

MING LI

☎ (+1) 407-881-4319 ✉ limingcv@gmail.com 🌐 [Personal Website](#) 🎓 [Google Scholar](#)

ABOUT ME

I am a Ph.D. candidate at the University of Central Florida and an intern at [ByteDance-Seed](#) (USA), working on reinforcement learning (RL) for multimodal generation. My Ph.D. research is advised by [Chen Chen](#), and my work at ByteDance-Seed is mentored by [Jie Wu](#) and [Rui Wang](#).

I am a core contributor (RL) to the **Seedance series** (1.5, 2.0, 2.5), [Seedream 5.0 Pro](#), and other production-oriented multimodal generation projects, such as streaming long-video generation and world models.

EXPERIENCE

Intern	ByteDance-Seed, USA	Mentors: Jie Wu , Rui Wang	01/2025 – Now
Intern	ByteDance Inc, USA	Mentor: Sijie Zhu	05/2024 – 12/2024
Intern	ByteDance Inc, China	Mentor: Jie Wu	01/2022 – 07/2023

EDUCATION

Ph.D.	Computer Science	University of Central Florida, USA	08/2023 – 12/2026
M.S.	Computer Science and Technology	Xiamen University, China	09/2020 – 06/2023
B.S.	Software Engineering	Hainan University, China	09/2016 – 06/2020

INDUSTRY PROJECTS

- **Seedance 2.5** [[Project](#)] 03/2026 – 06/2026
Owned reward model development for Seedance 2.5 video extension, covering both data construction and RM training pipeline optimization, which produced a high-accuracy reward model and, through RL training, significantly improved video extension quality over the SFT model. Also built Auto-Rubric, which automatically optimizes task-specific evaluation criteria from human-annotated data, giving more accurate rewards across diverse video generation tasks and further improving generation quality.
- **Seedream 5.0 Pro** [[Project](#)] 04/2026 – 06/2026
Operationalized the abstract concept of image vividness, a key perceptual quality dimension, into a structured set of measurable evaluation criteria, then applied Auto-Rubric to optimize the resulting rubric against a human-annotated benchmark and used the resulting reward in RL to improve vividness over the SFT baseline.
- **Seedance 2.0** [[Project](#) | [Paper](#)] 12/2025 – 02/2026
Led one of the RL workstreams end-to-end, which substantially improved reference-based audio-video generation quality: built the training code from scratch on top of the SFT codebase, optimized parallel training by distributing the paired forward passes across GPUs (~40% higher training efficiency), unblocked large-scale RL data construction by leveraging preemptible computation resources on other clusters, and explored the optimal data mixture and training strategy.
- **Seedance 1.5 Pro** [[Project](#) | [Paper](#)] 09/2025 – 12/2025
Implemented on-policy diffusion RL algorithms (FPO, DiffusionNFT, AWM, and GRPO) on Seedance 1.5 Pro, which improved the video and audio reward scores and yielded noticeably better generation quality over the SFT baseline, verifying the effectiveness of on-policy RL on a production-scale video generation model.
- **RewardDance: Reward Scaling in Visual Generation** [[Paper](#)] 01/2025 – 09/2025
Trained and scaled reward models for visual generation from 1B to 26B, demonstrating that RM scaling brings substantial gains on both Seedream 3.0 (text-to-image) and Seedance 1.0 (text-to-video), with up to +49% improvement in human preference win rate. Combined with ReFL training, these scaled reward models became a key factor in Seedance 1.0's advantage over competing models.

SELECTED PUBLICATIONS (COMPLETE LIST...)

In total: 20+ papers, 1000+ citations, 1000+ GitHub stars.

- **Seedance 2.0: Advancing Video Generation for World Complexity**
Team Seedance (Ming Li as core contributor, RL)
[[ByteDance-Seed Tech Report](#) | [Project](#) | [Paper](#)]
- **LoL: Longer than Longer, Scaling Video Generation to Hour**
Justin Cui, Jie Wu, Ming Li, Tao Yang, Xiaojie Li, Rui Wang, Andrew Bai, Yuanhao Ban, Cho-Jui Hsieh
[[CVPR'26](#) | [Paper](#)]
- **ViPO: Visual Preference Optimization at Scale**
Ming Li, Jie Wu, Justin Cui, Xiaojie Li, Rui Wang, Chen Chen
[[ICLR'26](#) | [Project](#) | [Paper](#)]
- **Self-Forcing++: Towards Minute-Scale High-Quality Video Generation**
Justin Cui, Jie Wu, Ming Li, Tao Yang, Xiaojie Li, Rui Wang, Andrew Bai, Yuanhao Ban, Cho-Jui Hsieh
[[ICLR'26](#) | [Project](#) | [Paper](#)]
- **Learning from Noisy Preferences: A Semi-Supervised Learning Approach to Direct Preference Optimization**
Xinxin Liu, Ming Li, Zonglin Lyu, Yuzhang Shang, Chen Chen
[[ICLR'26](#) | [Project](#) | [Paper](#)]
- **Seedance 1.5 Pro: A Native Audio-Visual Joint Generation Foundation Model**
Team Seedance (Ming Li as core contributor, RL)
[[ByteDance-Seed Tech Report](#) | [Project](#) | [Paper](#)]
- **CPO: Condition Preference Optimization for Controllable Image Generation**
Zonglin Lyu, Ming Li, Xinxin Liu, Chen Chen
[[NeurIPS'25](#) | [Project](#) | [Paper](#) | [Code](#)]
- **SuperEdit: Rectifying and Facilitating Supervision for Instruction-based Image Editing**
Ming Li, Xin Gu, Fan Chen, Xiaoying Xing, Longyin Wen, Chen Chen, Sijie Zhu
[[ICCV'25](#) | [Paper](#) | [Code](#) | [GitHub Stars 100+](#)]
- **RewardDance: Reward Scaling in Visual Generation**
Jie Wu, Yu Gao, Zilyu Ye, Ming Li, Liang Li, Hanzhong Guo, Jie Liu, Zeyue Xue, Xiaoxia Hou, Wei Liu, Yan Zeng, Weilin Huang
[[ByteDance-Seed Tech Report](#) | [Paper](#)]
- **Multi-Reward as Condition for Instruction-based Image Editing**
Xin Gu, Ming Li, Libo Zhang, Fan Chen, Longyin Wen, Tiejian Luo, Sijie Zhu
[[ICLR'25](#) | [Paper](#) | [Code](#)]
- **ControlNet++: Improving Conditional Controls with Efficient Consistency Feedback**
Ming Li, Taojiannan Yang, Huafeng Kuang, Jie Wu, Zhaoning Wang, Xuefeng Xiao, Chen Chen
[[ECCV'24](#) | [Project](#) | [Paper](#) | [Code](#) | [GitHub Stars 500+](#)]
- **AlignDet: Aligning Pre-training and Fine-tuning in Object Detection**
Ming Li, Jie Wu, Xionghui Wang, Chen Chen, Jie Qin, Xuefeng Xiao, Rui Wang, Min Zheng, Xin Pan
[[ICCV'23](#) | [Project](#) | [Paper](#) | [Code](#) | [GitHub Stars 100+](#)]
- **FreeSeg: Unified, Universal and Open-Vocabulary Image Segmentation**
Jie Qin, Jie Wu, Pengxiang Yan, Ming Li, Ren Yuxi, Xuefeng Xiao, Yitong Wang, Rui Wang, Shilei Wen, Xin Pan, Xingang Wang
[[CVPR'23](#) | [Paper](#) | [Code](#) | [GitHub Stars 100+](#)]

AWARDS AND HONORS

- Graduate Artificial Intelligence Initiative (Aii) Doctoral Fellowship, the University of Central Florida. 2023
- ORCGS Doctoral Fellowship, the University of Central Florida. 2023
- Championship of CVPR 2023 Long-form Video Understanding and Generation Academic Challenge. 2023
- Championship of CVPR 2022 Accessibility Vision and Autonomy Academic Challenge. 2022
- China National Inspirational Scholarship. 2017 & 2018