

# Minsu Kim

minsu.kim@mila.quebec | min-su@kaist.ac.kr

## RESEARCH EXPERTISE

---

- Post-training & controllable-generation of Diffusion Model/LLM/LMM
- Structured Probabilistic Inference with DRL
- Safe and trustworthy AI
- Combinatorial Search/Optimization with DRL
- Biological and Chemical Discovery using Generative Models

## WORK EXPERIENCE

---

### Postdoctoral Fellow

*Mila - Quebec AI Institute*

Host: Prof. Yoshua Bengio

- Search-based trustworthy reasoning in LLM
- Supported by CIFAR Safe AI Catalyst Fellowship via Yoshua Bengio

6/2025 – Current

*Canada-Korea Hybrid*

### Research Intern

*Mila - Quebec AI Institute*

Supervisor: Prof. Yoshua Bengio

- Search-based trustworthy reasoning in LLM
- Controlling generative models using latent diffusion sampling

1/2025 – 6/2025

*Canada-Korea Hybrid*

### Postdoctoral Fellow

*KAIST*

Host: Prof. Sungjin Ahn, Prof. Sungsoo Ahn

- KAIST-Mila Prefrontal Research Center collaboration

3/2025 – Current

*Canada-Korea Hybrid*

### Research Intern

*Mila - Quebec AI Institute*

Supervisor: Prof. Yoshua Bengio

- Generative Flow Networks (GFlowNets) for LLM/LMM fine-tuning
- GFlowNets for biological and chemical discovery
- Off-policy training of diffusion samplers

12/2023 – 5/2024

*Montreal, Canada*

## EDUCATION

---

### Doctor of Philosophy (Ph.D.)

*Korea Advanced Institute of Science and Technology*

- Major: Industrial and System Engineering
- Advised by Prof. Jinkyoo Park
- Thesis: Off-policy training methods for probabilistic agent in combinatorial space  
(*KAIST Presidential Best Ph.D. Thesis Award*)

3/2022 – 2/2025

*Daejeon, South Korea*

### Master of Science (M.S.)

*Korea Advanced Institute of Science and Technology*

- Major: Electronic and Electrical Engineering
- Advised by Prof. Joungho Kim
- Thesis: Novel reinforcement learning methods for routing problems on discrete space,  
Two-DIMM-per-Channel (2DPC) and PAM-4 interconnection

3/2020 – 2/2022

*Daejeon, South Korea*

- Major: Math and Computer Science (Dual Degree)
- Undergraduate research participation (URP) under supervision of Prof. Jinwoo Shin at 2019/06 - 2019/12, in *Graduate School of AI, KAIST*
- Thesis: Neural local search for travelling salesman problem

## NEW WORKS (\*: EQUAL CONTRIBUTION)

- 
- [1] **Minsu Kim\***, Jean-Pierre Falet\*, Oliver E. Richardson, Xiaoyin Chen, Moksh Jain, Sungjin Ahn, Sungsoo Ahn, and Yoshua Bengio. “Search-Based Correction of Reasoning Chains for Language Models”. In: *Under review at NeurIPS* (2025). URL: <https://arxiv.org/abs/2505.11824>.
  - [2] Minkyu Kim, Kiyoun Seong, Dongyeop Woo, Sungsoo Ahn, and **Minsu Kim**. “On scalable and efficient training of diffusion samplers”. In: *Under review at NeurIPS* (2025). URL: <https://arxiv.org/abs/2505.19552>.
  - [3] Xiaoyin Chen, Jiarui Lu, **Minsu Kim**, Dinghuai Zhang, Jian Tang, Alexandre Piché, Nicolas Gontier, Yoshua Bengio, and Ehsan Kamalloo. “Self-Evolving Curriculum for LLM Reasoning”. In: *Under review at NeurIPS* (2025). URL: <https://arxiv.org/abs/2505.14970>.
  - [4] Brian R Bartoldson, Siddharth Venkatraman, James Diffenderfer, Moksh Jain, Tal Ben-Nun, Seanie Lee, **Minsu Kim**, Johan Obando-Ceron, Yoshua Bengio, and Bhavya Kailkhura. “Trajectory Balance with Asynchrony: Decoupling Exploration and Learning for Fast, Scalable LLM Post-Training”. In: *Under review at NeurIPS* (2025). URL: <https://arxiv.org/abs/2503.18929>.
  - [5] Gyubin Lee, Bao N Nguyen Truong, Jaesik Yoon, Dongwoo Lee, **Minsu Kim**, Yoshua Bengio, and Sungjin Ahn. “Adaptive Cyclic Diffusion for Inference Scaling”. In: *Under review at NeurIPS* (2025).
  - [6] Dongyeop Woo, **Minsu Kim**, Minkyu Kim, Kiyoun Seong, and Sungsoo Ahn. “Energy-based generator matching: A neural sampler for general state space”. In: *Under review at NeurIPS* (2025). URL: <https://arxiv.org/abs/2505.19646>.

## PUBLICATIONS (\*: EQUAL CONTRIBUTION)

- 
- [1] Hyeonah Kim, **Minsu Kim**, Taeyoung Yun, Sanghyeok Choi, Emmanuel Bengio, Alex Hernández-García, and Jinkyoo Park. “Improved Off-policy Reinforcement Learning in Biological Sequence Design”. In: *International Conference on Machine Learning (ICML)* (2025).
  - [2] Siddharth Venkatraman, Mohsin Hasan, **Minsu Kim**, Luca Scimeca, Marcin Sendera, Yoshua Bengio, Glen Berseth, and Nikolay Malkin. “Outsourced diffusion sampling: Efficient posterior inference in latent spaces of generative models”. In: *International Conference on Machine Learning (ICML)* (2025).
  - [3] **Minsu Kim\***, Sanghyeok Choi\*, Taeyoung Yun, Emmanuel Bengio, Leo Feng, Jarrod Rector-Brooks, Sungsoo Ahn, Jinkyoo Park, Nikolay Malkin, and Yoshua Bengio. “Adaptive teachers for amortized samplers”. In: *International Conference on Learning Representation (ICLR)* (2025).
  - [4] **Minsu Kim\***, Sanghyeok Choi\*, Jiwoo Son, Hyeonah Kim, Jinkyoo Park, and Yoshua Bengio. “Ant Colony Sampling with GFlowNets for Combinatorial Optimization”. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)* (2025).
  - [5] Seanie Lee, **Minsu Kim**, Lynn Cherif, David Dobre, Juho Lee, Sung Ju Hwang, Kenji Kawaguchi, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, et al. “Learning Diverse Attacks on Large Language Models for Robust Red-teaming and Safety Tuning”. In: *International Conference on Learning Representation (ICLR)* (2025).

- [6] Seonghwan Seo, **Minsu Kim**, Tony Shen, Martin Ester, Jinkyoo Park, Sungsoo Ahn, and Woo Youn Kim. “Generative Flows on Synthetic Pathway for Drug Design”. In: *International Conference on Learning Representation (ICLR)* (2025).
- [7] Siddarth Venkatraman\*, Moksh Jain\*, Luca Scimeca\*, **Minsu Kim\***, Marcin Sendera\*, Mohsin Hasan, Luke Rowe, Sarthak Mittal, Pablo Lemos, Emmanuel Bengio, et al. “Amortizing Intractable Inference in Diffusion Models for Vision, Language, and Control”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [8] Nayoung Kim, Seongsu Kim, **Minsu Kim**, Jinkyoo Park, and Sungsoo Ahn. “MOFFlow: Flow Matching for Structure Prediction of Metal-Organic Frameworks”. In: *Under review at ICLR* (2024).
- [9] Nayoung Kim, **Minsu Kim**, Sungsoo Ahn, and Jinkyoo Park. “Decoupled Sequence and Structure Generation for Realistic Antibody Design”. In: *Transactions on Machine Learning Research (TMLR)* (2024).
- [10] Hyeonah Kim, **Minsu Kim**, Sanghyeok Choi, and Jinkyoo Park. “Genetic-guided GFlowNets for Sample Efficient Molecular Optimization”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [11] Marcin Sendera, **Minsu Kim**, Sarthak Mittal, Pablo Lamos, Luca Scimeca, Jarrod Rector-Brooks, Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. “Improved Off-policy Training of Diffusion Samplers”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [12] Hyosoon Jang, Yunhui Jang, **Minsu Kim**, Jinkyoo Park, and Sungsoo Ahn. “Pessimistic Backward Policy for GFlowNets”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [13] Hyeonah Kim, **Minsu Kim**, Sungsoo Ahn, and Jinkyoo Park. “Symmetric Replay Training: Enhancing Sample Efficiency in Deep Reinforcement Learning for Combinatorial Optimization”. In: *International Conference on Machine Learning (ICML)* (2024).
- [14] **Minsu Kim\***, Joohwan Ko\*, Taeyoung Yun\*, Dinghuai Zhang, Ling Pan, Taeyoung Yun, Woorchang Kim, Jinkyoo Park, Emmanuel Bengio, and Yoshua Bengio. “Learning to Scale Logits for Temperature-Conditional GFlowNets”. In: *International Conference on Machine Learning (ICML)* (2024).
- [15] **Minsu Kim**, Taeyoung Yun, Emmanuel Bengio, Dinghuai Zhang, Yoshua Bengio, Sungsoo Ahn, and Jinkyoo Park. “Local Search GFlowNets”. In: *International Conference on Learning Representation (ICLR), Spotlight Presentation* (2024).
- [16] Hyosoon Jang, **Minsu Kim**, and Sungsoo Ahn. “Learning Energy Decompositions for Partial Inference in GFlowNets”. In: *International Conference on Learning Representation (ICLR), Oral Presentation* (2024).
- [17] Jiwoo Son\*, **Minsu Kim\***, Sanghyeok Choi, Hyeonah Kim, and Jinkyoo Park. “Equity-Transformer: Solving NP-hard Min-max Routing Problems as Sequential Generation with Equity Context”. In: *AAAI Conference on AI (AAAI)* (2024).
- [18] **Minsu Kim**, Federico Berto, Sungsoo Ahn, and Jinkyoo Park. “Bootstrapped Training of Score-Conditioned Generator for Offline Design of Biological Sequences”. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2023).
- [19] Haeyeon Kim\*, **Minsu Kim\***, Federico Berto, Joungho Kim, and Jinkyoo Park. “DevFormer: A Symmetric Transformer for Context-Aware Device Placement”. In: *International Conference on Machine Learning (ICML)* (2023).

- [20] Jiwoo Son\*, **Minsu Kim\***, Hyeonah Kim, and Jinkyoo Park. “Meta-SAGE: Scale Meta-Learning Scheduled Adaptation with Guided Exploration for Mitigating Scale Shift on Combinatorial Optimization”. In: *International Conference on Machine Learning (ICML)* (2023).
- [21] Federico Berto\*, Chuanbo Hua\*, Junyoung Park\*, **Minsu Kim**, Hyeonah Kim, Jiwoo Son, Haeyeon Kim, Joungho Kim, and Jinkyoo Park. “RL4CO: an Extensive Reinforcement Learning for Combinatorial Optimization Benchmark”. In: *NeurIPS 2023 Workshop: New Frontiers in Graph Learning, Oral Presentation* (2023).
- [22] **Minsu Kim**, Junyoung Park, and Jinkyoo Park. “Sym-nco: Leveraging symmetricity for neural combinatorial optimization”. In: *Advances in Neural Information Processing Systems (NeurIPS)* 35 (2022), pp. 1936–1949.
- [23] Keeyoung Son, Seongguk Kim, Hyunwook Park, Taein Shin, Keunwoo Kim, **Minsu Kim**, Boogyo Sim, Subin Kim, Gapyeol Park, Shinyoung Park, et al. “Thermal and Signal Integrity Co-Design and Verification of Embedded Cooling Structure With Thermal Transmission Line for High Bandwidth Memory Module”. In: *IEEE Transactions on Components, Packaging and Manufacturing Technology* 12.9 (2022), pp. 1542–1556.
- [24] Hyeonah Kim, **Minsu Kim**, Changhyun Kwon, and Jinkyoo Park. “Neural Coarsening Process for Multi-level Graph Combinatorial Optimization”. In: *NeurIPS 2022 Workshop: New Frontiers in Graph Learning*. 2022.
- [25] Hyunwook Park, **Minsu Kim**, Seongguk Kim, Keunwoo Kim, Haeyeon Kim, Taein Shin, Keeyoung Son, Boogyo Sim, Subin Kim, Seungtaek Jeong, et al. “Transformer network-based reinforcement learning method for power distribution network (PDN) optimization of high bandwidth memory (HBM)”. In: *IEEE Transactions on Microwave Theory and Techniques* 70.11 (2022), pp. 4772–4786.
- [26] **Minsu Kim**, Jinkyoo Park, and Joungho Kim. “Learning collaborative policies to solve NP-hard routing problems”. In: *Advances in Neural Information Processing Systems (NeurIPS)* 34 (2021), pp. 10418–10430.
- [27] **Minsu Kim**, Hyunwook Park, Keeyoung Son, Seongguk Kim, Haeyeon Kim, Jihun Kim, Jinwook Song, Youngmin Ku, Jounghyu Park, and Joungho Kim. “Imitation Learning for Simultaneous Escape Routing”. In: *2021 IEEE 30th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2021, pp. 1–3.
- [28] Kyungjune Son, **Minsu Kim**, Hyunwook Park, Daehwan Lho, Keeyoung Son, Keunwoo Kim, Seongsoo Lee, Seungtaek Jeong, Shinyoung Park, Seokwoo Hong, et al. “Reinforcement-learning-based signal integrity optimization and analysis of a scalable 3-d x-point array structure”. In: *IEEE Transactions on Components, Packaging and Manufacturing Technology* 12.1 (2021), pp. 100–110.
- [29] Haeyeon Kim, Hyunwook Park, **Minsu Kim**, Seonguk Choi, Jihun Kim, Joonsang Park, Seongguk Kim, Subin Kim, and Joungho Kim. “Deep reinforcement learning framework for optimal decoupling capacitor placement on general PDN with an arbitrary probing port”. In: *2021 IEEE 30th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2021, pp. 1–3.
- [30] Seonguk Choi, **Minsu Kim**, Hyunwook Park, Keeyoung Son, Seongguk Kim, Jihun Kim, Joonsang Park, Haeyeon Kim, Taein Shin, Keunwoo Kim, et al. “Sequential Policy Network-based Optimal Passive Equalizer Design for an Arbitrary Channel of High Bandwidth Memory using Advantage Actor Critic”. In: *2021 IEEE 30th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2021, pp. 1–3.
- [31] Keeyoung Son, Seongguk Kim, **Minsu Kim**, Daehwan Lho, Keunwoo Kim, Hyunwook Park, Gapyeol Park, and Joungho Kim. “Signal integrity analysis of high speed channel considering thermal

- distribution”. In: *2021 IEEE 30th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2021, pp. 1–3.
- [32] Hyunwook Park, Jihun Kim, **Minsu Kim**, Keunwoo Kim, Boogyo Sim, Daehwan Lho, Taein Shin, Keeyoung Son, Jinwook Song, Youngmin Ku, et al. “Crosstalk-included PAM-4 Worst Eye Diagram Estimation Method for High-speed Serial Links”. In: *2021 IEEE 30th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2021, pp. 1–3.
- [33] Joonsang Park, **Minsu Kim**, Seongguk Kim, Keeyoung Son, Taein Shin, Hyunwook Park, Jihun Kim, Seonguk Choi, Haeyeon Kim, Keunwoo Kim, et al. “Deep Reinforcement Learning-based Pin Assignment Optimization of BGA Packages considering Signal Integrity with Graph Representation”. In: *2021 IEEE 30th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2021, pp. 1–3.
- [34] **Minsu Kim**, Hyunwook Park, Seongguk Kim, Keeyoung Son, Subin Kim, Kyunjune Son, Seonguk Choi, Gapyeol Park, and Joungcho Kim. “Reinforcement learning-based auto-router considering signal integrity”. In: *2020 IEEE 29th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2020, pp. 1–3.
- [35] Jihun Kim, Hyunwook Park, **Minsu Kim**, Seongguk Kim, Seonguk Choi, Keeyoung Son, Joonsang Park, Haeyeon Kim, Jinwook Song, Youngmin Ku, et al. “PAM-4 based PCIe 6.0 Channel Design Optimization Method using Bayesian Optimization”. In: *2021 IEEE 30th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE. 2021, pp. 1–3.
- [36] Daehwan Lho, Hyunwook Park, Seongguk Kim, Taein Shin, Keunwoo Kim, Kyungjune Son, Hyungmin Kang, Boogyo Sim, Keeyoung Son, **Minsu Kim**, et al. “Deep Neural Network-based Lumped Circuit Modeling using Impedance Curve”. In: *2020 IEEE Electrical Design of Advanced Packaging and Systems (EDAPS)*. IEEE. 2020, pp. 1–3.
- [37] Keunwoo Kim, Hyunwook Park, Daehwan Lho, **Minsu Kim**, Keeyoung Son, Kyungjune Son, Seongguk Kim, Taein Shin, Seonguk Choi, and Joungcho Kim. “Deep reinforcement learning-based through silicon via (TSV) array design optimization method considering crosstalk”. In: *2020 IEEE Electrical Design of Advanced Packaging and Systems (EDAPS)*. IEEE. 2020, pp. 1–3.

## HONORS & AWARDS

---

<b>KAIST Jang Yeong Sil Fellowship Award</b> <i>KAIST</i> A prestigious award from KAIST recognizing early-career postdoctoral researchers with outstanding potential.	4/2025 <i>Daejeon, South Korea</i>
<b>KAIST Presidential Best Ph.D. Thesis Award</b> <i>KAIST</i> Thesis: “Off-policy Training Methods for Probabilistic Agent in Combinatorial Space”	2/2025 <i>Daejeon, South Korea</i>
<b>Google Conference Scholarship</b> <i>International Conference on Learning Representations 2024</i> Paper: “Local Search GFlowNets”	5/2024 <i>Vienna, Austria</i>
<b>Qualcomm Innovative Fellowship Award</b> <i>Qualcomm Innovative Fellowship Korea (QIFK) 2023</i> Paper: “Sym-NCO: Leveraging Symmetricity for Neural Combinatorial Optimization”	11/2023 <i>Seoul, South Korea</i>
<b>Best Paper Award (coauthor)</b> <i>DesignCon 2022</i> Paper: “Deep Reinforcement Learning-based Channel-flexible Equalization Scheme: An Application to High Bandwidth Memory”	4/2023 <i>Santa Clara, United States</i>
<b>Best Paper Award (coauthor)</b> <i>DesignCon 2022</i> Paper: “Imitate Expert Policy and Learn Beyond: A Practical PDN Optimizer by Imitation Learning”	4/2023 <i>Santa Clara, United States</i>
<b>Scholar Award</b> <i>Neural Information Process System 2022</i>	12/2022 <i>New Orleans, United States</i>
<b>Best Paper Award (1st author)</b> <i>DesignCon 2021</i> Paper: “Neural Language Model Enables Extremely Fast and Robust Routing on Interposer”	4/2022 <i>Santa Clara, United States</i>
<b>Best Student Paper Award (coauthor)</b> <i>IEEE Electrical Design of Advanced Packaging and Systems 2020</i> Paper: “Deep Reinforcement Learning-based Interconnection Design for 3D X-Point Array Structure Considering Signal Integrity”	4/2022 <i>Remote</i>

## ACADEMIC REVIEWER

---

Conference on Neural Information Processing System (NeurIPS)	2022 – 2025
International Conference on Machine Learning (ICML)	2023 – 2025
International Conference on Learning Representation (ICLR)	2025
IEEE Transactions on Neural Networks and Learning Systems (TNNLS)	2024
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)	2024
International Conference on Artificial Intelligence and Statistics (AISTATS)	2023 - 2024

AAAI Conference on AI (AAAI)	2025
International Joint Conference on Artificial Intelligence (IJCAI)	2024 - 2025
Learning on Graphs Conference (LoG)	2023 – 2024
ICML workshop on structured probabilistic inference and generative modeling (SPIGM)	2023 – 2024
ICLR workshop on Frontiers in Probabilistic Inference (FPI)	2025

## TALKS

<b>“Search-based Corrections of Reasoning Chains for LMs”</b>	5/2025
<i>Rising Star Invited Talk at MBZUAI Trustworthy LLM Workshop</i>	<i>Abudabi, UAE</i>
<b>“Enhancing RL for Generative Models”</b>	11/2024
<i>KAIST-Mila Prefrontal Research Center</i>	<i>Remote</i>
<b>“Amortizing Intractable Inference in Diffusion model”</b>	5/2024
<i>Mila - Quebec AI Institute</i>	<i>Montreal, Canada</i>
<b>“On Diffusion Models for Amortized Inference”</b>	2/2024
<i>Mila - Quebec AI Institute</i>	<i>Montreal, Canada</i>
<b>“Symmetric Neural Combinatorial Optimization”</b>	11/2023
<i>Qualcomm AI Research (Fellowship Finalist Talk)</i>	<i>Seoul, South Korea</i>
<b>“Recent Trends for Generative Flow Networks”</b>	11/2023
<i>POSTECH Graduate School of AI, ML lab</i>	<i>Pohang, South Korea</i>
<b>“ML-based Offline Design Method for Biological Sequences”</b>	11/2023
<i>Samsung AI Forum (Invited Poster Session)</i>	<i>Suwon, South Korea</i>
<b>“Local Search GFlowNets”</b>	10/2023
<i>Mila - Quebec AI Institute</i>	<i>Remote</i>
<b>“Deep Learning for Combinatorial Optimization”</b>	8/2022
<i>KAIST Graduate School of Data Science</i>	<i>Daejeon, South Korea</i>
<b>“DRL Application for 2DPC and PAM4 Interconnection”</b>	8/2021
<i>NARA Institute of Science and Technology (NAIST)</i>	<i>Remote</i>
<b>“DRL for Multi-net Routing at Chip-to-chip interconnection design”</b>	4/2021
<i>Samsung Electronics</i>	<i>Remote</i>
<b>“DRL-based Auto-router for 3D Integration”</b>	7/2020
<i>Samsung Electronics</i>	<i>Remote</i>
<b>“Interposer Routing by Deep Reinforcement Learning”</b>	5/2020
<i>LG Electronics</i>	<i>Seoul, South Korea</i>

## REFERENCES

---

Yoshua Bengio

*A.M Turing Award Recipient* for founding of “Deep Learning”, *AAAI Fellow*

Full Professor of Department of Computer Science and Operations Research

University of Montreal

yoshua.bengio@mila.quebec

Jinkyoo Park

Associate Professor of Industrial System Engineering and Graduate School of AI

Korea Advanced Institute of Science and Technology (KAIST)

jinkyoo.park@kaist.ac.kr

Joungho Kim

*IEEE Fellow* for the contribution of “Modeling signal and power integrity in 3D integrated circuits.”

Full Professor of Electrical Engineering

Korea Advanced Institute of Science and Technology (KAIST)

joungho@kaist.ac.kr