

A selection of Northwestern Engineering papers at ICML 2025, NeurIPS 2025, ICLR 2026, and ICRA 2026

International Conference on Machine Learning (ICML) 2025

[In-Context Deep Learning via Transformer Models](#) (Weimin Wu, Maojiang Su, Jerry Yao-Chieh Hu, Zhao Song, Han Liu), International Conference on Machine Learning (ICML), 2025.

[Fast and Low-Cost Genomic Foundation Models via Outlier Removal](#) (Haozheng Luo, Chenghao Qiu, Maojiang Su, Zhihan Zhou, Zoe Mehta, Guo Ye, Jerry Yao-Chieh Hu, Han Liu), International Conference on Machine Learning (ICML), 2025.

[Nonparametric Modern Hopfield Models](#) (Jerry Yao-Chieh Hu, Bo-Yu Chen, Dennis Wu, Feng Ruan, Han Liu), International Conference on Machine Learning (ICML), 2025.

[In-Context Learning as Conditional Associative Memory Retrieval](#) (Weimin Wu, Teng-Yun Hsiao, Jerry Yao-Chieh Hu, Wenxin Zhang, Han Liu), International Conference on Machine Learning (ICML), 2025

[Sparse-pivot: Dynamic correlation clustering for node insertions](#), Mina Dalirrooyfard, Konstantin Makarychev, Slobodan Mitrović, ICML 2025 (Spotlight)

[You Get What You Give: Reciprocally Fair Federated Learning](#)

Aniket Murhekar, Jiaxin Song, Parnian Shahkar, Bhaskar Ray Chaudhury, Ruta Mehta
ICML 2025.

[The Sample Complexity of Online Strategic Decision Making with Information Asymmetry and Knowledge Transportability](#)

J Hu, R Ai, H Zhong, X Chen, L Wang, Z Wang, Z Yang

[BRiTE: Bootstrapping Reinforced Thinking Process to Enhance Language Model Reasoning](#)

H Zhong, Y Yin, S Zhang, X Xu, Y Liu, Y Zuo, Z Liu, B Liu, S Zheng, H Guo, L Wang, M Hong, Z Wang
International Conference on Machine Learning

[Reward-Augmented Data Enhances Direct Preference Alignment of LLMs](#)

S Zhang, Z Liu, B Liu, Y Zhang, Y Yang, Y Liu, L Chen, T Sun, Z Wang
International Conference on Machine Learning

[An Instrumental Value for Data Production and its Application to Data Pricing](#)

R Ai, B Lyu, Z Wang, Z Yang, H Xu
International Conference on Machine Learning

[Directly Forecasting Belief for Reinforcement Learning with Delays](#)

Qingyuan Wu*, Yuhui Wang*, Simon Sinong Zhan*, Yixuan Wang, Chung-Wei Lin, Chen Lv, Qi Zhu, Jürgen Schmidhuber, Chao Huang
ICML 2025.

Conference on Neural Information Processing Systems (NeurIPS) 2025

[Investigating Hallucinations of Time Series Foundation Models through Signal Subspace](#)

Analysis (Yufeng Zou, Zijian Wang, Diego Klabjan, Han Liu), Advances in Neural Information Processing Systems (NeurIPS), 2025.

[MetaFind: Scene-Aware 3D Asset Retrieval for Coherent Metaverse Scene Generation](#)

(Zhenyu Pan, Yucheng Lu, Han Liu), Advances in Neural Information Processing Systems (NeurIPS), 2025.

[Pareto-optimal energy alignment for designing nature-like antibodies](#)

(Yibo Wen, Chenwei Xu, Jerry Yao-Chieh Hu, Han Liu), Advances in Neural Information Processing Systems (NeurIPS), 2025.

[High-order flow matching: Unified framework and sharp statistical rates](#)

(Maojiang Su*, Jerry Yao-Chieh Hu*, Yi-Chen Lee*, Ning Zhu, Jui-Hui Chung, Shang Wu, Zhao Song, Minshuo Chen, Han Liu), Advances in Neural Information Processing Systems (NeurIPS), 2025.

[Attention Mechanism, Max-Affine Partition, and Universal Approximation](#)

(Hede Liu, Jerry Yao-Chieh Hu, Zhao Song, Han Liu), Advances in Neural Information Processing Systems (NeurIPS), 2025.

[Dynamic Algorithm for Explainable k-medians Clustering under \$\ell_p\$ Norm](#)

Konstantin Makarychev, Ilias Papanikolaou, Liren Shan
NeurIPS 2025 (Spotlight).

[Guarantees for Alternating Least Squares in Overparameterized Tensor Decompositions](#)

Dionysis Arvanitakis, Vaidehi Srinivas, Aravindan Vijayaraghavan
NeurIPS 2025 (Spotlight).

[Pin the Tail on the Model: Blindfolded Repair of User-Flagged Failures in Text-to-Image Services](#)

Gefei Tan, Ali Shahin Shamsabadi, Ellen Kolesnikova, Hamed Haddadi, Xiao Wang
NeurIPS 2025

[RHYTHM: Reasoning with Hierarchical Temporal Tokenization for Human Mobility](#)

Haoyu He*, Haozheng Luo*, Yan Chen, Qi Ryan Wang. In the Thirty-Ninth Annual Conference on Neural Information Processing Systems (NeurIPS 2025).

[Towards a Translative Model of Sperm Whale Vocalization](#)

Paradise, Orr and Chen, Liangyuan and Muralikrishnan, Pranav and Flores, Hugo and Pardo, Bryan and Diamant, Roe and Gruber, David and Gero, Shane and Goldwasser, Shafi
Advances in Neural Information Processing Systems (NeurIPS 2025)

[Learning from A Single Markovian Trajectory: Optimality and Variance Reduction](#). Zhenyu Sun, Ermin Wei NeurIPS 2025

[On the Existence and Complexity of Core Stable Data Exchanges](#). Jiaxin Song, Pooja Kulkarni, Parnian Shahkar Bhaskar Ray Cahudhury NeurIPS 2025

[MAESTRO: Adaptive Sparse Attention and Robust Learning for Multimodal Dynamic Time Series](#). Payal Mohapatra, Yueyuan Sui, Akash Pandey, Stephen Xia, Qi Zhu. NeurIPS 2025. (Spotlight).

International Conference on Learning Representations (ICLR) 2026

[In-Context Algorithm Emulation in Fixed-Weight Transformers](#) (Jerry Yao-Chieh Hu, Hude Liu, Jennifer Yuntong Zhang, Han Liu), International Conference on Learning Representations (ICLR), 2026.

[Sci2Pol: Evaluating and Fine-tuning LLMs on Scientific-to-Policy Brief Generation](#) (Weimin Wu, Alexander C Furnas, Eddie Yang, Gefei Liu, Akhil Pandey Akella, Xuefeng Song, Dashun Wang, Han Liu), International Conference on Learning Representations (ICLR), 2026.

[MetaSpatial: Reinforcing 3D Spatial Reasoning in VLMs for the Metaverse](#) (Zhenyu Pan, Han Liu), International Conference on Learning Representations (ICLR), 2026.

[Understanding Transformers for Time Series Forecasting: A Case Study on MOIRAI](#) (Dennis Wu, Yihan He, Yuan Cao, Jianqing Fan, Han Liu), International Conference on Learning Representations (ICLR), 2026.

[Theory of Space: Can Foundation Models Construct Spatial Beliefs Through Active Perception?](#) Pingyue Zhang*, Zihan Huang*, Yue Wang *, Jieyu Zhang*, Letian Xue, Zihan Wang, Qineng Wang, Keshigeyan Chandrasegaran, Ruohan Zhang, Yejin Choi, Ranjay Krishna, Jiajun Wu, Li Fei-Fei, Manling Li
ICLR 2026

[ENACT: Evaluating Embodied Cognition with World Modeling of Egocentric Interaction](#) Qineng Wang*, Wenlong Huang*, Yu Zhou, Hang Yin, Tianwei Bao, Jianwen Lyu, Weiyu Liu, Ruohan Zhang, Jiajun Wu, Li Fei-Fei, Manling Li
ICLR 2026

[Spatial Mental Modeling from Limited Views](#)

Qineng Wang*, Baiqiao Yin*, Pingyue Zhang, Jianshu Zhang, Kangrui Wang, Zihan Wang, Jieyu Zhang, Keshigeyan Chandrasegaran, Han Liu, Ranjay Krishna, Saining Xie, Jiajun Wu+, Li Fei-Fei+, Manling Li+
ICLR 2026

[ROSETTA: Constructing Code-Based Reward from Unconstrained Language Preference](#)

Sanjana Srivastava*, Kangrui Wang*, Yung-Chieh Chan*, Tianyuan Dai, Manling Li, Ruohan Zhang, Mengdi Xu, Jiajun Wu, Li Fei-Fei
ICLR 2026

[Activation Steering for LLM Alignment via a Unified ODE-Based Framework](#)

Hongjue Zhao*, Haosen Sun*, Jiangtao Kong, Xiaochang Li, Qineng Wang, Liwei Jiang, Qi Zhu, Tarek F. Abdelzaher, Yejin Choi, Manling Li+, Huajie Shao+ (equal advising)
ICLR 2026

[Your Language Model Secretly Contains Personality Subnetworks](#)

Ruimeng Ye, Zihan Wang, Zinan Ling, Yang Xiao, Manling Li, Xiaolong Ma, Bo Hui
ICLR 2026

[Weak-to-Strong Generalization with Failure Trajectories](#)

Ruimeng Ye, Zihan Wang, Yang Xiao, Zinan Ling, Manling Li, Bo Hui
ICLR 2026

[MoMaGen: Generating Demonstrations under Soft and Hard Constraints for Multi-Step Bimanual Mobile Manipulation](#)

Chengshu Li*, Mengdi Xu*, Arpit Bahety*, Hang Yin*, Yunfan Jiang, Huang Huang, Josiah Wong, Sujay Garlanka, Cem Gokmen, Ruohan Zhang, Weiyu Liu, Jiajun Wu, Roberto Martín-Martín, Li Fei-Fei.
ICLR 2026

[Explaining and Improving Information Complementarities in Multi-Agent Decision-making](#). Ziyang Guo, Yifan Wu, Jason Hartline, Jessica Hullman. ICLR 2026

[Homeostatic Adaptation of Optimal Population Codes under Metabolic Stress](#). Yi-Chun Hung, Gregory W. Schwartz, Emily A. Cooper, Emma Alexander. ICLR 2026

[FROST: Filtering Reasoning Outliers with Attention for Efficient Reasoning](#). Haozheng Luo, Zhuolin Jiang, Md Zahid Hasan, Yan Chen, Soumalya Sarkar. In the Fourteenth International Conference on Learning Representations (ICLR 2026).

[Belief-Based Offline Reinforcement Learning for Delay-Robust Policy Optimization](#). Simon Sinong Zhan*, Qingyuan Wu*, Philip Wang, Frank Yang, Xiangyu Shi, Chao Huang, Qi Zhu. ICLR 2026.

[TimeSliver: Symbolic-Linear Decomposition for Explainable Time Series Classification](#). Akash Pandey*, Payal Mohapatra*, Wei Chen, Qi Zhu+, Sinan Keten+. ICLR 2026.

[Local Linear Attention: An Optimal Interpolation of Linear and Softmax Attention For Test-Time Regression](#)

Y Zuo, Y Yin, Z Zeng, A Li, B Zhu, Z Wang
International Conference on Learning Representations

[Learning to Reason as Action Abstractions with Scalable Mid-Training RL](#)

S Zhang, D Yu, Y Feng, B Jin, Z Wang, J Peebles, Z Wang
International Conference on Learning Representations

[Beyond Markovian: Reflective Exploration via Bayes-Adaptive RL for LLM Reasoning](#)

S Zhang, Y Wang, Y Liu, T Liu, P Grabowski, E Ie, Z Wang, Y Li
International Conference on Learning Representations

[Are Transformers Able to Reason by Connecting Separated Knowledge in Training Data?](#)

Y Yin, Z Wang
International Conference on Learning Representations

International Conference on Robotics and Automation (ICRA) 2026

[Dream2Flow: Bridging Video Generation and Open-World Manipulation with 3D Object Flow](#)

Karthik Dharmarajan, Wenlong Huang, Jiajun Wu, Li Fei-Fei, Ruohan Zhang†
ICRA, 2026

[IMPASTO: Integrating Model-Based Planning with Learned Dynamics Models for Robotic Oil Painting Reproduction](#)

Yingke Wang, Hao Li, Yifeng Zhu, Koven Yu, Ken Goldberg, Li Fei-Fei, Jiajun Wu, Yunzhu Li, Ruohan Zhang
ICRA, 2026