

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 Scientist, KRAFTON Inc.

2021–2025

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.

- [ThinkSafe: self-generated safety alignment for reasoning models](#)
Seanie Lee, Sangwoo Park, Yumin Choi, Gyeongman Kim, Minki Kang, Jihun Yun, Dongmin Park, **Jongho Park**, Sung Ju Hwang. *Under Review*.
- [Lookahead unmasking elicits accurate decoding in diffusion language models](#)
Sanghyun Lee, Seungryong Kim, **Jongho Park**, Dongmin Park. *Under Review*.
- [Effective test-time scaling of discrete diffusion through iterative refinement](#)
Sanghyun Lee, Sunwoo Kim, Seungryong Kim, **Jongho Park**, Dongmin Park. *Under Review*.
- [Alignment as distribution learning: your preference model is explicitly a language model](#)
Jihun Yun*, Juno Kim*, **Jongho Park**, Jongha Jon Ryu, Junhyuck Kim, Jaewoong Cho, Kwang-Sung Jun. *FoPT workshop at COLT 2025. Under Review*.

Publications

- [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.
- [Can large language models develop strategic reasoning? Post-training insights from learning chess](#)
Dongyoon Hwang, Hojoon Lee, Jaegul Choo, Dongmin Park, **Jongho Park**.
ScalR workshop at COLM 2025.
- [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.
- [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).
- [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.
- [First-order stochastic optimization with oblivious noise](#)
Ilias Diakonikolas, Sushrut Karmalkar, **Jongho Park**, Christos Tzamos (α - β).
NeurIPS 2023.
- [GLM regression with oblivious corruptions](#)
Ilias Diakonikolas, Sushrut Karmalkar, **Jongho Park**, Christos Tzamos (α - β).
COLT 2023.
- [Prompted LLMs as chatbot modules for long open-domain conversation](#)

Gibbeum Lee*, Volker Hartmann*, **Jongho Park***, Dimitris Papailiopoulos, Kangwook Lee.
Findings of ACL 2023.

9. [ReLU regression with Massart noise](#)
Ilias Diakonikolas, **Jongho Park**, Christos Tzamos (α - β).
NeurIPS 2021.
10. [On robust mean estimation under coordinate-level corruption](#)
Zifan Liu*, **Jongho Park***, Theodoros Rekatsinas, Christos Tzamos.
ICML 2021.
11. [Literary event detection](#)
Matthew Sims, **Jongho Park**, David Bamman.
ACL 2019.

Awards

Kwanjeong Educational Foundation Fellowship	<i>2019–2021</i>
CS Summer Research Assistantship Awarded to outstanding first-year graduate students at UW-Madison.	<i>Summer 2020</i>
CS Departmental Scholarship	<i>2019–2020</i>

Teaching

Head Teaching Assistant , UW-Madison Intro to Algorithms (COMPSCI 577)	<i>Spring 2021</i>
Teaching Assistant , UW-Madison Advanced Algorithms (COMPSCI 787)	<i>Fall 2020</i>
Intro to Algorithms (COMPSCI 577)	<i>Spring 2020</i>
Intro to Discrete Math (COMPSCI 240)	<i>Fall 2019</i>
Undergraduate Student Instructor , UC Berkeley Efficient Algorithms and Intractable Problems (CS 170)	<i>2018–2019</i>
Intro to Discrete Math and Probability (CS 70)	<i>Summer 2017</i>

Service

Conference Reviewer : ACL, AISTATS, NeurIPS, ICML, ICLR, STOC	
Berkeley PhD EECS Admissions Reviewed applications to Berkeley EECS for the Berkeley AI Theory Group.	<i>2026</i>
Berkeley Student Committee for Faculty Hiring Participated in student interviews for EECS faculty candidates.	<i>2026</i>

Mentoring

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

Languages

English (native), Korean (native)
Python (PyTorch), L^AT_EX

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