

AI SUMMER SCHOOL 2020

3 - 7 AUGUST, 2020

SINGAPORE

<https://aisummerschool2020.aisingapore.org/>



AI SINGAPORE

PLATINUM SPONSORS



GOLD SPONSOR



AI INDUSTRY PARTNER



CONTENTS

PROGRAMME

- [Day 1](#)
- [Day 2](#)
- [Day 3](#)
- [Day 4](#)
- [Day 5](#)

SPEAKERS

- [Keynote Speakers](#)
- [Speakers](#)

[ORGANISING COMMITTEE](#)

[SPONSORS](#)

PROGRAMME - DAY 1

9:00 Opening & Welcome

Speakers: Assoc. Prof Stefan Winkler, Assoc. Prof Bryan Low

9:15 Gaussian Processes: Theory & Applications

Speaker: Dr Trọng-Nghĩa Hoàng

In this talk, I will give an overview of Gaussian processes (GPs) which provide a principled non-parametric approach to model prior probability distribution over random functions. In particular, I will first focus on the fundamental construction of a Gaussian process and how it can be used to formulate a Bayesian framework regression/classification. Second, I will discuss practical matters regarding learning and performing inference with Gaussian processes in large data settings. This will cover a spectrum of approximation methods unified under the same umbrella of variational inference. To this end, I will also discuss how Gaussian process and its approximations can be further extended to some of the existing trends in machine learning such as active learning, deep learning and federated learning.

10:45 Low-resource Machine Learning

Speaker: Asst. Prof Joey Tianyi Zhou

Machine learning has experienced a strong growth in recent years, due to increased dataset sizes and computational power, and to advances in algorithms that has made great achievements. However, the state-of-art machine learning methods especially deep models, which are trained over a large number of high-quality data with known labels from standard benchmarks, tend to be generic. Unfortunately, in many tasks, it is often expensive to annotate or even difficult to collect sufficient training data to build machine learning-based information systems. Another challenge is coming from the limitations of hardware. Currently machine learning methods especially deep learning based models are computationally -hungry. During the last decades, we have seen how data is migrating, first from on-premises to cloud data centres and now, from cloud to the "edge" points closer to the source, where it is being generated. Edge AI devices operate with tight resource budgets such as memory, power, and computing horsepower. AI technology with high-end GPUs for training and running large neural networks are not suitable for edge AI. The objective of this proposal to tackle the low-resource challenges such as low-resource data or constrained computing resources to solve many important problems in real-world applications.

12:00 Job Hunting in AI (for PhD students)

Speaker: Dr Yiluan Guo

Yiluan and his colleagues will share their experiences in looking for an AI-related job in industry, particularly as a fresh PhD on CV/ML. He will also invite some of his friends to join this session. They will be working as research scientists in companies or staying in academia as postdoc or faculty. This session will be focusing on sharing and social chat.

13:00 Unconference Session

Unconference sessions are designed to give participants an opportunity to break into small discussion groups, explore different topics based on their interests, and interact more informally.

14:00 Monitoring model performance in production systems

Speaker: Han Qiao

This session will cover some of the technical challenges in productionising machine learning, specifically looking at how model performance be monitored continuously after deployment. You will learn practical techniques in detecting feature and inference drifts, and how they can be applied at scale using Prometheus metrics. A live demo will be shown to guide you through instrumenting your training and model serving code, and crafting PromQL for real time analysis. We will also present the latency impact of this approach using anonymised data from model servers with ~100 inference requests per second.

15:00 Keynote 1:

AI @ Scale – Trends and Lessons Learnt from Large-scale Machine Learning Projects

Speaker: Dr Wee Hyong Tok

In this session, you will learn the key trends in machine learning and deep learning, grounded on practical lessons from working on AI projects with some of the world's largest Fortune 500 companies, on evolving their AI ideas from proof of concept to production systems. You will get a glimpse into how AI is used across the industry, and the technical challenges that need to be addressed to successful deployment of AI solutions. You will learn how AI is used to solve some of the world's toughest challenges in sustainability. Join Wee Hyong for a conversation about how the ingenuity of researchers, data scientists and ML engineers are using AI to create innovative products and changing the world.

16:00 Ethics and Governance of AI

Speaker: Prof Leong Tze Yun

In the context of relevant AI application examples, we will examine, discuss, and develop new strategies to rein in AI to make it fair, accountable, transparent, and beneficial to humankind. We will propose new approaches to educate the next generation of AI Scientists, AI Engineers, AI Managers, and AI Users. We will also design new guidelines to design, implement, use, and monitor “good” AI applications. This module is a first step toward addressing the grand challenge of Responsible AI: We will highlight the fundamental principles underlying the ethical, governance, and regulatory considerations in AI. We will also examine how these principles would guide us in an engineering approach to harness the power of AI to benefit society and economy, especially during and after the COVID-19 pandemic.

17:30 End of Day 1

PROGRAMME - DAY 2

9:00 Workshop 1:
NVIDIA Fundamentals of Deep Learning for Computer Vision
Speakers: Dr Teo Tee Hui / Mr Nicholas Walker

This workshop dives into deep learning techniques for a range of computer vision tasks, including training and deploying neural networks. Students will learn how to implement common deep learning workflows such as Image Classification and Object Detection, experiment with data, training parameters, network structure, and other strategies to increase performance and capability, and finally deploy networks to start solving real-world problems.

9:00 Workshop 2:
Reinforcement Learning: DinerDash challenge
Speaker: Dr Akshay Narayan

This workshop provides an interactive avenue for participants to learn about Reinforcement Learning (RL). Diner Dash is a game where a single waiter makes complex decisions on customer seating arrangements, taking orders, serving food and many others. In small groups, participants will test out RL baselines and compete for the highest score in the Diner Dash simulator.

18:00 End of Day 2

PROGRAMME - DAY 3

9:00 Transformer Models: Recent Advances and New Perspectives

Speaker: Dr Yi Tay

Transformers, characterized by dense self-attention mechanisms, has proven itself as a superior inductive bias across a myriad of applications. Transformers are highly pervasive in modern deep learning research, gaining widespread adoption across the fields of natural language processing, computer vision, and reinforcement learning. In some fields like natural language understanding, Transformers have become completely indispensable and staple. This lecture gives a comprehensive overview of the recent architectural advances of Transformer models. I will broadly discuss the evolution of this class of models across recent years such as the Sparse Transformer, Reformer, Routing Transformer, Funnel Transformers, Compressive Transformers, Transformer-XL, etc. Last but not least, I will present some new perspectives on the Transformers architecture, specifically discussing the recently proposed Synthesizer architecture and its implications for Transformer model research.

10:45 Trustworthy Machine Learning

Speaker: Asst. Prof Reza Shokri

As machine learning algorithms have become an influential component of critical decision making processes, the major question is whether they are trustworthy. Can we trust ML algorithms to perform accurately in noisy and adversarial environments, and be robust against adversarial data? Can we trust ML systems to have access to sensitive data during training and inference? Are there specific privacy risks of using machine learning models? Can we establish trust in black-box models, by providing interpretable predictions and explaining their decisions? Is it ethical to make use of machine learning algorithms? Can we trust them to be fair with respect to different individuals and groups? In this talk, I will discuss these fundamental concerns, and present various technical problems in this domain.

12:00 Unconference Session

Unconference sessions are designed to give participants an opportunity to break into small discussion groups, explore different topics based on their interests, and interact more informally.

13:00 Explainable AI: Lecture Session and Hands-on

Speaker: Asst. Prof Alexander Binder

Introduction on explainable AI briefly and presentation of LRP as one approach to explain classification decisions. It will show uses several cases of explainable AI beyond mere auditing, such as bias identification, and give an example of models to which explainable AI has been applied to.

In the tutorial part, the participants will use runnable LRP code for ResNets and Densenet-121 in pytorch. This will introduce the participants to three aspects of implementation. Firstly, they will be introduced to canonization of deep neural nets, showing them how to deal with implementation invariance arising from convolution-batchnorm and batchnorm-relu-convolution chains and how to cope with bias terms. Secondly, they can see the general implementation mechanisms for LRP methods using custom torch.autograd. Function classes which extend to other explainability methods.

The final goal of the tutorial part is to implement guided backprop (easy) and an LRP-method (more advanced) for a pytorch layer like convolution or a pooling layer which is not included in the above code.

16:00 Deep generative Model: Recent Advances and Applications

Speaker: Asst. Prof Ngai Man Cheung

Recent years have seen remarkable advances in deep generative models such as auto-regressive networks, variational autoencoders (VAEs), generative adversarial networks (GANs). Deep generative models become effective mechanisms in learning high-dimensional, complex data distributions and creating novel samples and high-resolution images. In this talk, I will first review fundamental ideas of deep generative models with a focus on GANs. Then, I will discuss recent advances: stabilizing the training and improving synthetic sample diversity, incorporating style transfer, training GAN with a single image, and applications ranging from image-to-image translation to revealing private training data from a trained mode

18:00 End of Day 3

PROGRAMME - DAY 4

9:00 Keynote 2:

AI for Education: Towards Personalized Learning at Scale

Speaker: Prof Chengxiang Zhai

The recent growth of online education has opened up many exciting opportunities for applying Artificial Intelligence (AI) techniques to transform education to make it both more affordable and more effective. The increasingly available educational data, including both educational content and the human-system interaction data, makes it possible to use machine learning to train intelligent pedagogical agents to (partially) replace human instructors, thus scaling up learning and making education more affordable. Assistive AI techniques can also be developed to augment human intelligence and enable personalized learning, helping instructors improve effectiveness of teaching and learners improving efficiency of learning. In this talk, I will present some of our recent work in these directions, including automating assessment of complex assignments using machine learning and topic models, understanding and analysing behaviours of learners using generative probabilistic models, mining learner interaction data to discover difficult concepts, and scaffolding exploratory learning by linking lecture slides to form a Web of Slides (WOS) and mining explanations of concepts.

10:00 Panel on Applied AI: From Academia to Industry (partnered with SGInnovate)

Speakers: Joumana Ghosn and Frédéric Laurin

With the rise of the digital economy, countries are increasingly powered by more Artificial Intelligence (AI) as the technology permeates every industry and much of human activity. AI can boost efficiencies, performance and productivity in various ways. And in a country like Singapore with a small and ageing workforce, Singapore has to tap on AI and automation to preserve its competitive advantage over other economies.

At SGInnovate, much of our work is to connect Singapore with the global AI R&D ecosystem to learn how we can better harness this technology for next-generation intelligence and decision making. Canadian research institute, Mila, joins us in this session to share how they have been translating their research to real-world applications.

Founded in 1993 by Turing Award Laureate, Professor Yoshua Bengio, Mila's mission is to be a global pole for scientific advances that inspires innovation and the development of AI for the benefit of all. Its industry-oriented team of applied machine learning experts aims to tackle some of the hardest AI projects where the industry might not already have a direct solution. This talk will discuss Mila's collaborative approach of working with its partner organisations and some of its past and current AI use cases.

11:15 SGInnovate Sharing Session

Speaker: Ms Shiangker Wong

11:30 Use of Technology, AI and Data Science in Transforming Healthcare in Singapore

Speaker: Prof Robert Morris

We will describe MOHT's approach to transformation of healthcare in Singapore, with emphasis on how Data Science and AI are addressing chronic diseases and mental health. A perspective will be given on how AI can best help management of COVID-19 and several projects will be described.

12:00 Poster Session

Summer school attendees share their work via pre-recorded videos and discuss with other participants.

13:00 Recommender Systems: A Hands-on Tutorial in Python

Speaker: Assoc. Prof Hady Wirawan Lauw

With pervasive digitization of our everyday lives, we face an increasing number of options, be it in which product to purchase, which movie to watch, which article read, which applicant to interview, etc. Driven by necessity, product and service providers rely on recommender systems to help narrow down the options to those most likely of interest to a target user. This tutorial provides a primer to recommendation algorithms, covering foundational methods based on matrix factorization, with a significant hands-on component conducted in Python using the Cornac recommender system library.

16:00 Strategic Interaction in Multi-Agent Systems

Speaker: Assoc. Prof Arunesh Sinha

A fundamental aspect of AI is multi-agent interaction. Such interaction is present in many scenarios that includes security situations, market interaction, etc. Often multi-agent interaction is studied through the lens of game theory. In this talk, I will walk through basics of game theory and some real-world applications to security problems. I will also point out limitations of classic game theory and how modelling and learning agent behaviour provide a way around stringent assumptions in classic game theory.

17:00 End of Day 4

PROGRAMME - DAY 5

8:30 Artificial Intelligence in Medical Imaging

Speaker: [Dr Jeevesh Kapur](#)

This talk would provide an overview of the AI technology, its potential in the Imaging world, how it fits into the current and future workflow. The talk will cover examples and opportunities of AI in Radiological diagnosis, while discussing Implications and Concerns.

9:00 Entrepreneurship in AI

Speaker: [David Quail](#)

An anti-pattern in many start-ups is to start with technology and then look to solve a problem with it. This "hammer seeking the nail" approach is exacerbated with deep technology companies who own vast amounts of intellectual property they've built over the years and look for every which way to leverage this asset. The danger in this approach is that it can lead to fabricating problems that don't actually exist - or overstating how important of a problem it is - in an attempt to apply their technology. How do such technology companies seek to solve real problems while at the same time utilizing their IP? In this talk I'll look to share lessons learned at various times in my career - from founding an AI based company, to being an entrepreneur in residence at Stanford Research Institute and Samsung Research of America, and now at MEDO.ai.

9:30 AI model for predicting hospitalisation risk of kidney patients

Speaker: [Lim Tern Poh](#)

Come join AI Singapore (AISG) at this talk where we will share about how the development of a medical AI model helps to equip medical staff with knowledge to provide early intervention in kidney dialysis patients at risk of hospitalisation. Every patient generates vast amounts of medical data and we will also look at how this data can be used for modelling a predictor to guide the decisions of medical professionals.

10:45 When AI meets Healthcare

Speaker: [Dr Zeng Zeng](#)

With the rapid development of deep learning, more and more real applications have been implemented in healthcare domain that can help the clinicians to diagnose disease, design treatment, monitor the progress, and even carry out experiments to improve the therapeutics etc. Recently, in order to fight the highly contagious COVID-19, a lot of deep learning-based applications have emerged that help to improve the productivity significantly, such as COVID-19 detection based on CT scans and / or chest X-rays images. As can observe, data play a very important role in AI for healthcare and have countless representations in the format that lead to more complexities of deep learning architectures. In this talk, I will give some brief introduction for "AI meets healthcare" from the view of data dimensions and present various methods that can meet the requirements of different data types. I will start from 1-dimension data and use EEG data as an example, to illustration the main methods in the literature and state-of-the-art performance on the time-series data. Then, I will continue the introduction on 2-D, 3-D, and 4-D data, using X-Rays, CT scans, and multi-phase CT scans as examples, respectively. In order to address the scenarios that multiple different data exist, I will present multi-modality methods that can extract features from multiple datasets and then, merge the features for the improved performance. Through the talk, the audience can have a broad view of AI for healthcare and some basic ideas on the corresponding benchmarks in deep learning domain.

12:00 Poster Session

Summer school attendees share their work via pre-recorded videos and discuss with other participants.

13:00 Federated Learning

Speaker: Asst. Prof Yu Han

The rapid development of artificial intelligence (AI) has benefited from large-scale training data in real-world application. As AI becomes increasingly ubiquitous, nations are increasingly concerned about AI governance and privacy protection. They instituted new legislations such as the General Data Protection Regulation (GDPR) for these purposes. These new laws can potentially limit the development of AI in the long run. Federated Learning (FL), in which training happens where data are stored and only model parameters leave the data silos, can help AI thrive in the privacy- sensitive regulatory environment. FL involves self-interested owners of potentially sensitive local datasets collaboratively training machine learning models. In this way, end-users can become co-creators of AI solutions. In this course, I will introduce the basics of federated learning. In addition, I will highlight some recent advances in FL which aim to enable trust-based dynamic collaboration among data owners to take place.

15:00 Self-supervised Deep Learning

Speaker: Dr Junnan Li

Recent research has seen exciting breakthroughs in self-supervised learning, one of the most popular topics in AI. Different from supervised learning, self-supervised learning does not rely on labels. Instead, it aims to train deep models using unlabelled data which can be freely acquired. This talk will focus on recent developments in self-supervised learning for computer vision. Self-supervised learning algorithms can be broadly categorized into three types: (1) pre-text task, (2) contrastive learning, and (3) cluster assignment. I'll introduce the representative methods under each category. Specifically, I'll discuss Prototypical Contrastive Learning (PCL) which is proposed by Salesforce Research. PCL bridges contrastive learning with clustering, and achieves state-of-the-art performance for self-supervised learning.

16:00 Awards Presentation and Closing Ceremony

Speakers: Assoc. Prof Stefan Winkler and Assoc. Prof Bryan Low

16:30 End of Day 5 and AI Summer School 2020

KEYNOTE SPEAKER

[Wee-Hyong TOK](#)

Principal Data Science Manager, Microsoft Corporation



Wee Hyong is part of the Azure Global AI team at Microsoft. He has extensive leadership experience leading multi-disciplinary team of engineers and data scientists, working on cutting-edge AI capabilities that are infused into products and services. He is a tech visionary with a background in product management, machine learning/deep learning and working on complex engagements with customers. Over the years, he has demonstrated that his early thought-leadership white papers on tech trends have become reality, and deeply integrated into many products. His ability to strategize, and turn strategy to execution, and hunting for customer adoption has enabled many projects that he works on to be successful. He is continuously pushing the boundaries of products for machine learning and deep learning. His team works extensively with deep learning frameworks, ranging from TensorFlow, CNTK, Keras, and PyTorch.

Wee Hyong has worn many hats in his career - developer, program/product manager, data scientist, researcher, and strategist, and his range of experience has given him unique super powers to lead and define the strategy for high-performing Data and AI innovation teams. Throughout his career, he has been a trusted advisor to the C-suite, from Fortune 500 companies to startups.

Prior to his current role, Wee Hyong is part of the SQL Server team, responsible for product management for SQL Server Integration Services (SSIS), and Azure Data Factory, and has shipped multiple versions of SSIS from 2008 to 2013, and is part of the founding team responsible for the ideation of Azure Data Factory, and taking it from incubation to shipping as an Azure Service in 2015.

Wee Hyong is an affiliate professor with the University of Washington, where he teaches data science courses. He co-authored several books on artificial intelligence – including bestselling books on "Practical Automated Machine Learning", "Deep Learning on Azure", and "Predictive Analytics Using Azure Machine Learning".

KEYNOTE SPEAKER

[Cheng-Xiang ZHAI](#)

Donald Biggar Willett Professor, University of Illinois at Urbana-Champaign



ChengXiang Zhai is a Donald Biggar Willett Professor in Engineering of the Department of Computer Science at the University of Illinois at Urbana-Champaign (UIUC), where he is also affiliated with the Carl R. Woese Institute for Genomic Biology, School of Information Sciences, and Department of Statistics. He received a Ph.D. in Computer Science from Nanjing University in 1990, and a Ph.D. in Language and Information Technologies from Carnegie Mellon University in 2002. He worked at Clairvoyance Corp. as a Research Scientist and a Senior Research Scientist from 1997 to 2000.

His research interests are in the general area of intelligent information systems, including specifically intelligent information retrieval, data mining, machine learning and their applications, particularly intelligent education systems. He has published over 300 papers with high citations, and a textbook on text data management and analysis, which is used worldwide by many learners of the two MOOCs that he offered on Coursera. He served as Associate Editors for major journals in multiple areas including information retrieval (ACM TOIS, IPM), data mining (ACM TKDD), and medical informatics (BMC MIDM), and as Program Co-Chairs of ACM SIGIR 2009 and WWW 2015.

He is an ACM Fellow and received a number of awards, including ACM SIGIR Test of Time Paper Award (three times), the 2004 Presidential Early Career Award for Scientists and Engineers (PECASE), Alfred P. Sloan Research Fellowship, IBM Faculty Award, HP Innovation Research Award, and UIUC Campus Award for Excellence in Graduate Student Mentoring. He graduated 35 PhD students and over 50 Master students. More information about him and his work can be found from his homepage at <http://czhai.cs.illinois.edu/>.

SPEAKERS

[Alexander BINDER](#)

Assistant Professor, Singapore University of Technology and Design (SUTD)

Alexander (Alex) Binder obtained a Dr. rer. nat. degree at the department of computer science, Technical University Berlin in 2013. Before he held a Dipl.-Math. degree from the Humboldt University Berlin. Since 2007, he has been working for the THESEUS project on semantic image retrieval at Fraunhofer FIRST where he was the principal contributor to top five ranked submissions at ImageCLEF2011 and Pascal VOC2009 challenges. From 2012 to 2015 he worked on real-time car localization topics in the Automotive Services department (ASCT) of the Fraunhofer Institute FOKUS. From 2010 to 2015 he was with the Machine Learning Group at the TU Berlin. He likes to program in C++, python and enjoys pytorch. His research interests include computer vision, machine learning, explaining non-linear predictions and medical applications.

[Ngai-Man \(Man\) CHEUNG](#)

Associate Professor, Singapore University of Technology and Design (SUTD)

I received my Ph.D. degree in Electrical Engineering from University of Southern California (USC), Los Angeles, CA, in 2008. My Ph.D. research focused on image and video coding, and the work was supported in part by NASA-JPL.

From 2009-2011, I was a postdoctoral researcher with the Image, Video and Multimedia Systems group at Stanford University, Stanford, CA. From 2012-2018, I was an Assistant Professor with Singapore University of Technology and Design (SUTD).

I have also held research positions with Texas Instruments Research Center Japan, Nokia Research Center, IBM T. J. Watson Research Center, HP Labs Japan, Hong Kong University of Science and Technology (HKUST), and Mitsubishi Electric Research Labs (MERL). My research has resulted in 11 U.S. patents granted with several pending, 2 technology licensing and 1 start-up.

[Joumana GHOSN](#)

Director of the Applied Research Team, Mila

With more than 20 years of experience in deep learning, Joumana Ghosn is Director of the Applied Research Team at Mila. In 2002, she obtained a doctorate in computer science (machine learning) from the Université de Montréal, under the supervision of Professor Yoshua Bengio. She then developed an expertise in natural language understanding, particularly at Nuance Communications, where she worked with a team of deep learning researchers dedicated to the creation of virtual assistants and to computer-assisted clinical documentation improvement.

[Yiluan GUO](#)

PhD Student, Singapore University of Technology and Design (SUTD)

Yiluan Guo is currently a PhD student in SUTD. His research interests are mainly on signal processing and image processing. Before joining SUTD, he got his Master's degree and Bachelor Degree from University of Chinese Academy of Sciences (Institute of Microelectronics), 2015 and Jilin University, 2012, respectively.

[Qiao HAN](#)

Senior Software Engineer, BasisAI

[Trọng-Nghĩa HOÀNG](#)

Research Staff Member, MIT-IBM Watson AI Lab

I am a Research Staff Member at MIT-IBM Watson AI Lab, IBM Research, Cambridge, Massachusetts, US. Before IBM, I did my postdoc with Professor Jonathan How at the Laboratory for Information and Decision Systems (LIDS), MIT. I have been working as a Research Fellow at SeSaMe Centre, NUS since March, 2015. Before, I was working as SeSaMe as a Research Assistant (Jan, 2014 - Feb, 2015). I received my Ph.D. degree (Feb, 2015) from NUS in the areas of sequential decision making and machine learning. My advisor is Assistant Professor Low, Bryan Kian Hsiang.

I am also a member of the Multi-Agent Planning, Learning and Coordination Group (MapleCG), School of Computing, NUS. My current research interests include (a) statistical machine learning for big data, (b) adaptive sampling and active learning/sensing, and (c) multi-agent planning and reinforcement learning. My CV is available here.

[Jeevesh KAPUR](#)

Senior Consultant Radiologist, Medo.AI

Dr Jeevesh Kapur is a Senior Consultant radiologist at National University Hospital, Singapore. He was admitted as a Fellow to the Royal College of Radiologists, UK in 2004 and completed his fellowship in 2008 to specialise in Paediatric imaging at the Hospital for Sick Children, Toronto, Canada. He has a special interest in Paediatric Tumour Imaging and Paediatric and general Ultrasound including contrast enhanced ultrasound. He is also actively involved in undergraduate and postgraduate education and sits on several committees in the Radiology fraternity. He is also actively involved in the Asian and Oceanic Society of Pediatric Radiology and part of its executive board. He has devoted much of his time in volunteering his expertise to various children hospitals and workshops in the ASEAN region.

He is co-founder of MEDO.ai, a health-tech startup dual-headquartered in Singapore and Canada. Medo utilises Artificial intelligence-augmented 3D US technology to radically simplify use of US imaging for common and critical conditions. Jeevesh is also an executive member of Collaborative for US deep learning (CUDL), an international collaborative between leading clinicians and AI scientists.

[Frédéric LAURIN](#)

Director of Partnerships, Mila

I am a trained actuary with 10 years' experience in the finance world and start-up experience in Africa and North America. Insatiable lifelong learner, I developed a passion for AI, deep learning, machine learning and neuroscience as well as a strong personal dedication to the cause of mental health.

[Hady Wirawan LAUW](#)

Associate Professor, School of Information Systems, Singapore Management University (SMU)

I am an Associate Professor of Information Systems as well as the Director of BSc (Computer Science) Programme at Singapore Management University (SMU).

My research group Preferred.AI is active in artificial intelligence, data mining, and machine learning, focusing on Web mining, preference analytics and recommender systems.

Concurrently, I am holding National Research Foundation Fellowship, Lee Kong Chian Fellowship, and Educational Research Fellowship.

I also currently serve as the Chair of the Singapore Chapter of ACM SIGKDD.

Previously, I was a Scientist in A*STAR Institute for Infocomm Research and a postdoctoral researcher in Microsoft Research Silicon Valley, after earning PhD and BEng (1st Class) degrees from Nanyang Technological University (NTU).

[Tze-Yun LEONG](#)

Director, AI Technology, AI Singapore

Professor, School of Computing, National University of Singapore (NUS)

LEONG Tze Yun is Director of AI Technology at AI Singapore, a national programme on Artificial Intelligence. She is also Professor of Practice of Computer Science at the School of Computing, National University of Singapore. She is an elected Fellow of the American College of Medical Informatics (ACMI) and a founding Fellow of the International Academy of Health Sciences Informatics (IAHSI).

Tze-Yun received her SB, SM, and PhD degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology (MIT), USA. Her research interests include decision-theoretic artificial intelligence, cognitive modelling and causal reasoning, transfer and reinforcement learning, adaptive computing and biomedical and health informatics. She has over 150 international peer-reviewed publications. She has also served on editorial boards and program committees of leading international journals and conferences on artificial intelligence and biomedical informatics.

With both academic background and business experience, Tze-Yun regularly contributes to panels and committees that advise on research and development directions and education models in computer science, artificial intelligence and health informatics in Singapore and abroad. In

Singapore, she served as working committee member of the Information and Communication Technology and Media Master Plan 2025 in Singapore (2013-2015) and the Ministry of Education Autonomous University Digital Readiness workgroup (2019-2020). She is a Board member of the Health Sciences Authority (HSA) of Singapore, and a reviewer of the AI Ethics Body of Knowledge in Singapore. Internationally, she served as an advisory board member of the United Nations University Institute for Computing and Society (UNU-CS) (2012-2016), and vice president of the International Medical Informatics Association (IMIA) (2013-2016). She is currently a member of the World Health Organization (WHO) Expert Group on Ethics and Governance of Artificial Intelligence in Health, and participates in various WHO workgroups on ethical and regulatory considerations of digital technologies in Health, in general and in the context of the COVID-19 pandemic.

[Junnan Li](#)

Research Scientist, Salesforce Research Asia

Junnan Li received his B.Eng. (Electronic Engineering) from the University of Hong Kong, and the Ph.D. degree (Computer Science) from the National University of Singapore. His research interests are computer vision and deep learning.

[Tern Poh LIM](#)

Principal AI Consultant, AI Singapore

Tern Poh provides consulting services to enable customers to undertake the development and implementation of AI minimum viable models within their organisations. He is also on secondment to Singapore's National AI Office (NAIO) to provide his technical expertise on AI. Before these, he was on secondment to the Singapore's Government Technology Agency (GovTech) as the Data Science & AI Programme Manager to assist in the building of analytics capabilities across the whole-of-Government.

Tern Poh started his career as a trainee trader at Bunge before becoming the APAC Purchasing Manager (IT Innovation) at Procter & Gamble. Coming from a commercial background, he understands the importance of bridging engineers, clients, and stakeholders to translate analytics insights into impact at scale. He is excited about combining his AI knowledge and capability to create exponential value for businesses and society. Tern Poh is pursuing a Master of Computing (AI) at the National University of Singapore. He holds a Bachelor of Business Management (Finance & International Trading, Cum Laude) from Singapore Management University and Advanced Diploma (Biomedical Sciences) from Republic Polytechnic.

[Robert JT MORRIS](#)

Chief Technology Strategist, Ministry of Health, Office for Healthcare Transformation (MOHT)
Professor, National University of Singapore (NUS)

Dr. Robert JT Morris is Professor, National University of Singapore Yong Loo Lin School of Medicine and Chief Technology Strategist for the Ministry of Health Transformation (MOHT). He is also an Advisor to the National Research Foundation in the Prime Minister's Office. His current focus is on application of computer and information sciences, including AI, to transform healthcare.

From 2011-2017, he led all IBM Research's Global Laboratories, and built new laboratories in Brazil, Australia, Africa and Singapore. IBM Research earned 6 Nobel prizes, laying the basis for the information revolution. In 2010 Fortune and Money Magazine declared Robert the "Smartest Scientist in Technology" in their print magazines and on the web <https://ibm.biz/BdrbY4>.

During 2006-2011, Robert was VP Services Research, IBM TJ Watson Research Center, NY, helping start IBM's healthcare and IoT businesses, and earlier VP Assets Innovation, IBM Global Services. From 1999-2004, he was the VP & Director of IBM Almaden Research Center, Silicon Valley, where relational database technology, the disk drive, etc., were invented. Responsible for creation of Services Science, research for IBM's ThinkPad, and Storage Systems. Earlier he led advanced systems research at TJ Watson Center, including the Deep Blue chess machine, which defeated world champion Garry Kasparov in 1997. Originally from Australia, he began his career at Bell Labs working on early neural networks and the largest packet switch of the era (2STP).

Robert was Chairman, Bay Area Science and Innovation Consortium (BASIC) 2001-2005, comprising heads of major research institutions in Silicon Valley. Published over 60 scientific articles, awarded 11 patents. Editor IEEE Trans. on Computers and on many boards. Received the Australian-American Fulbright and Gowrie Fellowships, PICMET Leadership in Tech. Mgmt. Award, Fellow IEEE, IBM Acad. of Technology, Government University Industry Research Roundtable (National Academies).

[Akshay NARAYAN](#)

Lecturer, School of Computing, National University of Singapore (NUS)

I am a Lecturer at SOC-NUS. I am also working towards completing my PhD. My advisor is Prof. Leong, Tze Yun. I completed MTech from IIIT-Bangalore in 2012. Prior to MTech I worked with NetApp India Pvt., Ltd., for three years.

I am broadly interested in AI planning, sequential decision making and its applications. My current research is on transfer learning in reinforcement learning using model-based approaches. I have explored areas of multi-objective reinforcement learning and hierarchical reinforcement learning during the first couple of years of PhD.

Previously I have worked in the area of cloud computing. Some of the projects include smart metering and chargeback as applied to cloud systems, incorporating power awareness in cloud metering (master's thesis project), workload analysis for virtual machine sizing and location in private clouds (research internship with Infosys labs, India) and QoS monitoring and response in cloud systems.

[David QUAIL](#)

VP of Technology, Medo.AI

[Reza SHOKRI](#)

Assistant Professor, School of Computing, National University of Singapore (NUS)

Reza Shokri is a NUS Presidential Young Professor of Computer Science. His research focuses on trustworthy machine learning, quantitative analysis of data privacy, and design of privacy-preserving algorithms for practical applications, ranging from data synthesis to federated learning. He is a recipient of the NUS Early Career Research Award 2019 for working on trustworthy machine learning. He received the Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies in 2018, for his work on analysing the privacy risks of machine learning models, and was a runner-up in 2012, for his work on quantifying location privacy. He obtained his PhD from EPFL.

[Arunesh SINHA](#)

Assistant Professor, School of Information Systems, Singapore Management University (SMU)

Dr. Arunesh Sinha is Assistant Professor at SMU. Previously, he was Assistant Research Scientist in the Computer Science and Engineering Department at the University of Michigan. Prior to joining

University of Michigan, he was a postdoctoral scholar with Prof. Milind Tambe in the Computer Science Department at the University of Southern California. He received his Ph.D. from Carnegie Mellon University, where he was advised by Prof. Anupam Datta. He obtained his undergraduate degree in Electrical Engineering from IIT Kharagpur in India. He worked as a software engineer before starting his Ph.D. He has industry research experience in the form of internships at Microsoft Research, Redmond and Intel Labs, Hillsboro.

Dr. Sinha has conducted research at the intersection of security, machine learning and game theory. His interests lie in the theoretical aspects of multi-agent interaction, machine learning, security and privacy, along with an emphasis on the real-world applicability of the theoretical models. He was awarded the Bertucci fellowship at CMU in appreciation of his novel research. His paper was nominated for the best application paper at AAMAS 2016. He is the chair of two of the leading workshops at the intersection of AI and computer security: AIsec and AICS.

[Yi TAY](#)

Research Scientist, Google AI

I am a Research Scientist at Google AI based in Mountain View, California. I am working on Natural Language Processing and Machine Learning research.

Previously, I obtained my PhD in Computer Science from NTU, Singapore. I did my PhD under the A*Star Graduate Scholarship, supervised by Prof Hui Siu Cheung (@ NTU), Dr Luu Anh Tuan (previously @ A*Star, now at MIT CSAIL) and Dr Su Jian (@ A*Star, I2R) where I was affiliated with the Baidu I2R Research Center working on NLP research. My thesis was entitled Neural Architectures for Natural Language Understanding, which proposed neural models and achieved new state-of-the-art (at that time) across a spectrum of NLU problems such as reading comprehension, question answering and natural language inference. Notably, my thesis also won the outstanding PhD thesis

award (best thesis out of all graduating PhDs in the year 2020) at NTU SCSE. Before pursuing a PhD, I earned my Bachelor in Engineering (Computer Science) from NTU, Singapore in 2015.

I am fortunate to have won best paper award (runner-up) at WSDM 2020, along with several prizes at data science competitions such as CIKM Cup 2016 (1st place) and 2017 (4th place). I also won first place at CVPR VQA Challenge 2017.

[Tee-Hui TEO](#)

Senior Lecturer, Singapore University of Technology and Design (SUTD)

T. Hui Teo graduated with Master of Engineering and Ph.D. from National University of Singapore and Nanyang Technological University in 2000 and 2009 respectively in Electrical & Electronic Engineering. Since 1996, he was with Sharp, ST-Microelectronics, Intelligent Micro-Devices (Matsushita), and etc. as a senior Integrated Circuits (IC) designer, prior joining Institute of Microelectronics, Agency for Science, Technology and Research (A*STAR), Singapore as principle investigator in advanced IC design R&D. In 2010, he joined education sector for setting up both Analog and Digital IC design courses and laboratories for Technical University of Munich, Asia. He is currently with Singapore University of Technology and Design. His research interest are IC design, device characterization & modelling and design science education. T. Hui is a Senior Member of IEEE, and Fellow of IES (Institution of Engineers Singapore).

[Nicholas WALKER](#)

Senior Solution Architect, NVIDIA

Nicolas Walker is a Senior Solution Architect at NVIDIA. He looks after South East Asia customers around data centre and workstation solutions in the areas of High Performance Computing, Deep Learning, Virtualised Desktops and Professional Graphics. Before joining NVIDIA in February 2016, Nicolas held roles in IBM and Lenovo as solution architect focusing on enterprise infrastructure and HPC for the last 15 years. Before moving to Singapore, he was based in Italy, Scotland and Malaysia. He holds a BSc (Hons) in Software Engineering.

[Han YU](#)

Assistant Professor, School of Computer Science and Engineering, Nanyang Technological University of Singapore (NTU)

Dr Han Yu is a Nanyang Assistant Professor (NAP) in the School of Computer Science and Engineering (SCSE), Nanyang Technological University (NTU), Singapore. He has been a Visiting Scholar at the Department of Computer Science and Engineering, Hong Kong University of Science and Technology (HKUST) from 2017 to 2018. Between 2015 and 2018, he held the prestigious Lee Kuan Yew Post-Doctoral Fellowship (LKY PDF) at the Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY). Before joining NTU, he worked as an Embedded Software Engineer at Hewlett-Packard (HP) Pte Ltd, Singapore. He obtained his PhD from the School of Computer Science and Engineering, NTU in 2014. During his PhD study, he held the prestigious Singapore Millennium Foundation (SMF) PhD Scholarship. His research focuses on online convex optimization, ethical AI, federated machine learning and their applications in complex collaborative systems such as

crowdsourcing. He has published over 140 research papers in book chapters, leading international conferences and journals including AAAI, IJCAI, ASE, ACM MM, AAMAS, CIKM, ACM/IEEE Transactions, Proceedings of the IEEE, as well as Nature Research journals - npj Science of Learning, Scientific Data and Scientific Reports. He co-authored the book "Federated Learning" - the first monograph on the topic of federated learning. His research work has been recognized with 16 awards from conferences and journals.

[Zeng ZENG](#)

Senior Scientist, Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A*STAR)

Zeng Zeng received the PhD degree in electrical and computer engineering from the National University of Singapore, Singapore, in 2005, and the BS and MS degrees from the Huazhong University of Science and Technology, Wuhan, China, in 1997 and 2000, respectively. Currently, he works as a Senior Scientist in, I2R, A*STAR, Singapore. From 2011 to 2014, he worked as a senior research fellow with the National University of Singapore. From 2005 to 2011, he worked as an associate professor in Computer and Communication School, Hunan University, China. His research interests include distributed/parallel computing systems, data stream analysis, deep learning, multimedia storage systems, wireless sensor networks, and onboard fault diagnosis.

[Joey Tianyi ZHOU](#)

Scientist and AI Group Manager, Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A*STAR)

Adjunct Assistant Professor, National University of Singapore (NUS)

I am a scientist, PI and group manager with the Institute of High Performance Computing (IHPC), A*STAR, Singapore. I am leading the AI Group with more than 30 research staff members. I also hold an Adjunct Assistant Professor in National University of Singapore. Earlier, I worked as a Senior Research Engineer with Sony US Research Center, San Jose, USA. I obtained my Ph.D. degree from Nanyang Technological University (NTU), SG, in 2015. My research topics focus on Low-Resource Machine Learning and its applications.

ORGANISING COMMITTEE

[Stefan Winkler](#) (Chair)

Deputy Director, AI Singapore

Associate Professor, Department of Computer Science, National University of Singapore

[Bryan Low Kian Hsiang](#) (Co-chair)

Deputy Director, AI Singapore

Associate Professor, Department of Computer Science, National University of Singapore

[David Hsu](#)

Professor, Department of Computer Science, National University of Singapore

[Cuntai Guan](#)

Professor, School of Computer Science and Engineering, Nanyang Technological University

[Pradeep Reddy Varakantham](#)

Associate Professor, School of Information Systems, Singapore Management University

[Ngai-Man \(Man\) Cheung](#)

Associate Professor, Singapore University of Technology and Design

[Cheston Tan](#)

Programme Manager, Artificial Intelligence Initiative, A*STAR - Agency for Science, Technology and Research

[SGInnovate](#)

SPONSORS

[Sea](#)

Sea Limited (NYSE: SE) is a leading global consumer internet company founded in Singapore in 2009. Our mission is to better the lives of consumers and small businesses with technology. We operate three core businesses across digital entertainment, e-commerce, as well as digital payments and financial services, known as Garena, Shopee, and SeaMoney, respectively. Garena is a leading global online games developer and publisher. Shopee is the largest pan-regional e-commerce platform in Southeast Asia and Taiwan. SeaMoney is a leading digital payments and financial services provider in Southeast Asia.

[NSCC](#)

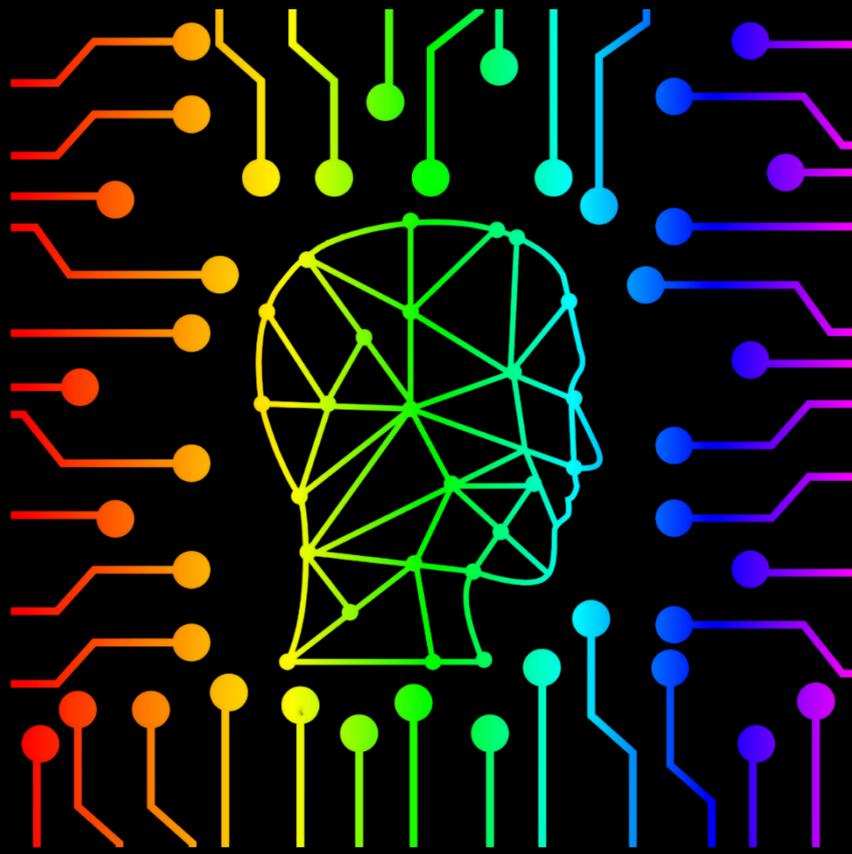
State-of-the-art petascale High Performance Computing (HPC), multi-petabyte data storage, multi-Gigabit speed global connectivity.

We enable users to solve scientific and technological problems. We stimulate industry to use high performance computing for problem solving, testing designs and advancing technologies. Our facility is linked by high bandwidth networks to provide high speed access to users everywhere.

NSCC Singapore was established in 2015 and manages Singapore's first national petascale facility with available high performance computing (HPC) resources. As a National Research Infrastructure funded by the National Research Foundation (NRF), we support the HPC research needs of the public and private sectors, including research institutes, institutes of higher learning (IHLs), government agencies and companies.

[Salesforce](#)

Founded in 1999, Salesforce pioneered the revolutionary idea of replacing traditional desktop CRM software with CRM in the cloud, making it accessible anytime from anywhere. Today, our innovative cloud platform is the world's #1 CRM solution.



AI SUMMER SCHOOL 2020



AI SINGAPORE

Time		AI Singapore AI Summer School 2020 Programme Schedule			
(Singapore)	Monday 3rd Aug	Tuesday 4th Aug	Wednesday 5th Aug	Thursday 6th Aug	Friday 7th Aug
08:30 - 08:45					
08:45 - 09:00					
09:00 - 09:15	Organizing Committee Welcome Message	Hands-on Workshops Workshop 1: Deep Learning for Computer Vision This workshop is designed to be a hands-on experience with Deep Learning (DL) using NVIDIA's graphical DL interface. While all of the projects we'll build focus on computer vision, the concepts are broadly applicable across many deep learning applications. The basic ideas of DL will be introduced followed by comprehensive hands-on to appreciate the DL requirement and steps.			
09:15 - 09:30	Dr Trong-Nghĩa Hoàng Gaussian Processes: Theory & Applications		Dr Yi Tay Transformer Models: Recent Advances and New Perspectives	Keynote by Prof Chengxiang Zhai AI for Education: Towards Personalized Learning at Scale SGInnovate panel with Mila Applied AI: From Academia to Industry	David Quail and Dr. Jeevesh Kapur Entrepreneurship in AI / AI in radiology
09:30 - 09:45					
09:45 - 10:00	Asst Prof Joey Tianyi Zhou Low-resource machine learning		Shiangker Wong SGInnovate Overview Dr Robert Morris AI and Data Science for Transforming Healthcare	Break	Dr Zeng Zeng When AI meets Healthcare
10:00 - 10:15					
10:15 - 10:30	Break				
10:30 - 10:45					
10:45 - 11:00					
11:00 - 11:15					
11:15 - 11:30					
11:30 - 11:45					
11:45 - 12:00					
12:00 - 13:00	Lunch Break / Dr Yiluan Guo Job hunting in AI for PhD students				
13:00 - 13:15		Workshop 2: Reinforcement Learning: DinerDash challenge This workshop provides an interactive avenue for participants to learn about Reinforcement Learning (RL). Diner Dash is a game where a single waiter makes complex decisions on customer seating arrangements, taking orders, serving food and many others. In small groups, participants will test out RL baselines and compete for the highest score in the Diner Dash simulator.			
13:15 - 13:30	Unconference Session				
13:30 - 13:45					
13:45 - 14:00	Han Qiao Monitoring model performance in production systems				
14:00 - 14:15					
14:15 - 1430					
14:30 - 14:45					
14:45 - 15:00	Break				
15:00 - 15:15	Keynote by Dr Wee Hyong Tok AI @ Scale – Trends and Lessons Learnt from Large-scale Machine Learning Projects				
15:15 - 15:30					
15:30 - 15:45					
15:45 - 16:00					
16:00 - 16:15					
16:15 - 16:30					
16:30 - 16:45	Prof Leong Tze Yun Ethics and Governance of AI				
16:45 - 17:00					
17:00 - 17:15					
17:15 - 17:30					
17:30 - 17:45					
17:45 - 18:00					