

HAN LIN

CONTACT ✉ hanlincs@cs.unc.edu 🏠 hl-hanlin.github.io 🎓 citation:1166 | h-index:12

INTERESTS Generative models, World models, Multimodal learning, Representation Learning
Theory-grounded algorithms for efficient Transformers

EDUCATION **University of North Carolina at Chapel Hill** 2023 - Exp. 2028
Ph.D. in Computer Science

- MURGe-Lab. Advised by Prof. [Mohit Bansal](#)

Columbia University 2021 - 2023
M.S. in Computer Science (Machine Learning Track) GPA: 4.13/4.33

- DVMM Lab. Advised by Prof. [Shih-Fu Chang](#)
- ROAM Lab. Advised by Prof. [Matei Ciocarlie](#) and Prof. [Shuran Song](#)

Relevant Courses: Learning Theory, Algorithms, Machine Learning, Unsupervised Learning, Bandits & Reinforcement Learning, Causal Inference, Computer Vision, Robotics Learning

Columbia University 2018 - 2020
M.S. in Financial Engineering GPA: 4.04/4.33

- Advised by Prof. [Krzysztof Choromanski](#) from Google Deepmind

Relevant Courses: Optimization, Combinatorial Optimization, Stochastic Models, Stochastic Calculus, Monte Carlo Methods, Statistical Inference, Bayesian Statistics, Graphical Models

Central University of Finance and Economics 2014 - 2018
B.S. in Financial Engineering GPA: 90.43/100

Relevant Courses: Linear Algebra, Mathematical Analysis, Probability, Statistics, Real Analysis, Numerical Methods, Stochastic Process, Differential Equations

PREPRINTS 1. **V-Co: A Closer Look at Visual Representation Alignment via Co-Denoising**
[[arxiv](#)|[project page](#)]

Han Lin, Xichen Pan, Zun Wang, Yue Zhang, Chu Wang, Jaemin Cho, Mohit Bansal

2. **Exploring MLLM-Diffusion Information Transfer with MetaCanvas** [[arxiv](#)
|[project page](#)]

Han Lin, Xichen Pan, Ziqi Huang, Ji Hou, Jialiang Wang, Weifeng Chen, Zecheng He, Felix Junfei-Xu, Junzhe Sun, Zhipeng Fan, Ali Thabet, Jaemin Cho, Mohit Bansal, Chu Wang

3. **AnchorWeave: World-Consistent Video Generation with Retrieved Local Spatial Memories** [[arxiv](#)|[project page](#)]

Zun Wang, **Han Lin**, Jaehong Yoon, Jaemin Cho, Yue Zhang, Mohit Bansal

4. **Error-Driven Scene Editing for 3D Grounding in Large Language Models**
[[arxiv](#)|[github](#)]

Yue Zhang, Zun Wang, **Han Lin**, Jialu Li, Jianing Yang, Yonatan Bitton, Idan Szpektor, Mohit Bansal

5. **Planning with Sketch-Guided Verification for Physics-Aware Video Generation** [[arxiv](#)|[github](#)]

Yidong Huang, Zun Wang, **Han Lin**, Dong-Ki Kim, Shayegan Omidshafiei, Jaehong Yoon, Yue Zhang, Mohit Bansal

6. **Training-free guidance in text-to-video generation via multimodal planning and structured noise initialization** [[arxiv](#)|[github](#)]

- PUBLICATIONS**
- 1. EPiC: Efficient Video Camera Control Learning with Precise Anchor-Video Guidance** [arxiv|project page]
Zun Wang, Jaemin Cho, Jialu Li, **Han Lin**, Jaehong Yoon, Yue Zhang, Mohit Bansal
Forty-Third International Conference on Machine Learning (**ICML**), 2026
 - 2. A Benchmark of Expert-Level Academic Questions to Assess AI Capabilities** [paper|project page]
Center for AI Safety, Scale AI and HLE Contributors Consortium
Nature, 2026
 - 3. Bifrost-1: Bridging Multimodal LLMs and Diffusion Models with Patch-level CLIP Latents** [arxiv|project page]
Han Lin, Jaemin Cho, Amir Zadeh, Chuan Li, Mohit Bansal
Advances in Neural Information Processing Systems (**NeurIPS**), 2025
 - 4. Ctrl-Adapter: An Efficient and Versatile Framework for Adapting Diverse Controls to Any Diffusion Model** [arxiv|project page]
Han Lin*, Jaemin Cho*, Abhay Zala, Mohit Bansal
International Conference on Learning Representations (**ICLR Oral, top 1.82%**), 2025
 - 5. VEDiT: Latent Prediction Architecture For Procedural Video Representation Learning** [arxiv]
Han Lin, Tushar Nagarajan, Nicolas Ballas, Mido Assran, Mojtaba Komeili, Mohit Bansal, Koustuv Sinha
International Conference on Learning Representations (**ICLR**), 2025
 - 6. DreamRunner: Fine-Grained Storytelling Video Generation with Retrieval-Augmented Motion Adaptation** [arxiv|project page]
Zun Wang, Jialu Li, **Han Lin**, Jaehong Yoon, Mohit Bansal
AAAI Conference on Artificial Intelligence (**AAAI**) 2025
 - 7. Fast Tree-Field Integrators: From Low Displacement Rank to Topological Transformers** [arxiv]
Krzysztof Choromanski, Arijit Sehanobish, Somnath Basu Roy Chowdhury, **Han Lin**, Avinava Dubey, Tamas Sarlos, Snigdha Chaturvedi
Advances in Neural Information Processing Systems (**NeurIPS**), 2024
 - 8. VideoDirectorGPT: Consistent Multi-scene Video Generation via LLM-Guided Planning** [arxiv|project page]
Han Lin, Abhay Zala, Jaemin Cho, Mohit Bansal
Conference on Language Modeling (**COLM**), 2024
 - 9. EnvGen: Generating and Adapting Environments via LLMs for Training Embodied Agents** [arxiv|project page]
Abhay Zala*, Jaemin Cho*, **Han Lin**, Jaehong Yoon, Mohit Bansal
Conference on Language Modeling (**COLM**), 2024
 - 10. DiagrammerGPT: Generating Open-Domain, Open-Platform Diagrams via LLM Planning** [arxiv|project page]
Abhay Zala, **Han Lin**, Jaemin Cho, Mohit Bansal
Conference on Language Modeling (**COLM**), 2024
 - 11. Efficient Graph Field Integrators Meet Point Clouds** [arxiv|github]
Krzysztof Choromanski*, Arijit Sehanobish*, **Han Lin***, Yunfan Zhao*, Eli Berger, Alvin Pan, Tetiana Parshakova, Tianyi Zhang, David Watkins, Valerii Likhoshesterov, Somnath Basu Roy Chowdhury, Avinava Dubey, Deepali Jain, Tamas Sarlos, Snigdha

Chaturvedi, Adrian Weller
International Conference on Machine Learning (ICML), 2023

12. **Supervised Masked Knowledge Distillation for Few-Shot Transformers** [arxiv|github]

Han Lin*, Guangxing Han*, Jiawei Ma, Shiyuan Huang, Xudong Lin, Shih-Fu Chang
Conference on Computer Vision and Pattern Recognition (CVPR), 2023

13. **Active Tactile Exploration for 3D Object Recognition** [arxiv|project page]

Jingxi Xu*, **Han Lin***, Shuran Song, Matei Ciocarlie
IEEE International Conference on Robotics and Automation (ICRA), 2023

14. **From Block-Toeplitz Matrices to Differential Equations on Graphs: Towards a General Theory for Scalable Masked Transformers** [arxiv|github]

Krzysztof Choromanski*, **Han Lin***, Haoxian Chen*, Tianyi Zhang, Arijit Sehanobish, Valerii Likhoshesterov, Jack Parker-Holder, Tamas Sarlos, Adrian Weller, Thomas Weingarten

International Conference on Machine Learning (ICML), 2022

15. **Hybrid Random Features** [arxiv|github]

Krzysztof Choromanski*, **Han Lin***, Haoxian Chen*, Yuanzhe Ma*, Arijit Sehanobish*, Deepali Jain, Michael Ryoo, Jake Varley, Andy Zeng, Valerii Likhoshesterov, Dmitry Kalashnikov, Vikas Sindhwani, Adrian Weller

International Conference on Learning Representations (ICLR), 2022

16. **Demystifying Orthogonal Monte Carlo and Beyond** [arxiv|github]

Han Lin*, Haoxian Chen*, Tianyi Zhang, Clement Laroche, Krzysztof Choromanski
Advances in Neural Information Processing Systems (NeurIPS), 2020

* Co-First Authors, Equal Contribution

RESEARCH EXPERIENCE

Google Research Collaboration on 3D LLMs 2025 - now

Advised by: *Yonatan Bitton, Idan Szpektor*

- Error-Driven Scene Editing for 3D Grounding in Large Language Models
- Evaluating Geometric Reasoning of Vision-Language Models

Google Research Collaboration on Efficient Graph Networks 2022 - 2024

Advised by: *Krzysztof Choromanski, Jack Parker-Holder, Somnath Basu Roy Chowdhury, Avinava Dubey, Tamás Sarlós*

- Fast Tree-Field Integrators: From Low Displacement Rank to Topological Transformers (NeurIPS 2024)
- Efficient Graph Field Integrators Meet Point Clouds (ICML 2023)
- From Block-Toeplitz Matrices to Differential Equations on Graphs: Towards a General Theory for Scalable Masked Transformers (ICML 2022)

Google Research Collaboration on Random Feature Algorithms 2019 - 2022

Advised by: *Krzysztof Choromanski, Deepali Jain, Michael S. Ryoo, Jake Varley, Andy Zeng, Dmitry Kalashnikov, Vikas Sindhwani*

- Hybrid Random Features (ICLR 2022)
- Demystifying Orthogonal Monte Carlo and Beyond (NeurIPS 2020)

Meta MSL (Movie Gen Team) 2025.5 - now

Research Scientist Intern

Advised by: *Chu Wang, Ji Hou, Jialiang Wang, Weifeng Chen, Zecheng He, Felix Xu, Junzhe Sun, Zhipeng Fan, Ali Thabet*

- Exploring VLM-Diffusion Information Transfer with MetaCanvas (under review)

Meta FAIR Lab (JEPA Team) 2024.5 - 2024.12

Research Scientist Intern

Advised by: *Koustuv Sinha, Tushar Nagarajan, Nicolas Ballas, Mido Assran, Mojtaba Komeili*

- VEDiT: Latent prediction architecture for procedural video representation learning (ICLR 2025)

MURGe-Lab, UNC 2023 - now

Research Assistant, Advised by Prof. *Mohit Bansal*

- Text-to-video generation, multimodal learning, and LLMs

DVMM Lab, Columbia University 2022 - 2023

Research Assistant, Advised by Prof. *Shih-Fu Chang*

- Supervised masked knowledge distillation for few-shot Transformers (CVPR 2023)

ROAM Lab, Columbia University 2022 - 2023

Research Assistant, Advised by Prof. *Matei Ciocarlie* and Prof. *Shuran Song*

- Active tactile exploration for 3D object recognition (ICRA 2023)

Cornell, Maryland, Max Planck Pre-doctoral Research School 2022

INDUSTRY EXPERIENCE **China Merchant Securities** 2020 - 2021

Option Market Making Quant Trader, Full Time

- Commodity options and futures trading and daily P&L attribution

TEACHING EXPERIENCE COMS 4231 Analysis of Algorithms Fall 2022

COMS 4732 Computer Vision II: Learning Spring 2022

COMS 4721 Machine Learning for Data Science Spring 2022

QMSS 5073 Machine Learning for Social Science Fall 2021

IEOR 4007 Optimization Models & Methods for FE Fall 2019

IEOR 4418 Transportation Analytics & Logistics Spring 2019

AWARDS CVPR 2023 Travel Award Jun 2023

CVPR community

Honorable Mention in Mathematical Context in Modeling Jan 2018

The Consortium for Mathematics and Its Applications (COMAP)

Scholarship for Excellence in Organization and Leadership Apr 2017

Central University of Finance and Economics

2nd Prize for Summer & Winder Social Practice Scholarship Oct 2015

Central University of Finance and Economics

SKILLS Python, C/C++, MATLAB, R, MySQL, L^AT_EX
PyTorch, TensorFlow, Large-Scale TPU/GPU Training

SERVICES **Workshop Organizer:** Any-To-Any Multimodal Learning Workshop, CVPR, 2026

Conference Reviewer: ICML 2022-2026, NeurIPS 2022-2026, ICLR 2024-2025, CVPR 2025-2026, ICCV 2025, ECCV 2026

Conference Volunteer: Robotics: Science and Systems (RSS) 2022

REFERENCE **Mohit Bansal** (mbansal@cs.unc.edu)
Department of Computer Science, UNC Chapel Hill
Relationship: PhD advisor, director of UNC MURGe-Lab.

Krzysztof Choromanski (kchoro@google.com)
Google Deepmind Robotics & Columbia University
Relationship: master's advisor and research collaborator.

Shih-Fu Chang (sc250@columbia.edu)
Columbia University
Relationship: master's research advisor, dean of Columbia Engineering.

Matei Ciocarlie (matei.ciocarlie@columbia.edu)
Columbia University
Relationship: master's research advisor, director of ROAM Lab at Columbia University.

Shuran Song (shuran@stanford.edu)
Stanford University
Relationship: master's research advisor, director of REAL Lab at Stanford University.

Koustuv Sinha (koustuvs@meta.com)
Meta
Relationship: internship advisor, research scientist at Meta FAIR Lab.

Chu Wang (wangchu@meta.com)
Meta
Relationship: internship advisor, research scientist at Meta MSL.