Yuning Wu

Carnegie Mellon University
Computational Design Laboratory
5000 Forbes Ave, Pittsburgh, PA, 15213

+1 (412)-626-2495
yuning.wu@outlook.com
https://ice-5.github.io/

Reinforcement Learning, Robot Learning, Human-Robot Interaction, Robot Prototyping

EDUCATION

Ph.D. in Computational Design, Carnegie Mellon University, Pittsburgh, PA 2024.9
Co-Advisors: Daniel Cardoso Llach (computational design) and Jean Oh (robotics)
Dissertation: *"A Human-Centered, Reinforcement Learning Driven Robotic Framework for Assisting Construction Workers"*

M.S. in Machine Learning, Carnegie Mellon University, Pittsburgh, PA 2021.5
Advisor: Katerina Fragkiadaki

M.Eng. & B.Eng., Architecture, Tsinghua University, Beijing, China 2017.7

WORK EXPERIENCE

Autodesk Research, San Francisco, CA 2021.5 - 2022.5

Research Intern and Collaborator, Autodesk AI Lab & Robotics Lab
Developed Inverse Reinforcement Learning (IRL) methods for contact-rich robotic manipulation, facilitating real-world applications such as autonomous and precise joint assembly of large-scale timber structures.

Carnegie Mellon University, Pittsburgh, PA 2020.8 - 2023.5

Graduate Researcher and Technical Lead
Led the full-cycle development of Husky, the "work companion robot", from concept to implementation, with support from Google, Autodesk, and the State of Pennsylvania.

Metascope Technology, Beijing, China 2017.12 - 2019.8

CTO and Co-founder
Led a multidisciplinary team of engineers and designers to develop an ML-based aesthetic profiling tool tailored for commercial architects and clients in China, benefiting an array of small- to medium-sized design firms and studios.

RESEARCH PROJECT

A Reinforcement Learning-Driven "Work Companion Rover" for Assisting Construction Workers [1] [2] [3]

Husky is an RL-driven "work companion rover" designed to be an on-site co-pilot for construction workers. By bearing heavy loads and autonomously transporting essential tools and materials, it closely supports workers in complex, dynamic, and labor-intensive real-world workspaces with adaptive social compliance.

- Designed and built a ROS-based, multi-sensor robot prototype from a basic chassis (scratch).
- Employed state-of-the-art reinforcement learning (RL) techniques, integrating attention-based motion prediction and value-based discrete action selection, to enhance worker-robot interaction comfort in complex indoor construction environments.
- Developed a novel alignment technique using the Plackett-Luce model for incremental adaptation of a pre-trained DRL model to context-specific worker behavioral patterns while mitigating policy shift in Sim2Real transfer.
- Utilized SoTA LiDAR Odometry and Mapping (LOAM) methods for unstructured site mapping and robot state estimation.
- Customized an efficient multimodal stack (LiDAR + multiple cameras) to robustly detect and track workers' movements in obstructive and cluttered construction environments.

Learning Dense Reward for RL-Based Contact-Rich Manipulation from Limited Human Demonstrations [4]

A dense reward—a continuous metric that quantifies task progress—can be efficiently learned from multimodal sensory data, including vision and force/torque (F/T) inputs, using a limited set of human demonstrations. This learned reward model resolves the need for expert-engineered rewards and can be effectively applied in training deep reinforcement learning models (e.g., PPO, TRPO, SAC) for contact-rich, precision-intensive manipulation tasks, such as lap joint assembly, USB insertion, and peg-in-hole.

- Designed a novel Temporal Variant Forward Sampling (TVFS) method, enhancing robustness and quality in augmenting data samples from 1-2 human demonstrations within a physical simulation environment (NVIDIA Isaac Sim, PyTorch, Mujoco).
- Introduced an end-to-end, self-supervised learning architecture that leverages the temporal relationships within sequential sample pairs.
- Evaluated among an array of short-horizon, precision-intensive tasks and proved generalizable.

A Distributed, Scalable RL-Driven Multi-Drone Framework for Collective Task Execution on Complex Geometric Structures [5]

Pixhawks is a fleet of drones powered by a reinforcement learning-based multi-agent collision avoidance algorithm, engineered for cooperative tasks around human workers on complex geometric structures. The system supports activities such as brick-laying and spray-coating on intricately shaped 3D structures (e.g., curved walls, domes, shells), showcasing a human-robot interactive framework capable of managing sophisticated, real-world designs and complexities.

- Developed a scalable, decentralized RL-based multi-agent collision avoidance algorithm, enhanced from MADDPG, to ensure efficient multi-drone operation around dynamic workers within limited space.
- Developed a multi-drone task execution and coordination workflow simulation environment using Unity and MLAgents.
- Customized a generalizable, curriculum-based Sim2Real pipeline using Unity MLAgents to bridge virtual training and real-world application.
- Conducted a proof-of-concept brick-laying experiment with physical drone prototypes to validate the framework in the real world.

RELATED PUBLICATION

1. Yuning Wu, Jiaying Wei, Jean Oh, and Daniel Cardoso Llach, "**Towards Human-Centered Construction Robotics: A Reinforcement Learning-Driven Companion Robot For Contextually Assisting Carpentry Workers**", in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Abu Dhabi, 2024. [PDF]
2. Yuning Wu*, Emek Erdolu*, Jiaying Wei, Jean Oh, and Daniel Cardoso Llach, "**Robot in the Loop: A Human-Centered Approach to Contextualizing AI and Robotics in Construction**", Construction Robotics, Springer Nature, 2024. (Accepted for publication). [PDF]
3. Yuning Wu, "**A Human-Centered, Reinforcement Learning Driven Robotic Framework for Assisting Construction Workers**", Ph.D. dissertation, Carnegie Mellon University, 2024. [PDF]
4. Yuning Wu, Jieliang Luo, and Hui Li, "**Learning Dense Rewards with Temporal Variant Self-Supervision**", in IEEE International Conference on Robotics and Automation (ICRA), Reinforcement Learning for Contact-Rich Manipulation Workshop, Philadelphia, PA, USA, 2022. [PDF][Code]
5. Zhihao Fang, Yuning Wu, Ammar Hassonjee, Ardavan Bidgoli, and Daniel Cardoso Llach, "**Towards an Architectural Framework for Distributed, Robotically-Assisted Construction: Using Reinforcement Learning to Support Scalable Multi-Drone Construction in Dynamic Environments**", In Proceedings of the 40th Annual Conference of the Association for Computer-Aided Design in Architecture (ACADIA), Philadelphia, 2020. [PDF]
6. Omer Akin and Yuning Wu, "**State Space Paradox of Computational Research in Creativity**", in the 19th Heron Island Conference on Computational and Cognitive Models of Creative Design (HI), 2019. [PDF]

OTHER PROJECT

Teaching AI Agent Spatial Sense and Awareness by Playing with LEGOs
This research trains an AI agent to iteratively build structures using LEGOs in a physical simulation environment by integrating Graph Neural Networks (GNN) for physical state inference and Monte Carlo Tree Search (MCTS) for optimal planning.

Can Vision Language Models (VLM) Acquire Deeper Spatial Sense?
While VLMs excel at processing and associating verbal and visual data, their potential in physical world interactions hinges on a deeper grasp of spatial and physical realities. This research extends the LEGO Player initiative, aiming to imbue VLMs with enhanced spatial understanding.

MoonRanger: An Autonomous Rover for Searching Water on The Moon
An autonomous lunar rover designed for exploring and analyzing ice and water on the moon's poles, featuring advanced navigational algorithms. Participated in system engineering advised by Prof. William (Red) Whittaker and Dr. Heather Jones, supported by NASA CLPS program.

HONOR AND AWARD

2023, 2024	**CMU Presidential Fellowship**. Carnegie Mellon University, TCS. Highest honors in CMU.
2021, 2022	**Manufacturing Futures Initiatives Fellowship.** Carnegie Mellon University.
2022	**Autodesk AI Merit Scholarship.** Autodesk AI Lab, Autodesk Research.
2021	**Manufacturing Pennsylvania Fellowship.** State of Pennsylvania. Carnegie Mellon University.
2020	**GSA-Provost GuSH Grant.** Carnegie Mellon University.
2019	**LEGO Scholarship.** LEGO Foundation.

RESEARCH TALK AND PRESENTATION

2023	**Cornell Tech, TCS Innovation Forum,** invited research talk.
2023	**Bosch Leadership and AI Team Research Exchange,** research talk and robot demonstration.
2022	**CMU Manufacturing Futures Institute Research Seminar**, invited research talk.
2022	**ICRA Reinforcement Learning for Contact-Rich Manipulation Workshop**, poster presentation.
2020	**ACADIA Automated Fabrication and Operation Session**, poster presentation.

SKILLSET

AI/ML: Pytorch, Numpy, Scipy, Docker, CUDA, HuggingFace
Robotics: ROS, ROS2, NVIDIA Omniverse Isaac, Unity ML-Agents, PyBullet, MuJoCo, robosuite
Programming: Python, C/C++, C#
3D Design & Modeling: Rhinoceros, AutoCAD, Autodesk Fusion 360

COURSE

(10-703) Deep Reinforcement Learning and Control, (10-716) Advanced Machine Learning, (16-720) Computer Vision
(10-605) Machine Learning with Large Datasets, (10-708) Probabilistic Graphical Models,
(10-725) Convex Optimization, (36-700) Probability and Mathematical Statistics, (16-861) Space Robotics

REFERENCE

Dr. Daniel Cardoso Llach, Associate Professor in Computational Design, Carnegie Mellon University
Dr. Jean Oh, Associate Professor in Robotics, Carnegie Mellon University
Dr. Jieliang (Rodger) Luo, Former Senior Principal Research Scientist, Autodesk Research AI Lab
Dr. Hui Li, Senior Principal Research Scientist, Autodesk Research Robotics Lab