1. Record Nr. UNINA9910510473003321

Autore	Chen Chang Wen
Titolo	Proceedings of the 28th ACM International Conference on Multimedia : October 12-16, 2020, Virtual Event, USA / / Chang Wen Chen
Pubbl/distr/stampa	New York : , : Association for Computing Machinery, , [2020] ©2020
Descrizione fisica	1 online resource (4814 pages) : illustrations
Collana	ACM Conferences
Disciplina	006.7
Soggetti	Multimedia systems
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Sommario/riassunto	With great pleasure, we welcome you to join ACM Multimedia 2020, the 28th ACM International Conference on Multimedia. ACM Multimedia is the premier international conference series in the area of multimedia within the field of computer science. Since 1993, ACM Multimedia has been bringing together worldwide researchers and practitioners from academia and industry to present their innovative research and to discuss recent advancements in multimedia. This year's conference, for the first time in the history, will be running as an online large-scale meeting due to COVID-19. Face-to-face gatherings have been replaced by all-hands live sessions, parallel live question-answering (Q&A) sessions, and presentations from the authors via pre-recorded videos. However, the distance needed to prevent the spreading of the virus does not cool down enthusiasm in participants, presenters, chairs, reviewers, and volunteers. Under today's challenging circumstance, we are pleased that we are still able to deliver the most cutting-edge research results, covering the latest findings in the field, for which the ACM Multimedia conference series is widely known. From some different perspectives, virtual conference is probably even better. All the videos can be previewed online before the start of main conference events, and within 10 days, you can replay them as many times as you would like to. During conference days, you can join and jump virtually

among as many Q&A sessions as possible. Furthermore, for each Q&A session, there will be a mirrored one after 10 hours for you to join in case your time zone is not suitable to participate for the first one. We would like to urge the participants to check the FAQ to plan your conference days in the best and most comfortable way. This year we are honored to welcome two distinguished keynote speakers. The opening keynote will be delivered by Dr. Shuicheng Yan, who is currently CTO of YITU Tech, and was previously a Professor at the National University of Singapore. He is Fellow of the Academy of Engineering, Singapore, IEEE Fellow, and IAPR Fellow. It is well-known that neural network architecture design is playing an important role in recent fast development of multimedia technology. In this talk, Shuicheng will introduce the research and development efforts in designing neural networks from two orthogonal perspectives: on how these models are bio-inspired and on how these models are more hardware-friendly to motivate the design of next-generation AI chips. These two lines of efforts are collaboratively enhancing the overall efficiency of multimedia systems. The second keynote is scheduled on the third main conference day and will be delivered by Prof. Klara Nahrstedt, Ralph and Catherine Fisher Professor in the Department of Computer Science Department, and Director of the Coordinated Science Laboratory in the Grainger College of Engineering at the University of Illinois at Urbana-Champaign. She is Fellow of ACM, IEEE and AAAS, and a member of the German National Academy of Sciences (Leopoldina Society). With the emergence of new 360-degree cameras, ambisonic microphones, and VR/AR display devices, more diverse multi-modal contents have become available. With such contents, the demand for the capability of streaming 360-degree videos to enhance users' 360-multimedia experience on mobile devices such as mobile phones and head-mounted displays has been growing substantially. The big issue for the mobile 360-multimedia delivery systems is the huge resource demand on the underlying networks and devices to deliver 360-multimedia content with high quality of experience. In this talk, Klara will discuss the research challenges in 360-degree video delivery systems in terms of large bandwidth, low latency, users' disorientation, and cyber-sickness, and the opportunities to solve these challenges through developing rate adaptation algorithms for tiled videos, view prediction algorithms, content navigation, enhancement of DASH streaming for 360-videos, and control of Quality of Experience (QoE). This year, we are also happy to introduce two industrial invited talks. The first one is by Itamar Friedman from Alibaba DAMO Academy and his topic is "Cloud Drive App - Closing the Gap Between AI Research and Practice", and will introduce his latest work related to highly-efficient image understanding, aiming to enable various novel methods (such as neural architecture search and advanced training techniques) to be practiced in Cloud Drive App use cases. The second one is by Dr. Dong Yu from Tencent AI Lab. Dong will introduce their recent work in building digital human. He will focus on their progresses in multi-modal text-to-speech synthesis and multi-modal voice separation and recognition. You will quickly notice that, unlike last year when all main conference papers are presented as posters, this year, videos and corresponding Q&A sessions are the centerpiece of the conference. If you are a professional researcher or a graduate student in multimedia, we suggest that you register the conference in advance. This will allow you to view the videos online, followed by remote interaction with authors in Q&A sessions via Zoom meetings. We believe this will be a highly effective channel of communication of research results and discussion of findings. All conference papers, not

just the oral papers, will be presented via videos. We hope that you take advantage of the Q&A sessions after studying as many as possible related papers to interact directly with the paper authors. Due to Covid-19, we faced unprecedented challenges, and had to cancel some arrangements in the middle and extended paper submission deadlines to leave maximal time for participated teams and minimal time for reviewers. Nonetheless, we still have received record-breaking number of 1,698 valid submissions for the main conference tracks (increased 81.4% compared with the number of last year) with an overall acceptance rate of 27.9% (474 papers), covering all areas and subareas of multimedia. In addition to the main track, the conference has tracks devoted to brave new ideas, reproducibility companion, open source competition, interactive arts, and to multimedia demos. As always, a particular highlight is the Open Source Software Competition. The conference also offers two panels. The first panel will be in the spirit of "Coping with Pandemics". The coronavirus pandemic has significantly changed the way people live, work, communicate and learn. "Opportunities and Challenges for AI Multimedia in the 'New Normal'" will be deliberated. The second panel aims at opening a discussion about "What Multimedia has to offer for our common digital future." Beyond the current coronavirus health crisis which is holding the world in breath, the world is also facing several global challenges from climate change to environmental deterioration, and from access to clean water and food to socioeconomic inequalities, to name just a few. Our future health system, global access to education, decent work, and reducing inequalities are just some of these goals where research in this field can contribute. Let us brainstorm these in this panel. No ACM Multimedia would be complete without Grand Challenges. This year, we organized eight Multimedia Grand Challenges, which were selected from a wide range of proposals based on the strength of proposals, relevance to the community, challenge continuity, diversity of topics, and organizers. The selected Grand Challenges are a great reflection of the multimodality nature of the multimedia problems, including two of them for image and six for video-related challenges. The topics range from activity recognition, vision & language understanding, medical applications, and metadata prediction, with organizers coming from diverse geolocations including Asia, Europe, and North America. On the tutorial and workshop days, eight tutorials provide the opportunity to learn about new directions in multimedia research, and stay up to date on the latest developments. We are also pleased to offer a diverse set of 11 workshops which cover a wide range of cutting-edge topics important for the field of multimedia. A conference such as ACM Multimedia is only made possible by tremendous expertise, dedication, and hours of hard work from a large number of people in this vibrant multimedia community.

Like every year in the past, we would like to extend a resounding word of thanks to everyone who have contributed to ACM Multimedia 2020. It is our great pleasure to collaborate with this year's outstanding organizing committee consisting of devoted volunteers, who have worked extremely hard to put together an impressive and interesting program. We would also like to thank everyone who served as Area Chairs for their hard work in guaranteeing the novelty and scientific rigor of the main-conference program. A special thanks goes to the committed reviewers for their great work in reviewing submissions across all tracks of the conference. And another special thanks goes to the virtual conference "online chairs" (originally the local chairs team), and the finance chair, who have resolved numerous unexpected difficulties in making our conference program available online to all

attendees. Finally, we would like to thank all the authors who submitted to ACM Multimedia. The heart of the conference is your research discoveries, which is what drives our field forward.