

Jeon Ho Kang

jeonhkang.com | jeonhoka@usc.edu | Los Angeles, CA

RESEARCH FOCUS

Robot Learning • Reinforcement Learning • Imitation Learning • Vision-Language-Action Models • Foundation Model Post-Training • Diffusion / Flow Policies • Sim-to-Real Robotic Manipulation

EDUCATION

University of Southern California

Ph.D. in Mechanical Engineering — Robotics & AI

M.S. in Mechanical Engineering — Robotics and Automation, GPA: 3.90/4.00

B.S. in Mechanical Engineering

Los Angeles, CA

Expected Dec 2026

May 2023

May 2022

RESEARCH & INDUSTRY EXPERIENCE

Amazon Robotics — Vulcan Stow, Applied Science

Applied Scientist II Intern

Aug 2026 – Present

Sunnyvale, CA

- Developing a **Real-to-Sim-to-Real robot learning framework** using **Newton and Isaac Lab** for contact-rich manipulation, targeting learned blade-insertion and space-creation behaviors.
- Developing **reinforcement-learning policies for contact-rich manipulation** and methods for transferring learned behaviors from reconstructed simulation environments to production robotic systems.
- Investigating simulation fidelity, contact dynamics, and policy robustness to reduce the **sim-to-real gap** for learned manipulation skills.

Amazon — Fauna, Applied Science

Applied Scientist II Intern

May 2026 – Aug 2026

Sunnyvale, CA

- Developed an **online RL post-training framework for Vision-Language-Action (VLA) foundation models**, reducing environment rollout requirements by approximately **6×** while maintaining strong task success across **LIBERO** and **ManiSkill-25**.
- Developed structured action-space exploration and policy-optimization methods for post-training pretrained robotic policies using **PPO / GRPO-style reinforcement learning**.
- Built scalable RL training and evaluation pipelines for **NVIDIA GR00T, $\pi 0.5$, and OpenVLA**, supporting parallel rollout collection and policy optimization on MuJoCo-based manipulation benchmarks.
- Optimized distributed multi-GPU training on **AWS** using **PyTorch Distributed**, improving training throughput, GPU utilization, and large-scale experiment efficiency.

Honda Research Institute

Research Scientist Intern

May 2025 – Aug 2025

Santa Clara, CA

- Investigated **Vision-Language Model reasoning** for long-horizon robotic decision making, studying instruction-following behavior, task-selection accuracy, and failure modes.
- Designed and implemented a **transformer-based imitation-learning pipeline** for multi-step robotic manipulation and integrated VLM-based task reasoning with learned low-level policies.
- Developed a classifier-conditioned observation module for **Diffusion Policy and ACT**, improving policy robustness over baseline methods in simulation and real-robot evaluations.
- Optimized PyTorch data loading and trajectory sequencing for large-scale robot demonstration datasets, achieving a **3×** **training-pipeline speedup** on multi-GPU hardware.

Realization of Robotic Systems Lab, University of Southern California

Graduate Research Assistant — Center for Advanced Manufacturing

Aug 2021 – May 2026

Los Angeles, CA

- Developed **multimodal imitation-learning methods for robotic manipulation** using vision, language, and force observations, with emphasis on diffusion policies and long-horizon contact-rich tasks.
- Developed **language-conditioned diffusion policies** with vision-language grounding and cross-attention for multi-task robotic disassembly without requiring object-level target annotations.
- Developed diffusion policies that fuse **vision and force observations** for compliant manipulation, achieving a **55% improvement** over a vision-only baseline.

- Built multimodal transformer and diffusion architectures in **PyTorch**, including language–vision cross-attention, temporal sequence modeling, and learned task-context representations.

SELECTED PUBLICATIONS

- **J.H. Kang**, I. Tamarkin, E. Niu, I. Novales, S.K. Gupta. *Vision–Language Grounded Task-Context-Aware Imitation Learning for Robotic Disassembly*. **IEEE Robotics and Automation Letters (RA-L)**.
- **J.H. Kang**, S. Joshi, R. Huang, S.K. Gupta. *Robotic Compliant Object Prying Using Diffusion Policy Guided by Vision and Force Observations*. **IEEE Robotics and Automation Letters (RA-L)**.
- **J.H. Kang**, S. Joshi, N. Dhanaraj, S.K. Gupta. *Task-Context-Aware Diffusion Policy with Language Guidance for Multi-task Disassembly*. **IEEE CASE**, 2025.
- **J.H. Kang**, N. Dhanaraj, S. Wadaskar, S.K. Gupta. *Using Large Language Models to Generate and Apply Contingency Handling Procedures in Collaborative Assembly*. **IEEE International Conference on Robotics and Automation (ICRA)**, 2024.
- **J.H. Kang**, N. Dhanaraj, O. Manyar, S. Wadaskar, S.K. Gupta. *A Task Allocation and Scheduling Framework to Facilitate Efficient Human-Robot Collaboration in High-Mix Assembly*. **ASME MSEC**, 2024. **Best Paper Award — Second Place**.

TECHNICAL SKILLS

Robot Learning:	Reinforcement learning, imitation learning, VLA post-training, diffusion policies, transformer policies, policy optimization, reward / value learning, sim-to-real
RL / ML:	PPO, GRPO, DPO, PyTorch, JAX, TensorFlow, Hugging Face Transformers, distributed training
Foundation Models:	Vision-Language Models, Vision-Language-Action Models, CLIP, multimodal transformers, cross-attention, fine-tuning and post-training
Robotics / Simulation:	Isaac Lab, Isaac Sim, Newton, MuJoCo, Gazebo
Systems:	Multi-GPU training, PyTorch Distributed, CUDA, AWS, Docker, Linux, Git
Programming:	Python, C/C++, Bash

HONORS & AWARDS

- **Best Conference Paper Award — Second Place**, ASME MSEC 2024
- NSF Travel Award, NAMRC / MSEC 2024