Foreword

Welcome to the ninth edition of the Proceedings of Machine Learning Research (PMLR) for the Asian Conference on Machine Learning (ACML). This volume contains 41 papers accepted to the Ninth Asian Conference on Machine Learning (ACML 2017) held in Seoul, Korea, from 15th to 17th November 2017. This volume continues the ACML tradition of having high-quality and original research papers in the area of machine learning following eight previous successful events held in China, Japan, Taiwan, Singapore, Australia, Vietnam, Hong Kong, and New Zealand 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. 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 proceedings include papers with authors based in Canada, Germany, France, Finland, Netherland, India, and the Northeast of the USA, as well as the Asia-Pacific region.

This year, we continued the tradition of having two cycles and double-blind review for the proceedings track; 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 reviewers comments to be addressed. In total there were 172 submissions to the conference tracks, of which 41 were accepted into the main program, for an acceptance rate of 23.8%. A strict double-blind reviewing process was enforced, and each paper was assigned with one meta-reviewer, and at least 3 reviewers. As the number of submissions increased, we formed the Program Committee (PC) with 190 PC members (reviewers) and 32 Senior PC members (meta-reviewers) to maintain review quality, and each PC was allocated no more than 3 papers in any one cycle, based on paper bidding, similarity of subject areas, and the Toronto Paper Machine System (TPMS) similarity scores between the submitted paper and the published papers of the Program Committee. Program Committee contributed with expert opinions, which were essential to the quality and standard of papers selected for the conference and, without them the conference would have not been possible. 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.

Following the new innovation of last year, this year ACML also ran an additional journal track. The journal track Co-Chairs Wee Sun Lee and Bob Durrant oversees the reviewing process of 23 submissions, out of which 6 papers are selected for publication in the Springer journal Machine Learning, for an acceptance rate of 26%. The journal track had a sequence of submission dates, which started earlier than the conference track schedule considering the longer process of journal review. For ACML this year the overall number of accepted papers, from both the journal and proceedings tracks, was 47 from 195 submissions for a 24.1% total acceptance rate.

All accepted papers from the two tracks received both an oral and poster presentation, and those accepted after the two cycles of the proceedings track are published in this

volume. The submissions covered a broad range of topics, including theoretical analyses, probabilistic models, large-scale machine learning, weakly-supervised/unsupervised learning, multi-view/multi-task/crowdsourced learning, deep learning, and applications to real world problems.

In addition to the submitted papers, we were very pleased to have two keynotes from leading experts, Bernhard Schölkopf (Max Planck Institute for Intelligent Systems) and Tom Dietterich (Oregon State University), and invited talks from Hang Li (Toutiao), and Eunho Yang (KAIST). In addition to the main program, we had a half day of tutorials and workshops prior to the main conference organized by Tutorial Co-Chairs Krikamol Muandet and Jihun Hamm, Workshop Co-Chairs Hung Bui and Jaesik Choi. Five tutorials were delivered on: High Dimensional Causation Analysis, Deep learning for Biomedicine, Statistical Relational Artificial Intelligence, Machine Learning for Industrial Predictive Analytics, and Distributed Convex Optimization. In parallel to the tutorials, we had three workshops: The First International Workshop on Machine Learning for Artificial Intelligence Platforms (MLAIP), The 2nd Asian Workshop on Reinforcement Learning (AWRL17), and 2017 Annual Korea AI Society Meeting. We thank all of the speakers and organizers for putting together such a fantastic program.

ACML Steering Committee Chair Zhi-Hua Zhou and Co-Chair Masashi Sugiyama provided valuable advice and support during the whole process. The Honorary Chairs Hyeran Byun and Byoung-Tak Zhang provided timely services on various issues which could make serious trouble otherwise. The General Co-Chairs Kee-Eung Kim and Masashi Sugiyama took care of the whole process of organization to make sure the event ran smoothly. Hsuan-Tien Lin and Ivor Tsang took the role Publicity Co-Chairs, and have done a great job for advertising the important conference dates and make many people in machine learning community know about ACML. Our special thanks go to the Local Arrangements Chair Seon Joo Kim for resolving many local issues regarding Venue.

Last but not least, a big thank you to all participants of ACML 2017 who made it such a great event!

November 2017

ACML 2017 Program Co-Chairs

Yung-Kyun Noh
Seoul National University, Republic of Korea
nohyung@snu.ac.kr

Min-Ling Zhang
Southeast University, China
zhangml@seu.edu.cn

Keynote Talks

Title: Causal Learning

Presenter: Bernhard Schölkopf (Max Planck Institute for Intelligent Systems, Germany)

Abstract:

In machine learning, we use data to automatically find dependences in the world, with the goal of predicting future observations. Most machine learning methods build on statistics, but one can also try to go beyond this, assaying causal structures underlying statistical dependences. Can such causal knowledge help prediction in machine learning tasks? We argue that this is indeed the case, due to the fact that causal models are more robust to changes that occur in real world datasets. We discuss implications of causality for machine learning tasks, and argue that many of the hard issues benefit from the causal viewpoint. This includes domain adaptation, semi-supervised learning, transfer, life-long learning, and fairness, as well as an application to the removal of systematic errors in astronomical problems.

Bio:

Bernhard Schölkopf's scientific interests are in machine learning and causal inference. He has applied his methods to a number of different application areas, ranging from biomedical problems to computational photography and astronomy. Bernhard has researched at AT&T Bell Labs, at GMD FIRST, Berlin, and at Microsoft Research Cambridge, UK, before becoming a Max Planck director in 2001. He is a member of the German Academy of Sciences (Leopoldina), and has received the J.K. Aggarwal Prize of the International Association for Pattern Recognition, the Max Planck Research Award (shared with S. Thrun), the Academy Prize of the Berlin-Brandenburg Academy of Sciences and Humanities, and the Royal Society Milner Award.

Title: Combining AI and Visualization to Manage Ecosystems

Presenter: Tom Dietterich (Oregon State University, USA)

Abstract:

As humans occupy virtually all of the planet, we must actively manage ecosystems in order to ensure their sustained functioning. Many ecosystem problems can be formulated as Markov Decision Problems (MDPs) in which the state transitions are provided by a simulator rather than by an explicit function. The two key technical challenges in solving such MDPs are (a) the state spaces are immense and (b) the simulators are very expensive. A third, more political, challenge (c) is that ecosystem management problems involve many stakeholders who often disagree about the objectives to be optimized. To address the first

problem, we employ search in a parameterized policy space. We have obtained excellent results using SMAC, a form of Bayesian Optimization, to find good policies. To address the second problem, we extend the method of Model-Free Monte Carlo (MFMC) to create a surrogate model. We reduce the size of the state space by factoring out exogeneous state variables. To address the third problem, we have built a visualization environment, MDPvis, that allows multiple stakeholders to modify the MDP reward function and explore the behavior of the system. We hope that the combination of visualization and rapid MDP solution will help multiple stakeholders arrive at consensus on how to manage complex ecosystems.

Bio:

Dr. Dietterich (AB Oberlin College 1977; MS University of Illinois 1979; PhD Stanford University 1984) is Professor Emeritus and Director of Intelligent Systems Research in the School of Electrical Engineering and Computer Science at Oregon State University, where he joined the faculty in 1985. Dietterich is one of the pioneers of the field of Machine Learning and has authored more than 130 refereed publications and two books. His research is motivated by challenging real world problems with a special focus on ecological science, ecosystem management, and sustainable development. He is best known for his work on ensemble methods in machine learning including the development of error-correcting output coding. Dietterich has also invented the MAXQ decomposition for hierarchical reinforcement learning.

Invited Talks

Title: Beyond Deep Learning: Combining Neural Processing and Symbolic Processing

Presenter: Hang Li (Toutiao, China)

Abstract:

Recently deep learning has brought significant breakthroughs to natural language processing. I will start the talk by summarizing the strengths and limitations of deep learning in natural language processing. I will then indicate that a hybrid approach combining neural processing (deep learning) and symbolic processing would be necessary and even more powerful for the tasks in natural language processing, particularly question answering. This is because the two paradigms are both advantageous and complementary; symbolic representations are easier to interpret and manipulate, while neural representations are more suitable for dealing with uncertainty in language and noise in data. Finally, I will introduce our recent effort toward developing neural symbolic processing, including building of a hybrid system for question answering from relational database and a hybrid system for generative question answering from knowledge base.

Bio:

Hang Li is director of Toutiao AI Lab, adjunct professors of Peking University and Nanjing University. He is an IEEE Fellow and an ACM Distinguished Scientist. His research areas include information retrieval, natural language processing, machine learning, and data mining. Hang graduated from Kyoto University in 1988 and earned his PhD from the University of Tokyo in 1998. He worked at NEC Research as researcher from 1990 to 2001, Microsoft Research Asia as senior researcher and research manager from 2001 to 2012, and chief scientist and director of Huawei Noahs Ark from 2012 to 2017. He joined Toutiao in 2017. Hang has published three technical books, and more than 120 technical papers at top international conferences including SIGIR, WWW, WSDM, ACL, EMNLP, ICML, NIPS, SIGKDD, AAAI, IJCAI, and top international journals including CL, NLE, JMLR, TOIS, IRJ, IPM, TKDE, TWEB, TIST. He and his colleagues papers received the SIGKDD08 best application paper award, the SIGIR08 best student paper award, the ACL12 best student paper award. Hang worked on the development of several products such as Microsoft SQL Server 2005, Office 2007, Live Search 2008, Bing 2009, Office 2010, Bing 2010, Office 2012, Huawei smartphones 2014 and Huawei smartphones 2017. He has 42 granted US patents. Hang is also very active in the research communities and has served or is serving top international conferences as PC chair, Senior PC member, or PC member, including SIGIR, WWW, WSDM, ACL, NACL, EMNLP, NIPS, SIGKDD, ICDM, IJCAI, ACML, and top international journals as associate editor or editorial board member, including CL, IRJ, TIST, JASIST, JCST.

Title: Beyond Gaussian/Ising Graphical Models

Presenter: Eunho Yang (KAIST, Korea)

Abstract:

Graphical models are the standard toolkit to model interactions between huge number of multiple random variables. Gaussian graphical models for continuous data and Ising (or discrete graphical models) for discrete data are the two most popular instances of undirected graphical models. In this talk, we will discuss limitations of these popular instances of pairwise graphical models with respect to two orthogonal directions: i) restriction on types of data following bell-shpaed Gaussian or categorical properties ii) lack of higher-order interactions due to computational overheads. We will then introduce more general graphical models beyond Gaussian and Ising graphical models, to overcome these limitations.

Bio:

Eunho Yang is an assistant professor at the School of Computing, KAIST. Before joining KAIST, he spent two years at IBM T.J. Watson Research Center as a Research Staff Member. He obtained his Ph.D. in 2014 from the university of Texas at Austin, and did M.S. and B.S from the Seoul National University, Korea in 2006 and 2004, respectively.

His research interests are in statistical machine learning in general with the special focuses on high-dimensional statistics. He is currently developing new theories and algorithms for graphical models and deep learning, with the applications of computational biology and medicine, etc.

Tutorials

Title: High Dimensional Causation Analysis

Presenter: Zhenjie Zhang, Ruichu Cai

Abstract:

Causation analysis is one of the most fundamental research topics in machine learning, which aims to identify causal variables linked to the effect variables from a group of sample in the high dimensional space. The result of causation analysis provides the key insight into the target problem domain, and potentially enable new technologies of genetic therapy in genomic domain and predictive maintenance in IoT domain. Different from existing regression and classification algorithms in machine learning by exploiting correlations among variables, e.g., random forest and deep learning, causation analysis is supposed to unveil the complete and accurate structure of causal influence between every pair of variables in the domain.

It raises extremely large challenges to both mathematical model and algorithm design, because of the exponential complexity growth by extending from correlational dependency to causal dependency. In last decade, huge efforts are devoted to a variety of research frontiers of causation analysis, generating interesting and impressive new understandings under completely different assumptions behind the underlying causal structure generation process. In this tutorial, we introduce the theoretical discoveries on new models, review the significance and usefulness of the new approaches, discuss the applicability of new algorithms on real world applications, and address possible future research directions.

Title: Deep learning for Biomedicine

Presenter: Truyen Tran

Abstract:

The ancient board game of Go, once predicted to remain unsolved for decades, is no longer an AI challenge. Equipped with deep learning, the program AlphaGo of DeepMind beaten a human champion 4 to 1. Indeed, deep learning has enjoyed many record-breaking successes in vision, speech and NLP and has helped boost a huge interest in AI from both academia and industry. Perhaps the next most important area for deep learning to conquer is biomedicine. With the obvious benefits to mankind and a huge industry, deep learning for biomedicine has recently attracted a great attention in both industry and academia. While we hold a great optimism for its success, biomedicine is new to deep learning and there are unique challenges yet to be addressed.

This tutorial consists of two parts. Part I briefly covers main ideas behind state-of-the-art deep learning theory and practice. Part II guides practitioners through designing deep architectures for biomedicine to best address the challenges, some unique to the field.

Title: Statistical Relational Artificial Intelligence

Presenter: Jaesik Choi, Hung Bui

Abstract:

An intelligent agent interacting with the real world will encounter individual people, courses, test results, drugs prescriptions, chairs, boxes, etc., and needs to reason about properties of these individuals and relations among them as well as cope with uncertainty.

Uncertainty has been studied in probability theory and graphical models, and relations have been studied in logic, in particular in the predicate calculus and its extensions. This book examines the foundations of combining logic and probability into what are called relational probabilistic models. It introduces representations, inference, and learning techniques for probability, logic, and their combinations.

This tutorial will provide a gentle introduction into the foundations of statistical relational artificial intelligence, and will realize this by introducing the foundations of logic, of probability, of learning, and their respective combinations.

Title: Machine Learning for Industrial Predictive Analytics

Presenter: Evgeny Burnaev, Maxim Panov

Abstract:

Approximation problems (also known as regression problems) arise quite often in industrial design, and solutions of such problems are conventionally referred to as surrogate models. The most common application of surrogate modeling in engineering is in connection to engineering optimization for predictive analytics. Indeed, on the one hand, design optimization plays a central role in the industrial design process; on the other hand, a single optimization step typically requires the optimizer to create or refresh a model of the response function whose optimum is sought, to be able to come up with a reasonable next design candidate.

The surrogate models used in optimization range from simple local linear regression employed in the basic gradient-based optimization to complex global models employed in the so-called Surrogate-Based Optimization (SBO). Aside from optimization, surrogate modeling is used in dimension reduction, sensitivity analysis, and for visualization of response functions. In this tutorial we are going to highlight main issues on how to construct and apply surrogate models, describe both state-of-the-art techniques and a few novel approximation algorithms, demonstrate the efficiency of the surrogate modeling methodology on several industrial engineering problems.

Title: Distributed Convex Optimization

Presenter: Jun Moon

Abstract:

Convex optimization is a special class of mathematical optimization, where both the objective function and the constraint sets are convex. A lot of optimization problems in engineering, economics and science can be formulated as convex problems, and there are various reliable and efficient algorithms to solve them.

In this tutorial, we consider convex optimization theory. The (tentative) topics for this tutorial are as follows;

- Convex sets and functions
- Unconstrained convex optimization
- Constrained convex optimization
- Distributed convex optimization
- Dynamic programming and linear quadratic control

Workshops

The First International Workshop on Machine Learning for Artificial Intelligence Platforms (MLAIP)

Organisers: Byoung-Tak Zhang, Sungroh Yoon, Dit-Yan Yeung, Sung Kim, Jaesik Choi, Jung-Woo Ha

Recently, several successful AI systems such as Amazon Alexa, Google Assistant, and NAVER X LINE Clova are developed based on AI-assistant platforms. These AI platforms contain several common technologies including speech recognition/synthesis, natural language understanding, image recognition, and dialog recommendation.

Building a successful MLAIP is a challenging mission because it requires a novel combination of heterogeneous machine learning models in a unified framework with efficient data processing. The goals of this workshop is to investigate and advance important topics in Machine Learning for AI Platforms (MLAIPs) further. In addition, we expect to provide the collaboration opportunities to researchers on ML theory on diverse application domains as well as industrial engineers.

The 2nd Asian Workshop on Reinforcement Learning (AWRL17)

Organisers: Tao Qin, Paul Weng, Yang Yu, Zongzhang Zhang

The Asian Workshop on Reinforcement Learning (AWRL) focuses on both theoretical foundations, models, algorithms, and practical applications. In the last few years, we have seen a growing interest in RL of researchers from different research areas and industries. We invite reinforcement learning researchers and practitioners to participate in this world-class gathering. We intend to make this an exciting event for researchers and practitioners in RL worldwide, not only for the presentation of top quality papers, but also as a forum for the discussion of open problems, future research directions and application domains of RL.

AWRL 2017 (in conjunction with ACML 2017) will consist of keynote talks (TBA), contributed paper presentations, and discussion sessions spread over a one-day period.

2017 Annual Korea AI Society Meeting

Organisers: Sun Kim, Hyeran Byun, Heejoon Chae, Seungjin Choi, Geun-Sik Jo, Sung-Bae Jo, Daijin Kim, Hyunjung Shin, Sungroh Yoon

Korea AI society is an organization with five special interest groups: SIG CVPR Pattern Recognition, SIG Machine Learning, SIG Artificial Intelligence, SIG Bio and Health, and SIG Brain Computer Interface. This is an annual joint meeting of five SIGs, introducing

recent research achievements in each SIG. This meeting is also an opportunity for meeting researchers in working in different research areas for potential collaboration.

Conference Organisation

Organising Committee

Honorary Chairs:

Hyeran Byun (Yonsei University, Korea) and Byoung-Tak Zhang (Seoul National University, Korea)

General Co-Chairs:

Kee-Eung Kim (KAIST, Korea) and Masashi Sugiyama (RIKEN/ The University of Tokyo, Japan)

Program Co-Chairs:

Yung-Kyun Noh (Seoul National University, Korea) and Min-Ling Zhang (Southeast University, China)

Journal Track Co-Chairs:

Wee Sun Lee (National University of Singapore, Singapore) and Bob Durrant (University of Waikato, New Zealand)

Workshop Co-Chairs:

Krikamol Muandet (Mahidol University/ Max Plank Institute, Thailand/ Germany) and Jihun Hamm (Ohio State University, USA)

Tutorial Co-Chairs:

Hung Bui (Adobe Research, USA) and Jaesik Choi (UNIST, Korea)

Publicity Co-Chairs:

Hsuan-Tien Lin (National Taiwan University, Taiwan) and Ivor Tsang (University of Technology Sydney, Australia)

Local Arrangements Chair:

Seon Joo Kim (Yonsei University, Korea)

Steering Committee

- Wray Buntine (Monash University, Australia)
- Thomas G. Dietterich (Oregon State University, USA)
- Tu-Bao Ho (Japan Advanced Institute of Science and Technology, Japan)
- Wee Sun Lee (National University of Singapore, Singapore)
- Chih-Jen Lin (National Taiwan University, Taiwan)
- Hiroshi Motoda (AFOSR/AOARD and Osaka University, Japan)
- Masashi Sugiyama (Co-Chair, RIKEN/The University of Tokyo, Japan)
- Zhi-Hua Zhou (Chair, Nanjing University, China)

Senior Program Committee

- Wray Buntine
- Marco Cuturi
- Stephen Gould
- Yuhong Guo
- Steven Hoi
- Sheng-Jun Huang
- Tomoharu Iwata
- Junmo Kim
- Young Jun Ko
- James Kwok
- Hoai An Le Thi
- Yuh-Jye Lee
- Sangkyun Lee
- Yu-Feng Li
- Hsuan-Tien Lin
- Shoude Lin
- Krikamol Muandet
- Alice Oh

- Sinno Pan
- Bernhard Pfahringer
- Dinh Phung
- Tao Qin
- Jesse Read
- Truyen Tran
- Grigorious Tsoumakas
- Liwei Wang
- Jun Xu
- Florian Yger
- Yang Yu
- Junping Zhang
- Kun Zhang
- Jun Zhu

Journal Track Committee

- Wray Buntine
- Steven Hoi
- Kee-Eung Kim
- Hang Li
- Yu-Feng Li
- Masashi Sugiyama
- Ichiro Takeuchi
- Ivor Tsang
- Lijun Zhang
- Jun Zhu

Sponsors

We would like to thank Korea Robotics Society, Seoul National University/Big Data Institute, KAIST/College of Engineering and International Relations Team, Yonsei University, AFOSR/AOARD/ITC-PAC, Amazon AWS, GIGA-BYTE Technology, Korea Tourism Organization, Google, NVIDIA, NAVER, RIKEN Center for Advanced Intelligence Project, Samsung SDS, and Seoul Metropolitan Government.