

# MINKI KANG

<https://nardien.github.io> ◊ [minkikang@kaist.ac.kr](mailto:minkikang@kaist.ac.kr) ◊ [linkedin.com/in/minki-kang-5aa7281bb](https://www.linkedin.com/in/minki-kang-5aa7281bb)

## EDUCATION

---

- KAIST**, Seoul, South Korea *Sep. 2021 - Feb. 2027 (expected)*  
Ph.D., Kim Jaechul Graduate School of AI, Advisor: Prof. [Sung Ju Hwang](#)
- KAIST**, Seoul, South Korea *Mar. 2020 - Aug. 2021*  
M.A., Kim Jaechul Graduate School of AI, Advisor: Prof. [Sung Ju Hwang](#)  
**Thesis:** Language model adaptation with neural mask generator and domain knowledge graph
- Technical University of Munich**, Munich, Germany *Mar. 2019 - Aug. 2019*  
Exchange Student, Electrical Engineering and Information Technology
- KAIST**, Daejeon, South Korea *Mar. 2015 - Feb. 2020*  
B.S., Electrical Engineering and Computer Science (double major)

## RESEARCH INTERESTS

---

My primary research interests lie in the development of efficient and specialized small language models (sLMs) that can perform expert-level problem solving under limited computational resources. I focus on model distillation [8, 12, 17], retrieval-augmented generation [1, 7, 12, 13], and tool-augmented reasoning [2]. A recent direction involves agent distillation [19], where agentic behaviors such as reasoning, decision-making, and tool use are distilled from LLM-based agents into smaller models having fewer parameters.

I am also interested in model specialization through domain adaptation [3, 7, 12] and knowledge injection [14] to equip models with domain-specific knowledge without extensive retraining. My work further includes graph representation learning [4, 6] and text-to-speech synthesis (TTS) [10, 11, 15, 16].

## CONFERENCE PUBLICATIONS

---

- [22] [SAGE: Shaping Anchors for Guided Exploration in RLVR of LLMs](#)  
Chanuk Lee, **Minki Kang**, Sung Ju Hwang  
International Conference on Machine Learning (**ICML**) 2026
- [21] [ACON: Optimizing Context Compression for Long-horizon LLM Agents](#)  
**Minki Kang**, Wei-Ning Chen, Dongge Han, Huseyin A. Inan, Lukas Wutschitz, Yanzhi Chen, Robert Sim, Saravan Rajmohan  
International Conference on Machine Learning (**ICML**) 2026  
Lifelong Agents (LLA) Workshop @ **ICLR** 2026
- [20] [T1: Tool-integrated Self-verification for Test-time Compute Scaling in Small Language Models](#)  
**Minki Kang**<sup>\*</sup>, Jongwon Jeong<sup>\*</sup>, Jaewoong Cho  
International Conference on Learning Representation (**ICLR**) 2026
- [19] [Distilling LLM Agent into Small Models with Retrieval and Code Tools](#)  
**Minki Kang**, Jongwon Jeong, Seanie Lee, Jaewoong Cho, Sung Ju Hwang  
Neural Information Processing System (NeurIPS) 2025, **Spotlight** (Top 3%)
- [18] [SafeRoute: Adaptive Model Selection for Efficient and Accurate Safety Guardrails in Large Language Models](#)  
Seanie Lee<sup>\*</sup>, Dong Bok Lee<sup>\*</sup>, Dominik Wagner, **Minki Kang**, Haebin Seong, Tobias Bocklet, Juho Lee, and Sung Ju Hwang  
Annual Meeting of the Association for Computational Linguistics (ACL) 2025 Findings
- [17] [HarmAug: Effective Data Augmentation for Knowledge Distillation of Safety Guard Models](#)  
Seanie Lee<sup>\*</sup>, Haebin Seong<sup>\*</sup>, Dong Bok Lee, **Minki Kang**, Xiaoyin Chen, Dominik Wagner, Yoshua Bengio, Juho Lee, and Sung Ju Hwang  
International Conference on Learning Representation (**ICLR**) 2025

- [16] [Stable-TTS: Stable Speaker-Adaptive Text-to-Speech Synthesis via Prosody Prompting](#)  
Wooseok Han\*, **Minki Kang\***, Changhun Kim, and Eunho Yang  
IEEE International Conference on Acoustics, Speech and Signal Processing (**ICASSP**) 2025
- [15] [Face-StyleSpeech: Enhancing Zero-shot Speech Synthesis from Face Images with Improved Face-to-Speech Mapping](#)  
**Minki Kang\***, Wooseok Han\*, and Eunho Yang  
IEEE International Conference on Acoustics, Speech and Signal Processing (**ICASSP**) 2025
- [14] [Latent Paraphrasing: Perturbation on Layers Improves Knowledge Injection in Large Language Models](#)  
**Minki Kang**, Sung Ju Hwang, Gibbeum Lee, and Jaewoong Cho  
Neural Information Processing System (**NeurIPS**) 2024
- [13] [Knowledge-Augmented Language Model Verification](#)  
Jinheon Baek, Soyeong Jeong, **Minki Kang**, Jong C. Park, and Sung Ju Hwang  
Conference on Empirical Methods in Natural Language Processing (**EMNLP**) 2023
- [12] [Knowledge-Augmented Reasoning Distillation for Small Language Models in Knowledge-Intensive Tasks](#)  
**Minki Kang**, Seanie Lee, Jinheon Baek, Kenji Kawaguchi, and Sung Ju Hwang  
Neural Information Processing System (**NeurIPS**) 2023
- [11] [ZET-Speech: Zero-shot adaptive Emotion-controllable Text-to-Speech Synthesis with Diffusion and Style-based Models](#)  
**Minki Kang\***, Wooseok Han\*, Sung Ju Hwang, and Eunho Yang  
**InterSpeech** 2023
- [10] [Grad-StyleSpeech: Any-speaker Adaptive Text-To-Speech Synthesis with Diffusion Models](#)  
**Minki Kang\***, Dongchan Min\*, and Sung Ju Hwang  
IEEE International Conference on Acoustics, Speech and Signal Processing (**ICASSP**) 2023
- [9] [Sparse Token Transformers with Attention Back Tracking](#)  
Heejun Lee, **Minki Kang**, Youngwan Lee, and Sung Ju Hwang  
International Conference on Learning Representation (**ICLR**) 2023
- [8] [Self-Distillation for Further Pre-training of Transformers](#)  
Seanie Lee, **Minki Kang**, Juho Lee, Sung Ju Hwang, and Kenji Kawaguchi  
International Conference on Learning Representation (**ICLR**) 2023
- [7] [KALA: Knowledge-Augmented Language Model Adaptation](#)  
**Minki Kang\***, Jinheon Baek\*, and Sung Ju Hwang  
Annual Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL**) 2022  
(oral presentation)
- [6] [Edge Representation Learning with Hypergraphs](#)  
Jaehyeong Jo\*, Jinheon Baek\*, Seul Lee\*, Dongki Kim, **Minki Kang**, and Sung Ju Hwang  
Neural Information Processing System (**NeurIPS**) 2021
- [5] [Learning to Perturb Word Embeddings for Out-of-distribution QA](#)  
Seanie Lee\*, **Minki Kang\***, Juho Lee, and Sung Ju Hwang  
Annual Meeting of the Association for Computational Linguistics (**ACL**) 2021
- [4] [Accurate Learning of Graph Representations with Graph Multiset Pooling](#)  
Jinheon Baek\*, **Minki Kang\***, and Sung Ju Hwang  
International Conference on Learning Representation (**ICLR**) 2021
- [3] [Neural Mask Generator: Learning to Generate Adaptive Word Maskings for Language Model Adaptation](#)  
**Minki Kang\***, Moonsu Han\*, and Sung Ju Hwang  
Conference on Empirical Methods in Natural Language Processing (**EMNLP**) 2020
- [2] [Episodic Memory Reader: Learning What to Remember for Question Answering from Streaming Data](#)  
Moonsu Han\*, **Minki Kang\***, Hyunwoo Jung, and Sung Ju Hwang  
Annual Meeting of the Association for Computational Linguistics (**ACL**) 2019 (oral presentation)
- (\*: equal contribution)

## WORKSHOP PUBLICATIONS & PREPRINTS

---

- [4] [Memory Transfer Learning: How Memories are Transferred Across Domains in Coding Agents](#)  
Kangsan Kim, **Minki Kang**, Taeil Kim, Yanlai Yang, Mengye Ren, Sung Ju Hwang  
Preprint
- [3] [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  
Preprint
- [2] [Rethinking Reward Models for Multi-Domain Test-Time Scaling](#)  
Dong Bok Lee\*, Seanie Lee\*, Sangwoo Park, **Minki Kang**, Jinheon Baek, Dongki Kim, Dominik Wagner, Jiongdiao Jin, Heejun Lee, Tobias Bocklet, Jinyu Wang, Jingjing Fu, Sung Ju Hwang, Jiang Bian, Lei Song  
Preprint
- [1] [Knowledge-Consistent Dialogue Generation with Knowledge Graphs](#)  
**Minki Kang**\*, Jin Myung Kwak\*, Jinheon Baek\*, and Sung Ju Hwang  
Knowledge Retrieval and Language Models (KRLM) Workshop @ ICML 2022  
(\*: equal contribution)

## EXPERIENCE

---

<b>NVIDIA</b> , Remote <i>Research Intern (Mentor: <a href="#">Byung-Kwan Lee</a>) / NVIDIA Research Taiwan</i>	<i>Feb. 2026 -</i>
<b>Microsoft</b> , Redmond, US <i>Research Intern (Mentor: <a href="#">Wei-Ning Chen</a>) / Microsoft 365 Research Team</i>	<i>Jun. 2025 - Sep. 2025</i>
<b>KRAFTON</b> , Seoul, South Korea <i>Part-time Researcher / Natural Language DL Team</i>	<i>Mar. 2025 - Jun. 2025</i>
<b>KRAFTON</b> , Seoul, South Korea <i>Applied Research Scientist (Full time regular as a technical research personnel) / Natural Language DL Team</i>	<i>Oct. 2023 - Mar. 2025</i>
<b>AITRICS</b> , Seoul, South Korea <i>Applied Research Scientist (Full time regular as a technical research personnel) / Virtual Human Team</i>	<i>Mar. 2022 - Oct. 2023</i>
<b>Machine Learning and Artificial Intelligence Lab</b> , KAIST <i>Research Intern (Advisor: Prof. <a href="#">Sung Ju Hwang</a>)</i>	<i>Mar. 2018 - Feb. 2020</i>
<b>Kakao</b> , Jeju, South Korea <i>Engineering Intern / Context Department (NLP)</i>	<i>Jan. 2018 - Feb. 2018</i>
<b>Intelligent Network Architecture and Distributed System Lab</b> , KAIST <i>Research Intern (Advisor: Prof. <a href="#">Dongsu Han</a>)</i>	<i>Jul. 2017 - Nov. 2017</i>

## PATENTS

---

- [2] **METHOD AND APPARATUS FOR LANGUAGE MODEL VERIFICATION**  
Sung Ju Hwang, Jinheon Baek, Soyeong Jeong, **Minki Kang**, and Jong C. Park  
Korean patent number: 10-2024-0036737, filed on Mar 15, 2024
- [1] **MEMORY-BASED REINFORCEMENT LEARNING METHOD FOR STORING OPTIONAL INFORMATION IN STREAMING DATA AND SYSTEM THEREFORE**  
Sung Ju Hwang, Moonso Han, **Minki Kang**, and Hyunwoo Jung  
US Patent: 11,875,119, filed on Feb 18, 2021

## SERVICES

---

## Reviewer

Annual Meeting of the Association for Computational Linguistics (ACL)	2022, 2023, 2024, 2025
International Conference on Learning Representation (ICLR)	2022, 2023, 2024, 2025, 2026
Neural Information Processing System (NeurIPS, NIPS)	2021, 2022, 2023, 2024, 2025
Conference on Empirical Methods in Natural Language Processing (EMNLP)	2021, 2022, 2023, 2024
International Conference on Machine Learning (ICML)	2021, 2022, 2023, 2024, 2025, 2026

## INVITED TALKS

---

<b>Amazon Science Seminar</b> , Amazon	<i>Mar. 2026</i>
Invited to present as the author of an accepted paper at NeurIPS 2025	
<b>AI Workshop 21/22</b> , KAIST	<i>Jan. 2022</i>
Participated in the poster presentation session with three papers [4, 5, 6].	
<b>Online Tech Talk</b> , Naver	<i>Dec. 2020</i>
Invited to present as the author of an accepted paper at EMNLP 2020, where I discussed the paper [3].	
<b>AI Tech Talk for NLU</b> , LG CNS	<i>Sep. 2019</i>
Invited to present as the author of an accepted paper at ACL 2019, where I discussed the paper [2].	