

Jikun Kang

Curriculum Vitae

McGill University


Mila - Québec AI Institute

☎ (+1) 5146928356

✉ luciferkonn@gmail.com

🌐 My Webpage: luciferkonn.github.io

🐙 Github in LinkedIn



Education

- 2017–2024 **PhD, School of Computer Science, McGill University**, Montreal, Canada.
Reinforcement Learning, Multi-agent Reinforcement Learning, Meta-Reinforcement Learning, LLM-based Reinforcement Learning
- 2014–2017 : **Master of Computer Software and Theory, Northeastern University**, Shenyang, China.
GPA : 3.63 / 4.00; Ranked: 10 / 180
- 2010–2014 : **Bachelor of Computer Science and Technology, Kunming University of Science and Technology**, Kunming, China.
GPA : 4.00/4.00; Ranked: 1 / 160

Publications

In Conference Proceedings

- 2024 Rongrong Wang, Duc Van Le, **Jikun Kang**, Rui Tan, , and Xue liu. Incentive temperature control for green colocation data centers via reinforcement learning. In *IEEE/ACM International Symposium on Quality of Service (IWQoS)*, 2024.
- 2024 **Jikun Kang**, Romain Laroche, Xindi Yuan, Adam Trischler, Xue Liu, and Jie Fu. Think before you act: Decision transformers with working memory. In *The International Conference on Machine Learning (ICML)*, 2024.
Explanation: This paper introduces a novel approach to add a memory module inside of transformer architecture.
- 2024 Qianxi Li, Yingyue Cao, **Jikun Kang**, Tianpei Yang, Xi Chen, and Matthew E. Taylor Jun Jin. Laffi: Leveraging hybrid natural language feedback for fine-tuning language. In *AAAI 2024 Human-Centric Representation Learning*, 2024.
Explanation: This paper leverages AI feedback to generate preferences and explanations for each data point, enhancing the alignment of LLMs.
- 2023 **Kang, Jikun**, Di Wu, Ju Wang, Ekram Hossain, Xue Liu, and Gregory Dedek. Multi-agent attention actor-critic algorithm for load balancing in cellular networks. In *ICC 2023-IEEE International Conference on Communications*, pages 5160–5165. IEEE, 2023.
- 2023 **Kang, Jikun**, Miao Liu, Abhinav Gupta, Christopher Pal, Xue Liu, and Jie Fu. Learning multi-objective curricula for robotic policy learning. In *Conference on Robot Learning*, pages 847–858. PMLR, 2023.
- 2022 **Kang, Jikun**, Ju Wang, Chengming Hu, Xue Liu, and Gregory Dudek. A generalized load balancing policy with multi-teacher reinforcement learning. In *GLOBECOM 2022-2022 IEEE Global Communications Conference*, pages 3096–3101. IEEE, 2022.
- 2021 Di Wu, **Jikun Kang**, Yitian Xu, Hang Li, Jimmy Li, Xi Chen, Dmitriy Rivkin, Jenkin Michael, Tasesop Lee, and Intaik Park. Load balancing for communication networks via data-efficient deep reinforcement learning. In *GLOBECOM*, pages 12–18. IEEE, 2021.

- 2021 **Jikun Kang**, Xi Chen, Di Wu, Yi Tian Xu, Xue Liu, Gregory Dudek, Taeseop Lee, and Intaik Park. Hierarchical policy learning for hybrid communication load balancing. In *ICC*, pages 1–6. IEEE, 2021.
- 2021 Chengming Hu, Xi Chen, Ju Wang, Hang Li, **Jikun Kang**, Yitian Xu, Xue Liu, Di Wu, Seowoo Jang, Intaik Park, and Gregory Dudek. Afb: Improving communication load forecasting accuracy with adaptive feature boosting. In *GLOBECOM*, pages 25–32. IEEE, 2021.
- 2020 Gupta Abhinav, Chakravoty Jhelum, **Jikun, Kang**, and Precup Doina. Multi-agent option critic framework. *Neurips RL*, 2020.

Journal Articles

- 2021 **Tianyi Chen, Shengrong Bu, Xue Liu, Jikun Kang, Yu F. Richard, and Han ZHu**, Peer-to-peer, transactive energy, blockchain power grid applications, In *IEEE Transactions on Smart Grid*.

Preprint

- 2025 **George Thomas, and Chan Alex J, Jikun Kang, Wenqi Wu, Filippos Christianos, Fraser Greenlee, Andy Toulis, Marvin Purtorab**, WebGames: Challenging General-Purpose Web-Browsing AI Agents, Arxiv: <https://arxiv.org/abs/2502.18356>.
Explanation: We introduce WebGames, a comprehensive benchmark suite designed to evaluate general-purpose web-browsing AI agents through a collection of 50+ interactive challenges.
- 2025 **Jikun Kang, Wenqi Wu, Filippos Christianos, Alex J. Chan, Fraser Greenlee, George Thomas, Marvin Purtorab, Andy Toulis**, LM2: Large Memory Models, Arxiv: <https://arxiv.org/abs/2502.06049>.
Explanation: The LM2 model integrates a memory module into the Transformer architecture to improve multi-step reasoning and information synthesis over long contexts. This enhancement leads to significant performance improvements in tasks requiring multi-hop inference and large-context question-answering, demonstrating the value of explicit memory in Transformer models.
- 2024 **Jikun Kang, Xin Zhe Li, Xi Chen, Amirreza Kazemi, Boxing Chen**, MindStar: Enhancing Math Reasoning in Pre-trained LLMs at Inference Time, Arxiv: <https://arxiv.org/abs/2405.16265v1>.
Explanation: In this work, we propose MindStar (M*), an inference-based method that improves reasoning in LLMs by formulating reasoning tasks as search problems. M* significantly enhances open-source models like Llama-2-13B and Mistral-7B, achieving comparable performance to larger models such as GPT-3.5 and Grok-1 while reducing model size and computational costs.

Working Experience

- September 2024 – **Founding member of technical staff (Research Team)**, *Convergence AI*, London, UK (remote).
- Present Focused on pre-training and post-training on memory-augmented language models and web agents.
- June 2023 – **AI Research Scientist**, *Noah's Ark Lab*, Montreal, Canada.
- August 2024 Focused on superalignment, improving reasoning abilities in LLM, RLHF, RLAIF, and foundation models for decision making in AI. I recently completed a project on MindStar, our latest research on large language model reasoning. By searching for the optimal reasoning path, we have successfully enhanced the reasoning capabilities of Llama2 to the level of Llama3. Our findings also demonstrate that open-source models like Mistral and Llama2 can achieve closed-source LLM reasoning performance with 13-25 times smaller model size. This method has been introduced as an important feature of Pangu 5.0 at HDC24.
- March 2020 – **AI Research Intern**, *Samsung AI Research Center*, Montreal, Canada.
- May 2023 Conducted research on AI for networking. During my internship, I published 5 papers in top networking conferences and secured 8 US patents, including 3 A1 patents and 5 A2 patents. A1 is the top patent ranking inside of Samsung.
- Apr.,2019 – **AI Research Intern**, *Learnable AI, Inc.*, Boston, US.
- Jul.,2019 Topics: AI for education

- Dec.,2018 – **AI Research Intern**, *Blaise Transit, Inc.*, Montreal, Canada.
Mar.,2019 Topics: AI for public transition
- Feb., 2017 – **Research Intern**, *Siemens China Research Institute*, Beijing, China.
Aug., 2017 Topics: VR for Factory 4.0

Awards

- 2024 NeurIPS 2024 Outstanding Reviewer
2024 ICML Conference Scholarship
2022 McGill University the GREAT Award
2017 – 2021 Merit scholarship program for foreign students.
2014 First Class Scholarship of Northeastern University.
2011 Provincial Government Scholarship.

Computer skills

Programming Languages Python, PyTorch, JAVA

Teaching Assistantship

- Fall,2018- Winter,2021: **COMP202: Foundations of Programming**, McGill University.
Winter,2019 **COMP535: Computer Networks**, McGill University.
Winter, 2023:
Winter,2018 : **COMP599: Advanced Topics in Computer Science**, McGill University.

Service

- 2025 **Reviewer of TMLR.**
2022-2025: **Reviewer of NeurIPS.**
2022-2025: **Reviewer of ICML.**
2022-2024 **Reviewer of ICLR.**
2023 **Reviewer of AISTATS.**

Invited Talk

- 2022: **Learning Multi-Objective Curricula for Robotics Policy Learning**, Represent for Ph.D. Student in 50th anniversary of the McGill School of Computer Science.