

Yi Wei

✉ wei577@purdue.edu 🌐 yiwei19.github.io 🔄 yiwei19 🌐 yiwei19

Education

Purdue University, West Lafayette, IN Expected 2028
B.S. in Computer Science and Mathematics, Minor in Electrical Engineering Technology GPA: 3.92/4.0
Xiamen University, Xiamen, China
Exchange at HKUST, 2024; transferred to Purdue, 2025

Publications

Fast Generative Grasping via Lie Group-Constrained MeanFlow

S. Talha Bukhari, Yi Wei, Ruiqi Ni, Zachary Kingston, Aniket Bera. Submitted to IEEE Transactions on Robot Learning (under review). arXiv:2608.26076

Research Experience

Undergraduate Researcher, Computational Motion, Manipulation, and Autonomy Lab 2026 – present

Purdue CoMMA Lab, West Lafayette, IN — advised by Prof. Zachary Kingston

Real-Time MeanFlow-Based Grasp Generation for Robot Manipulation

🔴 real-robot demo

- Co-authored GraspMF, a 5-pass generative model for diverse 6-DoF grasps; reached 87.4% / 71.7% seen / unseen success at 15.5 ms per 100 grasps, beating 4 baselines by up to 39×
- Built the complete Franka R3 demo and ran 180 real-robot trials across 5 methods and 3 objects; GraspMF achieved 93.3% success in 5 steps and 96.7% in 10 without retraining or domain adaptation
- Architected the end-to-end perception-to-action stack, integrating wrist-mounted RGB-D sensing, object segmentation, robot-frame grasp conversion, collision-aware VAMP planning and closed-loop execution
- Ran 8 filter/order ablations and progressed from 1 ArUco marker to 12×13 and 6×6 calibration boards, achieving 7.1 mm median error across 13 fused point-cloud views; traced a ~10 cm Z-axis sim-to-real offset to frame mismatch and retrained VAMP for FR3 after Panda constraints failed on hardware

Zero-Shot Vision-Language Planning on LIBERO

🔴 tiptop-libero

- Adapted TiPToP, a zero-robot-data vision-language planning system, to LIBERO, benchmarking structured task-and-motion planning against learned VLA control under a shared evaluation protocol
- Built the RGB-D perception-to-execution pipeline with hollow-object grasp recovery; fixed a 2.18 cm frame mismatch and GPU rollout failures, reaching 80% planning and 50% execution success on 10 tasks

Dual-System Vision-Language-Action Control on LIBERO

🔴 fast-in-slow-libero

- Adapted Fast-in-Slow, a dual-system VLA combining slow vision-language reasoning with fast robot control, as the learned-policy counterpart on the same LIBERO tasks and observations
- Built a standalone perception-to-control bridge; diagnosed 5 interface mismatches behind the initial 100% failure rate and identified synchronous slow/fast scheduling as the control-frequency bottleneck

World-Model Control with Q-Learning and Constrained Search

ongoing, project lead

- Leading a model-based reinforcement learning project that combines a learned latent world model, Q-learning and model predictive control for efficient robot action planning
- Designing feasibility-constrained trajectory sampling and controlled ablations against standard TD-MPC2 planning, targeting higher task success and lower planning latency

Skills

Languages:	Python, C++, C, Java
Robot systems:	Franka R3 (libfranka/FCI), ROS, MuJoCo, Isaac Gym, PyBullet
Planning:	cuRobo, cuTAMP, VAMP
Perception:	Orbbec Femto RGB-D, hand-eye calibration, ArUco, ICP, Open3D, OpenCV
ML systems:	PyTorch, HuggingFace, DeepSpeed, LoRA/QLoRA, Docker, Ray, FastAPI