Yinong He

Phone: 19121130886 | yinongh@andrew.cmu.edu | hyn-kulu.github.io

Education

Carnegie Mellon University — Pittsburgh, PA
Master of Science in Robotics (MSR) — *Aug. 2025 – May 2027*

- **Advisors:** Co-advised by Prof. David Held and Prof. Zackory Erickson
- **Related Courses:** Mathematical Foundations of Robotics (A), Mechanics of Manipulation (A), Computer Vision (A+), Introduction to Robot Learning (A)
- **Current Courses:** Legged Systems, Robotic Software Engineering (ROS2)

UM-SJTU Joint Institute — Ann Arbor, Michigan / Shanghai, China
B.S.E. in Data Science and Electrical and Computer Engineering — *Sept. 2021 – May 2025*

- Completed the program with two years at the University of Michigan and two years at Shanghai Jiao Tong University.
- **Related Courses:** Advanced Artificial Intelligence (A), Mathematical Foundations of Robotics (A), Large Language Models (A+), Convex Optimization in Control (A+), Introduction to Robotic Manipulation (A+), Deep Learning for Robot Perception (A), Introduction to Machine Learning (A), Data Structure and Algorithms (A+)

University of Michigan — Ann Arbor, Michigan
B.S.E in Data Science — *Sept. 2023 – May 2025*

Shanghai Jiao Tong University — Shanghai, China
B.S.E. in Electrical and Computer Engineering — *Sept. 2021 – Aug. 2023*

WORKING PAPER AND PUBLICATION

1. Learning Dexterous Manipulation with Quantized Hand State
   Ying Feng*, Hongjie Fang*, **Yinong He***, Jingjing Chen, Chenxi Wang, Zihao He, Ruonan Liu, Cewu Lu
   Accepted at ICRA 2026 (* indicates equal contribution) [Paper]

2. Implicit Contact Diffuser: Sequential Contact Reasoning with Latent Point Cloud Diffusion
   Zixuan Huang, **Yinong He***, Yating Lin*, Dmitry Berenson
   Accepted at ICRA 2025 (* indicates equal contribution) [Paper]
   Best Technical Contribution Award at Michigan AI Symposium

3. Teaching Embodied Reinforcement Learning Agents: Informativeness and Diversity of Language Use
   Yinong He*, Jiajun Xi*, Jianing Yang, Yinpei Dai, Joyce Chai
   Accepted at EMNLP 2024 Main Conference (* indicates equal contribution) [Paper]

Research Experience

R-PAD Lab & RCHI Lab — Jun. 2026 – Present
Advisor: David Held and Zackory Erickson — *Pittsburgh, PA*

- Developing a hierarchical vision-language-action (VLA) policy for long-horizon robotic manipulation, combining per-point geometric grounding with vision-language semantic reasoning for high-level subgoal generation.
- Designing a two-stage architecture in which a VLA-based high-level policy predicts geometry-aware manipulation subgoals, while a low-level diffusion policy generates executable 6-DoF robot delta actions for robust control.
- Investigating transferable intermediate representations by integrating 3D geometric affordances with semantic task understanding to improve long-horizon manipulation generalization.

R-PAD Lab & RCHI Lab — Sep. 2025 – Jun. 2026
Advisor: David Held and Zackory Erickson — *Pittsburgh, PA*

- Developed a multimodal diffusion policy for contact-rich insertion built upon Conditional U-Net, using LeRobot framework and Isaac Gym for policy training and setup. Designed dual ResNet18 visual encoders to fuse wrist RGB-D observations, aligned geometric depth representations and end-effector force embeddings for chunked 6-DoF action prediction.

- Designed a representation interface by projecting wrist and socket depth observations into canonical virtual views through depth registration and virtual camera projection, preserving plug–socket geometric alignment while enabling physically consistent SO(2) rotational augmentation.
- Built a complete sim-to-real RGB-D perception pipeline using wrist-mounted and external **ZED Mini** stereo cameras with the ZED SDK and FoundationStereo depth estimation. Developed realistic depth-domain augmentation including correlated edge noise, irregular depth-hole masking, Gaussian blur and random camera cropping to bridge simulation and real-world depth distributions.
- Implemented eye-in-hand and eye-to-side RGB-D calibration pipelines, enabling accurate coordinate transformation between robot base, wrist camera and external camera frames.
- Developed a high-precision robotic insertion system on a Franka Panda manipulator using PyBullet inverse kinematics, joint-space position control and force-guided execution, achieving sub-millimeter ($\sim$0.1 mm) insertion accuracy under tight-clearance industrial assembly settings.
- Built large-scale simulation datasets in Isaac Gym with procedural CAD asset generation, randomized insertion geometries, recovery trajectories and LeRobot-compatible demonstrations for diffusion policy few-shot adaptation.

- Designed a procedural geometry generation pipeline that automatically synthesizes hundreds of plug–socket pairs with randomized topology, rotational symmetry, manufacturing clearance and insertion tolerance, improving geometric diversity for sim-to-real policy transfer.
- Deployed the complete perception-learning-control stack on a real Franka Panda platform, integrating stereo vision, calibration, force sensing and diffusion policy inference for zero-shot generalization and few-shot adaptation on unseen industrial insertion tasks.
- Resulted in a first-author paper submitted to **CoRL 2026**.

NeoMatrix Research Internship May 2025 – Sep. 2025
Advisor: Cewu Lu *Shanghai, China*

- Developed a 3D diffusion-based visuomotor policy built upon **RISE**, consisting of a sparse 3D convolutional encoder (Minkowski sparse convolutions), sparse positional encoding, Transformer encoder-decoder, and a DDIM-based diffusion head for continuous arm-hand action generation from multi-view fused point clouds.
- Extended the RISE policy to dexterous manipulation by jointly predicting 6-DoF arm trajectories and high-DoF hand actions, enabling coordinated arm-hand control under contact-rich manipulation scenarios.
- Designed a structured hand action representation using a two-level **Residual VQ-VAE**, compressing high-dimensional dexterous hand poses into discrete latent codebooks while preserving grasp semantics and manipulation diversity.
- Proposed a PCA-based continuous relaxation over discrete VQ codewords, transforming hand token prediction into a continuous regression problem compatible with diffusion policies and eliminating optimization conflicts between action regression and discrete classification.
- Unified continuous arm trajectory prediction and quantized hand action generation within a single diffusion framework, allowing simultaneous optimization of arm localization and dexterous grasp synthesis in an end-to-end policy.
- Built a hybrid teleoperation system integrating Meta Quest 3, OyMotion GForce Pro motion gloves, dual RGB-D cameras, and a Flexiv robotic arm to collect synchronized real-world demonstrations for six dexterous manipulation tasks.
- Validated the proposed framework on six real-world dexterous manipulation benchmarks involving articulated object interaction, long-horizon manipulation and large-rotation tasks, achieving state-of-the-art performance over existing diffusion-policy baselines.
- Resulted in a co-first-author paper accepted at **ICRA 2026**.

Autonomous Robotic Manipulation Lab May 2024 – Apr. 2025
Advisor: Dmitry Berenson *Ann Arbor, MI*

- Developed a complete MuJoCo simulation environment for long-horizon contact-rich deformable manipulation, including cable routing and notebook folding tasks. Built scripted policies to collect over 15,000 trajectories for diffusion model training and evaluation.
- Designed a contact-aware implicit scene representation by extending Neural Descriptor Fields (NDF), using PointNet-based scene encoders, point-wise neural descriptors and SDF-gradient supervision to encode geometric and topological relationships between deformable objects and the environment.
- Developed a latent point-cloud diffusion planner by combining PointNet++ encoders, a VAE latent space, hierarchical latent diffusion models and reachability prediction to generate sequential contact-aware subgoals for long-horizon manipulation.

- Integrated diffusion-based planning with Model Predictive Path Integral (MPPI), enabling online receding-horizon trajectory optimization that tracks predicted contact subgoals while avoiding local minima in contact-rich manipulation.
- Deployed the complete planning framework on a real 7-DoF KUKA LBR iiwa robot using a Zivid 2 RGB-D camera and CDCPD deformable object tracking for sim-to-real cable routing across multiple cable types, demonstrating robust long-horizon contact reasoning.
- Resulted in a second-author paper accepted at **ICRA 2025**.

Situated Language and Embodied Dialogue Lab Aug. 2023 – Jun. 2024
Advisor: Joyce Chai, Professor in EECS Department *Ann Arbor, MI*

- Developed the **Language-Teachable Decision Transformer (LTDT)**, extending Decision Transformers with language-conditioned offline reinforcement learning for embodied agents capable of learning from natural-language feedback.
- Trained PPO policies for an embodied navigation task in simulation to collect expert data trajectories.
- Designed multimodal trajectory representations consisting of task descriptions, returns-to-go, states, actions, and Sentence-BERT language embeddings, enabling transformer-based policy learning from rich language feedback.
- Conducted large-scale empirical studies across four embodied RL benchmarks, demonstrating improved in-domain performance, few-shot adaptation to unseen tasks, and effective transfer under human language instructions.
- Performed comprehensive analyses on language diversity, feedback frequency, instruction informativeness, and robustness against adversarial language inputs, identifying when natural language most effectively improves offline RL agents.
- Resulted in a co-first-author paper accepted at **EMNLP 2024 Main Conference**.

Selected Projects

Reinforcement Learning Algorithms for Quadruped Locomotion Sep. 2025 – Dec. 2025
Course Project for 16-831 Robot Learning (Carnegie Mellon University) *Instructor: David Held*

- Built a reinforcement learning benchmark for the **Unitree Go2** quadruped using **Isaac Lab**, and systematically compared **PPO**, **SAC**, and **TD-MPC2** on rough-terrain velocity tracking with respect to training stability, sample efficiency, and locomotion performance.
- Implemented both **position-control** and **torque-control** action spaces, and analyzed how low-level control interfaces affect policy learning, balance recovery, contact-rich locomotion, and tracking accuracy.
- Designed and implemented **Hard-Sync GRPO**, a command-conditioned group-relative policy optimization framework that synchronizes velocity commands and terrain curricula across parallel environments to reduce policy-gradient variance without training a value network.
- Designed reward functions and evaluation metrics for velocity tracking, and performed comprehensive analyses of learning curves, failure modes, and control performance across different reinforcement learning algorithms.

State-Feedback Control Design with Sector-Bounded Nonlinearities Aug. 2024 – Dec. 2024
Course Project for ECE598 Convex Optimization in Control (University of Michigan) *Instructor: Peter Seiler*

- Developed state-feedback control theorems using advanced mathematical tools, including the Lyapunov Theorem, Circle Criterion, Schur Complement, and Linear Matrix Inequalities (LMI), leveraging Semi-Definite Programming (SDP) for convex optimization to ensure stability and performance under nonlinear sector-bounded dynamics.
- Optimized controllers with both H_2 performance for minimizing energy response and H_∞ performance for robust disturbance rejection.

Reasoning-Guided Video Generation for Robotic Manipulation Jan. 2024 – Apr. 2024
Course Project for EECS692 Advanced Artificial Intelligence. (University of Michigan) *Instructor: Joyce Chai*

- Applied the diffusion model for video generation, leveraged GPT-4V for interpreting the robots' behaviors and providing instructions, and fine-tuned an image-editing model for generating corrected subgoals.

Enhance Distilled Feature Field for Complex Language Query Jan. 2024 – Apr. 2024
Course Project for ROB498 Deep Learning for Robot Perception (University of Michigan) *Instructor: Xiaoxiao Du*

- Trained a distilled feature field with CLIP and NeRF for interpreting and rendering the scene.
- Implemented a pipeline leveraging GPT4 to decompose the complex language queries, analyze the spatial relationship, and locate the querying object by grounding to the feature field.

Empowering VLM with Spatial Reasoning Ability Jan. 2024 – Apr. 2024
Course Project for EECS498 Large Language Model (University of Michigan) *Instructor: Samet Oymak*

- Reproduced the SpatialVLM pipeline for augmenting data by lifting 2D images to 3D point clouds and extracting spatial relationships. Finetuned the LLaVA model with the generated data.
- Calibrated LLaVA's bias towards trusting the given spatial relationships.

Awards

Best Technical Contribution Award @ Michigan AI Symposium Oct. 2024

Dean's Honor List Apr. 2024

Dean's Honor List Dec. 2023

Silver Medal in University Physics Competition Nov. 2022

Shanghai Jiao Tong University Science and Technology Scholarship May 2023

Language Proficiency

- TOEFL: 113 (Reading: 30, Listening: 28, Speaking: 26, Writing: 29)
- GRE: 331 (Verbal: 161, Quant: 170)

Technical Skills

- **Programming:** Python, C++, Linux, Git
- **Machine Learning:** PyTorch, Hugging Face, OpenCV, Open3D
- **Simulation & Robotics:** Isaac Gym, MuJoCo, PyBullet, LeRobot, Franka Panda, Flexiv Rizon
- **Perception:** ZED SDK, FoundationStereo, RGB-D Vision, Point Cloud Processing, Camera Calibration