

Kangwook Lee, PhD

kangwooklee@krafton.com

- Personal webpage: <http://kangwooklee.com>
- Google Scholar: <https://scholar.google.com/citations?hl=en&user=sCEl8r-n5VEC&sortby=pubdate>

INDUSTRY EXPERIENCE

KRAFTON CAIO	<i>2026.02 – present</i>
Ludo Robotics CTO	<i>2026.02 – present</i>
KRAFTON Head of Deep Learning R&D	<i>2022.04 – 2026.01</i>

ACADEMIC APPOINTMENTS

University of Wisconsin-Madison Dept. of ECE and CS (by courtesy) Associate Professor (with tenure)	<i>2025.07 – 2026.01</i>
University of Wisconsin-Madison Dept. of ECE and CS (by courtesy) Assistant Professor	<i>2019.08 – 2025.06</i>
KAIST Research Assistant Professor Mentor: Prof. Changho Suh <i>In lieu of military service (June 2016 – June 2019)</i>	<i>2018.10 – 2019.06</i>
KAIST Postdoctoral Fellow Mentor: Prof. Changho Suh <i>In lieu of military service (June 2016 – June 2019)</i>	<i>2016.06 – 2018.09</i>
UC Berkeley Graduate Student Researcher	<i>2010.08 – 2016.05</i>

EDUCATION

University of California, Berkeley Ph.D. in Electrical Engineering and Computer Sciences Advisor: Prof. Kannan Ramchandran	<i>2016</i>
University of California, Berkeley M.S. in Electrical Engineering and Computer Sciences Advisor: Prof. Kannan Ramchandran	<i>2012</i>
KAIST B.S. in Electrical Engineering Advisor: Prof. Sae-Young Chung and Prof. Yung Yi Highest GPA (4.19/4.30) among all 800+ graduates across all departments, 2010	<i>2010</i>

RESEARCH INTERESTS

Virtual and real agents — theoretical/empirical analysis to understand how/why they work and improve them

AWARDS AND HONORS (SELECTED)

- Outstanding Paper Award, The 5th Wordplay Workshop @ EMNLP 2025, 2025.
- Fusion Fund Distinguished Scholar Network, Inaugural Member, 2025.
- **NSF CAREER Award**, 2024.
- Amazon Research Awards, 2024.
- Best Paper Award, The Federated Learning Systems (FLSys) Workshop @ MLSys 2023, 2023.
- ECE Grainger Faculty Scholarship Award, UW Madison ECE, 2022.
- **Young Investigator Grants Award**, The Korean-American Scientists and Engineers Association, 2022
- The Joint Communications Society/Information Theory Society Paper Award, IEEE, 2020.
- **The Outstanding Graduate Student Instructor Award**, UC Berkeley, 2016.
- Best Paper Award Finalist, IEEE MASCOTS 2013, 2013.
- **KFAS PhD Fellowship (Full funding for five years)**, 2010 - 2015.
- Highest GPA (4.19/4.30) among all 800+ graduates across all departments, KAIST, 2010.
- **Korea Talent Award (Presidential Award)**, KOFAC, 2009.
- Dugald C. Jackson Assistant Professorship, UW Madison, 2024.
- Brain Korea 21 Research Professor Fellowship, 2018.
- GCORE Post-doctoral Fellowship, 2017.
- Brain Korea 21 Post-doctoral Fellowship, 2016.

PUBLICATIONS

Preprints

- [Arxiv](#) Chungpa Lee, Jy-yong Sohn, and **Kangwook Lee**, “Fine-Tuning Without Forgetting In-Context Learning: A Theoretical Analysis of Linear Attention Models,” 2026.
- [Arxiv](#) Ziqian Lin, Shubham Kumar Bharti, and **Kangwook Lee**, “In-Context Learning with Hypothesis-Class Guidance,” 2025.
- [Arxiv](#) Erica Zhang, Ryunosuke Goto, Naomi Sagan, Jurik Mutter, Nick Phillips, Ash Alizadeh, **Kangwook Lee**, Jose Blanchet, Mert Pilanci, and Robert Tibshirani, “LLM-Lasso: A Robust Framework for Domain-Informed Feature Selection and Regularization,” 2025.
- [Arxiv](#) Yuchen Zeng, Shuibai Zhang, Wonjun Kang, Shutong Wu, Lynnix Zou, Ying Fan, Heeju Kim, Ziqian Lin, Jungtaek Kim, Hyung Il Koo, Dimitris Papailiopoulos, and **Kangwook Lee**, “ReJump: A Tree-Jump Representation for Analyzing and Improving LLM Reasoning,” 2025. [[Summary](#) · [Github](#)]
- [Arxiv](#) Chungpa Lee, Thomas Zeng, Jongwon Jeong, Jy-yong Sohn, and **Kangwook Lee**, “How to Correctly Report LLM-as-a-Judge Evaluations,” 2025. [[Summary](#) · [Github](#)]
- [Arxiv](#) Ozgur Guldogan, Jackson Kunde, **Kangwook Lee**, and Ramtin Pedarsani, “Multi-Bin Batching for Increasing LLM Inference Throughput,” 2024.
- [Arxiv](#) Liu Yang, Jifan Zhang, Joseph Shenouda, Dimitris Papailiopoulos, **Kangwook Lee**, and Robert D. Nowak, “PathProx: A Proximal Gradient Algorithm for Weight Decay Regularized Deep Neural Networks,” 2023.

Peer-reviewed Conference Papers

- ICLR'26 Kevin Galim, Ethan Ewer, Wonjun Kang, Minjae Lee, Hyung Il Koo, and **Kangwook Lee**, “Draft-based Approximate Inference for LLMs,” *The Fourteenth International Conference on Learning Representations (ICLR 2026)*, Rio de Janeiro, Brazil, Apr., 2026. [[Summary](#) · [Github](#)]
- ICLR'26 Dongmin Park, Minkyu Kim, Beongjun Choi, Junhyuck Kim, Keon Lee, Jonghyun Lee, Inkyu Park, Byeong-Uk Lee, Jaeyoung Hwang, Jaewoo Ahn, Ameya Sunil Mahabaleshwarkar, Bilal Kartal, Pritam Biswas, Yoshi Suhara, **Kangwook Lee**, and Jaewoong Cho, “Orak: A Foundational Benchmark for Training and Evaluating LLM Agents on Diverse Video Games,” *ICLR 2026*, Rio de Janeiro, Brazil, Apr., 2026. ▷ **Outstanding Paper Award @ EMNLP 2025 Wordplay Workshop** [[Summary](#) · [Github](#)]
- ICLR'26 Wonjun Kang, Kevin Galim, Seunghyuk Oh, Minjae Lee, Yuchen Zeng, Shuibai Zhang, Coleman Hooper, Yuezhou Hu, Hyung Il Koo, Nam Ik Cho, and **Kangwook Lee**, “ParallelBench: Understanding the Trade-offs of Parallel Decoding in Diffusion LLMs,” *The Fourteenth International Conference on Learning Representations (ICLR 2026)*, Rio de Janeiro, Brazil, Apr., 2026. [[Summary](#) · [Github](#)]
- EACL'26 Minjae Lee, Wonjun Kang, Byeongkeun Ahn, Christian Classen, Kevin Galim, Seunghyuk Oh, Minghao Yan, Hyung Il Koo, and **Kangwook Lee**, “TABED: Test-Time Adaptive Ensemble Drafting for Robust Speculative Decoding in LVLMs” *Findings of The 31st Conference of the European Chapter of the Association for Computational Linguistics (EACL 2026)*, Dubrovnik, Croatia, Mar., 2026.
- ICML'25 Zheyang Xiong, Ziyang Cai, John Cooper, Albert Ge, Vasilis Papageorgiou, Zack Sifakis, Angeliki Giannou, **spotlight** Ziqian Lin, Liu Yang, Saurabh Agarwal, Grigorios Chrysos, Samet Oymak, **Kangwook Lee**, and Dimitris Papailiopoulos, “Everything Everywhere All at Once: LLMs can In-Context Learn Multiple Tasks in Superposition,” *The Forty-Second International Conference on Machine Learning (ICML 2025)*, Vancouver, Canada, July, 2025.
- ICLR'25 Zheyang Xiong, Vasilis Papageorgiou, **Kangwook Lee**, and Dimitris Papailiopoulos, “From Artificial Needles to Real Haystacks: Improving Retrieval Capabilities in LLMs by Finetuning on Synthetic Data,” *Thirteenth International Conference on Learning Representations (ICLR 2025)*, Singapore, April, 2025. [[Summary](#) · [Github](#)]
- NeurIPS'24 Ying Fan, Yilun Du, Kannan Ramchandran, and **Kangwook Lee**, “Looped Transformers for Length Generalization,” *Thirteenth International Conference on Learning Representations (ICLR 2025)*, Singapore, April, 2025. [[Summary](#) · [Github](#)]
- ICML'25 Kevin Galim, Wonjun Kang, Yuchen Zeng, Hyung Il Koo, and **Kangwook Lee**, “Parameter-Efficient Fine-Tuning of State Space Models,” *ICML 2025*, Vancouver, Canada, Jul., 2025.
- ICLR'25 Dongmin Park, Sebin Kim, Taehong Moon, Minkyu Kim, **Kangwook Lee**, and Jaewoong Cho, “Rare-to-Frequent: Unlocking Compositional Generation Power of Diffusion Models on Rare Concepts with LLM Guidance,” *Thirteenth International Conference on Learning Representations (ICLR 2025)*, Singapore, April, 2025. [[Summary](#) · [Github](#)]
- ICLRW'25 Nayoung Lee, Ziyang Cai, Avi Schwarzschild, **Kangwook Lee**, and Dimitris Papailiopoulos, “Self-Improving Transformers Overcome Easy-to-Hard and Length Generalization Challenges,” *The Forty-Second International Conference on Machine Learning (ICML 2025)*, Vancouver, Canada, July, 2025.
- COLM'25 Liu Yang, Ziqian Lin, **Kangwook Lee**, Dimitris Papailiopoulos, and Robert Nowak, “Task Vectors in In-Context Learning: Emergence, Formation, and Benefit,” *Conference on Language Modeling (COLM 2025)*, Montreal, Canada, October, 2025.
- ICML'25 Thomas Zeng, Shuibai Zhang, Shutong Wu, Christian Classen, Daewon Chae, Ethan Ewer, Minjae Lee, Heeju **oral** Kim, Wonjun Kang, Jackson Kunde, Ying Fan, Jungtaek Kim, Hyung Il Koo, Kannan Ramchandran, Dimitris Papailiopoulos, and **Kangwook Lee**, “VersaPRM: Multi-Domain Process Reward Model via Synthetic Reasoning Data,” *The Forty-Second International Conference on Machine Learning (ICML 2025)*, Vancouver, Canada, July, 2025. [[Summary](#) · [Github](#) · [HuggingFace](#)]

- COLM'24 Yuchen Zeng*, Wonjun Kang*, Yicong Chen, Hyung Il Koo, and **Kangwook Lee**, “Can MLLMs Perform Text-to-Image In-Context Learning?,” *Conference on Language Modeling (COLM 2024)*, Philadelphia, PA, USA, October, 2024. [[Summary](#) · [Github](#)]
- ICLRW'24 Jongho Park, Jaeseung Park, Zheyang Xiong, Nayoung Lee, Jaewoong Cho, Samet Oymak, **Kangwook Lee**, and Dimitris Papailiopoulos, “Can Mamba Learn How To Learn? A Comparative Study on In-Context Learning Tasks,” *The Forty-first International Conference on Machine Learning (ICML'24)*, Vienna, Austria, July, 2024. [[Github](#)]
- ICLRW'24 Ziqian Lin and **Kangwook Lee**, “Dual Operating Modes of In-Context Learning,” *The Forty-first International Conference on Machine Learning (ICML'24)*, Vienna, Austria, July, 2024. [[Summary](#) · [Github](#)]
- NeurIPS'23 Sehyun Kwon, Jaeseung Park, Minkyu Kim, Jaewoong Cho, Ernest K. Ryu, and **Kangwook Lee**, “Image Clustering Conditioned on Text Criteria,” *Twelfth International Conference on Learning Representations (ICLR'24)*, Vienna, Austria, 2024. [[Summary](#) · [Github](#)]
- ICMLW'23 Liu Yang, **Kangwook Lee**, Robert D Nowak, and Dimitris Papailiopoulos, “Looped Transformers are Better at Learning Learning Algorithms,” *Twelfth International Conference on Learning Representations (ICLR'24)*, Vienna, Austria, 2024. [[Summary](#) · [Github](#)]
- UAI'24 Jy-yong Sohn, Dohyun Kwon, Seoyeon An, and **Kangwook Lee**, “Memorization Capacity for Additive Fine-Tuning with Small ReLU Networks,” *Fortieth Conference on Uncertainty in Artificial Intelligence (UAI'24)*, Barcelona, Spain, July, 2024.
- ICMLW'23 Nayoung Lee, Kartik Sreenivasan, Jason Lee, **Kangwook Lee**, and Dimitris Papailiopoulos, “Teaching Arithmetic to Small Transformers,” *Twelfth International Conference on Learning Representations (ICLR'24)*, Vienna, Austria, 2024. [[Summary](#) · [Github](#)]
- NeurIPS'23 Yuchen Zeng and **Kangwook Lee**, “The Expressive Power of Low-Rank Adaptation,” *Twelfth International Conference on Learning Representations (ICLR'24)*, Vienna, Austria, 2024. [[Summary](#) · [Github](#)]
- NeurIPS'23 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, **Kangwook Lee**, and Kimin Lee, “DPOK: Reinforcement Learning for Fine-tuning Text-to-Image Diffusion Models,” *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS'23)*, New Orleans, LA, Dec., 2023.
- ICLR'23 Ozgur Guldogan*, Yuchen Zeng*, Jy-yong Sohn, Ramtin Pedarsani, and **Kangwook Lee**, “Equal Improvability: A New Fairness Notion Considering the Long-Term Impact,” *The Eleventh International Conference on Learning Representations (ICLR'23)*, Kigali Rwanda, May, 2023.
- MLSys'23 Shenghong Dai, Bryce Yicong Chen, Jy-yong Sohn, S M Iftexharul Alam, Ravikumar Balakrishnan, Suman Banerjee, Nageen Himayat, **Kangwook Lee**, “FedGP: Buffer-based Gradient Projection for Continual Federated Learning,” *MLSys-FLSys 2023*, 2023. ▷ **Best Paper Award**
- ISIT'23 Yuchen Zeng, Hongxu Chen, and **Kangwook Lee**, “Federated Learning with Local Fairness Constraints” *The 2023 IEEE International Symposium on Information Theory (ISIT'23)*, Taipei, Taiwan, June, 2023.
- ICML'23 Yuji Roh, **Kangwook Lee**, Steven Euijong Whang, and Changho Suh, “Improving Fair Training under Correlation Shifts,” *The Fortieth International Conference on Machine Learning (ICML'23)*, Honolulu, HI, July, 2023.
- ICLRW'23 Angeliki Giannou*, Shashank Rajput*, Jy-yong Sohn, **Kangwook Lee**, Jason D. Lee, and Dimitris Papailiopoulos, “Looped Transformers as Programmable Computers,” *The Fortieth International Conference on Machine Learning (ICML'23)*, Honolulu, HI, July, 2023.
- CCNC'23 Shenghong Dai, S M Iftexharul Alam, Ravikumar Balakrishnan, **Kangwook Lee**, Suman Banerjee, and Nageen Himayat, “Online Federated Learning based Object Detection across Autonomous Vehicles in a Virtual World” *IEEE CCNC 2023, Demo Track*, Las Vegas, NV, Jan., 2023.

- ICML'23 Ying Fan and **Kangwook Lee**, "Optimizing DDPM Sampling with Shortcut Fine-Tuning," *The Fortieth International Conference on Machine Learning (ICML'23)*, Honolulu, HI, July, 2023.
- ACL'23 Gibbeum Lee, Volker Hartmann, Jongho Park, Dimitris Papailiopoulos, and **Kangwook Lee**, "Prompted LLMs as Chatbot Modules for Long Open-domain Conversation," *Findings of The 61st Annual Meeting of the Association for Computational Linguistics (Findings of ACL'23)*, Toronto, Canada, July, 2023.
- ISIT'22 Changhun Jo, Jy-yong Sohn, and **Kangwook Lee**, "Breaking Fair Binary Classification with Optimal Flipping Attacks" *The 2022 IEEE International Symposium on Information Theory (ISIT'22)*, Espoo, Finland, June, 2022.
- MLSys'22 Shenghong Dai, **Kangwook Lee**, and Suman Banerjee, "Dynamic Decentralized Federated Learning" *MLSys-CrossFL 2022*, 2022.
- ICML'22 Jy-yong Sohn, Liang Shang, Hongxu Chen, Jaekyun Moon, Dimitris Papailiopoulos, and **Kangwook Lee**, "Gen-Label: Mixup Relabeling using Generative Models," *ICML 2022*, 2022.
- MLSys'22 Yuchen Zeng, Hongxu Chen, and **Kangwook Lee**, "Improving Fairness via Federated Learning" *MLSys-CrossFL 2022*, 2022.
- NeurIPS'22 Tuan Dinh*, Yuchen Zeng*, Ruisu Zhang, Ziqian Lin, Michael Gira, Shashank Rajput, Jy-yong Sohn, Dimitris Papailiopoulos, and **Kangwook Lee**, "LIFT: Language-Interfaced FineTuning for Non-Language Machine Learning Tasks," *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS'22)*, New Orleans, LA, Nov., 2022.
- ICLR'22 Shashank Rajput, **Kangwook Lee**, and Dimitris Papailiopoulos, "Permutation-Based SGD: Is Random Optimal?," *The Tenth International Conference on Learning Representations (ICLR'22)*, Virtual Only, Apr., 2022.
- NeurIPS'22 Kartik Sreenivasan, Jy-yong Sohn, Liu Yang, Matthew Grinde, Aliot Nagle, Hongyi Wang, **Kangwook Lee**, and Dimitris Papailiopoulos, "Rare Gems: Finding Lottery Tickets at Initialization," *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS'22)*, New Orleans, LA, Nov., 2022.
- NeurIPS'22 Dohyun Kwon, Ying Fan, and **Kangwook Lee**, "Score-based generative modeling secretly minimizes the Wasserstein distance," *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS'22)*, New Orleans, LA, Nov., 2022.
- EMNLP'22 Tuan Dinh, Jy-yong Sohn, Shashank Rajput, Tim Ossowski, Yifei Ming, Junjie Hu, Dimitris Papailiopoulos, and **Kangwook Lee**, "Utilizing Language-Image Pretraining for Efficient and Robust Bilingual Word Alignment," *Findings of The 2022 Conference on Empirical Methods in Natural Language Processing (Findings of EMNLP 2022)*, Abu Dhabi, UAE, Dec., 2022.
- MLSys'21 Saurabh Agarwal, Hongyi Wang, **Kangwook Lee**, Shivaram Venkataraman, and Dimitris Papailiopoulos, "Accordion: Adaptive Gradient Communication via Critical Learning Regime Identification," *Fourth Conference on Machine Learning and Systems (MLSys'21)*, Virtual Only, April, 2021.
- ICML'21 Tuan Dinh and **Kangwook Lee**, "Coded-InvNet for Resilient Prediction Serving Systems," *The Thirty-eighth long oral International Conference on Machine Learning (ICML'21)*, Virtual Only, July, 2021.
- ICML'21 Changhun Jo and **Kangwook Lee**, "Discrete-Valued Latent Preference Matrix Estimation with Graph Side Information," *The Thirty-eighth International Conference on Machine Learning (ICML'21)*, Virtual Only, July, 2021.
- ICLR'21 Yuji Roh, **Kangwook Lee**, Steven Euijong Whang, and Changho Suh, "FairBatch: Batch Selection for Model Fairness," *The Ninth International Conference on Learning Representations (ICLR'21)*, Virtual Only, May, 2021.
- ICMLW'21 Jinwoo Jeon, Jaechang Kim, **Kangwook Lee**, Sewoong Oh, and Jungseul Ok, "Gradient Inversion with Generative Image Prior," *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS'21)*, Virtual Only, Dec., 2021.

- MVA'21 Hoon Kim*, **Kangwook Lee***, and Changho Suh, "Predicting Vehicle Collisions using Data Collected from Video Games" *Springer Machine Vision and Applications 2021*, 2021.
- NeurIPS'21 Yuji Roh, **Kangwook Lee**, Steven Euijong Whang, and Changho Suh, "Sample Selection for Fair and Robust Training," *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS'21)*, Virtual Only, Dec., 2021.
- NeurIPS'20 Hongyi Wang, Kartik Sreenivasan, Shashank Rajpu, Harit Vishwakarma, Saurabh Agarwal, Jy-yong Sohn, **Kangwook Lee**, and Dimitris Papailiopoulos, "Attack of the Tails: Yes, You Really Can Backdoor Federated Learning," *The Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS'20)*, Virtual Only, Dec., 2020.
- ICML'20 Yuji Roh, **Kangwook Lee**, Steven Euijong Whang, and Changho Suh, "FR-Train: A mutual information-based approach to fair and robust training," *The Thirty-seventh International Conference on Machine Learning (ICML'20)*, July, 2020.
- ECML'20 **Kangwook Lee**, Changho Suh, and Kannan Ramchandran, "Reprogramming GANs via Input Noise Design," *The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD'20)*, Ghent, Belgium, Sept., 2020.
- AAAI'19 Hoon Kim*, **Kangwook Lee***, Gyeongjo Hwang, and Changho Suh, "Crash to Not Crash: Learn to Identify Dangerous Vehicles using a Simulator," *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI'19)*, 2019.
- ISIT'19 **Kangwook Lee**, Hoon Kim, Kyungmin Lee, Changho Suh, and Kannan Ramchandran, "Synthesizing Differentially Private Datasets using Random Mixing," *The 2019 IEEE International Symposium on Information Theory (ISIT'19)*, Paris, France, July, 2019.
- NeurIPS'18 Kwangjun Ahn, **Kangwook Lee**, Hyunseung Cha, and Changho Suh, "Binary Rating Estimation with Graph Side Information," *The Thirty-second Conference on Neural Information Processing Systems (NeurIPS'18)*, 2018.
- ISIT'18 Hyeyeong Park, **Kangwook Lee**, Jy-yong Sohn, Changho Suh, and Jaekyun Moon, "Hierarchical Coding for Distributed Computing" *The 2018 IEEE International Symposium on Information Theory (ISIT'18)*, Vail, CO, June, 2018.
- Allerton'18 Jisang Yoon, **Kangwook Lee**, and Changho Suh, "On the Joint Recovery of Community Structure and Community Features" *The 55th Annual Allerton Conference on Communication, Control, and Computing*, 2018.
- SysML'18 **Kangwook Lee**, Kyungmin Lee, Hoon Kim, Changho Suh, and Kannan Ramchandran, "SGD on Random Mixtures: Private Machine Learning under Data-breach Threats" *SysML 2018*, 2018.
- ICLR'18 **Kangwook Lee***, Hoon Kim*, and Changho Suh, "Simulated+Unsupervised Learning With Adaptive Data Generation and Bidirectional Mappings," *The Sixth International Conference on Learning Representations (ICLR'18)*, 2018.
- ISIT'18 Tavor Baharav, **Kangwook Lee**, Orhan Ocal, and Kannan Ramchandran, "Straggler-proofing massive-scale distributed matrix multiplication with d-dimensional product codes" *The 2018 IEEE International Symposium on Information Theory (ISIT'18)*, Vail, CO, June, 2018.
- SysML'18 Jichang Chung, **Kangwook Lee**, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran, "Uber-Shuffle: Communication-efficient Data Shuffling for SGD via Coding Theory" *SysML 2018*, 2018.
- ICC'17 Kabir Chandrasekher, **Kangwook Lee**, Peter Kairouz, Ramtin Pedarsani, and Kannan Ramchandran, "Asynchronous and Noncoherent Neighbor Discovery for the IoT Using Sparse-Graph Codes" *IEEE International Conference on Communications (ICC'17)*, Paris, France, May, 2017.
- ISIT'17 **Kangwook Lee**, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran, "Coded Computation for Multicore Setups" *The 2017 IEEE International Symposium on Information Theory (ISIT'17)*, Aachen, Germany, July, 2017.

- ISIT'17 **Kangwook Lee**, Changho Suh, and Kannan Ramchandran, "High-Dimensional Coded Matrix Multiplication" *The 2017 IEEE International Symposium on Information Theory (ISIT'17)*, Aachen, Germany, July, 2017.
- ISIT'17 Kwangjun Ahn, **Kangwook Lee**, and Changho Suh, "Information-theoretic Limits of Subspace Clustering" *The 2017 IEEE International Symposium on Information Theory (ISIT'17)*, Aachen, Germany, July, 2017.
- Allerton'17 Geewon Suh, **Kangwook Lee**, and Changho Suh, "Matrix Sparsification for Coded Matrix Multiplication" *The 54th Annual Allerton Conference on Communication, Control, and Computing*, 2017.
- Allerton'16 Kwangjun Ahn, **Kangwook Lee**, and Changho Suh, "Community Recovery in Hypergraphs" *Allerton 2016*, 2016.
- ISIT'16 **Kangwook Lee**, Ramtin Pedarsani, and Kannan Ramchandran, "SAFFRON: Sparse-Graph Code Framework for Group Testing" *IEEE ISIT 2016*, Barcelona, Spain, Jul., 2016.
- ISIT'16 **Kangwook Lee**, Maximilian Lam, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran, "Speeding Up Distributed Machine Learning using Codes," *The 2016 IEEE International Symposium on Information Theory (ISIT 2016)*, Barcelona, Spain, July, 2016.
- ISIT'15 Ramtin Pedarsani, **Kangwook Lee**, and Kannan Ramchandran, "Capacity-Approaching PhaseCode for Low-Complexity Compressive Phase Retrieval" *The 2015 IEEE International Symposium on Information Theory (ISIT'15)*, Hong Kong, June, 2015.
- ISIT'15 Dong Yin, **Kangwook Lee**, and Kannan Ramchandran, "Fast and Robust Compressive Phase Retrieval with Sparse-Graph Codes" *The 2015 IEEE International Symposium on Information Theory (ISIT'15)*, Hong Kong, June, 2015.
- Allerton'15 **Kangwook Lee**, Ramtin Pedarsani, and Kannan Ramchandran, "On Scheduling Redundant Requests with Cancellation Overheads" *Allerton 2015*, 2015.
- Allerton'15 Ramtin Pedarsani, **Kangwook Lee**, and Kannan Ramchandran, "Sparse Covariance Estimation Based on Sparse-Graph Codes" *The 53th Annual Allerton Conference on Communication, Control, and Computing*, 2015.
- Allerton'14 Ramtin Pedarsani, **Kangwook Lee**, and Kannan Ramchandran, "PhaseCode: Fast and Efficient Compressive Phase Retrieval based on Sparse-Graph-Codes" *Allerton 2014*, 2014.
- ISIT'14 Nihar B. Shah, **Kangwook Lee**, and Kannan Ramchandran, "The MDS Queue: Analysing the Latency Performance of Erasure Codes" *IEEE ISIT 2014*, 2014.
- MASCOTS'13 **Kangwook Lee**, Lisa Yan, Abhay Parekh, and Kannan Ramchandran, "A VoD System for Massively Scaled, Heterogeneous Environments: Design and Implementation," *The IEEE 21st International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'13)*, 2013. **▷ Best Paper Award finalist**
- Allerton'13 Nihar B. Shah, **Kangwook Lee**, and Kannan Ramchandran, "When Do Redundant Requests Reduce Latency?" *Allerton 2013*, 2013.
- Allerton'12 **Kangwook Lee**, Hao Zhang, Ziyu Shao, Minghua Chen, Abhay Parekh, and Kannan Ramchandran, "An Optimized Distributed Video-on-Demand Streaming System: Theory and Design" *The 50th Allerton Conference on Communication, Control and Computing*, 2012.
- Asilomar'11 Sameer Pawar, Salim Rouayheb, Hao Zhang, **Kangwook Lee**, and Kannan Ramchandran, "Codes for a Distributed Caching based Video-On-Demand System" *Asilomar Conference on Signals, Systems, and Computers*, 2011.
- INFOCOM'11 Bruno Nardelli, Jinsung Lee, **Kangwook Lee**, Yung Yi, Song Chong, Edward Knightly, and Mung Chiang, "Experiment evaluation of optimal CSMA" *The 30th IEEE International Conference on Computer Communications (INFOCOM'11)*, 2011.

Peer-reviewed Workshop Papers (Machine Learning)

- ICLRW'26 Jongwon Jeong, Jungtaek Kim, and **Kangwook Lee**, "TAPE: Tool-Guided Adaptive Planning and Constrained Execution in Language Model Agents," ICLR 2026 Workshop on Agentic AI in the Wild [[Github](#)]
- NeurIPS'25 Ethan Ewer, Daewon Chae, Thomas Zeng, Jinkyu Kim, and **Kangwook Lee**, "ENTP: Encoder-only Next Token Prediction," NeurIPS 2025 (WCTD Workshop) **spotlight**
- ICMLW'25 Ying Fan, Fei Deng, Yang Zhao, Sahil Singla, Rahul Jain, Tingbo Hou, **Kangwook Lee**, Feng Yang, Deepak Ramachandran, and Qifei Wang, "Improvement-Guided Iterative DPO for Diffusion Models" ICML 2025 Workshop
- ICLRW'25 Minjae Lee, Wonjun Kang, Byeongkeun Ahn, Christian Classen, Minghao Yan, Hyung Il Koo, and **Kangwook Lee**, "In-batch Ensemble Drafting: Toward Fast and Robust Speculative Decoding for Multimodal Language Models" ICLR 2025 (SCOPE Workshop)
- ICMLW'25 Soyeon Choi, **Kangwook Lee**, Oliver Sng, and Joshua M. Ackerman, "Infected Smallville: How Disease Threat Shapes Sociality in LLM Agents," ICML 2025 Workshop [[Summary](#)]
- EMNLPW'25 Dongmin Park, Minkyu Kim, Beongjun Choi, Junhyuck Kim, Keon Lee, Jonghyun Lee, Inkyu Park, Byeong-Uk Lee, Jaeyoung Hwang, Jaewoo Ahn, Ameya Sunil Mahabaleshwarkar, Bilal Kartal, Pritam Biswas, Yoshi Suhara, **Kangwook Lee**, Jaewoong Cho, "Orak: A Foundational Benchmark for Training and Evaluating LLM Agents on Diverse Video Games," EMNLP 2025 (Wordplay Workshop)
- ICLRW'25 Nayoung Lee, Ziyang Cai, Avi Schwarzschild, **Kangwook Lee**, and Dimitris Papailiopoulos, "Self-Improving Transformers Overcome Easy-to-Hard and Length Generalization Challenges," ICLR 2025 Workshop on Scaling Self-Improving Foundation Models
- NeurIPS'25 Jungtaek Kim, Ziqian Lin, Thomas Zeng, Minjae Lee, Chungpa Lee, Jy-yong Sohn, Hyung Il Koo, and **Kangwook Lee**, "Transformers in the Dark: Navigating unknown search spaces via noisy feedback" NeurIPS 2025 (WCTD Workshop)
- ICLRW'24 Jongho Park, Jaeseung Park, Zheyang Xiong, Nayoung Lee, Jaewoong Cho, Samet Oymak, **Kangwook Lee**, and Dimitris Papailiopoulos, "Can Mamba Learn How To Learn? A Comparative Study on In-Context Learning Tasks" ICLR 2024 Workshop on ME-FoMo
- ICLRW'24 Ziqian Lin and **Kangwook Lee**, "Dual Operating Modes of In-Context Learning" ICLR 2024 Workshop on ME-FoMo
- NeurIPS'24 Ying Fan, Yilun Du, Kannan Ramchandran, and **Kangwook Lee**, "Looped Transformers for Length Generalization" NeurIPS 2024 Workshop on MATH-AI
- NeurIPS'24 Kevin Galim*, Wonjun Kang*, Yuchen Zeng*, Hyung Il Koo, and **Kangwook Lee**, "Parameter-Efficient Fine-tuning of State Space Models," NeurIPS 2024 Workshop on Fine-Tuning in Modern ML **oral**
- NeurIPS'24 Ying Fan, Steve Yadlowsky, Dimitris Papailiopoulos, and **Kangwook Lee**, "Transformers Can Learn Meta-skills for Task Generalization in In-Context Learning" NeurIPS 2024 Compositional Learning Workshop
- ICMLW'23 Joseph Shenouda, Rahul Parhi, **Kangwook Lee**, and Robert D Nowak, "A Representer Theorem for Vector-Valued Neural Networks" ICML 2023 Workshop on Duality Principles
- NeurIPS'23 Ziqian Lin, Yicong Chen, Yuchen Zeng, and **Kangwook Lee**, "Coded Prompts for Large Language Models" NeurIPS 2023 Workshop on R0-FOMO
- NeurIPS'23 Sehyun Kwon, Jaeseung Park, Minkyu Kim, Jaewoong Cho, Ernest K. Ryu, and **Kangwook Lee**, "Image Clustering Conditioned on Text Criteria" NeurIPS 2023 Workshop on R0-FOMO
- ICMLW'23 Liu Yang, **Kangwook Lee**, Robert D Nowak, and Dimitris Papailiopoulos, "Looped Transformers are Better at Learning Learning Algorithms" ICML 2023 Workshop on ES-FoMo
- ICLRW'23 Angeliki Giannou, Shashank Rajput, Jy-yong Sohn, **Kangwook Lee**, Jason D. Lee, and Dimitris Papailiopoulos, "Looped Transformers as Programmable Computers" ICLR 2023 Workshop on ME-FoMo

- ICLRW'23 Kartik Sreenivasan, Keon Lee, Jeong-Gwan Lee, Anna Lee, Jaewoong Cho, Jy-yong Sohn, Dimitris Papailiopoulos, and **Kangwook Lee**, “Mini-Batch Optimization of Contrastive Loss” ICLR 2023 Workshop on ME-FoMo
- NeurIPSW'23 Yuchen Zeng, Kristjan Greenewald, **Kangwook Lee**, Justin Solomon, and Mikhail Yurochkin, “Outlier-Robust Group Inference via Gradient Space Clustering” NeurIPS 2023 Workshop on DistShift
- ICMLW'23 Seongjun Yang, Gibbeum Lee, Jaewoong Cho, Dimitris Papailiopoulos, and **Kangwook Lee**, “Predictive Pipelined Decoding: A Compute-Latency Trade-off for Exact LLM Decoding” ICML 2023 Workshop on ES-FoMo
- NeurIPSW'23 Adam Rouhiainen, Michael Gira, Gary Shiu, **Kangwook Lee**, and Moritz Münchmeyer, “Super-Resolution Emulation of Large Cosmological Fields with a 3D Conditional Diffusion Model” NeurIPS 2023 Workshop on ML and the Physical Sciences
- ICMLW'23 Nayoung Lee, Kartik Sreenivasan, Jason Lee, **Kangwook Lee**, and Dimitris Papailiopoulos, “Teaching Arithmetic to Small Transformers” ICML 2023 Workshop on Neural Conversational AI
- NeurIPSW'23 Yuchen Zeng and **Kangwook Lee**, “The Expressive Power of Low-Rank Adaptation” NeurIPS 2023 Workshop on OPT
- NeurIPSW'23 Ruisu Zhang, Yicong Chen, and **Kangwook Lee**, “Zero-shot Improvement of Object Counting with CLIP” NeurIPS 2023 Workshop on R0-FOMO
- NeurIPSW'22 Liu Yang, Jifan Zhang, Joseph Shenouda, Dimitris Papailiopoulos, **Kangwook Lee**, and Robert D. Nowak, “A Better Way to Decay: Proximal Gradient Training Algorithms for Weight Decay” NeurIPS 2022 OPT Workshop
- NeurIPSW'22 Dongmin Park, Dimitris Papailiopoulos, and **Kangwook Lee**, “Active Learning is a Strong Baseline for Data Subset Selection” NeurIPS 2022 HITY Workshop
- ACLW'22 Michael Gira, Ruisu Zhang, and **Kangwook Lee**, “Debiasing Pre-Trained Language Models via Efficient Fine-tuning” ACL 2022 Workshop on LT-EDI
- AAAIW'22 Jichang Chung, **Kangwook Lee**, and Kannan Ramchandran, “Federated Unsupervised Clustering with Generative Models” AAAI 2022 FL Workshop
- ICMLW'22 Tuan Dinh, Daewon Seo, Zhixu Du, Liang Shang, and **Kangwook Lee**, “Improved Input Reprogramming for GAN Conditioning” ICML 2022 UpML Workshop
- AAAIW'22 Yuchen Zeng, Hongxu Chen, and **Kangwook Lee**, “Improving Fairness via Federated Learning” AAAI 2022 FL Workshop
- ICMLW'22 Nayoung Lee*, Shashank Rajput*, Jy-yong Sohn, Hongyi Wang, Aliot Nagle, Eric P. Xing, **Kangwook Lee**, and **oral** Dimitris Papailiopoulos, “Super Seeds: Extreme Model Compression by Trading Off Storage with Computation” ICML 2022 UpML Workshop
- ICMLW'21 Daewon Seo, Hongyi Wang, Dimitris Papailiopoulos, and **Kangwook Lee**, “Empirical Study on the Effective VC Dimension of Low-rank Neural Networks” ICML 2021 Overparameterization Workshop
- ICMLW'21 Jinwoo Jeon, Jaechang Kim, **Kangwook Lee**, Sewoong Oh, and Jungseul Ok, “Gradient Inversion with Generative Image Prior” ICML 2021 Federated Learning Workshop
- ICMLW'20 Jy-yong Sohn, **Kangwook Lee**, Jaekyun Moon, and Dimitris Papailiopoulos, “GAN-mixup: Augmenting Across Data Manifolds for Improved Robustness” ICML 2020 Workshop on Uncertainty & Robustness in Deep Learning
- NeurIPSW'19 Donghwa Kim, **Kangwook Lee**, and Changho Suh, “Improving Model Robustness via Automatically Incorporating Self-supervision Tasks” NeurIPS 2019 MetaLearn Workshop
- SysML'18 **Kangwook Lee**, Kyungmin Lee, Hoon Kim, Changho Suh, and Kannan Ramchandran, “SGD on Random Mixtures: Private Machine Learning under Data-breach Threats” SysML 2018
- ICMLW'17 **Kangwook Lee***, Hoon Kim*, and Changho Suh, “Crash to not crash: Playing video games to predict vehicle collisions” ICML 2017 Workshop on ML for Autonomous Vehicles
- W'17 **Kangwook Lee**, Jichang Chung, and Changho Suh, “Large-scale and Interpretable Collaborative Filtering for Educational Data” KDD 2017 Workshop

- NeurIPSW'17 Jichang Chung, **Kangwook Lee**, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran, "Uber-Shuffle" NeurIPS 2017 ML Systems Workshop
- NeurIPSW'16 **Kangwook Lee**, Jichang Chung, Youngmin Cha, and Changho Suh, "Learning Analytics: Collaborative Filtering or Regression With Experts?" NeurIPS 2016 Workshop on ML for Education
- NeurIPSW'15 **Kangwook Lee**, Maximilian Lam, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran, "Speeding Up Distributed Machine Learning using Codes" NeurIPS 2015 ML Systems Workshop

Peer-Reviewed Journal Papers

- TMLR Shenghong Dai, Jy-yong Sohn, Yicong Chen, S M Iftexharul Alam, Ravikumar Balakrishnan, Suman Banerjee, Nageen Himayat, and **Kangwook Lee**, "Buffer-based Gradient Projection for Continual Federated Learning," *Transactions on Machine Learning Research* [Github]
- TMLR Ethan Ewer, Daewon Chae, Thomas Zeng, Jinkyu Kim, and **Kangwook Lee**, "ENTP: Encoder-only Next Token Prediction," *Transactions on Machine Learning Research*
- TMLR Ruisu Zhang, Yicong Chen, and **Kangwook Lee**, "Improving CLIP Counting Accuracy via Parameter-Efficient Fine-Tuning," *Transactions on Machine Learning Research* [Github]
- TMLR Jaewoong Cho*, Kartik Sreenivasan*, Keon Lee, Kyunghoo Mun, Soheun Yi, Jeong-Gwan Lee, Anna Lee, Jy-yong Sohn, Dimitris Papailiopoulos, **Kangwook Lee**, "Mini-Batch Optimization of Contrastive Loss," *Transactions on Machine Learning Research*
- TMLR Seongjun Yang, Gibbeum Lee, Jaewoong Cho, Dimitris Papailiopoulos, **Kangwook Lee**, "Predictive Pipelined Decoding: A Compute-Latency Trade-off for Exact LLM Decoding," *Transactions on Machine Learning Research*
- PRD Adam Rouhiainen, Michael Gira, Moritz Münchmeyer, **Kangwook Lee**, and Gary Shiu, "Superresolution emulation of large cosmological fields with a 3D conditional diffusion model," *Physical Review D*
- JMLR Joseph Shenouda, Rahul Parhi, **Kangwook Lee**, and Robert D. Nowak, "Variation Spaces for Multi-Output Neural Networks: Insights on Multi-Task Learning and Network Compression," *Journal of Machine Learning Research*
- TIT'22 **Kangwook Lee**, Nihar B. Shah, Longbo Huang, and Kannan Ramchandran, "Addendum and Erratum to "The MDS Queue: Analysing the Latency Performance of Erasure Codes" " *IEEE Transactions on Information Theory* 2022
- TNNLS Won Joon Yun, MyungJae Shin, David Mohaisen, **Kangwook Lee**, and Joongheon Kim, "Hierarchical Deep Reinforcement Learning-based Propofol Infusion Assistant Framework in Anesthesia" *IEEE Transactions on Neural Networks and Learning Systems* 2022
- GetMobile'21 Suman Banerjee, Remzi Arpaci-Dusseau, Shenghong Dai, Kassem Fawaz, Mohit Gupta, **Kangwook Lee**, and Shivaram Venkataraman, "The Roaming Edge and its Applications" *ACM GetMobile* 2021
- T-IT Kwangjun Ahn*, **Kangwook Lee***, and Changho Suh, "Community Recovery in Hypergraphs," *IEEE Transactions on Information Theory*
- T-SP **Kangwook Lee**, Kabir Chandrasekher, Ramtin Pedarsani, and Kannan Ramchandran, "SAFFRON: Sparse-Graph Code Framework for Group Testing," *IEEE Transactions on Signal Processing*
- JSTSP'17 Kwangjun Ahn, **Kangwook Lee**, and Changho Suh, "Hypergraph Spectral Clustering in the Weighted Stochastic Block Model" *IEEE Journal of Selected Topics in Signal Processing* 2018
- TIT'18 **Kangwook Lee**, Maximilian Lam, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran, "Speeding Up Distributed Machine Learning Using Codes," *IEEE Transactions on Information Theory*, January 2018 ▷ **IEEE ComSoc/IT Society Paper Award, 2020**
- ToN **Kangwook Lee**, Ramtin Pedarsani, and Kannan Ramchandran, "On Scheduling Redundant Requests With Cancellation Overheads," *IEEE/ACM Transactions on Networking*

Allerton'14 Ramtin Pedarsani, Dong Yin, **Kangwook Lee**, and Kannan Ramchandran, "PhaseCode: Fast and Efficient Compressive Phase Retrieval based on Sparse-Graph-Codes," *The 52nd Annual Allerton Conference on Communication, Control, and Computing*

T-IT **Kangwook Lee**, Nihar Shah, Longbo Huang, and Kannan Ramchandran, "The MDS Queue: Analysing the Latency Performance of Erasure Codes," *IEEE Transactions on Information Theory*

TComm Nihar Shah, **Kangwook Lee**, and Kannan Ramchandran, "When Do Redundant Requests Reduce Latency?," *IEEE Transactions on Communications*

Invited Journal Papers

GetMobile Suman Banerjee, Remzi Arpaci-Dusseau, Shenghong Dai, Kassem Fawaz, Mohit Gupta, **Kangwook Lee**, and Shivaram Venkataraman, "The Roaming Edge and its Applications" *ACM GetMobile 2021*

Non Peer-reviewed Conference Abstracts

- Linipun Phuttitarn, Robert McDermott, Chuan-Hong Liu, **Kangwook Lee**, Liang Shang, and Daewon Seo, "Deep Neural Networks for High-fidelity Measurement of Multiqubit Circuits" *APS March Meeting 2022*
- Yuchen Zeng, Ziqian Lin, and **Kangwook Lee**, "On a bilevel optimization approach to fair classification" *2022 INFORMS Optimization Society Conference*

RESEARCH TRAINEES

Postdocs

- Jungtaek Kim, Nov. 2024 – present
- Jy-yong Sohn, 2021–2022
⇒ Assistant Professor @ Yonsei University, Korea
- Daewon Seo, 2020–2021
⇒ Assistant Professor @ DGIST, Korea

PhD Students

- Nayoung Lee (co-advised with Prof. Dimitris Papailiopoulos)
- Liu Yang (co-advised with Prof. Rob Nowak and Prof. Dimitris Papailiopoulos)
- Joseph Shenouda (co-advised with Prof. Rob Nowak)
- Thomas Zeng
- Jongwon Jeong
- Yuchen Zeng
⇒ Senior Researcher @ Microsoft Research
- Ziqian Lin
⇒ Research Scientist @ Google
- Ying Fan
⇒ Senior Researcher @ Microsoft Research
- Tuan Dinh
⇒ Postdoc @ UCSF
- Changhun Jo

MS Students

- Ruisu Zhang
⇒ Machine Learning Engineer @ WeRide
- Andrew Geng
⇒ Research Engineer @ IBM
- Liang Shang
⇒ PhD student @ UW Madison

Undergraduate Students

- Ethan Ewer
- Lynnix Zou
- Jackson Kunde, *Hilldale research fellowship*
⇒ Machine Learning Engineer @ Ohalo
- Bryce Chen, *Hilldale research fellowship*
⇒ PhD student @ University of Washington, Seattle
- Michael Gira, *Hilldale research fellowship*
⇒ Software Engineer @ Microsoft

Visiting Researchers

- Chungpa Lee (Yonsei University)
- Dosung Lee (Korea University)

INVITED TALKS AND PANELS

Foundation Models

- NVIDIA GTC 2026 Panel, Mar. 2026.
Charting a Course for the Next Decade of Gaming with AI
- Department Seminar, Seoul National University, Dec. 2025.
AI for Video Games
- AWS Research Day for UW-Madison, Oct. 2025.
Toward More Efficient and Useful LLM Agents
- SILO Seminar, UW-Madison, July 2025.
Generative Agents in Social Psychology and Video Gaming
- ICML 2025 Workshop on Tiny-Titans, July 2025.
Towards Principled Design of SLM Agents for Edge Devices
- Oh Lab, University of Washington, June 2025.
Dual Operating Modes of In-Context Learning
- Department Seminar, Korea University, May 2025.
Generative AI and AI Agents
- UCSC ECE Seminar, Apr. 2025.
Bridging Large Language Models and Classical Machine Learning: From LIFT to LLM-Lasso
- Invited lecture at Microbiology/Oncology 545, UW Madison, Apr. 2025.
Bridging Large Language Models and Classical Machine Learning: From LIFT to LLM-Lasso
- Helmholtz/ELLIS Workshop on Foundation Models in Science, Berlin, Mar. 2025.
Bridging Large Language Models and Classical Machine Learning: From LIFT to LLM-Lasso

- Invited talk at KFAS, Mar. 2025.
Impacts of AI on Researchers
- EnCORE Workshop on Theoretical Perspective on LLMs, Mar. 2025.
Beyond Decoder-Only Next Token Prediction
- ECE Grad Seminar, University of Pittsburgh, Feb. 2025.
Beyond Decoder-Only Next Token Prediction
- SILO Seminar, UW-Madison, Nov. 2024.
ENTP: Encoder-only Next Token Prediction
- 2024 SIAM Conference on Mathematics of Data Science, Atlanta, Oct. 2024.
Dual Operating Modes of In-Context Learning
- IFDS seminar, UW Madison, Sept. 2024.
ENTP: Encoder-only Next Token Prediction
- Johns Hopkins University CIS/MINDS seminar, Apr. 2024.
Theoretical Exploration of Foundation Model Adaptation Methods
- UW Madison Machine Learning Lunch Meetings, Mar. 2024.
Dual Operating Modes of In-Context Learning
- 58th CISS @ Princeton University, Mar. 2024.
A Probabilistic Framework for Understanding In-Context Task Learning and Retrieval
- 2024 Information Theory and Applications Workshop, San Diego, Feb. 2024.
The Expressive Power of Low-Rank Adaptation (LoRA)
- Foundations of Data Science, UCSD/NSF EnCORE, Feb. 2024.
Theoretical Exploration of Foundation Model Adaptation Methods
- CSP Seminar, University of Michigan, Dec. 2023.
Towards a Theoretical Understanding of Parameter-Efficient Fine-Tuning (and Beyond)
- Efficient ML workshop, Google Research, New York, Nov. 2023.
The Expressive Power of Low-Rank Adaptation (LoRA)
- Trust Perspectives in ML, Law, and Public Policy, IDEAL, Northwestern University, Oct. 2023.
Demystifying Large Language Models: A Comprehensive Overview
- Symposium in Honor of AI Pioneer Professor Leonard Uhr, UW Madison, Oct. 2023.
Demystifying Large Language Models: A Comprehensive Overview
- Distinguished Lectures in Microbiology, UW Madison, Oct. 2023.
GPT: Transforming Science, Engineering, and Beyond
- AI in Imaging and Medicine, UW Madison, Oct. 2023.
The Potential of Large Language Models in Imaging and Medicine
- ML4MI, UW-Madison, Sept. 2023.
Exploring Generative AI: An Introduction to Large Language Models and Diffusion Models
- MADLab monthly meeting, July 2023.
Modular Deep Learning Systems with Pretrained Components
- BarryFest, Madison, June 2023.
Unfolding the Magic of GPT: A Flipped Classroom Ode to Barry
- Innovation in Data Seminar, Early Warning, June 2023.
GPT: Transforming Science, Engineering, and Beyond
- The second annual Wisconsin Digital Symposium, May 2023.
Power and Possibility of Large Language Models (like ChatGPT)

- KSEA Distinguished Guest Series, May 2023.
GPT: Transforming Science, Engineering, and Beyond
- Keynote, Midwest Regional Conference, Mar. 2023.
Revolutionizing Science and Engineering through Language Models
- Keynote, CSL Student Conference, UIUC, Feb. 2023.
Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance
- Information Theory and Applications Workshop, San Diego, Feb. 2023.
Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance
- SILO Seminar, UW-Madison, Feb. 2023.
Theoretical Exploration of Foundation Model Adaptation Methods
- Information Theory and Data Science Workshop, Singapore, Jan. 2023.
Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance
- SILO Seminar, UW-Madison, Jan. 2023.
Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance
- Samsung Advanced Institute of Technology (SAIT), Aug. 2022.
LIFT: Language-Interfaced FineTuning for Non-Language Machine Learning Tasks
- Electronic & Information Research Information Center, July 2022.
LIFT: Language-Interfaced FineTuning for Non-Language Machine Learning Tasks
- Physics meets Machine Learning Seminar, UW-Madison, Jan. 2022.
A gentle introduction to new ideas in modern ML
- KRAFTON, June 2021.
Recent Trends of AI Research
- Furiosa.ai, Mar. 2021.
Recent Trends of AI Research

Machine Learning Fairness

- Fairness and Ethics in ML Seminars, AmFam, Aug. 2023.
Equal Improbability: A New Fairness Notion Considering the Long-Term Impact
- UW-Madison Law School, Law 915, Apr. 2023.
Recent advances in Trustworthy ML
- The 12th US-Mexico Workshop on Optimization and its Applications, Huatulco, Mexico, Jan. 2023.
On a Bilevel Optimization Approach to Fair Classification
- KAIST CS Colloquium, Sept. 2022.
Recent advances in Trustworthy ML
- Electrical Engineering Department, University of Southern California, Apr. 2022.
On a bilevel optimization approach to fair classification
- CCDC Seminar Series, UC Santa Barbara, Apr. 2022.
On a bilevel optimization approach to fair classification
- EE Colloquium Lecture Series, KAIST, Mar. 2022.
On Trustworthy and Scalable Machine Learning
- BLISS Seminar, UC Berkeley, Feb. 2022.
Improving Fairness via Federated Learning
- MLWiNS Tech Talks, INTEL, Feb. 2022.
Improving Fairness via Federated Learning

- KAIST International Symposium on AI and future Society, Dec. 2021.
Improving Fairness via Federated Learning
- WID All Hands, UW-Madison, Nov. 2021.
On Trustworthy Machine Learning
- AI+Society seminar, UW-Madison, Oct. 2021.
On Trustworthy Machine Learning
- KRAFTON Developer Connect, Oct. 2021.
On Trustworthy Machine Learning
- POSTECH, June 2021.
Information Theory and Coding for Trustworthy and Scalable Machine Learning
- "Shannon meets Turing" Colloquium, Seoul National University, May 2021.
Information Theory and Coding for Trustworthy and Scalable Machine Learning
- IFDS Ethics & Algorithms SIG, UC Santa Cruz, Apr. 2021.
Fairness in AI
- Korea Information and Communications Society, Feb. 2021.
Fairness in AI
- Machine Learning Ideas, Microsoft Research New England, Dec. 2020.
Fairness in AI
- SILO Seminar, UW-Madison, Nov. 2020.
Fairness in AI
- BLISS Seminar, UC Berkeley, Nov. 2020.
Fairness in AI
- The 11th International Conference on ICT Convergence, Oct. 2020.
Information Theory and Coding for Trustworthy and Scalable Machine Learning
- Air Force Research Laboratory, May 2020.
FR-Train: A mutual information-based approach to fair and robust training
- The Chaos and Complex Systems Seminar, UW-Madison, Feb. 2020.
Information Theory and Coding for Machine Learning at Scale
- SK T-Brain, Jan. 2020.
Information Theory and Coding for Machine Learning at Scale
- Furiosa.ai, Jan. 2020.
Information Theory and Coding for Machine Learning at Scale
- ECE, UW-Madison, Mar. 2019.
Information Theory and Coding for Machine Learning at Scale

Distributed Machine Learning

- National Information Society Agency, Daegu, Jan. 2018.
Speeding Up Distributed Machine Learning Using Codes
- DGIST, Daegu, Jan. 2018.
Speeding Up Distributed Machine Learning Using Codes
- Seoul National University, Dec. 2017.
Speeding Up Distributed Machine Learning Using Codes
- Naver, May 2017.
Speeding Up Distributed Machine Learning Using Codes

- Information Theory and Machine Learning Workshop, KAIST, May 2017.
Speeding Up Distributed Machine Learning Using Codes
- Samsung Electronics DMC R&D Center, June 2016.
Speeding Up Distributed Machine Learning Using Codes
- Information Theory and Applications Workshop, Feb. 2016.
Speeding Up Distributed Machine Learning Using Codes

Others

- KRAFTON Developer Conference 2022, Oct. 2022.
Deep Learning and Video Games
- American Family Mutual Insurance Company's Visiting Professor Series, Oct. 2022.
Mixup Relabeling using Generative Models
- American Family Insurance Visiting Professor Series, Sept. 2021.
Mixup Relabeling using Generative Models
- SILO Seminar, UW-Madison, Oct. 2019.
Binary Rating Estimation with Graph Side Information
- Samsung Electronics, Aug. 2019.
Learning with Simulated Data
- The 29th Joint Conference on Communications and Information, Korea, May 2019.
Binary Rating Estimation with Graph Side Information
- Korea Information and Communications Society, Apr. 2019.
Learning with Simulated Data
- Korea Information and Communications Society, Jan. 2019.
Machine Learning (Introduction and Advanced Topics)
- Kakao Brain, May 2018.
Binary Rating Estimation with Graph Side Information
- UC Berkeley BASiCS Seminar, Nov. 2017.
Binary Rating Estimation with Graph Side Information
- National Information Society Agency, Daegu, Nov. 2016.
Machine Learning (Introduction and Advanced Topics)
- Seoul National University, Jan. 2016.
Sub-linear Time Algorithms for Sparse Signal Recovery Based on Sparse-graph Codes
- IEEE Communication Theory Workshop, May 2015.
A VoD System for Massively Scaled, Heterogeneous Environments
- University of Seoul, May 2015.
A VoD System for Massively Scaled, Heterogeneous Environments
- KAIST, May 2014.
The MDS Queue: Analysing the Latency Performance of Codes
- DIMACS Workshop on Algorithms for Green Data Storage, Rutgers University, Dec. 2013.
When Do Redundant Requests Reduce Latency?
- IEEE International Conference on Big Data, Oct. 2013.
The MDS Queue: Analysing the Latency Performance of Codes

GRANTS

Total **\$3,231,832**

- (G10) 2024 – 2029 CAREER: Theory and Algorithms for Learning with Frozen Pretrained Models. Amount: \$584,029
- (G09) 2020 – 2026 Industry research grants from Amazon, FuriosaAI, and American Family Insurance. Amount: \$833,500
- (G08) 2023 – 2024 Fall Research Competition from UW Madison. Amount: \$44,266
- (G07) 2022 – 2024 Research Forward initiative funding from UW Madison. Amount: \$150,000
- (G06) 2022 – 2023 Fall Research Competition from UW Madison. Amount: \$44,266
- (G05) 2021 – 2023 Understanding and Reducing Inequality Initiative (URI) funding from UW Madison. Amount: \$97,981
- (G04) 2021 – 2022 Fall Research Competition from UW Madison. Amount: \$44,266
- (G03) 2021 IITP High Potential Global Training Program, supported by POSTECH. Amount: \$44,739
- (G02) 2020 – 2023 UW2020: WARF Discovery Initiative Award. Amount: \$498,785
- (G01) 2020 – 2023 NSF/Intel Partnership on Machine Learning for Wireless Networking Systems (MLWiNS). Amount: \$890,000

TEACHING EXPERIENCE

University of Wisconsin-Madison

- ECE 901 Advanced Topics in Large Language Models (Fall 2025)
 - 13 students enrolled
 - Course evaluation: 4.80/5.00
- ECE/ISYE 570 Ethics of Data for Engineers (Spring 2025, Spring 2024)
 - 11 students enrolled
 - Course evaluation: 4.60/5.00
- ECE/CS/ME 539 Introduction to Artificial Neural Networks (Fall 2024)
 - 88 students enrolled
 - Course evaluation: 3.67/5.00
- ECE 901 Theory of Deep Learning Algorithms and Architectures (Spring 2023)
 - 10 students enrolled
 - Course evaluation: Not available
- ECE/CS 561 Probability and Information Theory in Machine Learning (Fall 2022)
 - 53 students enrolled
 - Course evaluation: 4.15/5.00
- ECE/CS/ME 532 Matrix Methods in Machine Learning (Spring 2022, Fall 2020, Fall 2019)
 - 115 students enrolled
 - Course evaluation: 4.02/5.00
- ECE 204 Data Science & Engineering (Fall 2021)
 - 20 students enrolled
 - Course evaluation: 4.36/5.00
- ECE/CS 761 Mathematical Foundations of Machine Learning (Spring 2021, Spring 2020)
 - 40 students enrolled
 - Course evaluation: 4.25/5.00

KAIST

- *Not* allowed to teach while serving military duty (June 16 – June 19)

University of California, Berkeley

- EECS 126 Probability and Random Processes at UC Berkeley (Fall 2015, Fall 2014)
 - Head GSI
 - **The Outstanding Graduate Student Instructor Award (Top 10% GSIs)**

PROFESSIONAL ACTIVITIES

Journal Editor

- Action Editors, Transactions on Machine Learning Research (TMLR) (2026, 2025, 2024, 2023, 2022)

Journal Reviewer

- IEEE Transactions on Information Theory (2025, 2024, 2023, 2021, 2020, 2019, 2018, 2015, 2014)
- IEEE Journal on Selected Areas in Information Theory (2020)
- IEEE Journal of Selected Topics in Signal Processing (2018)
- IEEE Transactions on Signal Processing (2017)
- IEEE Internet of Things Journal (2017)
- IEEE/ACM Transactions on Networking (2019, 2018, 2017, 2016)
- IEEE Communications Letters (2019, 2018, 2016)
- IEEE Transactions on Services Computing (2017, 2016)
- IEEE Transactions on Parallel and Distributed Systems (2019, 2015)
- ACM Transactions on Mobile Computing (2016)
- ACM Transactions on Modeling and Performance Evaluation of Computing Systems (2018, 2017, 2016)
- Springer Machine Learning ECML-PKDD 2021 Special Issue (2021)

Conference Senior Program Committee Member

- NeurIPS Area Chair (2025, 2024, 2023, 2022, 2021)
- ICML Area Chair (2026, 2025, 2024, 2023)
- ICLR Area Chair (2026, 2025)
- COLM Area Chair (2026, 2025, 2024)
- MLSys Program Committee (2026, 2025, 2024, 2023, 2022, 2021, 2020)
- IJCAI Senior Program Committee (2021)

Conference Program Committee Member (Reviewer)

- ICML (2022, 2021, 2019, 2018)
- NeurIPS (2020, 2019, 2018, 2017)
- ICLR (2024, 2023, 2022, 2021)
- COLT (2025, 2024, 2023, 2022, 2021)
- AISTATS (2025, 2023, 2022, 2021, 2020)
- FAccT (2023, 2022)
- IEEE ISIT (2023, 2022, 2021, 2020, 2019, 2018, 2017, 2016, 2015, 2014)
- IJCAI-ECAI (2022)
- AAAI (2021)

- ECML-PKDD (2020)
- ACML (2019, 2018)
- ITW (2018)
- IEEE ICC (2021)
- IEEE INFOCOM (2015, 2013)
- IEEE Globecom (2016)

Workshop Program Committee Member

- NeurIPS 2022 OPT Workshop
- LT-EDI-2022
- ScaDL
- EMDL @ MobiCom 2020

Memberships

- Member, IEEE / IEEE Young Professionals / IEEE Information Theory Society, 2014 – present
- Life-time Member, Korean-American Scientists and Engineers Association (KSEA), 2019 – present

INDUSTRIAL EXPERIENCE (BEFORE 2016)

Lytmus, Inc., San Francisco, CA

– *Software Engineering Intern*

June 2013 – Sept. 2013

Samsung Electronics, Suwon, South Korea

– *Software Engineering Intern*

July 2009

LG Display, Gumi, South Korea

– *Software/Hardware Engineering Intern*

June 2008 – Aug. 2008

OUTREACH/SERVICE ACTIVITIES

- Invited reviewer/mentor, New In ML workshop at NeurIPS 2021, December 2021
- Invited speaker, Dream Lecture series, Korea Foundation for Advanced Studies, 2021 – present
 - The 99th invited dream lecture, Seoul, May 2022
 - Seonrin Internet High School, Seoul, May 2022
 - Chodang High School, Virtual, November 2021
 - Wonju Sahmyook middle School, Virtual, July 2021
 - Hyeongseok High School, Virtual, March 2021
- Invited speaker, Seoul Science High School, 2018 – present
 - Studying abroad seminar, Virtual, October 2021
 - Career seminar, Virtual, April 2021
 - Career seminar, November 2018
- Invited speaker/mentor, KAIST, 2012 – present
 - Career paths in EE, Virtual, December 2020
 - Freshman Seminar, October 2016

- Mentor of KAIST Presidential Fellows, 2012 – 2015
- Invited speaker, Young Engineers Honor Society of the National Academy of Engineering of Korea
 - Career Seminar, Virtual, May 2020
 - Open Seminar, August 2018
 - Career Seminar, May 2017
- Founder, Freedu (free online tutoring service for underprivileged students), 2014-2015
- Mentor, Dream Concert (offline mentoring service for underprivileged students), 2019, 2015, 2014

OTHER AWARDS AND HONORS

- NeurIPS 2018 Travel Grant, 2018
- Best Paper Award Runner-up, IEEE MASCOTS 2013.
- Chosen as One of 10 Korean Future Engineer Leaders, Korea Engineers' Club, 2009.
- First Prize (Minister of Knowledge Economy Award), Wearable Computer Contest, 2009.
- Exemplary Student Award, Korea Advanced Institute of Science and Technology, 2008.
- Dongbu Scholarship (Full funding), Dongbu Cultural Foundation, 2008-2010.
- National Scholarship for Engineering and Applied Science students (Full funding), KSAF, 2006-2008.
- Startup competitions/hackathons
 - 3rd Prize Winner, Bay BitHack 2014 (Bitcoin hackathon), Bitcoin Association of Berkeley
 - 3rd Prize Winner, Springboard 2012 (Startup weekend), TIDE Institute
 - 2nd Prize Winner, Springboard 2011 (Startup weekend), TIDE Institute