

Suhwan Choi

📍 Seoul, South Korea ✉ milkclouds00@gmail.com 📄 Suhwan Choi 🌐 milkclouds.work 📱 milkclouds

Summary

I am an undergraduate student majoring in Physics and Computer Science at Seoul National University. My main area of interest is approximating and imitating human behavior in terms of vision, language, and action modalities, utilizing end-to-end architectures and scalable training suites.

Experience

Senior Researcher

Feb 2024 – Present

Maum.ai

- Working on development of human-like agent capable of robotic navigation.
- Leading the project [SketchDrive](#) [🔗](#).
- Led Slurm-based DGX Cluster construction project.
- Developed a Notion-based workspace to enhance productivity and streamline entire company processes.

Machine Learning Engineer Intern

July 2023 – Jan 2024

Hyperconnect

- Worked on diffusion-based personalized profile image generation for real-world applications. The output images were conditioned not only on prompts but also on the face in the input image.
- Responsible for the entire process from task definition to optimized model deployment.

Publications

CANVAS: Commonsense-Aware Navigation System for Intuitive Human-Robot Interaction

1st author

Suhwan Choi*, Yongjun Cho*, Minchan Kim*, Jaeyoon Jung*, Myunchul Joe, Yubeen Park, Minseo Kim, Sungwoong Kim, Sungjae Lee, Hwiseong Park, Jiwan Chung, Youngjae Yu

Workshop on Open-World Agents at NeurIPS 2024 (**Outstanding paper Awards, 3%**)

ESREAL: Exploiting Semantic Reconstruction to Mitigate Hallucinations in Vision-Language Models

2nd author

Minchan Kim*, Minyeong Kim*, Junik Bae*, Suhwan Choi, Sungkyung Kim, Buru Chang
ECCV 2024

Education

Seoul National University

2021 – Present

B.S in Physics & Computer Science and Engineering

Awards

Google Codejam 2022

2022

- Round 3, 546th (awarded T-Shirt).

NAVER CLOVA AI RUSH 2022 [\[link\]](#) [🔗](#)

July 2022 – Sept 2022

- **3rd place** on Landmark Detection (3,000,000 KRW)
 - Classical one-object detection problem.
 - Used mmdetection (mmdet). Since CLOVA uses custom ML platform named NSML, built custom Docker image for NSML and implemented hook on mmdet to integrate mmdet with NSML training/testing pipeline.
- **2nd place** on Shopping User Embedding Extraction, Classification (7,000,000 KRW)

- 1.2M unlabeled data, 2k labeled training data, 10-class time-series classification problem.
- MLM pretraining on 1.2M unlabeled data and augmentation on 2k labeled training data was the key.

QHack Coding Challenge [\[link\]](#) [↗](#)

2023 and 2024

- Ranked **4th/793** teams in 2023, Ranked **3rd/618** teams in 2024.
- This is a contest implementing quantum algorithms, quantum machine learning, quantum chemistry, and brain-teasing puzzles.

2023 Quantum Hackathon [\[link\]](#) [↗](#)

2023

- **1st place, Minister of Science and ICT Award**
- Topic: Utilizing symmetry to solve variational quantum algorithm (quantum machine learning) efficiently.

Activities

Member

Deepest: SNU Deep Learning Society

Sept 2022 – Present

- Joined Season 12 Project: DeepMind papers reading, neural ODE study.
- Joined Season 13 Project: [pylixir](#) [↗](#) (solving Lost Ark's minigame with RL).
- Joined Season 14 Project: Parallel computing study, causal inference study, personalized voice generation project.

Software Maestro Trainee

2021

- Studied 3D vision and implemented AR game for iPad using AR Foundation of Unity
- In-game image: <https://bit.ly/3Vxw8Ls> [↗](#). Zombie is evading sofa because iPad realized furnitures, and occlusion is also working.

Open Source Contributions

- [Simple warning popup to prevent convoluted failure](#) [↗](#) in [transformers](#) [↗](#) (Stars: 138k).
- [Implemented a simple feature](#) [↗](#) in [diffusers](#) [↗](#) (Stars: 27.2k), which enables `num_images_per_prompt>1` in community pipeline.
- [Fixed deadlock issue](#) [↗](#) on [mmdetection](#) [↗](#) (Stars: 30.1k) and [mmsegmentation](#) [↗](#) (Stars: 8.5k).
 - Debugged deadlock issue which only occurs in distributed (multi-GPU) environment. The detailed error analysis report is on [GitHub Issue Comment](#) [↗](#).
- [Implemented label visualization](#) [↗](#) on [mmdetection3d](#) [↗](#) (Stars: 5.5k).
 - Solved bugs related to overall mmdet3d visualization process and implemented label visualization on point cloud.

Technical Blogging

- Writing personal tech blog [milkclouds.work](#) [↗](#), reviewing papers on ML/DL domain.
- Reviewed papers like [DeiT](#) [↗](#), [Transformers](#) [↗](#), [DETR](#) [↗](#), [ViT](#) [↗](#), [ConvMixer](#) [↗](#), [Perceiver](#) [↗](#), [PointMLP](#) [↗](#), [InfoNeRF](#) [↗](#), [NCA](#) [↗](#), [AttentionNeuron](#) [↗](#), [FNO](#) [↗](#), and so on.

Skills

Programming Languages: Proficient in: Python, C++, JAVA / Can use: JavaScript, TypeScript, Rust, ...

Technologies: Git, Linux system, PyTorch, huggingface lib(transformers, accelerate, diffusers, ...), ROS, ansible, HPC infrastructure, ...

Languages: Korean (native), English