Junhyuk So

Google Scholar

Email: junhyukso@postech.ac.kr
Mobile: +8210-2845-4886

Education

- **M.S/Ph.D - POSTECH (Advisor : Eunhyeok Park)** — Pohang, Korea
  Department of Computer Science and Enginnering — *2022 - 2027.02(Expected)*
- **B.S - University Of Seoul** — Seoul, Korea
  Department of Electrical and Computer Engineering — *2018 - 2022*

Research Interests

- **Efficient AI** Quantization, Speculative Decoding, Parallel Algorithm
- **Generative Models** Diffusion Models, Multimodal LLMs
- **Optimization** Numerical Methods, Sampling algorithms

Publiciations

[·] ML/AI top-tier conference, * co-first author

Under Review

- **Junhyuk So**, ... "Maximal Coupling Speculative Jacobi Decoding for Autoregressive Image Generation" (About training-free, plug and play, ∼10× loseless acceleration method for VLMs.)

Conference

2025

- [NeurIPS25] **Junhyuk So**, Chiwoong Lee, Shinyoung Lee, Jungseul Ok, Eunhyeok Park, "Improving Generative Behavior Cloning via Self-Guidance and Adaptive Chunking", **NeurIPS** 39th Conference on Neural Information Processing Systems, Dec, 2025
- [ICCV25] **Junhyuk So**, Junhyuk So, Juncheol Shin, Hyunho Kook, Eunhyeok Park, "Grouped Speculative Decoding for Autoregressive Image Generation", **ICCV** International Conference on Computer Vision, Oct, 2025
- [CVPR25] **Junhyuk So**, Jiwoong Shin, Chaeyeon Jang, Eunhyeok Park, "PCM : Picard Consistency Model for Fast Parallel Sampling of Diffusion Models", **CVPR** The IEEE/CVF Conference on Computer Vision and Pattern Recognition, June, 2025

2024

- [ECCV24] **Junhyuk So***, Jungwon Lee*, Eunhyeok Park, "FRDiff : Feature Reuse for Universal Training-free Acceleration of Diffusion Models", **ECCV** The 18th European Conference on Computer Vision , Oct, 2024

2023

- [NeurIPS23] **Junhyuk So***, Jungwon Lee*, Daehyun Ahn, Hyungjun Kim, Eunhyeok Park, "Temporal Dynamic Quantization for Diffusion Models" **NeurIPS** 37th Conference on Neural Information Processing Systems, Dec, 2023
- [NeurIPS23] **Junhyuk So***, Changdae Oh*, Yongtaek Lim, Hoyoon Byun, Minchul Shin, Kyungwoo Song, "Geodesic Multi-Modal Mixup for Robust Fine-Tuning" **NeurIPS** 37th Conference on Neural Information Processing Systems, Dec, 2023
- [CVPR23] Juncheol Shin*, **Junhyuk So***, Sein Park, Seungyeop Kang, Sungjoo Yoo, Eunhyeok Park, "NIPQ : Noise porxy based Integraded Psuedo Quantization " IEEE/CVF **CVPR** Computer Vision and Pattern Recognition Conference , June, 2023
- Eunchong Noh, **Junhyuk So**, Seunghwan Lee, "Machine-Learning based Optimal Design of a Wireless Power Transfer Coil for Battery-Powered Tram" ICPE 2023-**ECCE Asia** 11th International Conference on Power Electronics , May, 2023

2022

- Changdae Oh, Heeji Won, **Junhyuk So**, Taero Kim, Yewon Kim, Hosik Choi, Kyungwoo Song, "Learning Fair Representation via Distributional Contrastive Disentanglement" ACM **SIGKDD** International Conference on Knowledge Discovery Data Mining. July, 2022 (https://dl.acm.org/doi/10.1145/3534678.3539232)

2021

- [CODES21] Chanyoung Oh*, **Junhyuk So***, Sumin Kim*, Youngmin Yi, "Exploiting Activation Sparsity for Fast CNN Inference on Mobile GPUs" International Conference on Hardware/Software Codesign and System Synthesis **(CODES+ISSS)** Oct, 2021 (https://dl.acm.org/doi/abs/10.1145/3477008)

Journal [SCI]

2023

- **Junhyuk So**, Yongtaek Lim, Yewon Kim, Changdae Oh, Kyungwoo Song, "Robust Contrastive Learning With Dynamic Mixed Margin" **IEEE Access** June, 2023 (https://ieeexplore.ieee.org/abstract/document/10154052)

2021

- Chanyoung Oh*, **Junhyuk So***, Sumin Kim*, Youngmin Yi, "Exploiting Activation Sparsity for Fast CNN Inference on Mobile GPUs" ACM Transactions on Embedded Computing Systems (**TECS**) Oct, 2021 (https://dl.acm.org/doi/abs/10.1145/3477008)

Academic Services

- Reviewer : NeurIPS24-25, CVPR24-25, AAAI24-25, ICLR25-26

Talks

- Recent Topics on Image Generation Acceleration @ Squeezebits (June, 2025)

Skills

- Deep Learning, Machine Learning
- Pytorch, CUDA
- Python, C/C++, JAVA

Misc

- Homepage : https://junhyukso.github.io