

Research Vision: Lifelong Robot Autonomy and Embodied Intelligence

Robots and embodied AI agents are increasingly expected to operate persistently in environments, where tasks arise continuously and environmental conditions evolve over time. In such settings, intelligence is not solely measured by the ability to complete an assigned task, but by how actions shape the environment for what comes next. A robot that performs well on an isolated assignment may still degrade long-term performance if its decisions introduce clutter, obstruction, or inefficiencies that accumulate across time. Enabling embodied agents to act with sustained environmental awareness is therefore central to achieving robust, long-horizon autonomy.

Humans routinely carry out tasks in ways that limit undesirable side effects, or at least attempt to do so. When clearing a table, for instance, dirty cups are placed in or near the sink rather than on the nearest available surface. When loading a cabinet, objects are often arranged to preserve reachability and reduce obstruction. These behaviors are not tied to a single specified future task; instead, they reflect recurring patterns in how environments are used over time. They reveal an implicit capacity to act with awareness of likely future needs, even when those needs are not explicitly stated.

My research investigates how robots and embodied AI agents can acquire similar forward-looking capabilities. More broadly, I am interested in the principles and computational methods that support lifelong robot autonomy: how an agent can make decisions that not only solve the present problem, but also preserve or improve the environment for future interaction. This perspective shifts the focus of autonomy from short-term task completion to sustained performance across evolving streams of activity. It asks how embodied systems can become not merely reactive problem-solvers, but persistent, adaptive participants in shared spaces.

A central theme of my work is modeling the long-term value of environmental states. Rather than evaluating actions only by whether they achieve an immediate objective, I study how agents can estimate the downstream consequences of intermediate decisions under plausible future task distributions. These estimates can then guide planning and decision making toward behaviors that reduce cumulative effort, mitigate negative side effects, and maintain accessibility and organization over time which we refer to as *anticipatory planning*. Through this lens, planning becomes part of a broader lifelong autonomy problem: enabling embodied agents to reason over temporally extended consequences in environments that continue to matter after the current task has ended.

This research contributes to embodied intelligence by addressing a capability that will be essential for long-term deployment in homes, workplaces, warehouses, and care settings. In such environments, success depends not only on accomplishing tasks, but on doing so in a manner that supports what people and robots are likely to need next. By equipping embodied systems with anticipatory reasoning, my work aims to advance robots that behave more thoughtfully, collaborate more effectively with humans, and maintain environments in ways that improve long-term performance rather than degrade it.

Significance of My Research

The significance of this research lies in its focus on a foundational limitation in current embodied systems: the gap between short-horizon competence and long-horizon autonomy. Many existing methods can solve well-defined tasks effectively, yet they remain limited in their ability to account for how present actions influence future opportunities, constraints, and operating conditions. As robots move from controlled demonstrations toward persistent deployment in real environments, this limitation becomes increasingly consequential.

My work addresses this challenge by studying how embodied agents can reason about the future utility of the states they create. This is important not only for improving efficiency, but also for supporting safer, more reliable, and more human-compatible behavior. In logistics and warehouse settings, such capabilities can help robots organize inventory or place objects in ways that facilitate future retrieval and workflow. In assistive living and household environments, they can help maintain accessibility, reduce clutter, and prepare the environment for anticipated needs. In collaborative spaces, they can

support smoother interaction by reducing friction between immediate task completion and longer-term environmental usability.

More broadly, this research pushes toward a view of embodied AI in which autonomy is measured by continuity of performance over time. A capable embodied agent should not simply execute instructions; it should manage its actions in a way that preserves optionality, supports future activity, and contributes positively to the environments in which it operates. Advancing this capability is, in my view, an important step toward truly lifelong robot autonomy.

Potential Future Directions

Looking ahead, I am interested in pursuing several future directions that extend this research agenda beyond immediate anticipatory decision making and toward broader lifelong embodied intelligence.

First, I am interested in studying whether embodied agents can recognize and exploit opportunities to make incidental progress toward likely future objectives while executing their current task. In long-lived environments, many tasks share latent structure, environmental preconditions, or partially overlapping goals. This raises the question of whether an agent can do more than merely avoid harmful side effects: can it also create states that are incidentally beneficial for what may come next? For example, while clearing a table, a robot could pour leftover water from a cup into a nearby vase before placing the cup in the sink, incidentally supporting a future watering plant task. Such behavior reflects opportunistic value accumulation rather than isolated task completion. Investigating this direction would require new formulations of multi-task value, mechanisms for identifying beneficial overlap across tasks, and decision-making strategies that balance opportunistic progress with reliability on the immediate objective.

Second, I am interested in how embodied agents can form and adapt task expectations when future task distributions are not explicitly available. In real-world environments, future activities are shaped by context, routine, and common-sense patterns rather than fixed priors. Continual learning and planning offer a promising path, allowing agents to update task beliefs online as environments and usage patterns evolve. Language-grounded models and world knowledge can help generate task hypotheses from environmental context, interaction history, semantic cues, and object state. For example, chopped vegetables on a cutting board versus whole produce in storage imply different likely activities such as meal preparation versus ingredient organization; similarly, accumulating dirty dishes suggest cleaning tasks, while partially arranged table settings indicate imminent dining. Integrating such semantic and state-based signals into adaptive task-distribution models would enable agents to revise expectations, estimate long-term state value, and act anticipatorily in open-world settings where objectives are implicit, evolving, and user-specific.

Overall, my long-term goal is to develop embodied agents that do not merely respond competently in the moment, but instead act with sustained awareness of how present choices shape future performance. I believe this is a central challenge for lifelong robot autonomy and an essential ingredient for building embodied AI systems that are effective, adaptive, and genuinely useful in the environments they inhabit.