LINJI WANG
Fairfax, VA
LinkedIn
Email: joewwang@outlook.com
Mobile: 412-888-6071
linjiwang.com

EDUCATION

George Mason University — Fairfax, VA
PhD in Computer Science, Specialization in AI and Robotics — *Sep 2023 – May 2027 (Expected)*
- Research Focus: Generative AI and Reinforcement Learning for Robotic Systems

Carnegie Mellon University — Pittsburgh, PA
MSc in Mechanical Engineering — *Sep 2021 – May 2023*
- GPA: 3.94/4.0 (98.5%)
- Relevant Coursework: Machine Learning, Deep Learning, Computer Vision, Deep Reinforcement Learning & Control

RESEARCH EXPERIENCE

RobotiXX Lab, George Mason University — Fairfax, VA
Research Project: Grounded Curriculum Learning for Efficient Reinforcement Learning in Robotics — *Aug 2023 – Present*
- Engineered a novel dual-agent framework with a meta-learning-based teacher for adaptive curriculum in Embodied AI
- Implemented Variational Autoencoder to generate diverse, realistic training scenarios, boosting agent adaptability and performance by 30% in dynamic and unpredictable task environments
- Designed and integrated real-time task generation with large-scale parallelized RL training in IsaacGym, achieving a 6.8% increase in success rate and 50% improvement in sample efficiency across various robotic tasks
- Authored a research paper and submitted to IEEE International Conference on Robotics and Automation (ICRA) 2025

Computational Engineering and Robotics Lab, CMU — Pittsburgh, PA
Research Assistant: 3D AR Construction Scene Interaction and Inpainting via Deep Learning — *Jan 2022 – May 2023*
- Developed a real-time 3D AR scene interaction pipeline using Unity and C#, enabling dynamic object manipulation and physics simulation in construction environments, with 60 FPS performance on mobile devices
- Engineered an end-to-end deep learning pipeline for 3D scene inpainting, achieving 87% accuracy in large-scale occlusion completion tasks and reducing processing time by 40% through GPU optimization
- Designed and optimized a novel GAN architecture for high-resolution construction site image inpainting, enhancing texture realism by 35% over traditional computer vision methods while maintaining real-time performance

Biorobotics Lab, CMU — Pittsburgh, PA
Research Assistant: Recycled Paper Classification — *Sep 2021 – Dec 2021*
- Architected and trained a high-performance CNN model using PyTorch for recycled paper-grade classification, achieving 97% accuracy on a diverse dataset of 10,000+ images, surpassing previous benchmarks by 15%
- Engineered advanced data augmentation techniques, including adaptive mixing and style transfer, to enhance model generalization and robustness, improving performance on edge cases by 25% and reducing overfitting
- Developed a real-time image collection and processing pipeline using OpenCV and multiprocessing, optimized to handle 10 frames per second at 4K resolution, with a focus on low-latency inference (avg. 50ms per frame)

PROJECTS

Large-Scale Multi-Modal Mortality Prediction System — 11-785 Intro to Deep Learning, CMU
Course Project: Lead Developer — *Sep 2022 – Dec 2022*
- Engineered a scalable multi-modal survival analysis model integrating Vision Transformer (ViT) for processing large-scale medical imaging data with a BERT-based text encoder for clinical notes, coupled with a Cox Proportional Hazards model
- Implemented distributed training on a multi-GPU cluster to handle a dataset of 50,000+ patient records, including high-resolution radiography images and lengthy clinical text data
- Achieved a 22% improvement in concordance index (C-index) and a 18% increase in time-dependent AUC at 5 years, compared to state-of-the-art unimodal methods, validated using 5-fold cross-validation

TECHNICAL SKILLS

Programming Languages: Python, C++, C#, Java, SQL, MATLAB, R
AI Frameworks: PyTorch, TensorFlow, Keras, scikit-learn, OpenAI Gym
Cloud Computing: AWS, Google Cloud Platform
Robotic tools: ROS, Nvidia-IsaacGym, OpenCV, CUDA
Version Control: Git, Docker, Weight & Bias

PUBLICATIONS & PRESENTATIONS

Grounded Curriculum Learning — Submitted to ICRA 2025
3D Scene Interaction and Inpainting via Deep Learning — Poster Presentation, CMU MechE Symposium 2023