KARTIKAY MILIND PANGAONKAR

(480) 847-8346 | kpangaon@asu.edu | Linkedin

SUMMARY

Robotics engineer focused on learning-based control and policy generation for real-world robotic systems. Currently pursuing graduate research in robot learning, with hands-on experience developing reinforcement learning and diffusion-based policies and deploying them from simulation to physical platforms. Comfortable working across the full robotics stack, including perception, state estimation, and control, with strong experience in ROS, Python, and GPU-accelerated ML. Particularly interested in building robust, generalizable robotic systems that operate reliably outside controlled environments.

EDUCATION

Arizona State University **Aug 2024 - May 2026**
Master of Science, Robotics and Autonomous Systems *Tempe, AZ, USA*

- **GPA:** 3.92
- **Coursework:** machine learning, Reinforcement Learning Algorithms, Applied PDE, Advanced System modelling and control of robots, Autonomous Vehicle Engineering, Linear Algebra for Engineers, Machine Design

College Of Engineering Pune **Aug 2020 - May 2024**
Bachelor of Technology (B.Tech), Mechanical Engineering *Pune, India*

SKILLS

- **Robotics and Physics Engines**: MuJoCo, ROS1, ROS2, MoveIt, Isaac Gym, CoppeliaSim, SolidWorks, AutoCAD, ANSYS
- **Frameworks**: TensorFlow, Scikit, Keras
- **Languages**: Python, C, C++, Java, MatLAB, Rust, SQL
- **Libraries**: Stable Baselines, NumPy, JAX, Flax, Gymnasium, OpenCV
- **Infrastructure**: CUDA, Git, Linux, Bash, Docker, Kubernetes, Airflow, Terraform, Redis, Jupyter, GCP, AWS, REST APIs, CI/CD, Weights & Biases
- **Control Systems**: Model Predictive Control (MPC), PLC Programming, CAN
- **Robot Perception**: LiDAR, Sensor Fusion
- **ML & RL Algorithms**: RAG, RLHF, GRPO, SAC, PPO, LoRA, FSDP, Agentic AI, Imitation Learning

PROFESSIONAL EXPERIENCE

Logos Robotics Lab, Arizona State University **Mar 2025 - Present**
Research Associate

- Built teleoperation pipelines in ROS 1 and ROS 2 that let operators control robots with an auxiliary arm, a VR headset, and a SpaceMouse, enabling real-time manipulation of test platforms
- Designed a data-collection framework in Python that records synchronized camera images, joint states, and action commands to ROS bags, providing a clean dataset for training robot-manipulation models
- Trained robot-manipulation policies-including VLA, Diffusion Policy, and ACT on tabletop tasks, and conducted human-subject studies that demonstrated successful task completion with natural-language commands
- Designed a multimodal AI system that combines vision and language models to detect and correct distraction or obstruction failures in manipulation policies, reducing error occurrences during autonomous runs

IRIS Lab, Arizona State University **Nov 2024 - Dec 2024**
Research Associate

- Implemented 2D Gaussian splatting for object reconstruction using an Intel RealSense D435 camera mounted on a Franka Emika Panda robot; captured data and processed it with Python and OpenCV, reproduced the baseline methodology, and extended the pipeline to isolate target objects for downstream manipulation tasks

SKF (Swedish MNC in Ball Bearing and seals) **May 2023 - Jul 2023**
Intern, Research & Development *Pune*

- Engineered a complete process flow for full channel automation of HUB GEN 3 SKF automotive bearings, increasing production efficiency and consistency
- Designed and deployed an automated visual inspection system for taper HUB bearings, enhancing quality control and minimizing human error.

PUBLICATIONS

Factorizing Diffusion Policies for Observation Modality Prioritization

https://arxiv.org/html/2509.16830

- ICRA 2026 (Accepted), RSS 2025 Workshop
  - Introduced a factorized conditioning approach for diffusion policies to decouple and weight multiple observation modalities
  - Enabled task-dependent modality prioritization, improving performance in multimodal control settings
  - Demonstrated robustness to distribution shifts including distractors and sensor occlusions, with ~40% absolute gains
  - Showed improved data efficiency with ~15% gains in low-data regimes compared to standard joint-conditioning diffusion policies
  - Validated on both simulated benchmarks and real robotic platforms

StageCraft: Execution Aware Mitigation of Distractor and Obstruction Failures in VLA Models

https://arxiv.org/html/2603.20659v1

- (IROS 2026, Under Review)
  - Developed a training-free framework leveraging vision-language models (VLMs) for in-context reasoning to improve VLA policy robustness
  - Designed a system that analyzes rollout videos and success signals to identify and mitigate failure-inducing distractors and obstructions
  - Enabled adaptive environment modification without retraining underlying policies, improving robustness in unseen settings
  - Achieved ~40% performance improvement across real-world tasks and RLBench simulations
  - Built as a plug-and-play module compatible with existing VLA policies for sim-to-real deployment

Evaluating Language-Conditioned Robot Policies via human subject study (ongoing study)

- Evaluated the language instruction following capability of VLA models for free flow language instructions through crowdsourcing platforms like AMT.
- Conducting an ongoing human subject study to investigate ease of use for teleoperation systems like VR headset, Auxiliary arm.