Foreword

Welcome to the seventh edition of the JMLR Conference and Workshop Proceeding Series for the Asian Conference on Machine Learning. This volume contains 28 papers accepted to the Seventh Asian Conference on Machine Learning (ACML 2015) held in Hong Kong, from 20 to 22 November 2015. This volume continues the ACML tradition of having high-quality and original research papers in the area of machine learning following six previous successful events held in China, Japan, Taiwan, Singapore, Australia, and Vietnam respectively. 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. Submissions from other than the Asia-Pacific regions were also highly encouraged.

This year, we continued the tradition of having two cycles and double-blind review; each cycle had its own submission deadline. Papers that could not be accepted in the first cycle, but received positive feedback, were correctable and could be accepted after careful revision, were encouraged to resubmit in the second cycle, allowing the reviewer's comments to be addressed. After removing invalid submissions, there were 96 submissions, of which 28 were accepted into the main program, accounting for an acceptance rate of 29%. A strict double-blind reviewing process was enforced, and each paper was assigned with one meta-reviewer, and at least 3 reviewers. To maintain quality, each reviewer was allocated no more than 5 papers. In total, there were 30 senior program committee members and 76 program committee members to provide expert opinion and reviews. Their contributions were essential to the quality and standard of papers selected for the conference. Without their contribution, the conference would have not been possible. Their names are acknowledged in the following pages. Finally, the Program Co-chairs considered all the reviews and meta-reviews by senior program committee members to make the final decisions for the papers.

All accepted papers received both an oral and poster presentation, and are published in this volume. Following the tradition of previous ACML(s), this year's conference was also a single track. The submissions covered a broad range of topics, including theoretical analyses, probabilistic models, deep learning, and applications to real world problems.

In addition to the submitted papers, we were pleased to have two keynotes from leading experts, Xiaotie Deng (Shanghai Jiaotong University) and Eric Xing (Carnegie Mellon University), as well as three invited talks from Rong Jin (Alibaba), Ruslan Salakhutdinov (University of Toronto), and Dit-Yan Yeung (Hong Kong University of Science and Technology). In addition to main program, we had one day of tutorials and workshops prior to the main conference organized by the Tutorial and Workshops Chairs. Four tutorials were delivered " A New Look at the System, Algorithm and Theory Foundations of Distributed Machine Learning" by Qirong Hou, "Entity Search, Recommendation and Understanding" by Hao Ma, "Big Data Analytics: Optimization and Randomization" by Tianbao Yang and "Causal Discovery and Inference: Traditional Approach and Recent Advances" by Kun Zhang and Jiji Zhang. In parallel to the tutorials, we had two workshops: "Deep Learning" organized by Zhengdong Lu, Zheng Zhang, and Shuicheng Yan, and "Machine Learning in

China" organized by Yang Yu, Ping Luo, and Yu-Feng Li. We thank all the speakers and organizers for putting together a fantastic program.

ACML Steering Committee Chair Zhi-Hua Zhou provided valuable advice and support during the whole process. The General Chairs Hang Li and Irwin King took care of organizing issues and many other things to make sure the event ran smoothly. The Publication Chairs, Kaizhu Huang and Paul Pang have done a great job in producing the conference proceedings. Our special thanks go to the Local Co-Chairs Raymond Wong and Haiqing Yang, and the local teams whose contributions were indispensable in making the event run so well.

Last but not least, a big thank you to all participants of ACML 2015 who made it such a great event!

November 2015

ACML 2015 Program Co-Chairs

Geoffrey Holmes
Waikato University, New Zealand
geoff@waikato.ac.nz

Tie-Yan Liu
Microsoft Research, China
tyliu@microsoft.com

Keynote Talks

Title: Game Theoretic Understanding of Social Economic System Design

Presenter: Prof. Xiaotie Deng (Shanghai Jiaotong University)

Abstract:

We study competition, cooperation and coordination in system design over the Internet, focusing on strategic behaviour of participating agents in response to system algorithms and protocols of learning, regulating and controlling. We discuss algorithmic game theoretical methodology and approaches to address such challenges.

Bio:

Dr. Xiaotie Deng got his BSc from Tsinghua University, MSc from Chinese Academy of Sciences, and PhD from Stanford University. He is currently a Zhiyuan Chair Professor of Shanghai Jiaotong University. He taught in the past at University of Liverpool, City University of Hong Kong, and York University. Before that, he was an NSERC international fellow at Simon Fraser University. Deng's current research focuses on algorithmic game theory, with applications to Internet Economics. His works cover online algorithms, parallel algorithms, and combinatorial optimization. He is an ACM fellow for his contribution to the interface of algorithms and game theory.

Title: Distributed Machine Learning on Big Data

Presenter: Prof. Eric P. Xing (Carnegie Mellon University)

Abstract:

In many modern applications such as web-scale content extraction via topic models, genome-wide association mapping via sparse regression, and image understanding via deep neural networks, one needs to handle large-scale machine learning (ML) problems running on distributed system with multiple CPU/GPU cores or computer nodes. A key issue of both theoretical and engineering interest in building effective distributed systems are the so-called bridging models, which specify how parallel workers should be coordinated. In this talk, I discuss the unique challenges and opportunities in designing these models for ML, and present a number of new bridging models including the stale synchronous parallel (SSP) model, the structure-aware parallel (SAP) model that can speed up both data- and model-parallel ML programs at orders of magnitudes over the conventional models used in existing distributed systems, and enjoy provable correctness guarantee. I will introduce the Petuum framework built on such models for Distributed Machine Learning on Big Data that offers low cost and fast solutions to industry-scale problems in text modeling (topic model with 1M topics), personalized genome medicine (sparse regression on 100M dimensions), and computer vision (deep neural network with billions of parameters).

Bio:

Dr. Eric Xing is a Professor of Machine Learning in the School of Computer Science at Carnegie Mellon University, and the director of the CMU Center for Machine Learning and Health under the Pittsburgh Health Data Alliance. His principal research interests lie in the development of machine learning and statistical methodology; especially for solving problems involving automated learning, reasoning, and decision-making in high-dimensional, multimodal, and dynamic possible worlds in social and biological systems. Professor Xing received his Ph.D. in Computer Science from UC Berkeley. He is an associate editor of the Annals of Applied Statistics (AOAS), the Journal of American Statistical Association (JASA), the IEEE Transaction of Pattern Analysis and Machine Intelligence (PAMI), the PLoS Journal of Computational Biology, and an Action Editor of the Machine Learning Journal (MLJ), the Journal of Machine Learning Research (JMLR). He is a member of the DARPA Information Science and Technology (ISAT) Advisory Group, and a Program Chair of ICML 2014.

Invited Talks

Title: Making the Impossible Possible: Randomized Machine Learning Algorithms for Big Data

Presenter: Dr. Rong Jin (Alibaba)

Abstract:

We are continuing to encounter an explosive growth in data: the number of web pages grows from 300 million in 1997 to 50 billion in 2013; about 10 billion images are indexed by Google and 6 billion videos are indexed by YouTube; Alibabas ecommerce platform receives billions of requests on a daily basis. This data explosion poses a great challenge in data analysis. Randomized algorithms have attracted significant interests in the recent studies of machine learning, mostly due to its computational efficiency. But, on the other hand, the formal limitations of randomized algorithms have been established for various learning tasks, making them less effective in exploiting the massive amount of data that is available to computer programs. In this talk, I will discuss, based on two examples, how to overcome the limitation of randomized machine learning algorithms by exploiting either the side information or prior knowledge of data. We have shown, both theoretically and empirically, that with a slight modification, it is possible to dramatically improve the effectiveness of randomized algorithms for machine learning. I will also introduce the successful cases of applying randomized algorithms in Alibaba.

Bio:

Dr. Rong Jin is currently a Principle Engineer at Alibaba. He was a professor in the Department of Computer Science and Engineering at Michigan State University. His research is focused on statistical machine learning and its application to big data analysis. He has published over 200 technical papers, most in prestigious conferences (NIPS, ICML, KDD, SIGIR, CVPR, ICCV) and journals (TPAMI, JMLR, JML, and TKDD). Dr. Jin has served as area chair for NIPS 2013 and SIGIR 2008, and associate editor for TPAMI and ACM KDD. He received his Ph.D in Computer Science from Carnegie Mellon University in 2003, NSF Career Award in 2006 and the best student paper award from COLT in 2012.

Title: Recent Advances in Deep Learning: Learning Structured, Robust, and Multimodal Deep Models

Presenter: Prof. Ruslan Salakhutdinov (University of Toronto)

Abstract:

Building intelligent systems that are capable of extracting meaningful representations from high-dimensional data lies at the core of solving many Artificial Intelligence tasks, includ-

ing visual object recognition, information retrieval, speech perception, and language understanding.

In this talk I will first introduce a broad class of deep learning models and show that they can learn useful hierarchical representations from large volumes of high-dimensional data with applications in information retrieval, object recognition, and speech perception. I will then describe a new class of more complex models that combine deep learning models with structured hierarchical Bayesian models and show how these models can learn a deep hierarchical structure for sharing knowledge across hundreds of visual categories. Finally, I will introduce deep models that are capable of extracting a unified representation that fuses together multiple data modalities. I will discuss models that can generate natural language descriptions (captions) of images, as well as generate images from captions (using attention mechanism). I will show that on several tasks, including modelling images and text, video and sound, these models significantly improve upon many of the existing techniques.

Bio:

Dr. Ruslan Salakhutdinov received his PhD in machine learning (computer science) from the University of Toronto in 2009. After spending two post-doctoral years at the Massachusetts Institute of Technology Artificial Intelligence Lab, he joined the University of Toronto as an Assistant Professor in the Department of Computer Science and Department of Statistics.

Dr. Salakhutdinov's primary interests lie in statistical machine learning, Deep Learning, probabilistic graphical models, and large-scale optimization. He is an action editor of the Journal of Machine Learning Research and served on the senior programme committee of several learning conferences including NIPS and ICML. He is the recipient of the Early Researcher Award, Connaught New Researcher Award, Alfred P. Sloan Research Fellowship, Microsoft Research Faculty Fellowship, Google Faculty Research Award, and is a Fellow of the Canadian Institute for Advanced Research.

Title: Bayesian Deep Learning for Integrated Intelligence: Bridging the Gap between Perception and Inference

Presenter: Prof. Dit-Yan Yeung (Hong Kong University of Science and Technology)

Abstract:

While perception tasks such as visual object recognition and text understanding play an important role in human intelligence, the subsequent tasks that involve inference, reasoning and planning require an even higher level of intelligence. The past few years have seen major advances in many perception tasks using deep learning models. For higher-level inference, however, probabilistic graphical models with their Bayesian nature are still more powerful and flexible. To achieve integrated intelligence that involves both perception and inference, we have been exploring along a research direction, which we call Bayesian deep learning,

to tightly integrate deep learning and Bayesian models within a principled probabilistic framework. In this talk, some of our recent work on Bayesian deep learning with various applications in recommendation and representation learning will be presented.

Bio:

Dr. Yeung received his BEng degree in electrical engineering and MPhil degree in computer science from the University of Hong Kong (HKU), and PhD degree in computer science from the University of Southern California (USC) in 1989. He started his academic career in the same year as an assistant professor at the Illinois Institute of Technology (IIT) in Chicago. He then joined the Hong Kong University of Science and Technology (HKUST) where he is now a full professor in computer science and engineering.

Dr. Yeung's research interests include computational and statistical approaches to machine learning and artificial intelligence. He is also interested in applying machine learning techniques to computer vision and social computing.

Tutorials

Title: A New Look at the System, Algorithm and Theory Foundations of Distributed Machine Learning

Presenter: Dr. Qirong Ho (A*STAR and Singapore Management University)

Abstract:

The rise of Big Data has led to new demand for Machine Learning (ML) systems to learn complex models often with millions to billions of parameters that promise adequate capacity to analyze massive datasets and offer predicative functions thereupon. For example, in many modern applications such as web-scale content extraction via topic models, genome-wide association mapping via sparse structured regression, and image understanding via deep neural networks, one needs to handle BIG ML problems that threaten to exceed the limit of current architectures and algorithms. In this tutorial, we present a systematic overview of modern scalable ML approaches for such applications — the insights and challenges of designing scalable and parallelizable algorithms for working with Big Data and Big Model; the principles and architectures of building distributed systems for executing these models and algorithms; and the theory and analysis necessary for understanding the behaviors and providing guarantees of these models, algorithms, and systems.

We present a comprehensive, principled, yet highly unified and application-grounded view of the fundamentals and strategies underlying a wide range of modern ML programs practiced in industry and academia, beginning with introducing the basic algorithmic roadmaps of both optimization-theoretic and probabilistic-inference methods — two major workhorse algorithmic engines that power nearly all ML programs, and the technical developments therein aiming at large scales built on algorithmic acceleration, stochastic approximation, and parallelization. We then turn to the challenges such algorithms must face in a practical distributed computing environment due to memory/storage limit, communication bottleneck, resource contention, straggler, etc., and review and discuss various modern parallelization strategies and distributed frameworks that can actually run these algorithms at Big Data and Big Model scales, while also exposing the theoretical insights that make such systems and strategies possible. We focus on what makes ML algorithms peculiar, and how this can lead to algorithmic and systems designs that are markedly different from todays Big Data platforms. We discuss such new opportunities in algorithm, system, and theory on parallel machine learning, in real (instead of ideal) distributed communication, storage, and computing environments.

Bio:

Dr. Qirong Ho is a scientist at the Institute for Infocomm Research, A*STAR, Singapore, and an adjunct assistant professor at the Singapore Management University School of Information Systems. His primary research focus is distributed cluster software systems for Machine Learning at Big Data scales, with a view towards correctness and performance guarantees. In addition, Dr. Ho has performed research on statistical models for large-scale network analysis — particularly latent space models for visualization, community detection, user personalization and interest prediction — as well as social media analysis on hyperlinked documents with text and network data. Dr. Ho received his PhD in 2014, under Eric

P. Xing at Carnegie Mellon University's Machine Learning Department. He is a recipient of the 2015 KDD Dissertation Award (runner-up), and the Singapore A*STAR National Science Search Undergraduate and PhD fellowships.

Title: Entity Search, Recommendation and Understanding

Presenter: Dr. Hao Ma (Microsoft Research, USA)

Abstract:

Recent years have witnessed rapidly increasing interests on the research field of semantic search. Knowledge base powered entity search and recommendation experience has been widely adopted by major search engine companies. Entity search, recommendation and understanding techniques differ significantly from those in traditional search problems due to the introduction of knowledge base. The heterogeneity, semantic richness and large-scale nature of knowledge base make traditional approaches less effective. In this tutorial, we provide a detailed interdisciplinary introduction on how entity search and recommendation work and how various entity understanding methods based on Machine Learning, Natural Language Processing and Information Retrieval techniques could further improve entity related experience.

This tutorial consists of four major parts. In the first part, we give a brief introduction on entities and knowledge bases. We also show how we collect information from different data sources as well as how we infer users' interests on specific entities. In the second part, we demonstrate various entity search and recommendation applications we developed and productionized in Bing and Microsoft, including entity recommendation, natural language interpretation of recommendation, attribute ranking, carousel ranking, entity exploration, factoid answers, conversational search, semantic question and answering, etc. The architectures, challenges, and corresponding solutions on these systems will also be briefly introduced in the second part of the tutorial. The third part will give a deep dive on the algorithms that are related to entity search and recommendation, including basic non-personalized search and recommendation algorithms as well as recommendation models that tailor related entities to an individual search user's unique taste and preference. The fourth part of this tutorial will focus on how to further improve semantic recommendation and search experience by employing other entity understanding techniques, including entity linking to the knowledge bases, question and answering on the web documents, etc.

The tutorial will conclude by summarizing and reflecting back on the semantic search applications that users are experiencing on the Web and posit that what we have presented in the tutorial is just a tip of the iceberg to a whole area of exciting and dynamic research that is worthy of more detailed investigation for many years to come.

Bio:

Dr. Hao is a Researcher at the Internet Services Research Center, Microsoft Research at Redmond, USA. He obtained his Ph.D. in Computer Science at The Chinese University of Hong Kong. His research interests include Natural Language Processing, Machine Learning, Information Retrieval, Recommender Systems and Social Network Analysis. Most recently, Dr. Ma has been working on entity related research problems and applications. He designed the core learning algorithms that powered both Bing's and Microsoft's entity experience, including entity recommendation, attributes ranking, interpretation, exploration, carousel ranking, question and answering, etc. He has published more than 40 research papers in prestigious conferences and journals, including WWW, SIGIR, WSDM, CIKM, AAAI, TOIS, TKDE, TMM, TIST, etc. Some of his research output has been widely reported by popular News media, like MIT Technology Review, Search Engine Land, etc. Dr. Ma is also in the winning team that won the Microposts Entity Linking Challenge in WWW 2014.

Title: Big Data Analytics: Optimization and Randomization

Presenter: Prof. Tianbao Yang (University of Iowa)

Abstract:

As the scale and dimensionality of data continue to grow in many applications of data analytics (e.g., bioinformatics, finance, computer vision, medical informatics), it becomes critical to develop efficient and effective algorithms to solve numerous machine learning and data mining problems. This tutorial will focus on simple yet practically effective techniques and algorithms for big data analytics. In the first part, we plan to present the state-of-the-art large-scale optimization algorithms, including various stochastic gradient descent methods, stochastic coordinate descent methods and distributed optimization algorithms, for solving various machine learning problems. In the second part, we will focus on randomized approximation algorithms for learning from large-scale data. We will discuss i) randomized algorithms for low-rank matrix approximation; ii) approximation techniques for solving kernel learning problems; iii) randomized reduction methods for addressing the high-dimensional challenge. Along with the descriptions of algorithms, we will also present some empirical results to facilitate understanding and comparison between different algorithms.

Bio:

Dr. Tianbao Yang is currently an assistant professor at the University of Iowa (UI). He received his Ph.D. degree in Computer Science from Michigan State University in 2012. Before joining UI, he was a researcher in NEC Laboratories America at Cupertino (2013-2014) and a Machine Learning Researcher in GE Global Research (2012-2013), mainly focusing on

developing distributed optimization system for various classification and regression problems. Dr. Yang has board interests in machine learning and he has focused on several research topics, including large-scale optimization in machine learning, online optimization and distributed optimization. His recent research interests revolve around randomized algorithms for solving big data problems. He has published over 25 papers in prestigious machine learning conferences and journals. He has won the Mark Fulk Best student paper award at 25th Conference on Learning Theory (COLT) in 2012.

Title: Causal Discovery and Inference: Traditional Approach and Recent Advances

Presenter: Prof. Kun Zhang (Carnegie Mellon University) and
Prof. Jiji Zhang (Lingnan University)

Abstract:

Causality is a fundamental notion in science, and plays an important role in explanation, prediction, decision making, and control. Recently, interesting advances were made in machine learning for tackling some long-standing problems in causality, such as how to distinguish cause from effect given two random variables. On the other hand, causal models provide compact descriptions of the properties of data distributions, and it has been demonstrated that causal knowledge can facilitate various machine learning tasks. This tutorial talk aims to give a broad coverage of emerging approaches to causal discovery and causal inference from i.i.d data and from time series, with both theoretical and practical results, and related issues.

We start with the constraint-based approach to causal discovery, which relies on the conditional independence relationships in the data, and discuss its wide applicability as well as its drawbacks. We then talk about the identifiability of the causal structure implied by appropriately defined functional causal models; in particular, in the two-variable case, under what conditions (and why) is the causal direction between the two variables identifiable? We show that the independence between the noise and causes, together with appropriate structural constraints on the functional form, makes it possible. We will focus on the linear non-Gaussian causal model and the post-nonlinear causal model, due to their simplicity and generality, respectively. Next, we report some recent advances in causal discovery from time series. Assuming that the causal relations are linear with non-Gaussian noise, we focus on two problems which are traditionally difficult to solve, namely, causal discovery from subsampled data and that in the presence of confounding time series.

Finally, we address how causal knowledge is able to facilitate understanding and solving certain machine learning problems. We consider two learning problems–semi-supervised learning and domain adaptation–from a causal point of view, and discuss the implications of causal knowledge that help understand or solve the problems better. A number of open questions in the field of causal discovery and inference are also provided.

Bio:

Dr. Kun Zhang is an assistant professor in the philosophy department at Carnegie Mellon University (CMU). Before joining CMU, he was a senior research scientist at Max-Planck Institute for Intelligent Systems, Germany. He got his Ph.D from Chinese University of Hong Kong and then worked at University of Helsinki as a postdoctoral fellow. His main research interests include causal discovery, machine learning, and large-scale data analysis. He has served as a co-organizer of a series of workshops to foster interdisciplinary research in

Dr. Jiji Zhang is an associate professor of philosophy at Lingnan University. He got his PhD from Carnegie Mellon University in 2006, and subsequently taught at the California Institute of Technology before moving to Hong Kong in 2008. His research is interdisciplinary in nature, and centers on methodological, epistemological, and logical issues in causal inference and statistical inference. His work has been published in some venues in machine learning and statistics as well as those in philosophy.

Workshops

Deep Learning

Organisers: Zhengdong Lu, Zheng Zhang, Shuicheng Yan

The world has witnessed the resurgence of neural networks, or as now it is called, "Deep Learning", in the past years. This time, they came back not as promising methods that work on toy data, but as the driving force in the revolution of many application domains, staring with automatic speech recognition, then computer vision, and now natural language processing and probably symbolic artificial intelligence.

Deep learning has gone beyond just deep neural networks or particular approaches to classification and regression. It now stands for a new way of thinking many learning problems. Indeed, now we have end-to-end learning paradigm that directly connects vision and robot manipulation, or systems that aim to reason over complicated natural language facts, which previously required heavy feature engineering and cumbersome models.

This workshop aims to bring together researchers in deep learning, particularly in Asia, to share their research, view points, vision, and probably more importantly their questions and confusion. We are here to discuss not only how deep learning reshaped machine learning, but also how to reshape deep learning–it is after all still young, and again as people argued more than 20 years ago, with great great expectation.

Machine Learning in China (MLChina'15)

Organisers: Yang Yu, Ping Luo, Yu-Feng Li

During the past decade, machine learning researches in China have been growing in a blooming way. This is witnessed by the increasing number of works appeared in major machine learning related conferences and journals, and also numerous successful applications of machine learning techniques in major Chinese high-tech companies such as Huawei, Tencent, Baidu, Alibaba, etc. There are also many domestic machine learning conferences regularly held in China which attract a significant number of participants, such as the biennial Chinese Conference on Machine Learning, the annual Chinese Workshop on Machine Learning and Applications, the annual Chinese Vision and Learning Seminar, etc.

To take full advantage of the opportunity that ACML is to be held in Hong Kong, a dedicated full-day workshop for machine learning researchers and practitioners in China is organized. This workshop would be a great chance for sharing ideas and expertise among interested participants, encouraging students and junior researchers to get suggestions and advices from senior experts, and also fostering connections and possible collaborations between Chinese and International machine learning communities.

Conference Organisation

Organising Committee

General Co-chairs:

Irwin King (The Chinese University of Hong Kong, Hong Kong)
Hang Li (Huawei Noahs Ark Lab, Hong Kong)

Program Co-chairs:

Geoffrey Holmes (Waikato University, New Zealand)
Tie-Yan Liu (Microsoft Research Asia, China)

Local Arrangements Co-Chairs

Raymond Wong (Hong Kong University of Science and Technology, Hong Kong)
Haiqin Yang (The Chinese University of Hong Kong, Hong Kong)

Workshop Co-chairs:

Zhengdong Lu (Huawei Noahs Ark Lab, Hong Kong)
Zenglin Xu (University of Electronic Science and Technology of China, China)

Tutorial Co-chairs:

Yiu-Ming Cheung (Hong Kong Baptist University)
Chi-Sing Leung (City University of Hong Kong, Hong Kong)

Publication Co-Chairs:

Kaizhu Huang (Xian Jiaotong-Liverpool University, China)
Paul Pang (Unitec Institut of Technology, New Zealand)

Steering Committee

- Zhi-Hua Zhou, Chair (Nanjing University, China)
- Tu Bao Ho, Co-chair (Japan Advanced Institute of Science and Technology, Japan)
- Wray Buntine (National ICT Australia, Australia)
- Tom Dietterich (Oregon State University, USA)
- Wee Sun Lee (National University of Singapore, Singapore)
- Chih-Jen Lin (National Taiwan University, Taiwan)
- Bernhard Pfahringer (Waikato University, New Zealand)
- Masashi Sugiyama (The University of Tokyo, Japan)
- Takashi Washio (Osaka University, Japan)
- Geoff Webb (Monash University, Australia)
- Qiang Yang (Hong Kong University of Science and Technology, Hong Kong)

Senior Program Committee

- Wray Buntine (Monash University, Australia)
- Kevin Duh (NAIST, Japan)
- Stephen Gould (NICTA/Australian National University, Australia)
- Steven Hoi (Singapore Management University, Singapore)
- Hisashi Kashima (The University of Toyko, Japan)
- Kristian Kersting (Technical University of Dortmund, Germany)
- Kee-Eung Kim (KAIST, South Korea)
- James Kwok (Hong Kong University of Science and Technology, Hong Kong)
- Hoai An Le Thi (University of Metz, France)
- Yuh-Jye Lee (NTUST Taiwan, Taiwan)
- Ping Li (Rutgers University, USA)
- Hsuan-Tien Lin (National Taiwan University, Taiwan)
- Shoude Lin (National Taiwan University, Taiwan)
- Zhengdong Lu (Huawei Technologies, Hong Kong)
- Michael Lyu (The Chinese University of Hong Kong, Hong Kong)

- Alice Oh (KAIST, South Korea)
- Bernhard Pfahringer (University of Waikato, New Zealand)
- Dinh Phung (Deakin University, Australia)
- Tao Qin (Microsoft Research Asia, China)
- Volker Roth (University of Basel, Switzerland)
- Shirish Shevade (Indian Institute of Science, India)
- Masashi Sugiyama (The University of Tokyo,Japan)
- Dat Tran (University of Canberra, Australia)
- Truyen Tran (Deakin University, Australia)
- Liwei Wang (Peking University, China)
- Jun Xu (Chinese Academy of Sciences, China)
- Jieping Ye (Arizona State University, USA)
- Hwanjo Yu (POSTECH, South Korea)
- Junping Zhang (Fudan University, China)
- Min-Ling Zhang (Southeast University, China)

Technical Program Committee

- Aditya Menon (NICTA)
- Atsuyoshi Nakamura (Hokkaido University)
- Behrouz Madahian (University of Memphis)
- Ben Tan (HKUST)
- Bin Li (Wuhan University)
- Bob Durrant (University of Waikato)
- Chen Wang (Cornell University)
- Cho-Jui Hsieh (University of California)
- Christian Bauckhage (Fraunhofer IAIS and University of Bonn)
- Chun-Liang Li (Carnegie Mellon University)
- Dino Ienco (IRSTEA)

- Dongho Kim (University of Cambridge)
- DucSon Pham (Curtin University)
- Eibe Frank (University of Waikato)
- En-Hsu Yen (University of Texas)
- Florian Yger (University of Tokyo)
- Frank Wang (Academia Sinica)
- Gang Li (Deakin University)
- Hanchen Xiong (University of Innsbruck)
- Hongming Shan (Fudan University)
- Hung-Yi Lo (Academia Sinica)
- Ichigaku Takigawa (Hokkaido University)
- Ichiro Takeuchi (Nagoya Institute of Technology)
- Jaesik Choi (UNIST)
- Jiahui Chen (GE Global Research)
- Jiaqi Zhang (Peking University)
- Jiayu Zhou (Samsung Research America)
- Jinhui Yuan (Microsoft Research Asia)
- Jinoh Yu (POSTECH)
- Jinoh Oh (POSTECH)
- Junya Honda (University of Tokyo)
- Kaizhu Huang (Xi'an Jiaotong-Liverpool University)
- Khoat Than (Hanoi University of Science and Technology)
- Kohei Hatano (Kyushu University)
- Krzysztof Dembczynski (Poznan University of Technology)
- Lan Du (Macquarie University)
- Lei Yuan (Dow AgroSciences)
- Lei Wu (GE Global Research)
- Mahito Sugiyama (Osaka University)

- Makoto Yamada (Yahoo)
- Malik Magdon-Ismail (RPI)
- Marco Cuturi (Kyoto University)
- Mark Carman (Monash University)
- Masayuki Karasuyama (Nagoya Institute of Technology)
- Minsuk Chang (KAIST)
- Peilin Zhao (Rutgers University)
- Pinghua Gong (University of Michigan)
- Qian Sun (Arizona State University)
- Ran He (Chinese Academy of Sciences)
- Sangkyun Lee (TU Dortmund)
- Sheng-Jun Huang (Nanjing University of Aeronautics and Astronautics)
- Shinichi Nakajima (Nikon Corporation)
- Shuo Xiang (Arizona State University)
- Sunil Gupta (Deakin University)
- Taiji Suzuki (Tokyo Institute of Technology)
- Takafumi Kanamori (Nagoya University)
- Tam Le (Kyoto University)
- Trung Le (HCMC University of Pedagogy)
- Tsung-Ting Kuo (University of California)
- Vinh Nguyen (University of Melbourne)
- Wanli Ma (University of Canberra)
- Wei Zhang (Fudan University)
- Wei Wang (Nanjing University)
- Wei Chen (Microsoft Research)
- Xianghang Liu (UNSW)
- Xiaotong Yuan (Nanjing University of Information Science and Technology)
- Xuan-Hieu Phan (Vietnam National University)

- Xu-Ying Liu (Southeast University)
- Yangqiu Song (UIUC)
- Yanyan Lan (Institute of Computing Technology)
- Yashu Liu (Arizona State University)
- Yu-Feng Li (Nanjing University)
- Zenglin Xu (University of Science and Technology of China)
- Zheng Wang (Arizona State University)
- Zhiwen Yu (South China University of Technology)
- Zhi-Yong Liu (Chinese Academy of Sciences)

Sponsors

We are grateful to the Platinum Sponsor, Tencent, Gold Sponsors, Baidu and Huawei, and the K. C. Wong Education Foundation for the generous support. We are also thankful for the wonderful facilities provided by the Hong Kong Science and Technology Parks and The Chinese University of Hong Kong in hosting this event.