

SIGMM RECORDS

Table of Contents

1 Volume 3, Number 1

1 Editorial

1 An Interview with (Serious) Games Experts

7 Report from the Multimedia Grand Challenge 2010

8 SIGMM Education Column

9 MediaEval Multimedia Benchmark: Looking back on 2010 and looking forward to 2011

10 Report of ACM Multimedia Conference Review Committee

16 Call for Participation in the Multimedia Grand Challenge 2011

18 Award opportunities

18 SIGMM Award for Outstanding PhD Thesis in Multimedia Computing, Communications and Applications

19 SIGMM Award for Outstanding Technical Contributions to Multimedia Computing, Communications and Applications

20 PhD thesis abstracts

20 Andre Miede

20 Beatriz Soret

21 Lin Lin

23 Peter Knees

23 Vineeth N Balasubramanian

25 Event and publication reports

26 Calls for contributions

26 Calls for SIGMM Sponsored and Co-sponsored Events

26 Calls for Events held in cooperation with SIGMM

27 Other multimedia-related Events

28 Job Opportunities

28 Back matter

28 Notice to Authors

29 Impressum



Association for
Computing Machinery

Advancing Computing as a Science & Profession

ISSN 1947-4598
<http://sigmm.org/records>

SIGMM Records

Volume 3, Number 1, March 2011

Editorial

Dear Member of the SIGMM Community,

Welcome to the SIGMM Records in 2011!

This exciting issue features an interview again: you can read about commercial game developers' views on various questions of interest for multimedia researchers. Researchers hear about developers' view on serious games, the role of research communities for their work, and contributions that researchers can make to their work.

Then, you can read about MediaEval. A benchmarking initiative dedicated to evaluating new algorithms for multimedia access and retrieval, it was held in Pisa in 2010. With not only talks but many discussions about the results of the 2010 evaluation tasks, the audience discussed the progress in multimodal approaches to multimedia involving. Look at the article and you can see and hear interviews with MediaEval participants.

Furthermore, five PhD thesis summaries are included in this issue of the Records. The SIGMM Education Column features a course by Professor Wu Ja-Ling, who teaches multimedia security at National Taiwan University. You can read the report of ACM Multimedia Conference Review Committee; whether you intend to bid for ACM Multimedia or are just an attendee, you should read it in detail. The issue contains a report from the ACM Multimedia Grand Challenge 2010 and introduces the tasks that of the Grand Challenge 2011.

Last but definitely not least, you should look at the ACM SIGMM Best PhD thesis award. It's deadline has been extended - so you still have a chance to submit the most impressive multimedia-related thesis that was defended in 2010.

The Editors
Stephan Kopf
Viktor Wendel
Lei Zhang
Wei Tsang Ooi
Carsten Griwodz

An Interview with (Serious) Games Experts

Experts: Dr. Stefan Göbel (TU Darmstadt), Roman Schönsee (Ranj), Dr. Florian Stadlbauer (DECK 13), Avni Yerli (Crytek)

Summaries by Stefan Göbel (TU Darmstadt); interviews performed by Viktor Wendel (TU Darmstadt)

Video games (digital computer games, short: games) are very popular and successful, both as leisure activity and contemporary information and communication medium in the digital age, and as relevant economic factor and prospering market, not only in the creative industries. Games tackle a diversity of research aspects, e.g. Computer Graphics, AI, Storytelling, interfaces and sensorics, authoring and production, usability and user experience or other ICT and multimedia technologies. Game technology and game techniques are broadly used by other application domains apart from pure entertainment as well. Prominent examples in the field of Serious Games (games `more than fun') are games for health, persuasive games, advergaming or games for education and training, for instance in the form of multiplayer online games as tools to support collaborative learning settings.

Recently, numerous game-specific conferences - among others the Game Developers Conference (GDC), Future Play, Foundations of Digital Games (FDG), Games for Health or the GameDays - have been established to discuss underlying technology, methods and concepts of Games and Serious Games. Further, dedicated game aspects are discussed at well-known, high ranked academic conferences such as the ACM conferences on Computer Graphics and interactive techniques (SIGGRAPH) or about human factors in computing systems (CHI), the Artificial Intelligence for Interactive Digital Entertainment Conference (AIIDE), the international conference on computer supported learning (CSCL) or the European conferences on Technology-enhanced learning (EC-TEL) and Game-based learning (ECGBL).

Unfortunately, the Multimedia research community does not consider games and game technology in depth

so far - for instance, at ACM Multimedia 2010, only a few submissions (<10 papers out of 974 publications) dealing with game aspects have been proposed and presented. Therefore - also considering the ACM SIGGAME approach - the aim of this article is to strengthen the (doubtlessly existing) tight link between games (the game industry and game research) and multimedia technologies (the multimedia RTD community) and to encourage scientists and practitioners from both sides to communicate and profit from each other. A set of questions was prepared by the editors and sent to selected, renowned specialists in games and Serious Games in order to encourage and cultivate that process and cooperation. The focus of the interviews is set on (massively) multiplayer online games and its use for training and teaching purposes. We are grateful to the interviewees Roman Schönsee (Ranj), Dr. Florian Stadlbauer (DECK13), and Avni Yerli (Crytek) for their contributions to this article. We also thank Dr. Stefan Göbel for summarizing the responses of the experts enhanced by personal insights.



Roman Schönsee has a varied background in the creative and the research industry. Starting as a music composer and sound designer Roman worked for several years as a account executive in the market research industry. Roman is since 2007 with Ranj Serious Games and worked as Game Producer and Sound Designer on many education, health, marcom and business training projects.

Q1: Nowadays, the games industry represents one of the most prospering markets, not only in the creative industries. What do you think was the most significant change in the games industry during the last 20 years from a research/technology point of view?

Summary The game industry - originating in the non-academic world, without well-established development processes or

(academic) training and qualification programs - brought a set of high-end technologies, relevant not alone for pure entertainment, but also accepted and applied in other (serious games) application domains. The experts point out the broad spectrum and diversity of games and underlying technology, methods and concepts and emphasize the potential of cooperation between academia and industry in that field. For instance, a fruitful cooperation has been taken place in the context of the development of real-time, immersive 3D game environments.

Roman Schönsee

The most significant change is that games became grown-up and the industry is been taken serious. Other industries have adopted gaming and started to understand the strength of gaming technology and game techniques.

Florian Stadlbauer

I do not see one single technology, but rather a lot of small and incrementally further developed innovations. From a technical point of view it is the way how virtual worlds can be created and be shaped allowing players to completely immerse into them. Also how the "uncanny valley" could be overcome. I also see a lot of interesting developments in the field of A.I., or storytelling and arbitration of content, but I would not say that these developments are to be tied up to one single development step in gaming industry. The gaming industry itself evolves in a very iterative fashion. We are currently facing the release of one of the highest budgeted shooters, Crysis 2, which, compared to other shooters which have been released during the last years, has been further developed in many areas, but it is not a completely new invention. There were only few real inventions in the recent years where one could say "that was something never seen before". In how far Wii, Kinect and Wii Move help the (console) platform to reach new target audiences, which then can be utilized in the field of Serious Games, could be another issue.

Avni Yerli

20 years ago, the game companies were mostly solving problems by the process of trial and error and rarely looked for solutions in the academic

world. This is not surprising since both worlds usually had opposite goals in mind. With the introduction of faster hardware, the range of opportunities increased and both worlds saw the value in collaboration. As the game industry was approaching the "uncanny valley", they realized the value in the extensive research that was done in the past tens of years by the academic world and turned there for help. In turn, more and more researchers learned to appreciate the pragmatic approach that the game industry brought and invested more time into real-time solutions while sacrificing some precision and physical correctness. During the past years, the combination of both worlds has generated some amazing results and we expect the collaboration to tighten in the future.



Dr. Florian Stadlbauer, MBR, was born in Heidelberg in 1976. He studied Business Administration at the Universities Frankfurt and Munich and achieved his diploma in 2003. Since 2001 he is founder, partner, and Executive Director of the Deck13 Interactive GmbH. From 2003 to 2005 he is Consultant and Head of Marketing at the Netvolution GmbH. He graduated as "Master of Business Research" (MBR) at the University of Munich in 2006 and achieved a PhD in Business Administration at the Institute for Information Systems and New Media at the Munich University and at York Universities' Schulich School of Business in Toronto/Canada. Dr. Florian Stadlbauer held different courses about management in the games industry at the universities / colleges of Leipzig, Erfurt, Kassel, Darmstadt and Ilmenau. Since 2007 he is a jury member of the Gamesload Newcomer Award of the German Developer Award and a member of the board of the gameareaFRM e.V. since August 2009.

Q2: Which big research steps regarding (massive) multiplayer online games do you expect within the next 5 years?

Summary

The answers of the experts cover several RTD aspects ranging from bandwidth and processing power issues to game concepts such as 'free-to-play' browser and social games or 'seamless play' on multi-platforms anywhere at any time and the use of (M)MOG's in particular application domains, e.g. games for health or educational games. Hereby, according to performance issues and processing-intensive tasks in (M)MOG's, mechanisms of "cloud computing" are considered as promising approaches, which need further research investigations.

Roman Schönsee

I believe the health industry will look more closely into the possibilities to use multiplayer online games for convalescence, patient self-management, research and development. The health industry is changing on a global scale and demands new ways of thinking and working. New technologies and the possibility software combinations that include gaming as a component offer an infinite amount of possibilities that can help to find new solutions to these changes.

Florian Stadlbauer

I think that especially in the field of multiplayer online games, speaking of browser and social games which are being played by the general public and which have quite different numbers and growth rates than even WoW, two elements will become more important. First, the concept of "free-to-play" will be stronger established. Popular games already started to work on a "free-to-play" base. Second, for social and browser games, gameplay experience and game depth will play a more important role. The first generation of that type of games were kind of simple spread sheets regarding their visual representation. This circumstance is changing rapidly towards games which can compete with classical retail games in that aspect.

Avni Yerli

We have recently seen early examples of processor intensive tasks, like graphics rendering, happening in "the cloud" instead of on the player's

computer or gaming console. This allows for the developer to design modern looking games without having to worry about how big part of the target audience having powerful enough hardware to run it. The delay caused by the Internet connection can pose a slight problem for games based on reaction time, but the tactical nature of most massive multiplayer games, and the fact that they are already online, suggests that they can be a perfect fit for this technology. Run your game on your PC or your iPad, play at home or play on the bus - neither platform nor location matter. Another challenge requiring further research is inherent to massive titles - allowing thousands of players to coexist in the same space at the same time without overwhelming the server.



Avni Yerli, Managing Director, Crytek GmbH: After completing a degree in mechanical engineering in 1994, Avni Yerli started his business career as a project manager at a leading Bavarian planning and consulting company. Three years later, he became a Managing Director of a subsidiary of a foreign consulting company and, by early 1998, he was running his own business as an independent consultant. But, being an ardent gamer, it wasn't long before Avni joined his brothers Cevat and Faruk in founding the interactive entertainment development company Crytek in November 1999. As founding member and Managing Director, Avni is responsible for business development and strategic development. He currently holds advisory board positions with GDC Europe and ACGI (All-Russian Conference on Game Industry).

Q3: Which multimedia developments/innovations would you consider useful/helpful for the (massive) multiplayer online games industry in the future?

Summary Promising contributions are seen in the field of Human Computer Interfaces

and multimodal interfaces, addressing the potential of increased dimensions of (multi-touch) displays on (mobile) devices and natural, gesture-based interaction and control of computers and applications (cp. the Wii and Kinect technology provided by Nintendo and Microsoft).

Roman Schönsee

I am looking forward to the development in the area of multi touch as well as alternative ways of controlling computers. I believe we are just at the beginning of a new era and especially in the gaming industry new developments in these technologies will lead to many innovative and exiting game concepts and the way we interact with games in a multiplayer surrounding. The way I see it we go back to a more natural way of moving. Interfaces we are currently using like mouse's, game pads and keyboards force the user to interact in an unnatural way with the computer. New interfaces and technologies to interact with machines will bring gaming to a much broader audience. I think the Wii has already proven this and I also believe that this was just the beginning.

Florian Stadlbauer

I do not see a direct link between multiplayer games and multimedia innovations. A good multiplayer game can be developed using typical technologies, like for example SmartFox for browser games. Multimedia developments have a rather supporting function. I would not say that because I have a 3D technology, my multiplayer game will be better. Multimedia can only indirectly influence that, if at all.

Avni Yerli

The increasing dimensions of the displays on multimedia devices change how they can be used. The larger displays will help a transition of moving towards multi-platform games where you stay connected to the game world all the time, and where the nature of the interaction depends on the shape of the device you currently use.



Stefan Göbel holds a PhD in computer science from TU Darmstadt (TUD) and has long-term experience in Graphic Information Systems, Interactive Digital Storytelling, Edutainment applications and Serious Games. After five years work as researcher at Fraunhofer Institute for Computer Graphics, from 2002 to 2008 he was heading the Digital Storytelling group at the Computer Graphics Center in Darmstadt. In late 2008 he moved to TUD and is heading the Serious Gaming group at the Multimedia Communications Lab. Dr. Göbel is the author of more than 70 publications, associate professor for lectures and courses, initiator and host of the Serious Games conference and science meets business forum GameDays (since 2005), member of different program committees and serves as reviewer for the European Commission and jury member of the Serious Games Award.

Q4: How do you see the role of the serious games industry for the development and acceptance of new multimedia applications (games, simulations) - especially in classroom?

Summary Whilst the games and Serious Games experts agree on a general level and predict positive effects of using (serious) game technology and techniques for training, learning and teaching (in classrooms), on a practical level, the chances 'to make it happen' are seen slightly different based on experience achieved in projects with schools and teachers - which are still very sceptic in the use of games at school. This situation might be improved by the increasing number of 'digital native teachers' and - more relevant - by comprehensive, scientific validated best practice examples and sound evaluation studies, proving the benefit of games and game-based approaches for training and learning in terms

of learning support and competence development. From a technical point of view, the major challenge is to merge gaming and learning methodologies and technologies, e.g. concerning the integration of game engines and learning management systems or the implementation of mechanisms for personalization and adaptation. Second, cost-effective production of high-quality educational games are necessary in order to meet a) the high expectations of learners and teachers in terms of quality in graphics, sound and gameplay respectively pedagogic and didactic aspects and b) the financial constraints of schools. These aspects have been tackled in the scientific community for several years, e.g. the interdisciplinary research project 80Days, funded by the European Commission in FP7 in the field of Technology-enhanced Learning.

Roman
Schönsee

Schools in general hold tight to traditional learning methods. Although new technologies are introduced in schools, their full potential is still not used. Acceptance of serious gaming in the corporate sector will ultimately lead to an increased acceptance in the educational industry. Also a new generation of teachers will lead to an increased usage of technologies and software solutions such as serious games. In this context, a closer cooperation between the serious and non-serious games industry would certainly boost this process. Combined resources and commitment in technology and marketing activities could change a lot and lead to a broad acceptance. I also believe that such a development would not only increase the acceptance in classrooms but the acceptance of gaming in general and therefore being another step in the growing process of the industry.

Florian
Stadlbauer

In my opinion it does not play a role at all. I very much regret that, because I think that the game medium is an excellent instrument to teach content to pupils for which otherwise a lot of persuasiveness would be necessary to get them interested in it. We made that experience with our games at school. Variant approaches show that

playful learning is afflicted with various advantages and I think that schools slowly stop being afraid of games, although it is a very slow process. However, it is an ongoing process and the keyword "media competence" is a topic which gains in importance at school. I think Serious Games will enrich the teaching, not by playing fun games like shooters, but by using games especially developed for learning which loosen the seriousness of the learning situation, thus making learning easier. Though, we do not think that Serious Games are a solution for everything. The teacher will still be necessary, however maybe in some situations he/she takes a new role with the positive side effect of increasing the learning performance. Explanatory work is very important in this context. Also, good games have to be developed which make that possible. Here, industry is responsible because nobody will commission a Serious Game as long as the benefit is unknown. However, if industry is confident then they have to show how these instruments can be used in a meaningful way, i.e. by companion studies. If this happens, I am convinced that 'constant dripping wears away the stone' and such instruments will be deployed more and more.

Avni Yerli

The Serious Games Industry pushes gaming technology to be more practical for learning requirements. Although in many respects, learning tools do feature in gaming, under different names (e.g. multiplayer = group learning, replays = after action review), traditional learning management systems are being integrated into gaming engines to provide harder value for teaching effectiveness and the ability to design training and teaching within games to a greater level of effectiveness. Continuing success stories in the use of games in training should broaden the appeal of gaming. As more game engines are used and developed for training, existing and new techniques will be implemented that increase the effectiveness of such learning programmes. There is little doubt that experience can be of more value than theory - gaming is virtual experience for

learners. As that experience becomes more realistic, the benefits of learning via games will improve and soon will be the norm for most forms of training.

Q5: Do you think that the existing game engines / SDKs are robust enough to be used by non-specialists (for example, think of teachers as authors or game developers)?

Summary

At this point, the situation - and simultaneously relevant parameter about success or failure of game engines and software development kits - is characterised by the discrepancy between complexity/functionality and 'ease-of-use/usability' of authoring tools: Whereas Avni Yerli underlines the possibilities of game editors - offered together with games (e.g. the Sandbox of the Cry Engine) - both Roman Schönsee and Dr. Florian Stadlbauer emphasize the highly varying expectations and (subject matter dependent) necessities of individual teachers and their (typically limited) programming skills to create interactive learning scenarios.

Roman Schönsee

In my experience non specialists like authors and teachers have high demands on the flexibility and technical requirements from SDK's and engines. On the other hand especially these users often lack sufficient technical skills. Therefore it is very difficult for developers to take the needs of the individual SDK and engine users into account and still produce software that is robust and intuitive. When developing SDK's for non-specialists we usually make sure to satisfy these needs and still include enough functionality to ensure enough flexibility. I have seen tools for teachers that were too complex and frustrating for the user due to a poor interaction design. Examples like this increase the reservation towards the use of new technology, especially in the educational industry. Lucky enough in my experience examples like these are only exceptions.

Florian Stadlbauer

No, not at all. When nowadays a game is produced, only about 5% of the content (graphics, sound) can be taken from other games, the other 95% are created specifically for that game. However, that content is not produced

with game engines or editors, but with professional tools like 3D animation software. Those tools can be mastered only by experts. The games which can be "clicked-together" with game editors will always be very simple and not be able to provide a greater gameplay experience. I remember that there were a lot of game editors in the past, but all the games created with them were quite similar. If one could create both racing games and "jump n run" games with an editor, then all the racing games created with the editor would look quite similar and all the "jump n run" games as well. Another example from our company: We created a Serious Game for youths consisting of two parts, one part for youths in class 4 to 5 and another one for class 5 to 6. Regarding the fact that children in class 4 have a very different cognition compared to those in class 6, the requirements to the game are very different. Therefore, the two parts have to be shaped quite differently. A construction kit taking into account all variations for various target groups, thinking of children of different ages, seniors, or disabled people, would be very hard to handle and to control and would require a huge amount of assets.

Avni Yerli

Absolutely, generally modern game engines have grown up with the editors released with PC games for "Modding" communities. These are groups of amateur developers, often with no experience in game development, who teach themselves to use the tools for gaming very rapidly. Game Engine tools are also developed for modern, multi-discipline teams, aiming therefore at users with very different levels of skill. Of course, some processes and features of engines are complex and would still require specialist knowledge - for example, in animating a character in engine, but others are truly simple to use.

Q6: Summary - how can members of SIG Multimedia contribute to that research?

Summary The industrial experts expressly appreciate the involvement and cooperation with scientists settled in the multimedia community, especially with regard to the (prototypical)

development, implementation and testing of new, innovative concepts. In general, the situation is characterised by the typical chicken-and-egg problem: Industry needs proof of concept and evaluation studies proofing the (positive) effects (e.g. learning). Science is ready to deliver the proof of concept but in order to do so funding is required as well as realistic content and applied procedures and processes. Finally, it can be stated that cooperation between industry and science is necessary, so "let's work together".

Roman Schönsee

SIGMM offers a forum for those industries that are involved in this development. There will be a demand on specialized hardware, software solutions and research in terms of usage and medical benefit. Having a good network is a key factor in this development.

Florian Stadlbauer

I think that it is very important to clarify from a scientific point of view where learning results can be improved by a correct use of Serious Games. Another important field where science can help better than industry is the support of prototypical forms of innovative game-base learning methods. I think these are two very essential components which industry can pick up and turn into products.

Report from the Multimedia Grand Challenge 2010

Authors: Cees G.M. Snoek and Malcolm Slaney
by Cees Snoek

The recent ACM Multimedia conference in Florence, Italy covered the second edition of the Multimedia Grand Challenge. An event that is designed to bring commercial industry needs to the attention of researchers. Seven industrial partners led the Multimedia Grand Challenge by identifying issues they think are important to their business and worth further study. The industrial partners pointed out that addressing these challenges would open up new business opportunities and create a richer experience

for their users. In addition, they said they hoped that describing real-world challenges would allow researchers to focus on projects that have a better chance of success in the marketplace. The following 10 challenges were identified in the 2010 edition of the Multimedia Grand Challenge:

- CeWe: Automatic theme-identification of photo sets. Provide a user with a selection of photo-book styles accommodating the user's preferences for the images in the set, and their preferred structure.
- Google: Genre classification for video. Automatically classify user-generated videos and their metadata into genres.
- HP: High-impact visual communication. Create a high-impact collage (and its description) that uses a set of photos to convey information across cultural boundaries.
- Google: Interactive searching in personal diaries. Develop good schema, algorithms, and user interfaces that make multimedia diaries accessible.
- Radvision: Multimedia-content adaptation. Adapt, in real-time, the same multimedia content for different receiving devices, in a way that is perceptually optimal for users.
- Yahoo!: Novel image-understanding. Develop methods that move beyond simple image classification.
- Nokia: Photo location and orientation. Derive the exact camera poses (location and orientation) for photos that are lacking location annotation.
- Yahoo!: Robust video segmentation. Develop algorithms to automatically create narrative themes for a given video and present the content to end users in a search-engine experience.
- 3DLife: Sports-activity analysis in camera networks. Identify the limits in terms of 2D and 3D data extraction from a low-cost camera network for sports, such as tennis matches.
- Radvision: Videoconferencing experience. Develop new technologies and ideas to surpass the in-person meeting experience.

Together, the problems defined for the Multimedia Grand Challenge cover the full range of multimedia research. The partners contributed their problems, prize money, and data sets to the Multimedia Grand Challenge. Most importantly, they contributed judges, many of whom are directly responsible for multimedia products. In the end, 18 contributions were accepted that together covered all 10 industry problems.

All finalists for the Multimedia Grand Challenge presented their work to the panel of judges in a

conference-wide event. Each finalist had three minutes to present the idea in a format that mixed an elevator pitch with a show like American Idol. After an energetic session led by master of ceremony David Ayman Shamma, the industry partners decided on the three winners. First prize of US\$1,500 dollars went to Jana Machajdik, Allan Hanbury, and Julian Stöttinger who are affiliated with Vienna University of Technology and the Information Retrieval Facility, both in Austria. Machajdik presented their solution to the challenge on novel image-understanding methods. The winning idea is a system able to search for images with a requested emotional effect-such as images being cute, fearful, or sad. The second prize of US\$1,000 dollars went to Wei Song, Dian Tjondronegoro, and Ivan Himawan who are all affiliated with Queensland University of Technology, Australia. Tjondronegoro presented their solution to the challenge on multimedia content adaptation. The third prize of US\$500 dollars went to Julien Law-To, Gregory Grefenstette, Jean-Luc Gauvain, Guillaume Gravier, Lori Lamel, and Julien Despres who are affiliated with Exalead, LIMSI CNRS, and Vecsys Research, all in France. Law-To presented their solution to the robust video-segmentation challenge. While not all participants won prize money, everyone did go home with the message that solving real-world industry problems can be as challenging as research.

There will be another Multimedia Grand Challenge at ACM Multimedia 2011 in Scottsdale, Arizona. Hope to meet you there.

SIGMM Education Column

Authors: Wei Tsang Ooi and Wu Ja-Ling
URL: <http://www.sigmm.org/Education>
by Wei Tsang Ooi

Welcome to the SIGMM Educational Column, where we highlight interesting and noteworthy multimedia educational activities around the SIGMM community. In this issue's column, we wish to highlight a course on Multimedia Security taught in the National Taiwan University (NTU), by Professor Wu Ja-Ling.

Ensuring proper protection of digital media content is critical to the success of the modern media industry. Yet, only few existing courses focus on teaching the fundamental and the state-of-the-art in this topic. The course by Professor Wu is one of a few such courses.

The Multimedia Security course at NTU has a long history. The course was started back in 2000,

to fill in a gap in the multimedia-related course-maps of EECS college at NTU. Originally titled "Information Techniques for Intellectual Property Right Protection," the course is based on the reference book "Digital Watermarking" by Cox, Miller, and Bloom (Morgan Kaufmann Publisher, 1999) and covered digital watermarking, digital fingerprinting, steganography, and DRM.

The course syllabus expanded in 2007 to include information forensics and biometrics recognition. The course was subsequently renamed to "Multimedia Security."

The current (2011) incarnation of the course covers four major topics: (i) digital watermarking and digital fingerprinting; (ii) steganography and data hiding techniques, (iii) biometric-based security and secure biometric encryption; and (iv) privacy preserving multimedia encryption and search. The course is designed around a set of selected recent papers on these topics.

Students taking the class are required to implement one digital watermarking scheme and one data hiding scheme associated with the papers they read, and gave an oral presentation on the associated techniques. Students are also expected to write a survey paper on one chosen topic related to multimedia security, and complete a team-based (up to 3 person), research-oriented, final project. The students from the course have been consistently producing top quality projects. Some of the final projects have been extended and published in venues such as IEEE Transactions on Information Forensics and Security, Information Sciences, and IEEE Transactions on Multimedia.

The course is taken by about 20 graduate students every year, and is successful in piquing their interest in the topic of multimedia security. Up to now, there are at least six graduate students who have taken the course and graduated with a Ph.D. degree with a thesis on the topic of multimedia security. Feedback from students on the course has been positive, and the course has been consistently rated with a rating above 4.4 on a 5-point evaluation system.

The detailed syllabus, reading list, references, and homework for the 2010 version of the course can be found at the course website: <http://www.cmlab.csie.ntu.edu.tw/~ipr/mmsec2010/>

If you are teaching an multimedia-related course, we would like to list your course under the SIGMM Educational Portal (<http://www.sigmm.org/Education>). Please email the SIGMM Educational Committee at education@sigmm.org.

MediaEval Multimedia Benchmark: Looking back on 2010 and looking forward to 2011

Authors: Martha Larson
URL: <http://www.multimediaeval.org/>
by Martha Larson

MediaEval <http://www.multimediaeval.org> is a benchmarking initiative dedicated to evaluating new algorithms for multimedia access and retrieval. MediaEval puts emphasis on context and topic of video and aims at moving beyond what is explicitly depicted in the visual channel towards overall meaning and relevance to users. It attracts participants who are interested in multimodal approaches to multimedia involving, e.g., speech recognition, multimedia content analysis, user-contributed information (tags, tweets), viewer affective response, social networks, temporal and geo-coordinates. MediaEval was established in 2008 as VideoCLEF, a track of CLEF, and in 2010 became an independent benchmarking initiative.



Gerald Friedland (photo by Jaeyoung Choi)

The MediaEval 2010 workshop was held just before ACM Multimedia 2010 on Sunday, October 24, 2010 in a medieval convent in Pisa, Italy. The workshop was attended by 29 participants from nine countries and three continents. The day was dedicated to talks and discussion of the results that were achieved on MediaEval 2010 tasks. These results comprised algorithms that were developed for automatic annotation of video: prediction of geo-location, of user-contributed tags, of subject labels and of viewer-reported boredom.



Group photo from MediaEval 2010 (photo by Jaeyoung Choi)

The talks revealed the value of user-contributed metadata for automatic tag prediction. Co-occurrence information and information derived from external sources such as Wikipedia proved productive. Features derived from the visual and audio channel are also very useful, but remain more challenging to exploit. It will be the continued focus of future MediaEval benchmarks to encourage development of the most useful combinations of diverse and multimodal information sources.



Olivier Van Laere (photo by Cynthia Liem)

Highlights of MediaEval 2010 results will be presented in the special session "Automatic Tagging and Geo-Tagging in Video Collections and Communities" that will be held at ICMR 2011 (<http://www.icmr2011.org>), the First ACM International Conference on Multimedia Retrieval. ICMR 2011 will be held 17-20 April 2011 in Trento, Italy.

YouTube video: <http://www.youtube.com/embed/z1bjXwxkgBs>

Sign-up for MediaEval 2011 is open now through 31 May 2011 on the MediaEval website <http://www.multimediaeval.org>. The website also gives the full schedule for MediaEval 2011 and descriptions of the tasks that are offered this year. Sign-up for MediaEval is open to any group that wishes to participate. Participants receive task descriptions, data sets and supporting resources and develop algorithms that generate task results. When they receive the result evaluation from the MediaEval organizers, they write up a working notes paper and attend the MediaEval 2011 workshop. The MediaEval 2011 workshop is an official satellite event of Interspeech 2011. <http://www.interspeech2011.org> and will be held 1&2 September 2011, again in Pisa, Italy.

YouTube video: http://www.youtube.com/embed/sVYiV-4_UNE

MediaEval is coordinated by the EU FP7 PetaMedia Network of Excellence <http://www.petamedia.eu> and also by the ICT Labs of EIT <http://eit.ictlabs.eu/> and is made possible by the many projects, institutions and researchers that contribute to the organization of the individual tasks.

Report of ACM Multimedia Conference Review Committee

Authors: Chua Tat-Seng, Susanne Boll, Chang-Wen Chen, Balakrishnan Prabhakaran, Nicu Sebe, Hari Sundaram, Qi Tian
URL: <http://www.sigmm.org/>
by Mohan Kankanhalli

Highlights of Key Review Recommendations

1. ACM Multimedia Conference Steering Committee to play an active role in ensuring the high quality of the conferences, including the appointment of general and technical program co-chairs.

2. Stricter criteria for choice of Technical Program Committee (TPC) members.
3. Change the paper review system from the current track-based to area-based.
4. Increase the number of accepted quality papers from the current 15-17% to around 20%.
5. To encourage and facilitate fuller and well-rounded conference participation experience.
6. To conduct such major review once in five or six years.

1. Committee

In December 2009, ACM SIGMM Director of Conferences (Mohan Kankanhalli) formed a new committee to look into the organization and review of technical papers for ACM Multimedia and other conferences under the umbrella of ACM SIGMM.

1. Committee Chair: Chua Tat-Seng
2. Sub-committee on Conference Organization:
 - a. Boll, Susanne (Group Chair)
 - b. Chen, Chang-Wen
 - c. Prabhakaran, Balakrishnan
3. Sub-committee on Paper Review Process:
 - a. Sebe, Nicu (group Chair)
 - b. Sundaram, Hari
 - c. Tian, Qi

2. Terms of Reference

The committee will review two aspects of ACM Multimedia and related conferences:

1. the conference organization; and
2. the procedures for the management and review of papers for the conferences.

The Committee will only focus on the ACM Multimedia conference in this report.

3. Conference Organization

3.1 Establishment of the ACM Multimedia Conference Steering Committee (AMCSC)

We recommend that an ACM Multimedia Conference Steering Committee (AMCSC) be formed to oversee the planning and organization of future SIGMM conferences, including ACM Multimedia. The roles of AMCSC are to oversee the proper selection of key appointment holders and smooth execution for each conference, as well as to ensure the adherence of best practices as prescribed in the conference guidelines.

The AMCSC shall be chaired by the ACM SIGMM Director of Conferences, with members coming from

Chairs of recent ACM SIGMM Conferences. Our recommendation is for AMCSC to have five more members, to be selected from the main organizers of the three major ACM SIGMM conferences, including the ACM Multimedia. The maximum period of membership for each member is two terms of two years, with one-third of members being retired every two years.

3.2 Co-Chairs of the Organization and Technical Program Committee

These are important positions that represent the face and provide direction to the community. They should be filled by people who are well respected by the community, and who have strong interest in multimedia and are active participants of the conference.

One important gauge of the maturity of a conference is to look at the number of past chairs who are still active participants of the conference. This is in comparison with other equivalent top conferences such as ACM SIGIR, CVPR or ICCV.

The ACM Multimedia Conference Steering Committee (AMCSC) should take active role in the selection and appointment of these positions for the SIGMM conferences.

The General Chairs (and Technical Program Co-Chairs) should preferably not be involved again in same capacity with future ACM Multimedia conferences within the next four years. This is to allow for grooming of new leaders for the community.

The General co-Chairs need to make sure that program co-chairs for different tracks are working towards coordinated and relatively uniform milestones in the review process to avoid excessive delay of review procedure in any of the tracks. Mitigation measures need to be designed and properly taken by the Technical Program co-Chairs to avoid unexpected consequences during the review process.

3.3 Workshops

The role of workshops is to offer a forum for discussions of emerging and specialized topic, and should remain so.

The topics of workshop should be distinctive from the focuses of the main conference. In general, workshops should be used to nurture and grow new and emerging areas that do not currently have large participations. Once the topic of a workshop has matured, it should be transferred as a new area into the main conference.

Given the above role, we recommend that the workshop should not be allowed to run for more than 3 consecutive

years, and there should be no permanent "standing" workshops.

The workshops should be selected from competitive proposals with an acceptance rate of 50% or less and, under normal circumstance.

There should not be more than TEN (10) workshops per year.

3.4 Conference Registration and Participation

The aim of the conference is to encourage full and well-rounded conference experience for as many people as possible.

Package Registration:

- We should institute package registrations to encourage people to participate in all activities of the conference (including main conference sessions, tutorials and workshops) as part of the overall conference experience.
- This allows us to bring workshop participants to the conference and vice versa.
- It will also encourage more people to attend tutorials.

Student Participations:

- We should encourage student participation in all activities as they are the future of our community.
- We should provide good student participation package at affordable rate to encourage students to attend all activities (main conference sessions, workshops and tutorials)
- We should provide more student travel grants - SIGMM should help to coordinate this.
- The student registration package should include conference banquet, if it is available, in order to make the students feel welcome as part of the community.

3.5 The Doctoral Symposium Program

The SIG Multimedia executive committee and the conference will work towards mentoring and training of the next generation of multimedia researchers.

The current doctoral symposium program at the conference suffers from weak participation, partly because it is considered to be less prestigious by many faculty members who mentor the students. They are thus reluctant to send their students to attend only this part of the program.

The committee recommends that each year, the doctoral symposium program be tied to the SIG Multimedia dissertation award. We recommend that the top three candidates be invited to present their work.

This can be presented as a special session, with the award being given at the main awards ceremony. If possible, the travel and boarding expenses for the three students should be covered by the conference.

3.6 Archival of Conference Experience

Archives record past practices and experience, and provide valuable institutional memory to enable us to have a good understanding and maintenance of the best practices.

We should archive what is important, including, organizing committee, PC committee, participation list with number of participants in all activities, technical program, budget, sponsorship leads, and lessons learnt, etc.

In addition, video tapings of key note talks, award presentations, and other memorable conference activities should be archived too.

The "lessons learnt" and past best practices should be integrated into the "conference guidelines", to be maintained by SIGMM.

SIGMM should ensure that the conference chairs provide proper archival of key information within a certain period after the conclusion of the conference. The Archives could be made available to all SIGMMM members.

Other Implementation Details:

3.7 Conference Venue and Social Media Experience

The venue of the Conference needs to be carefully selected based on a number of criteria in order to make the conference a great meeting experience of the community. In particular, the cost of attending conference should not be so prohibitive that it discourages participation by students.

Preferably, conferences should be held in a university, where the cost is generally lower and the atmosphere is more conducive to technical conference as compared to top business hotels.

The conference should provide good Internet access to all participants, as well as fast and reliable Internet access for technical presentations and demonstrations.

Conference organizers are expected to provide social media experience for attendees such as Facebook, Twitter etc, where the attendees can share comments regarding technical paper presentations and demos, etc. This will be especially helpful given that the question-answering times are rather limited after the paper presentations. These social media blogs will be done

through a conference registration id to avoid flames and spams.

3.8 Food & Beverage

Conference lunch is an important and integral part of the conference and should be provided as part of the conference participation. Such lunches allow people to meet and network in a neutral environment. Without providing lunches, people tend to gather among their close groups to look for places for lunch.

All conference activities should be co-located in a main common area for teas and gatherings in order to ensure coherent conference atmosphere.

Drinks, tea and coffee should be available at all times of the conference. This avoids long queues during the coffee breaks and provides a conducive environment for those who want to continue their discussions even after the breaks.

The conference should try to integrate rather than segregate those with special dietary requirements, such as the vegetarians.

3.9 Student Volunteer Program

With the growing size of conference, local organizers might need to draw on the help of student volunteers to support the conference. Hence, the conference should introduce a student volunteer program as with other conferences such as SIGCHI. Student volunteers will receive free conference registration and in return they will commit to perform a certain hours of services to the conference.

Student volunteer program will also help to encourage greater student participation and integration, allowing local and foreign students to form valuable links for the benefit of the community in the long run.

4. Review Process of Technical Papers

4.1 Paper Review System

The renewal of topics and emphasis of the conference should be done in a more systematic manner, with greater involvement of the ACM Multimedia Conference Steering Committee.

ACM Multimedia should break with the current track-based system and adopt the area-based system. Instead of 4 broad tracks, which tend to be overlapping of often confusing to authors and reviewers, we should adopt, say, 10 or more focused areas each year. Each area will cover several closely related topics. It permits more flexibility in renewal of areas each year, and allows the submission and review of papers in each area to be more focused.

As a rule, we should actively look into growing one or more new areas each year, while consolidating existing areas. This includes dropping areas from previous years that have become less active with lower importance.

The conference chairs could call for and decide on new areas to be included for each year just before the call-for-papers is published. This process will involve the community in shaping the future of ACM Multimedia conferences.

There should be better coordination between the technical program committees of the main conference and the workshops on what new areas are to be introduced and in what forum. In general, new or emerging areas should be started as workshops before being integrated into the main conference.

4.2 Acceptance of More Quality Papers

Many ACM conferences are experiencing greater competition in paper acceptance, along with general increase in the number of researchers and submissions in many fields, including multimedia. We therefore recommend that ACM Multimedia Conference should accept more "quality" papers, raising from the current guideline of 15-17% to around 20%.

It will signal to the community that we are now adopting a more inclusive approach rather than an exclusive attitude to accepting quality works. It will help to reduce some randomness in paper acceptance among the high quality papers when the acceptance ratio is too low (at around 15%), and there will be fewer unhappy authors.

Higher acceptance, if done with good quality control, will also allow for a greater variety of good papers to be presented, and hence help the area to grow.

Most importantly, by accepting more papers, it will go towards reducing the number of good (but rejected) papers in circulation in the conference system. It will alleviate the growth in the number of paper submitted to all the related conferences, and hence reduce the overall review loads for researchers in the field. In this respect, ACM MM should take the lead.

4.3 Structure of Technical Program Committee (TPC)

4.3.1 Two-Tier Technical Program Committee

The adoption of area-based system means the implementation of 2-tier TPC, with the top-tier committee comprises Senior PC members (or Area Chairs) covering key areas to be emphasized in the conference; and a second tier committee (TPC Members) responsible for reviewing the papers in each areas. It is compulsory for Area Chairs to attend the TPC meeting.

4.3.2 Qualification of TPC members

Area Chairs should be established and active members of the community, with good publication records in the conference and in-depth understanding of the review process.

Regular TPC members should be active participants of the conference, who preferably have published ACM Multimedia papers before. This is where it is good to involve young researchers into the paper review process.

Balance of the PC: We should try to achieve a good balance among the TPC members with regards to the community, expertise, age, gender and nationality.

4.3.3 Roles of TPC members

Area Chairs recommend and recruit normal TPC members for each area; play mentoring role; ensure proper, thorough and timely review of each paper; summarize review of paper; and present the key aspects of papers at TPC meeting.

Area Chairs do not need to read all papers, but they must have good understanding of the merits of all papers, especially those borderline papers.

TPC members review the assigned papers and participate in the online discussion stage of paper review to help formulate the final decision for each paper.

4.3.4 Recruitment of PC members

Area Chairs are recommended by the TPC and General co-Chairs, with the approval of the ACM Multimedia Conference Steering Committee. Depending on the projected number of submitted papers, each area will recruit one or more Area Chairs.

TPC Members are recommended by a combination of TPC-co-Chairs and the Area Chair of respective areas.

4.4 Technical Contents of the Conference

4.4.1 Key Technical Areas to Focus

The areas should cover all major areas to be emphasized and nurtured in the field. As a good practice, we should consider the introduction of one or more new areas (though some may be small initially) each year.

4.4.2 Variety of Technical Programs

With the increased in number of submissions and level of participations, it is expected that the number of papers accepted, both long and short papers, will increase steadily over the years. While the primary

mode of presentation for long papers has traditionally been oral, it has become increasingly infeasible for oral presentation with the increased number of long papers accepted, unless more days or more tracks are added to the conference. Depending on the content, levels of interest, constraints of conference facilities and suitability for oral presentation, TPC committee will make decisions on which accepted long papers are to be oral and which are to be presented as posters. This possibility should be made clear to potential authors in the call-for-papers notices.

The committee feels that we should keep the current short paper track, which serves as a forum for papers describing work-in-progress or innovative new (breaking) ideas. It should not be viewed as a track for lower quality or rejected long papers. This view should be instituted in the review criteria.

The TPC will also make recommendations on whether a good-quality rejected long paper should be considered for short paper session (as it describes new innovative ideas) or workshops (as it touches on areas relevant to respective workshops).

For poster sessions (both long and short), it is recommended that an elevator pitch session should preferably be organized for authors to present their ideas within a short few- minute duration. This will allow the authors to present key ideas of their posters and audience to have quick overviews before deciding on which posters to explore in more details. Such sessions, which have increasingly becoming popular, also provide dynamism and life to the technical conference.

4.4.3 Relation with workshops

One idea tried at ACM MM 2010 is that during the TPC meeting, the committee will recommend some of the good but rejected long papers to be redirected to "relevant" workshops for consideration. It is a good idea to have some good papers to be presented in more focused meetings and hence increase the coverage of the conference.

- This means that the paper submission deadlines for Workshops should be set after the notification dates of the long papers in the main conference.
- However, in order to ensure that it works effectively, it might be necessary for overall workshop co-chairs of the conference to attend the TPC meeting.

4.5 Paper Review Process

4.5.1 Anonymity in Review

The norm of the review process should be double blind, where the reviewers and authors do not know each

other. We should not switch back to the single-blind review system at the discretion of conference organizer as was done in the past.

The practice of most top conferences is to disallow key decision makers such as the General co-Chairs, TPC co-Chairs, and Workshop co-Chairs to submit papers to their respective forums. This is to avoid conflicts of interests, both real and perceived. Since the conference is organized yearly, this should not present a big problem to key conference organizers.

4.5.2 Preferences/Bids

Area Chairs and TPC members should be given the option to state their preferences on papers that they could and would like to review. TPC co-chairs and Area Chairs should work to ensure the relevance of papers submitted to respective areas, and reject those papers that are clearly out of scope or of low quality (say too short and are badly formatted etc) to reduce the overall review loads.

4.5.3 Number of reviews needed

At least 3 reviews and a meta-review (to be done by the respective area chair) should be done for each paper. More reviews should be solicited if the recommendations are conflicting.

4.5.4 Level of details and quality of reviews:

As the aim of the review is to constructively criticize and provide guidance to help the authors improve their papers, hence sufficient details should be provided in the review. We should disallow reviews that are one-liner or are too general. To help new reviewers, guidelines and samples of good reviews should be provided. For example, see the excellent guideline provided in SIGCHI. (<http://chi2011.org/authors/chi-review-process.html>)

Length of the review must be substantial, and one approach is to prescribe a minimum of 200 words for review comments.

The review form should offer different sections that help the reviewer structure the comments and area chairs evaluate the overall reviews. Hence the review form should include sections on: Summary, Strengths and Weaknesses of the Paper (say at least 3 points each), Novelty and Contributions to the Community; and Expertise and Confidence of the Reviewer, etc.

4.5.5 Review loads

Area Chairs should handle 20 papers. As stated earlier, the role of area chairs is to manage the review process to ensure that all papers are properly reviewed and

commented. They are not expected to read most of the papers.

TPC members / reviewer should be expected to review up to 10 papers.

4.5.6 Rebuttal by authors on the reviews:

This has become a common (best) practice in many top technical conferences, and many authors do expect the opportunity to comment on the reviews. The rebuttal should take place during the paper discussion period and before the face-to-face TPC meeting.

4.6 Conduct of TPC Meeting

It should be compulsory for all the area chairs to attend the face-to-face or other form of synchronous TPC meetings. This ensures that all papers, especially the borderline papers, have someone who knows the contents of the paper well to argue for or against the paper.

The "compulsory" attendance of TPC meeting should be made clear during the invitation stage. Only those who are able to attend the TPC meeting should agree to be area chairs.

Optionally, as is done in the computer vision community, the conference may allocate budget to (partially) support the attendance of TPC meetings by the area chairs. This will make it easier for them to commit to attend the TPC meeting. One alternative is to organize the meeting in conjunction with one of the SIGMM conferences (such as the new ACM ICMR (International Conference on Multimedia Retrieval)), so as to facilitate the attendance of most area chairs.

Face-to-face vs. synchronous online meeting or Face-to-face meeting:

- The main problems of organizing such meetings are the cost of attendance, high possibility of absentees for various reasons; and may not be as effective as the majority of people are in the "wrong" time-zones. However, it provides a better feel of the sentiment, trends and overall quality of the review process.
- Virtual online meeting: The main advantages are the low cost; ease of ensuring all are present; and could conduct meeting in stages over a longer period with participation also by some regular TPC members. However, there is often a sense of loss as to the overall picture of the review process. This has been attempted in SIGIR and has not been found to be particularly effective and satisfactory.
- In practice, we should adopt the face-to-face system, but encourage TPC co-chairs to have separate virtual on-line meetings with various area chairs to iron out most problem papers before the main meeting.

4.7 Recognition of TPC Participations

To encourage ownership in reviews, we further recommend that the name of Area Chair be appeared in the first page of the accepted paper handled by him/her.

The conference should show appreciation and recognition to the best reviewers/PC members during the conference, through prizes presented during the banquet.

4.8 Keeping Track of Best practices

Summary of what works and what needs to be improved.

Handling of double submission - outright rejection of paper and banning the authors for 1 year (SIGIR's practice).

Black-listing of bad reviewers.

Rewards for good reviewers.

Call for Participation in the Multimedia Grand Challenge 2011

Authors: Gerald Friedland, Yohan Jin
by Gerald Friedland

Deadline: August 6th, 2011.

What problems do Yahoo, HP, Nokia, Technicolor, 3DLife, and other companies see in the future of multimedia?

The Multimedia Grand Challenge is a set of problems proposed by industry leaders, geared to engage the Multimedia research community in solving relevant, interesting, and challenging questions about the industry's 2-5 year horizon for multimedia. The Grand Challenge was first presented as part of ACM Multimedia 2009 and has established itself as a prestigious competition in the multimedia community.

This year's conference will continue the tradition with both ongoing as well as brand new challenges, including:

- **HP Challenge: High Impact Visual Communication**

Images can serve as a powerful communications vehicle, conveying a wealth of information as well as emotional impact. The color, composition, content, lighting and sharpness of an image all

contribute to a viewer's response to that image, and relative placement, scaling and orientation of a group of images in a collage adds an additional layer of richness and meaning to a page. These characteristics are used extensively by professionals on web sites, magazine covers and printed advertisements to draw attention, communicate a message and leave a lasting emotional impression.

Because images hold such power, people like to use their photos to tell their own stories. However, their end result often falls short since many people lack the skill and intuitive understanding needed to create a coherent visual story from their photos. In addition, a picture is worth a thousand words. How do we create a high impact picture that can convey information across cultural boundaries and find a thousand words that best describe such a picture? This grand challenge is to find a solution which can create a collage and generate a textual description that tells the story of the set of photos.

The system starts with a digital photo collection, such as photos taken during a vacation. It then analyzes the collection automatically using information from multiple sources such as image analysis, internet data sources, and EXIF tags. The result of the analysis is used to create the most appealing collage picture that best represents the original collection. In addition, it is also used to generate a description of the collage picture.

- **Technicolor Challenge: Precise Event Recognition and Description from Video Excerpts**

Visual search that aims at retrieving information on a specific object, person or place in a photo has recently witnessed a surge of activity in both academic and industrial worlds. It is for instance now possible to get precise information on a painting, a monument or a book by shooting it with a mobile phone. Such impressive systems rely on mature image description and matching technologies coupled with application-dependent databases of annotated images. They require however to circumscribe drastically the visual world to be queried (e.g., only paintings) and to organize a comprehensive information database accordingly.

This Grand Challenge aims at exploring tools to extend this search paradigm in different directions: (1) The query is an excerpt from a public event's video, that is a small video and/or audio chunk from a longer coverage; (2) There is no strong contextual prior, thus ruling out the use of a specialized database designed on purpose; (3) Sought output is a precise textual description of the audiovisual query "scene".

- **Nokia Challenge: Visual Landmark Recognition**

Mobile devices provide ubiquitous access to internet, and so to almost unlimited amounts of data. Finding the information relevant to you can be a time consuming task in itself. In the modern fast paced life, when you are on the go, coming up with suitable query terms and typing them on a virtual touch keyboard is simply too slow. Image recognition has been widely recognized as a potential novel way of accessing data relevant to your immediate surroundings: snap a picture of something and the system tells you about it.

Nokia and NAVTEQ together have created a dataset of street view data where individual buildings are identified. The dataset consists of 150k panoramic images aligned with a 3D city model consisting of 14k buildings obtained from footprint and elevation data. The images were labeled by projecting the 3D model into the panoramic images, computing visibility, and recording the identities of visible buildings.

- Yahoo! Video Challenge: Robust Automatic Segmentation of Video According to Narrative Themes

Video search today relies mostly on textual metadata that is associated with the video in terms of title, tags or surrounding page-text. This approach falls severely short by ignoring the richness of information within the video medium; an engine should ideally use this information to help a user search and navigate content. As video content explodes and user attention spans shrink, a next generation video search engine needs to provide users with the ability to search for sections within a video; allow users consume bits and pieces of a video that would be of interest to them; and let the users kill time during lunch breaks in creative ways. In addition, instead of offering just one thumbnail as representation for a whole video, it would be great to be able to partition a video into its constituent narrative themes and allow users to navigate through a video on a more granular level with better video surrogates.

The challenge to researchers in the multi-media community is to develop methods, techniques, and algorithms to automatically generate narrative themes for a given video, as well as present the content in an easy-to-consume manner to end-users in a search engine experience. Naturally, the themes that emerge depend entirely on the video itself - so the methods / algorithms have to be generic. Still, there could be approaches developed for certain types and genres of videos. For instance, one approach could be employed for sitcoms, sports content could have another, educational content could have another, etc.

- Yahoo! Image Challenge: Novel Image Understanding

There are over 200 billion images on the internet today and this collection continues to grow by leaps and bounds. Image search engines often only surface a portion of those images and often rely on the text surrounding an image on a webpage, or the image file's name. With the growing number of images on the Internet it is important to have the ability to organize and surface the images in the most efficient, meaningful way possible so that better images can be shown to searchers.

We want to move beyond simple image classification. Textual tags associated with an image often tell us that there is a tiger in an image. Not all images are labeled this way, but there are more than enough on any one subject to fill a search-result page.

People come to image-search engines for many reasons. Users type an average of 2.2 words, but their underlying request is much more subtle, often representing an information or entertainment need that would normally require a much longer and deeper query.

We need novel and useful ways to organize and structure image content. Can we sort celebrity pictures by their subject's age when the photo was taken? Or by their hair style? Can we discover how a logo has evolved over time? Can we organize pictures by their geographic location or the type of object? There are many ways to organize photos. What are the ways that are not obvious? What can we do better than we can do today?

Users would like a better fit between their information and entertainment requests and the content returned by a search. How can we better organize multimedia content to fit user's needs and desires?

- 3Dlife Challenge: Realistic Interaction in Online Virtual Environments

This challenge calls for demonstrations of technologies that support real-time realistic interaction between humans in online virtual environments. This includes approaches for 3D signal processing, computer graphics, human computer interaction and human factors. To this end, we propose a scenario for online interaction and provide a data set around this to support investigation and demonstrations of various technical components.

Consider an online dance class provided by an expert Salsa dancer teacher to be delivered via the web. The teacher will perform the class with all movements captured by a state of the art optical motion capture system. The resulting motion data will be used to animate a realistic avatar of the teacher

in an online virtual ballet studio. Students attending the online master-class will do so by manifesting their own individual avatar in the virtual dance studio. The real-time animation of each student's avatar will be driven by whatever 3D capture technology is available to him/her. This could be captured via visual sensing techniques using a single camera, a camera network, wearable inertial motion sensing, or recent gaming controllers such as the Nintendo Wii or the Microsoft Kinect. The animation of the student's avatar in the virtual space will be real-time and realistically rendered, subject to the granularity of representation and interaction available from each capture mechanism.

We therefore call for submissions of contributions to the ACM Multimedia 2011 Grand Challenge track.

The submissions should:

- Significantly address one of the challenges posted on the Grand Challenge web site.
- Depict working, presentable systems or demos.
- Describe why the system presents a novel and interesting solution .

Preference is given to results that are reproducible by the research community, e.g. where the data and the source code is made publicly available.

The submissions (4 pages) should be formatted according to ACM MM formatting guidelines. Based on the submission, the finalists will be selected by a committee consisting of academia and industry representatives. Finalist submissions will be published in the proceedings and presented in a special event during the ACM Multimedia 2011 conference in Scottsdale, AZ (USA). At the conference, the finalists will introduce their contribution shortly to the audience and take difficult questions from the judges. A team of judges and the attending crowd will select the top contributor and declare the winner of the Grand Challenge 2011. Special awards might be given to contributions that show an outstanding approach to the integration of multiple media or are based on a novel theoretical framework.

For more information visit the ACM Multimedia 2011 Grand Challenge website: <http://www.acmmm11.org/content-call-multimedia-grand-challenge-solutions.html>

Award opportunities

SIGMM Award for Outstanding PhD Thesis in Multimedia Computing, Communications and Applications

More informational: <http://www.sigmm.org/Awards/thesisaward>

Award Description

This award will be presented every year to a researcher whose PhD thesis has made contributions and has the potential of very high impact in multimedia computing, communication and applications. The goal will be to evaluate contributions towards advances in multimedia including multimedia processing, multimedia systems, multimedia network protocols and services, multimedia applications and interfaces. The award will recognize members of the SIGMM community and their research contributions in their PhD theses as well as the potential of impact of their PhD theses in multimedia area. The selection committee will focus on candidates' contributions as judged by innovative ideas and potential impact resulting from their PhD work. The award includes a \$500 honorarium, an award certificate of recognition, and an invitation for the recipient to receive the award at a current year's SIGMM-sponsored conference, the ACM International Conference on Multimedia (ACM Multimedia). A public citation for the award will be placed on the SIGMM website, in the SIGMM Records e-newsletter as well as in the ACM e-newsletter.

Funding

The award honorarium, the award plaque of recognition and travel expenses to the ACM International Conference on Multimedia will be fully sponsored by the SIGMM budget.

Nominations

Nominations will be solicited by April 18, 2011, with decision made by July 30, in time to allow the above recognition and award presentation at ACM Multimedia in that Fall (October/November). The PhD thesis to be nominated for the award must be deposited at the nominee's Academic Institution between January and December of the previous year of nomination. Nominations for the award must include:

1. A statement summarizing the candidate's PhD thesis contributions and potential impact, and justification of the nomination (two pages maximum);
2. PhD thesis (upload at: <http://sigmm.utdallas.edu:8080/drupal/>);
3. Curriculum Vitae of the nominee;
4. Three endorsement letters supporting the nomination including the significant PhD thesis contributions of the candidate. Each endorsement should be no longer than 500 words with clear specification of nominee PhD thesis contributions and potential impact on the multimedia field;
5. A concise statement (one sentence) of the PhD thesis contribution for which the award is being given. This statement will appear on the award certificate and on the website.

The PhD thesis should be uploaded to <http://sigmm.utdallas.edu:8080/drupal/>. You will receive a Paper ID after the submission. The other materials should be emailed to the committee chair with the title "[YourPaperID] PhD Thesis Award Submission". The nomination rules are:

1. The nominee can be any member of the scientific community.
2. The nominator must be a SIGMM member.
3. No self-nomination is allowed.

Selection Committee

Professor Svetha Venkatesh - Chair
Professor Dick Bulterman
Professor Abed El Saddik

SIGMM Award for Outstanding Technical Contributions to Multimedia Computing, Communications and Applications

More informational: <mailto:lienhart@informatik.uni-augsburg.de>

2011 Call for Nominations of Multimedia Researchers for the SIGMM Award for Outstanding Technical Contributions to Multimedia Computing, Communications and Applications

Deadline for Nominations: May 1, 2011

Submission of Nomination Material to Awards Committee Chair: Rainer Lienhart (lienhart@informatik.uni-augsburg.de)

This award is presented every year to a researcher who has made significant and lasting contributions to multimedia computing, communication and applications. Outstanding technical contributions through research and practice are recognized. Towards this goal, contributions are considered from academia and industry that focus on major advances in multimedia including multimedia processing, multimedia content analysis, multimedia systems, multimedia network protocols and services, and multimedia applications and interfaces. The award recognizes members of the community for long-term technical accomplishments or those who have made a notable impact through a significant technical innovation.

The selection committee focuses on candidates' contributions as judged by innovative ideas, influence in the community, and/or the technical/social impact resulting from their work. The award includes a \$1000 honorarium, an award certificate of recognition, and an invitation for the recipient to present a keynote talk at a current year's SIGMM-sponsored conference, the ACM International Conference on Multimedia (ACM Multimedia). A public citation for the award will be placed on the SIGMM website.

The award honorarium, the award certificate of recognition and travel expenses to the ACM International Conference on Multimedia is fully sponsored by the SIGMM budget.

Nominations are solicited by May 1, 2011, with decision made by June 30, 2011, in time to allow the above recognition and award presentation at ACM Multimedia 2011. Nominations for the award must include:

1. A statement summarizing the candidate's accomplishments, description of the significance of the work, and justification of the nomination (two pages maximum);
2. Curriculum Vitae of the nominee;
3. Three endorsement letters supporting the nomination including the significant contributions of the candidate. Each endorsement should be no longer than 500 words with clear specification of nominee contributions and impact on the multimedia field;
4. A concise statement (one sentence) of the achievement(s) for which the award is being given. This statement will appear on the award certificate and on the website.

The nomination rules are:

1. The nominee can be any member of the scientific community.
2. The nominator must be a SIGMM member.

3. No self-nomination is allowed.
4. Nominations that do not result in an award can be resubmitted.
5. The SIGMM elected officers as well as members of the Awards Selection Committee are not eligible.

PhD thesis abstracts

Andre Miede

Cross-Organizational Service Security - Attack Modeling and Evaluation of Selected Countermeasures



Challenging market dynamics and the rise of complex value networks require organizations to adjust their processes rapidly in order to stay competitive. Because many organizational processes are directly supported or even enabled by Information Technology (IT), a process is only as flexible as its underlying technological representation. The Service-oriented Architecture paradigm (SOA) offers means on both a technological and organizational level for the flexible integration of internal and external IT systems. Thus, services are used to assemble processes through service compositions, as well as across enterprise boundaries. Such cross-organizational service-based workflows lead to a global SOA which is often referred to as the "Internet of Services".

Just as any economic system requires security in order to function and to be accepted by its participants, the security of the involved IT systems, exchanged messages, and communication channels used has

to be ensured for cross-organizational service-based collaboration. Achieving and guaranteeing basic IT security goals such as confidentiality, authentication, authorization, non-repudiation, integrity, availability, and anonymity is a necessity in this context and an active topic, both in research and industry. The main tenor of current SOA security research is that conventional security measures are not effective enough in the SOA context. Furthermore, just equalizing SOA security with Web service security reduces SOA security requirements to Web service security standards and their configuration, which is an incomplete view. This thesis makes several contributions regarding the security of service-based systems: First, it is shown how a model of cross-organizational SOA concepts can be used for analyzing SOA elements regarding their impact on security. This is done by applying core IT security concepts, such as threats, vulnerabilities, etc., to the general elements of a cross-organizational SOA, such as loose coupling, composability, etc. Second, an analysis of attacks in the Internet of Services is performed by proposing an attack taxonomy for service-based systems and by modeling selected examples of service-specific attack classes. This goes beyond the current state-of-the-art regarding SOA attacks by taking into account more service-specific and business-oriented threats. The modeling of these attacks builds on a self-developed generic metamodel, that brings together the most important concepts of IT security and their relationships. It is shown, how assets, threats, vulnerabilities, risks, security goals, etc. relate to each other at the core of this metamodel and what the basic structure of countermeasures is. Third, an attack scenario of communication analysis that threatens relationship anonymity in the Internet of Services is further investigated, due to its system-inherent implications. With a particular focus on service compositions, a simulation-based evaluation of different attack models and scenarios offers insights regarding the anonymity of cross-organizational collaboration. Furthermore, the impact of using standard anonymity mechanisms on selected Quality of Service parameters is evaluated for Web services in real networks. The obtained results aim at identifying the limits of anonymity in the Internet of Services and at quantifying side-effects of using state-of-the-art countermeasures.

Advisor(s): Prof. Dr.-Ing. Ralf Steinmetz (Supervisor),
 Prof. Dr. Dr. h.c. Alexander Schill (Referee)
 SIG MM member(s): Ralf Steinmetz
 ISBN number: 978-3-86853-718-5
<http://www.dr.hut-verlag.de/9783868537185.html>

Beatriz Soret

Analysis of QoS parameters in fading channels based on the effective bandwidth theory



Providing Quality of Service (QoS) guarantees is an important challenge in the design of next generations of wireless networks. In particular, real-time services involving stringent delay constraints are expected to be increasingly popular among users of mobile equipments. In Rayleigh channels, the delay requirement is usually expressed in terms of a probabilistic delay constraint composed by two terms: the target delay and the probability of exceeding the target delay.


In this thesis, a variable-rate multiuser and multichannel system using adaptive modulation is addressed. Specifically, the tradeoff among information source, fading channel and delay is analyzed, based on the effective bandwidth theory. Within the effective bandwidth framework, expressions of the channel effective bandwidth function (also known as effective capacity) are obtained on the channel side. Several scenarios are addressed: uncorrelated and time-correlated flat Rayleigh channels and an OFDM system under a frequency-selective Rayleigh channels. The procedure to obtain these functions is generic and could be applied to other channel models and scenarios. The effective bandwidth theory makes feasible the analysis of the distribution tail of the delay. The percentile of the delay and the maximum information rate that can be transmitted over the channel under a target BER and a probabilistic delay constraint are evaluated. The delay suffered by certain information flow depends not only on the transmission rate but also on the distribution and self-correlation of the information process. Even in wired systems (constant rate channels) different distributions of the information process having the same

average rate will cause different delays. Indeed, the better conditions for the delay are obtained when the incoming user traffic is constant. For any other source process, the delay performance degrades. Besides, the correlation of the channel process, in time or in frequency, has also a negative impact on the delay. In the last part of the thesis, multiplexing of users over multiple shared fading channel is addressed. A new element comes up in this case: the scheduling algorithm. We calculate the maximum rate that each user can transmit by fulfilling a target BER and its own delay constraint, and under a given scheduling discipline. The analysis is done first in a single channel link and later on generalized to multiple shared channels employing OFDMA as multiplexing mechanism. Now it is not only the delay constraint and the channel and source process that influence the source rate, but also the discipline that rules the system. Three representative multiplexing algorithms are analyzed: Round Robin, Best Channel and Proportional Fair. The results make possible the comparison of the algorithms in terms of throughput, delay and fairness. In summary, this thesis shows the high sensitivity of the delay to the burstiness of the traffic, to the time or frequency correlation of the channel and to the scheduling discipline. The proposed procedure is generic and can be extended to other disciplines and traffic and channel models. Nevertheless, the effective bandwidth function of the source and the channel process cannot always be explicitly evaluated. For such cases, a semi-analytical strategy is also proposed.

Advisor(s): Dr. M. Carmen Aguayo-Torres (supervisor)
 SIG MM member(s): M. Carmen Aguayo-Torres
 ISBN number:
<http://www.ic.uma.es/repository/fileDownloader?rfname=e634eaca-a415-47e3-ae5d-8aca8799fc5a.pdf>

Lin Lin

Multimedia Data Mining and Retrieval for Multimedia Databases Using Associations and Correlations



With the explosion in the complexity and amount of pervasive multimedia data, there are high demands of multimedia services and applications in various areas for people to easily access and distribute multimedia data. Facing with abundance multimedia resources but inefficient and rather old-fashioned keyword-based information retrieval approaches, a content-based multimedia information retrieval (CBMIR) system is required to (i) reduce the dimension space for storage saving and computation reduction; (ii) advance multimedia learning methods to accurately identify target semantics for bridging the semantics between low-level/mid-level features and high-level semantics; and (iii) effectively search media content for dynamical media delivery and enable the extensive applications to be media-type driven.

This research mainly focuses on multimedia data mining and retrieval system for multimedia databases by addressing some main challenges, such as data imbalance, data quality, semantic gap, user subjectivity and searching issues. Therefore, a novel CBMIR system is proposed in this dissertation. The proposed system utilizes both association rule mining (ARM) technique and multiple correspondence analysis (MCA) technique by taking into account both pattern discovery and statistical analysis. First, media content is represented by the global and local low-level and mid-level features and stored in the multimedia database. Second, a data filtering component is proposed in the system to improve the data quality and reduce the data imbalance. To be specific, the proposed filtering step is able to vertically select features and horizontally prune instances in multimedia databases. Third, a new learning and classification method mining weighted

association rules is proposed in the retrieval system. The MCA-based correlation is used to generate and select the weighted N-feature-value pair rules, where the N varies from 1 to many. Forth, a ranking method independent of classifiers is proposed in the system to sort the retrieved results and put the most interesting ones on the top of the browsing list. Finally, a user interface is implemented in CBMIR system that allows the user to choose his/her interested concept, searches media based on the target concept, ranks the retrieved segments using the proposed ranking algorithm, and then displays the top-ranked segments to the user. The system is experimented with various high-level semantics from TRECVID benchmark data sets. TRECVID sound and vision data is a large data set, includes various types of videos, and has very rich semantics. Overall, the proposed system achieves promising results in comparison with the other well-known methods. Moreover, experiments that compare each component with some other famous algorithms are conducted. The experimental results show that all proposed components improve the functionalities of the CBMIR system, and the proposed system reaches effectiveness, robustness and efficiency for a high-dimensional multimedia database.

Advisor(s): thesis supervisor: Dr. Mei-Ling Shyu
 SIG MM member(s): Lin Lin
 ISBN number: 9781124241838
http://scholarlyrepository.miami.edu/oa_dissertations/434

Data mining, Database & Multimedia Research Group

<http://rvc.eng.miami.edu/>

The Data mining, Database & Multimedia (DDM) Research Group is located in the Department of Electrical and Computer Engineering at the University of Miami, Coral Gables, Florida, USA. Dr. Mei-Ling Shyu is the director.

The mission of the DDM research group is to perform leading edge research in data mining, multimedia database systems, multimedia data mining, multimedia networking, data integration, and network security. In support for excellence, the DDM research group receives funding from agencies such as the NSF, Naval Research Laboratory, NOAA (National Oceanic and Atmospheric Administration), Florida Office of Insurance Regulation, Florida Department of Insurance, and National Park Service.

Peter Knees

Text-Based Description of Music for Indexing, Retrieval, and Browsing



The aim of this PhD thesis is to develop automatic methods that extract textual descriptions from the Web that can be associated with music pieces. Deriving descriptors for music permits to index large repositories with a diverse set of labels and allows for retrieving pieces and browsing collections. The techniques presented make use of common Web search engines to find related text content on the Web. From this content, descriptors are extracted that may serve as

- labels that facilitate orientation within browsing interfaces to music collections, especially in a three-dimensional browsing interface presented,
- exing terms, used as features in music retrieval systems that can be queried using descriptive free-form text as input, and
- es in adaptive retrieval systems that aim at providing more user-targeted recommendations based on the user's searching behaviour for exploration of music collections.

In the context of this thesis, different extraction, indexing, and retrieval strategies are elaborated and evaluated. Furthermore, the potential of complementing Web-based retrieval with acoustic similarity extracted from the audio signal, as well as complementing audio-similarity-based browsing approaches with Web-

based descriptors is investigated and demonstrated in prototype applications.

Advisor(s): Dr. Gerhard Widmer (supervisor)
SIG MM member(s): Peter Knees
http://www.cp.jku.at/research/papers/knees_phdthesis_2010.pdf

Department of Computational Perception, Johannes Kepler University Linz, Austria


<http://www.cp.jku.at>

The Department of Computational Perception of the Johannes Kepler University Linz, Austria carries out basic and applied research in machine learning, pattern recognition, knowledge extraction, information retrieval, and generally Artificial and Computational Intelligence with a focus on intelligent audio (specifically: music) and image processing. Headed by Prof. Gerhard Widmer, it has become one of the world-leading research groups in Music Information Retrieval. Current music-related research directions comprise the recognition and transcription of musical dimensions such as beat, tempo, and pitch from audio recordings, real-time tracking of scores and vocals from live performances, automatic rendering of expressive piano performances, music retrieval and recommendation in collections of millions of songs, and the development of novel interfaces to music collections. In addition to signal-based music research, a focus is also put on Web-mining techniques to exploit contextually related information on music.

The Department of Computational Perception maintains close cooperation links with the Austrian Research Institute for Artificial Intelligence (OFAI), Vienna, and in particular with its Machine Learning, Data Mining, and Intelligent Music Processing Group (which is also headed by Prof. Gerhard Widmer).

Vineeth N Balasubramanian

Conformal Predictions in Multimedia Pattern Recognition



The field of multimedia pattern recognition is on a fundamental quest to design intelligent systems that can learn and behave the way humans do. One important aspect of human intelligence that has so far not been given sufficient attention in these fields is the capability of humans to hedge decisions. Humans can express when they are certain about a decision they have made, and when they are not. Unfortunately, machine learning techniques today are not yet fully equipped to be trusted with this critical task. This work seeks to address this fundamental knowledge gap. Existing approaches that provide a measure of confidence of a learning algorithm on a prediction such as those based on the Bayesian theory or the Probably Approximately Correct learning theory require strong assumptions or often produce results that are not practical or reliable. However, the recently developed Conformal Predictions (CP) framework - which is based on the principles of hypothesis testing, transductive inference and algorithmic randomness - provides a game-theoretic approach to the estimation of confidence with several desirable properties such as online calibration and generalizability to all classification and regression methods.

This dissertation builds on the theory of Conformal Predictions to compute reliable confidence measures that aid decision-making in real-world multimedia problems. The theory behind the CP framework guarantees that the confidence values obtained using this transductive inference framework manifest as the actual error frequencies in the online setting, i.e. they are well-calibrated. Further, this framework can be used with any classifier, meta-classifier or regressor (such as Support Vector Machines, k-Nearest Neighbors, Adaboost, ridge regression, etc). The key contributions of this dissertation (outlined below) are validated

on four problems from the domains of healthcare and assistive technologies: two classification-based applications (risk prediction in cardiac decision support and multimodal person recognition), and two regression-based applications (head pose estimation and saliency prediction in radiological images). The cost of errors in decision-making is often high in these application domains, and hence these problems are selected to validate the contributions. The key contributions of this work are summarized below: (1) Efficiency Maximization in Conformal Predictors: The CP framework has two important properties that define its utility: validity and efficiency. Validity refers to controlling the frequency of errors within a pre-specified error threshold. Also, since the framework outputs a set of possible predictions as the result, it is essential that the prediction sets are as small as possible. This property is called efficiency. Evidently, an ideal implementation of the framework would ensure that the algorithm provides high efficiency along with validity. However, this is not a straightforward task, and depends on the learning algorithm (classification or regression, as the case may be) as well as the non-conformity measure chosen in a given context. In this work, a novel framework to learn a kernel (or distance metric) that will maximize the efficiency in a given context has been proposed and validated on different risk-sensitive applications. (2) Conformal Predictions for Information Fusion: The CP framework ensures the calibration property in the estimation of confidence in pattern recognition. Most of the existing work in this context has been carried out using single classification systems or ensemble classifiers (such as boosting). However, there been a recent growth in the use of multimodal fusion algorithms and multiple classifier systems. A study of statistical approaches to combine p-values from multiple classifiers and regressors has been performed, which revealed the usefulness of quantile combination methods for calibrated confidence values in information fusion contexts. (3) Online Active Learning using Conformal Predictors: As newer data are encountered, it becomes essential to select appropriate data instances for labeling and updating the classifier to facilitate a continuously learning system. Using the p-values computed by the CP framework, a novel online active learning approach has been proposed and validated. This active learning method can also be extended to an information fusion setting, where there are multiple information sources or multiple modalities. The results obtained in this work demonstrate promise and potential in using these contributions to provide reliable measures of confidence in multimedia pattern recognition problems in real-world settings.

Advisor(s): Sethuraman Panchanathan
 SIG MM member(s): Jieping Ye, Baoxin Li, Vladimir Vovk
 ISBN number: 978-1-124-31019-0

<http://gradworks.umi.com/34/27/3427262.html>

Center for Cognitive Ubiquitous Computing, Arizona State University

<http://cubic.asu.edu>

The Center for Cognitive Ubiquitous Computing (CUbiC) at Arizona State University is an inter-disciplinary research center focused on human-centered multimedia computing in the domains of assistive, rehabilitative and healthcare technologies. CUbiC employs a transdisciplinary research approach, which includes computer scientists, cognitive scientists, psychologists, healthcare professionals, engineers, and designers, for solving the challenges in human-centered multimedia computing. Existing approaches have largely relied on the so-called "able" population to derive insights to shape the efforts towards achieving human-centeredness. In contrast, CUbiC has proposed a new archetype to human-centered multimedia computing inspired by the needs of individuals with disabilities. The study of sensory, motor, perceptual and cognitive disabilities helps us understand the subtleties of human capabilities and limitations, thereby necessitating the design of newer methodologies for data capture, information processing and multimodal delivery. This approach results in not only the design and development of innovative multimedia solutions for enriching the lives of individuals with disabilities/disorders, but is also valuable for gaining a deeper understanding towards realizing unique solutions for mainstream multimedia applications.

The focal application domains (assistive, rehabilitative and healthcare technologies) represent unique facets of human-machine interaction, which provide unique perspectives to our research. The healthcare domain primarily deals with how a disability or deficit, in the broader sense of the term, is diagnosed in a user (and further treated appropriately); the rehabilitative domain deals with how a technology is closely associated with the user for a temporary period of time to help the user overcome the disability and regain normalcy; and the assistive domain deals with how a technology is associated with a user for long periods of time (sometimes an entire lifetime) to support and enrich daily activities,

due to the presence of a chronic disability. This disability-inspired approach to multimedia computing has led to fundamental research advancements in various fields including multimodal sensing, signal processing, pattern recognition, machine learning, human-computer interaction and multimodal delivery. These advances have taken the shape of several projects under the umbrella of iCARE (information technology Centric Assistive and Rehabilitative Environments), including the Reader, Note Taker, Information Assistant, Environment Perception, Multimodal (audio and haptic) Interfaces, and the Interaction Assistant. Our work thus far has demonstrated that research centered on individuals with disabilities and deficits has far reaching implications for the general population and in advancing the core principles of human-centered multimedia computing.

Event and publication reports

MMSJ, Volume 16, Issue 6, November 2010

NewsBox=Editor-in-Chief: Thomas Plagemann
URL: <http://www.springer.de/>
Published: November 2010

Papers

- Pradeep K. Atrey, M. Anwar Hossain, Abdulmotaleb El Saddik and Mohan S. Kankanhalli: *Multimodal fusion for multimedia analysis: a survey*
- Aylin Kantarci: *Bandwidth-effective streaming of educational medical videos*
- Petros Belimpasakis and Anne Saaranen: *Sharing with people: a system for user-centric content sharing*

TOMCCAP, Volume 6, Issue 4, August 2010

NewsBox=Editor-in-Chief: Ralf Steinmetz
URL: <http://tomccap.acm.org/>
Sponsored by ACM SIG Multimedia
Published: November 2010

Papers

- Ralf Steinmetz: *Obituary to our dear friend professor Dr. Nicolas D. Georganas, PhD.*

- Thomas Haenselmann: *Foreword to the special issue on multimedia sensor fusion*
- Xiangyu Wang, Mohan Kankanhalli: *MultiFusion: A boosting approach for multimedia fusion*
- Girija Chetty, Matthew White: *Multimedia sensor fusion for retrieving identity in biometric access control systems*
- Gerald Friedland, Chuohao Yeo, Hayley Hung: *Dialocalization: Acoustic speaker diarization and visual localization as joint optimization problem*
- Abu Saleh Md Mahfujur Rahman, M Anwar Hossain, Abdulmotaleb El Saddik: *Spatial-geometric approach to physical mobile interaction based on accelerometer and IR sensory data fusion*
- Zhenyu Yang, Wanmin Wu, Klara Nahrstedt, Gregorij Kurillo, Ruzena Bajcsy: *Enabling multiparty 3D tele-immersive environments with ViewCast*
- Damien Marshall, Séamus Mcloone, Tomás Ward: *Optimizing consistency by maximizing bandwidth usage in distributed interactive applications*
- Long Vu, Indranil Gupta, Klara Nahrstedt, Jin Liang: *Understanding overlay characteristics of a large-scale peer-to-peer IPTV system*

Calls for contributions

Calls for SIGMM Sponsored and Co-sponsored Events

ACM Multimedia

Full paper Deadline: April 11, 2011
Event location: Scottsdale, AZ, USA
Event date: November 28-December 1, 2011
URL: <http://www.acmmm11.org/>

ACM Multimedia 2010 is the worldwide premier multimedia conference and a key event to display scientific achievements and innovative industrial products. The Conference offers to scientists and practitioners in the area of Multimedia plenary scientific and technical sessions, tutorials, panels and discussion meetings on relevant and challenging questions for the next years horizon of multimedia.

ACM Multimedia (Demos, Open Source)

Full paper Deadline: August 6, 2011
Event location: Scottsdale, AZ, USA
Event date: November 28-December 1, 2011
URL: <http://www.acmmm11.org/>

Calls for Events held in cooperation with SIGMM

ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC)

Full paper Deadline: April 10, 2011
Event location: Ghent, Belgium
Event date: Aug. 23-26, 2011
URL: <http://www.icdsc.org/>

Technological developments in imaging, processing, and networking have created an opportunity for multi-disciplinary approaches to applications based on vision. The extensive availability and use of cameras in various application domains calls for the study of new embedded and distributed processing systems and algorithms.

ACM Multimedia and Security Workshop (MM&Sec)

Full paper Deadline: May 23, 2011
Event location: Niagra Falls, NY, USA
Event date: September 29-30, 2011
URL: <http://www.mmsec11.com/>

The workshop's continuing objective is to explore research in areas of multimedia data security such as data protection, media forensics, covert channels and security issues in biometrics, as well as related issues in public policy and multimedia infrastructure in real world application.

International Conference on Advances in Mobile Computing & Multimedia (MoMM2010)

Full paper Deadline: July 15, 2011
Event location: Hue City, Vietnam
Event date: December 5-7, 2011
URL: <http://www.iivas.org/conferences/momm2011/>

The conference is meant for sharing original research results and practical development experiences from all mobile computing and multimedia related areas.

International Conference on Multimodal Interfaces (ICMI)

Full paper Deadline: May 13, 2011

Event location: Alicante, Spain
Event date: November 14-18, 2011
URL: <http://www.acm.org/icmi/2011/>

The conference is a forum for multidisciplinary research on multimodal human-human and human-computer interaction, interfaces, and system development. The conference focuses on theoretical and empirical foundations, component technologies, and combined multimodal processing techniques that define the field of multimodal interaction analysis, interface design, and system development.

2011 NetGames: International Workshop on Network and Systems Support for Games

Full paper Deadline: July 1 2011
Event location: Ottawa, Canada
Event date: October 6-7 2011
URL: <http://www.discover.uottawa.ca/netgames2011/>

NetGames brings together researchers and practitioners from both academia and industry to present the latest research results and challenges of today's networked games, and to understand their requirements and possibilities in order to enable the next generation of networked games. NetGames also provides industry keynote and panel discussions.

Other multimedia-related Events

IEEE International Symposium on Multimedia (ISM 2011)

Full paper Deadline: April 6, 2011
Event location: Dana Point, CA, USA
Event date: December 5-7, 2011
URL: <http://ism.eecs.uci.edu/>

This symposium is an international forum for researchers to exchange information regarding advances in the state-of-the-art and practice of multimedia computing, as well as to identify the emerging research topics and define the future of multimedia computing.

IEEE International Conference on Semantic Computing (ICSC 2011)

Full paper Deadline: May 3, 2011
Event location: Palo Alto, CA, USA

Event date: September 18-21, 2011
URL: <http://www.ieee-icsc.org/>

The conference builds on the success of the past ICSC conferences as an international forum for researchers and practitioners to present research that advances the state of the art and practice of Semantic Computing, as well as identifying emerging research topics and defining the future of the field.

IEEE International Workshop on Multimedia Signal Processing (MMSP 2011)

Full paper Deadline: April 15, 2011
Event location: Hangzhou, China
Event date: October 17-19, 2011
URL: <http://www.mmsp2011.org/>

The focus theme for this workshop is client-cloud multimedia computing. The cloud computers, clients and the network connecting them jointly constitute a new distributed computing environment with large scale computation, data and storage, which offers great opportunities for multimedia signals. Of particular interest is the case where mobile devices serve as the input and output console. Besides the existing challenges in the common client-cloud environment, multimedia computing has additional challenges in joint client-cloud computing.

ACM International Conference on Health Informatics (IHI)

Full paper Deadline: May 23, 2011
Event location: Miama, FL, USA
Event date: November 9-11, 2011
URL: <http://sites.google.com/site/web2011ihi/>

The conference is the premier international forum concerned with the application of computer science principles, information science principles, information technology, and communication technology to address problems in healthcare, public health, everyday wellness as well as the related social and ethical issues.

ACM/IFIP/USENIX 11th International Middleware Conference (Middleware)

Full paper Deadline: May 23, 2011
Event location: Lisboa, Portugal
Event date: Decemver 12-16, 2011
URL: <http://2011.middleware-conference.org/>

The Middleware conference is a forum for the discussion of innovations and recent advances in the design, implementation, deployment, and usage of middleware systems. Middleware is the software that resides between applications and the underlying architecture. The goal of middleware is to facilitate the development of applications by providing higher-level abstractions for better programmability, performance, scalability, security, and a variety of essential features. It is a rapidly evolving and growing field.

MediaEval Multimedia Benchmark Evaluation

Full paper Deadline: Register by May, 2011
 Event location: Pisa, Italy
 Event date: September 1-2, 2011
 URL: <http://www.multimediaeval.org/>

MediaEval is a benchmarking initiative that offers tasks promoting research and innovation on multimodal approaches to multimedia access and retrieval. The MediaEval benchmarking initiative offers tasks in multimedia access and retrieval including genre tagging, rich speech retrieval, spoken web search, affect task: violence detection, social event detection, and plating.

European Conference on Visual Media Production (CVMP)

Full paper Deadline: June 20, 2011
 Event location: London, UK
 Event date: November, 2011
 URL: <http://www.cvmp-conference.org>

The conference brings together production and post-production specialists from the worlds of film, broadcast and games with imaging and graphics researchers. It provides a European forum to discuss the latest research, advances and state-of-the-art industry practices.

Elsevier Journal of Visual Communication and Image Representation (JVCI): Special issue on Recent Advances on Analysis and Processing for Distributed Video Systems

Full paper Deadline: May 30, 2011
 Event date: Summer 2012
 URL: <http://www.elsevier.com/locate/jvci>
 CFP_JVCI_SI_Distributed_Video_Systems.pdf

Distributed video systems are of increasing importance in many applications, including surveillance, healthcare, entertainment, and unmanned area monitoring. There is an evolution from the static centric-based processing to dynamic collaborative computing and processing among distributed video processing nodes. This evolution is issuing new challenges. This special issue aims to bring together leading researchers and practitioners from around the world to present their latest research results and explore future directions in distributed video systems.

Job Opportunities

2-year Post-Doc/Research position in Audio and Multimedia scene analysis using multiple sensors

Employer: Institut Telecom - Telecom ParisTech
 Valid until: July 31, 2011
 More info: <http://service.tsi.telecom-paristech.fr/cgi-bin/announce/download.cgi?Aid=25&lang=en>

The researcher role will consist in participating to the collaborative integration activities of the Network of Excellence 3Dlife, and in conducting forefront research in audio/multimedia scene analysis using multiple sensors. A specific interest will include the development of innovative statistical fusion approaches capable of processing information on multiple semantic levels.

Back matter

Notice to Contributing Authors to SIG Newsletters

By submitting your article for distribution in this Special Interest Group publication, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish in print on condition of acceptance by the editor
- to digitize and post your article in the electronic version of this publication
- to include the article in the ACM Digital Library and in any Digital Library related services
- to allow users to copy and distribute the article for noncommercial, educational or research purposes

However, as a contributing author, you retain copyright to your article and ACM will refer requests for republication directly to you.

Impressum

Editor-in-Chief Carsten Griwodz, Simula Research Laboratory
Editors Stephan Kopf, University of Mannheim
 Viktor Wendel, Darmstadt University of Technology
 Lei Zhang, Microsoft Research Asia
 Wei Tsang Ooi, National University of Singapore

