arXiv:2611.18359v1 [cs.RO] 13 Nov 2026

Envisioning Agentic Robotics: A New Era of Robotic Systems

Grid Robotics

Abstract

The release of GPT-6 Astra marks the emergence of a new paradigm in robotics: *agentic robotics*. Within weeks, frontier models with agentic harnesses matched or exceeded task-specific policies on many embodied benchmarks. Instead of training a dedicated policy for each family of tasks, an agentic robotic system applies a foundation model to open-world embodied tasks. It operates in a closed loop: observing the environment, interpreting user intent, reasoning about the task, invoking skills and tools, verifying outcomes, and adapting its strategy. Despite these advances, the community has yet to converge on a shared definition, a reliable evaluation framework, or a research roadmap. We argue that *Embodied System 2*, realized by a foundation model together with an embodied harness, forms the decision-making core of agentic robotics. Its central challenge is to translate general-purpose reasoning into effective and verifiable actions under physical constraints such as limited reaction time, mechanical limits, unpredictable environment changes, and irreversible consequences. Motivated by this perspective, we propose an evaluation framework that measures both model capabilities and closed-loop system performance. Furthermore, we identify the principles that carry over from digital to physical agents and the challenges that must be addressed to achieve reliable physical autonomy. On top of the Embodied System 2 and the evaluation framework, we envision a research roadmap toward sustained autonomy: the ability to perform work safely and reliably with progressively broader responsibility, from executing individual primitives to fulfilling standing roles.

Website: https://grid-robotics-official.github.io/agentic-robotics-roadmap
Contact: contact@gridrobotics.space

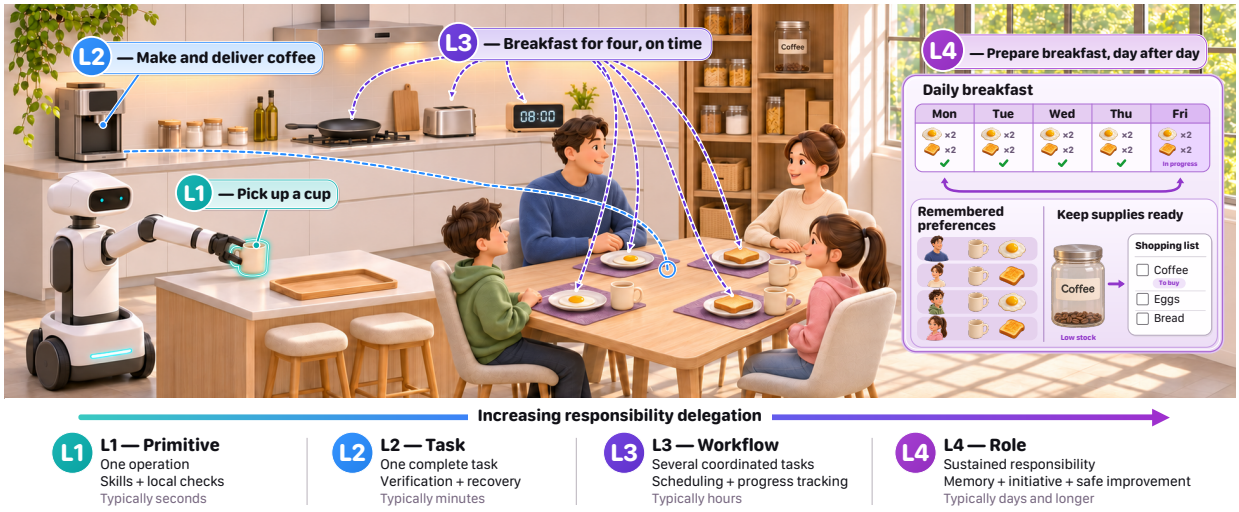


Figure 1 | **A vision for agentic robotics: from executing tasks to fulfilling roles.** A home breakfast scenario illustrates progressively broader responsibility entrusted to robots: from executing a primitive (L1), completing a task (L2), coordinating a workflow (L3), to fulfilling a standing role (L4).

Contents

1 Introduction **4**

2 Overview of Agentic Robotics **6**

2.1 Reference Architecture 6

2.2 Agentic Loop: From Intent to Verified Outcome 9

3 Embodied System 2: Foundation of Agentic Robotics **11**

3.1 Definition 11

3.2 S2 Foundation Model 13

3.3 Embodied Harness 15

3.4 Constraints Imposed by the Physical World 18

4 Evaluation Framework **20**

4.1 What to Evaluate 20

4.2 How to Evaluate 29

4.3 What Current Results Show 31

4.4 Limitations of Current Evaluation 35

5 Digital Agents versus Physical Agents **36**

5.1 What Digital Agents Can Do and Why 37

5.2 Which Assumptions Fail in the Physical World 38

5.3 What Transfers 41

5.4 What Must Be Rebuilt 41

6 Road Ahead **43**

6.1 Ladder of Delegable Work 44

6.2 Future Research Directions 45

7 Conclusion **48**

A Related Work **50**

A.1 Surveys and Position Papers 50

A.2 Language-Model-Based Robot Agents 50

A.3 Hierarchical and Dual-System Robot Foundation Models 52

A.4 Embodied Reasoning Models 53
A.5 Agent Harnesses 53
A.6 Memory and World Models 54
A.7 Benchmarks 54
A.8 Self-Improvement and Continual Learning 55

B Applications: What Agentic Robotics Changes **55**
B.1 Homes 56
B.2 Manufacturing and Logistics 56
B.3 Inspection 58
B.4 Healthcare and Laboratories 58
B.5 Agriculture and Mobility 58

C Contributors **59**

1. Introduction

A central goal of robotics is to enable robots to translate human intent into desired outcomes in the physical world. Early autonomous robots often integrated hand-engineered components for perception, planning, and control within modular or layered architectures (Nilsson, 1984; Gat et al., 1998). Since then, advances in deep learning have allowed robots to learn increasingly complex behaviors. Policies learned in simulation, mostly with deep reinforcement learning, have produced agile locomotion (Hwangbo et al., 2019) and versatile humanoid whole-body control by tracking human motion data (Liao et al., 2025; Luo et al., 2026c; Liao et al., 2026). However, these simulation-trained policies remain specialized to specific tasks and embodiments, limiting their ability to generalize across diverse and open-ended environments. To move beyond task-specific controllers and enable robots to follow human instructions, the field has increasingly turned toward models that unify perception, language understanding, and action generation within a single framework. Among them, vision-language-action (VLA) models, exemplified by the π series, combine vision-language pretraining with diverse robot data to learn general policies that map visual observations and language instructions to actions (Black et al., 2025a,b). More recently, world-action models (WAM) integrate video-based modeling of physical dynamics with robot policy learning (Ye et al., 2026a; Yuan et al., 2026b; Li et al., 2026c). Together, these advances have broadened robot capabilities. Yet reliable generalization remains a central challenge, as policy performance can still deteriorate when scenes or instructions change (Zhou et al., 2025c).

To address this generalization challenge, a major line of research focuses on scaling and diversifying training data for VLA and WAM policies (Zheng et al., 2026; Dyna Robotics, 2026a). These efforts draw on diverse data sources, including real-robot data (Black et al., 2025a), UMI demonstrations (Chi et al., 2024), egocentric human data (Zheng et al., 2026), and multimodal web data (Brohan et al., 2023). The release of GPT-6 Astra (OpenAI, 2026a,b) and its early applications to robotics offered a different perspective on this challenge. In an early evaluation, the model outperformed trained policies on many robotic benchmarks without task-specific fine-tuning (Zhang et al., 2026c). Subsequent work explored its use with learned policies and robotic tools (Galbot Team et al., 2026; Gao et al., 2026b). These results suggest that generalization can also be achieved by relying on a general-purpose model's knowledge and reasoning to organize and adapt behavior during execution. This perspective motivates a paradigm shift from *policy-centric robotics* to *agentic robotics*, by shifting the focus from learned action policies (Figure 2a) to an agent that pursues goals by coordinating skills and tools, evaluating outcomes, and adapting its decisions (Figure 2b).

Despite this progress, the community still lacks a shared understanding of agentic robotics or a reliable evaluation framework. The scope of the paradigm, its relationship to existing forms of robot autonomy, and the organization of its constituent functions remain insufficiently clarified (Qin et al., 2026a). Furthermore, differences in experimental conditions and system configurations also make it difficult to compare results across studies. For example, RoboDojo evaluations of Astra vary in the task-knowledge sources they use and the action budgets they permit (Zhang et al., 2026c; SenseNova et al., 2026); separately, Harness VLA shows substantial performance gains from execution memory when the planner and VLA are held fixed (Zhang et al., 2026f). All of these results underscore the need to establish what constitutes the agentic system, which capabilities are being assessed, and under what conditions comparisons are meaningful. These conceptual and methodological challenges call for a unified framework that clarifies the agentic workflow by specifying the functional roles of its components and how they interact. Such a framework can provide a shared conceptual foundation for characterizing agentic robotics, developing valid and comparable evaluations at the model and system levels, and formulating a research roadmap.

To address this, we propose a framework of agentic robotics that distinguishes three functional

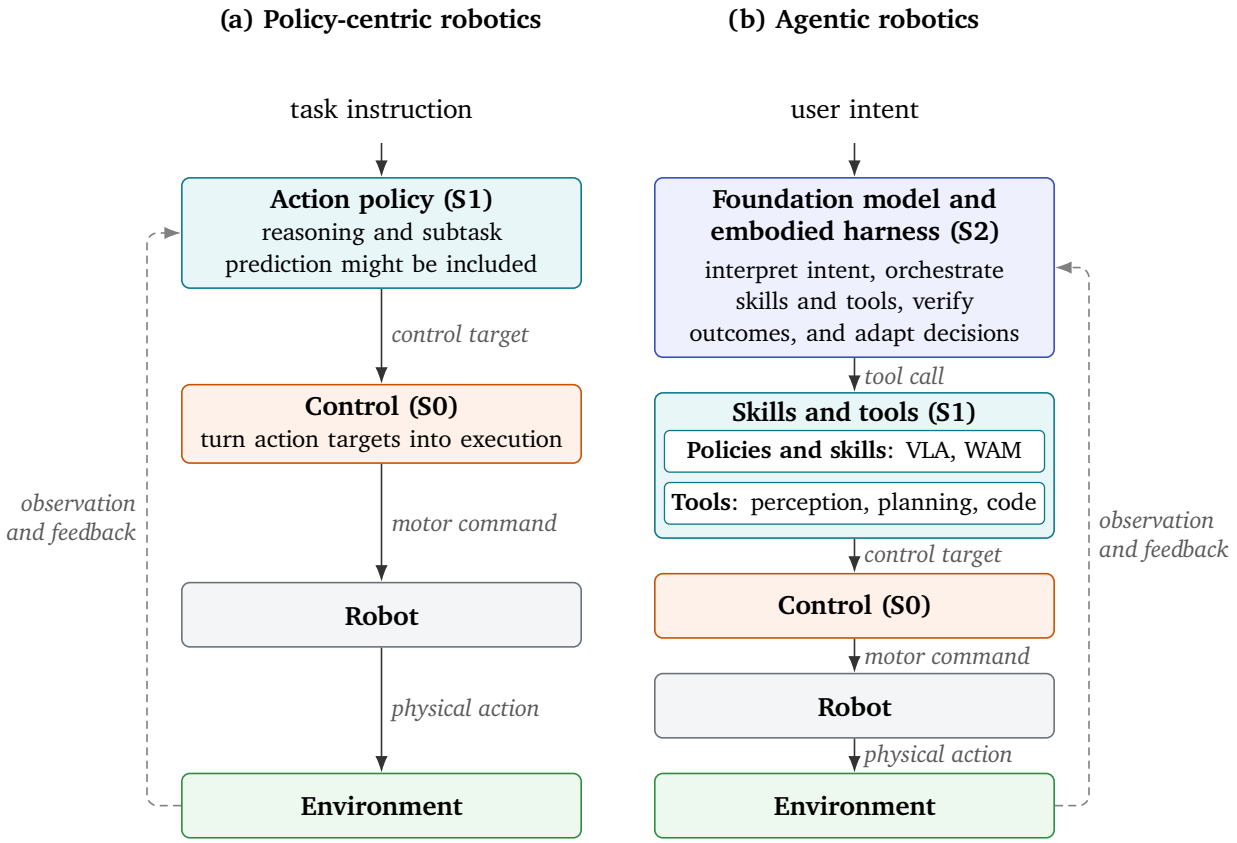


Figure 2 | **Two paradigms.** (a) An action policy maps the task instruction and observation to action targets. It might include reasoning, subtask prediction, or in-context adaptation. (b) S2 combines a foundation model with an embodied harness. Together, they interpret intent, orchestrate skills and tools, verify outcomes, and adapt subsequent decisions. S1 provides callable skills and tools. S0 performs body control. The S1 path is illustrative; S2 may also call S0 or external tools.

roles (Figure 3). System 2 (S2) interprets intent, reasons about goals and constraints, orchestrates skills and tools, and verifies outcomes to guide subsequent decisions. System 1 (S1) provides skills and callable modules, while System 0 (S0) supplies the body interface and real-time control. Specifically, S2 is the decision-making core of agentic robotics, realized by a foundation model together with an *embodied harness*. The foundation model supplies knowledge, understanding, and reasoning, while the harness organizes observations, maintains plans and memory, exposes skills and tools, and supports execution and verification. Learned policies, perception layers, motion planning modules, and generated programs are included in S1. S2 may orchestrate these capabilities or send action targets directly to S0, while retaining responsibility for verifying outcomes and adapting its decisions. With these roles defined, we investigate how the foundation model and harness should work together under physical constraints, and how to evaluate their contributions to closed-loop system performance.

Building on this division of responsibilities, we first explain how S2, S1, and S0 work together (Section 2). We then articulate S2 and examine its foundation model, its harness, and the physical constraints under which they operate (Section 3). Next, we develop an evaluation framework across nine dimensions that distinguishes model capabilities from whole system performance, and use it to assess the available evidence (Section 4). For a closed-loop system, the framework emphasizes independent outcome verification, transparent reporting, and evaluation while the environment continues to evolve. Furthermore, we examine which design principles carry over from digital agents

to physical systems and which assumptions must be reconsidered in the physical world (Section 5). These analyses inform a roadmap organized around four levels of delegable work, from executing primitives (L1), completing tasks (L2), coordinating workflows (L3), to fulfilling standing roles (L4) (Figure 1; Section 6). The roadmap connects this progression to four research directions, including continual improvement from physical experience. Finally, we conclude with several guidelines to verify and evaluate the progress of agentic robotics.

2. Overview of Agentic Robotics

We adopt a system-level view of agentic robotics: an agent interprets the user's intent, decides how to proceed, orchestrates skills and tools at run time, and verifies their physical effects against the intended goal. Feedback from verification guides subsequent decisions, enabling the system to adapt its actions, recover from failures, and seek clarification or assistance when needed. We first present a reference architecture that defines the responsibilities and interfaces of S2, S1, and S0. We then go through the agentic loop and examine three system-level properties: operation across multiple timescales, responsiveness during deliberation, and state persistence across tasks.

2.1. Reference Architecture

In general, an agentic robot requires three functional roles: S2 interprets intent, reasons about goals and constraints, orchestrates capabilities, and verifies outcomes; S1 provides skills and callable modules; and S0 provides the body interface and local real-time control. Note that this decomposition is functional: one model may serve several roles, and one role may span several models or modules. Figure 3 shows the three systems and their interfaces, together with the user, the robot body, the physical world, and a real-time event monitor. The systems typically run on different timescales, but timing alone does not distinguish them. We describe them from S2 downward, focusing on their responsibilities and interactions.

System 2: Intent, Reasoning, Orchestration, and Verification. S2 is the functional layer that interprets intent, reasons about goals and constraints to make decisions, orchestrates skills and tools accordingly, and verifies outcomes to adapt subsequent decisions toward the intended goals. It decides whether to act, gather information, ask for clarification, continue, retry, replan, or stop. When acting, it chooses a capability and specifies the call, which may invoke an S1 skill or a tool, or send an action target directly to S0. It checks the returned evidence against the intended goal before deciding on the next subtask. In hierarchical systems such as Hi Robot (Shi et al., 2025) and Gemini Robotics (Gemini Robotics Team, 2025a,b, 2026b), a high-level vision-language model contributes to the orchestration: it interprets instructions and feedback before delegating the next subtask to an action model.

In our framework, S2 is realized by a foundation model and an embodied harness. The *S2 foundation model* supplies knowledge, understanding, and reasoning, and proposes the next decision from the context it receives. The *embodied harness* supplies the runtime that assembles the model's context and turns its decisions into constrained, observable, and verifiable execution. Concretely, it has five core parts. A *context manager* assembles current observations, task progress, interaction history, and the state of the world, the user, and the body. A *tool interface* exposes S1 skills, permitted S0 calls, code execution, and other tools to S2. A *planner* maintains and revises task plans given the foundation model's proposals or dedicated planning routines. A *memory* stores, retrieves, updates, and removes information beyond the current context. A *verifier* judges from evidence whether a step achieved its subgoal and whether the task is complete. Existing systems implement some of these parts: Inner Monologue (Huang et al., 2023b) feeds scene and skill-success signals back into planning; SCOPER (Wei et al., 2026a) updates an evidence memory and verifies progress to guide recovery.

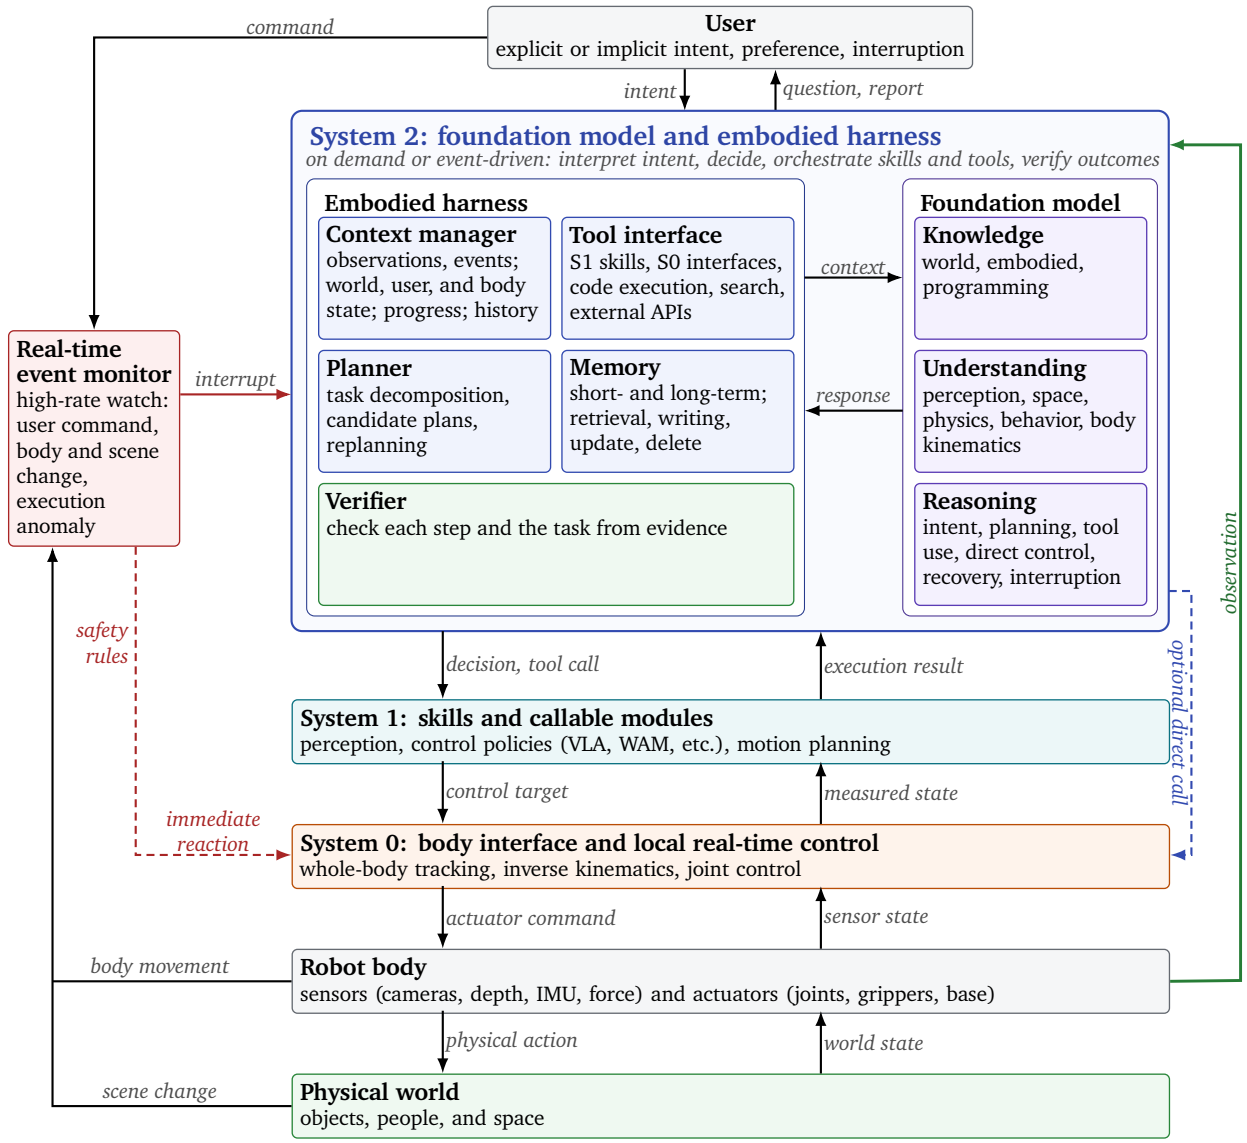


Figure 3 | **A reference framework for agentic robotics.** An Embodied S2 consists of a foundation model and an embodied harness, whose five core parts manage context, tools, planning, memory, and verification. S2 runs on demand or in response to events, whereas S1 and S0 execute at their own rates. A real-time event monitor outside the three systems watches the user, the robot body, and the physical world at a high rate. It interrupts S2 when an event may change the task and, under predefined safety rules, triggers an immediate reaction in S0 without waiting for S2. S2 may call S0 directly through the tool interface; S0 keeps its own real-time loop.

System 1: Skills and Callable Modules. S1 provides skills and callable modules that perform the functions S2 delegates, including action execution, perception, and motion planning. Specifically, *action execution* modules map observations to actions within the scope of a subtask, such as VLA policies (Black et al., 2025b,a; NVIDIA et al., 2025c), world-action models (WAMs) (Dyna Robotics, 2026b), locomotion policies, and navigation policies. *Perception* modules extract task-relevant estimates from sensor data through detection (Minderer et al., 2022), segmentation (Carion et al., 2025), pointing (Deitke et al., 2024), and pose estimation (Wen et al., 2024). *Motion planning* modules turn a goal into a feasible motion (Sundaralingam et al., 2023; Sundermeyer et al., 2021; Kim et al., 2025; Huang et al., 2025). A single call may combine several of these categories. For example, a grasp call may segment the object, select a grasp function, plan a path, and execute it. Programs written by S2 likewise serve as S1 skills when they execute delegated behavior (Yang et al., 2026j).

System 0: Body Interface and Local Real-Time Control. S0 is the embodiment-specific layer that executes control targets in a real-time loop. Examples are whole-body tracking and balance for a humanoid, gait tracking and stabilization for a quadruped, inverse kinematics and trajectory tracking for an arm, and chassis control for a mobile base. For humanoids, whole-body tracking is increasingly handled by general-purpose controllers—often called behavior foundation models—that are pretrained at scale (Yuan et al., 2025a), for example by learning to track large human-motion datasets (Luo et al., 2026c; Liao et al., 2026). S0 receives control targets from S1, or from S2 through a permitted action interface, and reports state estimates such as joint positions, contact, and pose. It maintains stability and enforces the body's control limits without waiting for S2 to deliberate.

Real-Time Event Monitor. A lightweight *real-time event monitor* runs alongside the three systems. At a high rate, it watches the user, the robot body, and the physical world for events that cannot wait for S2 to decide. Such events include a sudden command or interjection from the user, an unexpected body movement caused by a collision or fault, and a change in the scene such as a person entering the workspace. In practice, the monitor may use a small streaming vision-language model (VLM), rule-based detectors, or both. It interrupts S2 when an event may change the task. When an event demands an immediate reaction, it may apply its predefined safety rules and trigger a protective response in S0, such as stopping or slowing the current motion.

In general, commands flow downward and evidence flows upward across the three systems. In the downward direction, we divide the execution of a decision into three typical paths by how S2 reaches the body, as shown in Figure 4. In *layered orchestration*, S2 invokes S1 skills and tools, which produce the actions that S0 executes. In *direct control*, S2 bypasses S1 and controls S0 through an action interface, as several systems built on frontier models do (Zhang et al., 2026c; Yang et al., 2026j). Direct control does not place the foundation model inside the real-time control loop: S0 still runs its own local loop, and the action interface defines what S2 may send to it. In *code as policy*, S2 writes code that plays the role of S1 at run time (Liang et al., 2023); the code may execute several actions before S2 decides again. Note that these paths are not alternative architectures: one system may use all three within a single task. In the upward direction, S0 reports the state of the body. Each call should return the observations available at that point and an execution status, allowing S2 to assess the outcome. Outside this call chain, the event monitor signals S2 and S0 without being invoked.

The three systems also differ in how directly they depend on the body. S0 must account for the robot's dynamics and control interface, whereas S1 skills depend on the actions and observations that the body provides. An S2 strategy of looking in the drawer before the cabinet may transfer across bodies, but only if it remains appropriate for the new robot's capabilities, access constraints, and observations.

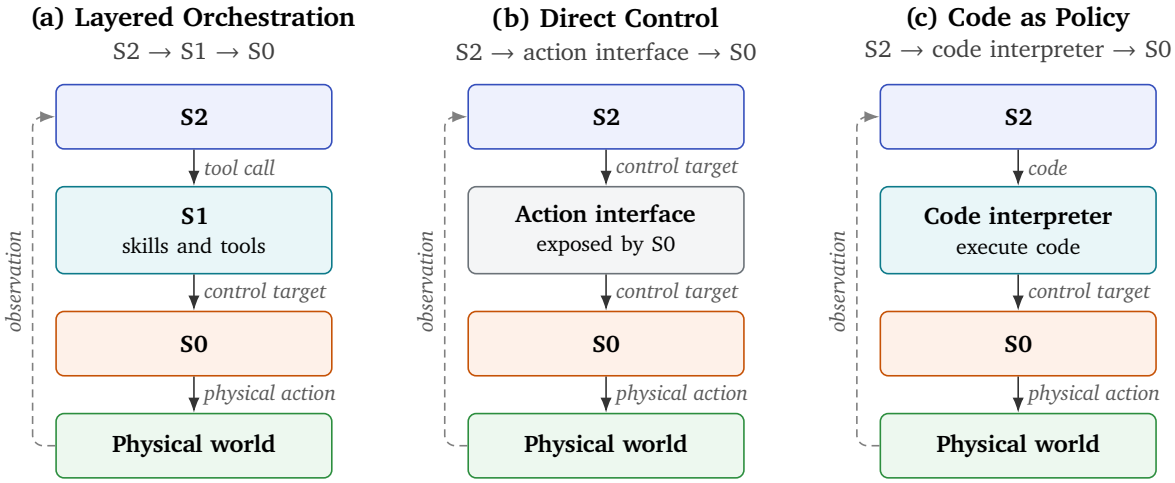


Figure 4 | **Three typical execution paths.** One system may use all three paths within a single task. For simplicity, the robot body is omitted. (a) Layered orchestration: S2 invokes S1 skills and tools, which produce control targets that S0 executes. (b) Direct control: S2 controls S0 directly through an action interface; S0 still runs its own real-time control loop. (c) Code as policy: S2 writes code that a runtime interpreter executes in the role of S1.

2.2. Agentic Loop: From Intent to Verified Outcome

Consider the request, "Find the red cup and bring it to the dining table." Given this request, S2 interprets it and plans four subtasks: locate, grasp, carry, and place the cup. It first invokes perception, and the harness records the cup's identity, location, and observation time. Then S2 invokes a grasp skill and checks the returned observation to make sure that the cup is held securely for transport. During transport, S1 and S0 maintain control while new observations inform S2's subsequent decisions. After placement, S2 verifies that the cup is supported by the table at the requested location. Each stage forwards state to the next, and the harness distinguishes a past event, "the cup was grasped", from a current state, "the cup is still held".

Suppose instead that the user interjects while the robot is grasping the red cup: "Not the red cup, the blue one." The event monitor detects the interjection and interrupts S2. S2 recognizes that the revised goal conflicts with the current action and requests that the grasp be paused or canceled. A cancellation request alone is insufficient: before dispatching a conflicting action, the harness must confirm that execution has stopped and that the body is in a suitable holding state. S2 then refreshes its observations and replans to retrieve the blue cup. Sometimes a skill reports success even though the action did not succeed. In this case, verification also catches unreported failures. For example, the grasping skill reports success on the blue cup, but the wrist camera shows that the cup is still on the shelf: the grasp was executed as commanded but missed. S2 attributes the miss to a grasp point that was occluded by the cup's handle. It therefore observes the cup from another side and retries with a different grasp rather than repeating the same call.

The example above traces six logical steps, which form a feedback cycle.

Step 1: Interpret Intent. The user expresses intent, which may be explicit ("put the mugs in the dishwasher") or implicit in preferences and context. S2 combines the request with the context assembled by the harness, such as a known preference that determines where a cup belongs. When the intent or the evidence is insufficient, a valid decision should be asking the user or observing again.

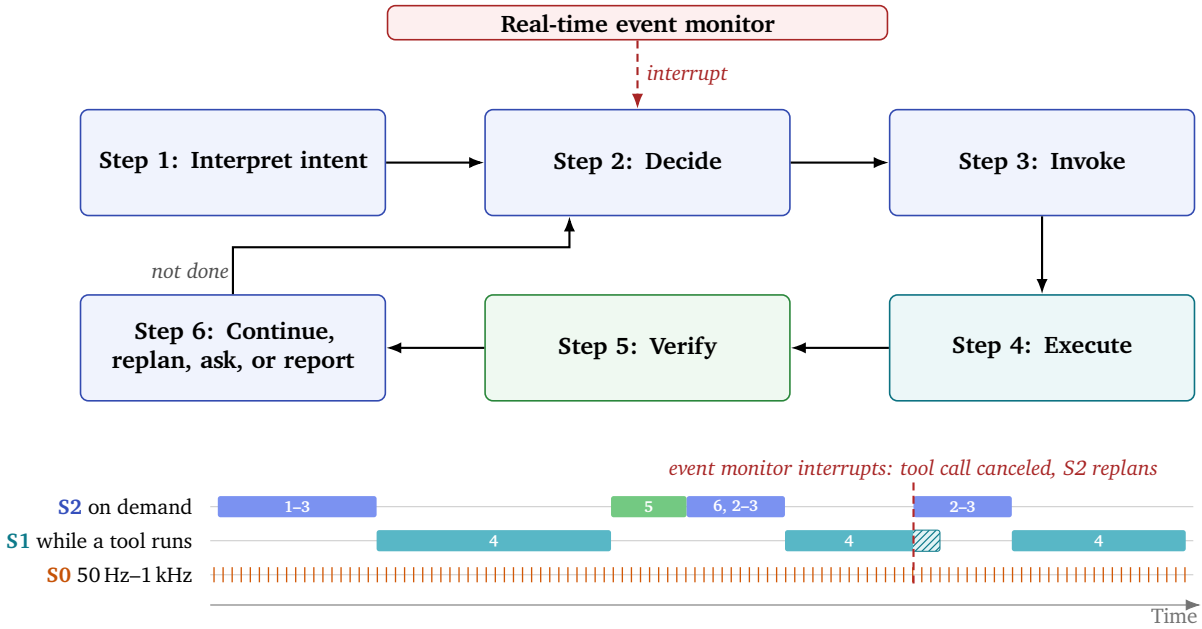


Figure 5 | **The agentic loop.** Top: one cycle comprises six steps. S2 interprets intent, reasons about goals and constraints to decide how to proceed, and invokes an S1 skill, an S0 interface, or an external tool as needed. The invoked capability executes under its contract. S2 then observes and verifies the outcome by answering four questions (whether the request was accepted, whether execution terminated, whether the physical effect holds, and whether the user's intent is satisfied). Based on the verification, it continues, recovers, asks, or reports. The real-time event monitor, which watches the user, the body, and the scene, can interrupt S2 at any point. Bottom: the same steps on a schematic time axis, labeled with their step numbers and colors. S2 deliberates on demand, S1 runs while a tool executes, and S0 runs continuously; after an interruption, S2 cancels the running tool call and replans while S0 keeps the body stable; the hatched segment marks the canceled execution.

Step 2: Decide. Given the current context and memory, S2 reasons about goals and constraints to decide how to proceed. It may select a subgoal or capability, specify a permitted control target, or supply a program.

Step 3: Invoke. S2 calls an S1 skill, a permitted S0 interface, or an external tool, with arguments that respect the tool's contract, such as reference frame, units, object identifiers, and preconditions.

Step 4: Execute. The invoked skill or tool carries out the call under its execution contract. Local control continues while S2 deliberates; a skill proceeds until it finishes, is paused or canceled, or a safety condition requires intervention. The event monitor continues to watch the user, the body, and the scene.

Step 5: Verify. S2 distinguishes four questions when assessing the outcome: was the request *accepted*, has the execution *terminated*, is there evidence that the intended physical *effect* holds, and does that effect satisfy the user's *intent*? Termination may indicate completion, failure, or cancellation; it does not by itself establish success. A successful tool return is only intermediate evidence: a grasp can be executed exactly as commanded and still miss the object.

Step 6: Continue, Replan, Ask, or Report. Completion is judged against the user's intent. If the task has not completed, S2 proceeds to the next subgoal or replans after a failure. If the task is complete, S2 reports; if the evidence is inconclusive, it observes again or asks the user.

Across the whole system, three properties should hold.

Multiple Timescales. The components typically operate on different timescales. S0 typically runs locally at roughly 50 Hz to 1 kHz (Luo et al., 2026c; Dyna Robotics, 2026b; Figure AI, 2026a). S1 control skills run continuously while a subtask executes, at rates that depend on the platform and the task. Some commercial stacks report rates from about 5 Hz to 200 Hz (Dyna Robotics, 2026b; Figure AI, 2025b; NVIDIA et al., 2025c). By contrast, perception and motion planning modules are invoked on demand and return results after their computation finishes. S2 does not run at a fixed rate: it could be invoked when a skill finishes or fails, when new observations arrive, when the event monitor reports an event, or when the user speaks. A frontier model in this role may take from several seconds to tens of seconds per decision (Xu et al., 2026d; Galbot Team et al., 2026), although this latency is an implementation property, not a defining feature of S2.

Continuous Responsiveness. While S2 is deliberating or a skill is executing, the physical world or the body state may change, and the user may revise an instruction. The event monitor should watch for these changes and remain responsive throughout deliberation and execution. Current systems achieve this responsiveness in different ways. Gemini Robotics-ER 2 treats a new prompt received during generation as an interruption of model reasoning and of pending tool calls (Google, 2026). MIRA combines streaming speech understanding with interruptible motion (Lin et al., 2026). ProAgent pairs a fast model that decides whether and when to intervene with a slower model that decides what to do (Yang et al., 2026g). In our architecture, the safety rules reside in the event monitor, so a protective stop can be requested in time.

Persistent State. The harness should keep a persistent state. We call one task an *episode* and a continuous period of use a *session*, which may contain many episodes. Within an episode, the harness tracks progress, object identities, and previous attempts. Across episodes and sessions, it maintains what the robot has learned about the environment, the user's preferences, and the body state. The record of the physical world must be reconstructed from observations and could become stale without notice, as when an object is moved. Persistent memory should therefore record when and how each item was obtained and allow revision, rather than treat a stored observation as a permanent fact.

3. Embodied System 2: Foundation of Agentic Robotics

The previous section described how S2 interacts with S1 and S0. In this section, we examine S2 itself. We first define S2 and give a criterion for identifying it in a robotic system. We then examine how the foundation model and embodied harness contribute to its eight capabilities and how physical constraints shape its operation.

3.1. Definition

We formally define Embodied System 2 as follows.

Definition (Embodied System 2). Embodied System 2 (Embodied S2) is the functional layer of an agentic robot that interprets intent, reasons about goals and constraints to make decisions, orchestrates skills and tools accordingly, and verifies outcomes to adapt subsequent decisions toward the intended goals.

In this paper, Embodied S2 refers specifically to an implementation in which a foundation model and an embodied harness jointly fulfill that role. Hereafter, we use Embodied S2 and S2 interchangeably and reserve *S2 foundation model* for the model inside Embodied S2 alone.

The above definition has three notable features. First, it is anchored in *intent* rather than in a

fixed task specification: the objective is a world state that satisfies the user. Reaching it may require clarifying an ambiguous request, inferring an unstated preference, or accommodating a change of mind. Second, it rests on *responsibility*: S2 makes high-level decisions, whereas S1 and S0 make action-level decisions within the scope that S2 delegates to them. Third, it includes the *harness*, because the state on which S2's decisions depend, namely that of the world, the user, and the body, is usually not supplied by the environment directly. This state must be assembled by the harness from partial, noisy, and aging observations.

Identifying S2. Labels such as System 2, embodied brain, planner, and orchestrator are applied inconsistently, so we identify S2 by functionality, not by name. A component, or a set of components, *acts as S2* if and only if it closes the loop of the definition, which requires two kinds of decision:

(i) *decisions on how to proceed*: interpreting the intent and, given the goals and constraints, either choosing the next subgoal and the skill or tool that should achieve it, together with the arguments of the call, or choosing to continue, retry, replan, ask, or stop; and
(ii) *decisions on outcomes*: judging from evidence whether the previous step, or the task as a whole, achieved its goal, so that the next decision can adapt accordingly.

Both kinds of decision must also be *identifiable from outside*: either the component's output interface exposes them, or a controlled replacement experiment reveals them. A component that makes only one kind of decision implements part of S2 but does not act as S2 on its own. For example, a planner that chooses and sequences skills without checking their effects orchestrates but never verifies. By contrast, a success detector makes neither kind of decision: it scores an observation, but whether that score settles the outcome of a step or of the task, and what follows from it, is decided elsewhere. State estimators in S0, perception modules in S1, world models, and the real-time event monitor likewise supply evidence without deciding what it means for the task: they inform S2 but do not act as S2. Because the criterion concerns function, it applies to any implementation, including classical planners and behavior trees.

Borderline Cases. In dual-system VLAs whose upper module emits only a latent representation, such as Helix (Figure AI, 2025b) and GR00T N1 (NVIDIA et al., 2025c), high-level decisions cannot be read from the interface. We therefore do not classify the upper module as S2 unless replacement experiments show otherwise. Conversely, a model that first emits a subtask and then the actions to achieve it, such as $\pi_{0.5}$ (Black et al., 2025a), or G0.5 with its chain-of-thought stream enabled (Liu et al., 2026e), makes an identifiable decision on how to proceed. Its subtask stage therefore implements part of S2, even though it shares weights with S1. As the authors of Hi Robot note, "the separation of roles at the model level is not fundamental to this design" (Shi et al., 2025).

Decision frequency is not part of the criterion. A model trained for the S2 role may decide about once per second (Shi et al., 2025), whereas a frontier model queried over a network may take tens of seconds per decision (Galbot Team et al., 2026). Either acts as S2 if it makes both kinds of decision. Speed determines whether an S2 is useful in a given setting, so we treat it as a physical constraint rather than as part of the definition. Nor is the granularity of the outputs part of the criterion. A component that emits end-effector targets or programs acts as S2 only if it also makes both kinds of decision, and an S2 that orchestrates skills remains S2 when it emits a pose directly (Figure 4).

Relation to Existing Definitions. Our definition builds on recent proposals that describe S2 as the embodied "brain" that selects skills and checks their effects (Luo et al., 2026b; Wei et al., 2026a), and on work that distinguishes a model from the runtime around it (Liang et al., 2026b; Huang et al., 2026c; Wang et al., 2026e). Our definition differs from these works in two respects: it provides an explicit, testable criterion, and it counts the harness as part of Embodied S2.

3.2. S2 Foundation Model

The S2 foundation model turns the context assembled by the harness into decisions. It may be a general frontier model used without embodied training, an embodied foundation model post-trained for embodied tasks, or the reasoning stage of a larger model that also produces actions. In each case, it inherits broad knowledge and reasoning from large-scale pretraining. However, its role differs from that of a chatbot or coding agent, since its decisions are carried out by a body in the physical world.

Doing the work that the definition assigns to S2 in the physical world requires eight abilities (Figure 6). Knowledge, understanding, and reasoning reside mainly in the foundation model; tool use, memory, cross-embodiment, safety, and efficiency arise from the foundation model and the embodied harness together. For each ability, we describe what it is, why S2 needs it, and how the foundation model's share of it can be cultivated.

Knowledge. Knowledge comprises the general facts, concepts, and procedures on which S2 draws. It matters because S2 must fill the gaps that a request leaves open, such as where a dish belongs or that a heavy box goes below a light one. *World knowledge* covers objects, everyday relations, and culture. *Embodied knowledge* covers materials, physical laws, tools and mechanisms, the preconditions and effects of actions, and the structure, sensing, and kinematics of robots. *Programming knowledge* covers languages, algorithms, and software interfaces, which empowers an S2 model to compose skills with code (Liang et al., 2023). Knowledge requires broad coverage, which only pretraining-scale data provides. World and programming knowledge are therefore inherited largely from pretraining on web text, images, and code. Embodied knowledge is sparser in these sources but needs the same breadth. It should therefore also be acquired during pretraining, by adding embodied data to the corpus at scale rather than leaving it to the smaller and narrower data of post-training. Suitable sources include egocentric and instructional videos that show how people handle objects and what their actions achieve, and robot trajectories paired with language, which tie actions to their preconditions and effects. For example, Magma includes human videos and robot data in its pretraining (Yang et al., 2025a). Others are technical documentation, such as robot specifications, kinematic models, and controller interfaces, and simulated data on physical properties and dynamics, such as mass, friction, stability, and deformation.

Understanding. Understanding is S2's grasp of the current situation: the state of the world, what the user wants, the state of the robot, and what cannot be inferred from the available observations. Understanding of the robot extends beyond its body to the capabilities beneath S2: which S1 skills exist, what their preconditions and failure modes are, and which S0 interfaces, reference frames, and limits apply. It differs from embodied knowledge: knowing how inverse kinematics works is knowledge, whereas judging whether *this* robot, in its current state, can reach a given pose is understanding. Understanding matters because every decision rests on it: an understanding error could propagate into planning and execution. Understanding of the robot lets S2 choose the level at which to act: whether to orchestrate a skill, control S0 directly, ask, or decline. Understanding of progress and completion, in turn, makes verification possible.

Reasoning. Reasoning is how S2 decides what to do next and revises that decision as evidence arrives. It encompasses inferring the intent behind an underspecified request, deciding when to ask for clarification, planning over long horizons, diagnosing failures, and handling interruptions. Reasoning is therefore central to S2's role. Two of these sub-abilities are particularly relevant to the execution paths in Figure 4: *skill selection and orchestration*, which decides which existing skills and tools to use and in what order (the layered-orchestration path), and *executable action inference*, which decides which action target to send directly when no suitable skill exists, without placing the model inside the real-time control loop (the direct-control path). Over a long task, diagnosis and recovery

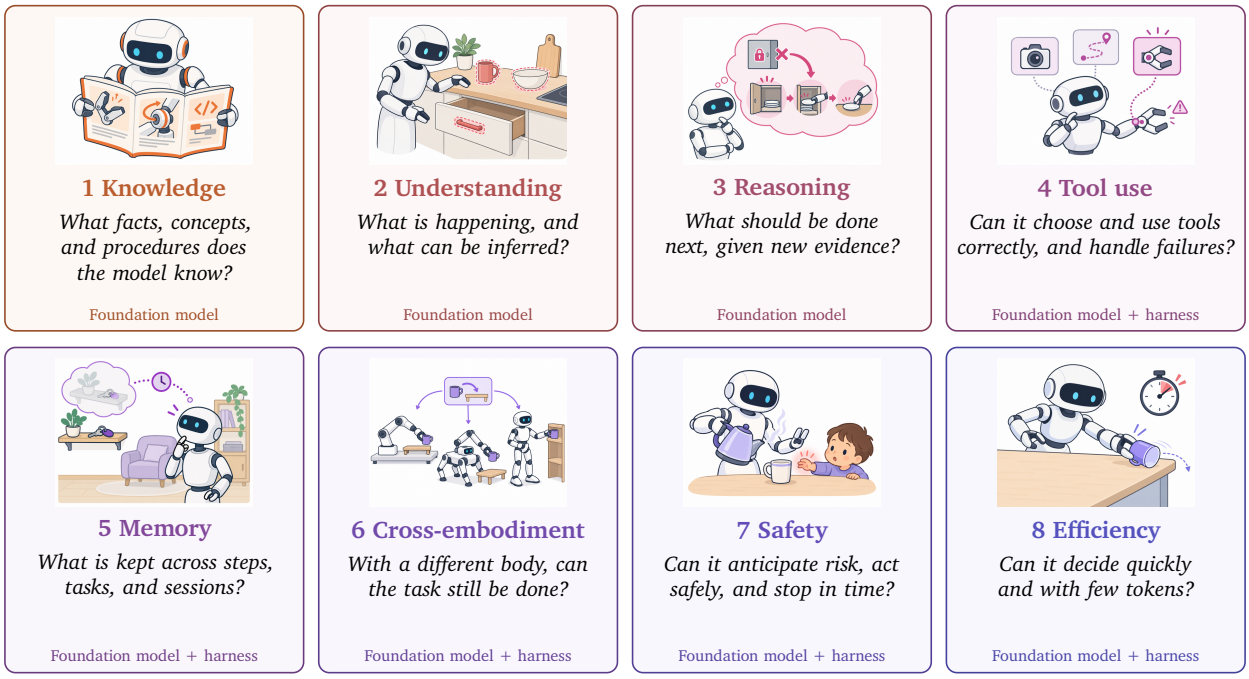


Figure 6 | **Core abilities of Embodied S2.** These abilities are necessary for Embodied S2 to perform successfully and reliably in the physical world; some reside in the foundation model, while others emerge from its integration with the embodied harness.

further limit the compounding of small failure rates (Section 2.2): S2 should change the conditions that caused a failure, such as the grasp point after a slip, rather than repeat the same call. Diagnosis can even drive learning without updating the model: by tracing each failure to its originating step, GPT-6 Astra learned manipulation skills on a real arm without demonstrations (Yu et al., 2026b).

Tool Use. Tool use is the ability to choose, parameterize, and combine the capabilities available to S2, namely perception, planning, action, and code, and to handle their failures. Each call involves three decisions: whether a call is needed at all; which arguments, in which reference frame and units, denote the intended physical target; and how far to rely on the result. Tool use matters because S2 acts on the world through calls. A sound decision executed with the wrong arguments still fails in execution, and sometimes a result can promise more than it delivers: a successful return from an inverse-kinematics solver does not guarantee that the robot can reach the solution without collision.

Memory. Memory is the ability to retain what the current observation no longer shows, to use it in a later decision, and to replace it when it becomes stale. A robot must remember objects and their attributes, locations, the order of completed steps, procedures demonstrated earlier, world state revealed by interaction, and the preferences of the people it serves. Memory matters because physical work extends beyond what a camera sees at any moment and acting on an outdated memory can cause failures (Liang et al., 2026a). The model's role is to decide what is worth recording and whether a memory still applies to the present situation; storing and retrieving information beyond the context window is the harness's role. The model learns these judgments from tasks in which earlier observations must be recalled and revised as the world changes.

Cross-Embodiment. Cross-embodiment is the ability to achieve a given objective when the body changes in appearance, morphology, end effector, action space, or dynamics. The same goal may require different paths or motions on different platforms. For example, a wheeled base may need a

longer route around obstacles that a legged robot can cross. It matters because the value of a general S2 lies in serving many bodies without being rebuilt for each. To serve a new body, the model must understand it, adapt the plans and action parameters to it, and reuse experience gained on other bodies without letting that experience interfere.

Safety. Safety is the ability to avoid harm along the entire action chain: recognizing a risk, predicting its consequences, choosing a cautious response, preventing harm before acting, monitoring during execution, recovering afterward, responding in time, and coordinating a stop. Specifically, *risk recognition* identifies hazards in the scene and in the request, such as hot liquid near a child's hand or a knife handed over blade first. *Consequence prediction* anticipates what an action would do in this particular scene. For example, pulling one box from a stack may topple the rest. A *prudent decision* chooses whether to proceed, proceed more cautiously, gather information, ask, or refuse, within the robot's validated capability limits, and prefers reversible or information-gathering actions under high uncertainty. *Prevention before action* checks a planned call against these predictions before it is dispatched. During execution, S2 monitors whether the situation still matches its assumptions and replans when conditions change–for instance, when a person enters the workspace. After a failure, *safe recovery* first returns the robot and the scene to a safe state and only then retries. A stop decided by S2 must be coordinated with S1 and S0 so that the running skill halts in a safe state. Safety is thus a property of decisions in the action loop rather than of answers about hazards. The model does not carry this responsibility alone. Hard limits on force and speed, protective stops, and permissions are enforced by the harness, S1, and S0 below the model, so that safety never rests on the model's judgment only.

Efficiency. Efficiency is the ability to make a good decision with low latency and with few tokens per decision. It matters because the physical world does not wait. A decision that arrives after the situation has changed is useless or even dangerous, and every token adds both latency and cost, which accumulate over hours of operation. Two ways to improve decision efficiency are faster decoding and distilling capable models into smaller ones that run on or near the robot.

A growing family of embodied foundation models, often called embodied reasoning models, develops the abilities described above by post-training a general vision-language backbone on embodied data. Examples include Gemini Robotics-ER (Gemini Robotics Team, 2025a,b), Cosmos-Reason1 (NVIDIA et al., 2025b), Embodied-R1 and R1.5 (Yuan et al., 2026c,d), RynnBrain (Dang et al., 2026; Li et al., 2026b), MiMo-Embodied (Hao et al., 2025), HY-Embodied (Tencent Robotics X et al., 2026), Vesta (Bjorck et al., 2026), ACE-Brain (ACE-Brain Team et al., 2026), and PhysBrain (DeepCybo Team et al., 2026). Whatever model is chosen, it decides from the context it is given and acts through the interfaces it is offered. Both are provided by the embodied harness, which we examine next.

3.3. Embodied Harness

Section 2.1 introduced the embodied harness through its five core parts; here we define its boundary. The harness comprises the components outside the foundation model that assemble the context for S2 and execute the decisions from S2. A recent survey similarly describes the harness as the "runtime control plane" that grounds observations, coordinates executable capabilities, monitors outcomes, and manages intervention when execution diverges from the plan (Huang et al., 2026c). Table 1 lists the five core components and three augmentations, together with the problem that each solves.

Context: World, User, and Body. A context manager assembles the model's current context, turning a heterogeneous, continuously changing physical situation into a bounded and structured input for the model. This context should cover at least three kinds of state. *World context* records objects, places, and their relations, together with when and how each fact was observed. *User context* records

Table 1 | **Components of an embodied harness.** The five core components appear in Figure 3; the three augmentations build on them.

Component	Problem addressed
Core components	
Context manager	Context from the world, the user, and the body arrives continuously and in large volume; it must be summarized and well managed.
Tool interface	Skills and tools must be exposed to the model at a suitable level of abstraction.
Planner	A long task cannot be planned once and executed blindly.
Memory	State must persist beyond the context window, and be updated.
Verifier	A completed call does not show that a step achieved its goal.
Augmentations	
Safety gate	Hazardous instructions and unsafe actions must be stopped before execution.
Budget control	Large models are slow and expensive.
Self-improvement	The system does not improve with use.

preferences, habits, and commitments; it changes at a different timescale and raises privacy questions that world context does not. *Body context* records the robot's geometry, kinematics, calibration, and current condition, together with descriptions of the S1 skills and S0 interfaces available to it. The model's understanding of the situation rests on these records. When the body changes, an explicit description of the new body, with its frames, kinematics, sensors, tools, and validated limits, lets S2 plan for a robot that it has not used before.

Memory beyond the Context Window. The entire interaction history of an agentic robot cannot simply reside in the model's context window due to limited capacity. For instance, a 128K-token context can hold only roughly three minutes of dense multimodal memory (Gopalakrishnan and Labenz, 2026). Within an episode, long context can nonetheless substitute for engineered memory, at a high cost in latency: a frozen GPT-6 Astra that kept an entire episode in context succeeded in 92.1% of the 800 test episodes of RoboMME (Dai et al., 2026) in a single run, averaging 18.3 s per model call with the simulator paused during each call (Jedlicki et al., 2026). This result shows that in-context memory can be viable within an episode, but its latency in this setting makes it unattractive as the sole memory mechanism for time-sensitive tasks. Across sessions, therefore, a memory module must store, retrieve, update, and invalidate information beyond the context window. Several harnesses make this module explicit: Thea maintains a persistent scene graph (Wang et al., 2026e), HarnessVLN combines hierarchical event memory with a persistent spatiotemporal graph (Chen et al., 2026c), and KnowBody maintains a revisable model of action-relevant body relations (Lou et al., 2026).

Execution Contracts. Since physical actions have side effects, the harness must know precisely what each call promises. Its tool interface exposes every tool with a schema, units, reference frames, and failure semantics, so that the model's choices of tool and arguments refer to well-defined physical quantities. Recent systems point to four contracts (Wang et al., 2026c; Deng et al., 2026; Huang et al., 2026c). An *observation snapshot* records when, where, and in which reference frame each observation was captured, and which calibration was in force. A *capability contract* specifies a skill's inputs and units, preconditions, expected effects, resources, time limits, behavior on cancellation, and the state it leaves for the next skill. An *invocation record* links each call to its task, subgoal,

attempt, inputs, and final status, so that retries and late results can be traced. An *evidence bundle* records the observations, controller states, and verifier outputs on which a judgment of success was based. Together, the contracts define a minimal lifecycle for every physical call: *propose* → *admit* → *dispatch* → *execute* → *confirm the terminal state* → *verify the effect* → *commit progress*. The lifecycle has explicit branches for rejection, failure, cancellation, and unknown outcomes. The branch for unknown outcomes is necessary since a timeout is not evidence that an action did not run. Before retrying, the harness should query the execution handle and observe the scene again; repeating a command that has in fact completed can execute a physical action twice. Cancellation passes through three distinct states as well: requested, accepted, and completed; the return of a cancel call does not show that the robot has stopped. Two systems illustrate this discipline: RegenHarness separates a model's proposal from a controller's termination and verified completion (Wang et al., 2026c), and DynaHarness records evidence under a shared physical execution contract (Deng et al., 2026). Skill boundaries require the same care, because a grasp can succeed in a way that obstructs the subsequent insertion. Handoffs can be made more reliable in two ways: by specifying explicit entry and exit conditions between skills (Xu et al., 2026c), and by restricting handoffs to configurations that the next policy is known to support (Huang et al., 2026b).

Safety outside the Model. Safety checks are required before execution. A dedicated safety gate is one possible implementation; the checks may instead be integrated into the tool interface or enforced by an independent execution layer. This check is the *admit* step of the lifecycle above. Hard limits on force and speed are enforced below the harness, by S1 and S0, and the real-time event monitor can trigger a protective stop in S0 according to predefined safety rules. A stop thus takes effect even if the model or the harness fails.

Three Loops. A harness supports three feedback loops. The *sensorimotor loop* determines how actions respond to physical state; it resides in S1 and S0. The *planning loop* determines how the next subgoal is chosen after each result; it is the main loop of S2. The *evolution loop* determines which capabilities the next run will use, such as skills, prompts, context rules, tools, or model weights revised from experience. A frozen model can still improve through the evolution loop, and a fixed program can still react through the sensorimotor loop. Conflating the three therefore obscures where a reaction or an improvement originates. Because the harness consists of code and data, S2 can also modify it: several recent self-improving robot agents revise tools, skills, context rules, or recovery policies rather than model weights (Tian et al., 2026; Deng et al., 2026; Lee et al., 2026a; Wang et al., 2026c; Lee et al., 2026b). Such changes require the same discipline as model training, namely held-out tasks, regression checks, and versioned rollback.

Harness Engineering. Many hard problems reside in the harness: multi-rate scheduling, cancellation that actually stops motion, consistent state across asynchronous components, and accountable records of what the robot did and why. Existing robotics infrastructure provides many building blocks for these problems, including action protocols with feedback and cancellation, middleware for real-time data flow (Macenski et al., 2022), and behavior trees with bounded recovery (Colledanchise and Ögren, 2018). Runtime governance layers add policy checks, capability admission, monitoring, rollback, and human override around the agent (Qin et al., 2026b; Liu et al., 2026d). The harness also carries the system's share of efficiency: it calls the model only when a decision is needed, compiles stable sub-processes into programs that run without the model, and delegates fast loops to S1 and S0. What remains is to connect these mechanisms to a model and to measure what each mechanism contributes.

Existing measurements already show that the harness substantially influences overall performance. For example, interfaces that let the model check and retry locally outperformed plain tool calling while using fewer model calls (Si et al., 2026). Verification can be decisive: in one ablation, with the model fixed, removing HumanCLAW's verifier nearly eliminated navigation success (Siyao et al.,

2026). The best level of abstraction for the interface also depends on the model: a skill-only interface raised GPT-5.6-Sol from 11% to 63% but lowered GPT-6 Astra from 92% to 61%; letting the model decide when to use skills recovered most of the loss (Han et al., 2026).

3.4. Constraints Imposed by the Physical World

Unlike digital environments, the physical world involves embodiment, continuous dynamics, and irreversible consequences (Section 5). Observations are partial and noisy, actions are delayed and constrained by actuator limits, and contact-rich interactions are governed by friction, compliance, and wear that are difficult to model exactly. Whatever model fills the S2 role, the physical world constrains how it may decide: decisions face deadlines, can be interrupted at any moment, must fit the robot's compute and cost budgets, and rest on aging, partial observations. Answer accuracy captures none of these constraints, which the following paragraphs discuss in turn.

Capped Scaling. For frontier models, capability has grown with the computation: more reasoning tokens generally yield better answers at the cost of longer latency (OpenAI, 2024). We call this regime *monotonic scaling*. A robot cannot follow it since the physical world imposes deadlines. A decision is useful primarily when it arrives before the situation that it addresses has changed, for instance before an object rolls off the table or a grasp slips. Embodied S2 therefore requires a different regime, which we call *capped scaling*: becoming more capable *within* a fixed decision deadline (Figure 7). What matters is not how long the model deliberates but how well the system decides within the time available. Humans meet physical deadlines through layered control (Todorov, 2004): grip-force corrections begin about 74 ms after an object starts to slip (Johansson and Westling, 1987), and the reflex responses that support recovery from a trip begin 60 to 140 ms after it (Eng et al., 1994). Deliberation handles only what cannot be handled faster. An S2 should likewise match its reasoning effort to the time available and delegate fast loops to S1 and S0. For the model, capped scaling means gaining capability within a fixed budget of tokens and time rather than by generating more tokens.

In terms of the S2 foundation model, capped scaling calls for *concise reasoning*. Much of the reasoning that people perform when acting physically is never verbalized, such as where to put a heavy box when packing or how to approach a fragile glass. We believe that such reasoning is

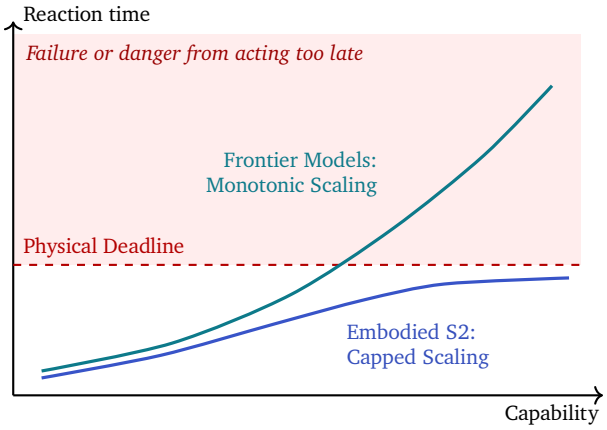


Figure 7 | **Capped scaling.** Frontier models scale monotonically: greater capability comes with longer latency. Embodied S2 must instead become more capable *within* a fixed decision deadline. A decision that arrives after the deadline is usually too late, even if it is correct: the situation has already changed, so acting on it could fail the task or create a hazard. Note that the curves are schematic: they illustrate the qualitative relationship, not measured data.

under-represented in the text and chain-of-thought data on which frontier models are trained, and that it does not require thousands of tokens. An S2 should therefore reason briefly and in grounded terms, using points, regions, relations, and short plans rather than long verbal deliberation. Concise reasoning helps meet deadlines; whether it also improves correctness on physical tasks remains an open question.

Figure 8 places current systems on a common scale to show how they compare with such deadlines. S0 loops run at about 50 Hz to 1 kHz, that is, every 1 to 20 ms (Luo et al., 2026c; Figure AI, 2026a). The control skills of S1 run at about 5 to 200 Hz, depending on the platform. The Dyna-2 policy updates its whole-body targets at about 5 Hz (Dyna Robotics, 2026b), $\pi_{0.5}$ produces actions at 50 Hz (Black et al., 2025a), and Helix's S1 runs at 200 Hz (Figure AI, 2025b, 2026a). For $\pi_{0.5}$, a median inference time of about 47 ms has been reported (as quoted in Yang et al., 2026d). The perception and motion planning modules of S1 have no fixed rate, because S2 invokes them on demand. Used directly as the S2 foundation model, GPT-6 Astra took roughly 40 to 8,000 times as long per decision as S1 control skills take per update (Figure 8). When issuing actions directly, it took 18.56 s per inference call (Yang et al., 2026d); when generating locomotion references, it took 39.86 s per call on average, and a 30-second walk required 250 calls with a paused simulator (Galbot Team et al., 2026). When the world does not pause, latency becomes a safety variable. RT-Safe (Xu et al., 2026d), a simulated urban navigation benchmark, keeps running while agents deliberate. On its hardest routes, Astra at its default reasoning effort took 7.7 s per decision, and models with longer response latency collided more often.

Interactivity. Language models typically complete a response before reading the next input. However, an S2 must be interactive: when the user says "stop" or the scene changes, it must abandon or revise its current line of reasoning. This requires a model that accepts new input during generation and resumes from partial state, and a harness that cancels in-flight tool calls. At least one commercial interface already supports such interruption: the streaming preview endpoint of Gemini Robotics-ER 2 cancels in-flight reasoning and tool calls when the user interjects (Google, 2026). Real-time chunking addresses a related problem in S1 by generating the next action chunk while the current one executes (Black et al., 2025c).

Deployability and Cost. Robots often operate with limited connectivity, and a single task may involve hundreds of decisions. Therefore S2 must run within the robot's compute and power budget or tolerate network latency and outages. The cost per task should be acceptable as well. Several open embodied foundation models intended for the S2 role have 2B to 30B parameters, and some include variants designed for edge deployment (Tencent Robotics X et al., 2026; Dang et al., 2026; Yuan et al., 2026d). Such smaller models are attractive in part because frontier models can be expensive in agentic settings: on one simulated benchmark, an agentic Astra policy incurred nearly 19 times the API fees of direct prompting (Yang et al., 2026d).

Fresh Observations and Partial Observability. While S2 deliberates, S1 and S0 keep acting and the world keeps changing, so the observation behind a decision has aged by the time the decision takes effect. S2 must therefore reason about the age of its information, request fresh observations when necessary, and express decisions in forms that remain valid under small changes. For instance, it should issue "place the cup on the shelf" rather than a fixed joint configuration. Partial observability compounds the problem: S2 must distinguish what it observes, what it infers, and what it does not know, and act on this uncertainty by looking again, asking, or abstaining.

These constraints do not imply that frontier models cannot serve as S2 foundation models; used with a harness, GPT-6 Astra already matches or exceeds trained policies on several simulated benchmarks (Section 4.3). Instead, the constraints imply that S2 must be judged not only by answer

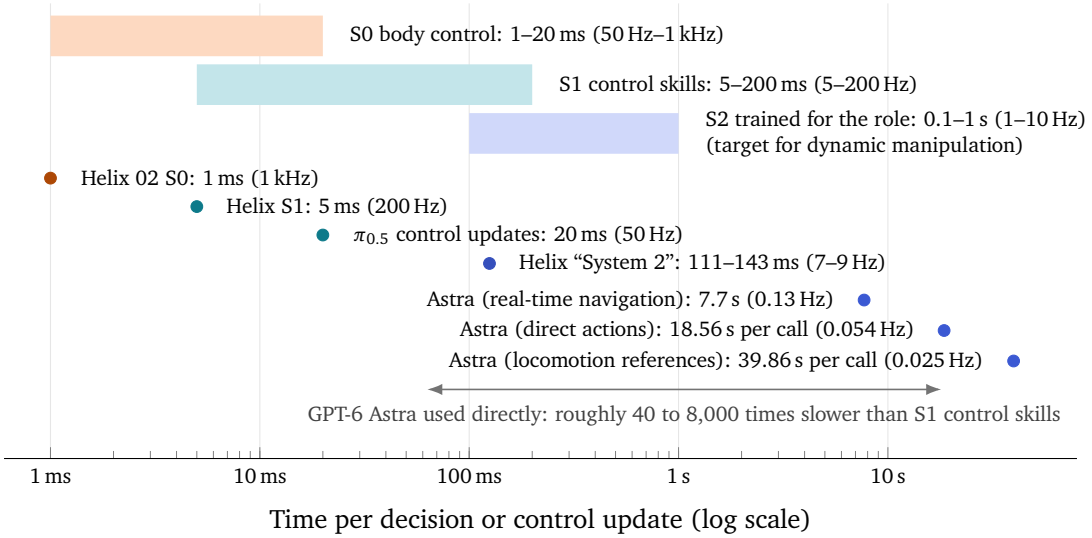


Figure 8 | **How long each layer takes per decision or control update.** All values are given as times, with the equivalent rate in parentheses (time = 1/rate). The S0 and S1 bands are indicative ranges. The S1 band covers control skills only, since the perception and planning modules of S1 are invoked on demand and have no fixed rate. S2 is event-driven rather than periodic, so we plot time per decision; the S2 band is a target for dynamic manipulation, not a requirement for every task. Helix's "System 2" is a vendor label: its output is a latent representation, so our criterion does not classify it as S2 (Section 3.1); we include it as the fastest upper-level module plotted. The Astra measurements come from different tasks and settings and are not directly comparable with one another.

accuracy but also by the quality of its decisions under the deadlines, interruptions, and budgets that the physical world imposes. This section has examined three aspects of Embodied S2: the foundation model's share of its abilities, the harness that supplies the rest, and the physical constraints on both. An evaluation must address all three: it must test the model on its own and the complete system in closed loop, and run the closed-loop tests on the world's clock rather than with the world paused. The next section develops such an evaluation framework.

4. Evaluation Framework

Two months after the release of GPT-6 Astra (OpenAI, 2026a), results in agentic robotics have become difficult to interpret. On RoboDojo alone, success rates reported for Astra across all 42 tasks range from 22% to 85%, depending on harness and protocol (Zhang et al., 2026c; SenseNova et al., 2026; Gao et al., 2026a). The difficulty has three sources. First, evaluations often do not state whether they measure a foundation model or a complete system. Second, they rarely disclose the full protocol, including observations, budgets, and demonstrations. Third, many pause the world while the model deliberates. This section asks how Embodied S2, a foundation model acting through an embodied harness, should be evaluated. It addresses what to evaluate, along the nine dimensions of Figure 9; how to evaluate; what current results show; and where current evaluation falls short.

4.1. What to Evaluate

We evaluate Embodied S2 at two levels. At the *model level*, the foundation model answers on its own, typically in a single turn without tools; these tests isolate knowledge, understanding, and the model-side components of reasoning and safety. At the *system level*, the model acts through a harness

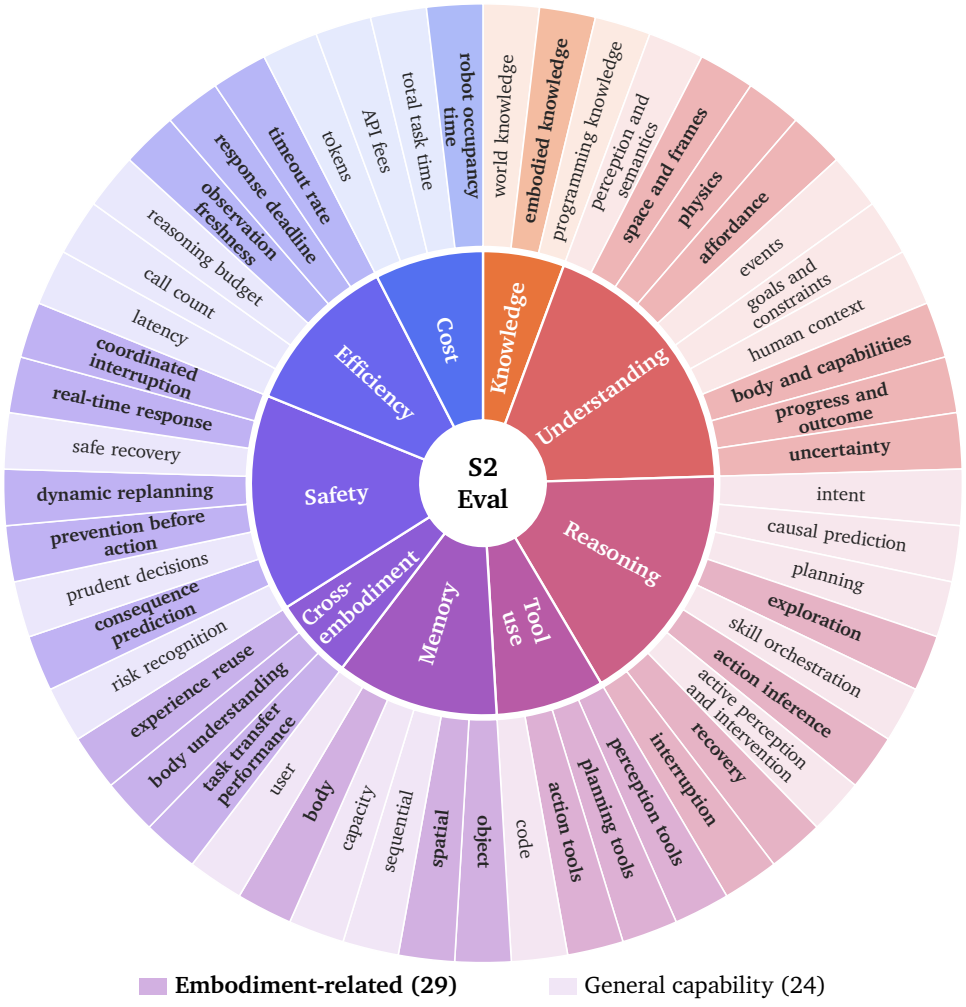


Figure 9 | **The S2 evaluation recipe at a glance: the 9 dimensions (inner ring) and their 53 sub-dimensions (outer ring).** Darker cells mark the 29 embodiment-related sub-dimensions: those specific to embodiment or strongly dependent on it; standard evaluations of general language models rarely cover these sub-dimensions.

and, where needed, invokes S1 skills and S0 interfaces; these tests measure closed-loop reasoning, tool use, memory, cross-embodiment transfer, safe execution, efficiency, and cost. The two levels answer different questions and provide complementary evidence about failure. A failure at the model level points to the model, whereas a failure that appears only at the system level points to the model, the harness, the interfaces, or their interaction. Comparing model-level and system-level results can help narrow down possible causes of failures. For example, replaying the observations and skill feedback that were available to the model at the point of failure can reveal whether the model would make the same decision again under controlled conditions. Within system-level evaluation, it also matters whether the harness is fixed across the models compared, supplied by each model's developer, or itself the object of study, because harness changes alone can move success a lot.

Across these two levels, Figure 9 organizes the evaluation into nine dimensions with 53 sub-dimensions: the eight abilities of Embodied S2 (Figure 6) and cost, which is not an ability but a measure of the resources the system consumes. The 29 *embodiment-related* sub-dimensions are specific to embodiment or depend strongly on it, and standard evaluations of general language models rarely cover them. The reasoning budget appears under efficiency rather than reasoning because it is a variable to be controlled and reported, not an ability in itself. Fulfillment of user intent is not a separate dimension but the end-to-end outcome that all dimensions serve. The dimensions play different roles and are not orthogonal. Knowledge, understanding, and reasoning describe what the foundation model can do. Tool use, memory, and cross-embodiment describe what the system achieves through its harness and across bodies. Safety constrains all of them. Efficiency describes how quickly and with how few tokens the system decides, and cost the resources it consumes. Results should therefore be reported as a profile across dimensions rather than summed into a single score.

Knowledge. Knowledge tests probe the model's facts, concepts, and procedures through self-contained, single-turn questions without tools. *World knowledge* spans objects and everyday relations, life science, matter, the Earth and the universe, and geography, history, and culture; CommonsenseQA (Talmor et al., 2019) is a candidate source. *Embodied knowledge* spans materials and physical properties, basic physical laws, object functions, tools and mechanical mechanisms, action preconditions and effects, robot structure and actuation, sensing and feedback, and coordinate systems and kinematic fundamentals. PIQA, NEWTON, and RoboCSKBench (Bisk et al., 2020; Wang et al., 2023; Töberg et al., 2025) cover parts of it. *Programming knowledge* spans language semantics, data structures and algorithms, software interfaces, and the writing of short functions from a specification. CodeMMLU (Nguyen et al., 2025) tests the first three, and the EvalPlus versions of HumanEval and MBPP (Liu et al., 2023c) test the last. In HumanEval and MBPP, the model completes a small, self-contained function from its signature and description, and unit tests score the result. On a robot, the same skill produces helpers such as a function that converts a detected bounding box into a grasp point. Results should identify the categories tested and the scoring protocol. General physical-commonsense scores do not establish coverage of robot structure, sensing, or kinematics, which require explicit probes. Offline scores also characterize knowledge only under the chosen protocol; closed-loop competence requires a separate system-level evaluation.

Understanding. Understanding tests ask whether an S2 foundation model can extract an accurate, task-relevant representation of the world, the user's goal, and the robot's own state from available observations. Unlike knowledge tests, they concern a particular physical situation. They cover multimodal perception, spatial and physical state estimation, affordance recognition, temporal understanding, task interpretation, recognition of human behavior, awareness of the robot's body and of the skills and interfaces available to it, progress estimation, and uncertainty assessment. Table 2 lists these abilities with an evaluation target and candidate benchmarks for each. Two of them deserve particular weight because they matter for closed-loop control and are weak in current models.

Understanding of the robot's own body and capabilities tells S2 whether to delegate to a skill or to act directly. Understanding of task progress and completion makes verification possible, yet success detectors are biased toward reporting success (Navasardyan et al., 2026). Static question answering also misses the temporal side of understanding: because observations age while S2 deliberates, a model must distinguish what it observed, and when, from what still holds now. Finally, tests must respect the boundary with reasoning: understanding determines *what is happening and what can be inferred*, whereas reasoning determines *what should be done next*. For example, recognizing that an object has slipped from the gripper tests understanding; identifying the cause and selecting a corrective action tests reasoning.

Evaluating understanding therefore requires separating state inference from downstream decision-making and execution. At the model level, tests should specify the observations, task instructions, robot-state information, and interaction history available to the model. Such specification matters most for progress estimation and uncertainty assessment, where the correct answer may depend on earlier observations or on information that remains unobservable. At the system level, task outcomes provide complementary evidence, but they cannot by themselves attribute failures to understanding, because planning, tool use, and execution may also contribute.

Together, these tests characterize the information quality on which S2 decisions depend. Whether the agent can use that information to select effective actions is the subject of reasoning evaluation.

Table 2 | **Understanding abilities, evaluation targets, and candidate benchmarks.** Abilities in bold are specific to embodiment or strongly depend on it, as in Figure 9.

Ability	Evaluation target	Candidate benchmark
Multimodal perception and semantics	Recognize objects, attributes, and task-relevant entities from multimodal observations.	RoboBench (Luo et al., 2026b)
Spatial relations and reference frames	Infer spatial relationships, relative positions, and geometry across viewpoints and frames.	MMSI-Bench (Yang et al., 2026f), VSI-Bench (Yang et al., 2025b)
Physical states and mechanisms	Identify physical properties, object states, support and contact relationships, and mechanisms relevant to interaction.	PhysBench (Chow et al., 2025)
Affordances and interaction feasibility	Identify physically feasible interactions and action regions given the observed scene and available capabilities.	RoboBench (Luo et al., 2026b)
Events and temporal structure	Identify events, their order, and temporal dependencies from observation sequences.	EgoTaskQA (Jia et al., 2022)
Tasks, goals, and stated constraints	Interpret expressed goals, desired states, and task constraints.	Embodied Agent Interface (Li et al., 2024b)
Human behavior and interaction context	Recognize human activities, intentions expressed through observable behavior, and the current stage of an interaction.	ProAct-75 (Zhu et al., 2026)
Embodiment, skills, and control interfaces	Infer the robot's body state, kinematics, and reachable workspace, the available S1 skills and S0 control interfaces, and their capabilities and constraints.	EQA-phys (Zhou et al., 2025b), HumanCLAW (Siyao et al., 2026)
Task progress and outcome state	Estimate completed steps, remaining progress, and whether an intended physical effect has occurred.	ContextProgress-Bench (Zhang et al., 2026a), FailBench (Navasardyan et al., 2026)
Observability, uncertainty, and knowledge limits	Distinguish observed facts from inferred or unavailable information and recognize when evidence is insufficient.	AbstainEQA (Wu et al., 2026)

Reasoning. Reasoning tests ask whether the S2 foundation model can select and revise decisions that advance the user's goal, given its understanding of the current state, the robot's capabilities, and the constraints of the physical environment. They cover intent inference, causal and counterfactual prediction, hierarchical planning, spatial exploration, skill selection and orchestration, executable action inference, active perception and intervention decisions, failure recovery, and handling of interruptions and rescheduling. Executable action inference, which corresponds to the direct-control path of Figure 4, needs a precise evaluation target. We treat it as a reasoning ability rather than as a separate policy dimension. The target is therefore whether the model produces physically executable, goal-consistent action decisions, not whether it can reliably control a robot as a complete policy. Reliable control also depends on continuous closed-loop operation, latency, controllers, and the harness, so it belongs to system-level evaluation.

Distinguishing reasoning competence from execution performance requires tests at both levels. At the model level, *controlled decision-making tests* assess whether the model selects feasible, goal-consistent actions under specified observations, capabilities, and constraints, without physical execution. At the system level, *closed-loop tests* additionally assess whether its decisions remain effective when actions are executed and their outcomes are observed. The distinction matters because task success depends not only on decision quality but also on the harness, tools, controllers, and physical environment. Table 3 summarizes these abilities and lists candidate benchmarks for each.

Table 3 | **Reasoning abilities, evaluation targets, and candidate benchmarks.** Abilities in bold are specific to embodiment or depend strongly on it, as in Figure 9.

Ability	Evaluation target	Candidate benchmark
Intent inference and disambiguation	Infer underspecified goals and decide whether clarification is needed before acting.	Ask-to-Act (Ramrakhya et al., 2025)
Causal, predictive, and counterfactual reasoning	Predict the consequences of actions, explain observed outcomes, and compare alternatives under hypothetical changes.	EgoTaskQA (Jia et al., 2022)
Hierarchical, long-horizon, and constrained planning	Decompose goals into feasible subgoals, order actions, and satisfy dependencies and constraints.	Embodied Agent Interface (Li et al., 2024b), LADDER (Petkov et al., 2026)
Spatial search and exploration planning	Select search locations, exploration strategies, and navigation subgoals to reach a target or reduce uncertainty.	GOAT-Bench (Khanna et al., 2024)
Skill selection and orchestration	Select, sequence, and coordinate the available skills and tools according to task dependencies, capability constraints, and expected effects.	RobotWorld (Yang et al., 2026j)
Executable action inference	Generate executable action targets directly, such as end-effector poses, waypoints, trajectories, joint targets, or skill parameters, that satisfy embodiment, geometric, and environmental constraints.	RoboDojo (Chen et al., 2026b; Zhang et al., 2026c), Robocurve (Menon et al., 2026), EmbodiedBench (Yang et al., 2025c)
Active perception and intervention decisions	Decide when additional observations are needed and whether, when, and how to intervene during an ongoing task.	RoboQuest (Liu et al., 2026c), ProAct-75 (Zhu et al., 2026)
Failure diagnosis, correction, and recovery	Identify the cause of a failed action, choose a corrective strategy, and restore progress without unnecessarily repeating the failure.	ViFailback-Bench (Zeng et al., 2026a), FailBench (Navasardyan et al., 2026), RESCUE (Ahmedt-Aristizabal et al., 2026)
Interruption, rescheduling, and collaboration	Revise plans after user interventions, environmental changes, or collaborator actions while preserving valid task progress.	Robix evaluation sets (Fang et al., 2025), PARTNR (Chang et al., 2025)

Evaluation gaps. Existing benchmarks provide useful measurements of reasoning-related capabilities, but several do not isolate the underlying decision. For instance, successful navigation does not necessarily establish effective exploration planning, and successful tool-based execution does not isolate skill selection and orchestration. Likewise, task success alone cannot show that one model generates better actions than another, because inverse kinematics, controllers, the action interface, and the observation contract also affect success. A more direct protocol would hold the state representation, available capabilities, and execution constraints fixed while scoring candidate decisions before execution. For failure recovery, existing benchmarks score detection, diagnosis, and correction guidance offline (Zeng et al., 2026a); controlled failure injection could measure the same steps during execution, together with recovery success and the additional time or actions required. Such tests would complement closed-loop benchmarks rather than replace them.

Tool Use. Tool use should be evaluated with the experimental setup held fixed: the same robot, sensors, controllers, and scene configuration, prepared in advance for every model compared. The unit of an ablation is the actual callable interface rather than a tool family, because a single grasping API may bundle perception, planning, and execution. Three kinds of experiment answer different questions and should not be conflated. Freezing the model and varying the tools measures the gain from provisioning. Freezing the tools and varying the model measures how well models use them. Freezing both and varying prompts, memory, or feedback measures the harness's effect.

The tools themselves fall into four types, each with its own evaluation target. *Perception* tools identify and segment targets, estimate spatial geometry, propose candidate grasps, and report contact; examples are SAM 3 (Carion et al., 2025), RGB-D projection and multi-view triangulation, Contact-GraspNet (Sundermeyer et al., 2021), and contact queries. Evaluation checks whether the model requests the estimate that the decision needs and uses the returned geometry in the correct reference frame. *Motion-planning* tools, such as planning backed by cuRobo (Sundaralingam et al., 2023) and waypoint execution, turn target poses or waypoints into trajectories within kinematic limits and modeled collision constraints. Evaluation checks whether the model supplies feasible targets and interprets planning failures. *Action* tools include learned-policy calls and direct-control commands. Learned-policy calls delegate a subtask to a VLA or WAM, such as $\pi_{0.5}$ (Black et al., 2025a), GR00T (NVIDIA et al., 2025c), or Cosmos3-Nano-Policy (NVIDIA, 2026). Evaluation checks whether the model delegates suitable subtasks and monitors, accepts, corrects, or overrides the resulting actions

Table 4 | **Tool-use benchmarks, evaluation targets, and corresponding scales.**

Benchmark	Evaluation target	Scale
CaP-Bench (Fu et al., 2026)	Writing and running programs that compose perception and control primitives, with various abstraction, perception, and execution-feedback settings	7 core manipulation tasks; 187 tasks in the full CaP-Gym release, from Robosuite, LIBERO-PRO (Zhou et al., 2025c) and BEHAVIOR (Li et al., 2024a)
EmbodiedBench (Yang et al., 2025c)	From high-level household planning to low-level navigation and manipulation	1,128 tasks in four environments
EmbodiedSkillBench (Han et al., 2026)	Whether packaged skills help general agents: native control, skill-only execution, and agent-directed reuse	250 skills; 100 perturbed LIBERO-PRO tasks
RobotWorld (Yang et al., 2026j)	Robot use across embodiments, with interaction budgets and executable success checks	84 tasks from 20 projects in five domains

where the interface allows. Direct-control commands for the end effector, the joints, and the gripper leave the target, the sequence, and the timing of observation to the model. Evaluation checks whether these commands respect the body's units, frames, and limits. *Code* tools allow the model to compose perception, planning, and action calls into executable programs, as in Code as Policies (Liang et al., 2023). Evaluation checks whether the generated code uses the available APIs correctly, handles feedback and failures, and terminates appropriately. The underlying planners and controllers may be shared across types, so a test must name the interface that the agent actually called.

The ability to use tools has three aspects, each scored separately. *Effectiveness and reliability* concern whether the model uses tools correctly, verifies their actual results, and adjusts its actions after failures. *Speed and efficiency* concern how long a task takes, how many calls it needs, and how quickly the model corrects a problem once it is detected. They matter because each call to a frontier model can take seconds (Figure 8). They should be reported even when success does not change, together with how tools are organized and scheduled, for example whether the agent keeps reasoning while a tool runs. *Learning and transfer* concern whether the model distills what it has learned about tools into skills that improve later tasks, as when frontier models author skills that later run without them (Bai et al., 2026). Measuring these aspects apart from provisioning gains requires paired tests that change one condition at a time, with the correct behavior and an independent scoring rule fixed in advance. Table 4 lists benchmarks that test tool use and the tool types that each exposes.

Memory. Memory evaluation asks whether a robot uses information that its current observation no longer contains. Inspired by MIKASA-Robo (Cherepanov et al., 2026a), we divide memory-intensive tasks into six classes. *Object memory* keeps information about objects that become temporarily unobservable, and *spatial memory* keeps object locations and a map of the environment. *Sequential memory* keeps temporally ordered information, and *memory capacity* concerns keeping several pieces of information at once. *Body memory* keeps information about the robot's own body, such as its kinematic configuration and action-relevant body relations; *user memory* keeps persistent information about the user, such as preferences and prior instructions. Table 5 lists five benchmarks that together cover tabletop manipulation, household interaction, and question answering. They score different outcomes: manipulation success (MIKASA-Robo (Cherepanov et al., 2026a), RoboMME (Dai et al., 2026), and RoboMemArena (Lei et al., 2026)), success on interactive tasks that require building and updating memory (EMem-Bench (Liang et al., 2026a)), or answers to questions about what the robot

Table 5 | **Memory benchmarks, evaluation targets, and corresponding scales.**

Benchmark	Evaluation target	Scale
MIKASA-Robo (Cherepanov et al., 2026b)	Object, spatial, sequential, and capacity memory	90 tasks in three horizon splits of 25 to 2,160 steps
RoboMME (Dai et al., 2026)	Temporal (counting), spatial (permanence), object (reference), and procedural (imitation) memory	16 tasks × 50 test episodes; 1,600 training demonstrations
RoboMemArena (Lei et al., 2026)	Multi-object transfer, occlusion, counting, and sequence, in long multi-stage manipulation	26 simulated and 5 real-world tasks
EMem-Bench (Liang et al., 2026a)	Passive observation, dynamic tracking, interaction failure, and experience generalization	2,554 episodes in 1,118 scenes
OpenEQA (Majumdar et al., 2024)	Questions about a recorded episode and questions answered by active exploration	More than 1,600 questions from 180 real-world environments

has seen (OpenEQA (Majumdar et al., 2024)). None of these benchmarks tests body or user memory.

Cross-Embodiment. Cross-embodiment evaluation asks whether a fixed objective remains achievable when the body changes. It has three aspects. *Task transfer performance* compares success rate, execution efficiency, and stability across the change of body, with the same objectives and matched scenes. *Body understanding* concerns whether S2 understands the new body's perception, motion, and manipulation capabilities and adjusts its action parameters and execution plans to them. For example, when a task moves from a small arm to a larger one such as a Franka Panda, aligning the end-effector target pose is not enough. Longer links change the trajectory of the whole arm, so S2 must check again for collisions with the environment and with the robot itself. *Experience reuse* concerns whether S2 carries over what still applies from its existing skills, revises the parts specific to the old body, and keeps outdated experience from interfering with execution on the new one. A transfer protocol should fix the S2 foundation model, pair the scenes across bodies, and report each embodiment separately. It should also disclose all exposure to the target body (demonstrations, adaptation data, tools, and controller changes) and compare only tasks that are feasible for both bodies. Each robot should be described by a versioned *capability card* that lists its frames and units, kinematic and geometric parameters, sensors, tools with their preconditions and postconditions, controller version, and validated capability limits. Table 6 summarizes five benchmarks evaluating multiple embodiments on tasks including navigation and manipulation.

Safety. Safety evaluation should cover the entire action chain, from risk recognition to coordinated interruption. The eight axes in Table 7 turn that chain into tests, each with candidate benchmarks. Existing benchmarks test three abilities directly (risk recognition, prudent decisions within capability limits, and prevention before action) and whether the postconditions of post-action handling hold. For example, IS-Bench (Lu et al., 2025) checks whether mitigation is performed before and after risk-prone steps. The other four are covered only partially, and executable protocols are still missing for delayed consequences, dynamic disturbances, recovery, cancellation, and takeover. Refusal, safe task completion, serious violations, and excessive refusal must be reported separately, because a refusal rate alone does not measure safe completion. Refusal must also be measured in the action loop, because refusal learned in text transfers only partially to action (Sun et al., 2026a). LLM-based safety judges alone do not suffice: in one study, an LLM judge missed nearly 83% of unsafe trajectories, labeling them safe or misidentifying the unsafe step (Zhan et al., 2025).

Table 6 | **Cross-embodiment benchmarks and what each can and cannot cover.**

Benchmark	Focus	Coverage
RobotWorld (Yang et al., 2026j)	Multi-body execution	Execution on many bodies; cannot isolate morphology, because objectives are not paired across bodies
Habitat-MAS (Chen et al., 2025b)	Heterogeneous capabilities and planning	Planning with robot capability descriptions; tests cooperation, not transfer
OmniNavBench (Sun et al., 2026c)	Navigation across bodies	The same instructions on wheeled, quadruped, and humanoid agents; navigation only
ZETA (Yan et al., 2026)	Tabletop manipulation transfer	Transfer to held-out target bodies; target-data exposure must be disclosed
DexJoCo-X (Jiang et al., 2026)	Dexterous-hand manipulation	Joint training across seven hands; transfer to an unseen hand is left to future work

Table 7 | **Safety abilities, from risk recognition to coordinated interruption, with candidate benchmarks and the coverage each provides.** Abilities in bold are specific to embodiment or depend strongly on it, as in Figure 9.

Ability	Candidate benchmark	Coverage
Risk recognition: identify risks involving objects, environments, and people	IS-Bench (Lu et al., 2025); HomeSafeBench (Yao et al., 2025)	Risks that emerge during household tasks; hazards reported during free exploration of a home
Consequence prediction: anticipate immediate, delayed, and counterfactual consequences	SafeMindBench (Chen et al., 2025c); RoboHarm (Sun et al., 2026a); ASIMOV-2.0-Injury (Jindal et al., 2025)	Factual, causal, and temporal constraints in planning; refusal on real arms; risk type and severity in text scenarios from injury reports, the safety effect of a candidate action, and the severity after it
Prudent decisions within capability limits: choose execution, verification, assistance, or refusal	ASIMOV-Agentic (Gemini Robotics Team, 2026a); SafeAgentBench (Yin et al., 2024)	Decisions under payload and reach limits and under uncertainty; refusal of hazardous tasks in planning
Prevention before action: check preconditions and use a safe action order	IS-Bench, mitigation before risky steps (Lu et al., 2025); DESPITE (Zhang et al., 2026b)	Mitigation performed in the correct order; dangerous plans among valid ones in symbolic planning
Monitoring and replanning: observe execution feedback and revise the plan after a deviation	SENTINEL (Zhan et al., 2025); RT-Safe (Xu et al., 2026d)	Formal checks of plans and trajectories; collisions in a world that keeps running. Timed faults and dynamic disturbances are missing
Post-action handling and recovery: remove residual hazards, recover, or terminate safely	IS-Bench, mitigation after risky steps (Lu et al., 2025)	Postconditions; protocols for incidents and recovery from failure are missing
Timely safety response: intervene effectively within a specified safety window	RT-Safe (Xu et al., 2026d); ASIMOV-2.0-Video (Jindal et al., 2025)	Safety events when the world keeps moving during inference, in simulation; predicted intervention times in videos. End-to-end timing on robots is missing
S2/S1 coordination and interruption: propagate constraints, cancel commands, stop, or transfer control	ASIMOV-Agentic (Gemini Robotics Team, 2026a)	Stop-or-continue decisions and the feasibility of VLA calls

Efficiency and Cost. Efficiency evaluation measures the latency from observation to action, the number of model calls, the freshness of the observations on which each decision rests (the observation age defined in Equation 1), and the rate of timeouts and missed deadlines, under a stated reasoning budget and response deadline. Output tokens per second measure decoding speed, not robot responsiveness. Latency should therefore include perception, context preparation, inference, network delay, tool execution, and controller response whenever they are on the critical path. Cost should be reported as tokens, API charges, total task time, and robot occupancy, per attempt and per success, including failures and retries. Reported costs differ by orders of magnitude, and cost rankings need not follow success rankings. In Astranomy, an agentic Astra session used 30.9M tokens ($36.40) per successful episode, against 160k tokens ($1.96) for direct prompting, although most of its input was served from cache (Yang et al., 2026d). On RoboQuest, an average episode with Astra took 10.4 minutes of simulated time but 53.8 minutes of wall-clock time and cost $11.29 (Liu et al., 2026c). On a single real six-axis arm, frontier coding agents took about 30 minutes per task that a person completes in 20 seconds (Artificial Analysis, 2026).

User Intent. For user intent, explicit instructions are the easy case; harder cases include vague requests, unstated preferences, learned habits, and changes of mind during a task. Evaluation should therefore score the final state against the user's intent rather than the literal instruction. It should also report the number of clarifying questions, adherence to stated and learned preferences, the handling of corrections and interruptions, and the number of human interventions. Existing benchmarks cover only parts of intent fulfillment. PRISM isolates implicit intent resolution, which was "a significant bottleneck for all model families" (Lim et al., 2026). TEACh requires acting on intent conveyed over several dialogue turns (Padmakumar et al., 2022), and Ask-to-Act scores clarification of ambiguous requests (Ramrakhya et al., 2025). MEMENTO finds that agents "can recall simple object semantics but struggle to apply sequential user patterns to planning" (Kwon et al., 2026). Two devices can be borrowed from digital agents: a hidden supervisor with a simulated user who gives multi-turn feedback "without leaking grading criteria" (Chen et al., 2026e), and the pass^k reliability metric (Yao et al., 2024). To our knowledge, no benchmark yet scores end-to-end intent fulfillment on real robots with real users across sessions.

4.2. How to Evaluate

The dimensions above specify what to measure. Comparable results also require a common way of measuring, which the same benchmark name does not guarantee. We therefore propose the following principles, several of which are adapted from the evaluation of digital agents (Yao et al., 2024; Kwa et al., 2025).

i. **Separate the model from the system.** Fix the harness when comparing foundation models, and fix the model when comparing harnesses. In both cases, match the embodiment, interfaces, observations, task instances, and budgets. Otherwise, label the comparison as one between complete systems, as CodeActionBench does (Lyu et al., 2026). What counts as "no tools" must be defined by the interface contract rather than by name: in EmbodiedSkillBench, for example, the "no skill" condition still acts through registered native end-effector and gripper tools (Han et al., 2026).

ii. **Treat the action interface as a controlled variable.** The same foundation model may call skills, emit end-effector targets, or output joint targets (Figure 4). These interfaces do not form a scale from easy to hard: each brings different execution priors, control burdens, failure modes, and degrees of freedom. Model comparisons should therefore fix the action abstraction level and report it. Its effect can be measured directly by holding the model, tasks, robot, sensors, environment, and budgets fixed while varying the interface. The comparison should include a

skill-only interface, a direct end-effector interface, and a hybrid in which the model decides when to call a skill and when to act directly. Such comparisons should report task success separately from action validity, collisions and constraint violations, the number of calls, physical execution time, inference latency, recovery, and total cost.

iii. **Run on a world clock.** Many simulated evaluations pause the world while the model deliberates (Jedlicki et al., 2026; Galbot Team et al., 2026; Yang et al., 2026j). Paused and live execution measure different abilities and must be reported separately. On matched hard routes of RT-Safe, live execution raised collisions by a factor of 12.3 but left task completion essentially unchanged. Safe success fell from 19.8% to 0.7%, and mean latency was positively associated with collisions across models (Xu et al., 2026d). Embodied First Exam separates a Standard setting, in which the agent receives sensor observations and the simulation pauses while it deliberates, from a Realtime setting, in which the environment continues to advance during decision-making. On 54 dynamic core tasks evaluated in both settings, GPT-6 Astra solved 23 Standard tests and 17 Realtime tests (Wang et al., 2026d). Live evaluations should report the rate of missed deadlines and the distribution of the *observation age*. The age is defined as

$$\Delta t_{\text{obs}} = t_{\text{effect}} - t_{\text{capture}}, \tag{1}$$

where t_{capture} is when the observation behind a decision was captured and t_{effect} when the resulting action takes effect. Because it includes perception, communication, queuing, and execution delays, observation age reflects the physical response more closely than model inference latency does. When sensors, computation, and actuators run on different devices, both times must be measured on synchronized clocks.

iv. **Disclose the observation contract.** State what the model and its tools receive: cameras and their history, the source of depth (a sensor, a prediction, or triangulation), privileged object state, success or progress signals, and any demonstrations or scene assets. Give the agent only information that a real robot could also obtain (OpenRUA Team, 2026). Without this disclosure, numbers from different studies cannot be compared.

v. **Verify completion independently.** Decide success with a checker that has access to ground truth that the system cannot observe, and never accept the model's own claim or a successful tool return as proof. Report the rate of false completion claims, because agents often declare success prematurely (Jedlicki et al., 2026; Liu et al., 2026c) and VLM judges are biased toward success (Navasardyan et al., 2026). The checker must also be inaccessible to the agent. In one study, agents read static object files that a physical robot could not read and located a hidden success check within a single session (Wang et al., 2026g). Hiding the verifier changes results substantially. When tool returns no longer revealed success, agents had to decide for themselves when an episode ended. Under this condition, the mean scores of two systems over 50 episodes on 10 tasks of RoboDojo fell from 83.2 to 63.5 and from 42.0 to 26.1 (OpenRUA Team, 2026).

vi. **Quantify uncertainty and reliability.** Report per-task results, the number of repetitions, and confidence intervals. A difference of a few percentage points on a handful of real-robot trials rarely supports a conclusion (TRI LBM Team et al., 2025; Kress-Gazit et al., 2024). Reliability across repeated runs, such as the probability that all k trials succeed, should complement the chance of at least one success (Yao et al., 2024). Rankings that aggregate benchmarks need the same care. The Robot Intelligence Index, for example, reports for each model the range within which its rank falls in 90% of resampled re-runs (Playground Research, 2026).

vii. **Control and sweep budgets.** Fix and report the reasoning effort, the budgets for calls, steps, and wall-clock time, and any retry rules, and sweep them into success–cost and success–latency curves where possible. In Astranomy, raising the effort from medium to extra-high improved success only from 72.7% to 74.5% at 2.6 times the latency and 2.2 times the cost (Yang et al., 2026d). Budgets must also match the benchmark's own limits: recounted under RoboDojo's

native step limit, one system's success rate fell from 47.1% to 1.0%, because 97 of its 99 successes came after that limit (OpenRUA Team, 2026). Improvement methods should be compared with retry baselines under the same budget, with development and test tasks kept apart (Wang et al., 2026j,g). For example, a skill developed on one task with the benchmark's verifier raised that task's score from 0 to 80 yet left two of four other tasks at baseline (OpenRUA Team, 2026). Results that use context about the test tasks should not be compared with zero-shot results (OpenRUA Team, 2026).

viii. **Validate on real robots.** Treat simulation as a screen rather than a verdict, and audit the shortcuts of shared simulated suites. In SIMPLER, simulated success tracked real success closely once control and visual gaps to the real setup were mitigated, and less closely when they were only partly mitigated (Li et al., 2025a). Policies trained in simulation can also fail badly on real robots (Ibarz et al., 2021). Real-robot evaluations should pair seeds and scenes and report the human assistance and environment resets they required.

ix. **Fix the conditions of a level claim in advance.** Some results are offered as evidence that a robot carries a level of delegated responsibility (Section 6.1). Such a result should fix beforehand the initial conditions, the permitted assistance and interruptions and how they are counted, and the evidence of completion. Completing a long script does not demonstrate delegation if a person must keep repairing the plan. The duration of the work should be measured in the time that a skilled person would need for it, as METR (Kwa et al., 2025) does for software agents. The robot's own wall-clock time and the human assistance it received should be reported separately, because a robot may need far longer than a person for the same work.

The checklist below collects the information that these principles require every report to disclose.

Reporting Checklist: The Harness Card. Every reported S2 result should be accompanied by a *harness card*. The card records (i) the foundation model, its version, and its reasoning budget, and (ii) the harness version, prompts, skill library, memory, and verifiers. It records (iii) the embodiment and the S1 and S0 interfaces, with the action abstraction and control rates, and (iv) the observation contract and any privileged information. It also records (v) the time semantics, that is, whether the world pauses during inference, and (vi) the number of runs and seeds, with confidence intervals. Finally, it records (vii) the rate and definition of human or rule-based interventions and the paths for cancellation and human takeover.

4.3. What Current Results Show

Table 8 collects published results for general models serving as the S2 foundation model. Most come from preprints posted within two months of Astra's release; we report them as published and have not reproduced them. The rows are grouped by what they show: competitiveness on narrower simulated suites, low success on broad or exploratory benchmarks, and sensitivity to the interface, live execution, and physical action. Read in light of the principles above, these and other results support seven conclusions, which we qualify by the strength of the evidence.

General models with harnesses are competitive on narrower simulated suites. On several simulated benchmarks, a general model with a harness matches or exceeds trained policies, particularly on tasks that reward semantic understanding, memory, or flexible composition (Gao et al., 2026b; Yang et al., 2026d; Si et al., 2026; Jedlicki et al., 2026; Petkov et al., 2026). On LADDER, for example, orchestration recovers part of the success that learned policies lose as tasks lengthen (Petkov et al., 2026). In several comparisons, however, the model received information that the policies did not, such as a written task recipe (Zhang et al., 2026c; Li et al., 2026e). The result is also model-specific:

Table 8 | **A snapshot of published results for general models serving as the S2 foundation model (October 2026).** Rows are grouped by the conclusions of Section 4.3 that they support. S1 lists the skills and task tools that the model invoked; *none* means that it sent end-effector, joint, or motion targets directly to low-level control. Settings differ, so rows are not comparable; several of the simulations pause while the model deliberates.

Benchmark	S2 foundation model	S1 components	Comparison
Competitive with trained policies on narrower simulated suites			
RoboDojo-Sim (Chen et al., 2026b; Zhang et al., 2026c)	GPT-6 Astra	None	22.48%; above all 40 public policies at the time
RoboDojo-Sim, official evaluation (Gao et al., 2026b)	GPT-6 Astra	Agent-written tools, mostly task-specific	46.9%; first on the leaderboard when posted
RoboLab-120 (Yang et al., 2026c,d)	GPT-6 Astra, direct / agentic session	None	66.3% / 74.8%; best leaderboard policy: 51.0%
LIBERO-PRO (Zhou et al., 2025c; Si et al., 2026)	GPT-6 Astra	Motion primitives, segmentation, and $\pi_{0.5}$, composed in code	94.0%; $\pi_{0.5}$: 53.0% (different evaluation setup)
RoboMME (Dai et al., 2026; Jedlicki et al., 2026)	GPT-6 Astra	None	92.1%; best leaderboard entry: 60.8%
LADDER (Petkov et al., 2026)	GPT-6 Astra	None	62.0%; learned-policy systems: at most 29.6%
Far from solved on broad or exploratory benchmarks			
RobotWorld (Yang et al., 2026j)	GPT-6 Astra	None	19.0%; 63 tasks solved by none of five models
RoboQuest (Liu et al., 2026c)	GPT-6 Astra	None	23.2%; fine-tuned $\pi_{0.5}$ almost never succeeds
HumanCLAW (Siyao et al., 2026)	Best of nine VLMs	Whole-body skills such as walking, turning, and climbing stairs	16.8%; failures arise from losing track of the body
Sensitive to the interface, live execution, and physical action			
EmbodiedSkillBench (Han et al., 2026)	GPT-6 Astra	None / selected skills only / skills on demand	92% / 61% / 85%; the interface can cap the model
RT-Safe, hard routes (Xu et al., 2026d)	Eight VLMs	None	Safe success 19.8% (paused) vs. 0.7% (live); collisions rise 12.3-fold in the live world
RoboHarm, real robot (Sun et al., 2026a)	Six frontier models	None	Refusal 72% (text) vs. 18% (robot); refusal in text transfers only partly to action

under the same action interface, a previous-generation model reached 0.88% on RoboDojo, whereas Astra reached 22.48% (Zhang et al., 2026c).

Broad benchmarks remain far from solved. When tasks span domains or demand active exploration, success remains low: the best model reached 19.0% on RobotWorld (Yang et al., 2026j), 23.2% on RoboQuest (Liu et al., 2026c), and 16.8% on HumanCLAW (Siyao et al., 2026). The high scores in Table 8 come from narrower suites. Recognition also does not imply completion: on HumanCLAW, the best model found the target in 64.9% of episodes but completed only 16.8% (Siyao et al., 2026).

Current weak points are physical and evidential. Failures cluster on two sides: physical action and perception, and the gathering and judging of evidence. On the physical side, major challenges remain in precise manipulation and contact-rich interaction (Gao et al., 2026b; Yu et al., 2026a), as well as dynamic control and locomotion; Astra, for example, solved none of the 11 locomotion tasks in RobotWorld (Yang et al., 2026j). Spatial understanding has improved but remains below human level. In the evaluation of DeepCybo Team et al. (2026), GPT-6 Astra at low reasoning effort scored 59.8 on VSI-Bench, against 45.4 for the best model in the original VSI-Bench evaluation and 79.2 for human evaluators on a 400-question subset (Yang et al., 2025b). On MMSI-Bench, the strongest proprietary model at publication answered about 40% of multi-image spatial questions correctly, against 97% for humans (Yang et al., 2026f). Body understanding is weak as well: in HumanCLAW, failures stem not from misrecognizing the target but from the model losing track of its own body (Siyao et al., 2026). On the evidential side, agents struggle to explore actively for task-relevant evidence (Liu et al., 2026c), to recognize what cannot be answered from observation, and to judge completion. On AbstainEQA, for example, the best frontier model recognized 42.79% of unanswerable questions, against 91.17% for humans (Wu et al., 2026). Agents also declare success prematurely (Jedlicki et al., 2026; Liu et al., 2026c); this was the most common failure of a frozen Astra policy on memory-dependent manipulation (Jedlicki et al., 2026). Automatic judges share this bias toward reporting success (Navasardyan et al., 2026).

Harness and protocol move results substantially. With the model held fixed, changes to verifiers, memory, interfaces, and post-call decision-making moved success rates by about 9 to 35 percentage points in four studies (Siyao et al., 2026; Wei et al., 2026a; Si et al., 2026; Ge et al., 2026). In a fifth, restricting agents to skills moved them by up to 52 points (Han et al., 2026). Changing only the interface illustrates the effect: with the same Astra planner, primitives, and call budget, Python cells raised success over tool calling from 63.1% to 71.7% (Si et al., 2026). Changing the model can have effects of similar size, as the RoboDojo comparison above shows. Studies that cross models with harness conditions remain small, with at most four models and a few conditions on the same tasks (Ge et al., 2026; Wei et al., 2026a; Han et al., 2026), so these figures show that both factors matter, but not which is generally larger. Combined differences in harness and protocol widen the spread further. For Astra on RoboDojo, one report gives 22.48% with one seed, a task recipe from the RoboDojo wiki, and direct end-effector control (Zhang et al., 2026c). Another gives 37.82% with the same wiki recipes plus experience reused from practice scenes and a higher reasoning effort (Li et al., 2026e). A third gives 46.9% in the official evaluation with a tool library built on the benchmark's own tasks (Gao et al., 2026b); its later version, evaluated by its authors on public test layouts, gives 69.9% with a small set of general tools and no task-specific input, and 85.0% with about a page of guidance per task (Gao et al., 2026a). A fourth gives 81.9% over five episodes per task, with in-context demonstrations and up to ten times the step budget (SenseNova et al., 2026). Embodied First Exam also finds substantial harness effects on a common set of 456 tasks: GPT-6 Astra solved 34% in Codex and 22% in Claude Code under its Standard setting (Wang et al., 2026d). Without a harness card, such numbers cannot be compared.

Tools add capability only beyond what the model can already do. More tools do not necessarily give better results. In EmbodiedSkillBench, restricting agents to skills selected for each task from a repository of 250 action, perception, and experience skills lifted GPT-5.6-Sol from 11% to 63% success but lowered GPT-6 Astra from 92% under native control to 61%. When Astra could decide for each skill whether to invoke, modify, or ignore it, it recovered to 85% (Han et al., 2026). Skills thus help an agent only as far as they exceed what it can already do. They can also reduce cost, even where they add no success. In the same benchmark, skill-only execution cut the token usage of the weaker agents by 58% to 85%. For Astra, which gained no success from skills, selective use still saved 14.1% of its tokens and 11.9% of its time relative to native control (Han et al., 2026). How tools are organized and scheduled matters as well. The motion, perception, skill, and episode tools of the uni-robot-agent harness share arguments, units, and feedback. In this harness, letting the agent keep reasoning while a tool runs shortened each of the five tasks tested, in one episode each, by 4.8% to 73.7%. Swapping two T-shaped blocks, for example, took 12.5 instead of 34.7 minutes, although the authors do not attribute the whole reduction to scheduling. Removing the harness's predefined skill library, by contrast, left an open-weight Qwen model with 96.6% of its score, so effective execution did not depend mainly on that library (Yuan et al., 2026a).

Paused worlds hide latency, cost, and risk. Frontier models take seconds to tens of seconds per decision (Figure 8), and many evaluations hide this delay by pausing the simulation. In a world that keeps running, latency is associated with collisions (Xu et al., 2026d); in one real-robot evaluation, deliberation dominated task time, which was about 100 times a person's (Artificial Analysis, 2026). Several designs therefore move work out of the model loop: program-level control (Si et al., 2026), VLA-first execution that consults the S2 foundation model only when needed (Sun et al., 2026b), and skills written offline by a frontier model that then run without it (Bai et al., 2026).

Safety learned in text transfers only partially to action. Pooled over the same six models, refusal fell from 72% of text requests to 18% of real-robot trials with end-effector control (Sun et al., 2026a). This drop highlights a substantial text-to-action refusal gap under the tested conditions, and a second study corroborates it (Grover, 2026). Safety must therefore be measured in the action loop, along the full chain from recognizing a risk to stopping in time.

What Is Not Yet Supported. Current evidence does not support cross-paper rankings, because interfaces, privileged information, budgets, and time semantics differ. Nor does it support claims of reliability, because many headline results come from single runs or few seeds. Aggregate rankings have begun to address the first problem. The Robot Intelligence Index, launched with 16 public simulated and real-robot benchmarks and extended to 22 by November 4, 2026, combines them into one rating per model using a weighted Bradley–Terry model. In this model, a benchmark's weight is the independence of its evaluators multiplied by the repeatability of its scores (Playground Research, 2026). The index itself shows, however, why a single rating cannot replace the profile we recommend: 13 of the 23 models that rank in the top three on some benchmark rank in the bottom half on another. Performance in the real world, in continuous time, and over long horizons also remains largely unmeasured, including memory across sessions and the fulfillment of user intent on real robots. Several further findings rest on a single study so far. One is that common embodied-reasoning benchmarks did not track closed-loop control (Yang et al., 2026b); another is that, when the world changes, memory fails mainly at replacing outdated state rather than at retaining history (Liang et al., 2026a). Finally, benchmarks saturate quickly. VLAs that exceeded 90% on LIBERO (Liu et al., 2023b) dropped to around 50% on LIBERO-PRO (Zhou et al., 2025c), yet an agentic system reached 94.0% on LIBERO-PRO's perturbed suites within weeks of Astra's release (Si et al., 2026). Evaluation must therefore move toward live execution, real robots, longer horizons, and the embodiment-specific sub-dimensions of Figure 9.

4.4. Limitations of Current Evaluation

Moving evaluation in these directions requires more than the principles above, because current evaluation remains limited even when they are followed. The most fundamental limitation is coverage: current evaluation spans too little responsibility, space, and time to test agentic robotics as we envision it. The others concern the validity of simulated results, comparability in continuous time, credible acceptance of complex tasks, and the cost of scaling real-world evaluation.

Coverage of Responsibility, Space, and Time. Table 9 lists the spatial scope and time span of benchmarks reviewed in this paper. Most of them test primitives and single tasks in a workcell or a room, within one episode. Building-scale scenes appear mainly in navigation and memory benchmarks, and site-scale settings in driving, aerial control, and city navigation. Evaluations beyond one episode use a continuous three-hour simulation, histories of earlier sessions, or learning across tasks. Responsibility is covered least: no benchmark we reviewed tests a workflow of interdependent tasks that share resources and deadlines, or a role sustained over days. How to construct benchmarks for workflows and roles at building and site scale, over days rather than minutes, remains open.

Validity of Simulated Results. Simulation is useful for S2 only to the extent that it preserves what S2 must respond to: realistic observations, informative feedback, failures, and changes in the environment. Simulators may have limited fidelity in modeling complex physical interactions like deformation,

Table 9 | **Spatial scope and time span of benchmarks reviewed in this paper.**

Benchmark	Spatial scope	Time span
LIBERO (Liu et al., 2023b)	Workcell	Single episodes
LIBERO-PRO (Zhou et al., 2025c)	Workcell	Single episodes
CALVIN (Mees et al., 2022)	Workcell	Single episodes[a]
RoboDojo (Chen et al., 2026b)	Workcell	Single episodes[b]
RoboTwin 2.0 (Chen et al., 2025d)	Workcell (dual-arm)	Single episodes
RoboLab (Yang et al., 2026c)	Workcell	Single episodes
VLABench (Zhang et al., 2025a)	Workcell	Single episodes
UniVTAC (Chen et al., 2026a)	Workcell (contact)	Single episodes
Bench2Dex (Yang et al., 2026h)	Workcell (two hands)	Single episodes
DexJoCo (Wang et al., 2026a)	Workcell (hands)	Single episodes
LADDER (Petkov et al., 2026)	Workcell	Single episodes
RoboMME (Dai et al., 2026)	Workcell	Single episodes
RoboMME-Interference (Rathi, 2026)	Workcell	After unrelated sessions[c]
RobotWorld (Yang et al., 2026j)	Workcell to outdoors	Interaction budgets
ALFRED (Shridhar et al., 2020)	Room (household)	Single episodes
RoboCasa365 (Nasiriany et al., 2026)	Room (kitchens)	Single episodes
RoboQuest (Liu et al., 2026c)	Room (kitchens)	Single episodes
MessyMem (Banwasi et al., 2026)	Room (kitchen)	Over three hours[d]
R2R (Anderson et al., 2018)	Building	Single episodes
R2R-CE (Krantz et al., 2020)	Building	Single episodes
HumanCLAW (Siyao et al., 2026)	Building (houses)	Single episodes
OmniNavBench (Sun et al., 2026c)	Building	Single episodes[e]
BEHAVIOR-1K (Li et al., 2024a)	Building (homes, offices)	Single long episodes
EMem-Bench (Liang et al., 2026a)	Building (homes)	One task after a history
EmbodiedCLBench (Huang et al., 2026a)	Household scenes	Learning, then test
RT-Safe (Xu et al., 2026d)	Site (city routes)	Continuous time

[a] Chains of five consecutive instructions. [b] 8 to 76 s of commanded motion (SenseNova et al., 2026). [c] Each task's memory holds its relevant prior demonstration followed by a controlled number of unrelated sessions. [d] 25 consecutive tasks without resetting memory or the environment. [e] About 49 s on average.

gear backlash, wear and fluid dynamics (Tobin et al., 2017), and on real robots, physical parameters such as motor strength, latency, and battery voltage change over time (Tan et al., 2018). Real-to-sim evaluation can track real-world rankings for some policies (Li et al., 2025a), but its effectiveness depends on the policy (Zhou et al., 2025e), and we know of no comparable validation for Embodied S2. Whether rankings of S2 foundation models and harnesses obtained in simulation hold on real robots, and which properties of a simulator decide this, remain open questions.

Comparability in Continuous Time. Many agentic evaluations pause the world while the model deliberates (SenseNova et al., 2026). One report estimates that the budgets of RoboDojo (Chen et al., 2026b) correspond to 8 to 76 seconds of commanded motion, against an average of 55 minutes of computation. It summarizes the effect: "The model reasons as though the world waits for it" (SenseNova et al., 2026). A fair comparison in a world that keeps moving must count each system's own reaction latency as part of its performance. Yet such a world does not repeat, so runs differ in timing as well as in decisions. How to control for timing variation across runs at a cost that real robots can bear is an open question.

Credible Acceptance of Complex Tasks. A robot's own judgment that a task is complete differs from external acceptance of the result. Automatic judges cannot yet provide such acceptance: VLM judges of robot task success have shown a systematic bias toward predicting success when the evidence is ambiguous (Navasardyan et al., 2026). False alarms are also costly: under a stop-and-reset model of intervention, even an idealized failure detector can afford a false-alarm rate of only 2.0% to 3.3% before unnecessary resets outweigh the savings from stopping failures early (Ahmedt-Aristizabal et al., 2026). Even an accepted result is hard to attribute to specific design choices. RobotWorld, for example, cautions that its five-model snapshot supports "matched-task observational comparisons, not controlled ablations" of history, code control, or reasoning budget (Yang et al., 2026j). Workflows and roles make acceptance harder still, because their acceptance criteria span many subgoals and depend on the user's intent. How to obtain acceptance judgments with known error rates for such work is open.

Scaling Real-World Evaluation. Because real-robot trials are usually slow and expensive, many reported real-robot results rest on few trials per task. This scarcity is especially consequential because results for the same task vary across runs and across the testers who prepare each trial (Yakefu et al., 2025). Distributed real-robot evaluation (Atreya et al., 2025), autonomous real-world evaluation (Zhou et al., 2025e), remote access to standardized hardware (Chen et al., 2026b; Yakefu et al., 2025), and third-party real-robot leaderboards (Robocurve, 2026) are promising steps toward evaluation at scale. For evaluating Embodied S2, one reported protocol offers a reasonable minimum standard (SenseNova et al., 2026): it sets a fixed wall-clock budget, stations an operator at the emergency stop for the entire trial, and requires a human to judge success from the physical outcome rather than from the model's own completion claim. The open question is how to obtain statistically reliable evidence as evaluation extends to larger spaces, longer time periods, and broader task responsibilities—settings in which each trial costs more.

5. Digital Agents versus Physical Agents

Agentic robotics inherits a central principle from digital agents: a foundation model can solve complex tasks by repeatedly observing an environment, invoking tools, and verifying the results. Its architecture, a foundation model in a loop with tools, context, and memory, draws on two lineages. One is digital agents, which have advanced rapidly on coding, computer use, and research tasks; the other is earlier work on robots driven by language models. The success of digital agents, however, rests on favorable properties of the digital world: interfaces that the model reads and writes directly,

outcomes that can be checked automatically, abundant data, and environments that are cheap to copy and reset.

Physical environments weaken each of these properties and add constraints of time, irreversibility, and embodiment. The challenge is therefore not to give a digital agent cameras and robotic tools, but to determine which assumptions no longer hold, which principles still transfer, and what must be rebuilt for reliable physical execution. Our answer is that the principles of digital agents largely transfer, but the assumptions that make them reliable do not. We first review what digital agents can do and why, and then take up these three questions in turn.

5.1. What Digital Agents Can Do and Why

The length of software tasks that frontier agents can complete has grown steadily. METR (Kwa et al., 2025) measures this progress as a *time horizon*: the duration, in human working time, of tasks that an agent completes with a given probability. The horizon has doubled roughly every seven months since 2019 (Kwa et al., 2025; METR, 2026a), and about every four months since 2023 on METR's revised task suite (METR, 2026a). By 2026, the best measured 50% horizons exceeded ten hours, and METR cautions that measurements above 16 hours are unreliable with its current task suite (METR, 2026b). These results measure capability on particular distributions of software tasks with checkable specifications, not reliable long-horizon autonomy in general. Reliable completion is harder: time horizons at 80% success are several times shorter than at 50% (Kwa et al., 2025; METR, 2026b).

A digital agent pairs a foundation model with a harness of tools, context management, memory, and an execution sandbox. Its results benefit from four conditions that software happens to satisfy. We call them *enabling conditions*: they recur across the systems behind much of the recent progress, and we treat them as design guidance. First, *the interface is native to the model*: state and actions take the form of the model's own inputs and outputs. State is text that the agent reads, and actions are text or code that it emits directly through an agent–computer interface (Yang et al., 2024). *Outcomes are generally verifiable*: compilers, tests, and program execution return objective signals quickly and at scale, signals that support iteration at test time (Yang et al., 2024) and reinforcement learning with verifiable rewards (Lambert et al., 2024; Guo et al., 2025). *Data are abundant*: public code, documentation, and discussion accumulated over decades give pretrained models broad knowledge of the domain; Codex, for example, was trained on code from 54 million public GitHub repositories (Chen et al., 2021). *Environments are cheap*: a sandbox with an executable runtime and unit tests can be copied, run in parallel, reset, and replayed, so that agents can be trained and evaluated on thousands of real tasks (Pan et al., 2025). The first two conditions shape every step of the agent's loop; the last two determine how much the agent can learn before deployment and how widely it can be tested. Wei (2025) links verification to learning through a *verifier's rule*: the ease of training AI to solve a task is proportional to how verifiable the task is.

Because these conditions are properties of a task setting rather than of programming per se, capability has spread to settings in which they could be established. These settings include software engineering (Jimenez et al., 2024) and professional work across many occupations (Patwardhan et al., 2025). They also include computer use, in which agents operate graphical interfaces in real operating systems (Xie et al., 2024), and visual coding, in which generated code is rendered and compared with a design or specification (Si et al., 2025; Xu et al., 2026e). Most recently, they include 3D content, in which agents write programs for modeling software and game engines (Gao et al., 2026c; Ye et al., 2026b). Most of these extensions weaken some of the conditions, and capability lags where they weaken. Computer use weakens the native interface by replacing text with pixels and clicks. Its time horizons are 40 to 100 times shorter than those in software and reasoning domains, although they are increasing at similar rates (Kwa and Cheng, 2025). In 3D content, verification

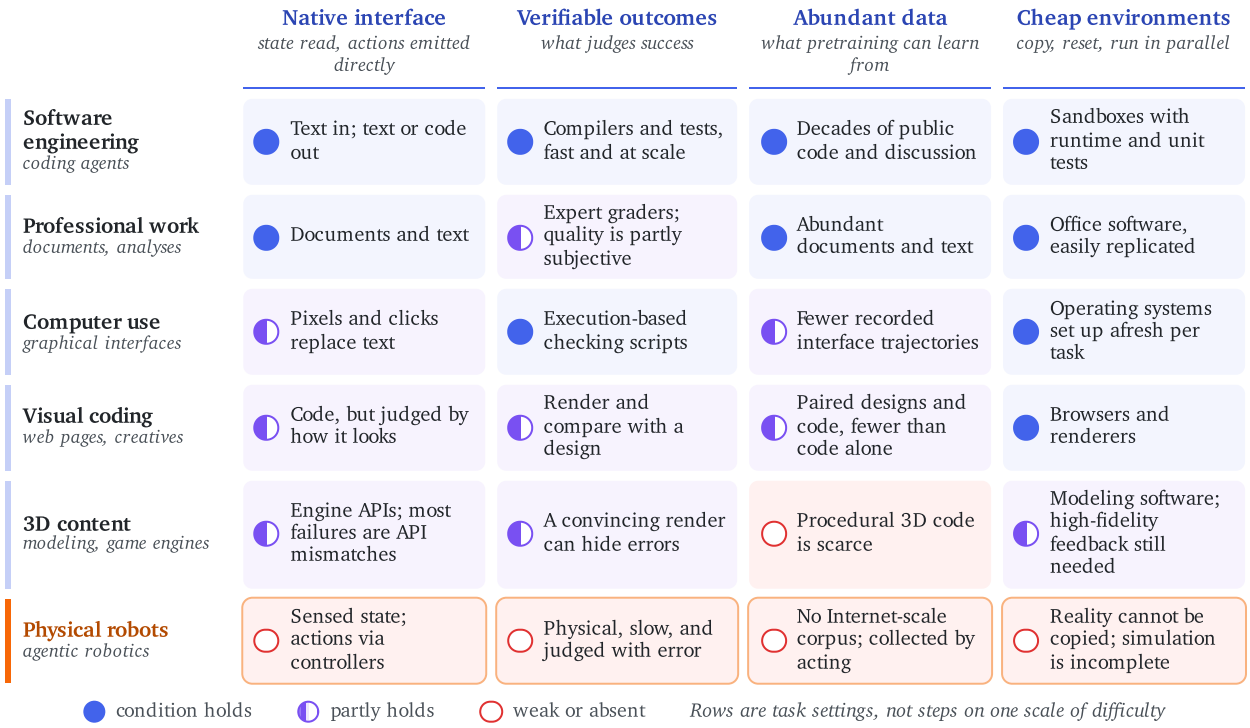


Figure 10 | **The four enabling conditions of digital agents, and how far each domain satisfies them.** The domains illustrate how the conditions vary across task settings rather than forming a strict progression of capability or difficulty.

weakens, because a convincing render can hide incorrect spatial relations (Ye et al., 2026b). Progress there depends on execution environments that return high-fidelity feedback and on more procedural data (Gao et al., 2026c). Figure 10 summarizes our assessment for each domain. The domains differ in inputs, outputs, and verification requirements and do not form a strict progression of capability or difficulty. Physical robots weaken all four conditions at once, as the next subsection shows.

5.2. Which Assumptions Fail in the Physical World

Each enabling condition enters digital-agent design as an assumption. This subsection asks which of these assumptions, and which further ones, fail in the physical world, and what problem each failure creates. Table 10 contrasts typical sandboxed coding agents, the setting in which digital agents are most mature, with physical agents along eleven aspects, grouped by the four enabling conditions and by the constraints that the physical world adds. Many of the differences trace back to two abilities that a coding sandbox provides and the physical world does not: reading the state of the world and judging the outcome of an action (Wang et al., 2026e). Not every digital task offers these abilities either: computer-use and web agents also face environments that change, hide state, or cannot be rolled back, and their limits foreshadow the physical setting.

State cannot be read directly, and actions do not take effect directly. Digital-agent design assumes that the agent reads the environment state, and acts on it, through documented interfaces. A coding sandbox offers such interfaces but a physical environment does not (Wang et al., 2026e). Digital agents already falter when their environment is less transparent: they perform worse on tasks without clear feedback loops (Kwa et al., 2025). For a robot, partial observation is the usual condition: it observes only a local, momentary slice of the world through noisy sensors and must assemble and

Table 10 | **Typical coding agents and physical agents compared.** Rows are grouped by the four enabling conditions and by the constraints that the physical world adds.

	Coding agent	Physical agent
Environment	A designed artifact: files, code, APIs	Objects, people, and spaces
Native interface		
Input and state	Directly readable files, logs, and tool outputs	Partial, noisy, occluded, and delayed sensing; contact and force are invisible to cameras; state must be constructed
Action interface	Tool calls, file edits, commands, tests	Skills, trajectories, and actuators
Verifiable outcomes		
Verification	Fast, reproducible automated tests	Costly, slow, hard to reproduce
Abundant data		
Data	Decades of public code and text; trajectories generated in sandboxes	No Internet-scale corpus; collected by operating robots; force and touch data scarcest
Cheap environments		
Training environment	Sandboxes that can be copied, reset, and run in parallel	Simulators faithful for rigid-body motion but not for all contact physics; real resets need time and often a person
Constraints added by the physical world		
Loop	Call a tool, then check the result	Continuous sensing and motion with online correction
Time	Can wait, pause, retry, or run faster	The world does not pause or speed up; low-level control must stay real-time
Failure and recovery	Version control, snapshots, and retries	Collisions and damage; physical state cannot be rolled back
Body and transfer	No body; easy reuse across environments	Bodies heat, wear, drift, and break; new embodiments need re-adaptation and re-validation
Main difficulty	Managing context, tools, dependencies, and progress over long tasks	Acting reliably under partial observability and physical constraints

maintain a model of the rest (Wang et al., 2026e). Physical outcomes also depend on quantities that cameras do not see, such as mass, friction, compliance, and contact force. Tasks such as insertion are difficult to complete robustly with vision alone (Chen et al., 2026a). Touch carries much of this information: people correct their grip from touch within tens of milliseconds (Section 3.4), and in simulated tasks a frontier multimodal model interpreted vision-based tactile images when given them (Yang et al., 2026a). Nor do actions take effect directly: they are skills, trajectories, and actuator commands whose meaning is set by controllers, reference frames, and units, and whose effects arrive with delay. The resulting problem is one of grounding: state must be estimated, together with its source, age, and uncertainty, and actions must carry explicit physical semantics.

Outcomes are not automatically verifiable. Digital-agent design usually assumes that outcomes can be checked automatically, quickly, and reproducibly. Even with such checks, verification is imperfect: agents asked to evaluate their own work tend to praise it confidently (Anthropic, 2026a), and

benchmark oracles can themselves be flawed (OpenAI, 2026c). Physical outcomes must instead be observed, judged, and sometimes re-measured, and both false successes and false failures are costly. A false success leaves a task unfinished or a hazard in place, and VLM-based judges are biased toward this error when the visual evidence is ambiguous (Navasardyan et al., 2026). By the verifier's rule (Wei, 2025), this scarcity of credible verification is itself a central obstacle to training physical agents. The resulting problem is to verify outcomes using evidence: each verdict must state which of the four questions of the agentic loop it answers (acceptance, termination, physical effect, or intent) and come from a calibrated verifier.

Data are not abundant. Digital-agent design assumes that pretraining has already covered the domain. Unlike natural language and two-dimensional vision, dexterous manipulation has no Internet-scale corpus (Hoque et al., 2025). Robot data are produced by operating robots, at the pace of the physical world. The 76,000 trajectories of DROID, about 350 hours of interaction, took 50 data collectors twelve months on three continents (Khazatsky et al., 2024). Teleoperation of this kind lacks scale and diversity, and synthetic data bring a gap to reality (Luo et al., 2025). Human video scales more easily, but large collections such as Ego4D (Grauman et al., 2022) lack native hand-pose annotations and action labels (Hoque et al., 2025), and video records no contact forces. Touch is scarcest: acquiring large-scale and reliable tactile data in the physical world remains costly (Chen et al., 2026a). The resulting problem is one of supply: abundant data must be manufactured by a data engine that turns the operation of robots into training data.

Environments are not cheap to copy and reset. Digital-agent design assumes that environments can be copied, reset, and run in parallel. Where simulation is faithful, the digital recipe returns: with thousands of simulated robots trained in parallel, a quadruped learns to walk on uneven terrain in about twenty minutes (Rudin et al., 2022). Where it is not, policies trained in simulation often fail to transfer to the real world (Tan et al., 2018), and simulation pipelines have only begun to cover complex physics such as deformable objects and fluids (Qin et al., 2026c). The alternative, practice in the real world, requires the scene to be reset after each attempt; this is typically handled by automated reset skills, with manual reset as a fallback when automatic reset fails (Li et al., 2026d). The cheap, parallel, and resettable environment of digital agents therefore exists only in part. The resulting problem is twofold: every physical experiment has a cost that must be budgeted, and results obtained in simulation must be validated against reality.

Assumptions beyond the Four Conditions. Three further assumptions of digital-agent design fail for reasons other than a missing condition: the physical world adds constraints of time, irreversibility, and embodiment. Web and computer-use agents meet milder forms of the first two, but the third has no close digital counterpart. Embodiment amplifies the first two and makes all three matters of safety and control.

- *Deliberation is not free.* While a robot deliberates, the world around it continues to evolve (Black et al., 2025c). A digital agent that takes a minute to respond inconveniences its user. A robot that takes a minute may act on a scene that no longer exists, with consequences for both performance and safety.
- *Errors cannot be undone.* Coding agents undo errors with version control, snapshots, and sandboxes (Anthropic, 2025b). Many physical actions cannot be undone (Wang et al., 2026e), and the world offers no general counterpart to restoring a software snapshot (Zhu and Cai, 2026). Safety must therefore be enforced before an action is taken, and below the model.
- *Bodies do not stay fixed.* Coding agents have no body to manage (Wang et al., 2026e). A robot's capabilities depend on its body, and that body's condition changes with use. As Sagar and Senanayake (2026) note, over hours of operation motors gradually heat up, current saturates near contact, and

voltage sags under load, so that the same command can produce a weaker, delayed, or noisier motion. Motor drivers reduce current near thermal limits and can leave closed-loop control beyond them (ODrive Robotics, 2026). Actuators can also switch off torque after overheating until they are rebooted (ROBOTIS, 2026). In simulated LIBERO (Liu et al., 2023b) tasks, one study perturbed only hardware and setup factors, such as camera mounting, calibration, actuation gain, control rate, and gripper pads. At the highest severity, the pooled success of the eight VLA models tested on all four suites fell from 93.2% to 50.9% (Cui et al., 2026).

5.3. What Transfers

Beneath the practices that make digital agents work lie principles about how an agent organizes its work, not about what its environment provides. Because they do not depend on the enabling conditions, five such principles carry over to physical agents and already shape our architecture.

- *Close the loop on evidence.* Agents are language models "using tools based on environmental feedback in a loop" and need ground truth from the environment at each step (Anthropic, 2024a). For a robot, the principle requires the agent to observe and assess an action's relevant effects before treating its intended outcome as established or making a decision that depends on that outcome. The agent may continue planning or monitoring while an action is in progress, provided that subsequent actions respect the execution and safety constraints.
- *Treat tools as contracts.* Tools are contracts between deterministic systems and non-deterministic agents (Anthropic, 2025d). Execution contracts extend this principle to physical calls. A capability contract, for example, states a skill's inputs and units, preconditions, expected effects, time limits, and behavior on cancellation, so that S2 can reason about a call before making it and interpret its result afterward.
- *Manage context explicitly, and keep durable state outside it.* Context is a finite resource with diminishing returns, managed through compaction, structured notes, and sub-agents (Anthropic, 2025a). Long-running agents also keep progress files (Anthropic, 2025b) and a session log outside the model's context window (Anthropic, 2026b). The embodied harness applies the same discipline to the world, the user, and the body: it keeps task progress, maps, and memories outside the model and supplies them when a decision requires them.
- *Separate doing from judging.* Long-running agents benefit from end-to-end tests, because models tend to mark work as complete without proper testing (Anthropic, 2025b). Separating the agent that performs the work from the agent that judges it also helps (Anthropic, 2026a). In the embodied harness, the verifier is the judge: it assesses each step and the task from evidence rather than from the model's claim that it has finished.
- *Package capabilities as reusable skills.* Digital agents load packaged skills when a task needs them (Anthropic, 2025c). An agent in Minecraft accumulated a library of skills written as code (Wang et al., 2024a). For robots, S1 skills and the programs that S2 writes play this role (Figure 4): a capability validated once can be invoked many times, and reuse also reduces the number of decisions that S2 must make at run time.

These principles transfer intact, but their digital implementations relied on the enabling conditions. The next subsection therefore turns to what physical agents must rebuild.

5.4. What Must Be Rebuilt

Where an enabling condition fails, a physical agent must rebuild it. Where the physical world adds a constraint, the agent must operate within it. Success at either cannot be assumed, so claims about

the agent must be validated in the running world. Figure 11 organizes these requirements in three layers: the foundations that the system must have, the constraints under which it must operate, and the evaluation that validates its claims.

Foundations. Four foundations replace the four enabling conditions.

- *Grounded state and action interfaces in place of a native interface.* On the state side, the harness maintains a belief state: a persistent and revisable estimate of world, user, and body state. The estimate is built from partial multimodal sensing, including touch and force when a task depends on contact. On the action side, each call passes through an execution contract. The contract states the call's reference frame, units, preconditions, capability limits, and expected effects. It defines the feedback reported during execution and the behavior on cancellation, and it enforces safety checks, permissions, and postconditions that the model cannot bypass. Because a physical call can be executed exactly as commanded and still fail to produce its effect, the contract separates the termination of a call from the verification of its effect.
- *Calibrated verifiers in place of exit codes.* Success detectors, progress estimators, and independent judges with known error rates make physical outcomes checkable. They serve both the loop at run time and the rewards for learning, so the quality of verification can become a bottleneck for how far a physical agent improves (Wan et al., 2026; Zhou et al., 2026b).
- *Embodied data engines in place of Internet-scale corpora.* Physical agents must produce their own data from teleoperation, human video, and simulation. Above all, they must produce it from deployment, in which autonomous attempts and human corrections across a fleet become training data (Wang et al., 2026i).
- *Simulation calibrated against reality in place of sandboxes.* Simulators and digital twins (Todorov et al., 2012) cover part of the need but not the deployed world (Wang et al., 2026e). They provide scale where their physics is faithful, and real-world experience corrects them where it is not. Autonomous resets make practice in the real world more affordable (Li et al., 2026d).

Runtime Constraints. The three constraints that the physical world adds govern how the system operates: time imposes deadlines and interruptions, irreversibility requires safety before execution, and a changing body requires adaptation.

- *Deadlines.* The agent treats latency as a hazard, decides within deadlines rather than deliberating at length, and delegates fast loops to S1 and S0. This division of labor mirrors the separation of fast control from slow reasoning in Helix (Figure AI, 2025b) and of fast interaction from background reasoning in Thinking Machines' interaction models (Thinking Machines Lab, 2026).
- *Interruptions.* The agent accepts new input while it reasons or acts. Before dispatching a conflicting action, it confirms the state of the calls in flight and of the body, because a cancellation request alone does not establish that execution has stopped.
- *Safety before execution.* Because many actions cannot be undone, the agent anticipates consequences, prefers reversible or information-gathering actions under uncertainty, and asks for help or refuses when needed.
- *Body-state adaptation.* The agent tracks heat, wear, and power, and adapts to them, for example by conditioning its actions on actuator telemetry (Sagar and Senanayake, 2026). It re-validates its skills after maintenance or a change of embodiment.

Evaluation. Claims about a physical agent must be validated under the conditions in which it operates.

- *Live execution.* Paused-world evaluations remain useful diagnostics because they separate the quality of a decision from its timing. Paused and live execution, however, measure different abilities,

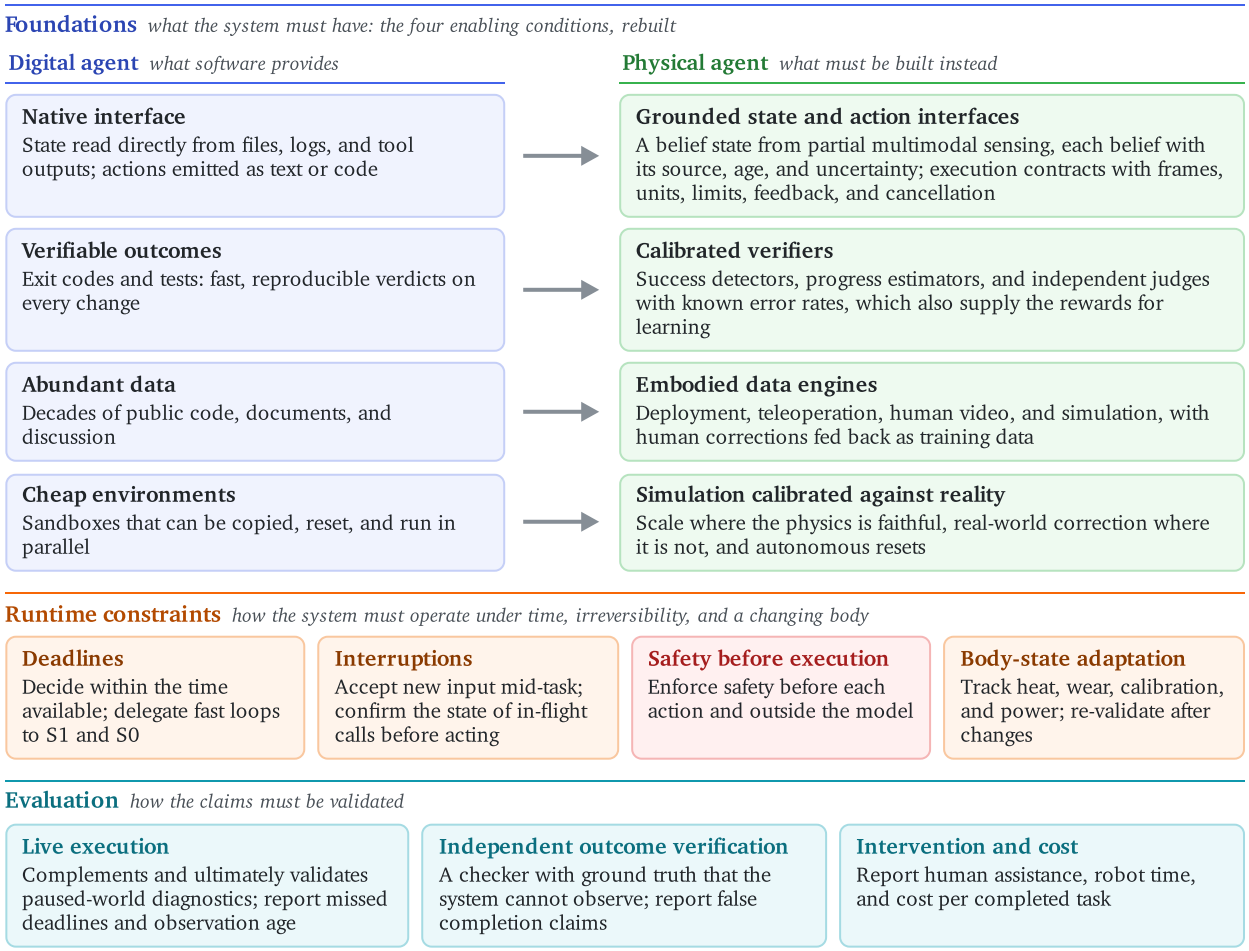


Figure 11 | **What physical agents must rebuild.** Foundations replace the four enabling conditions of Figure 10; runtime constraints, which follow from time, irreversibility, and a changing body, govern how the system operates; and evaluation in the running world validates its claims.

so live evaluation complements, and ultimately validates, paused-world results.

- *Independent outcome verification.* Success should be decided by a checker with access to ground truth that the system cannot observe.
- *Intervention and cost.* Each result should report the human assistance and interventions it required, the robot's time, and the cost per completed task.

These requirements distinguish Embodied S2 from a digital agent: they reshape both the capabilities required of S2 and the conditions under which those capabilities must be evaluated. Progress on digital-agent benchmarks, however rapid, therefore does not translate directly into the ability to delegate physical work; that ability must be measured and built in the physical world.

6. Road Ahead

Agentic robotics broadens the question by which progress in robotics is judged: from what a robot can do on command to what it can be entrusted to do on its own. General foundation models already supply much of the knowledge and reasoning that such trust requires. Turning these capabilities into dependable physical work, however, raises problems of timing, persistent state, verification,

adaptation, and learning from experience that no single model release will resolve. Where the field goes next will therefore depend on how it measures progress and which of these problems it addresses. This section proposes a ladder of delegable work as a measure of progress and places current systems on it. It then identifies research directions that climbing the ladder requires and closes with open questions about how agentic robotic systems should be built.

6.1. Ladder of Delegable Work

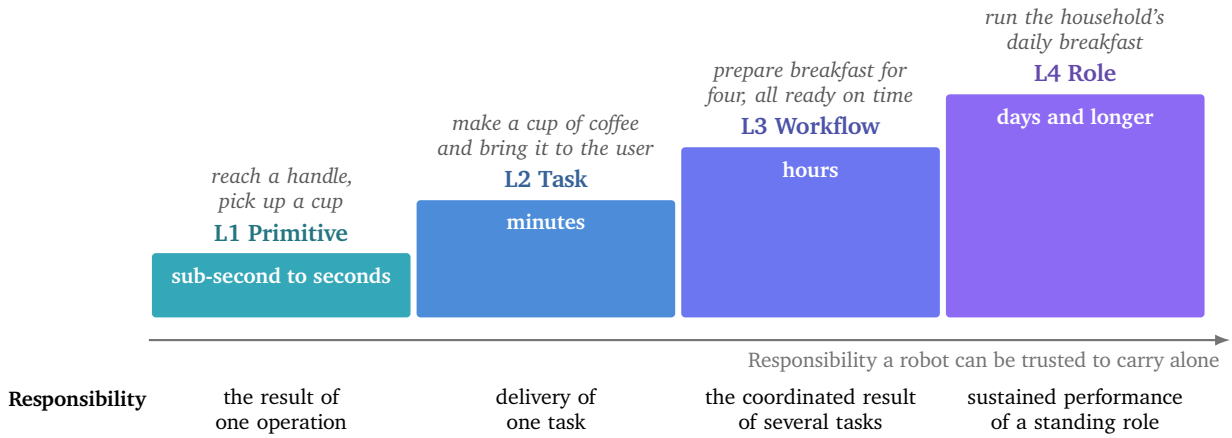


Figure 12 | **Ladder of delegable work.** Progress is measured by the responsibility that a robot can be trusted to carry alone. Each of the four levels includes those below it; typical durations are illustrative.

Reaching embodied intelligence that is genuinely useful calls for a high-level ladder that guides progress toward it. We propose such a ladder for agentic robotics: it measures progress by the scope of responsibility that a robot can reliably carry alone, independently of its body, domain, or implementation. The ladder has four levels (Figure 12). At L1, the *primitive* level, the robot executes a local skill, such as picking up a cup, and checks whether the intended result has been achieved. At L2, the *task* level, the robot coordinates skills to deliver a complete task, such as making and serving coffee: Embodied S2 decomposes the goal, selects actions, monitors execution, and responds to failures until the acceptance criteria are met. At L3, the *workflow* level, the robot coordinates interdependent tasks toward an overall goal, such as preparing breakfast for four: Embodied S2 manages timing, priorities, and shared resources, and revises its plan as needed. At L4, the *role* level, the robot sustains an agreed responsibility across successive tasks and workflows, such as managing daily breakfast around household schedules and preferences. This responsibility continues after individual workflows end, until it expires or is withdrawn. Each level builds on those below it, extending the scope over which the robot must connect reasoning, execution, and verification.

Spatial scope and complexity complement the ladder: within each level, progress means reliably fulfilling the delegated responsibility across broader environments and more demanding activities. As the operating scope expands from a workcell to a room, a building, or a site (Figure 13 and Table 11), robots must integrate observations across locations, remember relevant objects and events, and seek information beyond their immediate view. These demands, like those of complex activities, fall on Embodied S2, which maintains goals and context while guiding action through the loop of observation, reasoning, execution, and verification. We therefore expect the value of agentic robotics to grow with scope and complexity.

For current systems, they can execute primitives and complete tasks in unstructured environments, but have only begun to reach the workflow level, and have yet to reach the role level.

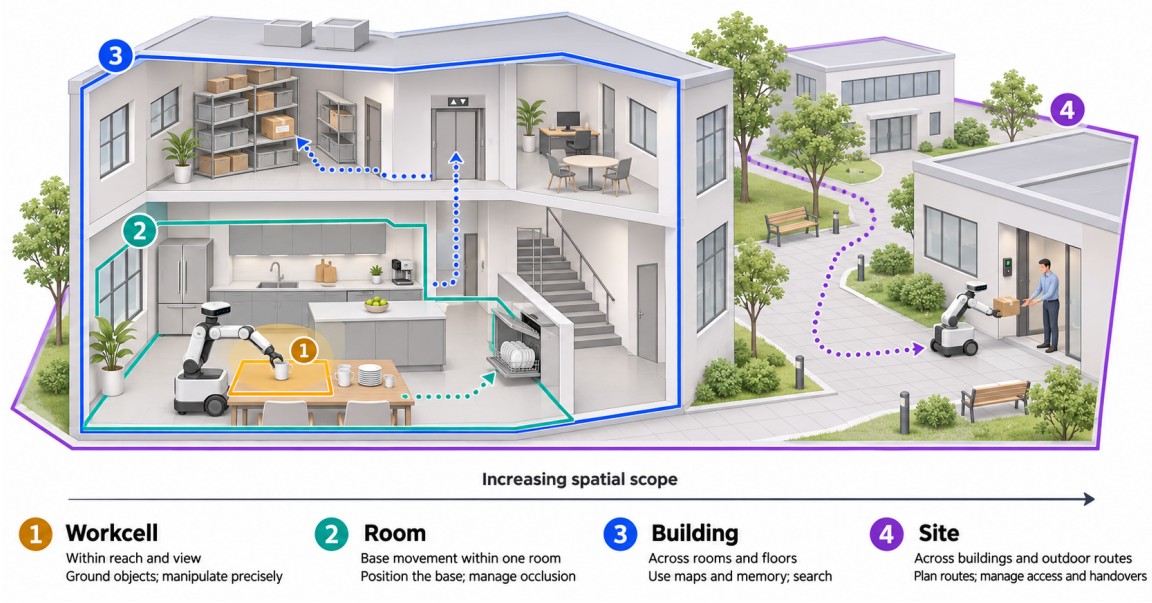


Figure 13 | **Spatial scope of delegated work, from a workcell within the robot's reach to a site with several buildings and outdoor routes.** The illustration is conceptual.

- *Primitives and tasks.* Reliability on narrow tasks is rising. Sunday reports 99.1% success in folding laundry in unseen homes with a single model checkpoint and no adaptation per home (Sunday Robotics, 2026). Dyna Robotics reports that Dyna-2 passed customer acceptance criteria at a rate of 87% at deployment sites whose data it had not seen (Dyna Robotics, 2026a). Longer tasks are now completed too: $\pi_{0.5}$ performs long-horizon cleaning in unseen homes (Black et al., 2025a), and Figure AI reports a four-minute task of unloading and reloading a dishwasher (Figure AI, 2026a).
- *Workflows.* The first workflow-level results are single-company demonstrations, such as the hour-long laundry-room workflow that Dyna Robotics reports for Dyna-2.1 (Dyna Robotics, 2026b). In simulation, a mobile manipulator with persistent memory has chained tasks for more than three hours (Banwasi et al., 2026).
- *Roles.* To our knowledge, no robot system in an unstructured environment has been reported to hold a role without human repair. In structured settings, an autonomous materials laboratory operated for 17 days (Szymanski et al., 2023), although the novelty of its products was later disputed (Leeman et al., 2024; Szymanski et al., 2026); sustained operation alone does not show that a system holds a role. Capabilities that a role relies on have begun to appear as harness prototypes that keep state outside the model. Some keep notes about a house that persist across tours and raise navigation success (Zhao et al., 2026b) or learn properties such as "locked" through interaction (Banwasi et al., 2026). Others carry user preferences across a program restart on a real robot (Wang et al., 2026l) or maintain a revisable model of the robot's own body (Lou et al., 2026). Remembering such information, however, is not the same as organizing work around it. It remains to be shown that a robot can use these capabilities to manage goals across tasks and sessions and to sustain a role.

6.2. Future Research Directions

Climbing the ladder, and doing so across larger spaces and longer periods, will require progress well beyond what current systems demonstrate. The evidence reviewed in this paper points to four

Table 11 | **Spatial scope of delegated work.** A claim of reaching a level should state the scope, duration, and complexity of the work.

Scope	Extent	What it adds for S2
Workcell	Within reach and view	Grounding objects and actions in view; contact and precision
Room	Within view, but the base must move	Where to place the base; reachability that changes as the robot moves
Building	Beyond what the robot can observe at once	Maps and memory of places and objects; search and exploration
Site	Several buildings and outdoor routes	Long routes; other people and robots; handover between areas

directions that need further research. They concern embodied reasoning that meets physical deadlines, control whose state and outcomes remain trustworthy across tasks and spaces, adaptation to new environments and bodies at low cost, and continual improvement from physical experience. The list is not exhaustive.

6.2.1. Embodied Reasoning Under Physical Constraints

S2 must understand the world, anticipate the consequences of its actions, and decide within limited time and computation, yet current S2 foundation models fall short in both speed and decision quality. With a frontier model as its S2 foundation model, S2 decides once every several to forty seconds (Xu et al., 2026d; Yang et al., 2026d; Galbot Team et al., 2026), far slower than the S1 skills it directs (Figure 8). More deliberation helps when the world waits (Gemini Robotics Team, 2025b) but hurts when it moves (Xu et al., 2026d). Decisions remain weak where embodiment matters most: judging success from ambiguous evidence (Navasardyan et al., 2026), anticipating the consequences of actions (Sun et al., 2026a), and deciding when to stop exploring (Liu et al., 2026c).

We therefore advocate making capped scaling an explicit training and system objective: maximizing decision quality subject to a response deadline, evaluated without pausing the world and reported as success–latency–cost curves. Three questions follow. First, when is further reasoning worth its time, and when should the robot observe again or act at once? Second, which decisions must be made online, and which can be reused or turned into skills? Partial answers to both include routing each subtask to a planner of appropriate strength (Dao et al., 2026) and launching chains of primitives with one call (Xu et al., 2026a). A fast policy can also consult a frontier model only for recovery (Sun et al., 2026b), and authoring skills offline removes the slow model from the loop (Bai et al., 2026). Third, how can the task benefit per unit of time and computation be raised? Distillation into small models that run locally is one candidate (Izzo et al., 2026; Liu et al., 2026b; Baumann et al., 2026; Foundation Model, Li Auto Inc, 2026).

6.2.2. Persistent and Verifiable Control Across Tasks and Spaces

As responsibility extends to workflows and roles, and the operating scope to buildings and sites, the system must keep its state trustworthy while organizing and verifying the work. However, current harnesses are built and evaluated almost entirely on single tasks in a workcell or a room (Section 4.4).

Three questions stand out. The first is how to represent and maintain state that spans rooms, floors, and days, so that the robot knows what it still knows. Candidates include augmenting spatially

grounded scene graphs with outcomes learned through interaction (Banwasi et al., 2026) and recording a subgoal as complete only after evidence of completion (Zhao et al., 2026a). The second is how to coordinate the interdependent tasks of a workflow, which share information, resources, and deadlines. The third is how to divide judgment between the model and the execution system: which checks belong inside skills and verifiers, and which decisions must return to S2. Such a division can also keep a slow decision-maker compatible with continuous control. A precedent exists one level down, within S1: real-time chunking lets a VLA generate the next action chunk while the current one executes, remaining robust to inference delays of more than 300 ms (Black et al., 2025c). S2 decisions, however, take seconds rather than milliseconds.

6.2.3. Efficient Adaptation Across Environments and Embodiments

Deployment demands that capabilities transfer to new sites and embodiments with minimal configuration, data, and human effort. Yet current success remains highly dependent on the specific embodiment (Section 4.1). We advocate three lines of work. The first is reuse of knowledge and experience: in-context adaptation and skill libraries of validated repairs reduce what must be learned anew (Lu et al., 2026), and compiled skills can be frozen and carried to other agents (Bai et al., 2026). The second is capability interfaces, such as capability cards and device descriptions, that let S2 plan for a body it has not used before. The third is skill adaptation: in one study, a single policy, conditioned on examples of commanded actions and resulting motions, adapted in context to unseen controller modes and gains (Yang et al., 2026e).

Fitting a system to a body also includes choosing loop rates and the placement of computation, as deployed stacks do. Helix 02 runs whole-body control at 1 kHz and a visuomotor policy at 200 Hz (Figure AI, 2026a), beneath an upper module that ran at 7 to 9 Hz in the first Helix (Figure AI, 2025b). An on-robot model avoids the network round trips and outages that a cloud model must tolerate. Conversely, skills, sensory feedback, and the body should be designed so that S2 can plan, execute, and verify more easily. Examples include skills that report their terminal states with evidence, and touch sensing: adding tactile observations to the in-context demonstrations of an embodied task agent raised its average success on four simulated tasks from 15.75% to 24.00% (Wang et al., 2026k).

6.2.4. Continual Improvement from Physical Experience

A robot holding a role for weeks keeps meeting new tasks and changes in its environment and body that no training run anticipates. It must therefore learn from physical experience which object properties, user preferences, causes of failure, and skills are worth keeping. We distinguish three kinds of change. *Adaptation* adjusts behavior to the current task, environment, or feedback. *Continual improvement* retains a change to memory, tools, skills, the harness, or a policy, and verifies that it improves later work. *Recursive self-improvement* (RSI) improves the process that produces such changes, making subsequent improvement more effective.

Robots have long improved through autonomous data collection, learned rewards, and reinforcement learning, and DrEureka used language models to design reward functions and domain randomization for sim-to-real transfer (Ma et al., 2024b). What recent agentic systems add, mostly around frozen foundation models, is a loop that diagnoses failures from operating experience and revises the system, usually with a foundation model proposing the revisions (Tian et al., 2026; Lee et al., 2026b; Zhang et al., 2026e; Deng et al., 2026). RegenHarness specifies such a loop as an evidence-gated protocol (Wang et al., 2026c). This loop starts with attribution: deciding whether a failure came from perception, planning, a skill, or an interface, which part should change, and why the change should help future tasks. Mechanisms differ in what they change. Some carry lessons forward in context: retrospective, cross-stage failure attribution let GPT-6 Astra learn manipulation

skills on a real Franka arm without demonstrations or parameter updates (Yu et al., 2026b), and LEXP reuses lessons extracted from execution evidence (Zeng et al., 2026b). Others change tools, skills, and harness code around a frozen model (Lu et al., 2026; Tian et al., 2026; Deng et al., 2026). Still others change weights, adopting a fine-tuned policy only after verification (Zhang et al., 2026e). Since every physical trial costs robot time and may be irreversible, a system must also decide which experience to acquire next, a question current methods rarely address explicitly.

A change counts as an improvement only if it helps later work and preserves existing ability; it must also be retired once outdated. Helping later work depends on practice tasks that grow harder and can be verified automatically. Constructing simulated practice tasks from capabilities identified in offline data (Wang et al., 2026h) and practicing in self-built simulated twins (Rao, 2026) are first steps, although, to our knowledge, progressive difficulty has not yet been demonstrated physically. Preserving ability requires an adoption rule: in one study, accepting every proposed revision degraded a harness, whereas a champion–challenger rule improved held-out results (Wei et al., 2026b). Retirement matters because a replaced gripper can invalidate learned behavior (Luo et al., 2026a).

Gains do not yet accumulate reliably. In EmbodiedCLBench, experience from basic tasks improved advanced ones, but the improvement did not grow with more learning samples (Huang et al., 2026a). With equal total development time, repeated short attempts matched or outperformed longer learning sessions in some model–task combinations; which did better depended on whether a model reached its first success within one short attempt (Wang et al., 2026g). In a 123-round study of agentic skill discovery in simulation, accepted changes kept passing their tests while the target task never succeeded: "the number of accepted changes measures activity, not progress" (Wang, 2026). A credible claim of improvement should therefore report gains on held-out tasks and show no regression on earlier tasks. It should also separate the contributions of the S2 foundation model, the harness, and S1, under the controls of Section 4.2 and within applicable safety requirements.

Claiming RSI further requires showing that an improved procedure for selecting practice, proposing changes, and verifying them makes later improvement more effective, including on tasks beyond those used to develop it. A survey of physical RSI reserves its highest level for "independently verified recursive gain" and concludes that "the surveyed literature does not yet demonstrate Physical RSI as defined by our framework" (Zhou et al., 2026a). Demonstrating RSI requires reliable feedback from physical trials, whose success detection is itself error-prone (Navasardyan et al., 2026), and causal attribution of gains to the improved procedure rather than to more attempts.

Dependence among the Directions. The four directions depend on one another. Decisions within deadlines rely on state and feedback that remain trustworthy across tasks and spaces. Reasoning and control reach new environments and bodies only through adaptation, which relies on clear capability interfaces. Continual improvement and recursive self-improvement rely on all three: a system should retain only the changes that verified feedback shows to be better, and a retained change must pay off in the complete system rather than in an isolated component. Conversely, each retained improvement can make decisions faster, state more reliable, and adaptation cheaper. Progress in the other three directions thus accumulates through the fourth.

7. Conclusion

Agentic robotics organizes capabilities around user intent through a closed loop of reasoning, action, verification, and adaptation. Our reference framework distinguishes three complementary functional roles: System 2 (S2) for goal-directed reasoning and decision-making, System 1 (S1) for executable skills, and System 0 (S0) for low-level control. S2 can invoke S1 skills or issue commands directly to

S0. We refer to its realization through a foundation model and an embodied harness as *Embodied S2*, the decision-making core that connects general-purpose intelligence with embodied execution.

Four ideas underpin this perspective. First, foundation models make S2 a promising starting point for robot generalization. Their broad knowledge and reasoning capabilities can help robots understand unfamiliar goals and situations, organize available capabilities, and adapt through feedback. This offers a complementary path to improving action policies through larger and more diverse training datasets. Second, physical autonomy requires the model and embodied harness to jointly support effective, verifiable action under constraints of time, embodiment, uncertainty, and safety. These constraints must shape how systems reason, execute, and adapt. Third, meaningful evaluation must distinguish model capabilities from closed-loop system performance while examining how those capabilities contribute to reliable outcomes. Finally, the long-term objective is sustained autonomy with progressively broader responsibility, extending from individual primitives and tasks to coordinated workflows and standing roles.

Realizing this ambition requires robots to track goals and progress over time, coordinate interdependent activities, recover from failures, and adapt to changing environments and embodiments. Improvement from physical experience must carry over to subsequent work while preserving existing capabilities. The evidence reviewed here does not yet establish reliable autonomy across these broader and more persistent responsibilities, making sustained performance a central target for both system development and evaluation.

This work calls for three community priorities. First, develop reliable evaluation frameworks that complement simulation with real-robot testing and extend task coverage to higher levels of responsibility and broader spatial scopes. Evaluation should assess sustained performance under changing conditions, including recovery from failure. Second, improve reasoning efficiency for time-sensitive, safety-critical, and highly dynamic settings. When test-time scaling relies on longer deliberation, gains in reasoning quality may come at the cost of timely action. Research should therefore prioritize better decisions within the time available for action, with the longer-term goal of developing embodied-native reasoning that integrates perception, reasoning, and action in response to the dynamics and urgency of physical tasks. Third, standardize reporting of the model, harness, available capabilities, evaluation conditions, inference budgets, response times, and human assistance. Clear reporting supports reproducible comparisons and helps distinguish model capabilities from the support provided by the surrounding system. Together, these priorities can help turn promising results into credible progress toward sustained physical autonomy.

We envision a new era of robotic systems that understand human intent, adapt to changing circumstances, and reliably take on broader responsibilities over longer time horizons. Agentic robotics offers a system-level path toward this vision, bringing the general-purpose capabilities of foundation models into systems that people can trust with sustained work in the physical world.

A. Related Work

This appendix reviews the work on which our reference architecture, evaluation framework, and roadmap build. Each subsection ends with our position on that work.

A.1. Surveys and Position Papers

Foundation Models and Robotics. Bommasani et al. (2021) coined the term foundation models and discussed robotics among their capabilities. Subsequent surveys organized robotic uses of foundation models by component (perception, decision-making, and control) (Firoozi et al., 2025; Hu et al., 2023) and reviewed real-world applications in which foundation models replace specific modules of existing robot systems (Kawaharazuka et al., 2024). Surveys of vision-language-action models place high-level planners inside a VLA taxonomy, distinguishing monolithic from hierarchical designs (Ma et al., 2026; Shao et al., 2025). Others classify robot autonomy driven by language and vision-language models by the model's role (Salimpour et al., 2025), or argue that world models, including models of users, are central to embodied agents (Fung et al., 2025). Our System 1/System 2 vocabulary comes from cognitive science (Kahneman, 2011) and was popularized in machine learning by Bengio (2019).

Definitions Closest to Ours. Several 2026 works define agentic robots by function rather than architecture. Wang et al. (2026b) characterize agentic robotics by "this decision-making role of the FM, rather than any particular model or interface", where FM stands for foundation model. They code 2,343 systems by where foundation-model decisions occur: in task orchestration, motion generation, or improvement. Qin et al. (2026a) define agentic robots as systems in which agents autonomously plan and coordinate task execution using observations and execution feedback, and organize their survey around model, data, environment, and harness. Chen et al. (2026d) characterize embodied agents by five functions (perceive, anticipate, plan, act, verify) and find that only 5 of 64 benchmarks assess anticipation and only 3 assess verification. Zhu and Cai (2026) separate model, harness, and environment and conclude that action-interface expansion is documented more convincingly than robust completion, recovery, authorization, or independent verification. Harness-specific accounts include a survey that defines the embodied agentic harness as the runtime control plane between an agent's proposals and embodied execution (Huang et al., 2026c). They also include a roadmap from world-action models to embodied brains (Liang et al., 2026b) and a proposal that robot middleware is the harness layer (Lee et al., 2026c). Frameworks for self-improvement define ladders of their own, for robots (Zhou et al., 2026a) and for large reasoning models (Yang et al., 2026i). Table 12 relates our definition of Embodied S2 to the closest published definitions.

Our Position. Like these works, we define S2 by function: it interprets intent, decides in light of goals and constraints, orchestrates skills and tools, and verifies outcomes to adapt later decisions. We add a criterion that identifies S2 by its decisions on how to proceed and on outcomes, and we count the embodied harness as part of Embodied S2. Unlike the self-improvement ladders, our four levels measure delegated responsibility (Section 6.1), with stated conditions for each level claim (Section 4.2).

A.2. Language-Model-Based Robot Agents

Early Systems (2022–2025). The basic S2 loop is not new. Early work used language models to decompose tasks into steps (Huang et al., 2022). Later work grounded them in the affordances of learned skills (Ahn et al., 2022), composed them with other pretrained models through language (Zeng et al., 2023), and closed the loop with feedback from success detectors and humans (Huang et al., 2023b). Code became an interface for composing perception and control (Liang et al., 2023;

Table 12 | **Closest definitions in the literature and how ours relates to them.**

Work	What it says	What we adopt	What we add or change
RoboBench (Luo et al., 2026b)	"We refer to System 2 as the embodied brain, the cognitive core for decision-making in manipulation."	S2 as the decision-making core	Embodied S2 includes the harness, not only a multimodal language model; criterion for when a component acts as S2
From WAMs to Embodied Brains (Liang et al., 2026b)	The embodied brain is "a model-level target", while "the harness manages memory, routing, logging, verification, and capability calls".	Separation of model and harness	Harness counted inside Embodied S2; S1 and S0 are tools that S2 calls
Thea (Wang et al., 2026e)	"The world, however, withholds two abilities that software grants for free: reading the state of the world, and judging the outcome of an action."	World, user, and body state built and maintained by the harness	User context as a separate category; the harness, including its outcome judgments, counted as part of Embodied S2
SCOPER (Wei et al., 2026a)	"An embodied brain, or System 2, must turn an abstract goal into actions while discovering relevant scene conditions and checking whether executed skills achieved their intended effects."	Choosing the next skill and judging the last one	Harness treated as part of Embodied S2 rather than an add-on; persistence across episodes
Embodied agentic harnesses survey (Huang et al., 2026c)	Harness as the "runtime control plane" that grounds observations, coordinates capabilities, monitors outcomes, and manages intervention.	Harness responsibilities and contracts	Each ability divided between the model and the harness, each evaluated with the other held fixed
Agentic robotics survey (Wang et al., 2026b)	"The distinguishing feature is this decision-making role of the FM, rather than any particular model or interface." (FM: foundation model)	A decision-based criterion	Adjudication of vendor "S2" labels; two evaluation levels; ladder of delegable work
Towards Agentic Robots (Qin et al., 2026a)	Agents "autonomously plan and coordinate task execution using environmental observations and execution feedback"; System 2 planners, System 1 visuomotor policies, System 0 controllers.	The S2/S1/S0 vocabulary	Harness counted as part of Embodied S2; identifiability requirement
Harness Engineering for Physical AI (Lee et al., 2026c)	Robot middleware is the harness layer, mediating control, computing, and communication.	The harness bounds what the model sees and does	Different scope: our harness assembles context for S2's decisions and carries them out; the middleware described is infrastructure beneath it

Singh et al., 2023). Language-model planners were coupled to classical planners (Liu et al., 2023a) and geometric feasibility checks (Lin et al., 2023), grounded in 3D scene graphs (Rana et al., 2023), taught when to ask for help (Ren et al., 2023), and equipped to explain their failures (Liu et al., 2023d). Language models also wrote value maps and keypoint constraints for motion solvers (Huang et al., 2023a, 2025) and summarized user preferences (Wu et al., 2023). OK-Robot integrated open-vocabulary modules for mobile manipulation in homes (Liu et al., 2024). In Minecraft, a language-model agent followed an automatic curriculum and built a skill library of code (Wang et al., 2024a). In parallel, embodied multimodal models and VLAs moved reasoning into the policy (Driess et al., 2023; Mu et al., 2023; Brohan et al., 2023; Gemini Robotics Team, 2025a).

Agentic Robotic Systems of 2026. Frontier models that act through code and tools broadened the scope of this work. Frameworks study coding agents as robot controllers across levels of abstraction (Fu et al., 2026), let agents write and call their own tools (Ge et al., 2026), and compose primitives and VLA policies in interactive code (Si et al., 2026). HexaShell runs a thin harness with tools written and tested by optimizer agents (Gao et al., 2026b); a later version uses a small set of 14 general tools (Gao et al., 2026a). Others maintain persistent playbooks (Lee et al., 2026b). General-purpose agents drive physical robots directly (Jia et al., 2026) and learn in context from demonstrations (Cheng et al., 2026). One study argues that, given the right interface, coding and computer-use agents are in a sense already robot-control agents (Hu et al., 2026a). Orchestration frameworks connect agents to ROS 2 and to learned skills and policies (Li et al., 2026d; Cardenas et al., 2026; Galanti et al., 2026), and other methods reduce how often the expensive model is consulted (Sun et al., 2026b; Xu et al., 2026a).

Our Position. These systems implement all or part of the S2 loop through layered orchestration, direct control, or code as policy (Figure 4). We identify S2 by its decisions, not its interface, and treat verification, persistent context, and revision as part of S2. From these works we draw comparisons that fix the model or the harness, and we report cost and time semantics alongside success.

A.3. Hierarchical and Dual-System Robot Foundation Models

Many robot foundation models separate high-level reasoning from low-level action. Helix pairs a slow vision-language model with a fast visuomotor policy that receives a latent vector (Figure AI, 2025b), and Helix 02 adds a learned whole-body controller (Figure AI, 2026a). GR00T N1 couples a vision-language module with a diffusion action module trained end to end (NVIDIA et al., 2025c). In Hi Robot, a high-level VLM chooses the next step for a low-level policy (Shi et al., 2025). The π_0 model follows commands from a high-level VLM (Black et al., 2025b), and $\pi_{0.5}$ co-trains subtask prediction with action generation (Black et al., 2025a). Gemini Robotics 1.5 interleaves natural-language reasoning with actions and is paired in an agentic system with an embodied reasoning model, Gemini Robotics-ER 1.5, that acts as orchestrator (Gemini Robotics Team, 2025b). Other dual-system designs use value-guided sampling (Song et al., 2025), combine a generalist with a specialist policy (Bu et al., 2024), or embed the fast system inside the slow one (Chen et al., 2025a). Still others switch adaptively between reasoning and acting (Lin et al., 2025) or communicate through 2D paths (Li et al., 2025b) and editable trajectory traces (Lee et al., 2025). Unified models generate reasoning and actions with one decoder (Liu et al., 2026e) or within one tightly coupled mixture-of-experts model (Zhai et al., 2025). Recent work adds world-model-guided test-time search to subtask generation (Cai et al., 2026) and studies design principles for hierarchical VLAs (Hu et al., 2026b). One study finds that the executor of a hierarchical VLA underuses its plans (Xie et al., 2026).

Our Position. We identify S2 by its decisions, not by its place in a hierarchy: a module acts as S2 only if it makes identifiable decisions on how to proceed and on outcomes. An upper module that

emits only a latent representation may perform this function, but we do not classify it as S2 unless replacement experiments reveal these decisions. The subtask stage of a unified model, by contrast, makes an identifiable decision on how to proceed and so implements part of S2. That executors can underuse their plans suggests that S1's adherence to them should be measured.

A.4. Embodied Reasoning Models

A family of models aims to supply the embodied reasoning that S2 needs. The Gemini Robotics-ER line introduced the ERQA benchmark and has evolved from a reasoning model into an orchestrator of VLAs (Gemini Robotics Team, 2025a,b; Hansen and Xu, 2026). Cosmos-Reason1 trains physical common sense and embodied reasoning with supervised fine-tuning and reinforcement learning (NVIDIA et al., 2025b). Vesta combines a generalist model with a multimodal memory harness (Bjorck et al., 2026). RoboBrain models cover planning, affordance, and trajectory prediction (Ji et al., 2025; BAAI RoboBrain Team et al., 2025; Tan et al., 2026). Embodied-R1 uses pointing as an embodiment-agnostic interface, and Embodied-R1.5 adds a closed planner–grounder–corrector loop (Yuan et al., 2026c,d). Other models vary in scale and training recipe: RynnBrain (Dang et al., 2026; Li et al., 2026b), PhysBrain (DeepCybo Team et al., 2026), MiMo-Embodied (Hao et al., 2025), HY-Embodied (Tencent Robotics X et al., 2026), Robix (Fang et al., 2025), ACE-Brain (ACE-Brain Team et al., 2026), Pelican-VL (Zhang et al., 2025b), and the MolmoER backbone of MolmoAct2 (Fang et al., 2026). Related spatial-grounding models include SpatialVLM, RoboPoint, RoboRefer, and Magma (Chen et al., 2024; Yuan et al., 2025b; Zhou et al., 2025a; Yang et al., 2025a). Vlaser studies how VLM initialization affects VLA fine-tuning (Yang et al., 2026b). Progress estimation can be framed as temporal ordering of shuffled frames (Ma et al., 2025). Recent models also train on data with execution feedback, goal revisions, and failure recovery (Foundation Model, Li Auto Inc, 2026; Li et al., 2026a).

Our Position. Reports on these models mostly evaluate offline question answering, pointing, and spatial skills. We treat such a model as a candidate S2 foundation model whose value for the S2 loop must be shown at the system level.

A.5. Agent Harnesses

Digital Precedents. ReAct interleaved reasoning and acting (Yao et al., 2023), Toolformer taught models to call APIs (Schick et al., 2023), and Reflexion stored verbal feedback in episodic memory instead of updating weights (Shinn et al., 2023). Generative agents introduced memory streams with reflection and retrieval (Park et al., 2023), and CodeAct used executable Python as a unified action space (Wang et al., 2024b). SWE-agent showed that the agent–computer interface affects agent performance (Yang et al., 2024), and OpenHands provided an open sandboxed platform (Wang et al., 2025a). OSWorld benchmarked computer use (Xie et al., 2024). Computer-use models and the Model Context Protocol gave agents general ways to reach software (Anthropic, 2024b,c). As a caution, automatic harness evolution can fail to outperform simple test-time scaling at a matched budget (Wang et al., 2026j).

Physical Harnesses. Thea adapts the coding-agent harness to robots, adding a scene graph as context and exit codes that report action outcomes (Wang et al., 2026e). Runtime governance layers add policy checks, admission, monitoring, rollback, and human override (Qin et al., 2026b). PhyAgentOS separates execution termination from semantic completion (Liu et al., 2026d), and in one study of LLM-controlled devices, scoped capability tokens rejected every capability-escalation call that JSON-schema validation let through (Yang, 2026). Harnesses with execution contracts and verification include DynaHarness, RegenHarness, SCOPER's harness, and a harness for safe robot control (Deng

et al., 2026; Wang et al., 2026c; Wei et al., 2026a; Xu et al., 2026b). Programs serve as evolvable interfaces (Wang et al., 2026f), critics run on separate timescales (Ding et al., 2026), and frozen VLAs become retryable primitives (Zhang et al., 2026f). Navigation and persistence harnesses validate planner proposals (Chen et al., 2026c), support lifelong sessions (Zhao et al., 2026b), and carry user and scene context across sessions (Wang et al., 2026l).

Our Position. We count the harness, that is, the non-model components that assemble the context for S2's decisions, carry them out, and verify their outcomes, as part of Embodied S2 rather than as a layer around the model. Harness cards disclose its configuration so that system-level results can be attributed.

A.6. Memory and World Models

Explicit world memories for robots include open-vocabulary 3D scene graphs (Gu et al., 2024), retrieval-augmented spatiotemporal memory (Anwar et al., 2025), memory snapshots for exploration (Yang et al., 2025d), and combined long- and short-term memories (Wang et al., 2025b). Policies can instead carry memory internally (Shi et al., 2026; Torne et al., 2026). User and long-term context remains less developed: in one study, agents struggled to apply sequential user patterns to planning (Kwon et al., 2026). Recent harnesses keep memories of people and preferences (Jin et al., 2026) and record the outcomes of interactions in the scene graph (Banwasi et al., 2026). In one manufacturing study, an external world-state model improved state tracking (Tang et al., 2026). Predictive world models have progressed from learned simulators (Ha and Schmidhuber, 2018; Hafner et al., 2025) to generative interactive environments (Bruce et al., 2024; Google DeepMind, 2025), physical-AI platforms (NVIDIA et al., 2025a; NVIDIA, 2026), and self-supervised video world models used for planning (Assran et al., 2025). They increasingly enter the decision and evaluation loop, both for imagination-based search (Zhang et al., 2026d) and for evaluating policies (Gemini Robotics Team et al., 2025). Criteria have also been proposed for when their verdicts can count as evidence (Oefinger et al., 2026).

Our Position. We treat cross-session world, user, and body state as a harness responsibility: each item records when and how it was obtained and can be revised. World models supply evidence for S2 but make neither kind of S2 decision, so they do not act as S2.

A.7. Benchmarks

Embodied benchmarks range from instruction following in simulated homes (Shridhar et al., 2020) and suites of manipulation tasks (Liu et al., 2023b; Zhou et al., 2025c; Li et al., 2024a) to real-to-sim evaluation (Li et al., 2025a) and crowdsourced real-world evaluation (Atreya et al., 2025). Benchmarks for embodied agents and S2 include EmbodiedBench, the Embodied Agent Interface, VLABench, RoboCerebra, VSI-Bench, RoboBench, and ASIMOV (Yang et al., 2025c; Li et al., 2024b; Zhang et al., 2025a; Han et al., 2025; Yang et al., 2025b; Luo et al., 2026b; Sermanet et al., 2025). Closed-loop benchmarks from 2026 include RoboDojo for manipulation policies (Chen et al., 2026b) and RobotWorld, SCOPER, LIBERO-Agent, and CodeActionBench for agents (Yang et al., 2026j; Wei et al., 2026a; Diao et al., 2026; Lyu et al., 2026). Embodied First Exam evaluates coding agents across diverse simulated robots and tasks, allowing them to use direct control, generated programs, or trained policies under different protocols (Wang et al., 2026d). Other studies test the reliability of success detectors, runtime monitors, and safety refusals under robot control (Navasardyan et al., 2026; Ahmedt-Aristizabal et al., 2026; Sun et al., 2026a). Section 4 discusses many of these benchmarks.

Our Position. Many benchmarks evaluate the model alone or the whole stack without stating which. Few fix S1 and vary S2 (Han et al., 2025; Wei et al., 2026a), and few score verification directly; exceptions include FailBench (Navasardyan et al., 2026), RESCUE (Ahmedt-Aristizabal et al., 2026), and the completion-claim accuracy of CodeActionBench (Lyu et al., 2026). We therefore evaluate at two levels, model and system, with a harness card for every result.

A.8. Self-Improvement and Continual Learning

Before 2025, robots improved through autonomous data collection and relabeling (Xiao et al., 2023; Bousmalis et al., 2023; Ahn et al., 2024; Zhou et al., 2025d) and through reward functions written by language models (Ma et al., 2024a). Weight-level improvement of S1 policies now uses learned progress signals (Ghasemipour et al., 2025), reinforcement learning from experience (Physical Intelligence et al., 2025), fleet learning during deployment (Wang et al., 2026i), and agentic loops that verify rollouts and refine policies (Liu et al., 2026a; Xiao et al., 2026). Orchestrators and policies can co-evolve (Zhang et al., 2026e). Around frozen foundation models, systems revise tools, skills, and harnesses (Tian et al., 2026; Wang et al., 2026h; Lu et al., 2026; Elmaaroufi et al., 2026; Lee et al., 2026a), and some co-evolve policy, curriculum, and judge (Zhou et al., 2026b). Some results are cautionary. In one simulation study, accepted changes kept passing their tests while the target task never succeeded (Wang, 2026). In another, accepting every proposed revision degraded a harness, whereas a champion–challenger rule improved it (Wei et al., 2026b). In a third, continual-learning gains did not grow with more samples (Huang et al., 2026a).

Our Position. Whereas S2 adapts to verified outcomes within a task, a claim of improvement asserts that a retained change helps later work. Such a claim should pass the three-part test of Section 6.2.4: gains on held-out tasks, no regression, and separate contributions of the S2 foundation model, the harness, and S1. This test is stricter than the standard of most current reports but weaker than the recursive-gain criterion of physical recursive self-improvement (Zhou et al., 2026a).

B. Applications: What Agentic Robotics Changes

Embodied S2 could expand the range of work that robots undertake and improve how reliably and efficiently they complete it. For each of several representative domains, we ask what earlier approaches already deliver, where they stop, what Embodied S2 could add, and how far reported deployments reach.

Two earlier approaches set the baseline. *Programmed automation*, from industrial arms to warehouse systems, executes engineered motions and rules with high speed and precision, but only where the environment and the task are fixed in advance. Each new product, layout, or exception requires engineering effort. *Learned policies*, including VLA models, generalize across objects and scenes within the skills on which they were trained, and now complete whole tasks in homes they have never seen (Black et al., 2025a; Figure AI, 2026b). Each new task, however, still requires demonstrations, and a policy does not by itself decide which task to do next, check that the result is correct, or recover when it is not. In reported deployments, these responsibilities fall to remote experts and operators: *human backup performs the role of S2*, a pattern that recurs across the domains below.

If Embodied S2 took over these responsibilities, three changes would follow. First, a robot could be given new work by instruction rather than by reprogramming or new demonstrations, which would make automation worthwhile for varied, low-volume tasks that never justified the engineering. Second, a robot could verify its own results and handle routine exceptions, so that one person could supervise many robots by exception instead of guiding one. Third, a robot could carry state across

tasks and days, a prerequisite for entrusting it with a workflow or a role rather than a single task. We state the reach of each domain on the ladder of delegable work (Section 6.1; Figure 1), from a single primitive (L1) through a task (L2) and a workflow (L3) to a standing role (L4). Because people still fill much of the role of S2, human interventions per task, or per shift for continuous operations, should become a standard deployment metric alongside reliability and execution efficiency. Table 13 summarizes the reports and our assessment for each domain.

B.1. Homes

Homes are the most numerous and the least structured of these settings. Single-purpose appliances such as robot vacuum cleaners work reliably because their task is fixed. General-purpose home robots are now emerging. Figure AI reports 56% whole-task success in 30 unseen homes in a company-run evaluation that counts any rollout requiring a safety intervention as a failure (Figure AI, 2026b). Consumer humanoids are offered for sale, although their shipping volumes have not been disclosed. According to its manufacturer, the 1X NEO is sold with remote supervision by a company expert for chores that the robot cannot yet perform alone (1X Technologies, 2026). An assessment by MIT Technology Review noted that the robot still needs a remote operator for most tasks (Condliffe, 2026). Reported home autonomy thus reaches L2, the task level, and remote operation bridges the rest.

Much of what the remote operator supplies is the work of S2. A request such as "get the living room ready for guests, but leave the children's puzzle on the table" leaves most objects, all steps, and the completion criterion unstated. Fulfilling it requires interpreting what the household wants, deciding where each item belongs, sequencing the work, checking the result, and responding to interruptions. Throughout, the robot must remember the home, its occupants, and their preferences while protecting their privacy. An Embodied S2 that handled such requests would change what a home robot is. It would no longer be an appliance for one chore or a teleoperated body, but a helper to which a household delegates an open list of chores. Remote experts would then approve exceptions rather than perform the work.

B.2. Manufacturing and Logistics

Factories and warehouses are where robots already work at large scale (International Federation of Robotics, 2026; Dresser, 2025), and where the limits of programmed automation are most visible. Industrial arms repeat engineered motions with high speed and precision inside fixtures and cells. This approach pays off in high-volume production but turns every change of product or layout into an engineering project. Humanoids and learned policies are now entering these settings. Figure AI reports that its Figure 02 deployment at BMW ran ten-hour shifts and contributed to the production of more than 30,000 vehicles (Figure AI, 2025a). Apptronik agreed to pilot its Apollo humanoid at Mercedes-Benz, runs data collection at customer sites including Mercedes-Benz, and uses the resulting fleet data to train Gemini Robotics models (Apptronik, 2024, 2026; Gemini Robotics Team, 2025a). Hyundai has announced plans to deploy Boston Dynamics' Atlas in its plants (Boston Dynamics, 2026b). Skild AI reports more than 60 paying customers ten months after its first commercial deployment; mobility accounts for 10% of its revenue, and the remaining 90% comes mainly from manipulation tasks (Pathak and Gupta, 2026). In logistics, Agility reports that its Digit robots have moved more than 100,000 totes at GXO (Agility Robotics, 2025; GXO Logistics and Agility Robotics, 2026). Amazon reports operating more than a million robots coordinated by what it describes as a fleet foundation model (Dresser, 2025), and its tactile picking robot asks human workers to take over items that it cannot handle (Davies, 2025). All of these figures are company-reported.

Because each cycle in these settings is short and repetitive, Embodied S2 would contribute less

Table 13 | **Application domains: what earlier approaches deliver, what Embodied S2 could add, and the reach reported so far.** Levels of delegable work are L1 primitive, L2 task, L3 workflow, and L4 role; the level needed and the contributions of S2 are our assessment. A reported duration of operation is not a level reached: no report fixes in advance all the conditions that a level claim requires (Section 4.2). Multi-day laboratory runs show sustained operation in a restricted domain rather than a role held across open-ended work.

Domain	Earlier approaches and where they stop	What Embodied S2 could add	Reported reach and evidence; level needed
Home	Single-purpose appliances for fixed chores; learned policies complete single tasks in unseen homes; remote experts bridge longer chores	Interpreting open requests; sequencing chores; verification and recovery; user and home memory; interruption; privacy	Task; company-run evaluation in unseen homes; remote experts disclosed for some products. Needed: L3 to L4
Manufacturing	Programmed cells for fixed products, with every change an engineering project; humanoid pilots on short tasks	Per-cycle verification; exception handling; re-tasking from work instructions	Primitive to task, within shifts of up to ten hours; company reports; intervention counts not disclosed. Needed: L3 to L4
Logistics	Large fleets in structured warehouses; items that the picking robot cannot handle go to people	Exception handling; fleet orchestration	Primitive level per robot; fleet operation in structured warehouses; company reports; human takeover disclosed for picking. Needed: L3 to L4
Inspection	Preplanned routes; recorded data interpreted by people	Verified reading of physical state; follow-up on anomalies; persistent site model	Task to workflow per inspection round, in structured sites; company reports. Needed: L4
Healthcare and eldercare	Delivery between fixed points; eldercare only in trials	Varied requests; timing of intervention; long-term memory; safety	Task for hospital delivery (company report); eldercare only in trials. Needed: L4 for eldercare
Laboratories	Fixed protocols on automated equipment; autonomous runs in restricted domains	Choosing the next experiment; outcome verification; provenance	Sustained operation for 8 to 17 days in structured laboratories (peer-reviewed). Needed: L4
Agriculture and mobility	Machines that each perform one job; vehicles in restricted operating domains	Switching between jobs from observations; safety case; real-time S0 and S1	Primitive (agriculture; manufacturer reports); task to workflow per trip, in restricted domains (mobility; company reports). Needed: L3 to L4, in wider domains

within a cycle than around it. It could verify each cycle and handle exceptions such as a jam or an unfamiliar item, diagnosing the cause and then either retrying differently or handing over with a precise report. It could also orchestrate fleets and re-task robots when production changes. Re-tasking may bring the largest gain. An Embodied S2 that composed existing skills into a new sequence from a work instruction, and verified the first parts before running at volume, would turn a reprogramming project into a configuration step. Automation could then extend from high-volume lines to high-mix, low-volume work that fixed automation has rarely reached.

B.3. Inspection

Inspection is an early commercial application of embodied reasoning models, which are candidate S2 foundation models. It suits current models because it tolerates latency and rewards verified reading of physical state, and it illustrates the step from collecting data to making judgments. A conventional inspection robot repeats a preplanned route and records data for people to interpret. Boston Dynamics reports more than 1,500 Spot robots with customers (Boston Dynamics, 2026c) and has integrated Gemini Robotics-ER into its inspection software for reading gauges and interpreting site conditions (Graesser and Xu, 2026; Boston Dynamics, 2026a). Google reports a demonstration in which Gemini 4 Argon controls a real Spot after practicing the task in a simulated twin that it built itself (Rao, 2026). An Embodied S2 could connect reading and control by acting on what it reads. When a gauge is out of range, it could look again from another angle, check the related valve, and file a report with the evidence attached, instead of leaving every anomaly for a person to find in the recordings. At the role level, inspection would also require a persistent model of the site that is updated on every round.

B.4. Healthcare and Laboratories

Hospital delivery robots already navigate elevators and corridors autonomously (Diligent Robotics, 2024), but they carry items between fixed points. An Embodied S2 could turn such a robot into an assistant for varied requests, such as fetching a particular item from a supply room. Eldercare is the most demanding domain considered here. Delegating it to a robot would require responsibility at the role level (L4), long-term memory of a person, well-timed intervention, and the strongest privacy and safety protections. Current evidence is limited to trials.

Laboratory automation has long executed fixed protocols, and structured laboratories also host autonomous research campaigns that run continuously for days. The A-Lab synthesized 36 of 57 target compounds over 17 days of continuous operation (Szymanski et al., 2023), although the novelty of its products and the automated analysis that identified them were later disputed; the count given here is the corrected one (Leeman et al., 2024; Szymanski et al., 2026). A mobile robotic chemist ran 688 experiments over eight days (Burger et al., 2020). These runs show sustained operation in a restricted domain rather than a role held across open-ended work. A language-model agent, Coscientist, has also planned and executed chemistry experiments (Boiko et al., 2023). In open-ended work, an Embodied S2 could decide which experiment to run next given the results so far, verify that each step produced what it should, and record provenance. In laboratories, verification of outcomes, rather than manipulation, is likely to be the bottleneck.

B.5. Agriculture and Mobility

Agricultural robots operate commercially at the primitive level (L1): laser weeders work in 15 countries according to their manufacturer, which reports more than 200 growers across its laser-weeding and tractor-autonomy products (Carbon Robotics, 2026), but each machine performs one job. An Embodied S2 could let one platform scout a field, decide where weeding or harvesting is

needed, and switch between these jobs as conditions change. In mobility, robotaxis provide hundreds of thousands of paid rides per week in restricted operating domains (Korosec, 2026). Cars and drones are physical agents as well, and the same S2–S1–S0 organization applies to them. Their history also suggests how agentic robotics is likely to scale: domain by domain, with each operating domain widened only as evidence of safety accumulates.

Common Patterns. Across these domains, evidence of unaided work in unstructured settings is strongest at L2. Some demonstrations extend toward workflows, but their reports do not yet establish all the conditions needed for a reliable L3 claim. Longer operation has been reported only in structured or restricted domains, and wherever the fallback is disclosed, it is a person. The same three changes recur: work assigned by instruction rather than by engineering or new data, supervision by exception rather than one operator per robot, and responsibility extended from single tasks to workflows and roles. Each change would widen the range of work for which deploying a robot is worthwhile. The research directions in Section 6.2 address what delivering these changes would take.

C. Contributors

Qingfeng Lan*	Duo Li	Siyao Li	Yiheng Lyu
Shijian Lu	Zhenghao Qi*	Shizun Wang*	Xinchao Wang
Ming Yang*	Zhiqin Yang*	Zuhao Yang*	Tao Yu*

Contributors are listed in alphabetical order by surname, then by first name within each surname, read from left to right across each row. * denotes core contributors.

References

1X Technologies. NEO Home Robot: Product Page. https://www.1x.tech/neo, 2026.

ACE-Brain Team, Ziyang Gong, Haoming Gu, Zehang Luo, Tianyi Zhang, Tao Tao, Yixiao Chi, Zhe Liu, Lingsi Zhu, et al. ACE-Brain-0.5: A Unified Embodied Foundational Model for Physical Agentic AI, 2026.

Agility Robotics. Digit Moves Over 100,000 Totes in Commercial Deployment. https://www.agilityrobotics.com/content/digit-moves-over-100k-totes, November 2025.

David Ahmedt-Aristizabal, Abdelwahed Khamis, Mohammad Ali Armin, and Lars Petersson. RESCUE: Is This Alarm Worth Acting On? Counterfactual Evaluation of VLA Runtime Monitors, 2026.

Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances, 2022.

Michael Ahn, Debidatta Dwibedi, Chelsea Finn, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Karol Hausman, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, et al. AutoRT: Embodied Foundation Models for Large Scale Orchestration of Robotic Agents, 2024.

Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-Language navigation: Interpreting visually-grounded navigation instructions in real environments. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

Anthropic. Building Effective Agents. https://www.anthropic.com/engineering/building-effective-agents, 2024a.

Anthropic. Introducing Computer Use, a New Claude 3.5 Sonnet, and Claude 3.5 Haiku. https://www.anthropic.com/news/3-5-models-and-computer-use, 2024b.

Anthropic. Introducing the Model Context Protocol. https://www.anthropic.com/news/model-context-protocol, 2024c.

Anthropic. Effective Context Engineering for AI Agents. https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents, 2025a.

Anthropic. Effective Harnesses for Long-Running Agents. https://www.anthropic.com/engineering/effective-harnesses-for-long-running-agents, 2025b.

Anthropic. Equipping Agents for the Real World with Agent Skills. https://www.anthropic.com/engineering/equipping-agents-for-the-real-world-with-agent-skills, 2025c.

Anthropic. Writing Effective Tools for Agents—with Agents. https://www.anthropic.com/engineering/writing-tools-for-agents, 2025d.

Anthropic. Harness Design for Long-Running Application Development. https://www.anthropic.com/engineering/harness-design-long-running-apps, 2026a.

Anthropic. Scaling Managed Agents: Decoupling the Brain from the Hands. https://www.anthropic.com/engineering/managed-agents, 2026b.

Abrar Anwar, John Welsh, Joydeep Biswas, Soha Pouya, and Yan Chang. ReMEmbR: Building and Reasoning Over Long-Horizon Spatio-Temporal Memory for Robot Navigation. In *International Conference on Robotics and Automation*, 2025.

Apptronik. Apptronik and Mercedes-Benz Enter Commercial Agreement That Will Pilot Apptronik's Apollo Humanoid Robot in Mercedes-Benz Manufacturing Facilities. https://apptronik.com/news-collection/apptronik-and-mercedes-benz-enter-commercial-agreement, March 2024.

Apptronik. Welcome to Robot Park: Where Apptronik's Apollo Goes to Work. https://apptronik.com/news-collection/welcome-to-robot-park-where-apptroniks-apollo-goes-to-work, June 2026.

Artificial Analysis. Welcome to the Era of Embodied AI: Introducing AA-Robotics v0.1. https://artificialanalysis.ai/articles/aa-robotics, 2026.

Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Mojtaba Komeili, Matthew Muckley, Ammar Rizvi, Claire Roberts, et al. V-JEPA 2: Self-Supervised Video Models Enable Understanding, Prediction and Planning, 2025.

Pranav Atreya, Karl Pertsch, Tony Lee, Moo Jin Kim, Arhan Jain, Artur Kuramshin, Cyrus Neary, Edward S. Hu, Kanav Arora, Kirsty Ellis, Luca Macesanu, Matthew Leonard, Meedeum Cho, Ozgur Aslan, Shivin Dass, Tony Wang, Xingfang Yuan, Abhishek Gupta, Dinesh Jayaraman, Glen Berseth, Kostas Daniilidis, Roberto Martín-Martín, Youngwoon Lee, Percy Liang, Chelsea Finn, and Sergey Levine. RoboArena: Distributed Real-World Evaluation of Generalist Robot Policies. In *Proceedings of The 9th Conference on Robot Learning*, 2025.

BAAI RoboBrain Team, Mingyu Cao, Huajie Tan, Yuheng Ji, Xiansheng Chen, Minglan Lin, Zhiyu Li, Zhou Cao, Pengwei Wang, Enshen Zhou, et al. RoboBrain 2.0 Technical Report, 2025.

Jingwen Bai, Guochao Bai, Canbo Zhang, Shuangqi Li, Minchao Hu, Weichun Huang, and Shuang Liu. Compile Once, Run Cheap: Frozen, Portable Skills for Embodied Agents, and What Recursive Self-Improvement Adds, 2026.

Anuva Banwasi, William Muckelroy, III, Priya Sundaresan, Linfeng Zhao, Jeannette Bohg, and Cherie Ho. MessyMem: Learning-from-Doing Memory for Mobile Manipulation. Conference on Robot Learning (CoRL 2026); arXiv:2609.15976, 2026. URL https://arxiv.org/abs/2609.15976.

Nicolas Baumann, Liam Boyle, Pu Deng, Edoardo Ghignone, Boyang Sun, Marc Pollefeys, Luca Benini, and Michele Magno. LocalNav: Distilling Frontier VLMs and Embodied RL for On-Device Object Goal Navigation, 2026.

Yoshua Bengio. From System 1 Deep Learning to System 2 Deep Learning. Invited talk, Conference on Neural Information Processing Systems, 2019. URL https://neurips.cc/virtual/2019/invited-talk/15488.

Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. PIQA: Reasoning about Physical Commonsense in Natural Language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.

Johan Bjorck, Zhiqi Li, Yunze Man, Jing Wang, An-Chieh Cheng, Sifei Liu, Shihao Wang, Zhiding Yu, Abhishek Badki, Stan Birchfield, et al. Vesta: A Generalist Embodied Reasoning Model, 2026.

Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: a Vision-Language-Action Model with Open-World Generalization. In *Proceedings of The 9th Conference on Robot Learning*, 2025a.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A Vision-Language-Action Flow Model for General Robot Control. In *Robotics: Science and Systems XXI (RSS)*, Los Angeles, CA, USA, June 2025b.

Kevin Black, Manuel Y. Galliker, and Sergey Levine. Real-Time Execution of Action Chunking Flow Policies. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*, 2025c.

Daniil A. Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous Chemical Research with Large Language Models. *Nature*, 2023.

Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the Opportunities and Risks of Foundation Models, 2021.

Boston Dynamics. AIVI-Learning Is Now Powered by Google Gemini Robotics. https://bostondynamics.com/blog/aivi-learning-now-powered-google-gemini-robotics/, April 2026a.

Boston Dynamics. Boston Dynamics Opens Robotics Metaplant Application Center to Train Humanoid Robots for Manufacturing Tasks. https://bostondynamics.com/news/boston-dynamics-opens-robotics-metaplant-application-center-to-train-humanoid-robots-for-manufacturing-tasks/, September 2026b.

Boston Dynamics. Spot: Agile Mobile Robot for Dynamic Sensing and Industrial Inspection. https://bostondynamics.com/products/spot/, 2026c.

Konstantinos Bousmalis, Giulia Vezzani, Dushyant Rao, Coline Devin, Alex X. Lee, Maria Bauza, Todor Davchev, Yuxiang Zhou, Agrim Gupta, Akhil Raju, et al. RoboCat: A Self-Improving Generalist Agent for Robotic Manipulation. *Transactions on Machine Learning Research*, 2023.

Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control, 2023.

Jake Bruce, Michael Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative Interactive Environments. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024.

Qingwen Bu, Hongyang Li, Li Chen, Jisong Cai, Jia Zeng, Heming Cui, Maoqing Yao, and Yu Qiao. Towards Synergistic, Generalized, and Efficient Dual-System for Robotic Manipulation, 2024.

Benjamin Burger, Phillip M. Maffettone, Vladimir V. Gusev, Catherine M. Aitchison, Yang Bai, Xiaoyan Wang, Xiaobo Li, Ben M. Alston, Buyi Li, Rob Clowes, et al. A Mobile Robotic Chemist. *Nature*, 2020.

Xiaowei Cai, Yunuo Cai, Bingao Chen, Jingxiao Chen, Zhi Chen, Siyuan Feng, Tengyu Hou, Jingshun Huang, Han Jiang, Runkun Ju, et al. τ_0-VLA: a Hierarchical Robot Foundation Model with World-Model-Guided Test-Time Computation, 2026.

Carbon Robotics. Carbon Robotics: Leader in Physical AI for Agriculture. https://carbonrobotics.com/, 2026.

Irvin Steve Cardenas, Marcus Anthony Arnett, Natalie Catherine Yeo, Lucky Sah, and Jong-Hoon Kim. ROSClaw: An OpenClaw ROS 2 Framework for Agentic Robot Control and Interaction, 2026.

Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, et al. SAM 3: Segment Anything with Concepts, 2025.

Matthew Chang, Gunjan Chhablani, Alexander Clegg, Mikael Dallaire Cote, Ruta Desai, Michal Hlavac, Vladimir Karashchuk, Jacob Krantz, Roozbeh Mottaghi, Priyam Parashar, et al. PARTNR: A Benchmark for Planning and Reasoning in Embodied Multi-agent Tasks. In *International Conference on Learning Representations*, 2025.

Baijun Chen, Weijie Wan, Tianxing Chen, Xianda Guo, Congsheng Xu, Yuanyang Qi, Haojie Zhang, Longyan Wu, Tianling Xu, Zixuan Li, et al. UniVTAC: A Unified Simulation Platform for Visuo-Tactile Manipulation Data Generation, Learning, and Benchmarking, 2026a.

Boyuan Chen, Zhuo Xu, Sean Kirmani, Brian Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. SpatialVLM: Endowing Vision-Language Models with Spatial Reasoning Capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Hao Chen, Jiaming Liu, Chenyang Gu, Zhuoyang Liu, Renrui Zhang, Xiaoqi Li, Xiao He, Yandong Guo, Chi-Wing Fu, Shanghang Zhang, et al. Fast-in-Slow: A Dual-System Foundation Model Unifying Fast Manipulation within Slow Reasoning, 2025a.

Junting Chen, Checheng Yu, Xunzhe Zhou, Tianqi Xu, Yao Mu, Mengkang Hu, Wenqi Shao, Yikai Wang, Guohao Li, and Lin Shao. EMOS: Embodiment-aware Heterogeneous Multi-robot Operating System with LLM Agents. In *International Conference on Learning Representations (ICLR)*, 2025b.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating Large Language Models Trained on Code, 2021.

Ruolin Chen, Yinqian Sun, Jihang Wang, Mingyang Lv, Qian Zhang, and Yi Zeng. SafeMind: Benchmarking and Mitigating Safety Risks in Embodied LLM Agents, 2025c.

Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, et al. RoboTwin 2.0: A Scalable Data Generator and Benchmark with Strong Domain Randomization for Robust Bimanual Robotic Manipulation, 2025d.

Tianxing Chen, Yue Chen, Zixuan Li, Junyuan Tang, Kailun Su, Haoran Lu, Weijie Wan, Baijun Chen, Songling Liu, Haowen Yan, et al. RoboDojo: A Unified Sim-and-Real Benchmark for Comprehensive Evaluation of Generalist Robot Manipulation Policies, 2026b.

Yang Chen, Lirong Che, Zhenyu Huang, Wenbo Fu, Chuang Wang, Xu Cao, Daqi Liu, Yuzhe Yang, Jian Su, and Lan-Zhe Guo. HarnessVLN: Unifying Training-Free Embodied Navigation through an Agent Harness, 2026c.

Yanzhe Chen, Ziyi Yang, Jifeng Zhu, Qiming Huang, Ruihe An, Peiyao Xu, Hesen Yang, Runda Liu, Chang Gong, Zhijun Cao, et al. Multimodal Embodied Agents: A Unified Capability-centric Perspective from Computer-Use to Robot-Use, 2026d.

Zhekai Chen, Chengqi Duan, Kaiyue Sun, Bohao Li, Yuqing Wang, Manyuan Zhang, and Xihui Liu. UniClawBench: A Universal Benchmark for Proactive Agents on Real-World Tasks, 2026e.

Dongzhou Cheng, Taoran Yi, Ye Fang, Xingwu Zhang, Fan Feng, Yixuan Li, Gengxiong Zhuang, Rongze Wang, Shuai Yang, Wei Song, et al. In-Context Robot Learning with VLM Agents, 2026.

Egor Cherepanov, Nikita Kachaev, Alexey K. Kovalev, and Aleksandr I. Panov. Memory, Benchmark & Robots: A Benchmark for Solving Complex Tasks with Reinforcement Learning. In *International Conference on Learning Representations (ICLR)*, 2026a.

Egor Cherepanov, Nikita Kachaev, Aleksandr Panov, and Alexey Kovalev. MIKASA-robo-VLA: Benchmarking memory in VLA models for long-horizon manipulation. In *RoboPAD: Post-Training Adaptation of Robot Foundation Models*, 2026b.

Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots. In *Robotics: Science and Systems (RSS)*, 2024.

Wei Chow, Jiageng Mao, Boyi Li, Daniel Seita, Vitor Guizilini, and Yue Wang. PhysBench: Benchmarking and Enhancing Vision-Language Models for Physical World Understanding. In *International Conference on Learning Representations (ICLR)*, 2025.

Michele Colledanchise and Petter Ögren. *Behavior Trees in Robotics and AI: An Introduction*. Chapman & Hall/CRC Artificial Intelligence and Robotics Series. CRC Press, 2018.

Jamie Condliffe. AI Breakthroughs in Robotics Won't Change Your Life Any Time Soon. https://www.technologyreview.com/2026/10/08/1145923/ai-breakthroughs-in-robotics-wont-change-your-life-any-time-soon/, October 2026.

Andrea Yaoyun Cui, Hanyu Wang, Xinyi Wang, Bangzheng Li, Lin Zhang, Zhenggang Tang, Shyamal Buch, Dejia Xu, and Matthew Tancik. LIBERO-RIG: A Hardware-Setup robustness benchmark for Vision-Language-Action models, 2026.

Yinpei Dai, Hongze Fu, Jayjun Lee, Yuejiang Liu, Haoran Zhang, Jianing Yang, Chelsea Finn, Nima Fazeli, and Joyce Chai. RoboMME: Benchmarking and Understanding Memory for Robotic Generalist Policies. In *International Conference on Machine Learning (ICML)*, 2026.

Ronghao Dang, Jiayan Guo, Bohan Hou, Sicong Leng, Kehan Li, Xin Li, Jiangpin Liu, Yunxuan Mao, Zhikai Wang, Yuqian Yuan, et al. RynnBrain: Open Embodied Foundation Models, 2026.

Jadelynn Dao, Milan Ganai, Yasmina Abukhadra, Ajay Sridhar, Mozhgan Nasr Azadani, Katie Luo, Clark Barrett, Jiajun Wu, Chelsea Finn, and Marco Pavone. DIRECT: When and Where Should You Allocate Test-Time Compute in Embodied Planners? *Conference on Robot Learning*, 2026.

Alex Davies. Introducing Vulcan: Amazon's First Robot with a Sense of Touch. https://www.aboutamazon.com/news/operations/amazon-vulcan-robot-pick-stow-touch, May 2025.

DeepCybo Team, Yu Bin, Haipeng Cao, Zheng Chang, Kai Chen, Youning Chen, Kailin Deng, Yichao Du, Xiaotong Fu, Haoyang Ge, et al. PhysBrain 1.5: From Vision-Language Models to Physical Foundation Models, 2026.

Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and PixMo: Open Weights and Open Data for State-of-the-Art Vision-Language Models, 2024.

Haoyuan Deng, Jiebin Liu, Tengxiao Zhang, Langning Yan, Hongye Cao, and Ziwei Wang. DynaHarness: A Dynamic Physical Harness for Self-Evolving Robot Agents, 2026.

Zijie Diao, Yitong Chen, Sicheng Xie, Tianyi Lu, Wujian Peng, Guojin Zhong, Houze Xu, Ziyi Ye, Zuxuan Wu, and Yu-Gang Jiang. LIBERO-Agent: Evaluating General-Purpose Agents for Direct Embodied Manipulation, 2026.

Diligent Robotics. Moxi Reaches Milestone of 100,000 Autonomous Elevator Rides in Hospitals. https://www.diligentrobots.com/press/moxi-reaches-milestone-of-100000-autonomous-elevator-rides-in-hospitals, November 2024.

Xin Ding, Liang Mi, Mingzhe Huang, Zixuan Wang, Chao Zhang, Zixu Hao, Fu Chen, Xiangyu Li, Yikai Zheng, Yaoyu Guo, et al. Zetta ζ: An Efficient Closed-Loop Embodied Harness for Self-Evolving Physical Intelligence, 2026.

Scott Dresser. Amazon Launches a New AI Foundation Model to Power Its Robotic Fleet and Deploys Its 1 Millionth Robot. https://www.aboutamazon.com/news/operations/amazon-million-robots-ai-foundation-model, July 2025.

Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. PaLM-E: An Embodied Multimodal Language Model. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 2023.

Dyna Robotics. Dyna-2: A 1-Million-Hour Scaling Law for World-Action Models. https://www.dyna.co/dyna-2, August 2026a.

Dyna Robotics. Dyna-2.1: A Physical Agent for End-to-End Workflows. https://www.dyna.co/dyna-2.1, September 2026b.

Karim Elmaaroufi, Justin Svegliato, Sarunas Kalade, Graham Schelle, Sanjit A. Seshia, and Matei Zaharia. RHO: Your Coding Agent is Secretly a Roboticist, 2026.

Janice J. Eng, David A. Winter, and Aftab E. Patla. Strategies for recovery from a trip in early and late swing during human walking. *Experimental Brain Research*, 1994.

Haoquan Fang, Jiafei Duan, Donovan Clay, Sam Wang, Shuo Liu, Weikai Huang, Xiang Fan, Wei-Chuan Tsai, Shirui Chen, Yi Ru Wang, et al. MolmoAct2: Action Reasoning Models for Real-world Deployment, 2026.

Huang Fang, Mengxi Zhang, Heng Dong, Wei Li, Zixuan Wang, Qifeng Zhang, Xueyun Tian, Yucheng Hu, and Hang Li. Robix: A Unified Model for Robot Interaction, Reasoning and Planning, 2025.

Figure AI. F.02 Contributed to the Production of 30,000 Cars at BMW. https://www.figure.ai/news/production-at-bmw, November 2025a.

Figure AI. Helix: A Vision-Language-Action Model for Generalist Humanoid Control. https://www.figure.ai/news/helix, February 2025b.

Figure AI. Introducing Helix 02: Full-Body Autonomy. https://www.figure.ai/news/helix-02, January 2026a.

Figure AI. Helix 2.5: Zero-Shot 30-Home Generalization. https://www.figure.ai/news/helix-2-5-zero-shot-30-home-generalization, September 2026b.

Roya Firoozi, Johnathan Tucker, Stephen Tian, Anirudha Majumdar, Jiankai Sun, Weiyu Liu, Yuke Zhu, Shuran Song, Ashish Kapoor, Karol Hausman, et al. Foundation Models in Robotics: Applications, Challenges, and the Future. *The International Journal of Robotics Research*, 2025.

Foundation Model, Li Auto Inc. ME-VLM: A Unified VLM for Embodied Cognition and Agent Coordination, 2026.

Letian Fu, Justin Yu, Karim El-Refai, Ethan Kou, Haoru Xue, Huang Huang, Wenli Xiao, Guanzhi Wang, Dantong Niu, Fei-Fei Li, et al. CaP-X: A Framework for Benchmarking and Improving Coding Agents for Robot Manipulation, 2026.

Pascale Fung, Yoram Bachrach, Asli Celikyilmaz, Kamalika Chaudhuri, Delong Chen, Willy Chung, Emmanuel Dupoux, Hongyu Gong, Hervé Jégou, Alessandro Lazaric, et al. Embodied AI Agents: Modeling the World, 2025.

Liane Galanti, Dhruv Shah, and Tri Dao. Addressing the Orchestration Gap in Generalist Robots via Physical Agency, 2026.

Galbot Team, Xuchuan Chen, Xiaoqian Cheng, Yu Deng, Lihe Ding, Shaocong Dong, Xiangjun Gao, Haozhe Jia, Zekai Li, Zhoujian Li, et al. Systematically Exploring the Capabilities of GPT-6 Astra as Embodied Policies, 2026.

Hongcheng Gao, Jianshu Zeng, Alex Zhou, Shijia Ge, Xuan Shangguan, Yazhe Wang, Yexing Wan, Zelin Zheng, Jay Zhu, Xiping Li, Di Wu, Yongli Wang, Lingyu He, Zhiqi Jia, Sihang Wu, and Xiao He. HexaShell: A Universal Harness for Robot Coding Agents in Simulation and the Real World, 2026a.

Hongcheng Gao, Alex Zhou, Jay Zhu, Zelin Zheng, Shijia Ge, Xiping Li, Di Wu, Yazhe Wang, Jianshu Zeng, Xuan Shangguan, Yongli Wang, Yexing Wan, Lingyu He, Zhiqi Jia, Sihang Wu, and Xiao He. HexaShell: A Thin Harness for Exploring the Limits of VLM Agents as Robot Coding Agents, 2026b.

Yipeng Gao, Lei Shu, Genzhi Ye, Xi Xiong, Ameesh Makadia, Meiqi Guo, Laurent Itti, and Jindong Chen. 3DCodeBench: Benchmarking agentic procedural 3D modeling via code, 2026c.

Erann Gat, R Peter Bonnasso, Robin Murphy, et al. On three-layer architectures. *Artificial intelligence and mobile robots*, 1998.

Shijia Ge, Alex Zhou, Jianshu Zeng, Yexing Wan, Di Wu, Zelin Zheng, Yazhe Wang, Zhiqi Jia, Xuan Shangguan, Jay Zhu, et al. Make Code as Policy Great Again: Frontier Agents Write, Call, and Evolve Robot Tools, 2026.

Gemini Robotics Team. Gemini Robotics: Bringing AI into the Physical World, 2025a.

Gemini Robotics Team. Gemini Robotics 1.5: Pushing the Frontier of Generalist Robots with Advanced Embodied Reasoning, Thinking, and Motion Transfer, 2025b.

Gemini Robotics Team. Gemini Robotics 2: Safety Evaluations. https://storage.googleapis.com/deepmind-media/gemini-robotics/Gemini-Robotics-2-Safety.pdf, 2026a.

Gemini Robotics Team. Gemini Robotics 2.5 Accelerates Real-World Physical Intelligence, 2026b. URL https://deepmind.google/blog/gemini-robotics-2-5-accelerates-real-world-physical-intelligence/.

Gemini Robotics Team, Krzysztof Choromanski, Coline Devin, Yilun Du, Debidatta Dwibedi, Ruiqi Gao, Abhishek Jindal, Thomas Kipf, Sean Kirmani, Isabel Leal, et al. Evaluating Gemini Robotics Policies in a Veo World Simulator, 2025.

Seyed Kamyar Seyed Ghasemipour, Ayzaan Wahid, Jonathan Tompson, Pannag Sanketi, and Igor Mordatch. Self-Improving Embodied Foundation Models. In *Advances in Neural Information Processing Systems*, 2025.

Google. Robotics with Streaming (Gemini API documentation). https://aistudio.google.com/docs/robotics-streaming, 2026.

Google DeepMind. Genie 3: A New Frontier for World Models. https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/, August 2025.

Keerthana Gopalakrishnan and Nathan Labenz. One Brain, Any Body: Google DeepMind's Keerthana on Gemini Robotics 2, Cross-Embodiment & Humanoids. https://cognitiverevolution.ai/one-brain-any-body-google-deepmind-s-keerthana-on-gemini-robotics-2-cross-embodiment-humanoids/, 2026.

Laura Graesser and Peng Xu. Gemini Robotics-ER 1.6: Powering Real-World Robotics Tasks through Enhanced Embodied Reasoning. https://deepmind.google/blog/gemini-robotics-er-1-6, April 2026.

Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4D: Around the World in 3,000 Hours of Egocentric Video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

Arnav Grover. Refuses When Asked, Complies When Acting: A Runtime Safety Check for Robot Control Loops, 2026.

Qiao Gu, Alihusein Kuwajerwala, Sacha Morin, Krishna Murthy Jatavallabhula, Bipasha Sen, Aditya Agarwal, Corban Rivera, William Paul, Kirsty Ellis, Rama Chellappa, et al. ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

Daya Guo, Dejian Yang, Haowei Zhang, et al. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 2025.

GXO Logistics and Agility Robotics. GXO and Agility Robotics Expand Partnership to Accelerate Deployment of Humanoids in U.S. Operations. https://www.agilityrobotics.com/content/gxo-and-agility-robotics-expand-partnership-to-accelerate-deployment-of-humanoids-in-u-s-operations, October 2026.

David Ha and Jürgen Schmidhuber. World Models, 2018.

Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, 2025.

Songhao Han, Boxiang Qiu, Yue Liao, Siyuan Huang, Chen Gao, Shuicheng Yan, and Si Liu. RoboCerebra: A Large-scale Benchmark for Long-horizon Robotic Manipulation Evaluation. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2025.

Zhi Han, Zihan Guo, Chenxi Zeng, Zicai Cui, Yue Xin, Ming Zhou, Dongzhan Zhou, Weinan Zhang, and Yang Li. Embodied Skill Bench: Benchmarking How Well Agents with Embodied Skills Perform on Physical Tasks, 2026.

Steven Hansen and Peng Xu. Introducing Gemini Robotics ER 2. https://blog.google/innovation-and-ai/models-and-research/google-deepmind/gemini-robotics-er-2/, July 2026.

Xiaoshuai Hao, Lei Zhou, Zhijian Huang, Zhiwen Hou, Yingbo Tang, Lingfeng Zhang, Guang Li, Zheng Lu, Shuhuai Ren, Xianhui Meng, et al. MiMo-Embodied: X-Embodied Foundation Model Technical Report, 2025.

Ryan Hoque, Peide Huang, David J. Yoon, Mouli Sivapurapu, and Jian Zhang. EgoDex: Learning Dexterous Manipulation from Large-Scale Egocentric Video, 2025.

Hengyuan Hu, Jensen Gao, Priya Sundaresan, Satvik Sharma, Jeannette Bohg, and Dorsa Sadigh. VIA: Visual Interface Agent for Robot Control, 2026a.

Jiaheng Hu, Mohit Shridhar, Caden Lu, Dhruv Shah, Hao-Tien Lewis Chiang, Jie Tan, and Annie Xie. What Matters in Orchestrating Robot Policies: A Systematic Study of Hierarchical VLA Agents, 2026b.

Yafei Hu, Quanting Xie, Vidhi Jain, Jonathan Francis, Jay Patrikar, Nikhil Keetha, Seungchan Kim, Yaqi Xie, Tianyi Zhang, Hao-Shu Fang, et al. Toward General-Purpose Robots via Foundation Models: A Survey and Meta-Analysis, 2023.

Jie Huang, Yanan Chen, Ruixun Liu, Kaichen He, Xinyi He, Rui Huang, Zilong Zheng, Zhenliang Zhang, and Yiwu Zhong. EmbodiedCLBench: Evaluating Continual Learning for Self-Evolving Embodied Agents, 2026a.

Jinbang Huang, Zhiyuan Li, Yuanzhao Hu, Ran Qi, Yixin Xiao, Mark Coates, Tongtong Cao, Zhanguang Zhang, and Yingxue Zhang. RoboHarness: Memory-Driven Orchestration of Heterogeneous Robot Policies for Long-Horizon Planning, 2026b.

Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.

Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. VoxPoser: Composable 3D Value Maps for Robotic Manipulation with Language Models. In *Proceedings of the 7th Conference on Robot Learning (CoRL)*, 2023a.

Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Tomas Jackson, Noah Brown, Linda Luu, Sergey Levine, Karol Hausman, and Brian Ichter. Inner Monologue: Embodied Reasoning through Planning with Language Models. In *Proceedings of The 6th Conference on Robot Learning*, 2023b.

Wenlong Huang, Chen Wang, Yunzhu Li, Ruohan Zhang, and Li Fei-Fei. ReKep: Spatio-Temporal Reasoning of Relational Keypoint Constraints for Robotic Manipulation. In *Proceedings of The 8th Conference on Robot Learning*, 2025.

Yu Huang, Yue Chen, Zijiang Yang, and Gary Ding. Embodied Agentic Harnesses for Physical Intelligence: A Survey, September 2026c.

Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 2019.

Julian Ibarz, Jie Tan, Chelsea Finn, Mrinal Kalakrishnan, Peter Pastor, and Sergey Levine. How to Train Your Robot with Deep Reinforcement Learning: Lessons We Have Learned. *The International Journal of Robotics Research*, 2021.

International Federation of Robotics. Five Million Robots now Operate in Factories Globally: World Robotics 2026 Report by IFR Released. https://ifr.org/ifr-press-releases/news/five-million-robots-now-operate-in-factories-globally, September 2026.

Riccardo Andrea Izzo, Cristiano Battistini, Gianluca Bardaro, and Matteo Matteucci. Multimodal Behavior Tree Generation: A Small Vision-Language Model for Robot Task Planning, 2026.

Jack B. Jedlicki, Amish Sethi, and Heng Yang. Can a Frontier VLM Serve as a Robot Policy for Memory-Dependent Manipulation?, 2026.

Yuheng Ji, Huajie Tan, Jiayu Shi, Xiaoshuai Hao, Yuan Zhang, Hengyuan Zhang, Pengwei Wang, Mengdi Zhao, Yao Mu, Pengju An, et al. RoboBrain: A Unified Brain Model for Robotic Manipulation from Abstract to Concrete. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.

Baoxiong Jia, Ting Lei, Song-Chun Zhu, and Siyuan Huang. EgoTaskQA: Understanding Human Tasks in Egocentric Videos. In *Advances in Neural Information Processing Systems (NeurIPS) 35, Datasets and Benchmarks Track*, 2022.

Mengzhao Jia, Yang Lin, Xixin Zhang, Zhihan Zhang, Xiaobai Liu, and Meng Jiang. Agent as Policy for Robotic Manipulation, 2026.

Xiangwei Jiang, Yao Mu, Lixin Duan, and Wen Li. DexJoCo-X: Benchmarking Action Representations for Multi-Hand Dexterous Manipulation, 2026.

Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. SWE-bench: Can Language Models Resolve Real-World GitHub Issues? In *International Conference on Learning Representations*, 2024.

Weiqi Jin, Peijun Tang, Kuncheng Luo, Baifu Huang, Binyan Sun, Haotian Yang, Shangjin Xie, and Jianan Wang. A Glimpse into Long-term Physical Coexistence with Intelligent Robots, 2026.

Abhishek Jindal, Dmitry Kalashnikov, R. Alex Hofer, Oscar Chang, Divya Garikapati, Anirudha Majumdar, Pierre Sermanet, and Vikas Sindhwani. Can AI Perceive Physical Danger and Intervene?, 2025.

R. S. Johansson and G. Westling. Signals in tactile afferents from the fingers eliciting adaptive motor responses during precision grip. *Experimental Brain Research*, 1987.

Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, 2011.

Kento Kawaharazuka, Tatsuya Matsushima, Andrew Gambardella, Jiaxian Guo, Chris Paxton, and Andy Zeng. Real-World Robot Applications of Foundation Models: A Review. *Advanced Robotics*, 2024.

Mukul Khanna, Ram Ramrakhya, Gunjan Chhablani, Sriram Yenamandra, Theophile Gervet, Matthew Chang, Zsolt Kira, Devendra Singh Chaplot, Dhruv Batra, and Roozbeh Mottaghi. GOAT-Bench: A Benchmark for Multi-Modal Lifelong Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. DROID: A Large-Scale In-The-Wild Robot Manipulation Dataset. In *Robotics: Science and Systems XX (RSS)*, 2024.

Chung Min Kim, Brent Yi, Hongsuk Choi, Yi Ma, Ken Goldberg, and Angjoo Kanazawa. PyRoki: A Modular Toolkit for Robot Kinematic Optimization, 2025.

Kirsten Korosec. Waymo Is Scaling Fast: Here's What the Fleet Data Shows. https://techcrunch.com/2026/09/24/waymo-is-scaling-fast-heres-what-the-fleet-data-shows/, September 2026.

Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the Nav-Graph: Vision-and-Language navigation in continuous environments. In *European Conference on Computer Vision (ECCV)*, 2020.

Hadas Kress-Gazit, Kunimatsu Hashimoto, Naveen Kuppuswamy, Paarth Shah, Phoebe Horgan, Gordon Richardson, Siyuan Feng, and Benjamin Burchfiel. Robot Learning as an Empirical Science: Best Practices for Policy Evaluation, 2024.

Thomas Kwa and Vincent Cheng. How Does Time Horizon Vary Across Domains? https://metr.org/blog/2025-07-14-how-does-time-horizon-vary-across-domains/, 2025.

Thomas Kwa, Ben West, Joel Becker, Amy Deng, Katharyn Garcia, Max Hasin, Sami Jawhar, Megan Kinniment, Nate Rush, Sydney Von Arx, et al. Measuring AI Ability to Complete Long Software Tasks. In *Advances in Neural Information Processing Systems*, 2025.

Taeyoon Kwon, Dongwook Choi, Hyojun Kim, Sunghwan Kim, Seungjun Moon, Beong-woo Kwak, Kuan-Hao Huang, and Jinyoung Yeo. Embodied Agents Meet Personalization: Investigating Challenges and Solutions Through the Lens of Memory Utilization. In *International Conference on Learning Representations*, 2026.

Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, et al. Tulu 3: Pushing frontiers in open language model Post-Training, 2024.

Hyun Jung Lee, Jungtaek Kim, Jongwon Jeong, Tae-Eui Kam, Donghyun Kim, and Yong Jae Lee. EMHO: EMbodied Agent Harness Optimization via Experience Traces, 2026a.

Jason Lee, Jiafei Duan, Haoquan Fang, Yuquan Deng, Shuo Liu, Boyang Li, Bohan Fang, Jieyu Zhang, Yi Ru Wang, Sangho Lee, et al. MolmoAct: Action Reasoning Models that can Reason in Space, 2025.

Junhoo Lee, Injun Baek, Seungyeon Kim, Suhyun Jeon, Minkyu Kim, Baekseung Kim, and Nojun Kwak. RobotUse: Allocating Computation, Context, and Decisions, 2026b.

Sanghoon Lee, Jiyeong Chae, and Kyung-Joon Park. Harness Engineering for Physical AI: Robot Middleware Is the Harness Layer, 2026c.

Josh Leeman, Yuhan Liu, Joseph Stiles, Scott B. Lee, Prajna Bhatt, Leslie M. Schoop, and Robert G. Palgrave. Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis. *PRX Energy*, 2024.

Huashuo Lei, Wenxuan Song, Huarui Zhang, Jieyuan Pei, Jiayi Chen, Haodong Yan, Han Zhao, Pengxiang Ding, Zhipeng Zhang, Lida Huang, et al. RoboMemArena: A Comprehensive and Challenging Robotic Memory Benchmark, 2026.

Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Wensi Ai, Benjamin Martinez, et al. BEHAVIOR-1K: A Human-Centered, Embodied AI Benchmark with 1,000 Everyday Activities and Realistic Simulation, 2024a.

Gangyi Li, LuYao Wang, Jingfu Sun, Jingyi Yuan, Hang Liu, Hao Zong, Yujia Hu, Tingting Gao, Jianyong Wang, and Xu Yang. JoyInside-VLM-Agent: Product-Conditioned Multimodal Policy Learning for Embodied Agents, 2026a.

Kehan Li, Bohan Hou, Minghao Zhu, Tianyi Zhang, Zesen Cheng, Zhikai Wang, Sicong Leng, Xin Li, Xiao Lin, Biying Yao, et al. RynnBrain 1.1: Towards More Capable and Generalizable Embodied Foundation Model, 2026b.

Lin Li, Qihang Zhang, Yiming Luo, Shuai Yang, Ruilin Wang, Fei Han, Mingrui Yu, Zelin Gao, Nan Xue, Xing Zhu, et al. Causal World Modeling for Robot Control, 2026c.

Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen, Tony Lee, Li Erran Li, Ruohan Zhang, et al. Embodied Agent Interface: Benchmarking LLMs for Embodied Decision Making. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2024b.

Ruiying Li, Yunlang Zhou, YuYao Zhu, Kylin Chen, Jingyuan Wang, Sukai Wang, Kongtao Hu, Minhui Yu, Bowen Jiang, Zhan Su, et al. RoboClaw: An Agentic Framework for Scalable Long-Horizon Robotic Tasks, 2026d.

Wenyi Li, Yizhe Chi, Youjie Zheng, Chunyu Liu, Xiaoqiu Wang, Mingju Gao, Zheng Jiang, Deyao Hong, Calvin Xiao, Kaisen Yang, et al. EmbodiedEvo: Self-Evolving Coding Agent For Physical Intelligence, 2026e.

Xuanlin Li, Kyle Hsu, Jiayuan Gu, Oier Mees, Karl Pertsch, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating Real-World Robot Manipulation Policies in Simulation. In *Proceedings of The 8th Conference on Robot Learning*, 2025a.

Yi Li, Yuquan Deng, Jesse Zhang, Joel Jang, Marius Memmel, Caelan Garrett, Fabio Ramos, Dieter Fox, Anqi Li, Abhishek Gupta, and Ankit Goyal. HAMSTER: Hierarchical Action Models for Open-World Robot Manipulation. In *International Conference on Learning Representations*, 2025b.

Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as Policies: Language Model Programs for Embodied Control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

Lizhou Liang, Xinyu Zhong, Miao Pan, Xiaohe Zhou, Xuanyu Liu, Qinfeng Li, Peng Li, Jintao Chen, Xuhong Zhang, and Wenqi Zhang. EmbodiedMemory-Bench: Benchmarking Embodied Memory for Long-Horizon Embodied Tasks, 2026a.

Yuanzhi Liang, Xufeng Zhan, Haibin Huang, Chi Zhang, and Xuelong Li. From World Action Models to Embodied Brains: A Roadmap for Open-World Physical Intelligence, 2026b.

Qiayuan Liao, Takara E. Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C. Karen Liu. BeyondMimic: From Motion Tracking to Versatile Humanoid Control via Guided Diffusion, 2025.

Qiayuan Liao, Yiyang Shao, Junfeng Long, Kevin Zakka, Roberto Horowitz, Pieter Abbeel, and Koushil Sreenath. SuperMimic: Scaling humanoid Whole-Body control with one differentiable world model, 2026.

Yunn Kang Lim, Pengzhan Sun, Ziyi Bai, Xun Xu, Angela Yao, Xulei Yang, and Shijie Li. PRISM: Planning and Reasoning with Intent in Simulated Embodied Environments, 2026.

Fanqi Lin, Ruiqian Nai, Yingdong Hu, Jiacheng You, Junming Zhao, and Yang Gao. OneTwoVLA: A Unified Vision-Language-Action Model with Adaptive Reasoning, 2025.

Kevin Lin, Christopher Agia, Toki Migimatsu, Marco Pavone, and Jeannette Bohg. Text2Motion: From Natural Language Instructions to Feasible Plans. *Autonomous Robots*, 2023.

Lijian Lin, Ye Zhu, Fan Zhang, Yunfei Liu, Baofeng Li, Xianwen Zeng, Jianan Wang, and Yu Li. MIRA: Real-Time Full-Duplex Human-Robot Interaction for Embodied Companions, 2026.

Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. LLM+P: Empowering Large Language Models with Optimal Planning Proficiency, 2023a.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking Knowledge Transfer for Lifelong Robot Learning. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2023b.

Chenzhi Liu, Yue Zhang, Jiehong Lin, Jianan Wang, Bo Wang, Zhongrui Wang, and Xiaojuan Qi. MobiAgent: Dual-Loop Recursive Policy Self-Improvement for Long-Horizon Mobile Manipulation, 2026a.

Haowen Liu, Xirui Li, Shaoxiong Yao, Peng Shi, Tianyi Zhou, Jia-Bin Huang, Furong Huang, and Jiayuan Mao. Guava: Distilling Frontier VLMs into a Compact Agent through a Robotic Manipulation Harness, 2026b.

Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023c.

Peiqi Liu, Yaswanth Orru, Jay Vakil, Chris Paxton, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. Demonstrating OK-Robot: What Really Matters in Integrating Open-Knowledge Models for Robotics. In *Robotics: Science and Systems (RSS)*, 2024.

Renhang Liu, Navonil Majumder, Tej Deep Pala, and Soujanya Poria. RoboQuest: Generalist Physical Agents that Search, Inspect and Test, 2026c.

Yang Liu, Weixing Chen, Xinshuai Song, Tao Pu, Siwen Mo, Yongjie Bai, Zihao Chen, Qianran Sun, Liruo Zhong, Ying Shen, et al. PhyAgentOS: A Self-Evolving Operating System for Embodied Agents with Decoupled Cognitive Planning and Physical Execution, 2026d.

Yicheng Liu, Zibin Dong, Baijun Ye, Tianyuan Yuan, Tao Jiang, Anqi Yang, Shicheng Cao, Haonan Liu, Yue Sun, Zihan Guo, et al. G0.5: One Autoregressive Stream for Robot Reasoning and Action, 2026e.

Zeyi Liu, Arpit Bahety, and Shuran Song. REFLECT: Summarizing Robot Experiences for Failure Explanation and Correction. In *Proceedings of The 7th Conference on Robot Learning (CoRL)*, 2023d.

Zeyu Lou, Yanhong Zeng, Yong Wang, and Chenyang Si. Know Your Body: A Harness for Direct and Self-Improving Robot Control with VLMs, 2026.

Runyu Lu, Yubo Wu, Ethan Kou, Haotian Lin, Letian Fu, Wenli Xiao, Ajay Mandlekar, Yinzhen Xu, Guanya Shi, Ken Goldberg, et al. ASPIRE: Agentic Skills Discovery for Robotics, 2026.

Xiaoya Lu, Zeren Chen, Xuhao Hu, Yijin Zhou, Weichen Zhang, Dongrui Liu, Lu Sheng, and Jing Shao. IS-Bench: Evaluating Interactive Safety of VLM-Driven Embodied Agents in Daily Household Tasks, 2025.

Hao Luo, Yicheng Feng, Wanpeng Zhang, Sipeng Zheng, Ye Wang, Haoqi Yuan, Jiazheng Liu, Chaoyi Xu, Qin Jin, and Zongqing Lu. Being-H0: Vision-Language-Action Pretraining from Large-Scale Human Videos, 2025.

Meng Luo, Zechen Bai, Shanqing Xu, Yanlin Li, Kaixin Li, Mengchen Qian, Weifeng Lu, Ran Zhang, Yingtian Zou, Ziyang Luo, et al. Toward Recursive Self-Improvement in Physical AI, 2026a.

Yulin Luo, Chun-Kai Fan, Menghang Dong, Jiayu Shi, Xiangju Mi, Mengdi Zhao, Bo-Wen Zhang, Cheng Chi, Jiaming Liu, Gaole Dai, et al. RoboBench: A Comprehensive Evaluation Benchmark for Multimodal Large Language Models as Embodied Brain. In *European Conference on Computer Vision (ECCV)*, 2026b.

Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Fernando Castañeda, Sirui Chen, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, et al. SONIC: Supersizing Motion Tracking for Natural Humanoid Whole-Body Control. *Science Robotics*, 2026c.

Yiheng Lyu, Xueying Jiang, Wenhao Li, Shijian Lu, and Gongjie Zhang. CodeActionBench: Evaluating Agentic Code-as-Policy for Embodied Manipulation, 2026.

Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-Level Reward Design via Coding Large Language Models. In *International Conference on Learning Representations*, 2024a.

Yecheng Jason Ma, William Liang, Hung-Ju Wang, Yuke Zhu, Linxi Fan, Osbert Bastani, and Dinesh Jayaraman. DrEureka: Language model guided Sim-To-Real transfer. In *Robotics: Science and Systems XX (RSS)*, 2024b.

Yecheng Jason Ma, Joey Hejna, Chuyuan Fu, Dhruv Shah, Jacky Liang, Zhuo Xu, Sean Kirmani, Peng Xu, Danny Driess, Ted Xiao, Osbert Bastani, Dinesh Jayaraman, Wenhao Yu, Tingnan Zhang, Dorsa Sadigh, and Fei Xia. Vision Language Models are In-Context Value Learners. In *International Conference on Learning Representations*, 2025.

Yueen Ma, Zixing Song, Yuzheng Zhuang, Jianye Hao, and Irwin King. A Survey on Vision-Language-Action Models for Embodied AI. *IEEE Transactions on Neural Networks and Learning Systems*, 2026.

Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. Robot Operating System 2: Design, Architecture, and Uses in the Wild. *Science Robotics*, 2022.

Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, et al. OpenEQA: Embodied Question Answering in the Era of Foundation Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. CALVIN: A Benchmark for Language-Conditioned Policy Learning for Long-Horizon Robot Manipulation Tasks. *IEEE Robotics and Automation Letters*, 2022.

Achu Menon, Sravanthi Machcha, Sabrina Zou, Tzu Kit Chan, and Jay Chooi. GPT-6 Astra on Robotic Manipulation. Robocurve, 2026. URL https://openai.robocurve.org/gpt-6-astra/.

METR. Time Horizon 1.1. https://metr.org/blog/2026-01-29-time-horizon-1-1/, 2026a.

METR. Task-Completion Time Horizons of Frontier AI Models. https://metr.org/time-horizons/, 2026b.

Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple Open-Vocabulary Object Detection with Vision Transformers. In *European Conference on Computer Vision (ECCV)*, 2022.

Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. EmbodiedGPT: Vision-Language Pre-Training via Embodied Chain of Thought. In *Advances in Neural Information Processing Systems*, 2023.

Soroush Nasiriany, Sepehr Nasiriany, Abhiram Maddukuri, and Yuke Zhu. RoboCasa365: A Large-Scale Simulation Framework for Training and Benchmarking Generalist Robots. In *International Conference on Learning Representations (ICLR)*, 2026.

Zaruhi Navasardyan, Tatul Danielyan, and Hrant Davtyan. FailBench: How Reliable are VLMs at Judging Robot Task Success?, 2026.

Dung Manh Nguyen, Thang Chau Phan, Nam Hai Le, Thong T. Doan, Nam V. Nguyen, Quang Pham, and Nghi D. Q. Bui. CodeMMLU: A Multi-Task Benchmark for Assessing Code Understanding & Reasoning Capabilities of CodeLLMs. In *International Conference on Learning Representations (ICLR)*, 2025.

Nils J. Nilsson. Shakey the Robot. Technical Note 323, Artificial Intelligence Center, SRI International, Menlo Park, CA, April 1984. URL https://www.sri.com/wp-content/uploads/2021/12/629.pdf.

NVIDIA. Cosmos 3: Omnimodal World Models for Physical AI, 2026.

NVIDIA, Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, et al. Cosmos World Foundation Model Platform for Physical AI, 2025a.

NVIDIA, Alisson Azzolini, Junjie Bai, Hannah Brandon, Jiaxin Cao, Prithvijit Chattopadhyay, Huayu Chen, Jinju Chu, Yin Cui, et al. Cosmos-Reason1: From Physical Common Sense To Embodied Reasoning, 2025b.

NVIDIA, Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi "Jim" Fan, Yu Fang, Dieter Fox, et al. GR00T N1: An Open Foundation Model for Generalist Humanoid Robots, 2025c.

ODrive Robotics. Hardware Configuration: Temperature Regulation. https://docs.odriverobotics.com/v/latest/manual/hardware-config.html, 2026.

Christian Oefinger, Finn Rasmus Schäfer, Korbinian Moller, Mattia Piccinini, and Johannes Betz. Validate the Dream Before You Trust Its Verdict: Admissibility for World-Model Simulators, 2026.

OpenAI. Learning to Reason with LLMs. https://openai.com/index/learning-to-reason-with-llms/, September 2024.

OpenAI. GPT-6 Astra: A New Generation of Intelligence. https://openai.com/index/gpt-6-astra/, 2026a.

OpenAI. GPT-6 Astra System Card. https://deploymentsafety.openai.com/gpt-6-astra, 2026b.

OpenAI. Why We No Longer Evaluate SWE-bench Verified. https://openai.com/index/why-we-no-longer-evaluate-swe-bench-verified, 2026c.

OpenRUA Team. Behind the RoboDojo Boom: Revisiting Embodied Evaluation of Robot-Use Agents. https://openrua.ai/revisiting-rua, 2026.

Aishwarya Padmakumar, Jesse Thomason, Ayush Shrivastava, Patrick Lange, Anjali Narayan-Chen, Spandana Gella, Robinson Piramuthu, Gokhan Tur, and Dilek Hakkani-Tur. TEACh: Task-Driven Embodied Agents That Chat. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.

Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with SWE-Gym. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*, 2023.

Deepak Pathak and Abhinav Gupta. The Hidden Pillar of Robotics. https://www.skild.ai/blogs/skild-crosses-100m-arr, September 2026.

Tejal Patwardhan, Rachel Dias, Elizabeth Proehl, Grace Kim, Michele Wang, Olivia Watkins, Simón Posada Fishman, Marwan Aljubeh, Phoebe Thacker, Laurance Fauconnet, et al. GDPval: Evaluating AI model performance on Real-World economically valuable tasks, 2025.

Petko Petkov, Giuliano Albanese, Nikolay Nikolov, Luc Van Gool, Jonathan Tremblay, and Danda Pani Paudel. LADDER: A Long-Horizon Robot Manipulation Benchmark, 2026.

Physical Intelligence, Ali Amin, Raichelle Aniceto, Ashwin Balakrishna, Kevin Black, Ken Conley, Grace Connors, James Darpinian, Karan Dhabalia, Jared DiCarlo, et al. $\pi^*_{0.6}$: a VLA That Learns From Experience, 2025.

Playground Research. Introducing the Robot Intelligence Index, 2026. URL https://www.withplayground.ai/blog/robot-intelligence-index.

Chonghan Qin, Dongming Wu, Kaituo Feng, Jiaxu Wang, Xinlei Yu, Peiwen Sun, Yijia Fan, Xilin He, Bingao Chen, Weitong Lian, et al. Towards Agentic Robots: Model, Data, Environment, and Harness, 2026a.

Xue Qin, Simin Luan, John See, Zeyd Boukhers, Cong Yang, and Zhijun Li. Harnessing Embodied Agents: Runtime Governance for Policy-Constrained Execution, 2026b.

Yikai Qin, Yifei Deng, Mingjian Liang, Wenxuan Song, Zepeng Lin, Zhiyi Jiang, Jiajun Fu, Qiao Sun, Huashuo Lei, Xicheng Gong, et al. EmbodiedSmith: Scaling Embodied Data through Recursive Self-Improvement Flywheel in Simulation, 2026c.

Ram Ramrakhya, Matthew Chang, Xavier Puig, Ruta Desai, Zsolt Kira, and Roozbeh Mottaghi. Grounding Multimodal LLMs to Embodied Agents that Ask for Help with Reinforcement Learning, 2025.

Krishan Rana, Jesse Haviland, Sourav Garg, Jad Abou-Chakra, Ian Reid, and Niko Suenderhauf. SayPlan: Grounding Large Language Models using 3D Scene Graphs for Scalable Robot Task Planning. In *Proceedings of The 7th Conference on Robot Learning (CoRL)*, 2023.

Kanishka Rao. Gemini 4 Argon for Robotics. https://developers.googleblog.com/gemini-4-argon-for-robotics/, October 2026.

Soumil Rathi. RoboMME-Interference: Benchmarking Robot Memory Under Interference, 2026.

Allen Z. Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, et al. Robots That Ask For Help: Uncertainty Alignment for Large Language Model Planners. In *Proceedings of The 7th Conference on Robot Learning (CoRL)*, 2023.

Robocurve. Robocurve: Real-World Evaluations of Physical AI. https://robocurve.org/, 2026.

ROBOTIS. DYNAMIXEL MX-106T/R(2.0) e-Manual: Temperature Limit(31) and Shutdown(63). https://emanual.robotis.com/docs/en/dxl/mx/mx-106-2/, 2026.

Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to Walk in Minutes Using Massively Parallel Deep Reinforcement Learning. In *Proceedings of the 5th Conference on Robot Learning*, 2022.

Som Sagar and Ransalu Senanayake. Test-Time adaptation of manipulation policies under actuator degradation, 2026.

Sahar Salimpour, Lei Fu, Kajetan Rachwał, Pascal Bertrand, Kevin O'Sullivan, Robert Jakob, Farhad Keramat, Leonardo Militano, Giovanni Toffetti, Harry Edelman, et al. Towards Embodied Agentic AI: Review and Classification of LLM- and VLM-Driven Robot Autonomy and Interaction, 2025.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language Models Can Teach Themselves to Use Tools. In *Advances in Neural Information Processing Systems*, 2023.

SenseNova, Joint Research Team, Zhongang Cai, Ziang Cao, Tailai Chen, Han Chong Chua, Linjun Dai, Hanming Deng, Hokin Deng, Haiwen Diao, et al. Has GPT-6-Astra Achieved Embodied Intelligence?, 2026.

Pierre Sermanet, Anirudha Majumdar, Alex Irpan, Dmitry Kalashnikov, and Vikas Sindhwani. Generating Robot Constitutions & Benchmarks for Semantic Safety. In *Proceedings of The 9th Conference on Robot Learning (CoRL)*, 2025.

Rui Shao, Wei Li, Lingsen Zhang, Renshan Zhang, Zhiyang Liu, Ran Chen, and Liqiang Nie. Large VLM-based Vision-Language-Action Models for Robotic Manipulation: A Survey, 2025.

Hao Shi, Bin Xie, Yingfei Liu, Lin Sun, Fengrong Liu, Tiancai Wang, Erjin Zhou, Haoqiang Fan, Xiangyu Zhang, and Gao Huang. MemoryVLA: Perceptual-Cognitive Memory in Vision-Language-Action Models for Robotic Manipulation. In *International Conference on Learning Representations*, 2026.

Lucy Xiaoyang Shi, Brian Ichter, Michael Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James Tanner, Anna Walling, Haohuan Wang, Niccolo Fusai, et al. Hi Robot: Open-Ended Instruction Following with Hierarchical Vision-Language-Action Models. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language Agents with Verbal Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 2023.

Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

Chenglei Si, Yanzhe Zhang, Ryan Li, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. Design2Code: Benchmarking multimodal code generation for automated Front-End engineering. In *Proceedings of NAACL 2025*, 2025.

Ruiyang Si, Jianxin Bi, Shunyu Yang, Rui Ni, Wenbo Huang, Qiang Wang, Shulong Jiang, Yintianrun Zhang, Feixiang Ruan, Kewei Zhang, et al. Fewer Tokens, Better Action: GPT-6 Astra Robot Agents with 14% Higher Success Rate but 65% Fewer Tokens, 2026.

Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. ProgPrompt: Generating Situated Robot Task Plans using Large Language Models. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

Li Siyao, Jiawei Gu, Shuai Liu, Kairui Hu, Zekun Li, Linjie Li, Chengcheng Tang, Po-Chen Wu, Ivan Shugurov, Lingni Ma, et al. HumanCLAW: Can Vision-Language Models Act Through a Body?, 2026.

Haoming Song, Delin Qu, Yuanqi Yao, Qizhi Chen, Qi Lv, Yiwen Tang, Modi Shi, Guanghui Ren, Maoqing Yao, Bin Zhao, et al. Hume: Introducing System-2 Thinking in Visual-Language-Action Model, 2025.

Edward Sun, Christopher Leet, Sravanthi Machcha, Achyutha Menon, Tzu Kit Chan, Yuchen Cui, and Jay Chooi. Refuses in Chat, Complies with Arms: LLM Refusal Does Not Survive Robot Control, 2026a.

Jianli Sun, Chengxiang Li, and Yonglin Tian. Make π0.5 Learn to Call Astra for Reliable Robot Manipulation, 2026b.

Samson Sun, Tianyi Yang, Tengyue Wang, Yikai Xue, Zhengjie Xu, Lingming Zhang, Qichen Zhang, Chao Liang, and Zhipeng Zhang. Beyond Isolation: A Unified Benchmark for General-Purpose Navigation, 2026c.

Balakumar Sundaralingam, Siva Kumar Sastry Hari, Adam Fishman, Caelan Garrett, Karl Van Wyk, Valts Blukis, Alexander Millane, Helen Oleynikova, Ankur Handa, Fabio Ramos, et al. cuRobo: Parallelized Collision-Free Minimum-Jerk Robot Motion Generation, 2023.

Sunday Robotics. ACT-2 Preview: Generalizing Reliability. https://www.sunday.ai/blog/act-2-preview, July 2026.

Martin Sundermeyer, Arsalan Mousavian, Rudolph Triebel, and Dieter Fox. Contact-GraspNet: Efficient 6-DoF Grasp Generation in Cluttered Scenes. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

Nathan J. Szymanski, Bernardus Rendy, Yuxing Fei, Rishi E. Kumar, Tanjin He, David Milsted, Matthew J. McDermott, Max Gallant, Ekin Dogus Cubuk, Amil Merchant, et al. An Autonomous Laboratory for the Accelerated Synthesis of Inorganic Materials. *Nature*, 2023.

Nathan J. Szymanski, Bernardus Rendy, Yuxing Fei, Rishi E. Kumar, Tanjin He, David Milsted, Matthew J. McDermott, Max Gallant, Ekin Dogus Cubuk, Amil Merchant, et al. Author Correction: An Autonomous Laboratory for the Accelerated Synthesis of Inorganic Materials. *Nature*, 2026.

Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Volume 1*, 2019.

Huajie Tan, Enshen Zhou, Zhiyu Li, Yijie Xu, Yuheng Ji, Xiansheng Chen, Cheng Chi, Pengwei Wang, Huizhu Jia, Yulong Ao, et al. RoboBrain 2.5: Depth in Sight, Time in Mind, 2026.

Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-Real: Learning Agile Locomotion For Quadruped Robots. In *Robotics: Science and Systems XIV (RSS)*, 2018.

Guoqin Tang, Qingxuan Jia, Gang Chen, Tong Li, Zeyuan Huang, Zihang Lv, and Ning Ji. VLM-DEWM: Dynamic External World Model for Verifiable and Resilient Vision-Language Planning in Manufacturing, 2026.

Tencent Robotics X, HY Vision Team, Xumin Yu, Zuyan Liu, Ziyi Wang, He Zhang, Yongming Rao, Fangfu Liu, Yani Zhang, et al. HY-Embodied-0.5: Embodied Foundation Models for Real-World Agents, 2026.

Thinking Machines Lab. Interaction Models: A Scalable Approach to Human-AI Collaboration. https://thinkingmachines.ai/blog/interaction-models/, 2026.

Wenqing Tian, Zeyu Zhang, Zhaocheng Liu, Fengwei Liu, Qiang Liu, and Liang Wang. PhysEvo: Astra Can Act, Let It, 2026.

Jan-Philipp Töberg, Svenja Kenneweg, and Philipp Cimiano. RoboCSKBench: Benchmarking Embodied Commonsense Capabilities of Large Language Models. In *2025 22nd International Conference on Ubiquitous Robots (UR)*, 2025.

Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.

Emanuel Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 2004.

Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A Physics Engine for Model-Based Control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.

Marcel Torne, Karl Pertsch, Homer Walke, Kyle Vedder, Suraj Nair, Brian Ichter, Allen Z. Ren, Haohuan Wang, Jiaming Tang, Kyle Stachowicz, et al. MEM: Multi-Scale Embodied Memory for Vision Language Action Models, 2026.

TRI LBM Team, Jose Barreiros, Andrew Beaulieu, Aditya Bhat, Rick Cory, Eric Cousineau, Hongkai Dai, Ching-Hsin Fang, Kunimatsu Hashimoto, Muhammad Zubair Irshad, et al. A Careful Examination of Large Behavior Models for Multitask Dexterous Manipulation, 2025.

Yang Wan, Xihang Yue, Zhirui Liu, Ziyuan Chu, Shuxun Wang, Yuhan Chen, Xiaonan Jiang, Xukun Zhu, Yubo Dong, and Linchao Zhu. No Free Checker: A Survey of Verifiers for Robot Policies, 2026.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An Open-Ended Embodied Agent with Large Language Models. *Transactions on Machine Learning Research*, 2024a.

Hanwen Wang, Weizhi Zhao, Xiangyu Wang, Siyuan Huang, He Lin, Boyuan Zheng, Rongtao Xu, Gang Wang, Yao Mu, He Wang, et al. DexJoCo: A benchmark and toolkit for Task-Oriented dexterous manipulation on MuJoCo, 2026a.

Jiaming Wang. What Stops Recursive Self-Improvement in Robotics? Lessons from 123 Rounds of Agentic Skill Discovery, 2026.

Jiaming Wang, Yuhua Jiang, Zhengcheng Shen, and Harold Soh. A Survey of Agentic Robotics: Toward Continual Self-Improvement, 2026b.

Kailin Wang, Haoxiang Jie, Yaoyuan Yan, Zhiyou Heng, and Zhaosong Li. RegenHarness: A Robot Agent Harness with Evidence-Gated Recursive Self-Improvement, 2026c.

Kangrui Wang, Qineng Wang, Changhe Chen, Pingyue Zhang, Chi Gui, Feishi Wang, Canyu Chen, Wenli Xiao, Xiaoyu Huang, Tonghe Zhang, Linjie Li, Zihan Wang, Minyeong Kim, Yiyou Sun, Mingkai Deng, Jinyu Hou, Qi Zhu, Han Liu, Eric P. Xing, Huaxiu Yao, Zhiwen Fan, Yuejiang Liu, Dawn Song, Pieter Abbeel, Li Fei-Fei, Jiajun Wu, Yuke Zhu, Guanya Shi, Ruohan Zhang, and Manling Li. Embodied first exam, 2026d.

Qi Wang, Tianyi Wang, Chengyang Li, Shikun Ban, Yurun Chen, Yizhong Ge, Jason Qin, Chengtai Li, and Wentao Zhu. Towards the Harness of Embodied Agents, 2026e.

Xin Wang, Wenhao Wu, Menghao Zhang, Zhi Wang, Kun Shao, Jian Luan, Yang Li, Qing Li, Shangding Gu, Huichi Zhou, et al. HarnessPAI: An Evolving Harness for Physical AI, 2026f.

Xingyao Wang, Yangyi Chen, Lifan Yuan, Yizhe Zhang, Yunzhu Li, Hao Peng, and Heng Ji. Executable Code Actions Elicit Better LLM Agents. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024b.

Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, et al. OpenHands: An Open Platform for AI Software Developers as Generalist Agents. In *International Conference on Learning Representations*, 2025a.

Xinyi Wang, Bangzheng Li, Hanyu Wang, Lin Zhang, Andrea Yaoyun Cui, Zhenggang Tang, Shyamal Buch, and Dejia Xu. SEAR: Self-Evolving Agentic Robotics, 2026g.

Yen-Jen Wang, Haozhe Jiang, Shuying Deng, Haoru Xue, Weirui Ye, Rocky Duan, Nika Haghtalab, S. Shankar Sastry, Pieter Abbeel, and Haozhi Qi. Reconstruct, Practice, Go Real: Guided Self-Improvement for Embodied Agents, 2026h.

Yi Wang, Xinchen Li, Pengwei Xie, Pu Yang, Buqing Nie, Yunuo Cai, Qinglin Zhang, Chendi Qu, Jeffrey Wu, Jianheng Song, et al. Learning While Deploying: Fleet-Scale Reinforcement Learning for Generalist Robot Policies, 2026i.

Yi Ru Wang, Jiafei Duan, Dieter Fox, and Siddhartha Srinivasa. NEWTON: Are Large Language Models Capable of Physical Reasoning? In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023.

Yike Wang, Huaisheng Zhu, Zhengyu Hu, Yige Yuan, Zhengyu Chen, Shakti Senthil, Hannaneh Hajishirzi, Yulia Tsvetkov, Noah A. Smith, Pradeep Dasigi, et al. Rethinking the Evaluation of Harness Evolution for Agents, 2026j.

Yubang Wang, Zezheng Huai, Yifei Zhang, Jinhao Wu, Yitong Chen, Haozhe Zhang, Sixian Li, Jingjing Gong, Zhenyang Lin, Ziyang Chen, et al. Embodied task agents can perform fine manipulation through tactile In-Context learning, 2026k.

Zixuan Wang, Bo Yu, Junzhe Zhao, Wenhao Sun, Sai Hou, Shuai Liang, Xing Hu, Yinhe Han, and Yiming Gan. KARMA: Augmenting Embodied AI Agents with Long-and-short Term Memory Systems. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025b.

Zixuan Wang, Yuqi Liu, Yuxin Chen, Sengqiao Yang, Jinhui Ye, Pengguang Chen, Changsheng Lu, Shu Liu, Bei Yu, and Jiaya Jia. CarryOn: A Dual-Loop Harness for Personalized and Reliable Robot Assistance, 2026l.

Jason Wei. Asymmetry of Verification and Verifier's Rule. https://www.jasonwei.net/blog/asymmetry-of-verification-and-verifiers-law, 2025.

Yuxi Wei, Ye Fang, Shuai Yang, Tong Wu, Tai Wang, and Xiaojuan Qi. Towards a Systematic Study of Embodied System 2 over Full Task Execution Cycles, 2026a.

Zhijie Wei, Ferris Tan, Jinghui Wang, and Hao Zhou. Scale and Selection: What Makes Automatic Harness Evolution Work for Visual-Interface Robot Agents, 2026b.

Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. FoundationPose: Unified 6D Pose Estimation and Tracking of Novel Objects. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Jimmy Wu, Rika Antonova, Adam Kan, Marion Lepert, Andy Zeng, Shuran Song, Jeannette Bohg, Szymon Rusinkiewicz, and Thomas Funkhouser. TidyBot: Personalized Robot Assistance with Large Language Models. *Autonomous Robots*, 2023.

Tao Wu, Chuhao Zhou, Guangyu Zhao, Haozhi Cao, Yewen Pu, and Jianfei Yang. When Robots Should Say "I Don't Know": Benchmarking Abstention in Embodied Question Answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.

Ted Xiao, Harris Chan, Pierre Sermanet, Ayzaan Wahid, Anthony Brohan, Karol Hausman, Sergey Levine, and Jonathan Tompson. Robotic Skill Acquisition via Instruction Augmentation with Vision-Language Models. In *Robotics: Science and Systems (RSS)*, 2023.

Wenli Xiao, Jia Xie, Tonghe Zhang, Haotian Lin, Letian Fu, Haoru Xue, Jalen Lu, Yi Yang, Cunxi Dai, Zi Wang, et al. ENPIRE: Agentic Robot Policy Self-Improvement in the Real World, 2026.

Chuanliang Xie, Boyu Ma, Gen Li, Yizhou Liu, Houwang Chen, Xinyu Zhou, and Jianfei Yang. Fast Plans, Faithful Actions: Closing the Planning-Execution Gap in Hierarchical Vision-Language-Action Models, 2026.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments. In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.

Bingxin Xu, Wenrui Bao, Xinxin Liu, Jiawei Zhou, Yuzhang Shang, and Emilio Ferrara. ChunkAct: Action Chunking for Agentic Robotics, 2026a.

Bingxin Xu, Yuzhang Shang, Zhen Dong, and Emilio Ferrara. Coding Agents with Harness for Safe Robot Control, 2026b.

Bingxin Xu, Yuzhang Shang, and Emilio Ferrara. Don't Drop the BATON: Long-Horizon Robot Manipulation via Agentic Subtask Exploration and Transition-aware Memory, 2026c.

Tianruo Rose Xu, Jiawei Ren, Yichi Yang, Zhaoxu Zheng, and Lianhui Qin. RT-Safe: Benchmarking Agent Safety in Real-Time Embodied Environment, 2026d.

Yajiao Xu, Jin Zhang, Jiangbo Ai, Tao Jiang, Mo Xu, Lina Huang, and Chengfu Huo. CommerceVibe: Learning to design E-Commerce creatives as executable visual code via Dual-Feedback reinforcement learning, 2026e.

Adina Yakefu, Bin Xie, Chongyang Xu, Enwen Zhang, Erjin Zhou, Fan Jia, Haitao Yang, Haoqiang Fan, Haowei Zhang, Hongyang Peng, et al. RoboChallenge: Large-scale Real-robot Evaluation of Embodied Policies, 2025.

Mi Yan, Wenhao Zhang, Zhiqi Zhang, Yu Peng, Tangxinyu Wang, Lingfei Zhai, Jiayi Su, Shengliang Deng, Lin Peng, Yaowei Liu, et al. ZETA: A Controlled Study of Zero-Shot Cross-Embodiment VLA Transfer for Tabletop Manipulation, 2026.

Chen Yang, Linzhe Shi, Hang Zhang, Mu Xu, Jiansheng Fan, and Chen Wang. GPT-6 Astra Is a Native Visuo-Tactile Policy for Contact-Rich Manipulation, 2026a.

Dongxu Yang. Device Context Protocol: A Compact, Safety-First Architecture for LLM-Driven Control of Constrained Devices, 2026.

Ganlin Yang, Tianyi Zhang, Haoran Hao, Weiyun Wang, Yibin Liu, Dehui Wang, Guanzhou Chen, Zijian Cai, Junting Chen, Weijie Su, et al. Vlaser: Vision-Language-Action Model with Synergistic Embodied Reasoning. In *International Conference on Learning Representations*, 2026b.

Jenai Xuning Yang, Rishit Dagli, Alex Zook, Hugo Hadfield, Ankit Goyal, Stan Birchfield, Fabio Ramos, and Jonathan Tremblay. RoboLab: A High-Fidelity Simulation Benchmark for Analysis of Task Generalist Policies. In *Robotics: Science and Systems (RSS) XXII*, 2026c.

Jenai Xuning Yang, Ashna Khetan, and Shun Zhang. Astranomy: An Empirical Study of GPT-6 Astra as a Native Robot Policy. https://research.nvidia.com/labs/srl/projects/astranomy/, 2026d.

Jianwei Yang, Reuben Tan, Qianhui Wu, Ruijie Zheng, Baolin Peng, Yongyuan Liang, Yu Gu, Mu Cai, Seonghyeon Ye, Joel Jang, et al. Magma: A Foundation Model for Multimodal AI Agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025a.

Jihan Yang, Shusheng Yang, Anjali W. Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in Space: How Multimodal Large Language Models See, Remember, and Recall Spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025b.

John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. SWE-agent: Agent-Computer Interfaces Enable Automated Software Engineering. In *Advances in Neural Information Processing Systems*, 2024.

Jonathan Yang, Chelsea Finn, and Dorsa Sadigh. Cross-Controller transfer via In-Context behavior adaptation, 2026e.

Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, et al. EmbodiedBench: Comprehensive Benchmarking Multi-modal Large Language Models for Vision-Driven Embodied Agents. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025c.

Sihan Yang, Runsen Xu, Yiman Xie, Sizhe Yang, Mo Li, Jingli Lin, Chenming Zhu, Xiaochen Chen, Haodong Duan, Xiangyu Yue, et al. MMSI-Bench: A Benchmark for Multi-Image Spatial Intelligence. In *International Conference on Learning Representations (ICLR)*, 2026f.

Yuncong Yang, Han Yang, Jiachen Zhou, Peihao Chen, Hongxin Zhang, Yilun Du, and Chuang Gan. 3D-Mem: 3D Scene Memory for Embodied Exploration and Reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025d.

Zhenhua Yang, Jiafeng Liang, Xingyu Zhou, Yi Jiang, Xiayu Cao, Zui Chen, Luqi Wang, Haiyang Sun, Yufei Xiang, Chang Li, et al. Proactive Physical AI through Continuous Observation and Timely Intervention, 2026g.

Zhenjie Yang, Yideng Zhang, Dongjie Zhang, Chenyu Jiang, Xianshuai Liu, Yufeng Li, Zuhao Ge, Xingyu Jiao, Zheng Zhang, Kaiyu He, et al. Bench2Dex: Benchmarking Visuo-Tactile bimanual dexterous manipulation across dexterous hands, 2026h.

Zhiqin Yang, Jingwen Fu, Yuhan Liu, Hengyu Liu, Yonggang Zhang, Kainan Cao, Zizhuo Zhang, Chenxin Li, Ruibin Yuan, Jiahao Pan, et al. Scaling Large Reasoning Models beyond Human Supervision: A Path toward Superintelligence, 2026i.

Zhiqin Yang, Chenxin Li, Xiaomeng Hu, Yibin Liu, Weidong Huang, Jiankai Sun, Haitao Li, Zijian Wu, Yuzhi Huang, Fanding Huang, et al. RobotWorld: Benchmarking Multimodal Agents for Robot Use Across Diverse Tasks and Embodiments, 2026j.

Jiashu Yao, Haoyu Wen, Siyuan Gao, Yuhang Guo, Zeming Liu, and Heyan Huang. HomeSafeBench: Benchmarking Embodied Vision-Language Models in Free-Exploration Home Safety Inspection, 2025.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing Reasoning and Acting in Language Models. In *International Conference on Learning Representations*, 2023.

Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. τ-bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains, 2024.

Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, et al. World Action Models are Zero-shot Policies, 2026a.

Xiaokang Ye, Siddhant Hitesh Mantri, Zimeng Chen, Edward Zhang, Zhaoxu Zheng, Yuanheng Li, Yizhao Chen, Tianyang Huang, and Lianhui Qin. Code4Scene: Benchmarking coding agents for constructing and editing 3D scenes, 2026b.

Sheng Yin, Xianghe Pang, Yuanzhuo Ding, Menglan Chen, Yutong Bi, Yichen Xiong, Wenhao Huang, Zhen Xiang, Jing Shao, and Siheng Chen. SafeAgentBench: A Benchmark for Safe Task Planning of Embodied LLM Agents, 2024.

Haibao Yu, Jiaqi Meng, Zijian Zhang, Bowen Jing, Lei Yang, Yue Hu, Chenchen Ge, Yiming Hou, Chenbo Xin, Dawen Ye, et al. LLM4Robot-Bench: How Far Can GPT-6 Astra Go in Robotic Manipulation? A Hierarchical, Paired Sim-Real Benchmark, 2026a.

Yue Yu, Bin Zhu, Junhui Li, Jiayu Wang, Shijie Zhou, and Jingjing Chen. ReTrace: Cross-Stage Failure Attribution for Self-Improving Robotic Manipulation, 2026b.

Jiarui Yuan, Yuxin Zuo, Xiaoyu Huang, Wenbo Zhang, Kwanyoung Park, Junlin Yang, Tianxiang Zhang, Ruoqu Chen, Yutao Ouyang, Bingxiang He, et al. Democratizing robot-use agents: A unified harness with open models. *arXiv preprint arXiv:2611.01560*, 2026a.

Mingqi Yuan, Tao Yu, Wenqi Ge, Xiuyong Yao, Huijiang Wang, Jiayu Chen, Bo Li, Wei Zhang, Wenjun Zeng, Hua Chen, et al. A survey of behavior foundation model: Next-Generation Whole-Body control system of humanoid robots, 2025a.

Tianyuan Yuan, Zibin Dong, Yicheng Liu, and Hang Zhao. Fast-WAM: Do World Action Models Need Test-time Future Imagination?, 2026b.

Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. RoboPoint: A Vision-Language Model for Spatial Affordance Prediction in Robotics. In *Proceedings of The 8th Conference on Robot Learning*, 2025b.

Yifu Yuan, Haiqin Cui, Yaoting Huang, Yibin Chen, Fei Ni, Zibin Dong, Pengyi Li, Yan Zheng, Hongyao Tang, and Jianye Hao. Embodied-R1: Reinforced Embodied Reasoning for General Robotic Manipulation. In *International Conference on Learning Representations*, 2026c.

Yifu Yuan, Yaoting Huang, Xianze Yao, Yutong Li, Shuoheng Zhang, Linqi Han, Pengyi Li, Jiangeng Sun, Wenting Jia, Zhao Zhang, et al. Embodied-R1.5: Evolving Physical Intelligence via Embodied Foundation Models, 2026d.

Andy Zeng, Maria Attarian, Brian Ichter, Krzysztof Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aveek Purohit, Michael Ryoo, Vikas Sindhwani, et al. Socratic Models: Composing Zero-Shot Multimodal Reasoning with Language. In *International Conference on Learning Representations*, 2023.

Xianchao Zeng, Xinyu Zhou, Youcheng Li, Jiayou Shi, Tianle Li, Liangming Chen, Lei Ren, and Yong-Lu Li. Diagnose, Correct, and Learn from Manipulation Failures via Visual Symbols. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026a.

Yihan Zeng, Yisu Li, Kaixin Huang, Marc WuDunn, Amelia Kovacs, Mikael Harseno, Kimmy Lin, Porter Hunley, Fred Shentu, Philipp Wu, et al. Learning from physical Experience: Context-Space adaptation for robotic manipulation, 2026b.

Andy Zhai, Brae Liu, Bruno Fang, Chalse Cai, Ellie Ma, Ethan Yin, Hao Wang, Hugo Zhou, James Wang, Lights Shi, et al. Igniting VLMs toward the Embodied Space, 2025.

Simon Sinong Zhan, Philip Wang, Yao Liu, Yiyan Peng, Zinan Wang, Qineng Wang, Zhian Ruan, Xiangyu Shi, Xinyu Cao, Frank Yang, et al. SENTINEL: A Multi-Level Formal Framework for Safety Evaluation of Foundation Model-based Embodied Agents, 2025.

Jianshu Zhang, Keliang Wu, Chengxuan Qian, Xiyuan Yang, Ce Zhang, Ariel Tian, Anbang Liu, Haoran Lu, and Han Liu. ProgressCompass: Embodied Progress Reward Models Are Lost Without the Right Context, 2026a.

Shiduo Zhang, Zhe Xu, Peiju Liu, Xiaopeng Yu, Yuan Li, Qinghui Gao, Zhaoye Fei, Zhangyue Yin, Zuxuan Wu, Yu-Gang Jiang, et al. VLABench: A Large-Scale Benchmark for Language-Conditioned Robotics Manipulation with Long-Horizon Reasoning Tasks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025a.

Tao Zhang, Kaixian Qu, Zhibin Li, Jiajun Wu, Marco Hutter, Manling Li, and Fan Shi. Using Large Language Models for Embodied Planning Introduces Systematic Safety Risks, 2026b.

Wenbo Zhang, Kaixuan Wang, Yutao Ouyang, Xiaoyu Huang, Liyang Li, Kailun Su, Weiyang Jin, Wenhao Chai, Haotian Liang, Zhiyang Dou, et al. An Unexpected Robot Policy: Early Evaluations of GPT-6 Astra on RoboDojo and Beyond, 2026c.

Xiangcheng Zhang, Runhan Huang, and Yilun Du. World Action Planner: Generalizable Robot Decision-Making with Action-Conditioned World Models, 2026d.

Xilun Zhang, Maggie Wang, Erik Bauer, Hong-Xing Yu, Huang Huang, Jiajun Wu, and Marco Pavone. Co-Evolving Robot Orchestrators and Policies through Deployment, 2026e.

Yi Zhang, Che Liu, Xiancong Ren, Hanchu Ni, Shuai Zhang, Zeyuan Ding, Jiayu Hu, Hanzhe Shan, Zhenwei Niu, Zhaoyang Liu, et al. Pelican-VL 1.0: A Foundation Brain Model for Embodied Intelligence, 2025b.

Yixian Zhang, Huanming Zhang, Feng Gao, Xiao Li, Zhihao Liu, Yi Nie, Chunyang Zhu, Jiaxing Qiu, Yuchen Yan, Jiyuan Liu, et al. Harness VLA: Steering Frozen VLAs into Reliable Manipulation Primitives via Memory-Guided Agents, 2026f.

Jingjiao Zhao, Dong Jing, Mingyu Cao, Ruochuan Shi, Xiansheng Chen, Pengwei Wang, and Zhiwu Lu. EventSync: Event-Driven memory updates for History-Dependent manipulation, 2026a.

Xunyi Zhao, Jian Zhou, Sihao Lin, Gengze Zhou, Zerui Li, Xinyu Yan, Jiajun Liu, Anton van den Hengel, and Qi Wu. NavHarness: Towards Lifelong Embodied Navigation, 2026b.

Ruijie Zheng, Dantong Niu, Yuqi Xie, Jing Wang, Mengda Xu, Yunfan Jiang, Fernando Castañeda, Fengyuan Hu, You Liang Tan, Letian Fu, et al. EgoScale: Scaling Dexterous Manipulation with Diverse Egocentric Human Data, 2026.

Enshen Zhou, Jingkun An, Cheng Chi, Yi Han, Shanyu Rong, Chi Zhang, Pengwei Wang, Zhongyuan Wang, Tiejun Huang, Lu Sheng, et al. RoboRefer: Towards Spatial Referring with Reasoning in Vision-Language Models for Robotics. In *Advances in Neural Information Processing Systems*, 2025a.

Hai Zhou, Fan Yang, Zeyi Chen, Lejun Zhang, Chenyang Fu, Jianping Li, Jiayi Zhu, Xianyang Liu, Yingxuan Yang, Lili Meng, et al. Who Improves the Robot? Toward Physical Recursive Self-Improvement, 2026a.

Weijie Zhou, Manli Tao, Chaoyang Zhao, Haiyun Guo, Honghui Dong, Ming Tang, and Jinqiao Wang. PhysVLM: Enabling Visual Language Models to Understand Robotic Physical Reachability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025b.

Xueyang Zhou, Yangming Xu, Guiyao Tie, Yongchao Chen, Guowen Zhang, Duanfeng Chu, Pan Zhou, and Lichao Sun. LIBERO-PRO: Towards Robust and Fair Evaluation of Vision-Language-Action Models Beyond Memorization, 2025c.

Zewei Zhou, Rachel Luo, Yulong Cao, Chaowei Xiao, Chensheng Peng, Boyi Li, Thomas Tian, Zheng Lian, Yan Wang, Jiaqi Ma, et al. VeriFine: Scaling Verification for Self-Improvement in Embodied Reasoning, 2026b.

Zhiyuan Zhou, Pranav Atreya, Abraham Lee, Homer Rich Walke, Oier Mees, and Sergey Levine. Autonomous Improvement of Instruction Following Skills via Foundation Models. In *Proceedings of The 8th Conference on Robot Learning*, 2025d.

Zhiyuan Zhou, Pranav Atreya, You Liang Tan, Karl Pertsch, and Sergey Levine. AutoEval: Autonomous Evaluation of Generalist Robot Manipulation Policies in the Real World. In *Proceedings of The 9th Conference on Robot Learning (CoRL)*, 2025e.

Linsen Zhu and Mengqing Cai. From Language Models to World-Acting Systems: Progress and Limits of Agentic AI across Digital, Social, Virtual, and Physical Environments, 2026.

Xiaomeng Zhu, Fengming Zhu, Weijie Zhou, Ye Tian, Zhenlin Hu, Yufei Huang, Yuchun Guo, Xinyu Wu, Zhengyou Zhang, Fangzhen Lin, et al. ProAct: A Benchmark and Multimodal Framework for Structure-Aware Proactive Response. In *International Conference on Machine Learning (ICML)*, 2026.