Changyeon Lee

Email: cyeonlee@umd.edu

Research Interests

Multisensory Machine Learning / Embodied AI / Image · Video Understanding & Generation

Education

- **University of Maryland, College Park** — College Park, MD, United States
  M.S. in Computer Science (Advisor: Ruohan Gao) — *Sep 2025 - Present*
- **Yonsei University** — Seoul, Republic of Korea
  B.S. in Computer Science - Summa Cum Laude, Valedictorian — *Feb 2018 - Feb 2025*
  - Absence for mandatory military service (Aug 2020 - May 2022)
- **Chungnam Science High School** — Gongju, Chungnam, Republic of Korea
  Gifted School for Mathematics and Science - Early Honors Graduate — *Feb 2016 - Jan 2018*

Publications (* denotes equal contribution)

- **Deepfake-Eval-2024: A Multi-Modal In-the-Wild Benchmark of Deepfakes Circulated in 2024**
  Nuria Alina Chandra, Ryan Murtfeldt, Lin Qiu, Arnab Karmakar, Hannah Lee, Emmanuel Tanumihardja, Kevin Farhat, Ben Caffee, Sejin Paik, ***Changyeon Lee****, Jongwook Choi, Aerin Kim, Oren Etzioni*
  Under Review
- **Cost-Efficient and Effective Counter Unmanned Aerial System via Visual-Acoustic Sensing**
  Changyeon Lee*, Dongju Yu*, Soyeon Cho*, Dane W. Hindsley, Halaevalu F. Patterson, Megan A. Clecak, Eric T. Matson*
  IEEE International Conference on Robotic Computing 2024
- **DistilDIRE: A Small, Fast, Cheap and Light Diffusion Synthesized Deepfake Detection**
  Yewon Lim, **Changyeon Lee***, Aerin Kim, Oren Etzioni*
  ICML 2024 Workshop on Foundation Models in the Wild
- **The Tug-Of-War Between Deepfake Generation and Detection**
  Hannah Lee, ***Changyeon Lee****, Kevin Farhat, Lin Qiu, Steve Geluso, Aerin Kim, Oren Etzioni*
  ICML 2024 Workshop on Data-centric Machine Learning Research
- **Towards Interpretable Controllability in Object-Centric Learning**
  Jinwoo Kim, Janghyuk Choi*, Jaehyun Kang,* **Changyeon Lee****, Ho-Jin Choi, Seon Joo Kim*
  CVPR 2024 Workshop on Causal and Object-Centric Representations for Robotics

Work Experience

- **Multisensory Machine Intelligence Lab, University of Maryland** — College Park, MD, United States
  Graduate Researcher (Advisor: Ruohan Gao) — *Jun 2025 - Present*

- **Miraflow AI** — Kirkland, WA, United States
  Founding Machine Learning Engineer — *Mar 2024 - May 2025*
  - Developed an end-to-end lip-sync video generation service, automating the data collection and preprocessing of 3.7TB+ video-related data using optimized Bash scripting with multi-GPU and multiprocessing support on 8× A100 GPUs, achieving over 90% CPU/GPU utilization.
  - Modified model architecture for distributed training, resolving instability and increasing training speed by 1.5×.
  - Reduced lip-sync video inference time by 63.4% via multi-GPU/process parallelism, optimized computation scheduling, and low-latency operation replacements.
  - Implemented a robust, production-grade ML server using Sanic, enabling API-driven model training, inference, job control, and detailed status logging.
  - Designed a Kubernetes orchestration system for scalable ML workloads, reducing estimated cloud costs by up to 62.5% under typical workloads via efficient resource control.
  - Improved system reliability and user experience by handling erroneous data uploads and integrating user feedback (e.g., seamless job cancellation), in collaboration with front-end and back-end engineers.
  - Led research on efficient diffusion-based deepfake image detection (DistilDIRE) for TrueMedia.org, achieving 3.2× speedup via knowledge distillation without sacrificing accuracy.

- **IITP K-SW Square, Purdue University** West Lafayette, IN, United States
  Visiting Scholar (Advisor: Eric T. Matson) *Mar 2024 - Jun 2024*
  - Fully sponsored by the Ministry of Science and ICT of the Republic of Korea
  - Led research on a Counter-Unmanned Aerial System (CUAS), integrating audio-visual modalities for real-time drone detection and tracking.
  - Built robust audio classification and multi-output regression models leveraging inductive biases in architecture for efficient training on coarse, limited real-world data and low-latency inference on standard laptop hardware, achieving 90%+ accuracy in real-world tests.
  - Integrated a WebRTC-based data pipeline into the CUAS to stream audio and image data to models, enabling non-technical operators to access the platform on any device with a browser without additional software, with secure authentication via access keys.

- **Computational Intelligence & Photography Lab, Yonsei University** Seoul, Republic of Korea
  Undergraduate Researcher (Advisor: Seon Joo Kim) *Jul 2022 - Jun 2024*
  - Participated in *Deep View* Project (led by the Korean government) - Space-Time Video Super Resolution (STVSR).
    - Implemented a Transformer-based STVSR architecture enhanced with a MeMViT-inspired memory bank for attention key-value pairs across video clips, effectively expanding the temporal dimension and achieving comparable performance to CNN-based benchmarks while reducing model parameters by 75.3%.
  - Participated in 'Controllable Object Centric Learning' and 'Online Temporal Action Localization (On-TAL)' research.
  - Controllable Object-Centric Representation Learning (Controllable OCL)
    - Led a compositional image generation task to evaluate object-property binding in our OCL framework by extracting object representations (slots), manipulating their properties, and composing them onto a new canvas to generate novel scenes.
    - Synthesized datasets under varying conditions such as color, shape, and positions using the CLEVR generator to train and evaluate object property controllability in our proposed OCL framework.
    - Demonstrated the framework's scalability and robustness on the challenging ClevrTex benchmark by employing a MAE-pretrained ViT and an SRT decoder both scaled beyond those used in standard setups.
  - Online Temporal Action Localization
    - Led a user study via Amazon Mechanical Turk (AMT) by creating video samples using Matplotlib and FFmpeg, and developing the evaluation webpage on AMT to assess the effectiveness of our proposed model's action localization on untrimmed streaming videos, benchmarking against state-of-the-art methods through human evaluation.

- **Hyundai Development Company (HDC)** Seoul, Republic of Korea
  Data Scientist *May 2020 - Aug 2020*
  - Developed a multiple regression model using XGBoost to estimate apartment lotting-out prices from real transaction data, reducing MAPE by 5% relative to the baseline by designing a custom feature engineering pipeline leveraging XGBoost feature importance (drop-out, pick-up) and correlation analysis to optimize model inputs.

Honors & Scholarships

- Merit-Based Scholarship, Yonsei University (Spring 2019, Spring 2023, Fall 2023, Fall 2024)
- Merit-Based Scholarship, Jung-Hun Foundation (Fall 2023, Fall 2024)
- Highest Honors, Yonsei University (Fall 2023)
- Honors, Yonsei University (Spring 2018, Fall 2018)

Skills Summary

- Programming Languages: Python, C++, Java, SQL
- Frameworks & Libraries: PyTorch, TensorFlow, OpenCV, Scikit-learn, SciPy, XGBoost, Pymoo
- Tools & Technologies: Git, Docker, FFmpeg, Kubernetes, Oracle, PostgreSQL, MySQL