Min-hwan Oh

Updated: December 2025

Contact Information

Building 43 Room 509
Seoul National University
1 Gwanak-ro, Gwanak-gu
Seoul, South Korea, 08826

minoh@snu.ac.kr
https://minoh.io

Academic Appointment

Seoul National University, Seoul, South Korea
Graduate School of Data Science

Associate Professor — **September 2024–Present**

Assistant Professor — **September 2020–August 2024**

AI Institute of Seoul National University
Head of AI Technology Research Division — **October 2025–Present**

Research Interests

Sequential decision making under uncertainty, reinforcement learning, bandit algorithms, stochastic optimization, statistical machine learning

Education

Columbia University, New York, NY, USA

Ph.D., Operations Research / Data Science Specialization — **2020**

M.S., Operations Research (M.S.-Ph.D. Track) — **2016**

Advisor: Garud Iyengar / Co-advisor: Assaf Zeevi

Ph.D. Dissertation: *Sequential Decision Making with Combinatorial Actions and High-Dimensional Contexts*

- INFORMS George B. Dantzig Dissertation Award Finalist, 2020

Columbia University, New York, NY, USA

B.A., Mathematics & Statistics

Summa cum laude

Departmental Honors in Statistics

Phi Beta Kappa

Refereed Publications

51. **Exploration via Feature Perturbation in Contextual Bandits**.
S. Yi and M. Oh
Neural Information Processing Systems (**NeurIPS**), Spotlight, 2025.

50. **Infrequent Exploration in Linear Bandits**.
H. Lee and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2025.

49. **Preference-based Reinforcement Learning beyond Pairwise Comparisons: Benefits of Multiple Options**.
J. Lee, S. Yi, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2025.

48. **Tractable Multinomial Logit Contextual Bandits with Non-Linear Utilities**.
T. Hwang, D. Kim, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2025.

47. **Thompson Sampling for Multi-Objective Linear Contextual Bandit**.
S. Park, H. Ann, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2025.

46. **True Impact of Cascade Length in Contextual Cascading Bandits**.
H. Choi, J. Lee, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2025.

45. **Oracle-Efficient Combinatorial Semi-Bandits**.
J. Kim, M. Vojnovic, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2025.

44. **Revisiting Follow-the-Perturbed-Leader with Unbounded Perturbations in Bandit Problems**.
J. Lee, J. Honda, S. Ito, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2025.

43. **EUGens: Efficient, Unified and General Dense Layers**.
S.M. Kim, B. Kim, A. Sehanobish, S.B.R. Chowdhury, R. Kidambi, D. Shim, A. Dubey, S. Chaturvedi, M. Oh, and K. Choromanski
Neural Information Processing Systems (**NeurIPS**), 2025.

42. **AI Should Sense Better, Not Just Scale Bigger: Adaptive Sensing as a Paradigm Shift**.
E. Baek, K. Park, J. Ko, M. Oh, T. Gong, H.-S. Kim
Neural Information Processing Systems (**NeurIPS**), Position Paper, 2025.

41. **Improved Online Confidence Bounds for Multinomial Logistic Bandits**.
J. Lee and M. Oh
International Conference on Machine Learning (**ICML**), 2025.

40. **Combinatorial Reinforcement Learning with Preference Feedback**.
J. Lee and M. Oh
International Conference on Machine Learning (**ICML**), 2025.

39. **Practical and Near-Optimal Algorithm for Batched Linear Bandits**.
S. Yu and M. Oh
International Conference on Machine Learning (**ICML**), 2025.

38. **Symmetry-Aware GFlowNets**.
H. Kim, S. Lee, and M. Oh
International Conference on Machine Learning (**ICML**), 2025.

37. **Linear Bandits with Partially Observable Features**.
W. Kim, S. Park, G. Iyengar, A. Zeevi, and M. Oh
International Conference on Machine Learning (**ICML**), 2025.

36. **Experimental Design for Semiparametric Bandits**.
S. Kim, G. Kim, and M. Oh
Conference on Learning Theory (**COLT**), 2025.

35. **ADAM Optimization with Adaptive Batch Selection**.
K.Y. Kim and M. Oh
International Conference on Learning Representations (**ICLR**), 2025.

34. **Minimax Optimal Reinforcement Learning with Quasi-Optimism**.
H. Lee and M. Oh
International Conference on Learning Representations (**ICLR**), 2025.

33. **Lasso Bandit with Compatibility Condition on Optimal Arm**.
H. Lee, T. Hwang, and M. Oh
International Conference on Learning Representations (**ICLR**), 2025.

32. **Adversarial Policy Optimization for Preference-based Reinforcement Learning**.
H. Kang and M. Oh
International Conference on Learning Representations (**ICLR**), 2025.

31. **Dynamic Multi-product Selection and Pricing under Preference Feedback**.
J. Kim and M. Oh
International Conference on Learning Representations (**ICLR**), 2025.

30. **Nearly Minimax Optimal Regret for Multinomial Logistic Bandit**.
J. Lee and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2024.

29. **Randomized Exploration for Reinforcement Learning with Multinomial Logistic Function Approximation**.
W. Cho, T. Hwang, J. Lee, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2024.

28. **Local Anti-Concentration Class: Logarithmic Regret for Greedy Linear Contextual Bandit**.
S.J. Kim and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2024.

27. **Queueing Matching Bandits with Preference Feedback**.
J. Kim and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2024.

26. **Improved Regret of Linear Ensemble Sampling**.
H. Lee and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2024.

25. **Follow-the-Perturbed-Leader with Fréchet-type Tail Distributions: Optimality in Adversarial Bandits and Best-of-Both-Worlds**.
J. Lee, J. Honda, S. Ito, and M. Oh
Conference on Learning Theory (**COLT**), 2024.

24. **Demystifying Linear MDPs and Novel Dynamics Aggregation Framework**.
J. Lee and M. Oh
International Conference on Learning Representations (**ICLR**), 2024.

23. **Learning Uncertainty-Aware Temporally-Extended Actions**.
J. Lee, S.J. Park, Y. Tang, and M. Oh
AAAI Conference on Artificial Intelligence (**AAAI**), 2024.

22. **Mixed-Effects Contextual Bandits**.
K Lee, M.C. Paik, M. Oh, and G-S Kim
AAAI Conference on Artificial Intelligence (**AAAI**), 2024.

21. **Doubly Perturbed Task Free Continual Learning**.
B.H. Lee, M. Oh, and S.Y. Chun
AAAI Conference on Artificial Intelligence (**AAAI**), 2024.

20. **Bridging the Gap between Self-Rreport and Behavioral Laboratory Measures: A Real-time Driving Task with Inverse Reinforcement Learning**.
S.H. Lee, , M.S. Song, M. Oh, and W.Y. Ahn
Psychological Science, 2024.

19. **Cascading Contextual Assortment Bandits**.
H. Choi, R. Udwani, and M. Oh
Neural Information Processing Systems (**NeurIPS**), 2023.
• Minister of Science and ICT Paper Award, K-Data Science Conference, 2023

18. **Combinatorial Neural Bandits**.
T. Hwang, K. Chai, and M. Oh
International Conference on Machine Learning (**ICML**), 2023.

17. **Model-based Offline Reinforcement Learning with Count-based Conservatism**.
B. Kim and M. Oh
International Conference on Machine Learning (**ICML**), 2023.

16. **Semi-Parametric Contextual Pricing Algorithm using Cox Proportional Hazards Model**.
Y-G. Choi, G-S. Kim, Y. Choi, W. Cho, M.C. Paik, and M. Oh
International Conference on Machine Learning (**ICML**), 2023.

15. **Squeeze All: Novel Estimator and Self-Normalized Bound for Linear Contextual Bandits**.
W. Kim, M.C. Paik, and M. Oh
International Conference on Artificial Intelligence and Statistics (**AISTATS**) 3098–3124. 2023.

14. **Model-based Reinforcement Learning with Multinomial Logistic Function Approximation**.
T. Hwang and M. Oh
AAAI Conference on Artificial Intelligence (**AAAI**), 2023.

13. **Stochastic-Expert Variational Autoencoder for Collaborative Filtering**.
Y-S. Cho and M. Oh
Proceedings of the ACM Web Conference (**WWW**), 2482–2490, 2022.

12. **Personalized Federated Learning With Server-Side Information**.
J. Song, M. Oh, and H.S. Kim
IEEE Access, 10, 120245-120255, 2022.

11. **Sparsity-Agnostic Lasso Bandit**.
M. Oh, G. Iyengar, and A. Zeevi
International Conference on Machine Learning (**ICML**), 8271–8280, 2021.
• INFORMS Applied Probability Society Student Paper Award Finalist, 2020

10. **Multinomial Logit Contextual Bandits: Provable Optimality and Practicality**.
M. Oh and G. Iyengar
Proceedings of the 35th AAAI Conference on Artificial Intelligence (**AAAI**), 9205–9213, 2021.

9. **Crowd Counting with Decomposed Uncertainty**.
M. Oh, P. Olsen, and K.N. Ramamurthy
Proceedings of the 34th AAAI Conference on Artificial Intelligence (**AAAI**), 11799–11806, 2020.

8. **Thompson Sampling for Multinomial Logit Contextual Bandits**.
M. Oh and G. Iyengar
Neural Information Processing Systems (**NeurIPS**), 3145–3155, 2019.

7. **Sequential Anomaly Detection using Inverse Reinforcement Learning**.
M. Oh and G. Iyengar
Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (**KDD**). 1480–1490, 2019.
• Oral presentation in research paper track (top 9% of total submissions)

6. **Automatic event detection in basketball using Hidden Markov Models with energy based defensive assignment**.
S. Keshri, M. Oh, S. Zhang, and G. Iyengar
Journal of Quantitative Analysis in Sports 15(2), 141-153, 2019.

5. **Adaptive Pattern Matching with Reinforcement Learning for Dynamic Graphs**.
H. Kanezashi, T. Suzumura, D. Garcia-Gasulla, M. Oh, and S. Matsuoka
IEEE International Conference on High Performance Computing, Data, and Analytics (**HiPC**), 92–101, 2018.
• Best Paper Award winner

4. **Learning Graph Topological Features via GAN**.
W. Liu, H. Cooper, M. Oh, P.Y. Chen, S. Yeung, F. Yu, T. Suzumura, G. Hu
IEEE Access, 7, 21834–21843, 133600, 2019.
Preliminary version appeared at *Workshop on Implicit Generative Models, International Conference on Machine Learning* (ICML), 2017.

3. **Efficient "Shotgun" Inference of Neural Connectivity from Highly Sub-sampled Activity Data**.
D. Soudry, S. Keshri, P. Stinson, M. Oh, G. Iyengar, and L. Paninski
PLoS Computational Biology, 11 (10), e1004464, 2015.

2. **Graphical Model for Basketball Match Simulation**.
M. Oh, S. Keshri, and G. Iyengar
MIT Sloan Sports Analytics Conference, 2015.
• Finalist in Research Paper Competition (top 2% of total submissions)

Workshop Papers

1. **Directed Exploration in PAC Model-free Reinforcement Learning**.
M. Oh and G. Iyengar.
Preliminary version appeared at *Exploration in Reinforcement Learning Workshop, International Conference on Machine Learning (ICML), 2018.*
• 2nd place winner, 2018 INFORMS Annual Meeting Poster Competition

Honors and Awards

Top Area Chair **2025**
NeurIPS 2025

INFORMS JFIG Paper Competition Finalist **2025**
INFORMS Annual Meeting

Best Paper Award **2025**
Korea Artificial Intelligence Association Conference

Outstanding Paper Award **2025**
Korea Artificial Intelligence Association Conference

Amazon Research Award **2025**
Amazon Science

Minister of Science and ICT Award **2025**
Korea Data Science Conference Research Competition
Advisee: Seoh-won Yi

Chief of Staff of the R.O.K. Army Award **2024**
International Army Modeling and Simulation Education Conference

Outstanding Paper Award **2024**
Korea Artificial Intelligence Association Conference

Minister of Science and ICT Award **2023**
Korea Data Science Conference Research Competition
Advisee: Hyunjun Choi

KT Best Paper Award **2023**
Korea Artificial Intelligence Association Conference

Outstanding Paper Award **2023**
Korea Artificial Intelligence Association Conference

Outstanding Paper Award **2022**
Korea Artificial Intelligence Association Conference

Creative-Pioneering Researcher **2021**
Seoul National University

George B. Dantzig Dissertation Award Finalist **2020**
INFORMS Annual Meeting

Best Student Paper Award Finalist **2020**
INFORMS Applied Probability Society

QSR Data Challenge Award Winner **2020**
INFORMS Annual Meeting

NAVER Doctoral Fellowship **2020**
NAVER Corporation

CKGSB Doctoral Fellowship **2018–2020**
Columbia University

Outstanding Teaching Assistant Award **2020**
Columbia University

Professional Development Scholarship **2020**
Columbia University

AAAI Student Scholarship **2020**
AAAI Conference on Artificial Intelligence

KSEA-KUSCO Scholarship **2019**
Korean-American Scientists and Engineers Association

W. Edwards Deming Doctoral Fellowship **2018**
Columbia University

Best Paper Award, **2018**
IEEE International Conference on HiPC

2nd Place Winner **2018**
INFORMS Annual Meeting Poster Competition

Summa cum laude **2015**
Columbia University

Statistics Departmental Honors **2015**
Columbia University

Phi Beta Kappa Honor Society **2015**
Columbia University

John A. Northcott Scholarship **2012–2015**
Columbia University

Dean's Scholarship, Columbia University **2011**

Industry Experience

IBM T. J. Watson Research Center, Yorktown Heights, NY, USA

Computational and Statistical Learning Group at IBM Research AI

Summer Research Intern **May–August 2018**

Summer Research Intern **May–August 2017**

Teaching Experience

Seoul National University, Seoul, South Korea

Graduate School of Data Science

Machine Learning & Deep Learning for Data Science **Spring/Fall 2022, Spring/Fall 2023, Spring 2024, Fall 2025**

Data Science & Reinforcement Learning **Spring/Fall 2021, Spring 2022, Spring 2023, Spring 2024, Spring 2025**

Special Lecture in Data Science: Reinforcement Learning **Fall 2020**

Columbia University, New York, NY, USA

Columbia University Science Honors Program
Graduate Instructor

Graph Theory by Example **Spring 2020**

Department of Industrial Engineering and Operations Research
Teaching Assistant

IEOR 4720 — Deep Learning **Fall 2018**

IEOR 4650 — Business Analytics **Spring 2017, Spring 2018**
IEOR 4007 — Optimization Methods for FE **Fall 2017**
IEOR 4404 — Simulation **Fall 2016**
IEOR 3106/4106 — Stochastic Models **Fall 2015, Spring 2016**

Guest Lecturer
IEOR 4650 — Business Analytics **Spring 2020**
IEOR 4106 — Stochastic Models **Spring 2016**

Department of Mathematics
Teaching Assistant (as undergraduate)
MATH 4106 — Modern Analysis I **Fall 2014**
MATH 2010 — Linear Algebra **Spring 2014**
MATH 1202 — Calculus IV **Fall 2013**
MATH 1201 — Calculus III **Spring 2013**

Invited Talks & Conference Presentation

"Provably Efficient Exploration for Multinomial Logit Contextual Bandits"
Amazon AWS Responsible AI Seminar November 2025

"Efficient Exploration in Contextual Bandits and RLHF with Combinatorial Actions"
Physical AI International Forum November 2025

"Efficient Reinforcement Learning"
IAMSEC 2025 November 2025

"Minimax Optimal Dynamic Assortment Selection with Contextual Information"
INFORMS October 2025

"Contextual Bandits with Combinatorial Actions & Preference Feedback"
Conference of the Korean Artificial Intelligence Association August 2025

"Nearly Minimax Optimal Regret for Multinomial Logistic Bandit"
INFORMS Annual Meeting October 2024

"Recent Advances in Greedy Contextual Bandits"
Conference of the Korean Artificial Intelligence Association August 2024

"Lasso Bandit with Compatibility Condition on Optimal Arm"
Bandit & RL Theory Workshop, KAIST June 2024

"Contextual Bandits with Combinatorial Actions"
The Korean Statistical Society Conference July 2024

"Cascading Contextual Assortment Bandits"
KIISE Korea Software Congress December 2023
Conference of the Korean Artificial Intelligence Association November 2023

Korea Institute For Advanced Study November 2023
Korean Society for Industrial and Applied Mathematics November 2023

"Real-Time Decision Making: Combinatorial Contextual Bandits"
Yonsei University September 2023

"Randomized Exploration for Reinforcement Learning with Multinomial Logistic Transition Model"
KAIST, Graduate School of Data Science June 2023

"Efficient Exploration in Reinforcement Learning with Multinomial Logistic Function Approximation"
KAIST, Stochastic Analysis & Application Research Center May 2023

"Reinforcement Learning Algorithms for Efficient Exploration"
Seoul National University, AI Institute December 2022

"Model-based Reinforcement Learning with Multinomial Logistic Function Approximation"
KAIA & NAVER Joint Autumn Conference November 2022

"Randomized Exploration in Structured Reinforcement Learning"
International Conference on Ubiquitous Robots July 2022

"Deep Exploration for Reinforcement Learning in Feature Space"
UNIST, Graduate School of Artificial Intelligence August 2022
Seoul National University, Graduate School of Data Science May 2022

"Sparsity-Agnostic Lasso Bandit"
The Korean Statistical Society Conference June 2022
Seoul National University, Department of Statistics September 2021
KAIST, Stochastic Analysis & Application Research Center August 2021

"Randomized Exploration in Sequential Decision Making"
Seoul National University, Department of Industrial Engineering May 2022

"Multi-armed Bandits for Sequential Decision Making"
AI Korea 2021 September 2021

"Online Decision Making: from Contextual Bandit to Reinforcement Learning"
Conference of the Korean Artificial Intelligence Association July 2021

"Towards Real-Life Reinforcement Learning"
Korea Research Institute of Ships and Ocean Engineering January 2021
Data Science Seminar, Seoul National University August 2020

"Sparsity-Agnostic Lasso Bandit"
INFORMS Annual Meeting November 2020
Virtual Conference on Reinforcement Learning for Real Life June 2020

"Thompson Sampling for Multinomial Logit Contextual Bandits"

INFORMS Annual Meeting	November 2020
IFORS 2020 (postponed)	June 2020
NeurIPS 2019, Vancouver	December 2019
IBM Thomas J. Watson Research Center	November 2019
INFORMS Annual Meeting, Seattle	October 2019
INFORMS Workshop on Data Mining & Decision Analytics	October 2019

"Crowd Counting with Decomposed Uncertainty"

INFORMS Annual Meeting	November 2020
AAAI 2020, New York	February 2020
Deming Doctoral Fellowship Seminar, Columbia University	April 2019

"Multinomial Logit Contextual Bandits"

INFORMS Annual Meeting, Seattle	October 2019
MSOM Conference, Singapore	July 2019
ICML 2019, Long Beach	June 2019
RM&P Conference, Stanford University	June 2019
POMS Annual Conference, Washington D.C.	May 2019
Data Science Day, Columbia University	April 2019

"Sequential Anomaly Detection using Inverse Reinforcement Learning"

INFORMS Workshop on Data Science	October 2019
KDD 2019, Anchorage	August 2019

"Automatic Event Detection in Basketball using HMM with Energy based Defensive Assignment"

INFORMS Annual Meeting, Seattle	October 2019
POMS Annual Conference, Washington D.C.	May 2019
Data Science Society Seminar, Columbia University	April 2018
NESSIS, Harvard University	September 2017
IBM Thomas J. Watson Research Center	June 2017

"Directed Exploration in PAC Model-Free Reinforcement Learning"

INFORMS Annual Meeting, Phoenix	November 2018
Princeton Day of Optimization, Princeton University	September 2018
IBM Thomas J. Watson Research Center	August 2018
ICML 2018, Stockholm	July 2018

"Graphical Model for Basketball Match Simulation"

Data Science Day, Columbia University	April 2016
Sports Analytics Seminar, Columbia University	March 2016
Columbia EPIC Graduate Student Research Seminar	February 2016
MIT Sloan Sports Analytics Conference, Boston	February 2015

Academic & Professional Services

Program Committee & Reviewer — ICML 2026 (Area Chair), ICLR 2026 (Area Chair), NeurIPS 2025 (Area Chair), ICML 2025 (Area Chair), AISTATS 2025, ICLR 2025, NeurIPS 2024, ICML 2024, COLT 2024, ICLR 2024, NeurIPS 2023, ICML 2023, ICLR 2023, NeurIPS 2022, ICML 2022, ICLR 2022, AAAI 2022, NeurIPS 2021, ICML 2021, KDD 2021, AAAI 2021, NeurIPS 2020

Journal Reviewer — Operations Research, Management Science, Annals of Statistics, Journal of Machine Learning Research (JMLR)

Session Chair — INFORMS Annual Meeting 2024, INFORMS Annual Meeting 2022, INFORMS Annual Meeting 2021, INFORMS Annual Meeting 2019, INFORMS Workshop on Data Mining & Decision Analytics 2019

Organizer — NeurIPS 2025 ML x OR Workshop: Mathematical Foundations and Operational Integration of Machine Learning for Uncertainty-Aware Decision-Making

Postdoc & Student Advising

Taehyun Hwang	Postdoc Researcher
Donghyun Lee	Postdoc Researcher
Heesang Ahn	Ph.D. Student
Sungwoo Cho	Ph.D. Student
Woo Seong Cho	Ph.D. Student
Hyunjun Choi	Ph.D. Student
Byeongchan Kim	Ph.D. Student
Doyoon Jeon	Ph.D. Student
Kyo Won Jin	Ph.D. Student
Gyu Yeol Kim	Ph.D. Student
Hohyun Kim (Co-advised with Seunggeun Lee)	Ph.D. Student
Joongkyu Lee	Ph.D. Student
Seouhwon Lee	Ph.D. Student
Jihyeong Park	Ph.D. Student
Somangchan Park	Ph.D. Student
Sungwoo Park	Ph.D. Student
Sanghoon Yu	Ph.D. Student
Hyeongjun Yun	Ph.D. Student
Shingeun Bang	M.S. student
Chai Won Kim	M.S. student
Dahngoon Kim	M.S. student
Younghoon Shin	M.S. student
Seunggyu Song	M.S. student
Deokgyu Yoon	M.S. student

Alumni	Wooseok Jang	M.S. 2022 / NH Investment & Securities
	Seungjoon Park	M.S. 2023 / Samsung Research
	Sooyoun Park	M.S. 2023 / LG AI Research
	Seokjin Kim	Undergrad Intern 2023 / Ph.D. Student at Columbia Univ.
	Jisang Yu	M.S. 2024 / CPA at PwC
	Yongsang Yoon	M.S. 2024 / The Financial Supervisory Service
	Kyuwook Chai	M.S. 2024 / Patent lawyer at Haean
	Young Hun Kim	M.S. 2024 / Samsung Electronics
	Jihoo Park	M.S. 2024 / LG CNS
	Kyungbok Lee	Postdoc Researcher / UNC Chapel Hill
	Kyonghyun Min	M.S. 2024 / Doosan
	Doeon Lim	M.S. 2025 / ROK Army Officer
	Daesoon Kim	M.S. 2025 / POSCO
	Sanghee Jung	M.S. 2025 / Samsung Research
	Jung-hun Kim	Postdoc Researcher / ENSAE Paris
	Jongyeong Lee	Postdoc Researcher / KIST