Roshan Dhakal

Research Vision: Lifelong Robot Autonomy and Embodied Intelligence

My research vision is to develop embodied AI agents that can reason, adapt, and improve over long time horizons in persistent real-world environments. As robots move from isolated demonstrations toward continuous deployment in homes, workplaces, warehouses, factories, and assistive settings, intelligence can no longer be measured solely by immediate task completion. Embodied agents continuously reshape the environments in which they operate. Over time, these interactions alter object arrangement, accessibility, free space, navigation structure, and manipulation feasibility, gradually transforming environments into configurations that are either favorable or unfavorable for long-term operation.

A robot may complete a task successfully while leaving behind cluttered workspaces, blocked access paths, unstable object arrangements, or poorly organized storage configurations that accumulate into increasingly difficult operating conditions. Conversely, agents acting with long-horizon awareness may maintain environments that remain organized, navigable, accessible, and easier to operate within over extended deployments. The quality of environmental state therefore becomes an important component of embodied intelligence itself. Long-term autonomy is not only about solving tasks, but also about continuously maintaining environments in configurations that remain favorable for future interaction.

Humans naturally exhibit this type of behavior in everyday environments. Dirty dishes are moved toward the sink rather than left scattered across countertops. Objects inside cabinets are arranged to preserve reachability and avoid obstruction. Shared spaces are reorganized to reduce clutter and maintain usable workspace. Importantly, these behaviors are often not tied to a single explicit future objective. Rather, they emerge from an implicit understanding that environmental structure itself matters over time. Behaviors such as organization, tidiness, preparation, and environmental maintenance arise naturally from prolonged interaction with persistent environments.

My research investigates how embodied agents can acquire similar capabilities. I am interested in how robots can reason about the long-term quality of environmental state configurations and distinguish between states that remain favorable over time and those that gradually degrade future operation. In long-horizon settings, actions continuously affect accessibility, reachability, obstruction, collision risk, and manipulation difficulty. Some environmental configurations preserve navigable structure, maintain free space, reduce interference between objects, and remain favorable across many future interactions, while others accumulate clutter, introduce blockage, and increase long-term difficulty for navigation and manipulation.

This perspective shifts the focus of embodied intelligence beyond isolated task execution toward persistent environmental reasoning. Through this lens, behaviors such as organization, decluttering, accessibility preservation, and preparation emerge not as explicitly programmed objectives, but as natural consequences of reasoning over long-term environmental quality. Future tasks remain an important part of this process, since the environments created by present actions directly influence what becomes easy or difficult later. More broadly, however, the challenge is enabling embodied agents to continuously maintain environmental configurations that remain stable, usable, and favorable across extended deployments.

My Ph.D. research has explored these problems through long-horizon embodied decision making in rearrangement and manipulation domains. In these settings, robots must reason over both symbolic and geometric structure while accounting for how local decisions alter long-term environmental quality. For example, when organizing objects inside a cabinet, a robot should not only place objects successfully, but also preserve reachability and avoid creating obstruction for later manipulation. Similarly, in navigation among movable obstacles, relocating an obstacle to the nearest feasible location may solve an immediate navigation problem while gradually degrading accessibility for future tasks. These scenarios illustrate a broader challenge in embodied AI: local decisions continuously reshape the environment itself, and those changes accumulate over time.

Going forward, I am interested in pursuing a broader research agenda centered on persistent environmental reasoning, self-improving embodied agents, grounded world models, and long-horizon agentic systems. In particular, I am interested in developing embodied agents that combine structured reasoning, semantic understanding, environmental memory, and continual adaptation to operate robustly within evolving real-world environments.

Research Direction 1: Persistent Environmental Reasoning

A central direction of my future research is understanding how embodied agents can reason about the long-term quality of environmental structure. In persistent environments, robots continuously modify object arrangement, accessibility, navigation structure, and free space. These changes may accumulate gradually, transforming environments into states that are increasingly favorable or increasingly degraded for future operation.

I am interested in developing agents that can recognize these long-term environmental consequences while interacting continuously with evolving real-world environments. This includes reasoning about clutter accumulation, accessibility preservation, obstruction avoidance, manipulation feasibility, and environmental organization over extended deployments. Rather than evaluating actions only through immediate task completion, I am interested in agents that reason about how actions alter long-term environmental usability.

A concrete scenario is a household robot operating continuously within a kitchen environment. Over days or weeks of deployment, the robot repeatedly performs tasks such as clearing dishes, organizing groceries, retrieving objects, and rearranging shared workspaces. A short-horizon agent may gradually leave behind cluttered countertops, obstructed cabinets, and increasingly disorganized storage structure. In contrast, an agent reasoning over long-term environmental quality may maintain cleaner workspace configurations, preserve reachability within cabinets, reduce object interference, and continuously reorganize environments into states that remain favorable over time.

Research Direction 2: Semantic Reasoning and Opportunistic Anticipatory Behavior

Real-world environments contain rich semantic and contextual structure that influences how environmental configurations evolve over time. Humans naturally use commonsense reasoning to maintain favorable environmental structure, often performing small opportunistic actions that improve long-term usability of the environment. For example, while clearing a table, a person may place dishes near the sink, reorganize cluttered objects into appropriate locations, or clear obstructed workspace even when those actions are not part of the immediate objective. These behaviors emerge from an implicit understanding of recurring environmental patterns and long-term environmental utility.

I am interested in studying how embodied agents can acquire similar forms of semantic and opportunistic reasoning. In particular, I am interested in how foundation models, predictive world models, and learned environmental representations can help robots recognize recurring patterns within environments and reason about how environments are likely to evolve over time. Large language models and vision-language models provide access to broad semantic and commonsense knowledge that may help embodied agents interpret contextual cues, infer latent environmental structure, and recognize favorable or unfavorable environmental configurations.

For example, chopped vegetables on a cutting board, cookware on a stove, and objects scattered across a countertop may indicate active meal preparation and increasing workspace clutter. A robot reasoning over these contextual cues may reorganize workspace opportunistically while performing another task, preserving free space and accessibility before clutter accumulates further. Similarly, repeated observation of blocked cabinet access or cluttered navigation routes may allow an agent to recognize recurring environmental failure patterns and proactively reorganize the environment to reduce future obstruction.

More broadly, I am interested in grounded embodied systems that combine semantic reasoning with physical understanding of accessibility, manipulation feasibility, navigation structure, and environmental organization. This direction investigates how embodied agents can move beyond isolated task execution toward commonsense environmental reasoning that continuously maintains favorable environmental structure over long deployments.

Research Direction 3: Self-Improving Embodied Agents and Learned Environmental Priors

Persistent deployment creates opportunities for embodied agents to improve through experience. A robot operating continuously within a home, warehouse, or laboratory repeatedly encounters similar layouts, workflows, environmental degradation patterns, and organizational failures. Over time, the agent should learn which environmental configurations remain favorable, which arrangements repeatedly introduce obstruction, and which anticipatory behaviors preserve long-term environmental quality.

I am interested in studying how predictive models, reinforcement learning, and other deep learning approaches can help embodied agents acquire long-horizon environmental reasoning through interaction. Reinforcement learning may allow agents to discover environmental behaviors that reduce long-term obstruction, preserve accessibility, and maintain favorable workspace structure across extended deployments. Predictive models may help agents estimate how environmental structure evolves over time and identify actions that prevent gradual degradation of environmental quality.

For example, a robot repeatedly organizing a kitchen cabinet may gradually learn placement strategies that preserve reachability and reduce interference between frequently used objects. A navigation robot operating in cluttered environments may learn obstacle relocation behaviors that preserve navigable free space over long horizons. More broadly, I am interested in how embodied agents can accumulate learned environmental priors that improve long-term reasoning through persistent interaction with real-world environments.

Research Direction 4: Multi-Agent and Human-Centered Environmental Coordination

Persistent environments are often shared by multiple robots and humans whose interactions continuously reshape environmental structure. A decision that appears locally effective for one agent may degrade accessibility, usability, or navigability for others. This becomes especially important in warehouses, hospitals, factories, and assistive living environments where many agents interact within shared physical spaces over long periods of time.

I am interested in studying how embodied agents can reason about environmental modifications in shared settings. For example, a warehouse robot relocating inventory may unintentionally obstruct another robot's future retrieval path. Similarly, a household robot reorganizing objects for one user may reduce accessibility for another occupant. These environments require agents to reason not only about isolated task completion, but also about how environmental changes influence future operation for multiple humans and agents simultaneously. This direction investigates how robots can coordinate through persistent environmental structure itself. More broadly, I am interested in embodied systems that maintain environments that remain interpretable, navigable, and usable across many interacting agents over extended deployments.

Long-Term Vision

My long-term goal is to develop embodied AI agents that act with sustained awareness of how present actions reshape future. I believe the next generation of embodied AI systems will require persistent world models, environmental memory, long-horizon reasoning, continual adaptation, and grounded understanding of environmental structure. More broadly, I believe intelligence in embodied systems should not be measured solely by immediate task completion, but also by the ability to continuously maintain environments that remain organized, accessible, navigable, and favorable for long-term operation. This perspective connects long-horizon reasoning, embodied intelligence, self-improving systems, and grounded agentic AI through a common challenge: enabling embodied agents to continuously shape and maintain environments that remain favorable across prolonged real-world interaction.