



The 11th ASIAN-AUSTRALASIAN CONFERENCE ON PRECISION AGRICULTURE (ACPA 11)
October 14-16, 2025 | Chiayi, TAIWAN

SPEAKER BIODATA FORM

Title: Dr

First Name: Nancy

Last Name: Chen

Position: Lead Principal Investigator, Programme Head

Organization: A*STAR (Agency for Science, Technology, and Research)

Country: Singapore



Brief career highlights (less than 250 words):

Dr. Nancy F. Chen is an ISCA Fellow (2025), AAIA Fellow (2025), and A*STAR Fellow (2023), and was recognized with the Asian Women Tech Leaders Award (2025). She is also a 2025 inductee of IEEE Eta Kappa Nu (HKN), the honor society of IEEE recognizing outstanding engineers.

At A*STAR, Dr. Chen leads the Multimodal Generative AI group and the AI for Education Programme. Dr. Chen is a serial best paper award winner across major conferences - including ICASSP, ACL, EMNLP, MICAAI, COLING, APSIPA, SIGDIAL and EACL – her research spans applications in education, healthcare, neuroscience, social media, security and forensics. Dr. Chen's multimodal, multilingual technologies have led to commercial spin-offs and adoption by Singapore's Ministry of Education.

Dr. Chen is an active international research leader, having served as Program Chairs for top-tier AI conferences such as NeurIPS (#1 on google scholar), ICLR (#2) and delivering international keynote addresses across the globe. She is also an IEEE SPS Distinguished Lecturer (2023-2024), inspiring the broader research community.

Dr. Chen's strengths in strategic vision are reflected in her service to the global STEM (Science, Technology, Engineering and Math) community, including as an APSIPA Board-of-Governor (2024–2026), ISCA Board Member (2021-2024), and honoree of Singapore's 100 Women in Tech (2021). She has won professional awards from IEEE, Microsoft, P&G, UNESCO, L'Oréal, National Institute of Health (USA).

Dr. Chen has long advised government and commercial entities on AI and technology, beginning during her time at MIT Lincoln Laboratory while pursuing her PhD at MIT and Harvard.

The 11th Asian-Australasian Conference on Precision Agriculture (ACPA 11)
October 14-16, 2025, Chiayi, Taiwan

Embodied Agentic Artificial Intelligence for Precision Agriculture: Cross-Domain Experience from Multimodal Generative AI

Nancy F. Chen

Agency for Science, Technology, and Research (A*STAR), Singapore

nfychen@i2r.a-star.edu.sg

INTRODUCTION

My team develops inclusive, responsible, and multimodal AI technology across education, healthcare, and digital services grounded in our research in embodied agentic intelligence and large language models. I will share deployed examples from these domains and draw parallels to agriculture, where similar technical challenges persist, ranging from multimodal fusion for contextual reasoning, explainable AI for actionable insights, and data-efficient learning for adaptation and localization. While the application contexts differ, the underlying AI considerations are remarkably aligned, opening pathways for cross-domain innovation.

Keywords: Multisensory models, foundation models, resource-efficient learning, AI for society

METHODS AND MODELS

Why does agriculture need a new kind of intelligence? Because the challenges it faces – climate volatility, an aging workforce, fragile supply chains, and accelerating land degradation are complex, interconnected, are beyond the reach of any single tool or sector. Precision agriculture sits at the intersection of environmental sustainability, economic resilience, and social equity. It is no longer about just applying sensors or analytics — it’s about how we learn from farms dynamically and in context. Below we highlight technical approaches and real-world deployments from other domains that tackle challenges about multimodal understanding, data sparsity, and interpretability – offering promising directions for a new generation of precision agriculture solutions.

I. Multimodal Learning, Reasoning, and Grounding

As part of the Singapore National Multimodal Large Language Model Programme, we have developed MERaLiON (Multimodal Empathetic Reasoning and Learning in One Network)¹ — the first multimodal large language model tailored to Southeast Asia’s diverse linguistic and cultural landscape (Wang et al 2024, Wang et al 2025, He et al 2025). MERaLiON reasons over multiple streams of input, including ambient audio (e.g., environmental sounds like waterfalls, thunderstorms, or machinery), emotional cues (e.g., distress signals such as screams resulting from accidents or equipment failure), and textual content (e.g., Indonesian social media posts reporting crop loss or pest outbreaks). Multimodal representation learning

¹<https://www.meralion.ai/>; Open Source Model and Data <https://huggingface.co/MERaLiON>

is critical in contextualizing information from heterogeneous data types with varying quantity and quality in addition to temporal and spatial resolution. Such lessons derived from digital service applications (e.g. customer engagement) can also be used in precision agriculture.

II. Data-Efficient Learning

Data sparsity remains a fundamental challenge in generalizing AI research, particularly when most mainstream AI models are developed within a North American context. To illustrate how representation scaling can address this, we present a language learning use case: **SingaKids AI Tutor** (Liu et al 2025), a multimodal, multilingual agent designed to engage young children in Singapore in learning ethnic languages of Malay, Mandarin and Tamil. The goal is to strengthen Singapore’s multilingual advantage while honoring its multicultural heritage. Although all four national languages in Singapore – English, Mandarin, Malay and Tamil – are officially treated equally, the availability of machine-learning-suitable linguistic resources decline sharply in that order. This disparity contributes to the performance gaps in multilingual AI systems, disproportionately disadvantaging underrepresented groups, particularly Malay and Tamil learners. To address these challenges, we developed a suite of computational strategies including self-supervised representation learning (Liu et al 2021a, Liu et al 2021b), instruction fine tuning (Wang et al 2024, Lin et al 2025), and reward process synthesis (Jiao et al 2024). These methods enable us to scale data efficacy for advanced deep learning models while strategically injecting noise to enhance robustness under low resource conditions. We can apply similar approaches in adapting agriculture AI models used in Asia.

III. Intelligence-Informed Analysis for Decision-Making

We have computationally modeled clinical conversations and vital signs to interpret lifestyle needs of diabetes and heart-failure patients, revealing links with downstream health outcomes (Liu et al 2023, Erdt et al 2024). These insights inform *Siu Dai*², a clinical-protocol-guided telehealth bot that engages patients between hospital visits and supports nurses in refining telehealth phone calls for enhanced chronic disease management and monitoring. This example demonstrates how structured understanding of unstructured, multimodal data can yield explainable, actionable insights for care delivery. This methodology is directly applicable to precision agriculture, where informal yet information-rich data—such as farmer extension calls, community radio broadcasts, or social media reports—often contain early indicators of equipment malfunction, pest outbreaks, or environmental stress. Such approaches enable transparent decision support systems that empower end users (e.g. smallholder farmers) and inform policy makers with interpretable summaries grounded in lived experience. These insights are not only model outputs but are traceable to conversational cues and patterns—facilitating trust, usability, and uptake in real-world, low-resource agricultural setting.

CONCLUSIONS

Precision agriculture cultivates more than just crops or livestock — it nurtures the intersection of planetary health, economic resilience, and human dignity. Achieving this vision calls for AI that learns from minimal and noisy data, listens across sensors and languages, and speaks clearly to those it serves. Our research experience in MERaLiON, SingaKids AI Tutor, and Siu Dai Telehealth Bot draws parallel technical solutions for cross-domain innovation and synergies.

² *Siu Dai* means “less sugar” in Singlish

REFERENCES

- Erdt M, Yusof SB, Chai L, Md Salleh SU, Liu Z, Sarim HB, Lim GC, Lim H, Suhaimi NFA, Yulong L, Guo Y, Ng A, Ong S, Choo BP, Lee S, Weiliang H, Oh HC, Wolters MK, Chen NF, Krishnaswamy P. Characterization of Telecare Conversations on Lifestyle Management and Their Relation to Health Care Utilization for Patients with Heart Failure: Mixed Methods Study, *Journal of Medical Internet Research* 2024
- Yingxu He, Zhuohan Liu, Geyu Lin, Shuo Sun, Bin Wang, Wenyu Zhang, Xunlong Zou, Nancy F. Chen, and AiTi Aw. 2025. MERaLiON-AudioLLM: Advancing Speech and Language Understanding for Singapore. *ACL* 2025.
- Fangkai Jiao, Chengwei Qin, Zhengyuan Liu, Nancy F. Chen, and Shafiq Joty. Learning Planning-Based Reasoning by Trajectories Collection and Process Reward Synthesizing, *EMNLP* 2024.
- G. Lin, B. Wang, Z. Liu, Nancy F. Chen. CrossIn: An Efficient Instruction Tuning Approach for Cross-Lingual Knowledge Alignment, *SUMEval Workshop*, *COLING* 2025. [Best Paper Award]
- Zhengyuan Liu, Ke Shi, and Nancy Chen. Coreference-Aware Dialogue Summarization. *SIGDIAL* 2021.
- Zhengyuan Liu and Nancy Chen. Controllable Neural Dialogue Summarization with Personal Named Entity Planning. *EMNLP* 2021.
- Zhengyuan Liu, Geyu Lin, Hui Li Tan, Huayun Zhang, Yanfeng Lu, Xiaoxue Gao, Stella Xin Yin, Sun He, Hock Huan Goh, Lung Hsiang Wong, and Nancy F. Chen. 2025. SingaKids: A Multilingual Multimodal Dialogic Tutor for Language Learning. *ACL* 2025
- Zhengyuan Liu, Siti Umairah Md Salleh, Hong Choon Oh, Pavitra Krishnaswamy, and Nancy Chen. Joint Dialogue Topic Segmentation and Categorization: A Case Study on Clinical Spoken Conversations, *EMNLP* 2023.
- Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, AiTi Aw, and Nancy Chen. 2024. SeaEval for Multilingual Foundation Models: From Cross-Lingual Alignment to Cultural Reasoning. *NAACL* 2024.
- Bin Wang, Geyu Lin, Zhengyuan Liu, Chengwei Wei, and Nancy Chen. CRAFT: Extracting and Tuning Cultural Instructions from the Wild. *Cross-Cultural Considerations in NLP workshop*, *ACL* 2024.
- Bin Wang, Chengwei Wei, Zhengyuan Liu, Geyu Lin, and Nancy F. Chen. 2024. Resilience of Large Language Models for Noisy Instructions. *EMNLP Findings* 2024.
- Bin Wang, Xunlong Zou, Geyu Lin, Shuo Sun, Zhuohan Liu, Wenyu Zhang, Zhengyuan Liu, Ai Ti Aw, and Nancy F. Chen. AudioBench: A Universal Benchmark for Audio Large Language Models. *NAACL* 2025.