Seungyoon Woo

+82-10-2927-0370 | seungyoon.woo@vision.snu.ac.kr | South Korea

EDUCATION

Seoul National University Sep. 2022 – Present
PhD in Artificial Intelligence Seoul, South Korea

- **Advisor**: Prof. Gunhee Kim (Vision & Learning Laboratory)
- **Research Interest**: AI safety, alignment, continual learning, and scalable oversight for large language models

PUBLICATIONS

COMET: Decoupled Distillation, Routing, and Capacity Control for Task-Agnostic Continual Vision–Language Learning
Seungyoon Woo, Keonhee Park, Gunhee Kim **NeurIPS 2026**

- Proposed COMET, a modular continual learning method for vision-language models that handles distillation, routing, and capacity control as separate components to reduce catastrophic forgetting without task labels at inference

Meta-Continual Learning of Neural Fields
Seungyoon Woo, Junhyeog Yun, Gunhee Kim **ICLR 2025**

- Defined the meta-continual learning setting for neural fields and combined a modular architecture with optimization-based meta-learning to reduce forgetting and speed up convergence
- Proposed FIM-NeRF, a Fisher information loss for NeRF with convergence and generalization guarantees; evaluated on image, audio, video, and view synthesis tasks across six datasets, including city-scale NeRF

EXPERIENCE

Intel Feb. 2021 – Jun. 2021
Machine Learning Engineering Intern Seoul, South Korea

- Built a TFDV-style data validation module in Python that cut dataset error detection time by ∼40% in pipelines used by 10+ internal teams
- Contributed 8+ pull requests to Datumaro, Intel's open-source dataset management framework
- Attended deep learning conferences on behalf of the ML team and briefed engineering teams on relevant work

Naver Jan. 2020 – Feb. 2020
Deep Learning Engineering Intern Seongnam, South Korea

- Built a transformer-based classifier (PyTorch, HuggingFace) to verify the authenticity of 50K+ user reviews and packaged it in Docker, cutting environment setup time by ∼70%
- Compared 5+ methods for a 7:1 class imbalance, including SMOTE, class weighting, and focal loss, and presented the results to 10+ people on two teams
- Reached 87% F1 on sentiment classification of user-generated text, a 19% gain over the baseline

BIZ Jul. 2019 – Dec. 2019
Student Researcher Seoul, South Korea

- Led a team of four that built and deployed a deep learning chatbot for financial services questions on AWS (EC2, S3)
- Designed the intent classifier with OOD detection and confidence calibration, reaching 87%+ accuracy over 15+ financial question categories
- Built a recommendation component for ambiguous queries based on calibrated uncertainty estimates, cutting OOD failures by ∼30%

PROJECTS

Market Price Volatility Prediction

Software Maestro, The Federation of Korean Information Industries May 2020 – Dec. 2020

- Built a deep learning model that predicts price volatility and detects distribution shift in seller behavior data
- Deployed the training and serving pipeline on AWS and calibrated the model's confidence scores
- Gave technical presentations on the model architecture and monitoring plan

NLP Model Training and Serving

NAVER BoostCamp - AI, Naver Connect Foundation Jun. 2021 – Dec. 2021

- Built training and inference pipelines for NLP classification and tuned batching to raise GPU training throughput
- Ran fine-tuning and knowledge distillation experiments to trade off serving latency against accuracy
- Tested models on OOD benchmarks and measured calibration (ECE) before deployment

SKILLS

Languages: Python, C++
DL Frameworks: PyTorch, TensorFlow, HuggingFace, JAX, CUDA
Research Areas: AI Safety, Continual Learning, OOD Detection, Calibration, Multimodal Learning, Scalable Oversight
Methods: PPO, Fine-Tuning, Distillation, Prompt Engineering
Infrastructure: AWS, Docker, Git