Foreword

Welcome to the eleventh edition of the Proceedings of Machine Learning Research (PMLR) for the Asian Conference on Machine Learning (ACML). This volume contains 78 papers accepted to the Eleventh Asian Conference on Machine Learning (ACML 2019) held in Nagoya, Japan, from 17th to 19th November 2019. It has been more than a decade since ACML started, and the conference has grown from strength to strength. This volume testifies to the quality of the conference with its collection of excellent original research papers in the area of machine learning. ACML aims at providing a leading international forum for researchers in machine learning and related fields to share their original research findings, new ideas and achievements. Despite originating in the Asia-Pacific region, ACML has become a worldwide conference: submissions from regions other than the Asia-Pacific were encouraged and this year the conference received submissions from authors based in Brazil, Canada, Chile, Denmark, Finland, France, Germany, Italy, Netherland, Norway, Russia, Sweden, Switzerland, United Kingdom, United States, as well as the Asia-Pacific region.

As in previous years, we had two rounds of submissions for the ACML conference track this year. Promising papers that could not be accepted in the first round were encouraged for re-submission in the second round, allowing reviewers' comments to be addressed. In total, there were 291 submissions to the conference track, of which 78 were accepted into the main program, for an acceptance rate of 26.8%. A strict double-blind reviewing process was enforced, and each paper was assigned to one meta-reviewer, and at least 3 reviewers. The Program Committee (PC) had 295 PC members (reviewers) and 46 Senior PC members (meta-reviewers). Papers were allocated based on paper bidding, similarity of subject areas, and the Toronto Paper Machine System (TPMS) scores. The Program Committee contributed their time and expert opinions, ensuring the high standard of papers selected for the conference. Without them the conference would have not been possible.

The journal track of ACML has also grown substantially. This year there were 52 submissions and 9 papers were accepted for publication in the Springer Machine Learning Journal, for an acceptance rate of 17.3%. Reviewing was overseen by the journal track Co-Chairs Kee-Eung Kim and Jun Zhu. The same program committee reviewed both the conference and journal track of ACML and it is thanks to their effort that we are able to maintain the high quality of the accepted papers. The overall number of accepted papers, from both the journal and conference tracks, was 87 from 343 submissions for a 25.4% total acceptance rate.

For the first time, ACML had two parallel sessions; all previous ACML conferences had a single track. All accepted papers from the two tracks were presented in either an oral or spotlight presentation. In addition, all papers also had a poster presentation. The submissions covered a broad range of topics, including Bayesian and probabilistic machine learning, deep learning, optimization, weakly-supervised learning, unsupervised learning, multi-label learning, multi-instance learning, adversarial learning, reinforcement learning, applications to computer vision and natural language processing, as well as application to

real world problems.

We are very pleased to have four invited talks from leading experts, Le Song (Georgia Institute of Technology), James T. Kwok (Hong Kong University of Science and Technology), Prateek Jain (Microsoft Research Lab Vigyan), and Tatsuya Harada (University of Tokyo). In addition to the main program, we had a full day of tutorials and workshops prior to the main conference organized by Tutorial Co-Chairs Yu-Feng Li and Emtiyaz Khan, and Workshop Co-Chairs Taesup Moon and Mathieu Blondel. Five tutorials were delivered on: Deep Learning for Natural Language Processing and Computer Vision, Towards Noisy Supervision: Problems, Theories, and Algorithms, Machine Learning for Manufacturing, Deep Learning for Conversational AI, and Text Representation Learning and Compositional Semantics. In parallel to the tutorials, we had four workshops: ACML2019 workshop on Machine Learning for Trajectory, Activity, and Behavior (ACML-TAB2019), ACML2019 RIKEN AIP workshop, ACML'19 Weakly-supervised Learning Workshop, and ACML 2019 Workshop on Statistics & Machine Learning Researchers in Japan. We thank all of the speakers and organizers for putting together such a fantastic program.

ACML Steering Committee Chair Masashi Sugiyama, together with the Steering Committee, provided invaluable advice for the organization of the conference. The General Co-Chairs Geoff Holmes and Ichiro Takeuchi coordinated the entire event and ensured that everything was well organized. The Publicity Chair Marco Cuturi ensured that the conference was well publicised, the Publication Chair Shinichi Nakajima handled the preparation and publication of the proceedings, and the Web Chair Hiroyuki Hanada ensured all web matters are appropriately handled. Special thanks to the Local Arrangements Co-Chairs Kota Matsui, Noriaki Hashimoto, Yuta Umezu for handling all the local arrangements issues.

Without the help of all the people involved in organizing the conference, the event would not have been possible. Finally, we would like to thank all the participants of ACML 2019, for making the event a success!

November 2019

ACML 2019 Program Co-Chairs

Wee Sun Lee
National University of Singapore, Singapore
leews@comp.nus.edu.sg

Taiji Suzuki
The University of Tokyo, Japan
taiji@mist.i.u-tokyo.ac.jp

Invited Talks

Title: Graph Neural Networks for Representation Learning and Symbolic Reasoning

Presenter: Le Song (Georgia Institute of Technology)

Abstract:

Graphs and hypergraphs are prevalent in many real world applications arising from online social platforms, recommendation systems, knowledge bases, computational biology and materials science. How to represent such graph data to capture their similarities or differences? How to integrate graph data with other sources of data in representation learning? How to combine deep learning with symbolic reasoning? How to better design algorithms over graphs? I will present a graph neural network framework for addressing these challenges based on the idea of embedding message passing algorithms into function spaces, and learning these algorithms from data. In large scale applications involving molecule design, recommendation system and knowledge reasoning, this graph neural network framework consistently achieves the-state-of-the-art results, in terms of accuracy, model size and scalability. Graph neural networks also appear to be a very good tool to advance AI to the next stage, which can combine deep learning with symbolic reasoning.

Bio:

Le Song is an Associate Professor in the Department of Computational Science and Engineering, College of Computing, an Associate Director of the Center for Machine Learning, Georgia Institute of Technology, and also a Principal Engineer of Ant Financial, Alibaba. He received his Ph.D. from University of Sydney in 2008, and then conducted his post-doctoral research in the Department of Machine Learning, Carnegie Mellon University, between 2008 and 2011. Before he joined Georgia Institute of Technology in 2011, he was a research scientist at Google. His principal research direction is machine learning, especially nonlinear models, such as kernel methods and deep learning, and probabilistic graphical models for large scale and complex problems, arising from artificial intelligence, network analysis, computational biology and other interdisciplinary domains. He is the recipient of the NIPS'17 Materials Science Workshop Best Paper Award, the Recsys'16 Deep Learning Workshop Best Paper Award, AISTATS'16 Best Student Paper Award, IPDPS'15 Best Paper Award, NSF CAREER Award'14, NIPS'13 Outstanding Paper Award, and ICML'10 Best Paper Award. He has also served as the area chair or senior program committee for many leading machine learning and AI conferences such as ICML, NeurIPS, AISTATS, AAAI and IJCAI, and the action editor for JMLR and IEEE TPAMI.

Title: Compressing deep networks

Presenter: James T. Kwok (Hong Kong University of Science and Technology)

Abstract:

Deep networks are powerful but computationally expensive. This hinders deployment to small computing devices such as cell phones and the internet of things. To alleviate this problem, a popular approach is to quantize the deep network weights to a small number of bits. In this talk, we consider a number of issues related to network quantization. First, we show that the effect of quantization on the loss can be directly formulated into the optimization problem. Morevoer, the exploding gradient problem can become more severe in training quantized LSTMs. By using the popularly used weight/layer/batch normalization, we show theoretically and empirically that the gradient magnitude can be stabilized. Besides, communication overhead is a major bottleneck in distributed deep network training. To improve communication efficiency, besides using weight quantization, we propose a general distributed compressed SGD scheme which compresses the gradients both to and from workers. With these techniques, empirical results show that the resultant network can significantly reduce its size without sacrificing performance, and runs much faster in a distributed environment.

Bio:

Prof. Kwok is a Professor in the Department of Computer Science and Engineering, Hong Kong University of Science and Technology. He received his B.Sc. degree in Electrical and Electronic Engineering from the University of Hong Kong and his Ph.D. degree in computer science from the Hong Kong University of Science and Technology. Prof. Kwok served/is serving as an Associate Editor for the IEEE Transactions on Neural Networks and Learning Systems, Neurocomputing and the International Journal of Data Science and Analytics. He has also served as Program Co-chair of a number of international conferences, and as Area Chairs in major machine learning and AI conferences. He is an IEEE Fellow.

Title: Resource Efficient ML in 2KB of RAM

Presenter: Prateek Jain (Microsoft Research Lab Vigyan)

Abstract:

Several critical applications require ML inference on resource-constrained devices, especially in the domain of Internet of Things like smartcity, smarthouse etc. Furthermore, many of these problems reduce to time-series classification. Unfortunately, existing techniques for time-series classification like recurrent neural networks are very difficult to deploy on the tiny devices due to computation and memory bottleneck. In this talk, we will discuss two new methods FastGRNN and EMI-RNN that can enable time-series inference on devices as small as Arduino Uno that have 2KB of RAM. Our methods can provide as much as 70x speed-up and compression over state-of-the-art methods like LSTM, GRU, while also providing strong theoretical guarantees.

This talk is based on joint works with Manik Varma, Harsha Simhadri, Kush Bhatia, Don Dennis, Ashish Kumar, Aditya Kusupati, Manish Singh, and Shishir Patil.

Bio:

Prateek Jain is a senior principal researcher at Microsoft Research India. He is also an adjunct faculty member at the Computer Science department at IIT Kanpur. He received his PhD in Computer Science from University of Texas at Austin and his B.Tech. in Computer Science from IIT Kanpur. His research interests are in resource-constrained machine learning, high-dimensional statistics, and non-convex optimization. He has served on several senior program committees for top ML conferences and also won ICML-2007, CVPR-2008 best student paper awards.

Title: Visual Recognition from Limited Supervised Data

Presenter: Tatsuya Harada (University of Tokyo)

Abstract:

Training deep neural networks from limited supervised data for constructing an accurate prediction model is one of the crucial tasks in visual recognition problems. In this talk, we introduce domain adaptation methods for both classification and generative models that transfer knowledge in a label rich domain to a label scarce domain. We also present a new learning method using between-class examples to train DNNs and boost a classification performance from limited data. Besides, we will briefly introduce various topics that we are working on in our team.

Bio:

Tatsuya Harada is a Professor in the Research Center for Advanced Science and Technology at the University of Tokyo. His research interests center on visual recognition, machine learning, and intelligent robot. He received his Ph.D. from the University of Tokyo in 2001. He is also a team leader at RIKEN AIP and a vice director of Research Center for Medical Bigdata at National Institute of Informatics, Japan.

Tutorials

Title: Deep Learning for Natural Language Processing and Computer Vision

Presenter: Yoshimasa Tsuruoka (University of Tokyo), Yoshitaka Ushiku (Omron Sinic X)

Abstract:

The field of natural language processing (NLP) has witnessed major advancements in recent years thanks to deep learning technology. The accuracy of basic NLP tasks has improved drastically and the high-level tasks that were formerly possible only through complex combination of customized algorithms are now possible via end-to-end learning with relatively simple neural network architectures. This tutorial will cover the basics of deep learning-based NLP and some representative NLP applications including machine translation, summarization and question answering. It will also cover some recent research efforts for building NLP systems in resource-poor settings, such as unsupervised neural machine translation and large-scale pretraining and adaptation of deep learning models.

Title: Towards Noisy Supervision: Problems, Theories, and Algorithms

Presenter: Ivor W Tsang (University of Technology Sydney), Bo Han (Riken)

Abstract:

As dataset sizes grow bigger, it is laborious and expensive to obtain clean supervision. As a result, the volume of noisy supervision becomes enormous, e.g., crowdsourcing and single-label corruption. Unfortunately, noisy supervision harms the performance of most learning algorithms, and sometimes even makes existing algorithms break down. Recently, there are a brunch of theories and approaches proposed to deal with noisy data. In this tutorial, we summarize the foundations and go through the most recent noisy supervision techniques. By participating the tutorial, the audience will gain a broad knowledge of noisy-supervised learning from the viewpoint of statistical learning theory, and detailed analysis of typical algorithms and frameworks.

Title: Machine Learning for Manufacturing

Presenter: Ryo Yoshida (Institute of Statistical Mathematics), Koji Tsuda (University of Tokyo)

Abstract:

The ability of machine learning (ML) models trained on massive amounts of data has reached or even outperformed humans in intellectually demanding tasks across various fields. As such, ML has received considerable attention in manufacturing to reap substantial time and cost savings in many potential applications in industry and science. In this tutorial, we introduce a set of ML technologies that would be a key driver to the next frontier of creative design and manufacturing. The primary objective is to identify a set of parameters, such as the structure of materials and process parameters for the devise manufacturing, such that resulting response variables meet arbitrary given requirements. In general, a ML workflow consists of two steps; the first step is to build a prediction model that describes forwardly the response variables as a function of the input parameters, and this forward model is inverted to the backward one. With the given backward model conditioned by the design target, a set of parameters that exhibits the desired response is computationally explored. In this tutorial, some outstanding successes of ML are demonstrated with examples from materials science. The topics cover novel applications of deep learning technologies for designing materials structures or synthetic routes, transfer learning for overcoming a limited supply of materials data, Bayesian optimization frameworks that integrate ML and data from computer experiments such as the first principles calculation, and so on. The first speaker, Prof. Tsuda at the University of Tokyo, shows some outstanding progresses made by state-of-the art ML technologies for inorganic solid-state materials. The second speaker, Prof. Yoshida at the Institute of Statistical Mathematics, presents emerging new technologies for creative design and discovery of new functional molecules.

Title: Deep Learning for Conversational AI

Presenter: Yun-Nung (Vivian) Chen (National Taiwan University)

Abstract:

In the past decade, conversational systems have been the most prominent component in today's virtual personal assistants. The classic dialogue systems have rather complex and/or modular pipelines. The advance of deep learning technologies has recently risen the applications of neural models to dialogue modeling. However, how to build a conversational AI that can satisfy users' needs is still challenging. Hence, this tutorial is designed to focus on an overview of the conversational system development while describing most recent research for building dialogue systems and summarizing the challenges, in order to allow researchers to study the potential improvements of the state-of-the-art conversational AI.

Title: Text Representation Learning and Compositional Semantics

Presenter: Aleksandr Drozd (Riken), Anna Rogers (University of Massachusetts (Lowell))

Abstract:

Natural language processing is a fast-growing field, with a rapid evolution of approaches and models. In the last 6 years, we have come a long way from word embeddings to contextualized representations to pre-trained transformers, with numerous success stories for NLP system performance on question answering, text classification, machine translation and other tasks. Despite the successes, we are still very far from reliable verbal reasoning, and one of the unresolved issues is semantic compositionality. It is not only a practical challenge, but also a theoretical one, as there is still no consensus on what a compositional representation of morphologically complex word, phrase or a sentence should be like. This tutorial provides an introduction to both state-of-the-art NLP models and aspects of linguistic theory in which they are explicitly or implicitly grounded, particularly compositionality. We will overview of the latest proposals for representing words, sentences, and texts, as well as the discussion of interpretable components in meaning representations. In addition, we will discuss some of the problems with the current evaluation methodology and frequently used benchmarks.

Workshops

ACML 2019 workshop on Machine Learning for Trajectory, Activity, and Behavior (ACML-TAB2019)

Organisers: Toru Tamaki (Hiroshima University), Keisuke Fujii (Nagoya University), Tsubasa Hirakawa (Chubu University)

Recent advances in sensing technology have made it possible to collect vast amounts of trajectories, activities and behavior data from humans, animals, and vehicles. Smart devices and visual tracking are used to capture the data of players in the sports scene and vehicles in the city for skill assessment or resource allocations. Small GPS and acceleration loggers collect behavioral data from animals in the wild, such as birds and bats, to better understand the ecology of animals. Therefore, machine learning techniques have been developed to recognize, analyze, and predict the trajectory, activity, and behavior of various targets.

This workshop provides a place for engineers, computer scientists, biologist, and neuroscientists to discuss machine learning and related methods for trajectory, activity, and behavior data collected from various sources, such as humans, animals, insects, and automobiles. The topics of interest include, but are not limited to:

- Machine learning, time series analysis, data mining, and knowledge extraction for trajectory/activity/behavior data
- Modeling, collecting, data preparation and labeling for trajectory/activity/behavior data
- Systems and applications of monitoring and recognition systems for trajectory/activity/behavior data
- Localization, recognition, prediction, and visualization for trajectory/activity/behavior data

ACML 2019 RIKEN AIP workshop

Organisers: Miao Xu (RIKEN), Hiromi Arai (RIKEN), Masashi Sugiyama (RIKEN / University of Tokyo)

The Center for Advanced Intelligence Project (AIP), RIKEN was founded in 2016 as a research center for the MEXT-AIP Project. We have fully started our research activities in 2017, mainly in our newly opened Nihonbashi Office. In the Center for AIP, we set up three research groups:

- Generic Technology Research Group
- Goal-Oriented Technology Research Group

- Artificial Intelligence in Society Research Group

Together with various companies, universities, research institutes and projects, we are tackling the following five activities: development of fundamental technology, acceleration of scientific research, the solution to societal problems, analysis of ethical, legal and social issues of AI, development of AI researchers and data scientists. This RIKEN AIP Workshop@ACML19 aims at introducing the forefront machine learning researches conducted at RIKEN AIP to ACML participants, and incurring fruitful discussions between researchers inside and outside of AIP about existing key progress and promising new directions. Topics of this workshop cover a variety of theoretical and applied researches, including deep learning, reinforcement learning, weakly-supervised learning, Bayesian optimization, AI for social good, AI for disaster prevention. Our workshop invites leading researchers from all three research groups and also hosts poster presenters with a wide range of perspectives and interests.

ACML 2019 Weakly-supervised Learning Workshop

Organisers: Bo Han (RIKEN), Gang Niu (RIKEN), Quanming Yao (4Paradigm/Hong Kong University of Science and Technology), Giorgio Patrini (DeepTrace), Aditya Menon (Google AI), Clayton Scott (University of Michigan), Masashi Sugiyama (RIKEN/University of Tokyo)

The focus of this workshop is three classical types of weak supervision: incomplete supervision, inexact supervision and inaccurate supervision. Specifically, incomplete supervision considers a subset of training data given with ground-truth labels while the other data remain unlabeled, such as semi-supervised learning and positive-unlabeled learning. Inexact supervision considers the situation where some supervision information is given but not as exacted as desired, i.e., only coarse-grained labels are available. For example, if we are considering to classify every pixel of an image, rather than the image itself, then ImageNet becomes a benchmark with inexact supervision. Besides, multi-instance learning belongs to inexact supervision, where we do not exactly know which instance in the bag corresponds to the given ground-truth label. Inaccurate supervision considers the situation where the supervision information is not always the ground-truth, such as learning with noisy labels.

Moreover, this workshop covers two emerging types of weak supervision: cross-domain supervision and imperfect demonstration. Cross-domain supervision considers the situation where the supervision information is scarce or even non-existent in the current domain but can be possibly derived from other domains. Examples of cross-domain supervision appear in zero-/one-/few-shot learning, where external knowledge from other domains is usually used to overcome the problem of too few or even no supervision in the original domain. Imperfect demonstration considers the situation for inverse reinforcement learning and imitation learning, where the agent learns with imperfect or non-expert demonstrations. For example, AlphaGo learns a policy from a sequence of states and actions (expert demonstra-

tion). Even if an expert player (human or agent) wins a game, it is not guaranteed that every action in the sequence is optimal.

This workshop will discuss the fundamental theory of weakly-supervised learning. Although theories of statistical weakly-supervised learning already exist, extending these results for deep weakly-supervised learning is still a challenge. Besides, this workshop also discusses on broad applications of weakly-supervised learning, such as weakly-supervised object detection (computer vision), weakly-supervised sequence modeling (natural language processing), weakly-supervised cross-media retrieval (information retrieval), and weakly-supervised medical image segmentation (healthcare).

ACML 2019 Workshop on Statistics & Machine Learning Researchers in Japan

Organisers: Kota Matsui (RIKEN), Yuta Umezu (Nagoya Institute of Technology)

In recent years, machine learning and data science have attracted explosive interest in Japan. These elemental technologies are spreading into fields that have been developed independently from ML such as medicine and material science, so interest from not only academia and tech companies but also other industries such as manufacturers becomes very high. Consequently, the number of participants in the Japanese domestic machine learning conference is substantially increasing like the top conferences.

This workshop aims to facilitate exchanges between all those interested in ML and DS in Japan and Asian experts who are active in these fields. The WS would be a great chance for sharing ideas and expertise among participants, encouraging industries or researchers in other fields to get suggestions and advices from ML experts, and also fostering connections and possible collaborations between Japanese and International ML communities.

Conference Organisation

Organising Committee

Honorary Chairs:
Takashi Washio (Osaka University, Japan) and Masashi Sugiyama (RIKEN / University of Tokyo, Japan)

General Co-Chairs:
Geoff Holmes (University of Waikato, New Zealand) and Ichiro Takeuchi (Nagoya Institute of Technology / RIKEN, Japan)

Program Co-Chairs:
Wee Sun Lee (National University of Singapore, Singapore) and Taiji Suzuki (University of Tokyo / RIKEN, Japan)

Journal Track Co-Chairs:
Kee-Eung Kim (KAIST, Korea Republic) and Jun Zhu (Tsinghua University, China)

Workshop Co-Chairs:
Taesup Moon (Sungkyunkwan University, Korea Republic) and Mathieu Blondel (NTT, Japan)

Tutorial Co-Chairs:
Yu-Feng Li (Nanjing University, China) and Emtiyaz Khan (RIKEN, Japan)

Publicity Chair:
Marco Cuturi (Google, France)

Publication Chair:
Shinichi Nakajima (Technische Universität Berlin, Germany)

Local Arrangements Co-Chairs:
Kota Matsui (RIKEN, Japan), Noriaki Hashimoto (Nagoya Institute of Technology, Japan), and Yuta Umezu (Nagoya Institute of Technology, Japan)

Web Chair:
Hiroyuki Hanada (RIKEN, Japan)

Steering Committee

- Wray Buntine (Monash University, Australia)
- Chih-Jen Lin (National Taiwan University, Taiwan)
- Wee Sun Lee (Co-Chair, National University of Singapore, Singapore)
- Hang Li (Toutiao, China)
- Geoffrey Holmes (University of Waikato, New Zealand)
- Kee-Eung Kim (KAIST, Korea Republic)
- Hiroshi Motoda (AFOSR/AOARD and Osaka University, Japan)
- Masashi Sugiyama (Chair, RIKEN/The University of Tokyo, Japan)
- Zhi-Hua Zhou (Chair, Nanjing University, China)

Senior Program Committee

- Wray Buntine (Monash University, Australia)
- Seungjin Choi (POSTECH, Korea)
- Marco Cuturi (ENSAE/CREST, France)
- Khan Emtiyaz (RIKEN, Japan)
- Stephen Gould (Australian National University, Australia)
- Quanquan Gu (University of California, Los Angeles)
- Yuhong Guo (Carleton University, Canada)
- Kohei Hatano (Kyusyu University, Japan)
- Steven Hoi (Singapore Management University, Singapore)
- Chenping Hou (National University of Defense Technology, China)
- Qinghua Hu (Tianjin University, China)
- Xiaolin Hu (Tsinghua University, China)
- Minlie Huang (Tsinghua University, China)
- Sheng-Jun Huang (Nanjing University of Aeronautics and Astronautics, China)
- Takafumi Kanamori (Tokyo Institute of Technology, Japan)
- Junmo Kim (KAIST, Korea Republic)
- James Kwok (The Hong Kong University of Science and Technology, Hong Kong)

- Yanyan Lan (Institute of Computing Technology, CAS, China)
- Yu-Feng Li (Nanjing University, China)
- Hsuan-Tien Lin (National Taiwan University, Taiwan)
- Shou-De Lin (National Taiwan University, Taiwan)
- Zhouchen Lin (Peking University, China)
- Shinichi Nakajima (Technische Universität Berlin, German)
- Yung-Kyun Noh (Seoul National University, Korea Republic)
- Alice Oh (KAIST, Korea Republic)
- Takayuki Okatani (Tohoku University/RIKEN AIP, Japan)
- Takayuki Osogami (IBM Research - Tokyo, Japan)
- Sinno Pan (Nanyang Technological University, Singapore)
- Bernhard Pfahringer (University of Waikato, New Zealand)
- Tomas Pfister (Google, USA)
- Dinh Phung (Deakin University, Australia)
- Tao Qin (Microsoft Research Asia, China)
- Joseph Salmon (Université de Montpellier, France)
- Hang Su (Tsinghua University, China)
- Yasuo Tabei (RIKEN AIP, Japan)
- Toru Tamaki (Hiroshima University, Japan)
- Ke Tang (Southern University of Science and Technology, China)
- Truyen Tran (Deakin University, Australia)
- Ivor Tsang (University of Technology Sydney, Australia)
- Grigorios Tsoumakas (Aristotle University of Thessaloniki, Greece)
- Shunji Umetani (Osaka University/RIKEN, Japan)
- Makoto Yamada (RIKEN AIP, Japan)
- Yang Yu (Nanjing University, China)
- Junping Zhang (Fudan University, China)
- Kun Zhang (Carnegie Mellon University, USA)
- Min-Ling Zhang (Southeast University)

Sponsors

We would like to thank Denso, NEC, Calm Island, Nvidia, Japan Digital Design, Quantumblack, SoftBank, Toshiba, Toyota, and Yahoo Japan.