

EDUCATION

University of California, Santa Barbara

Ph.D. in Computer Science

Santa Barbara, USA

Expected 2030

The University of Liverpool

B.Sc. in Information and Computing Science

Liverpool, UK

Expected 2025

Major GPA: 4.0/4.0 Overall GPA: 3.86/4.0

PUBLICATIONS

* Equal contribution.

Agentic Systems & Multimodal Reasoning

StreamMind: Proactive Video Agents for Evolving User Intent

Chengzhi Liu, Yuxiang Lai, Yuzhe Yang, Yanqing Liu, Jiheng Hou, Yue Cao, Haonan Wang, Ruixi Fan, Lan Wang, Sheng Liu, Songyou Peng, Yale Song, Kevin Qinghong Lin, Xin Eric Wang, Jui-Hsin Lai, Youbao Tang

Preprint 2026

Proactive video agents that adapt to evolving user intent.

WorldMemArena: Evaluating Multimodal Agent Memory Through Action–World Interaction

Chengzhi Liu, Yuzhe Yang, Xiao Pu, Yepeng Liu, James Zou, Sheng Liu, Yuheng Bu, Songyou Peng, Xin Eric Wang

NeurIPS 2026 [PDF]

Evaluating multimodal agent memory through action–world interaction.

Survive or Collapse: The Asymmetric Roles of Data Gating and Reward Grounding in Self-Play RL

Sophia Xiao Pu, Zhaotian Weng, Chengzhi Liu, Jayanth Srinivasa, Gaowen Liu, William Yang Wang, Xin Eric Wang

NeurIPS 2026 [Paper]

Studying the asymmetric roles of data gating and reward grounding in self-play RL.

Agents’ Last Exam

Yiyu Sun et al. (including Chengzhi Liu)

NeurIPS 2026 [Paper]

Benchmarking agents on long-horizon, economically valuable tasks.

Auditing Agent Harness Safety

Chengzhi Liu*, Yichen Guo*, Yepeng Liu, Yuzhe Yang, Qianqi Yan, Xuandong Zhao, Wenyue Hua, Sheng Liu, Sharon Li, Yuheng Bu, Xin Eric Wang

Preprint 2026 [Project]

Auditing the safety risks of agent harnesses.

Reasoning Within the Mind: Dynamic Multimodal Interleaving in Latent Space

Chengzhi Liu*, Yuzhe Yang*, Yue Fan, Qingyue Wei, Sheng Liu, Xin Eric Wang

CVPR 2026 [Project]

Dynamic multimodal reasoning through latent-space interleaving.

Presenting a Paper is an Art: Self-Improvement Aesthetic Agents

Chengzhi Liu*, Yuzhe Yang*, Kaiwen Zhou, Zhen Zhang, Yue Fan, Yannan Xie, Peng Qi, Xin Eric Wang

ICLR 2026 · ICLR 2026 RSI Workshop Spotlight [Project]

Self-improving agents for aesthetically driven presentation generation.

Multimodal Safety & Alignment

More Thinking, Less Seeing? Assessing Amplified Hallucination

Chengzhi Liu*, Zhongxing Xu*, Qinyue Wei, Juncheng Wu, James Zou, Xin Eric Wang, Yuyin Zhou, Sheng Liu
NeurIPS 2025 [Project]

Assessing how extended reasoning can amplify hallucinations.

Seeing Far and Clearly: Mitigating Hallucinations in MLLMs

Chengzhi Liu*, Feilong Tang*, Zhongxing Xu*, Hu Ming, Zile Huang, Jionglong Su, Imran Razzak, Zongyuan Ge
CVPR 2025 (Oral) [PDF]

Mitigating hallucinations in multimodal large language models.

Multimodal Situational Safety

Chengzhi Liu*, Kaiwen Zhou*, Xuandong Zhao, Anderson Compalas, Dawn Song, Xin Eric Wang
ICLR 2025 [Paper]

Evaluating safety under multimodal situational context.

The Hidden Risks of Large Reasoning Models: A Safety Assessment of R1

Kaiwen Zhou*, Chengzhi Liu*, Xuandong Zhao, Jayanth Srinivasa, Gaowen Liu, Dawn Song, Xin Eric Wang
AAACL 2025

Assessing safety risks in large reasoning models.

Intervening Anchor Token: Decoding Strategy in Alleviating Hallucinations

Feilong Tang*, Zile Huang*, Chengzhi Liu, Qiang Sun, Dawn Song, Harry Yang, Ser-Nam Lim
ICLR 2025 [Paper]

Intervening on anchor tokens during decoding to reduce hallucinations.

Multimodal Learning & Healthcare

Incomplete Modality Disentangled Representation for Ophthalmic Disease

Chengzhi Liu*, Zile Huang*, Feilong Tang, Yu Tian, Zhongxing Xu, Yalin Zheng, Yanda Meng
AAAI 2025 (Oral) [Paper]

Disentangled representations with incomplete modalities for ophthalmic disease.

RESEARCH EXPERIENCE

Google Cloud AI

Student Researcher

Jun. 2026–Sep. 2026

ACADEMIC SERVICE

Conference Reviewer: AAAI 2026; ICLR 2026; AAAI-AIA 2026; CVPR 2026; ICML 2026; NeurIPS 2026.

Invited Talks: Canvas, aesthetic-driven automatic presentation generation (Oct. 2025); NLP Academic Exchange Platform, Multimodal Latent-Space Reasoning (Nov. 2025).