COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF MICHIGAN

# DISSERTATION DEFENSE

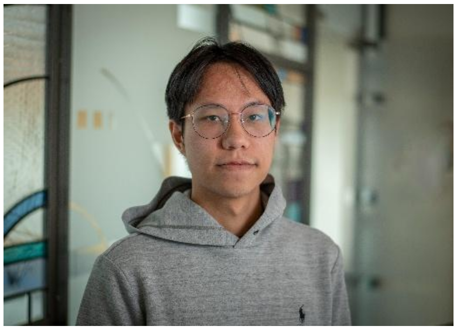

## Anthony Liu

## Leveraging Compositional Structure for Reinforcement Learning and Decision Making Problems

Monday, December 16, 2024
1:30pm – 3:30pm
Virtual – [Zoom]

**ABSTRACT:** Deep Learning approaches have made tremendous progress toward solving reinforcement learning and decision-making problems. However, current approaches still struggle with *long-horizon* tasks that require strong *generalization*. These are tasks that an agent must solve using many actions and may have situations where the agent must generalize its actions from prior experience. A dominant approach is to solve these tasks in a hierarchical manner: a high-level agent decomposes a task into multiple “subtasks” to be individually solved by a low-level agent, which specializes in solving these subtasks. The effectiveness of this approach relies on two key assumptions about the compositional structure of tasks: that tasks can be decomposed in a top-down way, and that subtasks can be re-used across tasks.
The thesis of this dissertation argues that **stronger assumptions** about compositional structure can be made to improve task learning efficiency and ability to generalize to new tasks.

In this dissertation, I present the following approaches for utilizing these assumptions, and experimental evidence that shows that these approaches improve learning performance on realistic benchmark tasks.
(1) I re-examine the top-down assumption, and find that learning policies for subtasks in isolation can lead to sub-optimal performance. I propose a novel hierarchical reinforcement learning framework that learns more optimally, which learns low-level policies that look ahead to multiple subtasks.
(2) Assuming that tasks follow *parameterized* structure (e.g. function-argument, action-object tasks), I show that we can improve the high-level agent's ability to efficiently learn and generalize through learning a novel structure: a parameterized subtask graph.
(3) Finally, assuming that tasks can be structured through *control flow* (e.g. solving tasks using code), I show how to use large language models (LLM) to write code to solve these tasks in an effective and error-free way.

In conjunction, these approaches show that embedding these strong assumptions about compositional structure can improve efficiency and generalization for long-horizon tasks through reinforcement learning and large language model approaches.

**CHAIR:** Prof. Honglak Lee