Soo Min Kwon

soominkwon0402@gmail.com | Google Scholar | soominkwon | soominkwon.github.io

EDUCATION

University of Michigan Sept. 2022 – May 2026
Ph.D., Electrical Engineering and Computer Science Ann Arbor, MI

- **Thesis**: "Deep Learning through Low-Dimensional Representations: Theory and Algorithms"
- **Advisor(s)**: Prof. Laura Balzano and Prof. Qing Qu

Rutgers University Sept. 2016 – May 2022
M.S., B.S., Electrical and Computer Engineering New Brunswick, NJ

- **Thesis**: "Optimization Problems with Low-Dimensional Tensor Structures"
- **Advisor**: Prof. Anand D. Sarwate

EXPERIENCE

Google July 2025 – Nov. 2025
Student Researcher New York, NY

- Hosted by Himanshu Jain and Ziteng Sun
- Developed CoDistill-GRPO, a novel variant of GRPO aimed at improving small language model performance via co-distillation training with larger language models
- Tested the effectiveness of CoDistill-GRPO on the Qwen and Llama model families and with models up to 8B parameters, demonstrating improvement in accuracy compared to GRPO and baselines on mathematical reasoning tasks

Amazon July 2024 – Nov. 2024
Applied Scientist Intern Seattle, WA

- Developed a causal inference framework using deep learning methods for the SCOT team that reduced variance estimates of average treatment effects by over 10%

LinkedIn May 2022 – Aug. 2022
Applied Research Scientist Intern Sunnyvale, CA

- Productionized a machine learning pipeline for the infrastructure team, reducing MAPE by over 15% in forecasting hardware needs for the next calendar year

SELECTED PUBLICATIONS († equal contribution)

[1] **Out-of-Distribution Generalization of In-Context Learning: A Low-Dimensional Subspace Perspective** [Paper]
S. M. Kwon†, *A. S. Xu*†, *C. Yaras, L. Balzano, Q. Qu*
AISTATS 2026; AISTATS 2026 CauScale Workshop **(Best Paper Award)**; DeepMath 2025 **(Oral)**

[2] **An Overview of Low-Rank Structures in the Training and Adaptation of Large Models** [Paper]
L. Balzano, T. Ding, B. D. Haeffele, ***S. M. Kwon****, Q. Qu, P. Wang, Z. Wang, C. Yaras* (α-β *Order)*
IEEE Signal Processing Magazine

[3] **Learning Dynamics of Deep Matrix Factorization Beyond the Edge of Stability** [Paper]
A. Ghosh†, ***S. M. Kwon***†, *R. Wang, S. Ravishankar, Q. Qu*
ICLR 2025

[4] **BLAST: Block-Level Adaptive Structured Matrices for Efficient DNN Inference** [Paper]
C. Lee, ***S. M. Kwon****, Q. Qu, H. Lee*
NeurIPS 2024

[5] **On the Relationship Between Small Initialization and Flatness in Deep Networks** [Paper]
S. M. Kwon, *L. Ding, L. Balzano, Q. Qu*
ICLR 2024 Workshop on Bridging the Gap Between Practice and Theory in Deep Learning

[6] **Efficient Compression of Overparameterized Deep Models through Low-Dimensional Learning Dynamics** [Paper]
S. M. Kwon, *Z. Zhang, D. Song, L. Balzano, Q. Qu*
AISTATS 2024

[7] **Solving Inverse Problems with Latent Diffusion Models via Hard Data Consistency** [Paper]
B. Song†, ***S. M. Kwon***†, *Z. Zhang, X. Hu, Q. Qu, L. Shen*
ICLR 2024 **(Spotlight, Top 5%)**

PREPRINTS

[1] **CoDistill-GRPO: A Co-Distillation Recipe for Efficient Group Relative Policy Optimization** [Paper]
S. M. Kwon, *Himanshu Jain, Ziteng Sun, Ananda Theertha Suresh, Sanjiv Kumar*

[2] **The Effect of Training Task Diversity on In-Context Learning through the Lens of Low-Dimensional Subspaces** [Paper]
S. M. Kwon, *Alec S. Xu, Can Yaras, Dogyoon Song, Laura Balzano, Qing Qu*

[3] **Dynamic Subspace Estimation from Undersampled Data using Grassmannian Geodesics** [Paper]
S. M. Kwon†, *C. Blocker*†, *H. Raja, J. Fessler, L. Balzano*

[4] **Decoupled Data Consistency for Solving General Inverse Problems with Diffusion Models** [Paper]
X. Li, ***S. M. Kwon***, *I. Alkhouri, S. Ravishankar, Q. Qu*

AWARDS & HONORS

Best Paper Award at the AISTATS 2026 CauScale Workshop 2026
Harvey G. and Joyce H. Behner Graduate Fellowship 2024
University of Michigan PhD Rackham Merit Fellowship 2023
Rutgers ECE Outstanding Master's Student Award 2022
Rutgers ECE Outstanding Teaching Assistant Award 2021
Rutgers ECE Departmental Leadership & Service Award 2020
Rutgers WINLAB GA/TA Grant 2020
Rutgers University Dean's List 2018 – 2020

INVITED TALKS

2025 INFORMS Annual Meeting 2025
Math Machine Learning Seminar @ MPI MIS + UCLA 2025