Sejoon Oh, Ph.D.

Senior AI Research Scientist @ Netflix
Santa Clara, CA | 404-889-1929 | sejun6431@gmail.com
sejoonoh.github.io | linkedin.com/in/sejoon-oh | Google Scholar

Research Summary

Senior AI Researcher specializing in **LLM post-training, agentic systems, personalization, and multimodal understanding**. At Netflix, I lead research on production-scale LLM systems combining multi-task supervised fine-tuning (SFT), reinforcement learning (RL), and semantic-ID representations for personalized content understanding and recommendation. My work has delivered 77% higher new-content recommendation accuracy, 22% improvement in content understanding, and 44% better accuracy in tool-augmented prediction, while accelerating training runtime by 2.6×. My research has also contributed to production systems associated with $10M+$ in year-over-year revenue impact and been published to NeurIPS and ICML workshops. Ph.D. in Computer Science from Georgia Tech, with additional research contributions in multimodal LLM safety and robustness.

Experience

Netflix Los Gatos, CA
Senior AI Research Scientist *June 2024 – Present*

LLM Post-training & Personalization:

1. Led post-training research for Netflix LLMs, combining multi-task supervised fine-tuning (SFT) and Group Relative Policy Optimization (GRPO) with semantic-ID representations to improve personalized recommendation and content understanding, while retaining language capabilities.
2. Developed the latest production model, achieving 77% higher new-content recommendation accuracy, 22% improvement in content understanding, and 2.6× faster training runtime.
3. Integrated semantic IDs into LLM-based personalization, with research accepted at the NeurIPS 2026 PALM Workshop.

Agentic Tool Calling:

1. Leading the development of tool-augmented LLM training and inference to improve prediction freshness and accuracy by leveraging Netflix's existing tools and data sources, including catalog information, ranking models, and user taste summaries.
2. Enabled LLM predictions to incorporate up-to-date tool outputs, achieving 44% better accuracy in new-content prediction while improving instruction following.

Multimodal Understanding:

1. Led research on personalized artwork prediction, using LLMs to predict optimal content thumbnails conditioned on user profiles and title information.
2. Demonstrated that LLMs post-trained with SFT and Direct Preference Optimization (DPO) using rich content metadata and user interaction signals can outperform existing non-LLM production models.
3. Published findings at the ICML 2026 Pluralistic Alignment Workshop.

User Intent Modeling (FM-Intent):

1. Invented a hierarchical multi-task foundation model to predict user session intent and capture evolving user interests.
2. Successfully productized the model across millions of users, driving measurable revenue impact.

Georgia Institute of Technology Atlanta, GA
Graduate Research Assistant *Aug 2019 – May 2024*

Multimodal LLMs and their Safety: Developed "UniGuard," a universal safety guardrail to defend Multimodal LLMs against jailbreak attacks (AAAI 2025). Demonstrated vulnerabilities in text-aware recommender systems via adversarial text rewriting (CIKM 2024).

Recommender Systems: Authored multiple papers on stabilizing recommendations (TKDD 2024) and rank list sensitivity (CIKM 2022). Investigated cross-modal projection in Multimodal LLMs (ACL 2024).

Netflix Los Gatos, CA
Machine Learning Research Intern *May 2023 – Aug 2023*

IntentRec: Preliminary work of FM-Intent. User intent prediction based on hierarchical multi-task learning.

The Home Depot Atlanta, GA
Data Science Research Intern *May 2021 – Aug 2021*

Real-time Recommendation: Researched and implemented "M2TRec," a metadata-aware multi-task transformer for session-based recommendations. Solved cold-start problems and improved real-time personalization accuracy.

Adobe Research San Jose, CA
Data Science Research Intern *May 2020 – Aug 2020*

Neural Tensor Completion: Developed influence-guided data augmentation techniques for neural tensor completion. Work resulted in a publication at CIKM 2021.

Education

Georgia Institute of Technology Atlanta, GA
Ph.D. in Computer Science (Advisor: Prof. Srijan Kumar) *Aug 2019 – May 2024*

Seoul National University Seoul, South Korea
B.S. in Computer Science and Engineering (Advisor: Prof. U Kang) *Mar 2012 – Aug 2018*

Recent Publications

Conferences

S. Oh, and Netflix Personalization Foundation Team, "Personalized LLM Recommendation with Semantic IDs: Autoregressive Beam Search vs. Multi-Token Prediction at Industry Scale." *Personalized, Aligned, Long-Term Memory for AI Systems Workshop at* ***NeurIPS***, 2026.

A. Nam, **S. Oh**, E. Kong, Y. Feng, M. Bhattacharya, "Modeling diverse preferences in movie artwork personalization with large language models." *Pluralistic Alignment Workshop at* ***ICML***, 2026.

S. Oh, Y. Jin, M. Sharma, et al. "UniGuard: Towards Universal Safety Guardrails for Jailbreak Attacks on Multimodal Large Language Models." *Deployable AI Workshop at* ***AAAI***, 2025.

S. Oh, M. Bhattacharya, Y. Feng, S. Lamkhede. "FM-Intent: Predicting User Session Intent with Hierarchical Multi-Task Learning." *Netflix Tech Blog*, 2025.

G. Verma... **S. Oh**, et al. "Cross-Modal Projection in Multimodal LLMs Doesn't Really Project Visual Attributes to Textual Space." **ACL**, 2024.

Honors & Awards

Kwanjeong Ph.D. Fellowship | $150K total funding, awarded to top Korean Ph.D. 2019 – 2024

Best Thesis Award | Seoul National University, 1st among all CSE undergrads 2018

Samsung Humantech Paper Award | Gold Prize (1st in Computer Science) 2018