Dr. Jiayu Chen

- **Short Bio:** Before joining HKU DASE, I was a postdoc at Carnegie Mellon University from 2024.05 to 2025.05. I obtained my Ph.D. degree in Industrial Engineering from Purdue University in 2024 and my Bachelor Degree from College of Engineering, Peking University in 2020.

- **Research Area:** Reinforcement Learning, Robotic Learning.

- **Lab Website:** https://agentic-intelligence-lab.org/

- **Email:** jiayuc@hku.hk

Offline Reinforcement Learning for Controllable Nuclear Fusion

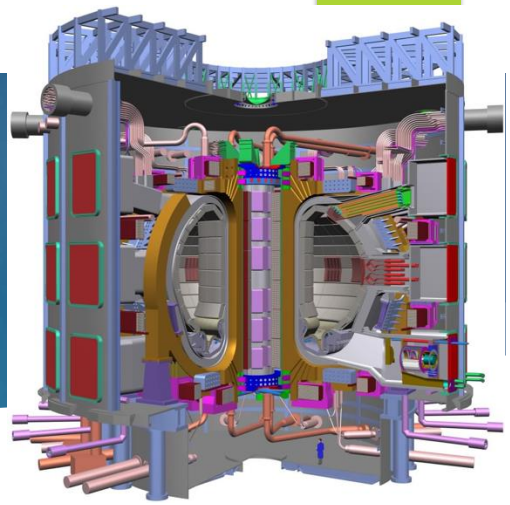

- **Background:** The complex physics underlying nuclear fusion makes the application of traditional control methods, such as Model Predictive Control (MPC), impractical. However, with access to extensive operational data from a real Tokamak device (DIII-D), we can leverage data-driven control approaches, specifically, offline reinforcement learning (RL), to develop effective control strategies. This project is a collaborative effort with Carnegie Mellon University and Princeton University.
- **Goals:** (1) Develop a benchmark for offline RL methods using fusion data. (2) Design robust control algorithms based on a fusion simulator.
- **Prerequisites:** (1) Basics of reinforcement learning; (2) Coding with PyTorch; (3) Comfort working with large codebases: Offline-RL-Kit-for-Nuclear-Fusion.

Continual Reinforcement Learning for Humanoid Robotic Control

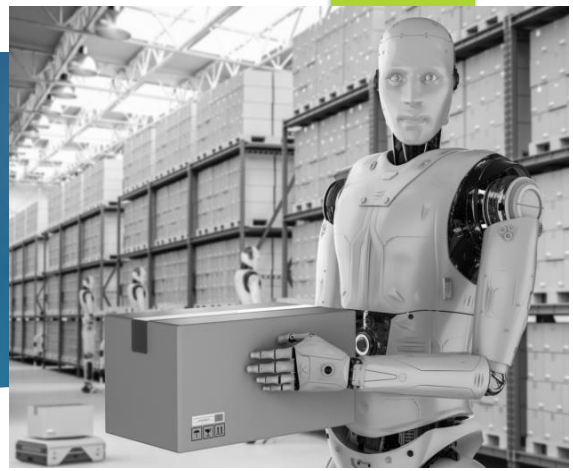

- **Background:** Next-level embodied intelligence should be capable of completing a series of tasks within a structured environment and adapting to new tasks through a few trials-and-errors or demonstrations. Achieving this requires a comprehensive and realistic simulator for the chosen scenario, a robust multi-task learning framework, and efficient continual RL algorithms.
- **Goals:** (1) Design an industrial scenario involving humanoid robots and develop a comprehensive simulator using Isaac Sim. (2) Propose a new multi-task learning framework suitable for simulation-based pretraining. (3) Develop a novel RL algorithm for efficient continual learning with robots.
- **Prerequisites:** (1) Basics of reinforcement learning; (2) Coding with PyTorch; (3) Experience with Issac Sim or ROS (optional).

Learning-based Ego-centric Evolution within a Multi-agent System

- **Background:** A likely next step in the development of AGI systems, such as large language models (LLMs), is enabling self-evolution within a multi-agent ecosystem. Achieving this vision requires multiple technical advancements, including—but not limited to—the following goals.
- **Goals:** (1) Develop a simulation platform that supports diverse task design, visualization, and efficient parallel and distributed training. (2) Design continual learning algorithms for multi-agent systems composed of pretrained models. (3) Empirically verify scaling laws within the ecosystem. (4) Explore ego-centric learning under sparse reward signals, potentially involving automatic information gathering and dynamic agent grouping, etc.
- **Prerequisites:** (1) Basics of (multi-agent) reinforcement learning; (2) Coding with PyTorch.