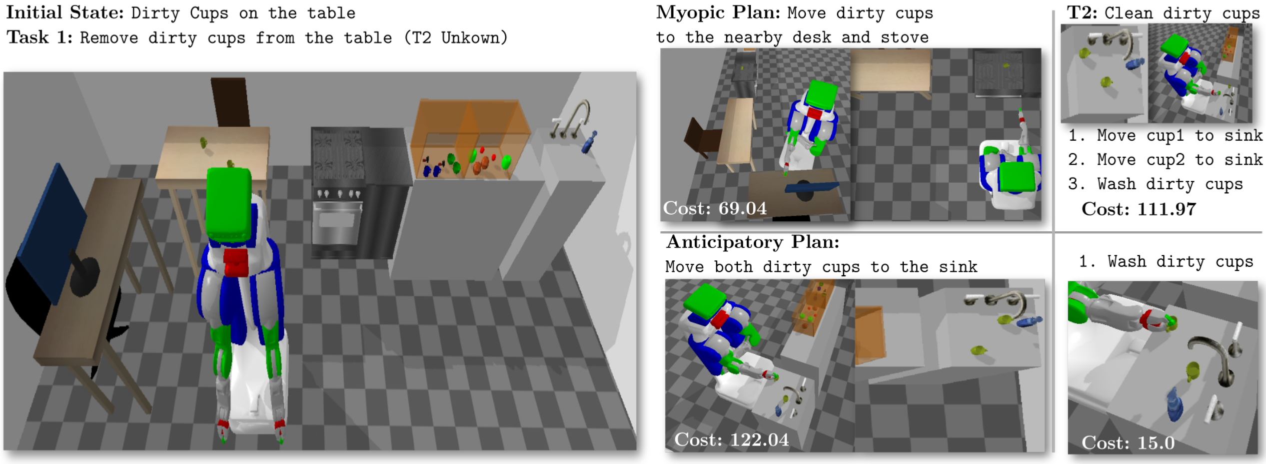

Figure 1: *Anticipatory vs. myopic planning.* Asked to remove dirty cups from the table (Task 1), a myopic planner moves them to the nearest free surfaces, completing the task cheaply but forcing costly rework when a later, initially unknown task (Task 2: clean the cups) arrives. An anticipatory planner instead moves both cups directly to the sink; the immediate action costs slightly more, but the environment it leaves behind makes the subsequent task far cheaper, lowering total long-horizon cost.

My research develops embodied AI agents that reason about how their present actions reshape the future. As robots move from isolated demonstrations toward continuous deployment in homes, warehouses, factories, and assistive settings, intelligence can no longer be measured by immediate task completion alone. Every interaction alters object arrangement, free space, and accessibility, and these changes accumulate: a robot may complete each task it is given while gradually leaving behind cluttered workspaces, blocked paths, and disorganized storage that make every subsequent task harder. Humans avoid this failure mode instinctively — moving dishes toward the sink rather than leaving them scattered, keeping warehouse aisles clear while shelving, leaving a caregiver's frequently used items within a resident's reach — not because each act serves an explicit future objective, but from an implicit understanding that environmental structure matters over time.

Most planners deployed today are myopic: they complete each task in isolation and, lacking any model of what will be asked next, unwittingly introduce side effects that impede future tasks as shown in myopic plan in Fig. 1. My central methodological stance is that this is not a reason to discard existing planners and design new ones from scratch. My Ph.D. work instead endows these planners with foresight, formalizing the missing capability as *anticipatory planning*: decision making that augments task objectives with learned estimates of the expected cost of future tasks, so that behaviors like decluttering, accessibility preservation, and preparation emerge as consequences of long-horizon reasoning rather than hand-coded rules. Concretely, a learned future-cost model guides which plan the underlying planner commits to, giving the planner anticipation without altering the guarantees it already provides. In my ICRA 2023 paper [1], I introduced this framework for long-lived agents, showing that learned cost models over future task distributions substantially improve long-term performance over myopic planners(Fig. 1). My IEEE RA-L 2026 paper [2] brings anticipatory planning to task and motion planning (TAMP) for rearrangement in persistent continuous-space environments, where a robot organizing objects (e.g., inside a household `cabinet`; Fig. 2) must preserve reachability and avoid creating obstructions for later retrieval. I have studied the same principle in navigation among movable obstacles (NAMO), where relocating a box to the nearest feasible spot solves the immediate problem while gradually blocking aisles. Across these domains, the unifying insight is that local decisions continuously reshape the environment, and a planner endowed with a model of future tasks can keep that environment in states that remain easy to work in.

Building on this foundation, my in-submission work further advances anticipatory TAMP in two ways. The first shows that anticipation is inherently integrated across abstraction levels: the long-term impact of a decision depends jointly on symbolic structure, such as which objects are moved and in what order, and geometric state, such as where those objects are placed [3]. The second advances anticipatory TAMP by revealing an emergent computational property of anticipatory cost estimates. Although these estimates are developed to reason about the long-term impact of actions in persistent deployments, they can also be used as a learned prior for solving a single TAMP problem in isolation. By guiding sampling toward configurations that are more promising under the anticipatory model, this prior reduces search and sampling computation without modifying the underlying planner objective [4].

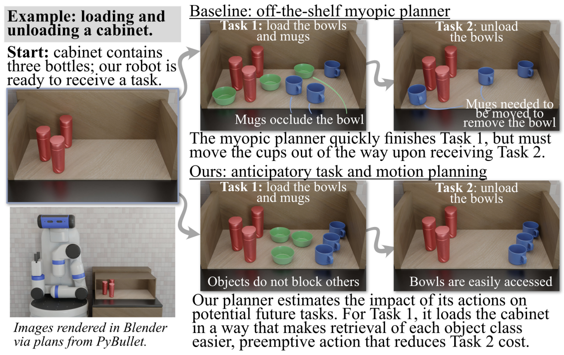

Figure 2: *Anticipatory* TAMP *in a cabinet (RA-L 2026).* A myopic planner loads the `bowls` and `mugs` wherever is quickest, so the `mugs` occlude the `bowls` and must be moved to unload them later. My anticipatory planner estimates how its placements affect plausible future tasks and loads the cabinet so each object class stays retrievable.

Research Direction 1: Foundation Models for Semantic World Models and Underspecified Tasks

Anticipatory planning requires a distribution over future tasks and a specification of goal states, and in the real world neither is given. Instructions such as "`tidy up the kitchen`" or "`put this away`" leave the goal, the relevant objects, and acceptable placements implicit, and the tasks an environment will demand next are never written down. Large language models (LLMs) and vision-language models (VLMs) offer exactly the commonsense and semantic priors needed to fill these gaps: interpreting contextual cues, inferring latent goals from vague instructions, and hypothesizing plausible future tasks from the current scene. Given "`put this away`" for a `mug`, an agent should combine visual context with environmental memory to select a reachable shelf near the coffee machine; observing chopped vegetables and cookware on a stove, it should infer that counter space will soon be needed and opportunistically clear it.

Foundation models, however, are not reliable planners: their outputs are ungrounded in geometry and physical feasibility, and a semantically sensible suggestion may be unreachable, unstable, or obstructive. My position is therefore not to plan with foundation models directly, but to use them to *build the components anticipatory planning depends on* — semantic world models, task-distribution priors, and goal specifications — while decision making remains with principled planners whose long-term consequences are verified rather than assumed. A VLM can predict how a kitchen is likely to be used; a model-based anticipatory planner then selects actions against that prediction. This separation keeps commonsense knowledge where it is strong (interpretation, prediction, hypothesis generation) and planning where it must be trustworthy, and it equips anticipatory planning for settings where objectives are implicit, evolving, and user-specific.

Research Direction 2: Anticipation at the Skill Level with Vision-Language-Action Models

Anticipation should not live only in high-level task planning; it must also be embedded in the low-level skills that physically realize plans, since *how* a robot grasps, places, or stacks often determines the environmental state it leaves behind more directly than *which* task it chose. A placement policy that sets a `pan` down with its handle protruding into a walkway succeeds at "`place the pan`" while creating a hazard; a grasp that takes a `mug` by its handle forecloses a clean handover; a stacking policy that buries a frequently used `bowl` completes the rearrangement while degrading future access. Vision-language-action models (VLAs) now make it possible to train such skills end-to-end, raising a question I find compelling: can anticipation be learned *inside* the skill itself?

I plan to train and fine-tune VLA policies with objectives that reflect downstream environmental quality — augmenting imitation and reinforcement learning losses with learned estimates of future accessibility, obstruction, and manipulation cost, and distilling anticipatory preferences discovered by high-level planners into low-level policies. This creates a two-way bridge between planning and modern robot learning: anticipatory planners provide long-horizon supervision signals for skills, and anticipatory skills reduce the burden on planners by making individual actions less environmentally harmful by default.

Research Direction 3: Self-Improving Agents that Learn Anticipatory Models from Deployment

Persistent deployment is not only the setting that demands anticipation — it is also the data source from which anticipation can be learned. A robot operating in the same home or warehouse for weeks repeatedly encounters the same layouts, routines, and organizational failures, and its own history records which configurations remained favorable and which repeatedly caused obstruction. I plan to develop agents that close this loop: using deployment experience to refine the learned future-cost models at the heart of anticipatory planning, to update task-distribution beliefs online from observed routines (e.g., learning that dishes accumulate before evening cleanup in a particular household and preparing the workspace accordingly), and to fine-tune the VLA skills of Direction 2 from the environmental successes and failures the agent itself observes. Because real future task distributions are never explicitly available, this continual, self-supervised adaptation is what makes anticipatory planning deployable beyond simulation, and it offers a concrete path toward self-improving embodied agents.

Research Direction 4: Multi-Agent Coordination and Benchmarks for Persistent Autonomy

Real deployments are shared: warehouses, hospitals, and homes contain multiple robots and humans whose interactions jointly reshape environmental structure, and a decision that is locally effective for one agent may degrade accessibility for others — one robot's convenient staging location is another's blocked retrieval path. I am interested in anticipatory planning for this multi-agent setting, where the future task distribution an agent must anticipate includes the tasks of *other* agents and occupants, and where coordination can occur through the persistent environmental structure itself: maintaining shared spaces that remain interpretable, navigable, and usable across many interacting agents. My in-submission work takes a first step in this direction: it builds upon anticipatory planning so that each robot acts *courteously*, accounting for how its own rearrangements affect the future tasks of other robots sharing the environment, and improves long-lived planning across the team rather than for any single agent in isolation [5].

Progress on all of these questions is currently limited by evaluation. Existing embodied AI benchmarks are episodic — the environment resets between tasks — so the long-term consequences of decisions are never measured, and an agent can score perfectly while exhibiting exactly the short-horizon behavior that undermines real deployments. I plan to build simulation benchmarks in which state persists across procedurally generated task streams over shared household and warehouse scenes, with cumulative metrics: total planning and execution cost over the stream, accessibility degradation over time, frequency of self-induced obstruction, and success-rate drift as deployments lengthen.

Long-Term Vision

My long-term goal is embodied agents that act with sustained awareness of how present actions reshape the future, combining persistent world models, environmental memory, continual adaptation, and grounded semantic understanding from foundation models around anticipatory planning as the decision-making core. Intelligence should be measured not by immediate task completion but by the ability to keep environments organized, accessible, and favorable for long-term operation — a challenge that unites anticipatory planning, foundation models, self-improving systems, multi-agent coordination, and rigorous benchmarking.

References

[1] R. Dhakal, M. R. H. Talukder, and G. J. Stein, "Anticipatory planning: Improving long-lived planning by estimating expected cost of future tasks," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

[2] R. Dhakal, D. M. Nguyen, T. Silver, X. Xiao, and G. J. Stein, "Anticipatory task and motion planning: Improved rearrangement in persistent continuous-space environments," *IEEE Robotics and Automation Letters (RA-L)*, vol. 11, no. 2, pp. 1850–1857, 2026.

[3] R. Dhakal and G. J. Stein, "Integrated symbolic–geometric anticipation for long-lived task and motion planning." In submission, 2026.

[4] R. Dhakal, M. R. Hossain Talukder, and G. J. Stein, "Anticipation as prior for efficient sampling in task and motion planning." In submission, 2026.

[5] M. R. Hossain Talukder, R. Dhakal, E. Phillips, and G. J. Stein, "Courteous anticipation: Improving long-lived planning for multiple robots in shared environments." In submission, 2026.