

Curriculum Vitae

Yuancheng Xu

Email: xuyuancheng0@gmail.com

Education

University of Maryland, College Park

Ph.D. in Applied Mathematics & Statistics, and Scientific Computation

2020-2025

Advisor: Furong Huang

Southern University of Science and Technology, China

2016-2020

B.S. in Mathematics and Applied Mathematics

GPA 3.94/4.00 (1/909), Summa Cum Laude (1%)

Work and Research Experience

- **Netflix**

June 2025 - Present

Research Scientist

- Building Netflix video foundation models
- Developing multimodal models with advanced understanding, reasoning, interaction, and generative capabilities

- **University of Maryland, College Park**

Sept 2021 - March 2025

Research Assistant

- Large language models alignment and safety, adversarial robustness, multimodal understanding and reasoning, and reinforcement learning

- **Netflix Eyeline Labs**

Sept 2024 - March 2025

Research Intern

- Motion control and 3D consistency in video generation model

- **JPMorgan Chase & Co**

June 2024 – Aug 2024

AI Research Intern

- Test-time alignment of large language model

- **Comcast**

June 2023 – Aug 2023

Machine Learning Graduate Research Intern

- Scene-text understanding via vision-language models

- **University of California, Los Angeles**

May 2019 – Sept 2019

Cross-disciplinary Scholars in Science and Technology (CSST) Program

- Tree-based methods for longitudinal analysis

- **New York University**

May 2018 – Sept 2018

Undergraduate Research Program

- Topics on computational neuroscience regarding working memory

Publications

-
1. Zhengming Yu, Li Ma, Mingming He, Leo Isikdogan, **Yuancheng Xu**, Dmitriy Smirnov, Pablo Salamanca, Dao Mi, Pablo Delgado, Ning Yu, Julien Philip, Xin Li, Wenping Wang, Paul Debevec. “DiffHDR: Re-Exposing LDR Videos with Video Diffusion Models” In *The European Conference on Computer Vision (ECCV)*, 2026.
 2. Koichi Namekata, Yash Kant, Zhizheng Liu, Ryan Burgert, **Yuancheng Xu**, Kuan Heng Lin, Emmett Steven, Julien Philip, Li Ma, Andrea Vedaldi, Paul Debevec, Ning Yu. “Go-with-the-Track: Video Compositing and Motion Control with Point Tracking” In *SIGGRAPH*, 2026.
 3. Kuan Heng Lin, Zhizheng Liu, Pablo Salamanca, Yash Kant, Ryan D Burgert, **Yuancheng Xu**, Koichi Namekata, Yiwei Zhao, Bolei Zhou, Micah Goldblum, Paul Debevec, Ning Yu. “Vista4D: Video Reshooting with 4D Point Clouds”. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.
 4. Udari Madhushani Schwag, Shayan Shabihi, Alex McAvoy, Vikash Schwag, **Yuancheng Xu**, Dalton Towers, Furong Huang. “PropensityBench: Evaluating Latent Safety Risks in Large Language Models via an Agentic Approach”. In *International Conference on Learning Representations (ICLR)*, 2026.
 5. **Yuancheng Xu**, Wenqi Xian, Li Ma, Julien Philip, Ahmet Taşel, Yiwei Zhao, Ryan Burgert, Mingming He, Oliver Hermann, Oliver Pilarski, Rahul Garg, Paul Debevec, Ning Yu. “Virtually Being: Customizing Camera-Controllable Video Diffusion Models with Volumetric Performance Captures”. In *SIGGRAPH Asia*, 2025.
 6. Baijiong Lin, Weisen Jiang, **Yuancheng Xu**, Hao Chen, Ying-Cong Chen. “Parm: Multi-objective Test-time Alignment via Preference-aware Autoregressive Reward Model”. In *International Conference on Machine Learning (ICML)*, 2025.
 7. Xiyao Wang, Zhengyuan Yang, Linjie Li, Hongjin Lu, **Yuancheng Xu**, Chung-Ching Lin, Kevin Lin, Furong Huang, Lijuan Wang. “Scaling Inference-time Search with Vision Value Model for Improved Visual Comprehension”. In *International Conference on Computer Vision (ICCV)*, 2025.
 8. **Yuancheng Xu**, Udari Madhushani Schwag, Alec Koppel, Sicheng Zhu, Bang An, Furong Huang, Sumitra Ganesh. “GenARM: Reward Guided Generation with Autoregressive Reward Model for Test-Time Alignment”. In *International Conference on Learning Representations (ICLR)*, 2025.

9. Ryan D Burgert, **Yuancheng Xu**, Wenqi Xian, Oliver Pilarski, Pascal Clausen, Mingming He, Li Ma, Yitong Deng, Lingxiao Li, Mohsen Mousavi, Michael S Ryoo, Paul Debevec, Ning Yu. “Go-with-the-Flow: Motion-Controllable Video Diffusion Models Using Real-Time Warped Noise Flow”. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Oral Presentation), 2025.
10. Michael-Andrei Panaitescu-Liess, Zora Che, Bang An, **Yuancheng Xu**, Pankayaraj Pathmanathan, Souradip Chakraborty, Sicheng Zhu, Tom Goldstein, Furong Huang. “Can Watermarking Large Language Models Prevent Copyrighted Text Generation and Hide Training Data?”. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2025. In *New Frontiers in Adversarial Machine Learning Workshop at Neural Information Processing Systems (NeurIPS)* (Best Paper Award), 2024.
11. **Yuancheng Xu**, Jiarui Yao, Manli Shu, Yanchao Sun, Zichu Wu, Ning Yu, Tom Goldstein, Furong Huang. “Shadowcast: Stealthy Data Poisoning Attacks Against Vision-Language Models”. In *Neural Information Processing Systems (NeurIPS)*, 2024.
12. Bang An, Sicheng Zhu, Ruiyi Zhang, Michael-Andrei Panaitescu-Liess, **Yuancheng Xu**, Furong Huang. “Automatic Pseudo-Harmful Prompt Generation for Evaluating False Refusals in Large Language Models”. In *Conference on Language Modeling (COLM)*, 2024
13. **Yuancheng Xu**, Chenghao Deng, Yanchao Sun, Ruijie Zheng, Xiyao Wang, Jieyu Zhao, Furong Huang. “Adapting Static Fairness to Sequential Decision-Making: Bias Mitigation Strategies towards Equal Long-term Benefit Rate”. In *International Conference on Machine Learning (ICML)*, 2024.
14. Bang An, Mucong Ding, Tahseen Rabbani, Aakriti Agrawal, **Yuancheng Xu**, Chenghao Deng, Sicheng Zhu, Abdirisak Mohamed, Yuxin Wen, Tom Goldstein, Furong Huang. “Benchmarking the Robustness of Image Watermarks”. In *International Conference on Machine Learning (ICML)*, 2024.
15. Xiyao Wang, Yuhang Zhou, Xiaoyu Liu, Hongjin Lu, **Yuancheng Xu**, Feihong He, Jaehong Yoon, Taixi Lu, Gedas Bertasius, Mohit Bansal, Huaxiu Yao, Furong Huang. “Mementos: A Comprehensive Benchmark for Multimodal Large Language Model Reasoning over Image Sequences”. In *Association for Computational Linguistics (ACL)*, 2024.
16. Mucong Ding, Bang An, **Yuancheng Xu**, Anirudh Satheesh, Furong Huang. “SAFLEX: Self-Adaptive Augmentation via Feature Label Extrapolation”. In *International Conference on Learning Representations (ICLR)*, 2024.
17. Xiaoyu Liu, Jiaxin Yuan, Bang An, **Yuancheng Xu**, Yifan Yang, Furong Huang. “C-Disentanglement: Discovering Causally-Independent Generative Factors under an Inductive Bias of Confounder”. In *Neural Information Processing Systems (NeurIPS)*, 2023.
18. **Yuancheng Xu**, Yanchao Sun, Micah Goldblum, Tom Goldstein, Furong Huang. “Exploring and Exploiting Decision Boundary Dynamics for Adversarial Robustness”. In *International Conference on Learning Representations (ICLR)*, 2023.

Awards

Dean’s fellowship, University of Maryland, College Park	2020,2021
Summa Cum Laude at Southern University of Science and Technology (10/1000)	2020
China National Scholarship (0.2%, Highest honor of Chinese undergraduate students)	2019