Ce Zhang

https://zhangce01.github.io/ zhangce1203@gmail.com (412) 844-1074

Education

- **Carnegie Mellon University, Pittsburgh, United States** *Aug 2025 - Aug 2028 (Expected)*
  Ph.D. in Robotics, School of Computer Science Research Advisor: Prof. Katia Sycara QPA: 4.16/4.00
  Courses: Mathematical Foundations for Robotics, Computer Vision, Independent Research.
- **Carnegie Mellon University, Pittsburgh, United States** *Aug 2023 - Dec 2024*
  M.Sc. in Machine Learning, School of Computer Science Research Advisor: Prof. Katia Sycara QPA: 3.96/4.00
  Courses: Introduction to Machine Learning, Convex Optimization, Probability Graphical Models, Probability and Mathematical Statistics, Machine Learning in Practice, Intermediate Deep Learning, Independent Research.
- **Southern University of Science and Technology (SUSTech), Shenzhen, China** *Aug 2019 - Jun 2023*
  B.Eng. in Communication Engineering (Summa Cum Laude), College of Engineering GPA: 3.91/4.00 Rank: 1/30
  Selected Courses: Data Structures and Algorithm Analysis (100), Linear Algebra (100), Introduction to Computer Programming (98), Artificial Intelligence (96), Probability and Statistics (96).

Work Experience

Research Scientest Intern, TikTok *May 2026 - Aug 2026*
Mentor: Ming Zhou *San Jose, United States*

- Advanced streaming video understanding with hierarchical memory and agentic reasoning for real-time video QA.
- Developed a foundation embedding model for product catalog retrieval, encoding rich shop metadata and descriptions for scalable similarity search.

Research Intern, Tencent AI Lab *Feb 2025 - Aug 2025*
Research Mentor: Dr. Kaixin Ma *Shenzhen, China*

- Introduced VScan, a two-stage training-free visual token reduction framework for LVLMs, achieving 2.91× inference speedup and 10× FLOPs reduction while retaining 95.4% of original performance across 16 benchmarks.
- Co-developed WebAggregator, a data synthesis framework shifting Deep Research agents from retrieval-centric to compositional reasoning; the resulting WebAggregator-32B surpasses GPT-4.1 and matches Claude-3.7-Sonnet on GAIA, WebWalkerQA, and XBench.

Research Assistant, Carnegie Mellon University *Aug 2023 - Dec 2024*
Research Advisor: Prof. Katia Sycara *Pittsburgh, United States*

- Developed a self-correcting mechanism for large vision-language models to enhance their real-world reliability.
- Enhanced vision-language models to generalize to out-of-distribution domains through few-shot/test-time adaptation.
- Designed a hierarchical inference approach for accurate scene graph generation under real-world corruptions.

Selected Publications (*Equal Contribution)

- **Ce Zhang**, Jinxi He, Yaqi Xie, Katia P. Sycara. LENS: Adaptive Spatio-Temporal Zooming for Keyframe Sampling in Long-Form Videos. Under review *European Conference on Computer Vision* (*ECCV*), 2026.
- **Ce Zhang***, Jinxi He*, Junyi He, Katia P. Sycara, Yaqi Xie. Evolving Contextual Safety in Multi-Modal Large Language Models via Inference-Time Self-Reflective Memory. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2026. Also presented at *ICLR 2026 Workshop on Lifelong Agents*. [PDF] [Project] [Code]
- Rui Wang, **Ce Zhang**, Jun-Yu Ma, Jianshu Zhang, Hongru Wang, Yi Chen, Boyang XUE, Tianqing Fang, Zhisong Zhang, Hongming Zhang, Haitao Mi, Dong Yu, Kam-Fai Wong. WebAggregator: Enhancing Compositional Reasoning Capabilities of Deep Research Agent Foundation Models. In *Annual Meeting of the Association for Computational Linguistics* (*ACL Main*), 2026. [PDF] [Code]
- Zhanpeng Luo*, **Ce Zhang***, Silong Yong, Cunxi Dai, Qianwei Wang, Haoxi Ran, Guanya Shi, Katia P. Sycara, Yaqi Xie. pySpatial: Generating 3D Visual Programs for Zero-Shot Spatial Reasoning. In *International Conference on Learning Representations* (*ICLR*), 2026. Also presented at *ICLR 2026 Workshop on Efficient Spatial Reasoning*. [PDF] [Project] [Code]
- **Ce Zhang**, Kaixin Ma, Tianqing Fang, Wenhao Yu, Hongming Zhang, Zhisong Zhang, Haitao Mi, Dong Yu. VScan: Rethinking Visual Token Reduction for Efficient Large Vision-Language Models. *Transactions on Machine Learning Research* (*TMLR*), 2025. [PDF] [Project] [Code]
- **Ce Zhang**, Zifu Wan, Simon Stepputtis, Katia Sycara, Yaqi Xie. Spectral-Aware Global Fusion for RGB-Thermal Semantic Segmentation. In *International Conference on Image Processing* (*ICIP*), 2025. [PDF]
- Zifu Wan*, **Ce Zhang***, Silong Yong, Martin Q. Ma, Simon Stepputtis, Louis-Philippe Morency, Deva Ramanan, Katia Sycara, Yaqi Xie. ONLY: One-Layer Intervention Sufficiently Mitigates Hallucinations in Large Vision-Language Models. In *International Conference on Computer Vision* (*ICCV*), 2025. [PDF] [Code]
- Zifu Wan, Yaqi Xie, **Ce Zhang**, Zhiqiu Lin, Zihan Wang, Simon Stepputtis, Deva Ramanan, Katia Sycara. InstructPart:

Task-Oriented Part Segmentation with Instruction Reasoning. In *Annual Meeting of the Association for Computational Linguistics* (*ACL Main*), 2025. Also presented at *AAAI 2024 Workshop on Public Sector LLMs*. [PDF] [Project] [Data]

- **Ce Zhang***, Zifu Wan*, Zhehan Kan, Martin Q. Ma, Simon Stepputtis, Deva Ramanan, Russ Salakhutdinov, Louis-Philippe Morency, Katia Sycara, Yaqi Xie. Self-Correcting Decoding with Generative Feedback for Mitigating Hallucinations in Large Vision-Language Models. In *International Conference on Learning Representations* (*ICLR*), 2025. Also presented at *NeurIPS 2024 Workshop on Responsibly Building the Next Generation of Multimodal Foundational Models*. [PDF] [Project] [Code]
- **Ce Zhang**, Simon Stepputtis, Katia Sycara, Yaqi Xie. Enhancing Vision-Language Few-Shot Adaptation with Negative Learning. In *IEEE/CVF Winter Conference on Applications of Computer Vision* (*WACV*), 2025. Also presented at *ICLR 2024 Workshop on Mathematical and Empirical Understanding of Foundation Models*. [PDF] [Code]
- **Ce Zhang**, Simon Stepputtis, Katia Sycara, Yaqi Xie. Dual Prototype Evolving for Test-Time Generalization of Vision-Language Models. In *Conference on Neural Information Processing Systems* (*NeurIPS*), 2024. Also presented at *ICML 2024 Workshop on Foundation Models in the Wild*. [PDF] [Project] [Code]
- **Ce Zhang**, Simon Stepputtis, Joseph Campbell, Katia Sycara, Yaqi Xie. HiKER-SGG: Hierarchical Knowledge Enhanced Robust Scene Graph Generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2024. Also presented at *NeurIPS 2023 New Frontiers in Graph Learning Workshop*. [PDF] [Project] [Code]
- Xueting Hu, **Ce Zhang**, Yi Zhang, Bowen Hai, Ke Yu, and Zhihai He. Learning to Adapt CLIP for Few-Shot Monocular Depth Estimation. In *IEEE/CVF Winter Conference on Applications of Computer Vision* (*WACV*), 2024. [PDF]
- Yi Zhang, **Ce Zhang**, Ke Yu, Yushun Tang, Zhihai He. Concept-Guided Prompt Learning for Generalization in Vision-Language Models. In *AAAI Conference on Artificial Intelligence* (*AAAI*), 2024. [PDF]
- **Ce Zhang**, Kailiang Wu, and Zhihai He. Critical Sampling for Robust Evolution Operator Learning of Unknown Dynamical Systems. *IEEE Transactions on Artificial Intelligence* (*IEEE TAI*), 2023. Also presented at *First Workshop on Out-of-Distribution Generalization in Robotics at CoRL 2023.* [PDF]
- Yi Zhang*, **Ce Zhang***, Zihan Liao, Yushun Tang, and Zhihai He. BDC-Adapter: Brownian Distance Covariance for Better Vision-Language Reasoning. In *British Machine Vision Conference* (*BMVC*), 2023. [PDF] [Project]
- Yushun Tang, **Ce Zhang**, Heng Xu, Shuoshuo Chen, Jie Cheng, Luziwei Leng, *et al.* Neuro-Modulated Hebbian Learning for Fully Test-Time Adaptation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2023. [PDF]
- Zhehan Kan, Shuoshuo Chen, **Ce Zhang**, Yushun Tang, *et al.* Self-Correctable and Adaptable Inference for Generalizable Human Pose Estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2023. [PDF]

</> Selected Research Experience

#1. Self-Correcting Decoding with Generative Feedback for Mitigating Hallucinations in LVLMs *Jun 2024 - Oct 2024*

- Introduced self-correcting Decoding with Generative Feedback (DeGF), a training-free algorithm that incorporates feedback from text-to-image generative models into the decoding process to mitigate hallucinations in LVLMs.
- Demonstrated generative models can provide self-feedback for mitigating hallucinations at response/token levels.
- Outperformed state-of-the-art approaches in effectively mitigating hallucinations in LVLMs across five benchmarks.

#2. Dual Prototype Evolving for Test-Time Generalization of Vision-Language Models *Feb 2024 - May 2024*

- Proposed Dual Prototype Evolving (DPE) that effectively accumulates task-specific knowledge from multi-modalities.
- Introduced and optimized learnable residuals for each test sample to align the prototypes across modalities.
- Outperformed state-of-the-art methods across 15 datasets while maintaining competitive computational efficiency.

Honors and Awards

- *Liang Zhao Fellowship*, RI Departmental PhD Fellowship, CMU — *Oct 2025*
- *Top 10 Summa Cum Laude Graduates* (highest distinction, top 1%), SUSTech — *Jun 2023*
- *Honorary Undergraduate Thesis* (top 5%), SUSTech — *Jun 2023*
- *Top 10 Undergraduate Graduates* (top 2%), College of Engineering, SUSTech — *May 2023*
- *National Scholarship* (top 0.2%), Ministry of Education of the People's Republic of China — *Nov 2022*
- *School Motto Scholarship Special Award* (top 1%), SUSTech — *Nov 2022*
- *Outstanding Teaching Assistant Award*, SUSTech — *Jan 2022 & Jun 2022 & Jan 2023*
- *The First Prize of Outstanding Student Scholarship* (top 5%), SUSTech — *Nov 2020 & Nov 2021 & Nov 2022*

Academic Service

- Journal Reviewer: *IJCV, IEEE TCSVT (>10), IEEE TIP, IEEE TMM, IEEE TAI, Neurocomputing, Knowledge-Based Systems, Pattern Recognition*
- Conference Reviewer, *NeurIPS '25 '26, ICLR '25 '26, CVPR '26, ICCV '25, AAAI '26, MM '26, WACV '25 (Outstanding Reviewer), ICASSP '25, BMVC '24 '25 '26 (Outstanding Reviewer), ICME '24 '25, IJCNN '25, AVSS '25*
- Teaching Assistant, *Linear Algebra @ SUSTech for 3 semesters (Fall 2021, Spring 2022, Fall 2022)*
- Teaching Assistant, *Computer Vision @ CMU, Fall 2026*