Kuan-Hao Huang

Assistant Professor
Department of Computer Science and Engineering, Texas A&M University
Room 219, Peterson Building, 435 Nagle St, College Station, TX 77843

Tel: (979) 458-2205
Email: khhuang@tamu.edu
Website: https://khhuang.me

Research Interests

My research focuses on natural language processing and machine learning, with the goal of developing trustworthy and generalizable AI systems. I work on fundamental challenges in AI robustness, reliability, safety, fairness, and adaptability, aiming to ensure that large language models and vision-language models are reliable and socially responsible.

Academic Appointments

Texas A&M University, College Station, TX, USA — Aug. 2024 – Present
Assistant Professor, Department of Computer Science and Engineering

University of Illinois Urbana-Champaign, IL, USA — Aug. 2023 – Jun. 2024
Postdoctoral Research Associate, Department of Computer Science

Education

University of California Los Angeles, CA, USA — Sep. 2018 – Jun. 2023
Ph.D. in Computer Science
Advisor: Kai-Wei Chang

National Taiwan University, Taipei, Taiwan — Sep. 2014 – Jun. 2016
M.S. in Computer Science and Information Engineering
Advisor: Hsuan-Tien Lin

National Taiwan University, Taipei, Taiwan — Sep. 2010 – Jun. 2014
B.S. in Computer Science and Information Engineering

Selected Awards

Outstanding Senior Area Chair, EMNLP 2026 — 2026
First Place, Black-Box Track of Aletheia's Quest AI Lie Detection Challenge — 2026
Second Place, White-Box Track of Aletheia's Quest AI Lie Detection Challenge — 2026
Google Research Scholar Award — 2025
Area Chair Paper Award (*Outstanding Paper in the Semantics Track*), ACL 2023 — 2023
Thesis Honorable Mention Award, Taiwanese Association for Artificial Intelligence — 2016
Fourth Place, KDD Cup — 2015
Second Place, ICASSP Signal Processing Cup — 2014
First Place, Track 1 of KDD Cup — 2013
First Place, Track 2 of KDD Cup — 2013

Conference Publications

C36. **Pop-Up Distractions Reveal Bag-of-Events Behavior in Video Large Language Models**
Oscar Chew*, Serhii Honcharenko*, Qian-Hui Chen, Patricia Lu, Dishant Zaveri, Khoa D. Doan, and Kuan-Hao Huang (*equal contribution)
Proceedings of the 20th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 2027

C35. **TRACES: Proactive Safety Auditing for Multi-Turn LLM Agents via Trajectory-State Modeling**
Jiaqian Li, Yanshu Li, Boxuan Zhang, Ruixiang Tang, and Kuan-Hao Huang
Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2026

C34. **Steering Vector Fields for Context-Aware Inference-Time Control in Large Language Models**
Jiaqian Li, Yanshu Li, and Kuan-Hao Huang
Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2026

C33. **ConstructCIE: A Dataset for Extracting Causal Information from Construction Accident Narratives**
Hung Nguyen, Jaehoon Lee, Namgyun Kim, and Kuan-Hao Huang
Findings of the Association for Computational Linguistics: EMNLP (EMNLP-Findings), 2026

C32. **PEPPER: Perception-Guided Perturbation for Robust Backdoor Defense in Text-to-Image Diffusion Models**
Oscar Chew*, Po-Yi Lu*, Jayden Lin, Kuan-Hao Huang, and Hsuan-Tien Lin (*equal contribution)
Findings of the Association for Computational Linguistics: EMNLP (EMNLP-Findings), 2026

C31. **Is CLIP Cross-Eyed? Revealing and Mitigating Center Bias in the CLIP Family**
Oscar Chew, Hsiao-Ying Huang, Kunal Jain, Tai-I Chen, Khoa D. Doan, and Kuan-Hao Huang
Proceedings of the Third Conference on Language Modeling (COLM), 2026

C30. **From Papers to Panoramas: Building Hierarchies of Scientific Literature at Scale**
Muhan Gao, Jash Shah, Weiqi Wang, Kuan-Hao Huang, and Daniel Khashabi
Proceedings of the Third Conference on Language Modeling (COLM), 2026

C29. **The First Drop of Ink: Nonlinear Impact of Distracting Information in Long-Context Reasoning**
Muhan Gao, Zih-Ching Chen, and Kuan-Hao Huang
Proceedings of the Forty-Third International Conference on Machine Learning (ICML), 2026

C28. **Embodied Agents Meet Personalization: Investigating Challenges and Solutions Through the Lens of Memory Utilization**
Taeyoon Kwon*, Dongwook Choi*, Sunghwan Kim, Hyojun Kim, Seungjun Moon, Beong woo Kwak, Kuan-Hao Huang, and Jinyoung Yeo (*equal contribution)
Proceedings of the Fourteenth International Conference on Learning Representations (ICLR), 2026

C27. **What's Missing in Vision-Language Models? Probing Their Struggles with Causal Order Reasoning**
Zhaotian Weng, Haoxuan Li, Xin Eric Wang, Kuan-Hao Huang, and Jieyu Zhao
Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 2026

C26. **Adaptive Helpfulness-Harmlessness Alignment with Preference Vectors**
Ren-Wei Liang*, Chin-Ting Hsu*, Chan-Hung Yu, Saransh Agrawal, Shih-Cheng Huang, Chieh-Yen Lin, Shang-Tse Chen, Kuan-Hao Huang, and Shao-Hua Sun (*equal contribution)
Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 2026

C25. **Contrastive Visual Data Augmentation**
Yu Zhou*, Bingxuan Li*, Mohan Tang*, Xiaomeng Jin, Te-Lin Wu, Kuan-Hao Huang, Heng Ji, Kai-Wei Chang, and Nanyun Peng (*equal contribution)
Proceedings of the Forty-Second International Conference on Machine Learning (ICML), 2025

C24. **Eliminating Position Bias of Language Models: A Mechanistic Approach**
Ziqi Wang, Hanlin Zhang, Xiner Li, Kuan-Hao Huang, Chi Han, Shuiwang Ji, Sham M. Kakade, Hao Peng, and Heng Ji
Proceedings of the Thirteenth International Conference on Learning Representations (ICLR), 2025

C23. **TextEE: Benchmark, Reevaluation, Reflections, and Future Challenges in Event Extraction**
Kuan-Hao Huang, I-Hung Hsu, Tanmay Parekh, Zhiyu Xie, Zixuan Zhang, Premkumar Natarajan, Kai-Wei Chang, Nanyun Peng, and Heng Ji
Findings of the Association for Computational Linguistics: ACL (ACL-Findings), 2024

C22. **Contextual Label Projection for Cross-Lingual Structure Prediction**
Tanmay Parekh, I-Hung Hsu, Kuan-Hao Huang, Kai-Wei Chang, and Nanyun Peng
Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2024

C21. **Event Detection from Social Media for Epidemic Prediction**

Tanmay Parekh, Anh Mac, Jiarui Yu, Yuxuan Dong, Syed Shahriar, Bonnie Liu, Eric Yang, Kuan-Hao Huang, Wei Wang, Nanyun Peng, and Kai-Wei Chang
Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2024

C20. **Understanding and Mitigating Spurious Correlations in Text Classification**
Oscar Chew, Hsuan-Tien Lin, Kai-Wei Chang, and Kuan-Hao Huang
Findings of the Association for Computational Linguistics: EACL (EACL-Findings), 2024

C19. **Learning Easily Updated General Purpose Text Representations with Adaptable Task-Specific Prefixes**
Kuan-Hao Huang, Liang Tan, Rui Hou, Sinong Wang, Amjad Almahairi, and Ruty Rinott
Findings of the Association for Computational Linguistics: EMNLP (EMNLP-Findings), 2023

C18. **Self-Augmentation Improves Zero-Shot Cross-Lingual Transfer**
Fei Wang, Kuan-Hao Huang, Kai-Wei Chang, and Muhao Chen
Proceedings of the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (AACL), 2023

C17. **ParaAMR: A Large-Scale Syntactically Diverse Paraphrase Dataset by AMR Back-Translation**
Kuan-Hao Huang, Varun Iyer, I-Hung Hsu, Anoop Kumar, Kai-Wei Chang, and Aram Galstyan
Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL), 2023
Area Chair Paper Award

C16. **TAGPRIME: A Unified Framework for Relational Structure Extraction**
I-Hung Hsu*, Kuan-Hao Huang*, Shuning Zhang, Wenxin Cheng, Premkumar Natarajan, Kai-Wei Chang, and Nanyun Peng (*equal contribution)
Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL), 2023

C15. **AMPERE: AMR-Aware Prefix for Generation-Based Event Argument Extraction Model**
I-Hung Hsu*, Zhiyu Xie*, Kuan-Hao Huang, Premkumar Natarajan, and Nanyun Peng (*equal contribution)
Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL), 2023

C14. **GENEVA: Benchmarking Generalizability for Event Argument Extraction with Hundreds of Event Types and Argument Roles**
Tanmay Parekh, I-Hung Hsu, Kuan-Hao Huang, Kai-Wei Chang, and Nanyun Peng
Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL), 2023

C13. **PIP: Parse-Instructed Prefix for Syntactically Controlled Paraphrase Generation**
Yixin Wan, Kuan-Hao Huang, and Kai-Wei Chang
Findings of the Association for Computational Linguistics: ACL (ACL-Findings), 2023

C12. **Unsupervised Syntactically Controlled Paraphrase Generation with Abstract Meaning Representations**
Kuan-Hao Huang*, Varun Iyer*, Anoop Kumar, Sriram Venkatapathy, Kai-Wei Chang, and Aram Galstyan (*equal contribution)
Findings of the Association for Computational Linguistics: EMNLP (EMNLP-Findings), 2022

C11. **DEGREE: A Data-Efficient Generation-Based Event Extraction Model**
I-Hung Hsu*, Kuan-Hao Huang*, Elizabeth Boschee, Scott Miller, Premkumar Natarajan, Kai-Wei Chang, and Nanyun Peng (*equal contribution)
Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2022

C10. **Multilingual Generative Language Models for Zero-Shot Cross-Lingual Event Argument Extraction**
Kuan-Hao Huang*, I-Hung Hsu*, Premkumar Natarajan, Kai-Wei Chang, and Nanyun Peng (*equal contribution)
Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL), 2022

C9. **Improving Zero-Shot Cross-Lingual Transfer Learning via Robust Training**
Kuan-Hao Huang, Wasi Uddin Ahmad, Nanyun Peng, and Kai-Wei Chang
Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2021

C8. **Disentangling Semantics and Syntax in Sentence Embeddings with Pre-trained Language Models**
James Y. Huang, Kuan-Hao Huang, and Kai-Wei Chang

Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2021

C7. **Generating Syntactically Controlled Paraphrases without Using Annotated Parallel Pairs**
Kuan-Hao Huang and Kai-Wei Chang
Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 2021

C6. **Generating Sports News from Live Commentary: A Chinese Dataset for Sports Game Summarization**
Kuan-Hao Huang, Chen Li, and Kai-Wei Chang
Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (AACL), 2020

C5. **JECL: Joint Embedding and Cluster Learning for Image-Text Pairs**
Sean T. Yang, Kuan-Hao Huang, and Bill Howe
Proceedings of the 25th International Conference on Pattern Recognition (ICPR), 2020

C4. **Examining Gender Bias in Languages with Grammatical Gender**
Pei Zhou, Weijia Shi, Jieyu Zhao, Kuan-Hao Huang, Muhao Chen, Ryan Cotterell, and Kai-Wei Chang
Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2019

C3. **Cost-Sensitive Reference Pair Encoding for Multi-Label Learning**
Yao-Yuan Yang, Kuan-Hao Huang, Chih-Wei Chang, and Hsuan-Tien Lin
Advances in Knowledge Discovery and Data Mining - 22nd Pacific-Asia Conference (PAKDD), 2018

C2. **A Novel Uncertainty Sampling Algorithm for Cost-Sensitive Multiclass Active Learning**
Kuan-Hao Huang and Hsuan-Tien Lin
Proceedings of the IEEE 16th International Conference on Data Mining (ICDM), 2016

C1. **Linear Upper Confidence Bound Algorithm for Contextual Bandit Problem with Piled Rewards**
Kuan-Hao Huang and Hsuan-Tien Lin
Advances in Knowledge Discovery and Data Mining - 20th Pacific-Asia Conference (PAKDD), 2016

JOURNAL PUBLICATIONS

J7. **Generate, Filter, Control, Replay: A Comprehensive Survey of Rollout Strategies for LLM Reinforcement Learning**
Rohan Surana, Gagan Mundada, Xunyi Jiang, Chuhan Wang, Zhenwei Tang, Difan Jiao, Zihan Huang, Yuxin Xiong, Junda Wu, Sheldon Yu, Xintong Li, Raghav Jain, Nikki Lijing Kuang, Sizhe Zhou, Bowen Jin, Zhendong Chu, Tong Yu, Ryan Rossi, Kuan-Hao Huang, Jingbo Shang, Jiawei Han, and Julian McAuley
Transactions on Machine Learning Research, 2026

J6. **FARM: Enhancing Molecular Representations with Functional Group Awareness**
Thao Nguyen, Kuan-Hao Huang, Ge Liu, Martin D. Burke, Ying Diao, and Heng Ji
Transactions on Machine Learning Research, 2026

J5. **Visually Descriptive Language Model for Vector Graphics Reasoning**
Zhenhailong Wang, Joy Hsu, Xingyao Wang, Kuan-Hao Huang, Manling Li, Jiajun Wu, and Heng Ji
Transactions on Machine Learning Research, 2025

J4. **Dynamic Principal Projection for Cost-Sensitive Online Multi-Label Classification**
Hong-Min Chu, Kuan-Hao Huang, and Hsuan-Tien Lin
Machine Learning (ECML PKDD Journal Track), 2019

J3. **Cost-Sensitive Label Embedding for Multi-Label Classification**
Kuan-Hao Huang and Hsuan-Tien Lin
Machine Learning (ECML PKDD Journal Track), 2017

J2. **Combination of Feature Engineering and Ranking Models for Paper-Author Identification in KDD Cup 2013**
Chun-Liang Li, Yu-Chuan Su, Ting-Wei Lin, Cheng-Hao Tsai, Wei-Cheng Chang, Kuan-Hao Huang, Tzu-Ming Kuo, Shan-Wei Lin, Young-San Lin, Yu-Chen Lu, Chun-Pai Yang, Cheng-Xia Chang, Wei-Sheng Chin, Yu-Chin Juan, Hsiao-Yu Tung, Jui-Pin Wang, Cheng-Kuang Wei, Felix Wu, Tu-Chun Yin, Tong Yu, Yong Zhuang, Shou-De

Lin, Hsuan-Tien Lin, and Chih-Jen Lin
Journal of Machine Learning Research, 2015
Extended first-place winner report of KDD Cup 2013 track 1

J1. **Effective String Processing and Matching for Author Disambiguation**
Wei-Sheng Chin, Yong Zhuang, Yu-Chin Juan, Felix Wu, Hsiao-Yu Tung, Tong Yu, Jui-Pin Wang, Cheng-Xia Chang, Chun-Pai Yang, Wei-Cheng Chang, Kuan-Hao Huang, Tzu-Ming Kuo, Shan-Wei Lin, Young-San Lin, Yu-Chen Lu, Yu-Chuan Su, Cheng-Kuang Wei, Tu-Chun Yin, Chun-Liang Li, Ting-Wei Lin, Cheng-Hao Tsai, Shou-De Lin, Hsuan-Tien Lin, and Chih-Jen Lin
Journal of Machine Learning Research, 2014
Extended first-place winner report of KDD Cup 2013 track 2

Other Publications

O8. **Look Before You Zoom: Adaptive Routing for the Resolution-Context Trade-off in Visual RAG**
Oanh N. Tran, Thanh Quoc Hung Le, Oscar Chew, Kuan-Hao Huang†, and Khoa D. Doan† (†jointly supervising)
ICML Workshop on Efficient Multimodal Question Answering (EMM-QA@ICML), 2026

O7. **Accident Causality Extraction from Construction Accident Reports Using a Large Language Model**
Jaehoon Lee, Taewoo Ko, Kuan-Hao Huang, Jeewoong Park, Jungho Jeon, Brian A. Anderson, and Namgyun Kim
Proceedings of ASCE International Conference on Computing in Civil Engineering (i3CE), 2026

O6. **The Impact of Low-Quality Accident Reports on the Performance of Construction Safety AI Models**
Jaehoon Lee, Taewoo Ko, Kuan-Hao Huang, JeeWoong Park, and Namgyun Kim
Proceedings of ASCE International Conference on Computing in Civil Engineering (i3CE), 2025

O5. **SHA256 at SemEval-2025 Task 4: Selective Amnesia – Constrained Unlearning for Large Language Models via Knowledge Isolation**
Saransh Agrawal and Kuan-Hao Huang
Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval), 2025

O4. **ARMADA: Attribute-Based Multimodal Data Augmentation**
Xiaomeng Jin, Jeonghwan Kim, Yu Zhou, Kuan-Hao Huang, Te-Lin Wu, Nanyun Peng, and Heng Ji
EMNLP Workshop on Advancing Natural Language Process for Wikipedia (WikiNLP@EMNLP), 2024

O3. **MIRACLE: An Online, Explainable Multimodal Interactive Concept Learning System**
Ansel Blume, Khanh Duy Nguyen, Zhenhailong Wang, Yangyi Chen, Michal Shlapentokh-Rothman, Xiaomeng Jin, Jeonghwan Kim, Zhen Zhu, Jiateng Liu, Kuan-Hao Huang, Mankeerat Sidhu, Xuanming Zhang, Vivian Liu, Raunak Sinha, Te-Lin Wu, Abhay Zala, Elias Stengel-Eskin, Da Yin, Yao Xiao, Utkarsh Mall, Zhou Yu, Kai-Wei Chang, Camille Cobb, Karrie Karahalios, Lydia Chilton, Mohit Bansal, Nanyun Peng, Carl Vondrick, Derek Hoiem, and Heng Ji
Proceedings of the 32nd ACM International Conference on Multimedia: Technical Demos and Videos Program (ACMMM-Demo), 2024

O2. **Zero-Shot Cross-Lingual Sequence Tagging as Seq2Seq Generation for Joint Intent Classification and Slot Filling**
Fei Wang, Kuan-Hao Huang, Anoop Kumar, Aram Galstyan, Greg Ver Steeg, and Kai-Wei Chang
EMNLP Workshop on Massively Multilingual Natural Language Understanding (MMNLU@EMNLP), 2022

O1. **A Comparative Survey of Deep Active Learning**
Xueying Zhan, Qingzhong Wang, Kuan-Hao Huang, Haoyi Xiong, Dejing Dou, and Antoni B. Chan
NeurIPS Workshop on Human in the Loop Learning (HiLL@NeurIPS), 2022

Preprints

P7. **Re-Search: Dissecting LLM Ideation Through a Two-Player Questioning Game**
Muhan Gao, Hangxiao Zhu, Xinqi Su, Yu Zhang, Daniel Khashabi†, and Kuan-Hao Huang† (†jointly supervising)
arXiv preprint arXiv:2611.14487, 2026

P6. **Not Every Mention Is a Claim: Rethinking Object Hallucination Detection in Vision-Language Models**
Hsin-Ling Hsu*, Neeraja Kirtane*, and Kuan-Hao Huang (*equal contribution)

arXiv preprint arXiv:2610.23610, 2026

P5. **Are We Really Measuring Forgetting? Pitfalls in Evaluating Machine Unlearning for Multimodal LLMs**
Oscar Chew*, Po-Yi Lu*, Yi-Jie Cheng, Kunal Jain, Hsuan-Tien Lin, Zhengzhong Tu, and Kuan-Hao Huang (*equal contribution)
arXiv preprint arXiv:2610.19613, 2026

P4. **ChronoCulture: Evaluating Cultural Awareness Across Time in Large Language Models**
Rusali Saha, Nitya Khurana, Ruihong Huang, and Kuan-Hao Huang
arXiv preprint arXiv:2610.20472, 2026

P3. **TwinICL: Diagnosing Multimodal In-Context Learning through Paired Counterfactuals**
Zihan Xue, Po-Yi Lu*, Serhii Honcharenko*, Zih-Ching Chen, Hsuan-Tien Lin, Nanyun Peng, I-Hung Hsu, and Kuan-Hao Huang (*equal contribution)
arXiv preprint arXiv:2609.15028, 2026

P2. **When Correct Demonstrations Hurt: Rethinking the Role of Exemplars in In-Context Learning**
Chenghao Qiu, Chunli Peng, Yufeng Yang, Kuan-Hao Huang, and Yi Zhou
arXiv preprint arXiv:2605.26350, 2026

P1. **Language Steering for Multilingual In-Context Learning**
Neeraja Kirtane and Kuan-Hao Huang
arXiv preprint arXiv:2602.02326, 2026

ADVISING

Ph.D. Students

Taeyoon Kwon, Computer Science, TAMU	2026 – Present
Neeraja Kirtane, Computer Science, TAMU	2026 – Present
Muhan Gao, Computer Science, TAMU	2025 – Present
Oscar Chew, Computer Science, TAMU	2025 – Present
Chan-Wei Hu, Computer Science, TAMU *(Rotating Student)*	2024 – 2024

M.S. Students

Hung Nguyen, Computer Science, TAMU	2026 – Present
Junggeun Do, Computer Science, TAMU	2025 – Present
Bhaskar Ruthvik Bikkina, Computer Science, TAMU	2025 – Present
Sicong Liang, Computer Science, TAMU *(Next: Ph.D. Student at UCF)*	2025 – 2026
Rusali Saha, Computer Science, TAMU *(Next: Associate in Software Group at Global Prior Art)*	2025 – 2026
Dishant Zaveri, Computer Science, TAMU *(Next: Technical Associate at Jefferies Group)*	2024 – 2026
Dheeraj Mudireddy, Data Science, TAMU *(Next: Data Scientist at Amazon)*	2025 – 2025
Rahul Baid, Computer Science, TAMU *(Next: Senior Software Engineer at Mubama Group)*	2024 – 2025
Saransh Agrawal, Data Science, TAMU *(Next: AI Engineer at Saptarishis Software)*	2024 – 2025

Undergraduate Students

Serhii Honcharenko, Computer Science, TAMU	2026 – Present
Kunal Jain, Computer Science, TAMU *(Next: Software Engineer at Google)*	2025 – 2026

Visiting Students

Yi-Jie Cheng, M.S. Student at Columbia	2026 – Present
Ting-Chen Yang, Undergrad at UC Berkeley	2026 – Present
Cheng-Yu Wang, M.S. Student at USC	2026 – Present
Xingrun Chen, Undergrad at UCLA	2026 – Present
Qian-Hui Chen, M.S. Student at NTU	2026 – Present
Jiongran Wang, Ph.D. Student at TAMU	2026 – Present
Hsin-Ling Hsu, Undergrad at NCCU	2025 – Present
Po-Yi Lu, Ph.D. Student at NTU	2025 – 2026

Jiaqian Li, M.S. Student at Brown 2025 – 2026
Anyi Wang, M.S. Student at LMU Munich *(Next: Ph.D. Student at TAMU)* 2025 – 2026
Neeraja Kirtane, M.S. Student at UIUC *(Next: Ph.D. Student at TAMU)* 2025 – 2026
Zihan Xue, Undergrad at UCLA *(Next: Ph.D. Student at Northeastern)* 2025 – 2026
Zhaotian Weng, M.S. Student at USC *(Next: Ph.D. Student at UC Santa Barbara)* 2024 – 2025

Ph.D. Students (On Committee)

Chenghao Qiu, Computer Science, TAMU *(Advisor: Yi Zhou)* Present
Zhaonan Li, Computer Science, ASU *(Advisor: Ben Zhou)* Present
Mu-Ruei Tseng, Computer Science, TAMU *(Advisor: Ricardo Gutierrez-Osuna)* Present
Lianhao Zhou, Computer Science, TAMU *(Advisor: Shuiwang Ji)* Present
Yuyang Bai, Computer Science, TAMU *(Advisor: Yu Zhang)* Present
Siyuan Yang, Computer Science, TAMU *(Advisor: Zhengzhong Tu)* Present
Jaehoon Lee, Construction Science, TAMU *(Advisor: Namgyun Kim)* Present
Hangxiao Zhu, Computer Science, TAMU *(Advisor: Yu Zhang)* Present
Xiangbo Gao, Computer Science, TAMU *(Advisor: Zhengzhong Tu)* Present
Xingyu Su, Computer Science, TAMU *(Advisor: Shuiwang Ji)* Present
Chan-Wei Hu, Computer Science, TAMU *(Advisor: Zhengzhong Tu)* Present

M.S. Students (On Committee)

Pranavi Deshpande, Electrical Engineering, TAMU *(Advisor: I-Hong Hou)* Present
Arunima Chowdhury, Computer Science, TAMU *(Advisor: Ruihong Huang)* Present
Dhawal Jiten Chaudhari, Computer Science, TAMU *(Advisor: Xin Li)* Present
Pranati Tyagi, Computer Science, TAMU *(Advisor: James Caverlee)* Present
Mihir Santosh Sanjay, Computer Science, TAMU *(Advisor: Shreyas Kumar)* Present
Duc Nguyen, Computer Science, TAMU *(Advisor: Yu Zhang)* Present
Yi Wen, Computer Science, TAMU *(Advisor: Meng Xia)* 2026
Mingyang Wu, Computer Science, TAMU *(Advisor: Zhengzhong Tu)* 2026
Michael Norman, Computer Science, TAMU *(Advisor: James Caverlee)* 2025
Divij Bajaj, Computer Science, TAMU *(Advisor: Ruihong Huang)* 2025

TEACHING

Instructor, Texas A&M University, College Station, TX

CSCE 471: Natural Language Processing Spring 2027
CSCE 689: Special Topics in Advanced Natural Language Processing Fall 2026
CSCE 638: Natural Language Processing: Foundation and Techniques Spring 2026
CSCE 689: Special Topics in Trustworthy Natural Language Processing Fall 2025
CSCE 638: Natural Language Processing: Foundation and Techniques Spring 2025
CSCE 689: Special Topics in Trustworthy Natural Language Processing Fall 2024

SELECTED TUTORIALS AND TALKS

Tutorial on Large Language Models

Tutorial, AI Bridge Summit, Texas A&M University Jan. 2026

Towards Robust and Trustworthy Large Language Models: Issues and Mitigation Strategies

Tutorial, The 17th Asian Conference on Machine Learning (ACML) Dec. 2025

Challenges of Multimodal AI Foundation Models in Scientific Discovery

Guest Lecture, DATA 440/691, The College of William & Mary Apr. 2026

Controlling Large Language Models via Inference-Time Steering

Guest Lecture, DATA 405/603, The College of William & Mary Oct. 2026
Invited Talk, Computer Engineering and Systems Group Seminar, Texas A&M University Sep. 2026

Invited Talk, Rice Workshop on Large Language Models, Rice University — Mar. 2026

Toward Robust and Reliable Large Language Models

Invited Talk, Graduate Seminar, Texas A&M University — Feb. 2026
Guest Lecture, CS 6501, University of Virginia — Oct. 2025
Invited Talk, Data Science Colloquium, The College of William & Mary — Oct. 2025
Invited Talk, AI Seminar Series, Rice University — Oct. 2025
Invited Talk, Institute of Information Science, Academia Sinica — Jun. 2025
Invited Talk, National Tsing Hua University — Jun. 2025
Invited Talk, National Taiwan University — Jun. 2025

PROFESSIONAL SERVICES

Organizing Committee

Texas NLP Symposium — 2026

Editorial Roles

Action Editor, ACL Rolling Review (ARR) — 2024 – Present
Action Editor, Transactions on Machine Learning Research (TMLR) — 2026 – Present

Conference Senior Area Chair

Meeting of the Association for Computational Linguistics (ACL) — 2027
Conference on Empirical Methods in Natural Language Processing (EMNLP) — 2025 – 2026

Conference Area Chair

Meeting of the Association for Computational Linguistics (ACL) — 2024 – 2026
Conference on Empirical Methods in Natural Language Processing (EMNLP) — 2024
Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL) — 2025 – 2027
Conference of the European Chapter of the Association for Computational Linguistics (EACL) — 2026 – 2027
International Conference on Machine Learning (ICML) — 2026 – 2027
Conference on Neural Information Processing Systems (NeurIPS) — 2025 – 2026
International Conference on Learning Representations (ICLR) — 2026 – 2027

Conference Reviewer

Meeting of the Association for Computational Linguistics (ACL) — 2021 – 2023
Conference on Empirical Methods in Natural Language Processing (EMNLP) — 2021 – 2023
Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL) — 2022 – 2024
Conference of the European Chapter of the Association for Computational Linguistics (EACL) — 2023 – 2024
Conference on Language Modeling (COLM) — 2024 – 2026
International Conference on Machine Learning (ICML) — 2020 – 2025
Conference on Neural Information Processing Systems (NeurIPS) — 2021 – 2024
International Conference on Learning Representations (ICLR) — 2021 – 2025
AAAI Conference on Artificial Intelligence (AAAI) — 2022 – 2025

Journal Reviewer

Computational Linguistics (CL) — 2025 – Present
Journal of Machine Learning Research (JMLR) — 2026 – Present
Transactions on Machine Learning Research (TMLR) — 2024 – 2026

Handbook Assistant

Conference on Empirical Methods in Natural Language Processing (EMNLP) — 2018

Grant-Proposal Review Panelists

NSF CISE Directorate — 2025 – 2026

Faculty Advisor

Texas A&M AI Safety (Student Organization) — 2026 – Present