital rights management, multimedia databases, content analysis, content-based retrieval, multimedia storage, multimedia networking, media processing, entertainment, compression, hypermedia authoring, and innovative applications.

An award will be given for the **best technical demonstration**. Demonstrators will be provided with space and access to network. Participants will be required to provide their own computing equipment and any additional network, display, haptic, sensor, etc., hardware needed for the demonstration.

## Important Dates

<b>1 June</b>	<b>Technical Demo Submission Deadline (Firm Deadline)</b>
<b>22 June</b>	<b>Technical Demo Acceptance Notification</b>
<b>29 June</b>	<b>Camera-ready papers</b>

## Submission Instructions

### What to submit:

- + A short paper (two page maximum) describing your technical demo. Please use the standard ACM format (see <http://www.acm.org/sigs/pubs/proceed/template.html>)
- + In addition to the two page short paper, you may optionally submit one or more of the following multimedia documents that describe or illustrate the demonstration:
  - a Microsoft Powerpoint presentation
  - a Macromedia Flash presentation, or
  - an video clip (not more than 3 minutes long) in either Windows Media or Quicktime format

### How to submit:

Please submit your demonstration proposal by sending email with the short paper and other supporting materials as attachments to [acmmm-demo@cis.fiu.edu](mailto:acmmm-demo@cis.fiu.edu)

## Contacts

For any questions regarding demos, please email the co-chairs:

- [Baochun Li](#) (U. Toronto)

- **Raju Rangaswami** (FIU)

# Doctoral Symposium

[[Note: Click here for information on student conference participation grants](#)]

## Call For Doctoral Symposium

The Doctoral Symposium is an opportunity for students involved in the preparation of a PhD in any area of Multimedia to interactively discuss their research issues and ideas with senior researchers, receive constructive feedback from members of the research community and expose themselves as up and coming multimedia researchers. During the Doctoral Symposium, selected students will present their thesis topic, the work they have performed so far and the results that they have obtained. They will also reveal the difficulties, problems and questions that they encounter in the continuation of their work and can ask for comments from the audience.

## Important Dates

**1 June** Short Paper/Interactive Art Program Short Paper/Open Source/Doctor Program/Demo Proposal Submission Deadline (**Firm Deadline**)

## Submission Instructions

Interested students should submit 2-3 page abstracts describing their research using the standard paper submission template for the conference. The abstract should be in Portable Document Format (PDF) and formatted in two-column conference style. Please see the ACM proceedings template available at <http://www.acm.org/sigs/pubs/proceed/template.html>. Please email your proposal abstract to the Doctoral Symposium Chair.

## Contacts

For any questions regarding the Doctoral Symposium please email the chair:

- [Reza Rejaie](#) (University of Oregon)

# Video Program

Video program videos are available at <http://amg.cs.pdx.edu/>

## Call for Video Demonstrations

We invite you to make a video for ACM Multimedia 2006. This forum is an ideal way to demonstrate your tool, system, application or brave new idea without having to bring the equipment for a "live" demo. Even if you have never made a video before, we encourage you to consider this opportunity for Video Submissions at ACM Multimedia 2006.

## Important Dates

10 June	<b>NEW!</b> Submission deadline for video figures associated with a paper ( <b>Extended</b> )
10 June	<b>NEW!</b> Submission deadline for video demonstrations, video visions and video figures for short papers ( <b>Extended</b> )
28 June	Notification of Acceptance for all categories
TBD	Submission of final video version and submission of camera-papers

## Video Submissions Category

There are three categories of video submissions:

- **Video demonstrations** (8 minutes) - gives a self-explanatory demo of a tool, system, or application. In addition we would like to encourage media artists to submit work that exemplifies new trends in video art or describes an interactive multimedia art application
- **Video visions** (8 minutes) - describes a new brave idea of multimedia technology.
- **Video figures** (3 minutes) - accompanies a regular or short paper submission and serves as an illustration.

Accepted video demonstrations and video visions are presented during the video demonstration session, whereas video figures are presented together with the paper during the allocated session. Videos will be published on a dedicated web site at the SIGMM website (see <http://www.sigmm.org/apache/video2004/> - for previous years). The best video demonstration and video vision will get a special podium.

## Video Submissions Format

Videos in the categories demonstration and vision need to be accompanied with a 2 page description of the video content using the ACM template. The descriptions of the selected videos will become part of the conference proceedings. The format for videos is **MPEG-2**, although we will accept other digital formats for the review process, such as **MPEG-1**, **MOV** or **AVI**. **Please note: for the final video version MPEG-2 is obligatory.**

## Submission Instructions

- For all submissions please indicate the type: demonstration, vision or short.
- For all categories: Prepare video for review in MPEG-2 format (preferred - other possible formats: mpeg1, mov or avi). Make the video available at a stable URL. The URL must be accessible at least until the 1st of July 2006. Please make the URL known to the video program chair, [wuchi@cs.pdx.edu](mailto:wuchi@cs.pdx.edu).

- For those submitting a video figure: send the following information to [wuchi@cs.pdx.edu](mailto:wuchi@cs.pdx.edu) by JUNE 1, 2006:
  - + Title of paper
  - + Paper track (if a video figure)
  - + Authors
  - + Affiliations
  - + Address
  - + Primary contact
- For those submitting a video demonstration or video vision: Prepare a 2-page PDF description (using the ACM template). Note, there is no double-blind review applied at the video track. Submit the paper along with the above information to [wuchi@cs.pdx.edu](mailto:wuchi@cs.pdx.edu).

### **Contacts**

For any questions regarding the Video Program, please email the chair:

- **Wuchi Feng** (Portland State University)

# Open Source Competition

The winner of the **ACM Multimedia 2006 Open Source Competition** is

**CLAM** (<http://clam.iaa.upf.edu>)  
by **Xavier Amatriain, Pau Arumi, and David Garcia**

CLAM is a comprehensive open-source C++ framework for researching audio and music through rapid cross-platform prototyping of audio applications. It offers an abstract model for audio systems and includes a repository of processing algorithms and data types, as well as a number of tools such as audio and MIDI input/output. These features can be exploited to build cross-platform applications or to build rapid prototypes to test signal processing algorithms.

Congratulations to Xavier, Pau, and David!

---

## Call for Open Source Software Competition

**The open-source software competition is a relatively recent and highly successful addition to the ACM Multimedia program. 2006 will be our third year in running the competition.**

The competition celebrates the invaluable contribution of researchers who advance the field by providing the community with implementations of codecs, middleware, frameworks, toolkits, libraries, and other multimedia software.

To qualify, software must be provided with source code and licensed in such a manner that it can be used free of charge in academic and research settings. For the competition, the software will be built from the sources. Submissions should be made in the form of a compressed zip or tar archive file and must include the following:

- A text file named **README.txt** which summarizes the purpose of the software and its overall design and use.
- A text file named **INSTALL.txt** which includes any and all instructions on building and installing the software and any system requirements. The judging committee will make a reasonable effort to build any submitted software. However, if we are unable to make the software run given a reasonable effort, we must unfortunately exclude it from the competition. Therefore, complete and clear build and install instructions are a crucial component of any submission.

Additionally, authors are highly encouraged to provide as much documentation as possible, including examples of how the provided software might be used. Entries may be made available to the research community via the ACM Multimedia SIG website and the judging committee may employ feedback from the community at large to help determine the winning entry. In order to encourage participation, non-winning entries are encouraged to submit again in the future. Student-led efforts are particularly encouraged. Authors of the winning entry will be recognized formally at ACM Multimedia 2006, awarded a prize (to be announced), and invited to demonstrate their software as part of the conference demonstration program.

## Submission Instructions

People interested in submitting an entry to the open source software competition should submit 1-2 pages containing the following information via the [online submission system](#).

- Title of submission.
- Names and affiliations of authors (indicate students).
- Brief overview of submission (i.e., purpose, intended audience, main features, etc.)

- URL for a compressed tar or zip archive file with all source code and documentation.

If you have questions about what can or should be submitted, please feel free to address these questions to the competition chair **Scott Brandt** and I will do my best to answer. We are excited to be organizing this innovative and creative addition to the ACM Multimedia program. Please help us make this venture a continuing success by submitting your entry to the competition. We look forward to your participation.

### Important Dates

- 2 June** **NEW!** Open Source Submission Deadline (**Extended**)
- 1 July** Notification of Acceptance
- 1 August** Camera Ready Write-up Due

### Contacts

For any questions regarding the open source competition, please email the chair:

- **Scott Brandt**

# Panels

## Call For Panels

The panels at Multimedia 2006 will provide a forum venturous position statements and vivid debates of controversial, unresolved, or emerging issues related to all aspects of multimedia.

Special focus will be set on the "multi" part of the "multimedia" with its specific challenges. With the panels we would like to stimulate discussions on challenges and research trends. The conference will host two panels, one each on the first and the second day of the conference.

## Important Dates

**17 April** Panels Submission Deadline (**Firm Deadline**)

## Submission Instructions

The submission deadline for panel proposals is April, 17th 2006. Please send a 1-page description of the panel including the title, participants, and a short description of what ideas will be covered. The list of participants should include an indication of whether the person has agreed to be a panel member. You should email your panel proposal including CV's of the panelist in PDF to the panel co-chairs.

Best panel proposals will be selected by the panel co-chairs supported by a Panel Selection Board comprised of members from the ACM MM 2006 Program Committee.

## Contacts

For any questions regarding panels, please email the panel co-chairs:

- **Wei-Ying Ma** (Microsoft Research, Asia)
- **Larry Rowe** (UC, Berkeley)

## Multimedia Foundational Sciences / Brave New Topics

The ACM International Conference on Multimedia (ACM MM-2006) is seeking to explore foundational sciences of multimedia research in the form of Brave New Topics Sessions at this year's conference.

We solicit proposals for special sessions that examine basic research forming the core multimedia research discipline. We seek to emphasize those aspects that best differentiate multimedia research from its related and constituent disciplines, e.g., *image processing, pattern recognition, natural language processing, speech and signal processing, video technology, information theory*. By providing specific focus on the foundations of multimedia science, we wish to improve the ability of the multimedia community to work together in advancing the field.

Example topics include but are not limited to: **human-centered multimedia systems, multimedia semantic analysis, multi-modal systems, multimedia user experiences, multimedia in life sciences**.

Each Brave New Topics session shall be comprised of a coordinated set of papers managed by the session organizer. Proposals shall outline the scope of the session and the planned contributors. The session organizers, together with the Brave New Topics chairs, shall work together on organizing a final session.

The Brave New Topics session is being handled as a separate track at this year's conference. We expect to accommodate between **2** and **4** Brave New Topics sessions. Since acceptance is likely to be highly selective, we encourage you to submit a short session proposal well in advance of the final submission dates. We can provide feedback on your proposal or help bring groups together in expanded sessions.

### Important Dates

- |                    |  |
|--------------------|--|
| <b>15 May 2006</b> | <b>Submission deadline for brave session proposals</b> |
| <b>22 May 2006</b> | <b>Notification of acceptance</b>                      |

### Submission Instructions

- **What to submit:** A session proposal should include
  - + a title
  - + a one-page abstract describing the scope of the session
  - + names and bios of the organizers
  - + names of the participants
  - + abstracts for each of the proposed papers
- **Where to submit:** Please submit your proposal by sending email with the abstract and other supporting materials as attachments to the brave new topic co-chair.

### Contacts

For any questions regarding the Foundational Sciences Brave New Topics program, please email the co-chairs:

- **John R. Smith** (IBM T. J. Watson Research Center, USA)
- **Jonathan Foote** (FX Palo Alto)

# Interactive Arts Program

The ACM Multimedia Interactive Art Program seeks to bring together the arts and multimedia communities to create the stage to explore, discuss, and push the limits for the advancement of both multimedia technology through the arts, and the arts through multimedia technology.

For more information on submissions to the Interactive Arts Program: Long Papers, Short Papers, and Exhibition - please visit the "[Interactive Arts Program](#)" website.

## Important Dates

- 1 June** Long/short papers and exhibition submission deadline (**Extended**)
- 1 July** Notification of acceptance
- 1 August** Camera-ready papers

## Contacts

For any questions regarding full papers, please email the chairs and curators:

- Alejandro Jaimes (Fuji Xerox, Japan)
- [George Legrady](#) (UCSB)
- [Lonce Wyse](#) (Institute for Infocomm Research & NUS, Singapore)

# ACM MM 2006 Tutorials

## **Title: Peer-to-Peer Multimedia Applications**

Full-day Tutorial at ACM Multimedia 2006, on Oct. 23, Monday

Speaker:

Jin Li, Microsoft Research

Intended Audience: Professionals, researchers and students interested in building P2P multimedia applications.

Overview:

In both academia and industry, peer-to-peer (P2P) applications have attracted great attention. Peer-to-peer file sharing applications, such as Napster, Gnutella, Kazaa, BitTorrent, Skype and PPLive, have witnessed tremendous success among end users. And the uses of peer-to-peer network for multimedia streaming, conferencing, gaming, file backup, information retrieval is on the rise. Recent statistics suggests that P2P traffic accounts for as much as 70% of Internet traffic. The purpose of the tutorial is to examine issues associated with the successful building and deployment of a P2P multimedia application. The technologies discussed can be applied to P2P file sharing, P2P conference, P2P media streaming, P2P VoIP, and P2P storage applications. We start by examining two popular P2P applications, BitTorrent and Skype. The study of the two P2P applications helps us to understand the design principles of P2P applications in general. We then investigate a number of tools for building P2P multimedia applications, such as the overlay network, the scheduling algorithm, the erasure resilient coding, and NAT/firewall traversal. Finally, we move on to critical deployment decisions that make or break the P2P applications, such as P2P economy, security issues in P2P application, peer selection, monitoring and debugging utilities in P2P application.

Speaker's Bio:

Dr. Jin Li received his Ph.D. in electrical engineering from Tsinghua University (Beijing, China) in 1994. From 1994 to 1996, he served as a Research Associate at the University of Southern California (USC). From 1996 to 1999, he was a Member of the Technical Staff at the Sharp Laboratories of America (SLA), (Camas, WA), and represented the interests of SLA in the JPEG2000 and MPEG4 standardization efforts. He was a Researcher/Project Leader at Microsoft Research Asia (Beijing, China) from 1999 to 2000. He is currently a Senior Researcher at Microsoft Research Redmond. From 2000, Dr. Li has also served as an Adjunct Professor in the Electrical Engineering Department, Tsinghua University (Beijing, China). Dr. Li has personally built a number of P2P applications, such as P2P web hosting, P2P streaming and P2P distributed storage system. He was the driving force behind Microsoft's strategy and application development in the peer-to-peer area. He is the lead guest editor of the special issue of Content Storage and Delivery in Peer-to-Peer Networks for IEEE Trans. on Multimedia and the guest editor of the special issue of Advances in Peer-to-Peer Streaming Systems for IEEE Journal on Selected Areas in Communication. He has organized a special session on Peer-to-Peer Media Communication for MMSp 2005, and has co-organized the workshop of Advances in Peer-to-Peer Multimedia Streaming in ACM Multimedia 2005, and the workshop of Recent Advances in Peer-to-Peer Streaming in QShine 2006. He holds 15 issued US patents, with many more pending.

---

## **Title: Interactive Digital Television and Multimedia Systems**

Half-day Tutorial at ACM Multimedia 2006, on Oct. 23, Monday

Speakers:

Pablo Cesar, CWI Amsterdam

Konstantinos Chorianopoulos, AUEB, Athens

Intended Audience: The level of the tutorial will be from introductory to intermediate.

Tutorial website: <http://uitv.info/tutorials/acm-multimedia-06/>

Overview:

Interactive digital television is an emerging field with a high impact in our societies: it offers interactive services to the masses. This tutorial aims to establish a common framework by summarizing the most significant results in this multidisciplinary field. The review includes topics such as content distribution, system software of the receivers, and user interaction. In addition, we will discuss current commercial events such as the next generation of optical discs (e.g., blue-ray), BBC peer-to-peer service, and mobile television. Based on this discussion, we will formulate an agenda for further research. The agenda includes, for example, end-user enrichment of television content and social television. This half-day tutorial will provide the attendee a solid understanding of the technologies currently in use and an introduction of the open questions in the field.

Speaker's Bio:

Dr. Pablo Cesar holds a Ph.D. in interactive digital media (with the dissertation: A Graphics Software Architecture for High-End Interactive TV Terminals; 2005). Since 2000 he has participated in several iDTV-related projects. Currently, he is researching how TV viewers are becoming producers and distributors.

---

**Title: Flexible Modelling and Performance Debugging of Real-Time Embedded Multimedia Systems**

Half-day Tutorial at ACM Multimedia 2006, on Oct. 23, Monday

Speaker:

Samarjit Chakraborty  
Department of Computer Science  
National University of Singapore

Intended Audience: The level of the tutorial will be from introductory to intermediate.

Tutorial website: <http://www.comp.nus.edu.sg/~samarjit/mm06/>

Overview:

This tutorial is primarily intended for an audience with a background in traditional real-time multimedia applications, who would like to establish a connection to embedded systems design for the multimedia domain. It would be useful to researchers, students, multimedia application developers, as well as engineers interested in getting an overview of recent developments in the area of multimedia processing on embedded System-on-Chip (SoC) platform architectures. The emphasis will be on tools and techniques for modelling and analyzing such architectures from the perspective of implementing multimedia applications on them. After attending the tutorial, the audience should get an overview of common design issues in the embedded multimedia domain and techniques for dealing with them. These would include handling real-time and throughput constraints associated with continuous media-processing, power/performance analysis and runtime platform-management techniques. No background in embedded systems or VLSI design automation techniques will be assumed and the lectures will focus on a diverse set of applications, with detailed self-contained examples.

Speaker's Bio:

Samarjit Chakraborty is an Assistant Professor in the Department of Computer Science at the National University of Singapore. He obtained his Ph.D. from ETH Zurich in 2003. For his Ph.D. thesis, he received the ETH Medal and the European Design and Automation Association's Outstanding Doctoral Dissertation Award in 2004. Samarjit's research interests are primarily in the area of system-level design of real-time and embedded systems, with a focus on architectures for multimedia applications. He has served on the technical program committees of a number of conferences in the real-time embedded systems area, including CODES+ISSS, DATE, RTSS, ECRTS and RTCSA. He will also serve as the TPC Co-Chair of the 2006 IEEE Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia), which will be held in Seoul in October this year. In the last few years he has extensively worked on problems related to the topic of this tutorial and published both in real-time, as well as

electronic design automation conferences and journals. Recently, his paper entitled "Approximate VCCs: A New Characterization of Multimedia Workloads for System-level MpSoC Design" was nominated for a Best Paper Award at the 2005 ACM Design Automation Conference (DAC).

---

**Title: Computer Audition: An introduction and research survey**

Half-day Tutorial at ACM Multimedia 2006, on Oct. 23, Monday

Speaker:

Shlomo Dubnov

Department of Music, University of California, San Diego

Intended Audience: The level of the tutorial will be from introductory to intermediate.

Tutorial website: <http://music.ucsd.edu/~sdubnov/ComputerAudition.htm>

Overview:

Computer Audition is an interdisciplinary field of audio understanding by machine that combines engineering, information processing and artificial intelligence, cognitive science, music theory and artistic creativity. The tutorial will survey the research in computer audition, focusing on the semantic gap between human and machine levels of audio understanding. Features and methods of audio classification will be covered from analysis-synthesis models and low-level signal features to perceptual and general audio basis representations. Examples will be given from applications of recognition, sound description, score alignment, musical summarization, affect, aesthetics and style modeling.

In the last part of the tutorial Dr. Dubnov will introduce an anticipatory listening model that offers a framework for investigating the relations between methods of computer audition, style modeling and human experience when listening to music, such as emotional force and familiarity.

Speaker's Bio:

Shlomo Dubnov is an associate Professor in music technology at UCSD. Prior to this he served as researcher in Institute for Research and Coordination of Acoustics and Music (IRCAM) in Paris and was a senior lecturer in department of communication systems engineering in Ben-Gurion University in Israel. He holds PhD in Computer Science from Hebrew University and B.Mus in music composition from Rubin Academy in Jerusalem. His work on polyspectral analysis of musical timbre and his research on machine learning of musical style are widely acknowledged by the computer music community. He served as co-PI in several projects dealing with semantic analysis of audio, such as a recent EU sponsored "Semantic HiFi" project. Currently he is co-editing a book on "The Structure of Style: algorithmic approaches to understanding manner and meaning" and working on a textbook on semantic audio processing.

---

**Title: Data Mining and Information Retrieval in Time Series and Multimedia Databases.**

Half-day Tutorial at ACM Multimedia 2006, on Oct. 23, Monday

Speaker:

Eamonn Keogh, UC Riverside

Intended Audience: The level of the tutorial will be from introductory to intermediate.

Tutorial website: <http://www.cs.ucr.edu/~eamonn/tutorials.html>

Overview:

Time series and multimedia data are ubiquitous; large volumes of such data are routinely created in scientific, industrial, entertainment, medical and biological domains. Examples include gene expression data, X-rays, electrocardiograms, electroencephalograms, gait analysis, stock market quotes, space telemetry, microarrays, CAT Scans etc. Because such data is intrinsically real valued, most of the work on information retrieval of text has little utility for such datasets.

In this tutorial, Dr. Keogh will outline the state of the art in mining and indexing such data, with particular emphasis on data representations, dimensionality reduction and similarity measures.

Speaker's Bio:

Dr. Keogh's research interests are in Data Mining, Machine Learning and Information Retrieval. He has published more than 80 papers, including works in SIGIR, SIGMOD, SIGKDD, SIGGRAPH, VLDB, ICML, EDBT, PKDD, PAKDD, IEEE ICDM, IEEE ICDE, and SIAM SDM, conferences and in the TODS, DMKD, VLDB, KAIS and IJTAI journals. Several of his papers have won best paper awards. In addition he has won several teaching awards. He is the recipient of a 5-year NSF Career Award for Efficient Discovery of Previously Unknown Patterns and Relationships in Massive Time Series Databases and a grant from Aerospace Corp to develop a time series visualization tool for monitoring space launch telemetry.

Dr. Keogh has given well-received tutorials on time series, machine learning and data mining all over the world, and his papers have been referenced well over 2,000 times.

---

**Title: Recent developments in video compression standards and their impact on embedded platforms: from scalable to multi-view video coding**

Half-day Tutorial at ACM Multimedia 2006, on Oct. 23, Monday

Speaker:

Iole Moccagatta, Ph.D.  
Scientific Director Multimedia Group  
IMEC, Belgium

Intended Audience: The level of the tutorial will be from introductory to intermediate.

Overview:

This tutorial will focus on the impact and requirements of recent developments in video compression standards on embedded platforms. It will first present and discuss recent advances in standardized video coding technology, namely the Scalable Video Coding and the Multiple-view Video Coding, both currently developed collaboratively by ITU-T VCEG and ISO/IEC MPEG. The tutorial will then discuss the requirements imposed by these new video coding technologies on embedded platforms. Particular attention will be devoted on how to balance increase performance requirements and power consumption. In this context, the use of parallelization and dynamic (run-time) resource adaptation will be discussed as well.

Speaker's Bio:

Iole Moccagatta received her M.Sc. from the University of Pavia, Italy, in 1990, and her Ph.D. from the Swiss Federal Institute of Technology in Lausanne, Switzerland, in 1995, both in electrical engineering. From 1995 to 1997 she was a member of technical staff at the Texas Instruments DSPS R&D Center in Dallas, Texas. From 1997 to 2000 she held a similar position at the Rockwell Science Center Multimedia Group in Thousand Oaks, California. From 2000 to 2003 she was a staff design engineer with the Broadband Entertainment Division at LSI Logic (previously C-Cube Microsystems) in Milpitas, California. In 2000, she joined nVIDIA in Santa Clara, California, as a video architect of the Handheld Division. She is currently the scientific director of the Multimedia Group at IMEC, Belgium. She is the author or co-author of more than 25 technical papers and holds 4 patents in the field of image and video compression.

---

**Title: Multimedia Content Protection**

Half-day Tutorial at ACM Multimedia 2006, on Oct. 23, Monday

Speakers:

Dulce Ponceleon and Nelly Fazio

IBM Almaden Research Center  
San Jose, CA 95120-6099, USA

Intended Audience: The level of the tutorial will be from introductory to intermediate.

Tutorial website:

[http://www.almaden.ibm.com/cs/people/dulce/MultimediaContentProtectionACMtutorial\\_2006.htm](http://www.almaden.ibm.com/cs/people/dulce/MultimediaContentProtectionACMtutorial_2006.htm)

Overview:

Multimedia content protection is a controversial topic. Content owners want to protect their rights while consumers want flexible usage and seamless content flow. In this tutorial we cover from cryptography fundamentals, to history, standards, state-of-the-art approaches and live demonstrations.

Key topics covered in the tutorial include:

- o Legal, Technical and Business Components
- o DRM, Conditional Access, Media Content Protection, Electronic Distribution, etc
- o IP, Copyrights, Licensing
- o Piracy and its Categories
- o Confidentiality, Authentication, Data Integrity, Non-repudiation
- o Matrix-based: CPRM
- o Tree-based: NNL
- o Watermarking and Tracing Traitors
- o History of CP Systems
- o Standards: 4C and AACCS
- o The digital home
- o On-line and emerging models

We will demonstrate the following technologies and systems: CPRM Download, Kiosk Video Distribution, Digital Home Network, and Apple iPod/iTunes model.

Speaker Bios:

Dr. Ponceleon holds an M.S. and a Ph.D. degree in Computer Science from Stanford University. While at Apple Computer Inc. she contributed to compression technologies to QuickTime. She is currently in the Content Protection Competency Center, at the IBM Almaden Research Center. She contributed to the ISO MPEG-7 standards. She is an IBM technical representative in the Advanced Access Content System (AACCS), a content protection standard for the next generation of pre-recorded and recorded optical media for consumer use with PCs and CE devices. The 4C Entity has developed content protection standards for recordable and pre-recorded media (CPRM/CPM). Dr. Ponceleon is the Chair of the 4C Technical Group since 2004. She has co-authored several tutorials in multimedia information retrieval (SIGIR 2002, SIGIR 2005 and ICASPP 2006). This year she is organizing a workshop in content protection. She holds several patents and publications in content protection and related multimedia areas.

Nelly Fazio earned her M.Sc.(2003) and Ph.D.(2006) in Computer Science from New York University. During her studies, she also conducted research at Stanford University, Ecole Normale Supérieure (Paris, France) and Aarhus University (Denmark). In 2003, she was awarded the CIMS Sandra Bleistein prize, for "notable achievement by a woman in Applied Mathematics or Computer Science." Dr. Fazio's research focuses on Cryptography and its applications to digital content protection. She has also done research on Zero-Knowledge, Multi-party Computation and Access Control. She is currently a postdoctoral fellow at IBM Almaden Research Center.

---

**Title: Semantic Indexing and Retrieval of Video**

Speaker: Marcel Worring and Cees Snoek, University of Amsterdam

Intended Audience: The tutorial is especially meant for people who are new to the field (introductory) or who have started in this direction (intermediate).

Overview:

The semantic gap between the low level information that can be derived from the visual data and the conceptual view the user has of the same data is a major bottleneck in video retrieval systems. It has dictated that solutions to image and video indexing could only be applied in narrow domains using specific concept detectors, e.g., sunset or face. This leads to lexica of at most 10-20 concepts. The use of multimodal indexing, advances in machine learning, and the availability of some large, annotated information sources, e.g., the TRECVID benchmark, has paved the way to increase lexicon size by orders of magnitude (now 100 concepts, in a few years 1,000). This brings it within reach of research in ontology engineering, i.e. creating and maintaining large, typically 10,000+ structured sets of shared concepts. When this goal is reached we could search for videos in our home collection or on the web based on their semantic content, we could develop semantic video editing tools, or develop tools that monitor various video sources and trigger alerts based on semantic events. This tutorial lays the foundation for these exciting new horizons. It will cover:

- Basic video analysis techniques
- Different methods for semantic video indexing.
- Interactive access to the data.
- Semantic retrieval
- Evaluation of indexing and interactive access in TRECVID
- The challenges ahead and how to meet them.

#### Speaker Bios:

Marcel Worring is Associate Professor of Computer Science at the University of Amsterdam, The Netherlands. He is the chair of the IAPR TC12 on Multimedia and Visual Information Systems. He is co-chair of the Conference on Image and Video Retrieval (CIVR 2007), co-organizer of the First International Workshop on Image Databases and Multi Media Search (1996), the International Conference on Visual Information Systems (1999) and the Conference on Multimedia & Expo (ICME, 2005). He is guest editor of the special issue on Semantic Image and Video Indexing in Broad domains for IEEE Transactions on Multimedia (2007). He is leading the successful MediaMill team which has been participating from the beginning of the TRECVID benchmark.

Cees Snoek received the M.Sc. degree in business information systems (2000) and the Ph.D. degree in computer science (2005) both from the University of Amsterdam, The Netherlands, where is currently a senior researcher at the Intelligent Systems Lab Amsterdam. He was a Visiting Scientist at Informedia, Carnegie Mellon University, USA in 2003. Dr. Snoek is a lead architect of the award-winning MediaMill Semantic Video Search Engine, which obtained state-of-the-art performance in recent NIST TRECVID evaluations. He is the local chair of the 2007 International Conference on Image and Video Retrieval in Amsterdam.

## ACM Multimedia 2006 Fee Schedule

Note that **September 21** is the deadline for "Advance" registration. After this date, the "Late/On-Site" fees apply.

		Non-ACM Member (USD)	ACM Member (USD)	Full-time Student (USD)
<b>MAIN CONFERENCE</b>	Advance	\$650	\$580	\$280
	Late/On-Site	\$720	\$650	\$330
<b>Half-Day Tutorial</b>	Advance	\$120	\$100	\$75
	Late/On-Site	\$145	\$120	\$100
<b>Full-Day Tutorial</b>	Advance	\$220	\$180	\$120
	Late/On-Site	\$260	\$215	\$150
<b>One-day Workshop</b>	Advance	\$180	\$150	\$100
	Late/On-Site	\$200	\$170	\$120
<b>Two-day Workshop</b>	Advance	\$230	\$190	\$135
	Late/On-Site	\$270	\$225	\$165
<b>Package: Conference and one Half-Day Tutorial</b>	Advance	\$735	\$645	\$340
		Late/On-Site not available		
<b>Package: Conference and two Half-Day Tutorials</b>	Advance	\$830	\$725	\$380
		Late/On-Site not available		
<b>Package: Conference and one Full-Day Tutorial</b>	Advance	\$830	\$725	\$380
		Late/On-Site not available		
<b>Package: Conference and one One-Day Workshop</b>	Advance	\$790	\$690	\$360
		Late/On-Site not available		
<b>Package: Conference and one Two-Day Workshop</b>	Advance	\$830	\$720	\$390
		Late/On-Site not available		
<b>Package: Everything (full conference, tutorials, and workshops)</b>	Advance	\$880	\$765	\$420
		Late/On-Site not available		

Become an [ACM member](#)!

## Fess Parker's DoubleTree Resort Hotel

### Notes:

- Reservations must be made online by **September 21, 2006**.
- Cancellation policy: **72 hours** prior to arrival in order to avoid cancellation penalties
- A limited number of rooms for guests with special needs, in accordance with ADA, are available. Please contact the hotel reservations department at 805-564-4333 to confirm your arrangements.
- Rates published are for single/double occupancy. A \$10 extra person charge per person will apply for 3 or more guests sharing the same guestroom.
- There is a 12.09% per room per night tax.
- There is a Credit Card required for this reservation for guarantee purposes. The following credit cards are accepted: American Express, Carte Blanche, Diners Club, VISA, MasterCard, and JCB.
- Check-in is **3:00 pm**. Requests for early check-in will be handled based on hotel availability at the time of check-in. Should you need to confirm an earlier arrival than the hotel's standard check-in time, please contact the hotel directly.

## **ACM MM06 – Nearby Food and Dining**

There are several to each or purchase food within walking distance of the conference hotel. Here is a list of most of the nearby places, plus a couple that are a little farther but (according to the locals) well worth the walk. In addition, there are many more restaurants with a large selection of styles and price ranges a longer walk (or short drive or taxi ride) away, either on within a couple blocks of State Street, the main street of downtown Santa Barbara.

Go to <http://www.santabarbara.com/dining/> for a Santa Barbara dining guide, including restaurant reviews.

### **At the Doubletree Resort Hotel:**

Rodney's Steakhouse  
Café Los Arcos

### **A short walk from the hotel:**

Bistro 1111  
[Hotel Mar Monte]  
1111 E Cabrillo Blvd

East Beach Grill  
[Burgers, pancakes]  
1118 E Cabrillo Blvd

Fresco at the Beach  
[in the Santa Barbara Inn]  
901 E Cabrillo Blvd

Tri-County Produce  
[Farmer's market, snacks, etc.]  
335 S Milpas St

Arturo's Taqueria  
226 S Milpas St

The Habit  
[Burgers]  
216 S Milpas St

Carl's Jr.  
7 S Milpas St

Jack's Famous Bagels  
53 S Milpas St

Your Place  
[Thai]  
22 N Milpas St

Trader Joe's Market  
[Groceries]  
29 S Milpas St

Taqueria El Bajio  
129 N Milpas St

Little Caesar's Pizza  
21 N Milpas St

Cajun Kitchen Café  
217 N Milpas St

La Gloria Taqueria  
336 N Milpas St

Saigon In and Out  
[Vietnamese]  
318 N Milpas St

### **A bit further, but still walkable:**

La Super Rica Taqueria  
[A favorite local place]  
622 N Milpas St

Fish House  
[Nice seafood restaurant]  
101 E Cabrillo Blvd



## **Become an ACM MM 2006 Corporate Supporter!**

The conference organizers warmly invite your support of ACM Multimedia 2006. Corporate support publicizes your organization's interest in and commitment to the multimedia field. An event with international visibility, ACM MM 2006 is the perfect place to inform leaders and students in the field, as well as the interested general public, about your company's activities and products.

Contact conference co-chair **Matthew Turk** or **Klara Nahrstedt** to arrange your company's support of ACM MM 2006:

Email: [mturk@cs.ucsb.edu](mailto:mturk@cs.ucsb.edu) or [klara@cs.uiuc.edu](mailto:klara@cs.uiuc.edu)

The conference recognizes two levels of contribution: Corporate Benefactors and Corporate Sponsors.

### **Sponsors (\$2,000 to 4,999 U.S.)**

In recognition of your support at the Sponsor level:

- Your name and logo will be included on conference announcements and the program.
- Your company's name, logo, brief description, and link will be included in the conference web site, maintained at <http://mmdb.ece.ucsb.edu/acmmm06/>.
- We will announce your company's contribution at the opening and closing sessions of the conference.
- You will have the opportunity to include material and give-always with the conference material for each attendee.
- You will receive one free registration for someone from your company to attend the conference.
- You will be invited to give a demonstration of your choice at the ACM MM 2006 demo night.

### **Benefactors (\$5,000 U.S. or greater contribution)**

Corporate Benefactors receive all the benefits of sponsors as well as the following additional acknowledgements:

- Your name will be listed in all the material as a Benefactor. You will receive special acknowledgement as a benefactor in the conference announcements and publications.
- A conference event, such as a coffee break, luncheon, or reception, will be named after your company. Event naming is based on the amount of the contribution and otherwise on a first-come-first-served basis.
- You will receive one additional free registration for a second individual from your company to attend the conference.

<b>General</b>
<a href="#">Home</a>
<a href="#">News</a>
<a href="#">Sponsors</a>
<b>Contributors</b>
<a href="#">Important Dates</a>
<a href="#">Full Papers</a>
<a href="#">Short Papers</a>
<a href="#">Tutorials</a>
<a href="#">Workshops</a>
<a href="#">Panels</a>
<a href="#">Foundational Sciences</a>
<a href="#">Doctoral Symposium</a>
<a href="#">Open Source Contest</a>
<a href="#">Demos</a>
<a href="#">Interactive Arts Program: Long-Short Papers and Exhibition</a>
<a href="#">Video Program</a>
<a href="#">Paper Preparation and Submission</a>
<b>Organization</b>
<a href="#">Contacts</a>
<a href="#">Organizing Committee</a>
<a href="#">Program Committees</a>
<b>Conference Programs</b>
<a href="#">Main Program</a>
<a href="#">Interactive Art Program</a>
<a href="#">Workshops</a>
<a href="#">Brochure</a>
<b>Participants</b>
<a href="#">Welcome</a>
<a href="#">Registration</a>
<a href="#">Accommodation</a>
<a href="#">Conference Venue</a>
<a href="#">Arrival Instructions</a>
<a href="#">Social Events</a>
<b>Santa Barbara Information</b>
<a href="#">General Information</a>
<a href="#">Things to do in Santa Barbara</a>
<a href="#">Sight Seeing in Santa Barbara</a>
<b>Miscellaneous</b>
<a href="#">Local Arrangement Wiki</a>

## Program Committee for Short Papers

### Content Track

- **Belle Tseng** NEC Labs America
- Tsuhan Chen (CMU)
- Tat-Seng Chua (National University of Singapore)
- Chitra Dorai (IBM T.J. Watson Research Center)
- Andreas Girgensohn (FX Palo Alto Laboratory)
- Mei Han (NEC Labs America)
- Alexander Hauptmann (Carnegie Mellon University)
- Howard Wing Ho Leung (City University of Hong Kong)
- Ching-Yung Lin (IBM T.J. Watson Research Center)
- Simon Moncrieff (Curtin University)
- Chandrasekhar Narayanaswami (IBM T. J. Watson Research Center)
- Yi Wu (Intel Corporation)

### Systems Track

- **Nalini Venkatasubramanian** (University of California, Irvine, Program Co-Chair)
- Ludmilla Cherkasova (Hewlett-Packard Laboratories)
- Magda El Zarki (University of California, Irvine)
- Ahsan Habib (Siemens TTB Center)
- P I Halvorsen (University of Oslo)
- Qi Han (Colorado School of Mines)
- Mohan Kumar (University of Texas Arlington)
- Andreas U. Mauthe (Lancaster University)
- Shivajit Mohapatra (Motorola Research Laboratories)
- Milena V Radenkovic (University of Nottingham)

### Applications Track

- **Brian Bailey** (University of Illinois)
- Brett Adams (Curtin University of Technology)
- John Adcock (FX Palo Alto Laboratory)
- Barbara Barry (MIT Media Lab)
- Matt Cooper (FX Palo Alto Laboratory)
- Berna Erol Ricoh (California Research Center)
- Jim Gemmell (Microsoft Research)
- Vera Goebel (University of Oslo)
- Xian-Sheng Hua (Microsoft Research Asia)
- Andruid Kerne (Texas A&M Interface Ecology Lab)
- Qiong Liu (FX Palo Alto Laboratory)
- Shivajit Mohapatra (Motorola Labs)
- Frank Nack (CWI, Amsterdam)
- Balakrishnan Prabhakaran (University of Texas at Dallas)
- Lloyd Rutledge (CWI, Amsterdam)
- Tele Tan (Curtin University of Technology)
- Zhen When (IBM T.J. Watson Research)
- Pavel Korshunov (National University of Singapore)
- Andrea Cavallaro (Queen Mary, University of London)
- Andreas Schrader (University of Lebeck)
- Nicolas Roussel (LRI & INRIA Futurs)
- Yasuo Aiki (Kobe University)
- Kevin Curran (University of Ulster)
- Mor Naaman (Yahoo! Research Berkeley)



SPECIAL INTEREST GROUP on MULTIMEDIA



Website Administrator: **Gang Wu**