Jongho Park

jjhpark@berkeley.edu jongho.me

Education

University of California, Berkeley *Ph.D, EECS, 2025–*
Advisor: Jason D. Lee, Song Mei

University of Wisconsin-Madison *M.S., Computer Sciences, 2021*
Advisor: Christos Tzamos

University of California, Berkeley *B.A., Computer Science & Mathematics (high honors), 2019*

Work Experience

Research Intern, Microsoft Research, AI Frontiers *May 2026–Aug 2026*
Research Scientist, KRAFTON Inc. *2021–2025*
Team Lead, Core Research Team *2024–2025*
Worked primarily on scalable algorithms for LLMs at Krafton AI.
Advised research on quantization, in-context learning, RL finetuning.
Led efficient deployment of generative agents ("Smart ZOI") onto InZOI in collaboration with NVIDIA.
Projects: Smart ZOI, Long-term memory
Fulfilled alternative mandatory military service, 11/2021–11/2024.

Preprints

*α-β denotes alphabetical order. * denotes equal contribution.*

1. Scaling Discovery through Test-time Communication
   Jongho Park, Vasilis Kontonis, Shivam Garg, Akshay Krishnamurthy, Eric Horvitz, Dimitris Papailiopoulos.
   Under Review.
2. Effective test-time scaling of discrete diffusion through iterative refinement
   Sanghyun Lee, Sunwoo Kim, Seungryong Kim, **Jongho Park**, Dongmin Park.
   Structured Probabilistic Inference & Generative Modeling workshop at ICML 2026. Under Review.
3. Beyond RLHF: a unified theoretical framework of alignment
   Jihun Yun*, Juno Kim*, **Jongho Park**, Jongha Jon Ryu, Junhyuck Kim, Jaewoong Cho, Kwang-Sung Jun.
   Foundations of Post-training workshop at COLT 2025. Under Review.

Publications

1. ThinkSafe: self-generated safety alignment for reasoning models
   Seanie Lee, Sangwoo Park, Yumin Choi, Gyeongman Kim, Minki Kang, Jihun Yun, Sung Ju Hwang, Dongmin Park, **Jongho Park**. *NeurIPS 2026.*
2. Looped Diffusion Language Models
   Sanghyun Lee, Chunsan Hong, Seungryong Kim, Jonghyun Lee, **Jongho Park**, Dongmin Park. *NeurIPS 2026.*
3. On-policy self-distillation via prompt optimization
   Jongho Park, Donghyun Lee, Matei Zaharia, Jason D. Lee.
   Reinforcement Learning from World Feedback workshop at ICML 2026.
4. Lookahead unmasking elicits accurate decoding in diffusion language models
   Sanghyun Lee, Seungryong Kim, **Jongho Park**, Dongmin Park. *ICML 2026.*
5. Not all bits are equal: scale-dependent memory optimization strategies for reasoning models
   Junhyuck Kim, Ethan Ewer, Taehong Moon, **Jongho Park**, Dimitris Papailiopoulos.
   ICLR 2026. Efficient Reasoning workshop at NeurIPS 2025.
6. Task diversity shortens the ICL plateau
   Jaeyeon Kim, Sehyun Kwon, Joo Young Choi, **Jongho Park**, Jaewoong Cho, Jason D. Lee, Ernest K. Ryu.
   TMLR 2025.
7. Lexico: extreme KV cache compression via sparse coding over universal dictionaries
   Junhyuck Kim, **Jongho Park**, Jaewoong Cho, Dimitris Papailiopoulos.
   ICML 2025. ICLR Workshop on Sparsity in LLMs (spotlight).
8. Can Mamba learn how to learn? A comparative study on in-context learning
   Jongho Park, Jaeseung Park, Zheyang Xiong, Nayoung Lee, Jaewoong Cho, Samet Oymak, Kangwook Lee, Dimitris Papailiopoulos.
   ICML 2024.

9. Can large language models develop strategic reasoning? Post-training insights from learning chess
   Dongyoon Hwang, Hojoon Lee, Jaegul Choo, Dongmin Park, **Jongho Park**.
   Test-time Scaling and Reasoning Models workshop at COLM 2025.
10. First-order stochastic optimization with oblivious noise
    Ilias Diakonikolas, Sushrut Karmalkar, **Jongho Park**, Christos Tzamos (α-β).
    NeurIPS 2023.
11. GLM regression with oblivious corruptions
    Ilias Diakonikolas, Sushrut Karmalkar, **Jongho Park**, Christos Tzamos (α-β).
    COLT 2023.
12. Prompted LLMs as chatbot modules for long open-domain conversation
    Gibbeum Lee*, Volker Hartmann*, **Jongho Park***, Dimitris Papailiopoulos, Kangwook Lee.
    Findings of ACL 2023.
13. ReLU regression with Massart noise
    Ilias Diakonikolas, **Jongho Park**, Christos Tzamos (α-β).
    NeurIPS 2021.
14. On robust mean estimation under coordinate-level corruption
    Zifan Liu*, **Jongho Park***, Theodoros Rekatsinas, Christos Tzamos.
    ICML 2021.
15. Literary event detection
    Matthew Sims, **Jongho Park**, David Bamman.
    ACL 2019.

Awards

Kwanjeong Educational Foundation Fellowship *2019–2021*

CS Summer Research Assistantship *Summer 2020*
Awarded to outstanding first-year graduate students at UW-Madison.

CS Departmental Scholarship *2019–2020*

Teaching

Head Teaching Assistant, UW-Madison
Intro to Algorithms (COMPSCI 577) *Spring 2021*

Teaching Assistant, UW-Madison
Advanced Algorithms (COMPSCI 787) *Fall 2020*
Intro to Algorithms (COMPSCI 577) *Spring 2020*
Intro to Discrete Math (COMPSCI 240) *Fall 2019*

Undergraduate Student Instructor, UC Berkeley
Efficient Algorithms and Intractable Problems (CS 170) *2018–2019*
Intro to Discrete Math and Probability (CS 70) *Summer 2017*

Service

Conference Reviewer: ACL, AISTATS, COLM, ICML, ICLR, NeurIPS, STOC

Berkeley PhD EECS Admissions *2026*
Reviewed applications to Berkeley EECS for the Berkeley AI-Theory Group.

Berkeley Student Committee for Faculty Hiring *2026*
Participated in student interviews for EECS faculty candidates.

Mentoring

Jaden Park, Krafton AI research intern ($\rightarrow$ UW-Madison CS PhD) *2023–2024*
Junhyuck Kim, Krafton AI *2023–2025*
Sanghyun Lee, Krafton AI research intern ($\rightarrow$ KAIST AI PhD) *2025–2026*

Languages

English (native), Korean (native)
Python (PyTorch), LaTeX

References

Jason D. Lee — Associate Professor, UC Berkeley
Song Mei — Assistant Professor, UC Berkeley

Kangwook Lee — Associate Professor, UW-Madison; Krafton AI
Dimitris Papailiopoulos — Associate Professor, UW-Madison; Microsoft Research
Christos Tzamos — Associate Professor, University of Athens
Ilias Diakonikolas — Professor, UW-Madison