

Industry Day @ MMAAsia 2025

Theme: Industrial x-Models: Foundations Meet Applications

Date & Time: 10:00 – 18:30, December 10 (Wednesday), 2025

Room: Ballroom

"x-Models" stands for the foundation models including Large Language Model (LLM), Vision Language Model (VLM), Large Multimodal Model (LMM), Small Language Model (SLM), Large Action Model (LAM), etc.

Industry Program

8:30 – 9:00	Opening Ceremony
9:00 – 10:00	Keynote 1: Towards Artificial General Intelligence: Key Challenges and Trends <u>Speaker:</u> Prof. Zhiyuan Liu (<i>Tsinghua University, China</i>)

10:00 – 12:30 *Industry Session 1: Foundations of Industrial x-Models*

Room: Ballroom

Session Chair: Dr. Jianquan Liu (*Director, NEC Corporation, Japan*)

In this session, invite 4 speakers from those industry companies in APAC countries, which developed their own x-Models, to introduce their foundation models.

10:00 – 10:30	<u>Talk 1:</u> MCP-Universe: Connecting Large Language Models to the Real World with Model Context Protocol Servers <u>Speaker:</u> Dr. Ziyang Luo (<i>Research Scientist, Salesforce AI Research, Singapore</i>)
10:30 – 11:00	<u>Talk 2:</u> Language, Speech, and Beyond: The Marco Models for Multilingual Language and Audio Intelligence <u>Speaker:</u> Dr. Chenyang Lyu (<i>Staff Researcher, AI Business Group of Alibaba, China</i>)
11:00 – 11:30	Coffee Break
11:30 – 12:00	<u>Talk 3:</u> cotomi Act: Building Web Agents that Outperform Humans <u>Speaker:</u> Dr. Masafumi Oyamada (<i>Chief Scientist, NEC Corporation, Japan</i>)
12:00 – 12:30	<u>Talk 4:</u> The ILMU Journey: Malaysia's First Multimodal LLM <u>Speaker:</u> Mr. Chee Mun Foong (<i>CEO of YTL AI Labs, Malaysia</i>)
12:30 – 14:00	Lunch

14:00 – 17:30 ***Industry Session 2: Applications of Industrial x-Models***

Room: Ballroom

Session Chair: Dr. Jianquan Liu (Director, NEC Corporation, Japan)

In this session, invite 6 speakers from those industry companies in APAC countries, which applied x-Models for different business domains, to introduce their applications.

14:00 – 14:30	<u>Talk 1:</u> Sarashina: Building Japanese-Native LLMs from Scratch <u>Speaker:</u> Dr. Yuyang Dong (Chief Research Scientist, SB Intuitions Corp., Japan)
14:30 – 15:00	<u>Talk 2:</u> Spark LLM and Its Application: from Multilingual Intelligent Speech to AI4S <u>Speaker:</u> Dr. Xin Li (VP of AI Research Institute, iFLYTEK, China)
15:00 – 15:30	<u>Talk 3:</u> From Hype to Hospitals: Real-World Lessons on Deploying AI Medical Scribes in Thailand's Public Healthcare <u>Speaker:</u> Dr. Chiraphat Boonnag (Lead AI Healthcare Specialist, Looloo Health, Thailand)
15:30 – 16:00	Coffee Break
16:00 – 16:30	<u>Talk 4:</u> x-Models in Manufacturing: Impressive, But Are They Ready? <u>Speaker:</u> Dr. Chih-Fan Hsu (Senior Data Scientist, Inventec Corporation, Taiwan)
16:30 – 17:00	<u>Talk 5:</u> Sovereign & Shared: Frugally Scalable Multilingual–Multimodal AI for Bharat <u>Speaker:</u> Dr. Ganesh Ramakrishnan (Principal Investigator, BharatGen, India)
17:00 – 17:30	<u>Talk 6:</u> Building the Last Line of Defense: AI-Powered Content Moderation in the AI Era <u>Speaker:</u> Ms. Qian Wu (Senior AI Scientist, YiDun AI Lab of NetEase, China)

17:30 – 18:30 ***Industry Session 3 (Panel Discussion):***

Industrial x-Models: Foundations Meet Applications – Opportunities & Challenges

Room: Ballroom

In this panel, the moderator and panelists will discuss the opportunities and challenges with audiences from different perspectives of industrial x-Models in terms of their foundations and applications.

17:30 – 18:30	<u>Moderator:</u> Dr. Jianquan Liu (Director, NEC Corporation, Japan) <u>Panelists:</u> (1) Dr. Masafumi Oyamada (Chief Scientist, NEC Corporation, Japan) (2) Dr. Xin Li (VP of AI Research Institute, iFLYTEK, China) (3) Mr. Chee Mun Foong (CEO of YTL AI Labs, Malaysia) (4) Dr. Ganesh Ramakrishnan (Principal Investigator, BharatGen, India) (5) Dr. Chiraphat Boonnag (Lead AI Healthcare Specialist, Looloo Health, Thailand)
---------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

KEYNOTE I

Towards Artificial General Intelligence: Key Challenges and Trends

Prof. Zhiyuan Liu (*Tsinghua University, China*)

ABSTRACT

In early 2025, DeepSeek V3 and R1 released by China's DeepSeek Team attracted global attention. This demonstrates that Chinese teams can maximize the utilization of limited resources through algorithmic innovations despite constraints from computing power bottlenecks, achieving advanced model capabilities with low cost and high quality. These two models reflect two main themes in the development of artificial intelligence: enhancing model intelligence through algorithmic innovations to expand the capability spectrum of AI; achieving continuous improvement in the capability density of models, rapid reduction in training and application costs, and high-quality development of AI through synergistic innovation across data, models, and computing power—laying the foundation for AI inclusivity. This talk will explore the development trends of large models moving towards artificial general intelligence from these two perspectives.

BIOSKETCH

Liu Zhiyuan is a Tenured Associate Professor in the Department of Computer Science at Tsinghua University. His main research interests cover large language models, knowledge graphs, and social computing. He has published over 200 papers in international journals and conferences on AI with over 67k citations and has released series of MiniCPM and MiniCPM-o/v models specifically designed on devices. He was listed as an Elsevier China Highly Cited Scholar for five consecutive years (2020–2024) and in MIT Technology Review's 35 Innovators Under 35 (China).



Talk #IS1-1

Title: MCP-Universe: Connecting Large Language Models to the Real World with Model Context Protocol Servers

Speaker: Dr. Ziyang Luo (*Research Scientist, Salesforce AI Research, Singapore*)

Abstract: Model Context Protocol (MCP) is rapidly redefining how large language models interface with external systems, enabling real-world tool use across diverse environments. Yet current evaluations for MCP-enabled agents remain superficial, overlooking long-horizon reasoning, unfamiliar tool spaces, and dynamic task requirements. This talk introduces MCP-Universe, the first rigorous benchmark for assessing LLMs through interaction with real MCP servers across six practical domains, including navigation, repository management, financial analysis, 3D design, browser automation, and web search. MCP-Universe evaluates agents via execution-based metrics, covering format adherence, static correctness, and real-time dynamic verification, while exposing two emerging challenges: context explosion from multi-step planning and performance degradation on previously unseen tools. Experimental results reveal large performance gaps among frontier models, with GPT-5, Grok-4, and Claude-4.0-Sonnet achieving only 43.72%, 33.33%, and 29.44% average accuracy, and enterprise agents failing to outperform standard ReAct baselines. The talk concludes by presenting our open-source evaluation framework with UI support, enabling seamless integration of future MCP servers and agent architectures, and laying the foundation for standardized measurement in the evolving MCP ecosystem.

Biosketch: Ziyang Luo is a Research Scientist at Salesforce AI Research, specializing in AI agents and foundation models for computer-use. He has extensive experience with large-scale LLMs, contributing to the open-source project WizardLM, and the first commercial-scale diffusion LLM, Mercury. He has also held research positions at Microsoft Research, Alibaba DAMO Academy, and the National University of Singapore. Ziyang has published over 20 papers at premier AI conferences, including ICLR, CVPR, ICCV, ACL, and EMNLP, and received his PhD in Computer Science from Hong Kong Baptist University.



Talk #IS1-2

Title: Language, Speech, and Beyond: The Marco Models for Multilingual Language and Audio Intelligence

Speaker: Dr. Chenyang Lyu (*Staff Researcher, AI Business Group of Alibaba, China*)

Abstract: As multilingual AI systems rapidly evolve, language technologies are expanding beyond text to encompass a richer spectrum of spoken and auditory intelligence. In this talk, I will introduce the Marco models—our family of multilingual language and audio systems designed to advance robust, scalable intelligence across diverse linguistic settings. The talk begins with an overview of Marco-LLM, a large-scale multilingual text model that serves as a foundational component across our research efforts. I will then present the core progress of our spoken-language research, highlighting how these systems address challenges such as multilingual data scarcity, cross-accent generalization, expressive speech modeling, and long-context audio understanding. Beyond the speech domain, I will discuss how our work integrates into the broader Marco ecosystem. Together, these efforts form a cohesive research trajectory aimed at building

next-generation multilingual language and audio intelligence. This talk outlines the conceptual framework, key technical advances, and future opportunities for expanding the Marco models across modalities and linguistic frontiers.

Biosketch: Dr. Chenyang Lyu is a staff researcher and a tech lead at Alibaba's AI Business Group, where he focuses on speech LMs and multilingual LLMs. He leads/co-leads and participates in the research and development of multiple Marco-series LLMs projects including Marco-LLM, Marco-Voice, Marco-o1 and many others at Alibaba's AI Business Group. He was previously a researcher at Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), focusing on multilingual and multimodal large language models. He obtained his Ph.D. in Natural Language Processing from Dublin City University's ML-Labs, following a Bachelor of Engineering from Northeastern University. He previously held visiting positions at Tencent AI Lab, National Institute of Informatics (NII), Huawei Noah's Ark Lab, and IBM Research-China. He has published over 40 papers in top conferences like NeurIPS, ACL and EMNLP, his GPT4Video work was nominated for Best Paper at ACM-MM 2024. The open-source projects he has led have garnered over 4k GitHub stars, and he has served as area chairs and reviewers for multiple top conferences including ICLR, NeurIPS, ACL, etc. His work has been recognized with awards such as the German DAAD Alnet Fellowship and the 2023 Irish AI Youth of the Year Award, and has been featured by media outlets including Ireland's national broadcaster RTÉ and Sator.



Talk #IS1-3

Title: cotomi Act: Building Web Agents that Outperform Humans

Speaker: Dr. Masafumi Oyamada (*Chief Scientist, NEC Corporation, Japan*)

Abstract: Recent progress in foundation models—such as LLMs and VLMs—has opened the door to agents that can operate computers by invoking external tools. Techniques like Chain-of-Thought have further expanded the problem-solving abilities of these models. In this context, NEC has developed cotomi Act, an autonomous agent that performs end-to-end web operations. Notably, cotomi Act is the first system to surpass human success rates on WebArena, a leading benchmark for web agents. This talk introduces cotomi Act and the core technologies that enable its performance.

Biosketch: Dr. Masafumi Oyamada is Chief Scientist at NEC Corporation, where he leads Large Language Model research and development. Since receiving his Ph.D. from the University of Tsukuba in 2018, he has worked in the interdisciplinary area of machine learning, data engineering, and computational linguistics. His research includes work on probabilistic modeling for entity-relationship data (ICDM17), approaches to annotating tabular data (AAAI19), and tabular data search (ICDE21, SIGIR22, VLDB23). His recent work focuses on language models and structured data processing, including knowledge extraction from language models (EMNLP21), entity matching with question answering models (PAKDD23), organizational large language models (BigData23), bias analysis in retrieval-augmented generation (EMNLP Findings23), and applications of large language models to tabular data (EMNLP24).



Talk #IS1-4

Title: The ILMU Journey: Malaysia's First Multimodal LLM

Speaker: Mr. Chee Mun Foong (*CEO of YTL AI Labs, Malaysia*)

Abstract: ILMU is a sovereign multimodal large language model developed to capture Malaysia's linguistic and cultural diversity, spanning Bahasa Malaysia, Manglish, Malaysian Mandarin, and multimodal signals. ILMU underpins Ryt AI, an agentic LLM framework deployed in Ryt Bank to support natural-language execution of core financial operations within regulated settings. To assess Malay-language competence, ILMU is evaluated on MalayMMLU, a 24k-question benchmark covering 22 curriculum-aligned subjects. The results demonstrate ILMU's capability in Malay reasoning and its suitability as a foundation for high-stakes, domain-specific applications.

Biosketch: Foong Chee Mun has made significant contributions to artificial intelligence (AI) and technology over a 25-year career, beginning as a founding member of Simulex Inc., where he worked on behavior prediction for the US Department of Defense and Fortune 500 companies. He then co-founded MoneyLion Inc., serving as CTO and applying AI to assist middle-class Americans, which was later listed on the NYSE. In MoneyLion, Chee Mun founded a technology and AI center in Kuala Lumpur, leading over 300 professionals and transforming Malaysia into an innovation hub, earning him recognition as one of the top 25 CTOs in fintech in 2019 and 2020. Currently, as the CEO of YTL AI Labs, he aims to advance Malaysia's consumer and technology sectors through AI and fintech, showcasing Malaysia's potential in global technology leadership.



Talk #IS2-1

Title: Sarashina: Building Japanese-Native LLMs from Scratch

Speaker: Dr. Yuyang Dong (*Chief Research Scientist, SB Intuitions Corp., Japan*)

Abstract: Large language models, exemplified by ChatGPT, are transforming how people live and work. While most frontier progress has centered on English, an increasing number of companies and research institutions based in Japan are developing foundation models specifically for Japanese. SB Intuitions, launched by SoftBank Corp., operates one of the largest GPU clusters in Japan (around 10,000 GPUs) and is pursuing a long-term vision: to build powerful, general-purpose Japanese LLMs. In this talk, we present Sarashina, a family of Japanese-native large language models trained fully from scratch, and we share practical lessons from pretraining, post-training, and reinforcement learning.

Biosketch: Dr. Yuyang Dong is a Chief Research Scientist at SB Intuitions, Japan. He leads the development of the Sarashina, a series of Japanese-Native LLMs and build from scratch. He obtained the Ph.D. degree from the University of Tsukuba (2019). He has a wide range of research interests including Database, NLP, LLM and VLM. He won the Best Paper Award in DEXA 2016, and published over 20 papers at database, AI, LLM and vision field, including ICDE, VLDB, AAAI, SIGIR and EMNLP. He also won the DBSJ Kambayashi Young Researcher Award in 2022.



Talk #IS2-2

Title: Spark LLM and Its Application: from Multilingual Intelligent Speech to AI4S

Speaker: Dr. Xin Li (*VP of AI Research Institute, iFLYTEK, China*)

Abstract: Since the release of ChatGPT, Large Language Models have pioneered the forefront of academic research and driven transformations in the industry sector. This report will first introduce the R&D milestones and latest technological advancements of the Spark Large Language Model—independently developed by iFLYTEK based on non-NVIDIA computing infrastructure. It will then elaborate on the high-speed digital human interaction technology and multilingual Large Language Models developed on the basis of this foundational model, as well as their respective industry applications. Furthermore, building on the above context, the report will provide an in-depth overview of iFLYTEK's technological layout and application practices in the field of AI for Science. Finally, in light of current technological trends, the report will offer insights into the future development directions of Large Language Models.

Biosketch: Xin Li, Ph.D. and the senior engineer, is the vice president of the AI Research Institute and the head of the R&D department of iFLYTEK. He obtained his Ph.D. degree from and served as postdoctoral researcher and associate professor at the University of Science and Technology of China (USTC), and was a visiting scholar at the University of Technology Sydney (UTS). He is also a researcher at the National Key Laboratory for Cognitive Intelligence, the life member of International Communication Association (ICA), the senior member of China Computer Federation (CCF), the member of the Executive Committee of the CCF's Big Data Committee, the member of council in the China Association of Standardization (CAS), the deputy director of Brain-Computer Interface and Brain-inspired Intelligence Special Committee of CAS, and the vice chairman of the System and Industry Application Group in Brain-Computer Interface Alliance (BCIA). He is also a standing director of Anhui Artificial Intelligence Society and a founding editor of the Journal of Natural Language Processing and received "The Young 30" honor of Brain Science and Brain-like Intelligence, KSEM, CIKM, SDM best paper/runners up award. He is mainly responsible for the research and application of artificial intelligence technologies, including cognitive neuroscience and scientific intelligence (AI4S). He has led and participated in over 10 research projects, including the strategic leading project of the CAS, 2030 Program and the key research and development programs of the Ministry of Science and Technology of China, along with several funds of Natural Science Foundation of China. He has published more than 60 papers in top international academic conferences and well-known journals and filed over 60 patents.



Talk #IS2-3

Title: From Hype to Hospitals: Real-World Lessons on Deploying AI Medical Scribes in Thailand's Public Healthcare

Speaker: Dr. Chiraphat Boonnag (*Lead AI Healthcare Specialist, Looloo Health, Thailand*)

Abstract: Generative AI holds immense promises for reducing physician burnout, yet bridging the gap between foundation models and clinical reality remains a significant challenge. This is particularly true in the

resource-constrained, high-volume environment of Thailand's public hospitals. This talk explores the practical implementation of an AI-powered medical scribe system designed specifically for the Thai healthcare context. We will move beyond theoretical benchmarks to discuss the messy reality of deployment: handling ambient noise in crowded outpatient departments, managing Thai language nuances, and integrating with legacy Hospital Information Systems. By sharing quantitative data on time-saving efficiency and qualitative insights on clinician adoption, this session offers a blueprint for scaling GenAI solutions that are not just technically robust, but operationally viable in real-world public health settings.

Biosketch: Dr. Chiraphat Boonnag is a physician and Lead AI Healthcare Specialist at Looloo Health, Thailand, where he leads the development of "PresScribe," Thailand's first AI-powered medical scribe designed to alleviate physician burnout. Bridging the gap between clinical medicine and data science, his work focuses on implementing scalable digital health solutions in resource-constrained environments. Previously, he served as Co-PI for the "CHIVID" project, utilizing AI for remote patient monitoring during the COVID-19 pandemic, and developed the "Smoke Alert" IoT health platform. Dr. Boonnag holds a Doctor of Medicine from Chiang Mai University.



Talk #IS2-4

Title: x-Models in Manufacturing: Impressive, But Are They Ready?

Speaker: Dr. Chih-Fan Hsu (*Senior Data Scientist, Inventec Corporation, Taiwan*)

Abstract: x-Models have demonstrated remarkable generalization across vision, language, and multimodal tasks, prompting growing interest in their adoption within modern manufacturing. However, whether these models truly meet industrial-grade requirements remains an open question. This talk examines the key challenges faced in real manufacturing environments, focusing on the need for model adaptation, the importance of trustworthiness, and the difficulty of integrating models into field systems. Using case studies and empirical observations, we illustrate where x-Models excel, where they fall short, and which engineering constraints limit their deployment on production lines. We conclude by highlighting the most critical needs—such as data quality assessment, protection of proprietary factory data, and hybrid human–AI systems—that point to new research opportunities for bridging the gap between today's x-Models and the high standards required in manufacturing.

Biosketch: Chih-Fan Hsu is a Senior Data Scientist at Inventec Corporation. His research interests include Trustworthy AI, Computer Vision, Machine Learning, Virtual Reality, and Multimedia Systems. Dr. Hsu was a Postdoc at the University of California, Davis, National Tsing Hua University, and National Yang Ming Chiao Tung University (2020-2021), and a research assistant at the Academia Sinica (2014-2020). Dr. Hsu received his MS in Computer Science and Information Engineering from the National Taiwan Normal University (2010) and Ph.D. in Electrical Engineering from the National Taiwan University (2019).



Talk #IS2-5

Title: Sovereign & Shared: Frugally Scalable Multilingual–Multimodal AI for Bharat

Speaker: Dr. Ganesh Ramakrishnan (*Principal Investigator, BharatGen, India*)

Abstract: The movement for Sovereign AI is accelerating. Meeting its promise requires vertically integrated AI stacks—spanning data, models, and reasoning systems—that remain sovereign while adhering to shared scientific principles around which global research communities can coalesce. This talk presents BharatGen as a sovereign-yet-shared effort to make AI work for the many: creation of datasets, benchmarks, and models that natively support Indian languages, dialects, and code-mix across text, speech, and vision; data pipelines grounded in local realities; and frugal methods that reduce cost and lower barriers. We outline our journey to date across language infrastructure, efficient training and distillation, and early sector pilots. The R&D deep dive will draw from some of our recent work on cross-lingual knowledge distillation for low-resource languages, tokenization/phonetic design for code-mix robustness, or trustworthy document AI with visual grounding focusing on robustness under dialect/code-mix shift, and latency/cost trade-offs. We hope to inspire other Sovereign-AI efforts, especially in the low-resource ecosystems of the Global South and close by inviting international collaborations toward principled research to build people-serving AI.

Biosketch: Dr. Ganesh Ramakrishnan (<https://www.cse.iitb.ac.in/~ganesh/>) is currently serving as Bank of Baroda Chair Professor in Digital Entrepreneurship at the Department of Computer Science and Engineering, IIT Bombay. He completed his BTech and PhD both from the Department of CSE, IIT Bombay. His current areas of research include Large Language Models and Generative AI, human-assisted AI/ML, use of AI/ML in resource-constrained environments, learning with symbolic encoding of domain knowledge in ML and NLP, etc. His research has been featured in top conferences such as AAAI, ACL, NeurIPS, ICML and EMNLP, and he has served as the area chair for major conferences such as AAAI and ACL. Prof. Ganesh Ramakrishnan has been leading the Large Language Modeling initiatives for India at BharatGen funded by the Ministry of Electronics and Information Technology (MeitY) of India through the India AI Mission as well as by the Department of Science and Technology through the NM-ICPS program. He also engages extensively in Industrial collaborations such as IBM Research, Adobe and Google. In recognition of his leadership in the field of AI, he was recently honored with the recognition as one of the Top 30 Indian minds leading the AI revolution by Accel and forbesindia.com. For a long time, he has been focusing his energy on organizing relevant machine learning modules for resource-constrained environments into <https://decile.org/>. He has demonstrated the impact of such data-efficient machine learning in applications such as Video Analytics (<https://www.cse.iitb.ac.in/~vidsurv/>), an end-to-end machine translation eco-system (<https://www.udaanproject.org/>) and OCR (<https://www.cse.iitb.ac.in/~ocr/>) that are all used extensively, as well as in works in the making such as multi-modal analytics (<https://www.cse.iitb.ac.in/~malta/>). He has received the prestigious National Gold Award for eGovernance (Gold Award) in 2022, the Dr P.K. Patwardhan Award for Technology Development 2020 as well as an IIT Bombay Impactful Research Awards in both 2024 and in 2017. He has also received awards such as the IBM Faculty Award, Amazon Research Award, and awards from Google Research, Qualcomm, Adobe, Microsoft, etc. He also held the Institute Chair



Professorship between 2021 and until 2024, J.R. Isaac Chair at IIT Bombay between 2014 and 2016. Ganesh is very passionate about boosting the AI research eco-system for India and toward that, the research by him and his students as well as collaborators has resulted in startups that he has either jointly founded, has transferred technology to, or is mentoring. Ganesh has also served as the founding head of the Koita Centre for Digital Health at IIT Bombay (<https://www.kcdh.iitb.ac.in/>), the first of its kind in India. Owing to his leadership at the National Disease Modeling Consortium, he has been selected to be a disease and economic modeling expert member of Standing Technical Sub Committee (STSC) as well as an expert member of Standing Working Group – Immunization and Vaccine Research and Capacity Building (SWG-IVRCB) as part of NTAGI (National Technical Advisory Group on Immunization), Ministry of Health and Family Welfare, Govt of India.

Talk #IS2-6

Title: Building the Last Line of Defense: AI-Powered Content Moderation in the AI Era

Speaker: Ms. Qian Wu (*Senior AI Scientist, YiDun AI Lab of NetEase, China*)

Abstract: In the AI era, content moderation faces a growing “scissors gap”: harmful content is easier than ever to generate, while traditional discriminative models struggle to keep pace with rapidly evolving and increasingly complex threats. This talk explores how to apply GenAI in an efficient and controllable way within a content moderation system. We will demonstrate the critical roles GenAI plays at different stages of our visual content moderation system, including dataset construction, automated moderation workflows, and human–AI collaborative review. Through real-world production cases, we present a practical roadmap for building a trustworthy content moderation system that mitigates GenAI-driven risks and helps platforms balance compliance, user experience, and operational efficiency.

Biosketch: Qian Wu is a Senior AI Scientist at NetEase Yidun AI Lab. She draws on her research expertise to work closely with customers and engineering teams, contributing to the development of high-impact content moderation systems and solutions. Ms. Wu received her M.S. degree from Peking University. Her research spans computer vision, machine learning, and vision language models, with a particular focus on video understanding technologies for content moderation that are both effective and practical under real-world constraints. She has published several related peer-reviewed papers. Most recently, her work has centered on open-vocabulary, vision-based recognition algorithms for multimodal harmful content moderation.

