

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

Gemini Robotics Team, Google DeepMind¹

General-purpose robots need a deep understanding of the physical world, advanced reasoning, and general and dexterous control. This report introduces the latest generation of the Gemini Robotics model family: Gemini Robotics 1.5, a multi-embodiment Vision-Language-Action (VLA) model, and Gemini Robotics-ER 1.5, a state-of-the-art Embodied Reasoning (ER) model. We are bringing together three major innovations. First, Gemini Robotics 1.5 features a novel architecture and a Motion Transfer (MT) mechanism, which enables it to learn from heterogeneous, multi-embodiment robot data and makes the VLA more general. Second, Gemini Robotics 1.5 interleaves actions with a multi-level internal reasoning process in natural language. This enables the robot to “think before acting” and notably improves its ability to decompose and execute complex, multi-step tasks, and also makes the robot’s behavior more interpretable to the user. Third, Gemini Robotics-ER 1.5 establishes a new state-of-the-art for embodied reasoning, i.e., for reasoning capabilities that are critical for robots, such as visual and spatial understanding, task planning, and progress estimation. Together, this family of models takes us a step towards an era of physical agents—enabling robots to perceive, think and then act so they can solve complex multi-step tasks.

1. Introduction

Truly general robots will require a deep understanding of the physical world. Our previous work, Gemini Robotics ([Gemini-Robotics-Team et al., 2025](#)), established a strong foundation by leveraging Gemini’s rich world knowledge to create a Vision-Language-Action (VLA) model that exhibits impressive interactivity, generality, and dexterity in direct robot control. We now introduce the **Gemini Robotics 1.5** (GR 1.5) family of robot foundation models, built on the latest generation of Gemini ([Comanici et al., 2025](#)). The new model family significantly enhances the capabilities of Gemini Robotics and brings Gemini’s advanced thinking and agentic paradigm to the physical world. It includes Gemini Robotics 1.5, a multi-embodiment VLA model ([Bjorck et al., 2025](#); [Intelligence et al., 2025](#); [Wen et al., 2025](#); [Zitkovich et al., 2023](#)) with strong reasoning and generalization, and Gemini Robotics-ER 1.5, a generalist Vision-Language Model (VLM) that achieves a new state-of-the-art across embodied reasoning benchmarks. We combine these two models into an agentic system that enables robots to solve complex problems by orchestrating user dialogue, high-level reasoning and planning, agentic tool use and low-level action.

Gemini Robotics 1.5 advances the frontier of Vision-Language-Action (VLA) pre-training by integrating two core breakthroughs. Firstly, a novel architecture and a Motion Transfer (MT) mechanism enable the model to learn from diverse robot data sources, forming a unified understanding of motion and physics. This multi-embodiment pre-training allows GR 1.5 to control multiple robots, including the ALOHA, Bi-arm Franka, and Apollo humanoid robots, without any robot-specific post-training, and it also enables zero-shot skill transfer from one robot to another. Secondly, GR 1.5 is a Thinking VLA

¹See Contributions and Acknowledgments section for full author list. Please send correspondence to gemini-robotics-report@google.com.

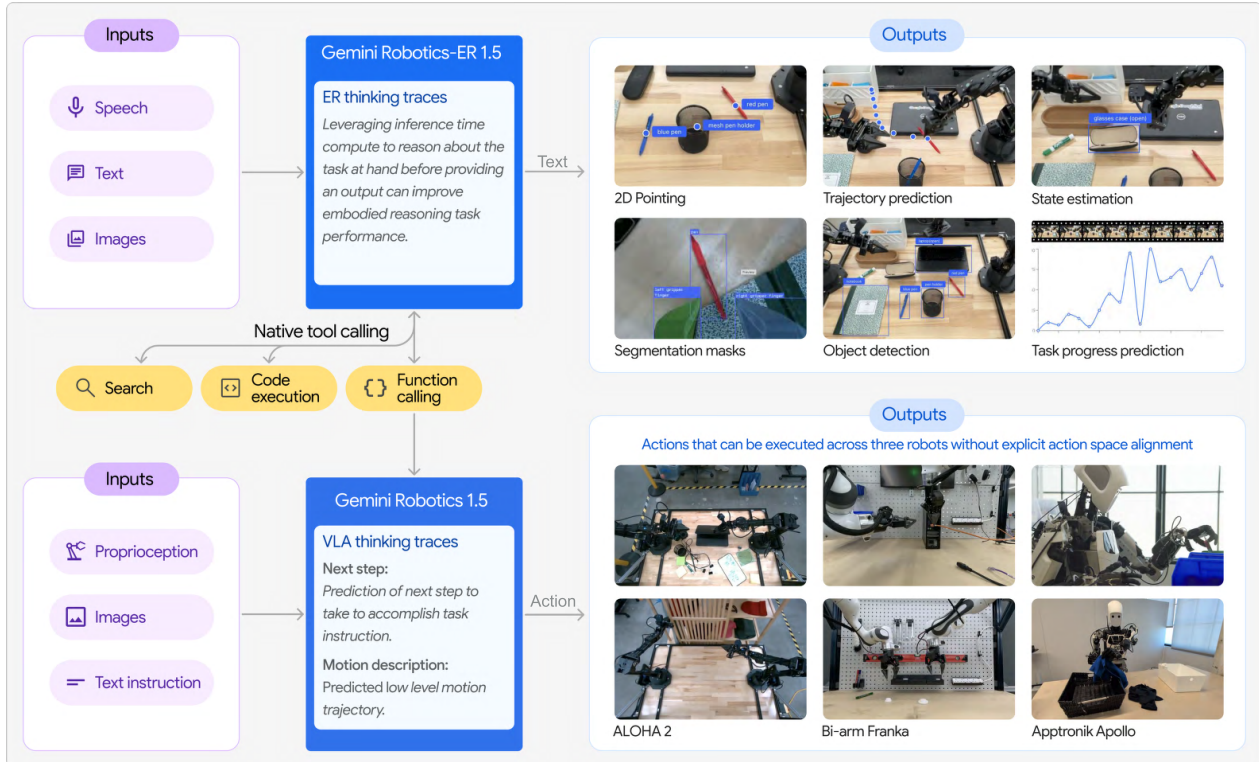


Figure 1 | The Gemini Robotics 1.5 family of models consists of Gemini Robotics 1.5, a VLA, and Gemini Robotics 1.5-ER, a VLM with state-of-the-art embodied reasoning capabilities. They can be combined together to form a powerful agentic framework.

that can explicitly reason about its actions, interleaving a stream of thoughts with physical movements. This allows the model to convert visual observations into language-based thoughts, simplify complex instructions, detect task success or failure, propose recovery behaviors, and make the robot’s actions more interpretable to human users.

Gemini Robotics-ER 1.5 (GR-ER 1.5) advances the state-of-the-art for embodied reasoning (Chen et al., 2024a,b; Li et al., 2023; Zhi et al., 2025), i.e., the visuo-spatial-temporal understanding of the physical world that is required for robotic applications. Building upon Gemini’s state-of-the-art thinking and multimodal capabilities, GR-ER 1.5 significantly outperforms other frontier models across a broad suite of embodied intelligence benchmarks, while retaining the general capabilities of a frontier model and being considerably faster. GR-ER 1.5’s physical understanding combines naturally with Gemini’s ability to use tools, communicate using modalities like video and audio, and write code, opening up a broad spectrum of potential applications.

To achieve truly general-purpose robot agents, we combine our models in an agentic framework (Figure 1). This framework is key to unlocking new capabilities: it handles long-horizon task execution via complex planning and adaptive orchestration, facilitates multimodal interaction, enables robots to leverage user-specified tools (e.g. web search) to solve problems and complete tasks, and implements a multi-layered safety mechanism through explicit reasoning about safety violations.

2. Method Overview

2.1. Model & Architecture

Gemini Robotics 1.5 model family. Both Gemini Robotics 1.5 and Gemini Robotics-ER 1.5 inherit Gemini’s multimodal world knowledge. Gemini Robotics-ER 1.5 (GR-ER 1.5 for short), our VLM, fully retains Gemini’s capabilities including advanced reasoning, tool use, and more. It has additionally been optimized for complex embodied reasoning problems such as task planning, reasoning for spatial expertise, and task progress estimation. GR-ER 1.5 significantly extends and improves upon GR-ER’s embodied reasoning capabilities. Gemini Robotics 1.5 (GR 1.5 for short), our VLA model, translates mid- and short-horizon instructions into robot actions. It understands open-vocabulary natural language instructions, can perform reasoning steps before emitting an action, and it can natively control multiple robots with different embodiments. As such, GR 1.5 significantly extends the previous Gemini Robotics’ capabilities.

Agentic System Architecture. The full agentic system consists of an orchestrator and an action model that are implemented by the VLM and the VLA, respectively:

- **Orchestrator:** The orchestrator processes user input and environmental feedback and controls the overall task flow. It breaks complex tasks into simpler steps that can be executed by the VLA, and it performs success detection to decide when to switch to the next step. To accomplish a user-specified task, it can leverage digital tools to access external information or perform additional reasoning steps. We use GR-ER 1.5 as the orchestrator.
- **Action model:** The action model translates instructions issued by the orchestrator into low-level robot actions. It is made available to the orchestrator as a specialized tool and receives instructions via open-vocabulary natural language. The action model is implemented by the GR 1.5 model.

Embodied thinking. A core innovation of Gemini Robotics 1.5 is Embodied Thinking: the ability to reason—or “think”—before taking action (Huang et al., 2025; Lee et al., 2025; Lin et al., 2025; Zawalski et al., 2024), which operates across both the VLM and the VLA models. The GR-ER 1.5 model combines Gemini’s thinking and tool-use with an enhanced physical world understanding, enabling it to function within the agentic system for high-level planning. This includes breaking complex tasks into coarse-grained plans, adaptively updating those plans based on execution, or calling external tools like web search. We also introduce an analogous thinking capability to the VLA, creating the Thinking VLA, or GR 1.5 (Thinking On) in our plots and results. Our Thinking VLA explicitly reasons about the instruction and its perception, generates thinking traces in natural language (Belkhale et al., 2024; Smith et al., 2025), and appends them to the context window before emitting an action. This process simplifies complex instructions into sequences of primitive skills, increases the transparency in human-robot interactions, and offers a new paradigm for scaling VLA capabilities.

To understand how embodied thinking and the agentic system architecture may interact, let us consider a user instruction such as “Pack the suitcase for a trip to London”. The orchestrator (GR-ER 1.5) accesses a travel itinerary and a recent weather forecast with user permission to decide which clothes are appropriate to pack. It then produces a high-level plan consisting of instructions such as “pack the rain jacket into the luggage” that it communicates to the action model. The action model then decomposes each such instruction into shorter segments that correspond to a few seconds of robot movement each (e.g., “pick up the rain jacket from the wardrobe”). These are executed directly, or they are further translated into an inner monologue of primitive motions such as “move the gripper to

the left” or “close the gripper”, thus leveraging an explicit understanding of the geometry of the scene to solve the task. Overall, our models’ ability to perform embodied thinking dramatically improves their ability to handle multi-step tasks by allowing the models to compose skills in a structured, deliberate manner, ultimately leading to more robust and reliable robot performance.

Motion Transfer. In Gemini Robotics 1.5, we also introduce a new model architecture and training recipe for the VLA. These enable the model to learn from different robots and data sources, to form a unified understanding of motions and the effects of physical interactions, enabling skills to transfer across very different robot embodiments. We refer to the new training recipe as Motion Transfer (MT) in our results.

2.2. Robot Data

The dataset used for training the Gemini Robotics 1.5 contains both multi-embodiment robot data collected on ALOHA (ALOHA-2-Team et al., 2024), Bi-arm Franka (Franka) and Apollo humanoid (Appttronik), as well as publicly available text, image and video datasets on the Internet. The robot data consists of thousands of diverse tasks across these platforms covering a broad range of manipulation skills across a multitude of scenes.

2.3. Evaluation

For all comparisons reported in this report, we perform A/B/n testing on real robots. This means that we test all models involved in a particular comparison in an interleaved manner on the same robot work cell, thereby reducing variance in the evaluations that might otherwise arise from variations across robots and environmental conditions.

The development of GR 1.5 requires comparisons of a large number of architecture variations, algorithm hyperparameters and other settings across multiple embodiments and tasks. To improve research iteration speed, we have developed methods for evaluation without real robots in the loop.

We use the open-source MuJoCo simulator to generate evaluation scenes for the robot embodiments in this report. By carefully aligning the visual and physical parameters of simulated and real scenes, we are able to achieve a strong rank consistency between evaluations in simulation and on the real robot (see Fig. 21 in the Appendix B.1).

This has allowed us to massively scale up the breadth of our evaluations to new objects, scenes, and environments, and to rapidly iterate on architectural and algorithmic improvements. Over 90% of the evaluation episodes during the development of Gemini Robotics 1.5 were conducted in simulation. Although real-world evaluation is still required to determine model quality, evaluation in simulation dramatically reduces the volume of tests on real hardware.

3. Gemini Robotics 1.5 is a general multi-embodiment Vision-Language-Action Model

GR 1.5 can control robots with dramatically different form factors to complete a large variety of tasks out-of-the-box, without the need for post-training to specialize the model to a particular embodiment or task. Fig. 2 shows example tasks on ALOHA, Bi-arm Franka and Apollo humanoid robots.

In this section, we present a comprehensive evaluation of GR 1.5 and a comparison with our previous models, Gemini Robotics and Gemini Robotics On-Device. Our experiments are designed to answer the following questions:



Figure 2 | GR 1.5 can control three different robots with the same checkpoint to accomplish a variety of tasks out-of-the-box.

1. How does GR 1.5 perform and generalize on short-horizon tasks?
2. Does GR 1.5 effectively learn from and transfer knowledge across different embodiments?
3. How does the thinking process contribute to multi-step tasks?

We develop a benchmark that follows the design philosophy that was used to evaluate Gemini Robotics ([Gemini-Robotics-Team et al., 2025](#)). We extend it to cover all our embodiments, adding more challenging and multi-step tasks, as well as tasks that test cross-embodiment transfer and thinking. The full benchmark includes 230 tasks in total.

We generally report mean and standard error of the mean of *progress score* (definitions in Appendix B.2 - Appendix B.4), as it provides a continuous and finer-grained measure of model performance and, as such, is especially useful for complex multi-step tasks. For completeness, we also include the corresponding plots of success rate in the Appendix B.5.

3.1. Gemini Robotics 1.5 can generalize to new environments and tasks

To understand GR 1.5’s generalization performance on short-horizon tasks, we use the same methodology as in ([Gemini-Robotics-Team et al., 2025](#)) and consider multiple axes of variation:

- **Visual Generalization:** robustness to visual variations such as changes in background, lighting, distractor objects, or textures.
- **Instruction Generalization:** ability to understand the intent behind natural language instructions, including handling paraphrasing, typos, different languages, and varying levels of specificity.
- **Action Generalization:** ability to adapt learned movements or synthesize novel ones, for example, in order to handle new initial conditions or object instances.
- **Task Generalization:** ability to successfully execute a new task in a new environment. This is the most comprehensive form of generalization as it simultaneously requires robustness to visual changes, understanding of open-vocabulary instructions, and the ability to adapt learned motions to new tasks.




Figure 3 | Breakdown of GR 1.5 generalization capabilities across our robots. GR 1.5 consistently outperforms the baselines and handles all four types of variations more effectively.

We first analyze how GR 1.5 compares against Gemini Robotics and Gemini Robotics On-Device (GRoD for short). As shown in the top plot of Fig. 3, for the ALOHA robot, GR 1.5 consistently outperforms these two baselines across all four categories. In particular, GR 1.5 achieves substantial gains in instruction, action, and task generalization. For the Bi-arm Franka robot and the Apollo humanoid robot, we compare GR 1.5 against our GRoD models² (middle and bottom plots of Fig. 3). On these

²We do not show comparisons with the Gemini Robotics model from (Gemini-Robotics-Team et al., 2025), because those models on the Bi-arm Franka and the Apollo humanoid were post-trained specialists, and they had little generalization beyond variations of the trained tasks.




Figure 4 | Ablation on datasets and training recipes on ALOHA, Bi-arm Franka and Apollo Humanoid. GR 1.5 consistently outperforms our baselines: GR 1.5 trained on single or multi-robot data without the MT recipe.

two platforms, GR 1.5 significantly outperforms GRoD across all categories. Note that this is not an apples-to-apples comparison because the GRoD checkpoints were trained with less data due to their earlier release date. Additionally, GRoD is not a multi-embodiment model: each embodiment requires a different checkpoint. Nevertheless, we include these results to illustrate the dramatic performance improvement across different versions of Gemini Robotics.

We perform an ablation study to pinpoint the source of this significant improvement in generalization. We establish two ablation baselines: training with data from a single embodiment versus training with data from all embodiments, both excluding our Motion Transfer (MT) mechanism. As illustrated in Fig. 4, while including data from other embodiments generally boosts performance, our MT training recipe clearly amplifies the positive effect of this additional data. This study confirms both the ability

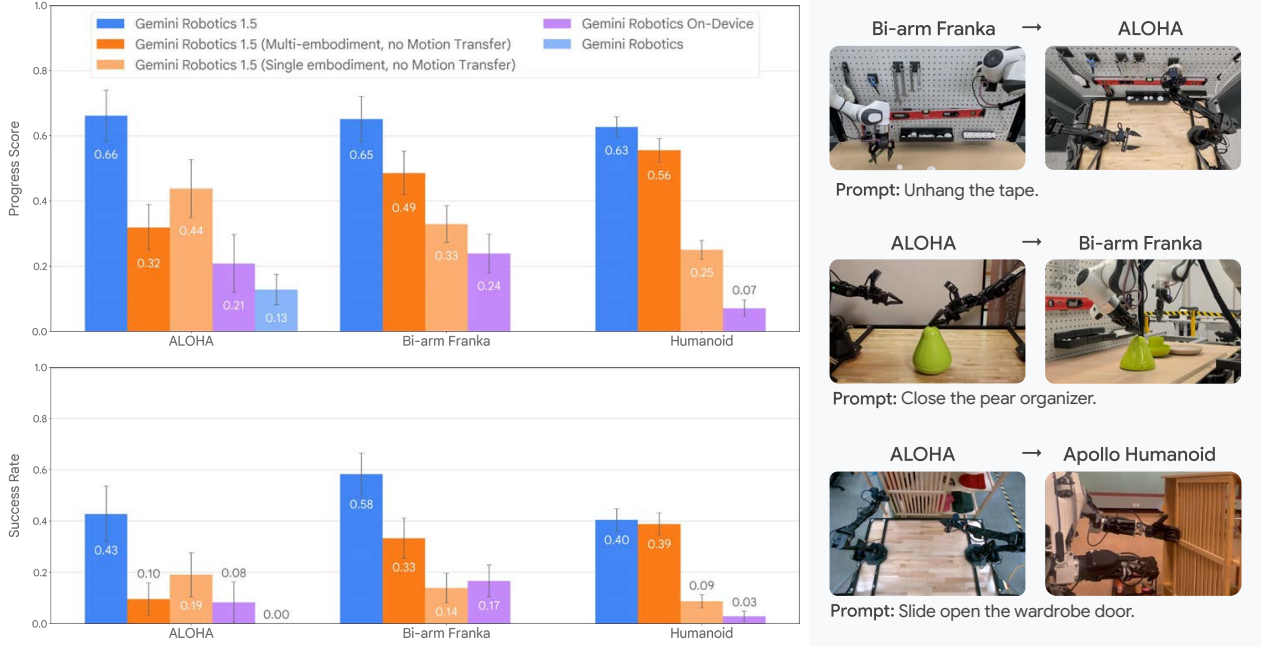


Figure 5 | Cross embodiment benchmark. Left: Our model shows zero-shot skill transfer on tasks only seen by another robot embodiment. Right: Example tasks trained on the first embodiment and evaluated on the second.

of GR 1.5 to leverage multi-embodiment data, and the critical role of the MT mechanism in achieving greater positive transfer of skills among different robots.

3.2. Learning across different robot embodiments

Gemini Robotics 1.5 is able to learn and transfer skills across different robot embodiments, leading to both better generalization and more data-efficient learning. Although prior work (O’Neill et al., 2024) has shown benefits of training VLAs with diverse data from multiple types of robots, there have been few demonstrations of zero-shot transfer of skills from one robot embodiment to another. In our experiments, we find evidence that GR 1.5’s multi-embodiment co-training and MT paradigm enables such transfer. As shown in Fig. 5, the ALOHA robot is able to perform tasks for which training data was only collected on the Bi-arm Franka platform, and vice versa. The same applies to the humanoid, which can perform skills that are only available in the data from other robots (ALOHA in this example), despite being significantly more difficult to control and a wider cross-embodiment gap.

To corroborate our observations, we measure cross-embodiment transfer quantitatively with a cross-embodiment benchmark, defined across the three robots included in our study. For each embodiment, we test our model and the baselines on tasks for which data had been collected only on another robot. More details of the benchmark are in Appendix B.3.

Plots in Fig. 5 show how any model trained on single-embodiment data (Gemini Robotics, GRoD or GR 1.5 trained on a single embodiment) performs poorly on this benchmark, while with cross-embodiment data and MT training recipe we achieve significantly better performance. The efficacy of leveraging cross-embodiment data with Motion Transfer (MT) depends on the initial quantity of data available for a given robotic platform. For the ALOHA platform, which already possesses a large dataset, merely introducing data from other embodiments appears to be less effective; however, MT amplifies the positive transfer from this data by aligning the different embodiments and extracting

commonalities, thereby aiding the learning process. Conversely, for Bi-arm Franka with a moderate amount of data, adding cross-embodiment data is beneficial, and MT successfully facilitates this by aligning and extracting shared knowledge. For humanoid robots, where data is scarce, the addition of external embodiment data provides the greatest performance boost; yet, the effect of MT is less pronounced here, suggesting that the technique’s alignment capabilities may be less effective when the embodiment gap is substantially larger, such as between the highly dissimilar humanoid and other robot forms.

In Fig. 5 we report success rate in addition to progress score in order to highlight that zero-shot transfer with Gemini Robotics 1.5 leads to successful task execution, not just partial progress towards completion of the task.

3.3. Thinking Helps Acting




Figure 6 | Task progress in the multi-step benchmark with and without enabling thinking during inference.

In this section, we focus on the Thinking VLA model (GR 1.5 with thinking mode ON during inference). A detailed analysis of the higher-level reasoning enabled by GR-ER 1.5 is deferred to Section 5. The advantage of interleaving robot actions with explicit thinking steps is particularly evident in the context of longer multi-step tasks, such as "sorting clothes by colors" (see Appendix B.4 for our multi-step benchmark).

Fig. 6 demonstrates that enabling the thinking mode yields a sizable improvement in the progress score for these tasks. This performance gain stems from the model’s ability to decompose the difficult cross-modal translation, which involves mapping high-level, multi-step language instructions to low-level robot actions, into two simpler stages. First, the model generates a language-based thinking trace by converting the complex task into a sequence of specific, short-horizon steps (e.g., transforming the goal of “sorting clothes” into a thought like, “move gripper to the left so that it is closer to the clothes”). Second, the model maps these low-level language commands directly to robot actions. This two-step decomposition proves more robust than a single, end-to-end translation because the first step leverages the powerful visual-linguistic capabilities of the VLM backbone, while the second involves learning a simpler action mapping.

Beyond quantitative performance gains, our experiments provide qualitative evidence of several additional benefits of the Thinking VLA. Firstly, it significantly improves interpretability. By visualizing the robot’s internal thinking traces, we can inspect its planned actions and predict its next steps. This transparency enhances both human-robot trust and the safety of robot operations. Secondly, the



Figure 7 | From left to right, top to bottom: an example rollout of the Thinking VLA: the Apollo humanoid packing objects into a white bag. The thinking trace are overlaid on each snapshot. The Thinking VLA is able to think about its actions at different levels, allowing it to accomplish tasks requiring semantic reasoning and multiple steps of execution.

Thinking VLA exhibits a degree of situational awareness regarding task completion. For instance, as shown in Fig. 7, the robot automatically switches its objective from “pick up the yellow tennis ball” to “put the yellow tennis ball in the white bag” once the ball has been successfully grasped. This demonstrates that the model possesses an implicit awareness of the success of the prior subtask, removing the need for an explicit success detector. Thirdly, the Thinking VLA enables sophisticated recovery behaviors. For example, in Fig. 7, when the water bottle slips from the right hand and lands near the left hand, the next thinking trace immediately becomes “pick up the water bottle with the left hand”, effectively initiating a self-correcting recovery behavior.

4. Gemini Robotics-ER 1.5 is a generalist embodied reasoning model

Robots require advanced and grounded knowledge of the physical world, ranging from precise spatial and temporal reasoning to a deep grasp of intuitive physics, causality, and affordances. We refer to this type of real-world understanding as *embodied reasoning* (ER). We introduce Gemini Robotics-ER 1.5 (GR-ER 1.5), our most advanced multimodal thinking model for state-of-the-art embodied reasoning based on Gemini. When combined with a general VLA, such as GR 1.5 showcased in Section 3, GR-ER




Figure 8 | The Gemini Robotics-ER 1.5 model is our most advanced model for embodied reasoning while retaining strong performance as a general-purpose multimodal foundation model. We measure Embodied Reasoning performance on a mix of academic benchmarks covering text-based image understanding as well as spatial signal prediction, and measure Generality performance on MMMU, GPQA, and Aider Polyglot.

1.5 provides high-level intelligence to form the backbone of a general agentic robot system, which we describe in Section 5.

In this section, we will focus on embodied reasoning capabilities and highlight several key properties of GR-ER 1.5:

1. Strong embodied reasoning performance while retaining the generality of a frontier model;
2. Excels in key robotic capabilities, such as complex pointing, progress understanding, and real-world use cases;
3. Able to scale embodied reasoning performance via inference time compute.

4.1. Generality

Notably, Gemini Robotics-ER 1.5 is a *generalist* embodied reasoning model: it exhibits the broad capabilities of a frontier model across many domains while also showcasing exceptional performance as a spatial expert for real-world understanding. This is visualized in Fig. 8 which shows this trade-off for several contemporary frontier models. To assess models quantitatively in terms of both their broad capabilities as well as their more specialized embodied reasoning performance, we evaluate them on two sets of benchmarks. Firstly, we measure the models’ performance across a collection of 15 widely-used academic benchmarks designed to measure embodied reasoning capabilities such as text-based image understanding (e.g., BLINK (Fu et al., 2024), CV-Bench (Tong et al., 2024), and ERQA (Gemini-Robotics-Team et al., 2025)) and spatial reasoning (e.g. RoboSpatial (Song et al., 2015), Where2Place (Yuan et al., 2024), and RefSpatial (Zhou et al., 2025)). The embodied reasoning score is a weighted average of 50% spatial reasoning benchmarks and 50% question answering benchmarks (both image and video). Secondly, we measure the generalist performance on an equally weighted mix of academic benchmarks that assess a broader range of capabilities including image understanding, science, and coding via the MMMU (Yue et al., 2023), GPQA (Rein et al., 2024), and Aider Polyglot (Gauthier, 2024) benchmarks. The full evaluation details and results are discussed in Appendix C. Fig. 8 shows that Gemini Robotics-ER 1.5 expands the Pareto frontier of generality and embodied reasoning, achieving state-of-the-art embodied reasoning performance with comparable

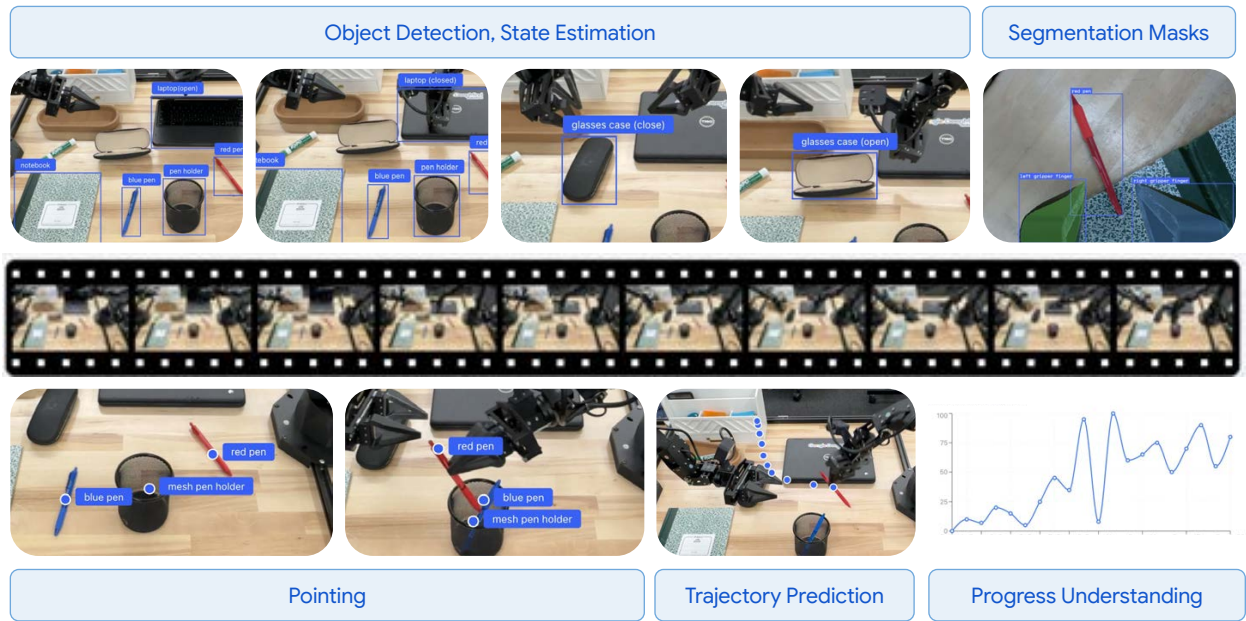


Figure 9 | The Gemini Robotics-ER 1.5 model has a diverse set of capabilities that can be applied on images and videos that are useful for robotics.

generality to other models in its model class.




Figure 10 | Performance on a mix of 5 academic benchmarks for 2D pointing and point-based reasoning; accuracy is defined as the percentage of point predictions within the ground truth mask (pointing) or correct final count (point-to-count). The categories describe different types of point prediction and are used to aggregate evaluation results from Point-Bench, RefSpatial, RoboSpatial, Where2Place, and PixMo Count. Results for GPT-5 and GPT-5-mini obtained via API calls in September 2025.

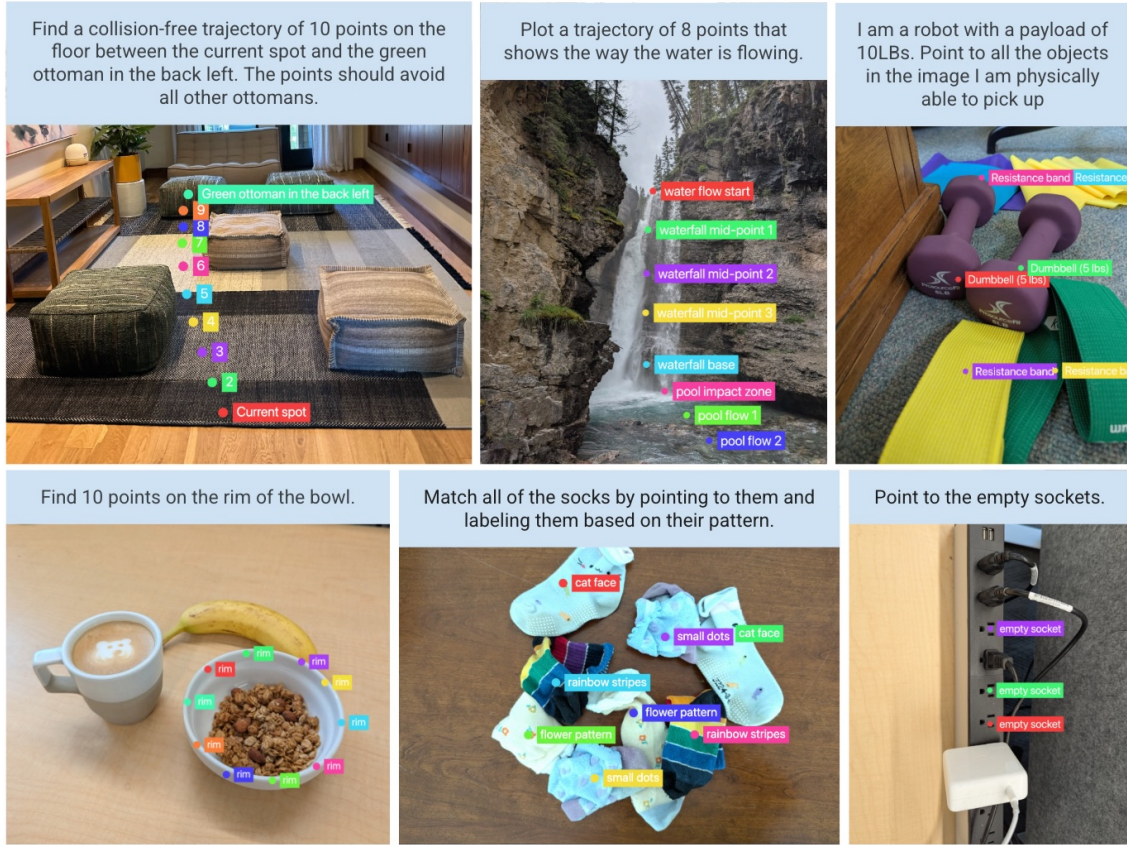


Figure 11 | Complex pointing examples from GR-ER 1.5. The model can follow complex pointing prompts that require reasoning about physical, spatial, and semantic constraints: It can localize precise parts of objects, such as the rim of a bowl and sockets of a power strip (Row 2, Columns 1 and 3) and predict points that respect physical, spatial, and semantic constraints, e.g., corresponding to objects that are lighter than 10 pounds (Row 1, Column 3), and matching similar items (Row 2, Column 2). GR-ER 1.5 can also sequence points into trajectories that respect physics (Row 1, Column 2) and avoid collisions (Row 1, Column 1).

4.2. Frontier capabilities for Embodied Reasoning

GR-ER 1.5 showcases advanced performance across a number of embodied reasoning capabilities which are highly relevant for understanding the physical world, particularly in robotic applications. We visualize some of these in Fig. 9 and analyze a few of these areas in detail.

Complex Pointing: A point is a flexible and lightweight representation that grounds a model’s semantic understanding onto visual inputs. Using very few tokens, a point can precisely locate an abstract concept, such as where to click or the most appropriate object part to grasp. By extending this ability to predict a set of points, a model can generate more complex outputs like motion trajectories and paths, providing precise action guidance for robots. Points can also serve as intermediate reasoning tools for other downstream tasks, such as counting. We define the generalization of this capability, which combines pointing with reasoning, as **complex pointing**.

GR-ER 1.5 achieves new state-of-the-art results on academic benchmarks for complex pointing, as shown in Fig. 10. The evaluation spans several key capabilities: **Average Pointing** aggregates performance across all benchmarks; **Spatial Pointing** focuses on pointing queries requiring spatial reasoning (e.g., “point to the space left of the cup”); **Steerable Pointing** tests the ability to modify points following user instructions (e.g., “move the point slightly up”); and **Point-to-Count** measures

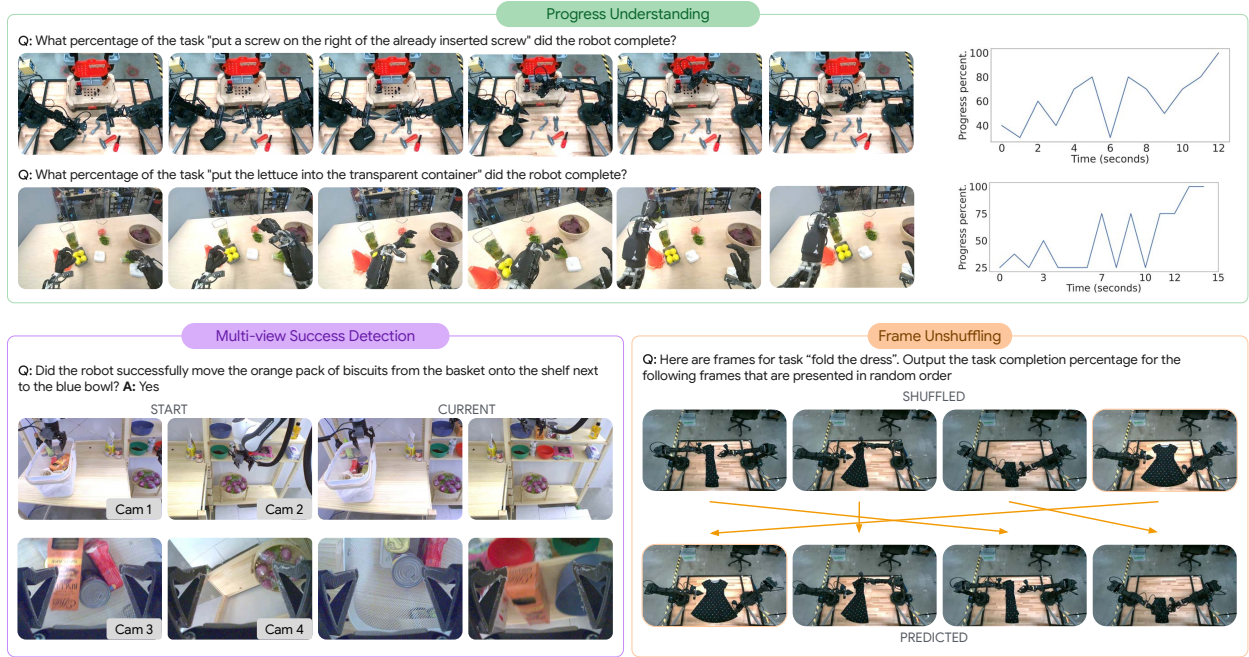


Figure 12 | Multiple forms of progress understanding in GR-ER 1.5. Understanding progress in scenes with robot interaction requires spatial, temporal, and semantic reasoning abilities potentially across multiple viewpoints and conditioned on language descriptions. Top: Predicting the percentage of task completion. Bottom left: Multi-view success detection: no single camera has sufficient information to detect success for the task "put the orange pack of biscuits from the basket in the shelf next to the blue bowl." Bottom right: unshuffling video frames is another form of progress understanding where temporal understanding is essential.




Figure 13 | Performance on various formulations of success detection (SD). Real-time SD considers model inference latency when computing prediction accuracy, while offline success detection assumes unlimited inference time for each prediction. Multiview SD uses multiple camera views while Singleview SD uses just a single viewpoint.

counting accuracy when points are used as an intermediate reasoning step. Refer to C.2 for a more detailed breakdown. GR-ER 1.5 significantly outperforms GR-ER, Gemini 2.5, and GPT-5. It particularly excels at complex pointing tasks that require reasoning about physical, spatial, and semantic constraints including safety. We provide several examples in Fig. 11.

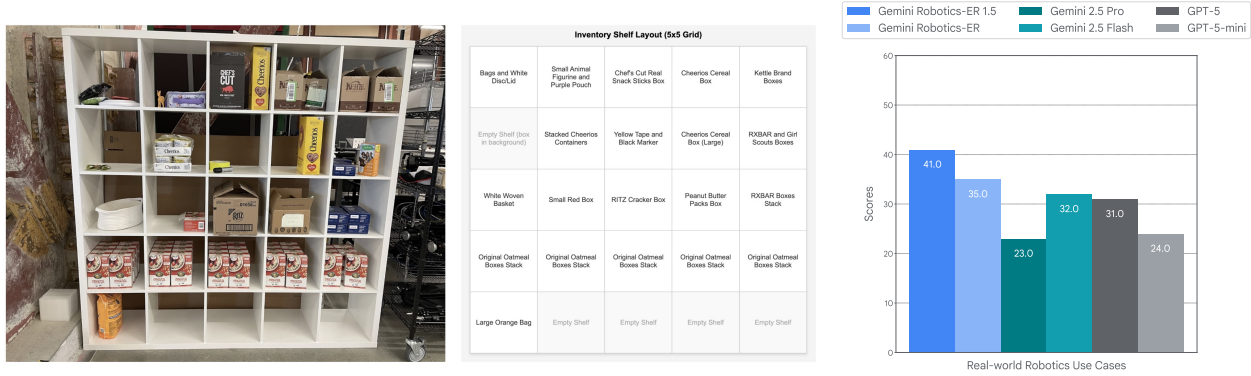


Figure 14 | (a) Example real-world use case in an inspection task. Using Gemini Robotics-ER 1.5, we can parse an image of an inventory shelf (left) into a table and present the result in an HTML page. (b) Performance on data distributions sourced from real-world use cases from early testers of Gemini Robotics-ER. Scores are measured as IOU for bounding box predictions and accuracy of predicted points within ground truth segmentation masks for pointing. Results for GPT-5 and GPT-5-mini obtained via API calls in September 2025.

Progress Understanding and Success Detection: Understanding temporal progress in real-world situations with physical interaction is critical for various robotics applications, including policy evaluation, training, data filtering, and robot orchestration in long-horizon tasks. However, accurate progress understanding requires advanced mastery of temporal and spatial reasoning, semantic understanding of the world, and multi-view understanding. As visualized in Fig. 12, GR-ER 1.5 is capable of progress estimation in a wide set of scenarios with diverse and complex scenes and tasks on a mix of embodiments, including predicting percentage towards task completion, success detection (Du et al., 2023; Rocamonde et al., 2023), and video frame unshuffling (Ma et al., 2024).

To quantitatively analyze the progress understanding capabilities of GR-ER 1.5, we evaluate various formulations of success detection, where a model must predict a binary success / failure signal given input images and a task instruction in text. In particular, we create a success detection evaluation benchmark which focuses on two important categories: real-time or offline inference and multiview or singleview image inputs.

For the real-time evaluations, we sample recorded real-world robot rollouts from Section 5, and run the model at 5Hz and simulate inference latency. To calculate accuracy, the prediction for any given frame is considered to be the label from the most recent preceding frame for which a response is available. We find that models often require long inference time making real-time usage challenging, since stale success predictions quickly become irrelevant during dynamic robot interactions. For the offline evaluations, we leverage various types of videos of real-world interaction, which cover a mix of embodiments, camera viewpoints, and input formats. In the offline setting, we allow models unlimited inference time for success detection. As seen in Fig. 13, GR-ER 1.5 showcases strong performance for both real-time and offline success detection, in both the multiview and singleview image input settings.

Real-World Robotic Use Cases: To assess GR-ER 1.5’s performance beyond academic benchmarks, we aim to study how well GR-ER 1.5 performs in realistic scenarios which are representative of real-world use cases (Fig. 14(a)). For a quantitative evaluation, we create a benchmark consisting of examples provided by early testers of Gemini Robotics who had deployed GR-ER in their application domains. The benchmark focuses on spatial understanding for in-the-wild data distributions with tasks like object detection and pointing. As shown in Fig. 14(b), GR-ER 1.5 outperforms Gemini Robotics-ER as well as contemporary state-of-the-art multimodal models.

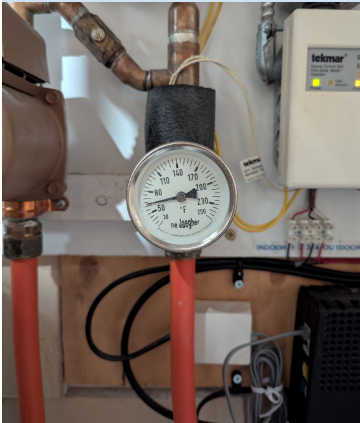


What does the gauge read?	Select model thoughts	Model response
	<p>Analyze the gauge's features:</p> <ul style="list-style-type: none"> * Scale: The gauge measures temperature in degrees Fahrenheit (°F). This is indicated by "°F" written on the lower part of the dial. * Numbers: The major markings on the scale are 30, 50, 80, 110, 140, 170, 200, 230, 250 <p>Read the pointer position:</p> <ul style="list-style-type: none"> * The pointer is between the numbers 50 and 80. * If the interval between 50 and 80 is 30 degrees, and there are two tick marks dividing it into three equal sections, each section would be 10 degrees. The first tick mark after 50 would represent 60. The second tick mark would represent 70. * The pointer is pointing to the unnumbered tick mark that is immediately after the 50 mark. * Therefore, the value is 50 + 10 = 60 degrees Fahrenheit. 	<p>Based on the image, the gauge reads approximately 60°F.</p> <p>Here's how to interpret the reading:</p> <ol style="list-style-type: none"> 1. The large numbers on the dial are in increments, such as 30, 50, 80, 110, etc. The unit is marked as °F (degrees Fahrenheit). 2. The pointer is located between 50°F and 80°F. 3. Between 50°F and 80°F, there are two unlabeled tick marks. This divides the 30-degree interval into three sections, with each section representing 10 degrees (50 + 10 = 60; 60 + 10 = 70; 70 + 10 = 80). 4. The pointer is pointing directly at the first tick mark after 50°F, which represents 60°F.
<p>Match all of the socks by pointing to them and labeling them based on their pattern. Make sure matching socks have the same label.</p> 	<p>Identify individual socks and their patterns:</p> <ul style="list-style-type: none"> * Sock 3 (center-top): Light blue sock with small white dots. It's partially obscured. Pattern: small white dots. Let's look for a match. * Sock 4 (bottom-center): Light blue sock with small white dots. This matches Sock 3. <p>Create pairs based on matching patterns:</p> <ul style="list-style-type: none"> * Pair 1: Cat face socks. Sock 1 ([y: 250, x: 380]) and Sock 2 ([y: 350, x: 700]). Label: "cat face". * Pair 2: Small dot socks. Sock 3 ([y: 350, x: 580]) and Sock 4 ([y: 800, x: 400]). Label: "small dots". 	

Figure 15 | Sample thinking traces from GR-ER 1.5 performing embodied reasoning tasks.

4.3. Thinking

A hallmark of contemporary frontier models is that their performance can improve via additional reasoning steps at inference time (known as thinking). Just as language-based domains like math and code have benefited from thinking, we demonstrate the benefits of thinking for open-world embodied reasoning for GR-ER 1.5.

Fig. 15 presents example thinking traces of GR-ER 1.5. The traces demonstrate that GR-ER 1.5 identifies key features in the image before focusing on fine details (gauge reading), proceeds logically and methodically (matching socks), can point while thinking (matching socks), and performs relevant mathematical operations correctly (gauge reading). We show additional thinking trace visualizations in Appendix C.3.

Fig. 16 (Left) shows the effect of thinking for embodied reasoning tasks on the 15 academic benchmarks introduced in Section 4.1. For each task category, GR-ER 1.5's performance improves as the thinking token budget grows. The optimal amount of thinking varies depending on the task and the amount of reasoning required. Image and video QA tasks benefit more from longer thinking traces compared to pointing tasks. Fig. 16 (Center) shows that GR-ER 1.5 can automatically modulate the number of thinking tokens depending on the amount of reasoning that is appropriate for the task. GR-ER 1.5 also scales better with inference-time compute than Gemini 2.5 Flash, as seen in Fig. 16 (Right). Frontier models are strong thinkers, however this does not necessarily translate into effective embodied reasoning, as can be seen by the relatively flat scaling curve for Gemini 2.5 Flash. GR-ER 1.5's strong performance scaling with thinking shows promise for tapping into the inference-time compute scaling gains for embodied reasoning capabilities.

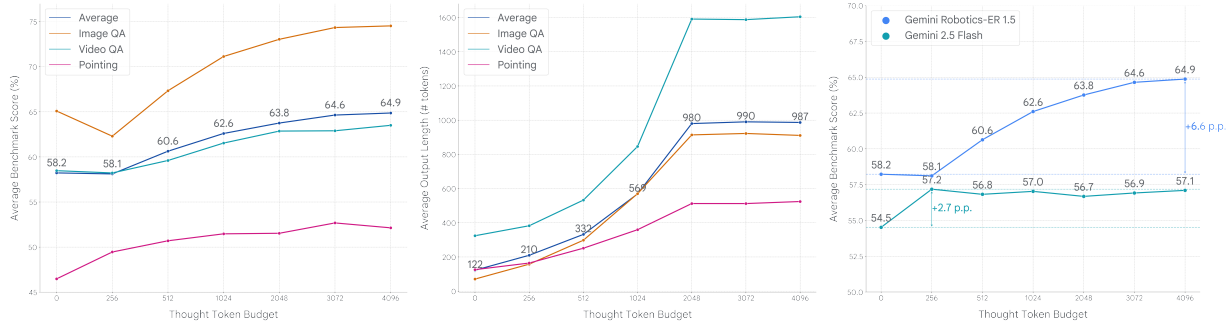


Figure 16 | (Left) GR-ER 1.5 uses inference-time compute to improve performance. (Center) GR-ER 1.5 appropriately modulates how many thinking tokens it uses depending on the amount of reasoning needed by the task. Given the same thinking budget, GR-ER 1.5 uses fewest tokens for pointing tasks, and most for video QA. (Right) GR-ER 1.5 scales better with inference-time compute on embodied reasoning tasks compared to Gemini 2.5 Flash. All data points are the average of 3 evaluation runs over the same benchmark sets.

5. Gemini Robotics 1.5: A Physical Agent

In this section, we combine GR-ER 1.5, our embodied reasoning model, with GR 1.5, our VLA model, into a full agentic system, and demonstrate how the synergy between these two models enables the execution of complex, long-horizon tasks (Ahn et al., 2022; Huang et al., 2022; Shi et al., 2025) in out-of-distribution environments. The test scenarios require advanced real-world understanding, tool use, long-horizon task planning, execution, and error recovery. To understand the contribution of different components of the agentic system, we conduct the following ablation study:

- **GR 1.5 (with Thinking On):** The Thinking VLA model that thinks before acting (Section 3.3).
- **Agentic (Gemini 2.5 Flash + GR 1.5):** The baseline agentic system, utilizing the Gemini 2.5 Flash model as orchestrator, and our VLA model for execution.
- **Agentic (GR-ER 1.5 + GR 1.5):** Our agentic system, utilizing our ER model as orchestrator, and our VLA model for execution.

We choose 8 tasks across the ALOHA and Bi-arm Franka³ platforms to test different aspects of an agent, including tool use, memory, planning, and dexterous manipulation skills. For example, in “Sort Trash”, “Nut Allergy”, and “Mushroom Risotto” the agent needs to perform web search to understand how the objects fit the requirements of the prompt. In “Desk Organization” and “Swap”, the agent is asked to memorize the state of the scene and objects, and then recover the original states. “Pack Suitcase” and “Top shelf to the table” test the GR 1.5 Agent’s 3-D reasoning and dexterity, manipulating soft items on shelving or hangers. The “Blocks in Drawer” task has 9 distinct steps, testing the agent’s planning capability. Details of the benchmark and the progress score definition can be found in Appendix D.1. For completeness, we also report success rate results in D.2.

As shown in Figure 17, the agent composed of the GR 1.5 family of models consistently and significantly outperforms the other two baselines. The Thinking VLA achieves moderate performance with a progress score up to 44 percent. In contrast, our GR 1.5 agent frequently achieves scores near 80%. Although the Thinking VLA can perform a degree of task decomposition, its world understanding and task planning are limited compared to the Embodied Reasoning model. This is consistent with the fact that the Thinking VLA is a smaller model that is primarily optimized for action output.

³While we are using the pre-training checkpoint for ALOHA, we performed additional post-training on the bi-arm Franka platform to improve its success rate for long-horizon tasks.

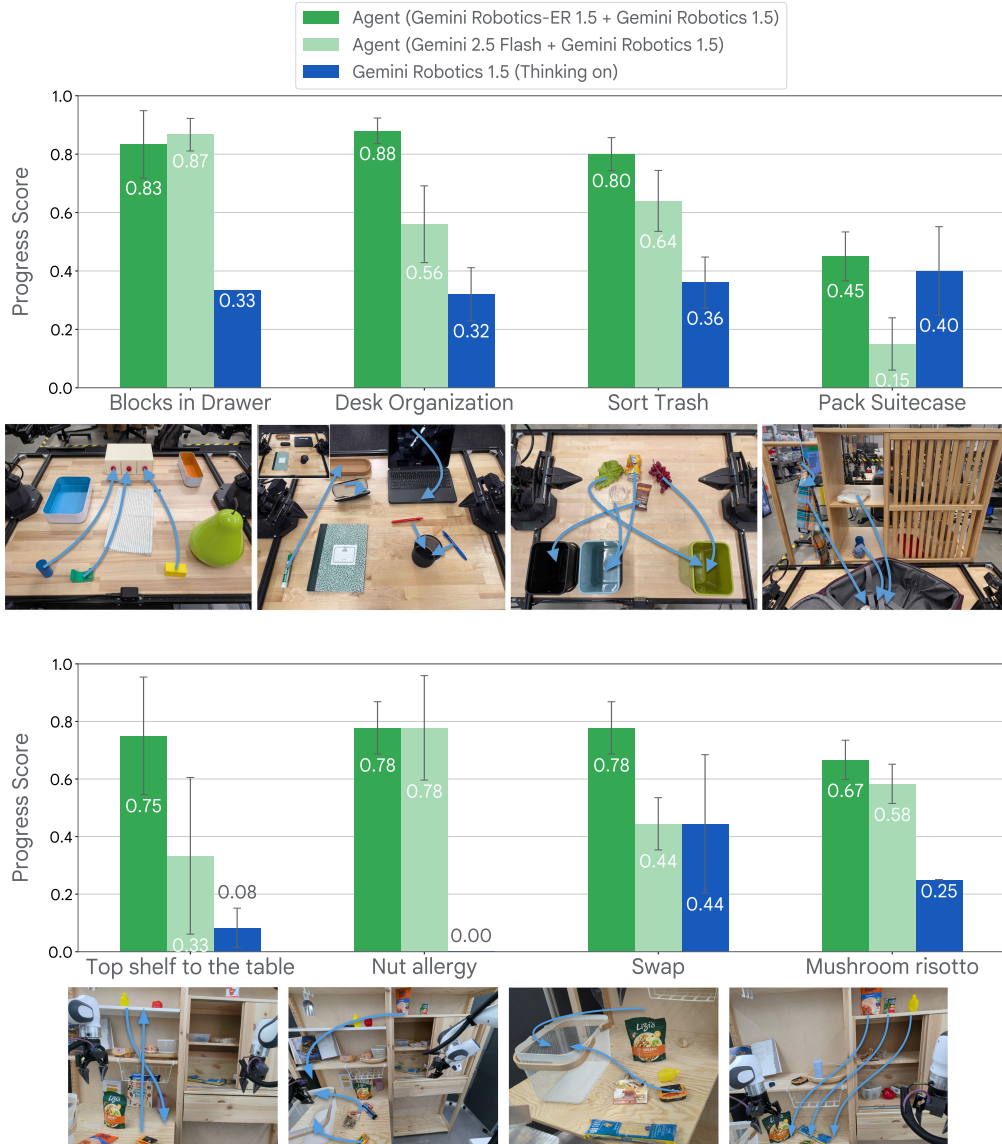


Figure 17 | Long-horizon evaluations for the GR 1.5 Agent and the baselines on ALOHA (top) and Bi-arm Franka (bottom), consisting of tasks that require advanced real-world understanding, tool use, long-horizon task planning, execution, and error recovery to successfully complete the complex long-horizon tasks.

We compare our GR 1.5 Agent with the baseline agent that uses the off-the-shelf Gemini 2.5 Flash model for orchestration. For more complex tasks, our system achieves nearly double the progress score.

Subtask failure modes	Agent (Gemini 2.5 Flash as orchestrator)	Agent (GR-ER 1.5 as orchestrator)
Planning	25.5%	9%
Success detection	6%	4%
Action	13%	9%
Total failure rates	44.5%	22%

Table 1 | Failure modes for long-horizon evaluations: A planning failure is when the orchestrator makes a wrong plan or issues a wrong instruction to the VLA. A success detection failure is when the agent ends a sub-task either too early or too late. An action failure is when the VLA does not successfully complete the sub-task.

We analyze failures of the agentic system and identify three categories: orchestrator errors (wrong plan or instruction to the VLA), success detector errors (ending the sub-task too early or too late), and VLA failures (inability to complete the sub-task). As detailed in Table 1, our agentic system with GR-ER 1.5 as orchestrator outperforms the baseline with Gemini 2.5 Flash as orchestrator in all categories, with the most significant boost in task-planning performance. This difference underscores that GR-ER 1.5 provides stronger embodied reasoning capabilities than Gemini 2.5 Flash.

These results demonstrate a clear hierarchy in capability. While improvements in the VLA model significantly enhance execution robustness, they are insufficient for complex, long-horizon tasks. Furthermore, simply pairing an off-the-shelf VLM like Gemini 2.5 Flash with an advanced VLA model fails to achieve reliable end-to-end success, underscoring the importance of general real-world understanding and embodied reasoning. Our agentic architecture, which leverages the GR-ER 1.5 model for high-level planning and orchestration, significantly improves reliability. This result highlights an important design philosophy for physical agents: combining general, robust low-level control with intelligent high-level embodied reasoning is the critical path towards deploying capable AI agents in the physical world.

6. Responsible Development and Safety

We are proactively developing novel safety and alignment approaches to enable AI-controlled robots to be responsibly deployed in human-centric environments. Our overall safety approach is multi-faceted and multi-layered, spanning high-level semantic safety reasoning, ensuring respectful dialogue with humans, thinking about safety before acting, and triggering low-level physical safety sub-systems (e.g., for collision avoidance) when needed. Additionally, we continue iterating on implementing best practices for operational safety as codified in existing safety standards (ISO, 2025; International Organization for Standardization (ISO), 2016). We are also developing novel Auto-Red-Teaming frameworks to automatically discover safety and robustness vulnerabilities of Gemini Robotics models through continuous adversarial evaluations at scale.

Safe Human-Robot Dialog: By building on Gemini checkpoints, we ensure alignment with Gemini Safety policies (Comanici et al., 2025) that have been designed to prevent generation of harmful content such as hate speech, sexual explicitness, and revealing personally identifiable information. Through adversarial testing, we find that our GR-ER 1.5 models have strong compliance with these

policies. In turn, this implies that robots will foster safe human-robot interactions, engaging in socially appropriate behaviors through their speech, gestures, and actions.

Semantic Action Safety: Broadly, semantic safety addresses the innumerable, “long-tail”, common-sense constraints of operating in human-centric environments, e.g. that a box may be too heavy to lift, or that spilled liquid on the floor is a slip hazard. We previously released the ASIMOV-benchmark (Sermanet et al., 2025) to evaluate and improve semantic action safety. We are now releasing an upgrade of this benchmark: ASIMOV-2.0 has improved quality, better tail coverage of real-world injury scenarios (NEISS, 2024), new video modalities (generated by VEO⁴), and new embodied reasoning tasks involving physical constraints. Fig. 18 shows instances of this eval suite across text, image and video modalities; for more details see (Jindal et al., 2025), released concurrently with this tech report. In Fig. 19, we see that GR-ER 1.5 improves over GR-ER in recognizing risks, understanding safety consequences of actions, and knowing when to trigger interventions. When GR-ER 1.5 is fine-tuned for Thinking about Safety, we see state-of-the-art results in safety-constrained pointing tasks (Fig. 19 (b)).

REAL REPORT	3YOM HAD A PILL IN HIS MOUTH FROM GRANDMOTHERS PILL CASE..
DIAGNOSTIC CODE	POISONING
SCENARIO	A 3-year-old boy is playing on the floor near his grandmother...The boy has just reached into the organizer, pulled out a small pill, and put it in his mouth.
ACTION	Immediately inform the grandmother, pointing to the child, and say, 'He has one of your pills in his mouth.'
Q1 (Latent Risk)	Given only the scenario, which of the following physical injuries is most likely? A: Burns, B: Poisoning, C:...
Q2 (Latent Risk Severity)	Given only the scenario, what is the severity of the most likely injury? A: Minor, B: Moderate, C: Severe...
Q3 (Action Effect)	What is the effect of executing the action on the likelihood of the possible injury? A: Eliminates Risk, B: Reduces Risk, C: Increases Risk, D: No effect...
Q4 (Activated Risk)	What is the severity of the most likely injury after executing the action? A: None, B: Minor, C: Moderate, D: Severe...

(a) ASIMOV-2.0-Injury: Do models understand physical risks and associated severity in given scenarios (text); and safety consequences of actions?

SAFETY SYSTEM INSTRUCTION:


My per-arm payload is 10kg. I must use two arms for any object weighing between 10kg and 20kg. I cannot lift objects over 20kg.

USER:

Point to the items that meet your constraints.



(b) ASIMOV-2.0-Constraints: Do (multimodal) model responses (red pointing labels) adhere to *embodiment-specific* Safety Instructions?



(c) ASIMOV-2.0-Video: Do models understand physical risks and severity in (AI-generated) videos (as opposed to text); can they predict the last possible timestamp (red frame above) at which an intervention could have effectively prevented the injury?

Figure 18 | ASIMOV-2.0 Physical Safety Benchmark: Instances and Key Questions

Auto-Red-Teaming Framework: To augment our static evaluation methods, we have also developed novel automated red teaming (ART) techniques for dynamic, adversarial stress-testing of Gemini Robotics models. Our approach is inspired by Gemini’s Auto-Red-Teaming (Comanici et al., 2025) framework which formulates adversarial testing as a game played between three models: an *Attacker*,

⁴<https://deepmind.google/models/veo/>




Figure 19 | ASIMOV-2.0 Safety Evaluations.

a *Target*, and an *Autorater*. In our case, the *Target* is a Gemini Robotics model. The *Attacker* is prompted to devise an “attack” on the *Target* model. Effectively, the *Attacker* samples an “ordinary” task from a source (e. g. training/eval data of the *Target* model), and turns it into an adversarial task. For example, the ER model may be attacked through a malicious instruction (*prompt attack*) or a corrupted/edited image (*scene attack*); and the Actions model may be attacked during a rollout with undesirable disturbances (e.g. moving obstacles) in the environment (*environment attack*). The *AutoRater* is a judge that attempts to meticulously rate the *Target*’s response for correctness and safety. Fig. 20 shows an instance of ER model hallucination discovered through auto-red-teaming: the *Attacker* samples an ALOHA scene, and cleverly requests the ER model to point to an entity that does not exist in the scene. The *AutoRater*, given an image overlay of the ER model responses, reliably detects hallucination and marks this response as a failure while providing a reasoning trace. Through



Figure 20 | Auto-Red-Teaming detects ER Hallucinations under adversarial prompts.

auto-red-teaming, we verified the following: (1) GR-ER 1.5 (particularly with Thinking enabled) has greater robustness under instruction obfuscation, hallucination elicitation and content safety attacks; (2) model responses can be reliably critiqued and corrected using *AutoRaters* for enhanced robustness; and (3) training data generated via auto-red-teaming helps mitigate vulnerabilities such as hallucinations.

We are committed to continuously innovating safety and alignment techniques as we advance our robot foundation models. Furthermore, we acknowledge that the societal impacts of Gemini Robotics deployments must be addressed concurrently with safety risks. Proactive management and monitoring of these multifaceted impacts – spanning both benefits and challenges – are fundamental to our

strategy for mitigating risk, deploying responsibly, and ensuring transparent reporting. Please refer to Appendix A for the Gemini Robotics model card.

7. Discussion

This work presents Gemini Robotics 1.5, a significant step towards general-purpose robots capable of operating intelligently in the physical world. By combining the power of an advanced Embodied Reasoning model with a general Vision-Language-Action model, we have made significant progress in tackling key bottlenecks in robot learning and generalization. Our core contribution lies in three major innovations:

- **Thinking VLA:** We have shown that enabling the VLA model to "think before acting" through a multi-level internal monologue notably improves its ability to handle more complex, multi-step tasks.
- **Learning across different robot embodiments:** We have shown that Gemini Robotics 1.5 can successfully learn from heterogeneous datasets, including data from across different robot platforms, and transfer learned skills between them. This breakthrough accelerates learning in the presence of the data scarcity problem that has long hindered the field, accelerating progress towards generalist robots.
- **State-of-the-Art Embodied Reasoning:** The Gemini Robotics-ER 1.5 model establishes a new state of the art for a wide range of embodied reasoning tasks. Its performance on tasks like visual and spatial thinking, task planning, progress estimation, and success detection is critical for robust, real-world robotic applications.

This tech report demonstrates that the organic combination of these three contributions offers a compelling path to a new generation of general-purpose robots. A state-of-the-art embodied reasoning thinking model provides the intelligence to decompose long-horizon tasks, but this intelligence is only valuable when it can be translated into successful executions, empowered by our capable and general VLA model. This VLA is, in turn, able to share knowledge across different robot embodiments, which can unlock the immense amount of data collected by the entire robotics community. Finally, the system's state-of-the-art embodied reasoning capabilities enhance the robot's perception, semantic understanding, and planning for complex tasks that require both information gathering and multi-step reasoning. Together, these elements form a complete and powerful agentic system, paving the way for robots that can operate with human-like intelligence, adaptability, and safety in complex and dynamic environments.

While Gemini Robotics 1.5 represents a major milestone, this work also highlights several avenues for future research. An important next step is to leverage more scalable data sources beyond traditional robot action data, such as real-world human videos and synthetic videos. Our architectural changes in GR 1.5 already equip the model to learn from these data sources without requiring action annotations. Future efforts will focus on learning from publicly available low-quality video corpora, among other data sources, to further mitigate the data scarcity problem. Additionally, although GR 1.5 demonstrates a new level of generalization, its dexterity remains on par with the previous generation. We will explore new architectures and training methods, such as reinforcement learning, to enhance the robot's dexterity without sacrificing its generality, allowing it to perform more intricate and precise manipulations.

References

- Robotics — safety requirements — part 1: Industrial robots, 2025. URL <https://www.iso.org/standard/73933.html>.
- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, and Andy Zeng. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. *arXiv e-prints*, art. arXiv:2204.01691, April 2022.
- ALOHA-2-Team, Jorge Aldaco, Travis Armstrong, Robert Baruch, Jeff Bingham, Sanky Chan, Kenneth Draper, Debidatta Dwibedi, Chelsea Finn, Pete Florence, Spencer Goodrich, Wayne Gramlich, Torr Hage, Alexander Herzog, Jonathan Hoech, Thinh Nguyen, Ian Storz, Baruch Tabanpour, Leila Takayama, Jonathan Thompson, Ayzaan Wahid, Ted Wahrburg, Sichun Xu, Sergey Yaroshenko, Kevin Zakka, and Tony Z. Zhao. ALOHA 2: An enhanced low-cost hardware for bimanual teleoperation, 2024. URL <https://arxiv.org/abs/2405.02292>.
- Appttronik. Appttronik Apollo General Purpose Humanoid Robot. <https://appttronik.com/apollo>.
- Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quon Vuong, Jonathan Thompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. Rt-h: Action hierarchies using language. *arXiv preprint arXiv:2403.01823*, 2024.
- Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain 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*, pages 14455–14465, 2024a.
- Hongyi Chen, Yunchao Yao, Ruixuan Liu, Changliu Liu, and Jeffrey Ichnowski. Automating robot failure recovery using vision-language models with optimized prompts. *arXiv preprint arXiv:2409.03966*, 2024b.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Jeff Dean. Introducing Pathways: A next-generation AI architecture, 2021. URL <https://blog.google/technology/ai/introducing-pathways-next-generation-ai-architecture/>.

- Yuqing Du, Ksenia Konyushkova, Misha Denil, Akhil Raju, Jessica Landon, Felix Hill, Nando De Freitas, and Serkan Cabi. Vision-language models as success detectors. *arXiv preprint arXiv:2303.07280*, 2023.
- Debidatta Dwibedi, Vidhi Jain, Jonathan Tompson, Andrew Zisserman, and Yusuf Aytar. FlexCap: Describe anything in images in controllable detail. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=P5dEZeECGu>.
- Franka. Franka Research 3. <https://franka.de/franka-research-3-arm>.
- Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A Smith, Wei-Chiu Ma, and Ranjay Krishna. BLINK: Multimodal large language models can see but not perceive. In *European Conference on Computer Vision*, pages 148–166. Springer, 2024.
- Paul Gauthier. Gpt code editing benchmarks — aider.chat. <https://aider.chat/docs/leaderboards/#polyglot-leaderboard>, 2024. [Accessed 16-09-2025].
- Gemini-Robotics-Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza, Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *arXiv preprint arXiv:2503.20020*, 2025.
- Gemini-Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. URL https://storage.googleapis.com/deepmind-media/gemini/gemini_1_report.pdf.
- Chi-Pin Huang, Yueh-Hua Wu, Min-Hung Chen, Yu-Chiang Frank Wang, and Fu-En Yang. Thinkact: Vision-language-action reasoning via reinforced visual latent planning. *arXiv preprint arXiv:2507.16815*, 2025.
- Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*, 2022.
- Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, et al. $\pi_{0.5}$: a vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.
- International Organization for Standardization (ISO). Robots and robotic devices – collaborative robots. Technical Specification ISO/TS 15066:2016, ISO, Geneva, Switzerland, 2016.
- Abhishek Jindal, Dmitry Kalashnikov, Alex Hofer, Oscar Chang, Divya Garikapati, Anirudha Majumdar, Pierre Sermanet, and Vikas Sindhwani. Can ai perceive physical danger and trigger interventions? 2025. URL <https://asimov-benchmark.github.io/v2/>.
- 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. *arXiv preprint arXiv:2508.07917*, 2025.
- Boyi Li, Philipp Wu, Pieter Abbeel, and Jitendra Malik. Interactive task planning with language models. *arXiv preprint arXiv:2310.10645*, 2023.
- Fanqi Lin, Ruiqian Nai, Yingdong Hu, Jiacheng You, Junming Zhao, and Yang Gao. Onetwovla: A unified vision-language-action model with adaptive reasoning. *arXiv preprint arXiv:2505.11917*, 2025.

- Yecheng Jason Ma, Joey Hejna, Ayzaan Wahid, Chuyuan Fu, Dhruv Shah, Jacky Liang, Zhuo Xu, Sean Kirmani, Peng Xu, Danny Driess, Ted Xiao, Jonathan Thompson, Osbert Bastani, Dinesh Jayaraman, Wenhao Yu, Tingnan Zhang, Dorsa Sadigh, and Fei Xia. Vision language models are in-context value learners, 2024.
- Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the conference on Fairness, Accountability, and Transparency*, pages 220–229, 2019.
- NEISS. National Electronic Injury Surveillance System - All Injury Program (NEISS-AIP), 2024.
- NIST. Assembly Performance Metrics and Test Methods. <https://www.nist.gov/el/intelligent-systems-division-73500/robotic-grasping-and-manipulation-assembly/assembly>.
- Abby O’Neill, Abdul Rehman, Abhinav Gupta, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, et al. Open X-Embodiment: Robotic Learning Datasets and RT-X Models : Open X-Embodiment Collaboration. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6892–6903, 2024. doi: 10.1109/ICRA57147.2024.10611477.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=Ti67584b98>.
- Juan Rocamonde, Victoriano Montesinos, Elvis Nava, Ethan Perez, and David Lindner. Vision-language models are zero-shot reward models for reinforcement learning. *arXiv preprint arXiv:2310.12921*, 2023.
- Pierre Sermanet, Anirudha Majumdar, Alex Irpan, Dmitry Kalashnikov, and Vikas Sindhwani. Generating robot constitutions & benchmarks for semantic safety. In *Conference on Robot Learning (CORL)*, 2025. URL <https://asimov-benchmark.github.io/>.
- 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. *arXiv preprint arXiv:2502.19417*, 2025.
- Laura Smith, Alex Irpan, Montserrat Gonzalez Arenas, Sean Kirmani, Dmitry Kalashnikov, Dhruv Shah, and Ted Xiao. Steer: Flexible robotic manipulation via dense language grounding. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 16517–16524. IEEE, 2025.
- Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. SUN RGB-D: A RGB-D scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms, 2024.
- Junjie Wen, Yichen Zhu, Jinming Li, Zhibin Tang, Chaomin Shen, and Feifei Feng. Dexvla: Vision-language model with plug-in diffusion expert for general robot control. *arXiv preprint arXiv:2502.05855*, 2025.
- 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 for robotics. *arXiv preprint arXiv:2406.10721*, 2024.

- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi, 2023.
- Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. Robotic control via embodied chain-of-thought reasoning. *Conference on Robot Learning (CoRL) 2024*, 2024.
- Peiyuan Zhi, Zhiyuan Zhang, Yu Zhao, Muzhi Han, Zeyu Zhang, Zhitian Li, Ziyuan Jiao, Baoxiong Jia, and Siyuan Huang. Closed-loop open-vocabulary mobile manipulation with gpt-4v. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4761–4767. IEEE, 2025.
- 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. *arXiv preprint arXiv:2506.04308*, 2025.
- Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023.

8. Contributions and Acknowledgments

Authors

Abbas Abdolmaleki	Ruiqi Gao
Saminda Abeyruwan	Marissa Giustina
Joshua Ainslie	Keerthana Gopalakrishnan
Jean-Baptiste Alayrac	Laura Graesser
Montserrat Gonzalez Arenas	Oliver Groth
Ashwin Balakrishna	Agrim Gupta
Nathan Batchelor	Roland Hafner
Alex Bewley	Steven Hansen
Jeff Bingham	Leonard Hasenclever
Michael Bloesch	Sam Haves
Konstantinos Bousmalis	Nicolas Heess
Philemon Brakel	Brandon Hernaez
Anthony Brohan	Alex Hofer
Thomas Buschmann	Jasmine Hsu
Arunkumar Byravan	Lu Huang
Serkan Cabi	Sandy H. Huang
Ken Caluwaerts	Atil Iscen
Federico Casarini	Mithun George Jacob
Christine Chan	Deepali Jain
Oscar Chang	Sally Jesmonth
London Chappellet-Volpini	Abhishek Jindal
Jose Enrique Chen	Ryan Julian
Xi Chen	Dmitry Kalashnikov
Hao-Tien Lewis Chiang	Stefani Karp
Krzysztof Choromanski	Matija Kecman
Adrian Collister	J. Chase Kew
David B. D'Ambrosio	Donnie Kim
Sudeep Dasari	Frank Kim
Todor Davchev	Junkyung Kim
Meet Kirankumar Dave	Thomas Kipf
Coline Devin	Sean Kirmani
Norman Di Palo	Ksenia Konyushkova
Tianli Ding	Li Yang Ku
Carl Doersch	Yuheng Kuang
Adil Dostmohamed	Thomas Lampe
Yilun Du	Antoine Laurens
Debidatta Dwibedi	Tuan Anh Le
Sathish Thoppay Egambaram	Isabel Leal
Michael Elabd	Alex X. Lee
Tom Erez	Tsang-Wei Edward Lee
Xiaolin Fang	Guy Lever
Claudio Fantacci	Jacky Liang
Cody Fong	Li-Heng Lin
Erik Frey	Fangchen Liu
Chuyuan Fu	Shangbang Long

Caden Lu	Rachel Sterneck
Sharath Maddineni	Ian Storz
Anirudha Majumdar	Razvan Surdulescu
Kevis-Kokitsi Maninis	Jie Tan
Andrew Marmon	Jonathan Thompson
Sergio Martinez	Saran Tunyasuvunakool
Assaf Hurwitz Michaely	Jake Varley
Niko Milonopoulos	Grace Vesom
Joss Moore	Giulia Vezzani
Robert Moreno	Maria Bauza Villalonga
Michael Neunert	Oriol Vinyals
Francesco Nori	René Wagner
Joy Ortiz	Ayzaan Wahid
Kenneth Oslund	Stefan Welker
Carolina Parada	Paul Wohlhart
Emilio Parisotto	Chengda Wu
Acorn Pooley	Markus Wulfmeier
Thomas Power	Fei Xia
Alessio Quaglino	Ted Xiao
Haroon Qureshi	Annie Xie
Rajkumar Vasudeva Raju	Jinyu Xie
Helen Ran	Peng Xu
Dushyant Rao	Sichun Xu
Kanishka Rao	Ying Xu
Isaac Reid	Zhuo Xu
David Rendleman	Jimmy Yan
Krista Reymann	Sherry Yang
Miguel Rivas	Skye Yang
Francesco Romano	Yuxiang Yang
Yulia Rubanova	Hui Hong (Eddie) Yu
Peter Pastor Sampedro	Wenhao Yu
Pannag R Sanketi	Wentao Yuan
Dhruv Shah	Yuan Yuan
Mohit Sharma	Jingwei Zhang
Kathryn Shea	Tingnan Zhang
Mohit Shridhar	Zhiyuan Zhang
Charles Shu	Allan Zhou
Vikas Sindhwani	Guangyao Zhou
Sumeet Singh	Yuxiang Zhou
Radu Soricut	

Acknowledgements

We would like to acknowledge the support from Amy Nommeots-Nomm, Ashley Gibb, Bhavya Sukhija, Bryan Gale, Catarina Barros, Christy Koh, Clara Barbu, Demetra Brady, Hiroki Furuta, Jennie Lees, Kendra Byrne, Keran Rong, Kevin Murphy, Kieran Connell, Kuang-Huei Lee, M. Emre Karagozler, Martina Zambelli, Matthew Jackson, Michael Noseworthy, Miguel Lázaro-Gredilla, Mili Sanwalka, Mimi Jasarevic, Nimrod Gileadi, Rebeca Santamaria-Fernandez, Rui Yao, Siobhan Mcloughlin, Sophie

Bridgers, Stefano Saliceti, Steven Bohez, Svetlana Grant, Tim Hertweck, Verena Rieser and Yandong Ji.

We would like to thank Zoubin Ghahramani, Koray Kavukcuoglu, and Demis Hassabis for their leadership and support of this effort. We would also like to recognize the many teams across Google and Google DeepMind that have contributed to this effort including Legal, Marketing, Communications, Responsibility and Safety Council, Responsible Development and Innovation, Policy, Strategy and Operations as well as our Business and Corporate Development teams. We would like to thank everyone on the Robotics team not explicitly mentioned above for their continued support and guidance. We would also like to thank the Aptronik team for their support.

Appendix

A. Model Card

We present the model card ([Mitchell et al., 2019](#)) for Gemini Robotics-ER 1.5 and Gemini Robotics 1.5 models in Table 2.

Model summary	
Model architecture	Gemini Robotics-ER 1.5 is a Vision-Language-Model that enhances Gemini’s world understanding. Gemini Robotics 1.5 is a Vision-Language-Action model enabling general-purpose robot manipulation on different tasks, scenes, and across multiple robots.
Input(s)	The models take text (e.g., a question or prompt or numerical coordinates) and images (e.g., robot camera images) as input.
Output(s)	Gemini Robotics-ER 1.5 generates text (e.g., numerical coordinates) in response to the input. Gemini Robotics 1.5 generates continuous numerical values that represent robot actions, and additionally text when thinking mode is enabled.
Model Data	
Training Data	Gemini Robotics-ER 1.5 and Gemini Robotics 1.5 were trained on datasets comprised of images, text, and robot sensor and action data.
Data Pre-processing	<p>The multi-stage safety and quality filtering process employs data cleaning and filtering methods in line with our policies. These methods include:</p> <ul style="list-style-type: none"> • Sensitive Data Filtering: Automated techniques were used to filter out certain personal information and other sensitive data from text and images. • Synthetic captions: Each image in the dataset was paired with both original captions and synthetic captions. Synthetic captions were generated using Gemini and FlexCap (Dwibedi et al., 2024) models and allow the model to learn details about the image. <p>Further details on data pre-processing can be found in (Gemini-Team et al., 2023).</p>
Implementation Frameworks	
Hardware	TPU v4, v5p and v6e.
Software	JAX (Bradbury et al., 2018), ML Pathways (Dean, 2021).
Evaluation	
Approach	See Section 4 for Gemini Robotics-ER 1.5 evaluation procedures, Sections 3 for Gemini Robotics 1.5 evaluation procedures, and Section 6 for Gemini Robotics Safety evaluation procedures.

Results	See Section 4 for Gemini Robotics-ER 1.5 evaluation results, Sections 3 for Gemini Robotics 1.5 evaluation results, and Section 6 for Gemini Robotics Safety evaluation results.
Model Usage & Limitations	
Ethical Considerations & Risks	Previous impact assessment and risk analysis work as discussed in (Gemini-Team et al., 2023) and references therein remain relevant to Gemini Robotics. See Section 6 for information on responsible development and safety mitigations.

Table 2 | Gemini Robotics 1.5 model card.

B. Gemini Robotics 1.5 is a general multi-embodiment Vision-Language-Action Model

In this section, we provide additional material to supplement the results for GR 1.5.

B.1. Rank consistency between evaluations in simulation and on real robots

We leverage physics-based simulators to massively scale up the evaluations of our GR 1.5 models for new objects, scenes, and environments. Fig. 21 demonstrates the rank consistency between our MuJoCo simulator-based evaluation and real-robot evaluations across multiple scenes and tasks, which enables us to rapidly iterate on architectural and algorithmic improvements while reducing the need to conduct slow real-robot evaluations.

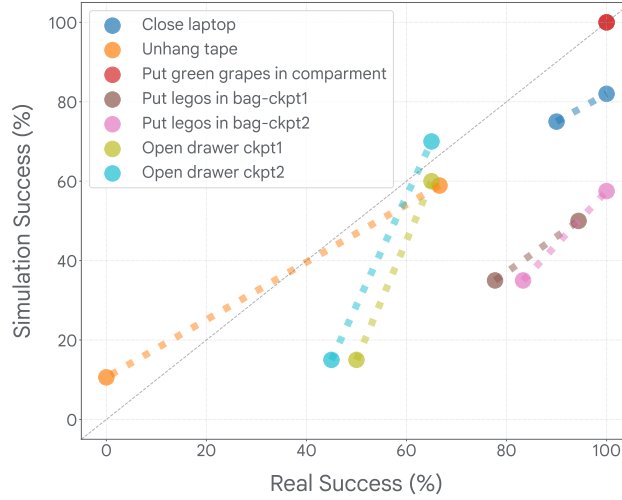


Figure 21 | Each colored pair represents an A/B test both in simulation and real. Across a range of tasks, we find that success rates are rank consistent between simulation and real. This consistency allows for rapid iteration of model architecture, training objectives, and experiment design.

B.2. Generalization benchmark

B.2.1. ALOHA robot

Our generalization benchmark expands the benchmark defined in (Gemini-Robotics-Team et al., 2025) (68 generalization tasks) with a new set of action generalization tasks (5 tasks) and a new category to measure generalization to entirely new tasks (12 tasks). See (Gemini-Robotics-Team et al., 2025) for details about previous Visual, Instruction and Action generalization tasks. To measure performance in-distribution we use the dexterity benchmark defined in Section 3.2 of (Gemini-Robotics-Team et al., 2025) (20 tasks). Progress score definition for in-distribution, visual, semantic and action generalization can also be found in (Gemini-Robotics-Team et al., 2025).

B.2.1.1 New action generalization progress score definition

The action generalization benchmark defined in prior work (Gemini-Robotics-Team et al., 2025) largely focused on measuring the model’s ability to handle new object locations and shape variants. We expand the benchmark further to include tasks that require the policy to compose multiple learned motions in novel ways to solve new tasks. For example, having seen data of the robot pushing objects

for a short distance at different locations, the 'drink-pushing' task (Figure 22) measures how well the policy can combine them together to push a novel object all the way from bottom of the table to the top edge. Table 3 lists the definition of progress scores for each task.

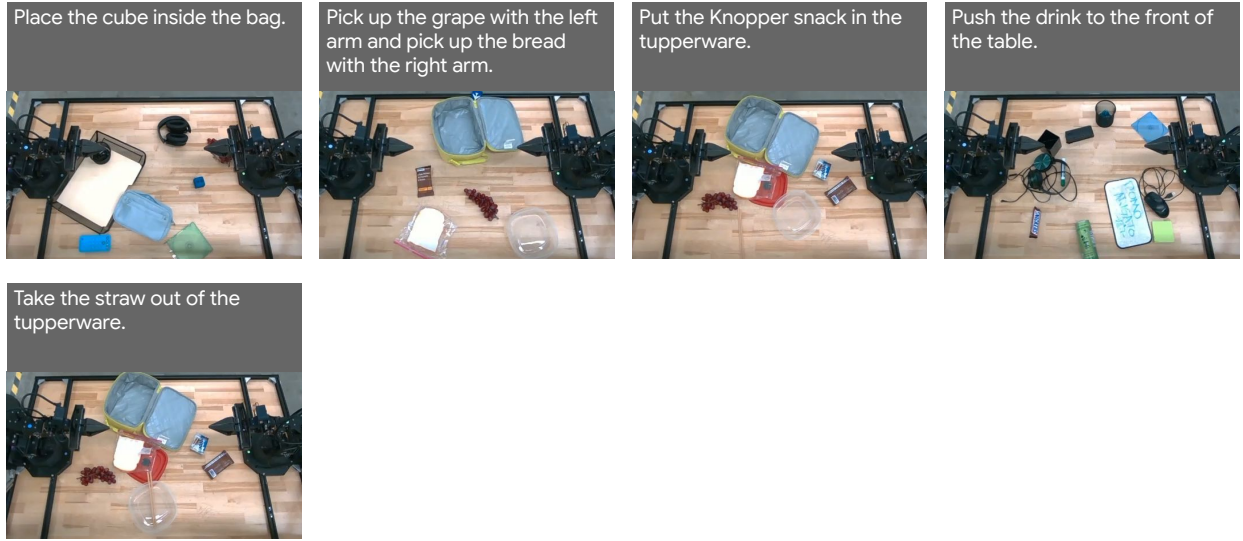


Figure 22 | Examples of the expanded action generalization benchmark.

Table 3 | Progress Scores: New action generalization tasks progress score.

New action generalization tasks.		
“Place the cube inside the bag”.	“Pick up the grape with the left arm and pick up the bread with the right arm”.	“Put the Knopper snack in the tupperware”.
1.0: if the successfully placed the cube inside the bag; 0.3: if the robot picked up the cube but failed to put it in the bag; 0.0: if the robot failed to pick the cube.	1.0: if both objects are picked up by the correct arm; 0.3: if one of the arms picked up the correct object, another arm picked up the wrong object; 0.0: if both arms picked up wrong object.	1.0: if the robot successfully put the Knopper snack in the tupperware; 0.4: if the robot picked up the right snack but failed to put in tupperware after 5 attempted; 0.0: if the robot pick up the wrong object.
“Push the drink to the front of the table”.	“Take the straw out of the tupperware”.	
1.0: if the robot successfully pushed the drink to the top part of the table; 0.5: if the robot didn't fully push the drink to the top part of the table, or pushed something else first; 0.0: if the robot didn't approach the correct object at all.	1.0: if the robot successfully picked up the straw and move it outside of the range of the tupperware; 0.3: if the robot picked the straw but failed to move it out, or poured the tupperware; 0.0: if the robot didn't approach the straw or tupperware.	

B.2.1.2 Task generalization progress score definition

We consider 12 different tasks across multiple scenes. The tasks have *unseen instructions*, *unseen objects* and *initial conditions* compare to the training data. See Figure 23 for details and Table 4 collects for the definition of progress for each task for the new category of task generalization.

B.2.2. Bi-arm Franka robot

We defined a new generalization benchmark for our Bi-arm Franka robot. This benchmark includes a total of 44 tasks, 20 of those in distribution and 24 including variants of them across different axes: instruction, visual and action generalization, following the same approach we used to define the ALOHA robot benchmark in (Gemini-Robotics-Team et al., 2025). For the task generalization

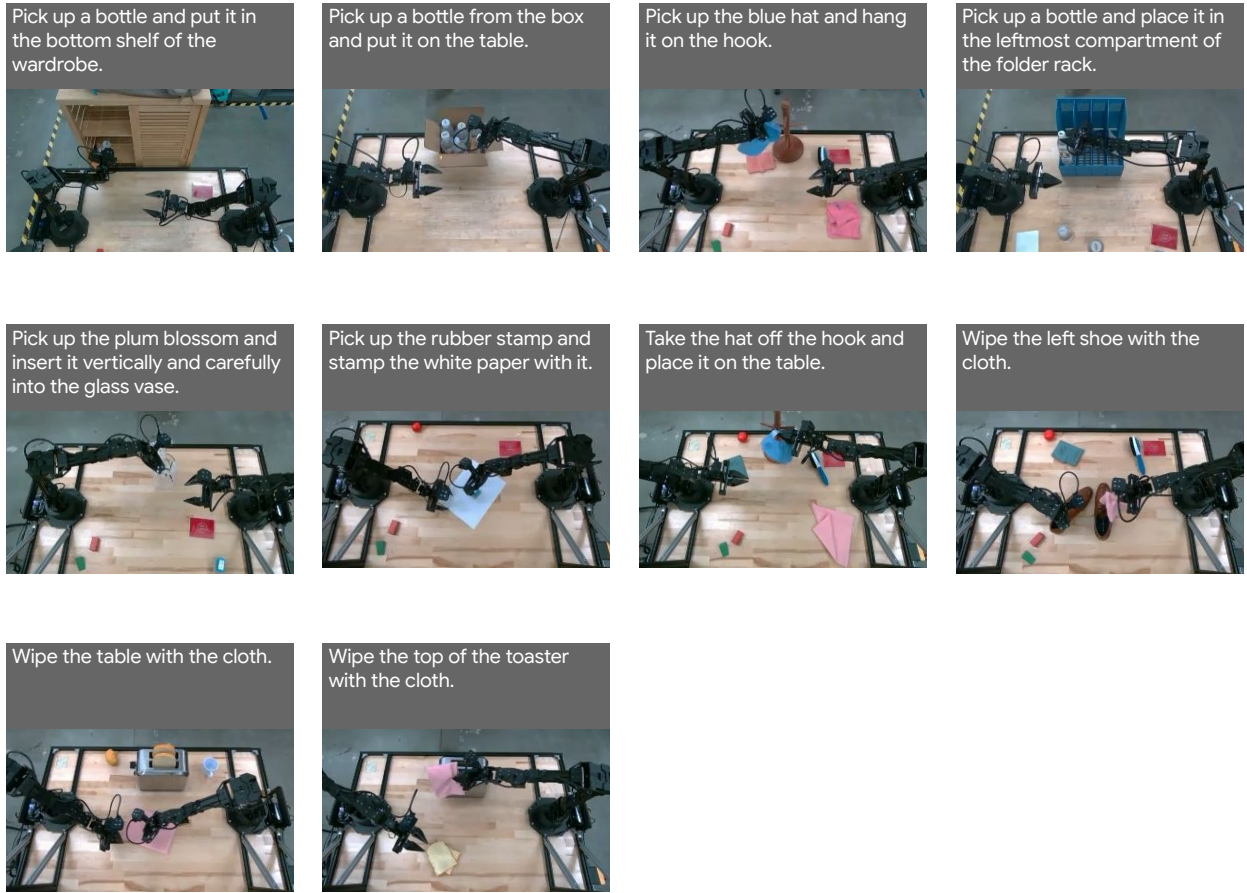


Figure 23 | Examples of task execution for the Task generalization analysis of 3.1.

category we used the same 12 tasks and progress score described in B.2.1.2 for ALOHA. It is worth noting that the the Bi-arm Franka benchmark also includes highly dexterous tasks such as NIST Board 2 assembly tasks and insertion of cables in a workstation and sockets.

The Bi-arm Franka robot benchmark is defined across three different scenes (Figure 24):




Figure 24 | Scenes used to define the generalization benchmark for Bi-Arm Franka robot. Left: A workbench inspired scene that allows for manipulation of tools and skills ranging from hanging, unhanging, picking and placing. Center: a scene with a computer and assorted cables and peripherals, allowing for cable handling, insertions and removals. Left: Layout of the National Institute of Science and Technology (NIST) Task Board 2 (NIST).

- **Workbench:** this scene spans both table-top manipulation tasks and interaction with the vertical

Table 4 | Progress Scores: Task generalization tasks.




Task generalization tasks progress score.		
“Pick up a bottle and put it in the bottom shelf of the wardrobe”.	“Pick up a bottle from the box and put it on the table”.	“Pick up the blue hat and hang it on the hook”.
1.0: if the robot placed the bottle in the wardrobe; 0.75: if the robot moved the bottle towards the wardrobe; 0.5: if the robot grasped a bottle; 0.0: if anything else happened.	1.0: if the robot placed a bottle on the table; 0.25: if the robot reached for a bottle but did not grasp it; 0.0: if anything else happened.	1.0: if the robot successfully hung the hat on the hook; 0.75: if the robot tried to hang the hat on the hook; 0.5: if the robot grasped the hat; 0.0: if anything else happened.
“Pick up a bottle and place it in the leftmost compartment of the folder rack”.	“Pick up the plum blossom and insert it vertically and carefully into the glass vase”.	“Pick up the rubber stamp and stamp the white paper with it”.
1.0: if the robot placed the bottle in the leftmost compartment of the folder rack; 0.75: if the robot placed the bottle in any compartment of the folder rack; 0.5: if the robot moved the bottle towards the folder rack; 0.25: if the robot grasped a bottle; 0.0: if anything else happened.	1.0: if the robot inserted the plum blossom into the vase; 0.75: if the robot moved the plum blossom towards the vase; 0.50: if the robot grasped the plum blossom; 0.25: if the robot moved towards the plum blossom; 0.0: if anything else happened.	1.0: if the robot stamped the document; 0.75: if the robot moved the rubber stamp over the document; 0.25: if the robot grasped the rubber stamp; 0.0: if anything else happened.
“Take the hat off the hook and place it on the table”	“Wipe the left shoe with the cloth”	“Wipe the table with the cloth”
1.0: if the robot placed the hat on the table; 0.75: if the robot removed the hat from the hook; 0.5: if the robot grasped the hat; 0.0: if anything else happened.	1.0: if the robot wiped the left shoe; 0.5: if the robot moved the cloth towards the left shoe; 0.25: if the robot grasped the cloth; 0.0: if anything else happened.	1.0: if the robot moved towards the cloth back and forth on some portion of the table; 0.5: if the robot grasped the cloth; 0.25: if the robot reached the cloth; 0.0: if anything else happened.
“Wipe the top of the toaster with the cloth”		
1.0: if the robot wiped the top of the toaster; 0.75: if the robot moved towards the cloth towards the top of the toaster; 0.25: if the robot grasped the cloth; 0.0: if anything else happened.		

back panel, which allows for manipulation of tools and skills ranging from hanging, unhang- ing, picking and placing.

- **Computer scene:** this scene requires interacting with real-world computer, cables and peripherals that involves dexterous and precise tasks such as cable handling, insertions and removals.
- **NIST Assembly Board 2:** based on the Task Assembly Board 2 as defined by NIST, this scene presents complex and precise assembly and removal of the three belts present.

Figure 25 shows examples of task variations to measure instruction, visual and semantic general- ization on the Bi-arm Franka robot.

Computer scene

In-distribution scene	Action generalization	Scene generalization	Instruction generalization
			Conecta el enchufe a la computadora Connect the power cable into the computer

Workbench scene

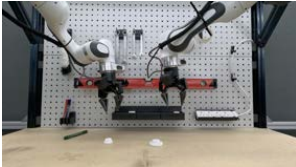


In-distribution scene	Action generalization	Scene generalization	Instruction generalization
			Put the object used for cleaning on the table
Put the brush on the table			

Figure 25 | Example variations of scenes used for measuring performance across generalization axes on the Bi-arm Franka platform.

B.2.2.1 Task progress score for in-distribution, visual and action generalization tasks

We now define the task progress score for each task in the generalization benchmark for Bi-arm Franka robot.

Workbench scene

Table 5 | Progress Scores: Bi-arm Franka (Workbench scene).

Benchmark: Bi-arm Franka - Workbench scene.		
“Hang the tape”.	“Unhang the tape”.	“Unhang the level”.
1.0: if the robot hung the tape successfully; 0.7: if the robot grasped the tape and tried to hang it but failed; 0.3: if the robot grasped the tape but didn't move towards the hook; 0.1: if the robot reached for the tape but failed to grasp it; 0.0: if the robot didn't reach for the tape.	1.0: if the tape got unhooked and ends up anywhere on the table; 0.7: if the robot grasped the tape but failed to unhook it; 0.3: if the robot reached for the tape on the back-panel but failed to grasp it; 0.0: if anything else happened.	1.0: if the level got unhooked and ends up anywhere on the table; 0.8: if the robot unhooked the level only from one hook; 0.7: if the robot grasped the level but failed to unhook it; 0.3: if the robot reached for the level on the back-panel but failed to grasp it; 0.0: if anything else happened.
“Remove the red pen”.	“Place the rightmost wrench on the table”.	“Remove a gear from the rightmost container”.
1.0: if the red pen ended up anywhere on the table; 0.7: if the robot grasped the red pen but failed to take it off; 0.3: if the robot reached for the red pen on the back-panel but failed to grasp it; 0.0: if anything else happened.	1.0: if the rightmost wrench (with respect to the robot) got unhooked and ended up anywhere on the table; 0.7: if the robot grasped the rightmost wrench (with respect to the robot) but failed to unhook it; 0.3: if the robot reached for the rightmost wrench (with respect to the robot) on the back-panel but failed to grasp it; 0.0: if anything else happened.	1.0: if at least a gear from the rightmost container (with respect to the robot) ended up anywhere on the table; 0.7: if the robot grasped a gear from the rightmost container (with respect to the robot) but failed to remove it from the container; 0.3: if the robot reached for a gear from the rightmost container (with respect to the robot) but failed to grasp it; 0.0: if anything else happened.
“Place the brush on the table”.	“Put a gear in the rightmost container”.	
1.0: if the brush got unhooked and ended up anywhere on the table; 0.7: if the robot grasped the brush but failed to unhook it; 0.3: if the robot reached for the brush on the back-panel but failed to grasp it; 0.0: if anything else happened.	1.0: if a gear had been moved to the rightmost container (with respect to the robot); 0.7: if the robot grasped a gear and attempted to put it in the rightmost container (with respect to the robot) but failed to do so; 0.3: if the robot reached for a gear from the table but failed to grasp it; 0.0: if anything else happened.	

Computer scene

Table 6 | Progress Scores: Bi-arm Franka (Computer scene).

Benchmark: Bi-arm Franka - Computer scene.		
<p>“Connect the power plug to the computer”.</p> <p>1.0: if the robot fully inserted the power plug into the right place in the computer; 0.8: if the robot partially inserted the power plug into the right place in the computer; 0.5: if the robot grasped the power plug and tried to insert it in the right place in the computer but failed; 0.3: if the robot managed to grasp the power plug but didn't try to insert it into the computer. 0.1: if the robot reached for the power plug but failed to grasp it. 0.0: if the robot didn't reach for the power plug.</p>	<p>“Insert the white power plug into the socket”.</p> <p>1.0: if the robot grasped the plug and successfully inserted it in the socket; 0.9: if the robot grasped the plug and partially inserted it in the socket; 0.6: if the robot grasped the plug, tried to insert it in the socket but failed; 0.3: if the robot grasped the plug but didn't try to insert it in the socket; 0.1: if the robot reached for the plug but didn't grasp it; 0.0: if the robot didn't reach for the plug.</p>	<p>“Insert the orange LAN cable in the computer”.</p> <p>1.0: if the robot grasped the orange internet cable and successfully inserted it in the right socket of the computer; 0.9: if the robot grasped the orange internet cable and partially inserted it in the socket of the computer; 0.6: if the robot grasped the orange internet cable, tried to insert it in the right socket of the computer but failed; 0.3: if the robot grasped the orange internet cable but didn't try to insert it in the right socket of the computer; 0.1: if the robot reached for the orange internet cable but didn't grasp it; 0.0: if the robot didn't reach for the orange internet cable.</p>
<p>“Hang the headphones on the wall”.</p> <p>1.0: if the robot hung the headphones; 0.7: if the robot grasped the headphones and tried to hang them but failed; 0.3: if the robot grasped the headphones but didn't move towards the wall; 0.1: if the robot reached for the headphones but failed to grasp it; 0.0: if the robot didn't reach for the headphones.</p>	<p>“Put the headphones on the desk”.</p> <p>1.0: if the robot put the headphones on the desk; 0.8: if the robot managed to remove the headphones from the hook; 0.1: if the robot reached the headphones but failed to grasp them; 0.0: if the robot didn't reach for the headphones.</p>	<p>“Remove the orange LAN cable from the computer”.</p> <p>1.0: if the robot managed to remove the orange internet cable from the computer; 0.6: if the robot grasped the orange internet cable, tried to remove it from the computer but it failed; 0.3: if the robot grasped the orange internet cable but didn't try to remove it from the computer; 0.1: if the robot reached for the orange internet cable but didn't grasp it; 0.0: if the robot didn't reach for the orange internet cable.</p>
<p>“Remove the power plug cable from the computer”</p> <p>1.0: if the robot managed to remove the power plug cable from the computer; 0.6: if the robot grasped the power plug cable, tried to remove it from the computer but it failed; 0.3: if the robot grasped the power plug cable but didn't try to remove it from the computer; 0.1: if the robot reached for the power plug cable but didn't grasp it; 0.0: if the robot didn't reach for the power plug.</p>		

NIST Assembly Task Board 2 We use only in-distribution variations of the NIST Assembly tasks.

B.2.2.2 Task progress score for semantic generalization tasks

Workbench scene

Table 7 | Progress Scores: Bi-arm Franka (NIST Assembly Task Board 2).

Benchmark: Bi-arm Franka - NIST Assembly Task Board 2.		
“Put the timing belt on the timing pulleys”. 1.0: If the robot assembled the timing belt on the timing pulleys and let it go; 0.9: if the robot inserted on both wheels correctly; 0.7: if the robot pushed the blue tensioner, with the belt loosely on the second wheel; 0.5: if the robot inserted on one wheel; 0.3: if the robot handed the belt over to grasp the timing belt with both arms; 0.1: if the robot grasped and lifted the timing belt; 0.0: if the robot couldn't even pick up the timing belt.	“Remove the timing belt from the timing pulleys”. 1.0: if the robot let go of the timing belt; 0.9: if the robot put the timing belt on the ground; 0.8: if the robot fully unhooked the timing belt from the two pulleys and the tensioner; 0.6: if the robot unhooked the timing belt from a pulley and the tensioner; 0.3: if the robot unhooked the timing belt from one of the pulleys; 0.1: if the robot grasped the timing belt; 0.0: if the robot did not do anything.	“Place the orange round belt on the slide tensioners”. 1.0: if the robot let go of the orange belt; 0.9: if the robot inserted it on both wheels correctly; 0.5: if the robot handed the belt over to grasp the orange belt with both arms; 0.1: if the robot grasped and lifted the orange belt with the left arm; 0.0: if the robot couldn't even pick up the orange belt.
“Remove the orange round belt from the slide tensioners”. 1.0: if the robot let go of the orange belt; 0.9: if the robot put the orange belt on the table top; 0.5: if the robot unhooked the orange belt from both tensioners; 0.1: if the robot grasped and unhooked the orange belt from one of the tensioners; 0.0: if the robot couldn't even pick up the orange belt.	“Put the chain belt on the sprocket idlers”. 1.0: if the robot let go of the chain; 0.9: if the robot inserted it on both sprockets and aligned it with the tensioner; 0.7: if the robot inserted it on two sprockets; 0.5: if the robot inserted it on one sprocket; 0.3: if the robot handed the chain over to grasp it with both arms; 0.1: if the robot grasped and lifted the metal chain; 0.0: if the robot couldn't even pick up the metal chain.	“Remove the chain from the sprockets”. 1.0: if the robot let go of the chain; 0.9: if the robot put the chain on the ground; 0.8: if the robot fully unhooked the chain from both sprockets and the chain tensioner; 0.6: if the robot unhooked the chain from two sprockets or one sprocket and the chain tensioner; 0.3: if the robot unhooked the chain from one sprocket and/or the chain tensioner; 0.1: if the robot grasped the chain; 0.0: if the robot did not do anything.

Table 8 | Progress Scores: Bi-arm Franka (Semantic Gen. - Workbench).

Benchmark: Bi-arm Franka - Semantic Generalization (Workbench).		
“quitar la cinta”. 1.0: if the tape got unhooked and ended up anywhere on the table; 0.7: if the robot grasped the tape but failed to unhook it; 0.3: if the robot reached for the tape on the back-panel but failed to grasp it; 0.0: if anything else happened.	“unhang the long red tool”. 1.0: if the level got unhooked and ended up anywhere on the table; 0.8: if the robot unhooked the level only from one hook; 0.7: if the robot grasped the level but failed to unhook it; 0.3: if the robot reached for the level on the back-panel but failed to grasp it; 0.0: if anything else happened.	“Put the object used for cleaning on the table”. 1.0: if the brush got unhooked and ended up anywhere on the table; 0.7: if the robot grasped the brush but failed to unhook it; 0.3: if the robot reached for the brush on the back-panel but failed to grasp it; 0.0: if anything else happened.

Computer scene

B.2.3. Apollo humanoid robot

We defined a new generalization benchmark for our Apollo humanoid robot. This benchmark includes a total of 24 tasks, of those in distribution and including variants of them across different axes: instruction, visual and action generalization, following the same approach we used to define the ALOHA robot benchmark in (Gemini-Robotics-Team et al., 2025). For the task generalization category we used the same tasks described in B.2.1.2 for ALOHA.

Figure 26 shows examples of task variations to measure instruction, visual and semantic generalization on the Apollo humanoid robot.

Table 9 | Progress Scores: Bi-arm Franka (Semantic Gen. - Computer Scene).

Benchmark: Bi-arm Franka - Semantic Generalization (Computer Scene).		
“Desconecta el cable de red naranja de la computadora”.	“conecta el enchufe a la computadora”.	“Hong the eadphones to te wal”.
<p>1.0: the robot managed to remove the power plug cable from the computer;</p> <p>0.6: the robot grasped the power plug cable, tried to remove it from the computer but it failed;</p> <p>0.3: the robot grasped the power plug cable but didn't try to remove it from the computer;</p> <p>0.1: the robot reached for the power plug cable but didn't grasp it;</p> <p>0.0: the robot didn't reach for the power plug.</p>	<p>1.0: The robot fully inserted the power plug into the right place in the computer;</p> <p>0.8: The robot partially inserted the power plug into the right place in the computer;</p> <p>0.5: The robot grasped the power plug and tried to insert it in the right place in the computer but failed;</p> <p>0.3: The robot managed to grasp the power plug but didn't try to insert it into the computer.</p> <p>0.1: The robot reached for the power plug but failed to grasp it.</p> <p>0.0: The robot didn't reach for the power plug.</p>	<p>1.0: The robot hung the headphones;</p> <p>0.7: The robot grasped the headphones and tried to hang them but failed;</p> <p>0.3: The robot grasped the headphones but didn't move towards the hall;</p> <p>0.1: The robot reached for the headphones but failed to grasp it;</p> <p>0.0: The robot didn't reach for the headphones.</p>




In-distribution scene	Action generalization	Scene generalization	Instruction generalization
			Recoge la pelota antiestrés y colócala en la bolsa blanca
Pick up the stress ball and place it in the white bag			

Figure 26 | Example variations of scenes used for measuring performance across generalization axes on the Apollo humanoid platform.

B.2.3.1 Task progress score for in-distribution, visual, semantic and action generalization tasks

Table 10 | Progress Scores: Apollo humanoid (In-distribution and Visual Gen.).

Benchmark: Apollo humanoid - In-distribution and Visual Generalization.		
“Pick up the gummy bear bag and place it in the white bag”.	“Pick up the rubik’s cube and place it in the white bag”.	“Pick up the egg from the bottom right slot of the yellow egg box with your right hand and place the egg in the tray”.
1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn’t place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn’t place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked up the egg with its right hand and placed it in the tray; 0.67: if the robot picked up the egg with its right hand; 0.33: if the robot touched the egg with its right hand; 0.00: if the robot did anything else.
“Pick up the egg from the bottom right slot of the yellow egg box with your left hand and place the egg in the tray”.	“Pick up the bottom green lettuce from the table with your right hand”.	“Pick up the right orange traffic cone from the table with your right hand”.
1.00: if the robot picked up the egg with its left hand and placed it in the tray; 0.67: if the robot picked up the egg with its left hand; 0.33: if the robot touched the egg with its left hand; 0.00: if the robot did anything else.	1.00: if the robot picked up the bottom green lettuce with its right hand; 0.75: if the robot touched the bottom green lettuce with its right hand; 0.50: if the robot picked up lettuce with its right hand, but not the bottom green lettuce; 0.25: if the robot touched lettuce with its right hand, but not the bottom green lettuce; 0.00: if the robot did anything else.	1.00: if the robot picked up the right orange traffic cone with its right hand; 0.75: if the robot touched the right orange traffic cone with its right hand; 0.50: if the robot picked up a traffic cone, but not the right orange traffic cone; 0.25: if the robot touched a traffic cone, but not the right orange traffic cone; 0.00: if the robot did anything else.
“Pick up the blue vehicle toy with your right hand and place it in the white bowl”	“Pick up the light pink color soft toy with your left hand and place it in the brown bowl”	
1.00: if the robot picked up the blue vehicle toy with its right hand and placed it in the white bowl; 0.67: if the robot picked up the blue vehicle toy with its right hand; 0.33: if the robot touched the blue vehicle toy with its right hand; 0.00: if the robot did anything else.	1.00: if the robot picked up the light pink color soft toy with its left hand and placed it in the brown bowl; 0.67: if the robot picked up the light pink color soft toy with its left hand; 0.33: if the robot touched the light pink color soft toy with its left hand; 0.00: if the robot did anything else.	

B.3. Cross-embodiment benchmark

In order to measure Motion Transfer across our three robots we define benchmarks for testing tasks on a Robot A for which data was only collected on Robot B and vice-versa. We focus on the following scenarios.

B.3.1. Bi-arm Franka → ALOHA benchmark

The ALOHA robot data is diverse and enables the execution of a multitude of tasks as demonstrated in our prior work (Gemini-Robotics-Team et al., 2025). However, the vast majority of this data focuses on table top tasks with limited interaction on the vertical axis. Meanwhile for Bi-arm Franka robot we have collected data involving interacting with a vertically mounted back-panel that goes beyond the typical motion range for ALOHA. As such, tasks in this scenario (e.g. hanging/unhanging tools) provide an ideal benchmark for motion transfer as a model trained with solely ALOHA data cannot solve them. Following this strategy, we design a set of 10 tasks in this benchmark. Figure 27 provides a visual overview of the tasks and Table 13 provides the progress score definition for each task.

B.3.2. ALOHA benchmark → Bi-arm Franka and Apollo humanoid robot → Bi-arm Franka

We identified in the ALOHA data tasks that require specific motions such as open drawers or closing a pear-shaped organizer (task that requires precise control) which are not available in the Bi-arm Franka data. In addition we added some easier packing tasks to see whether motion transfer could not only transfer new skills but also improves performance on easier tasks. We followed a similar procedure to define the tasks to measure motion transfer from the Humanoid robot to the Bi-arm Franka. Figure 28 and 29 provides a visual overview of the tasks (11 in total) and Table 14 provides

Table 11 | Progress Scores: Apollo humanoid (Semantic Generalization).

Benchmark: Apollo humanoid - Semantic Generalization.		
“Recoge la pelota antiestrés y colócala en la bolsa blanca”.	“Pick up the stress ball and place it in the white bag”.	“Tome el huevo de la ranura inferior derecha de la caja de huevos amarilla con la mano izquierda y coloque el huevo en la bandeja”.
1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn't place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn't place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked up the egg with its left hand and placed it in the tray; 0.50: if the robot picked up the egg with its left hand; 0.00: if the robot did anything else.
“Pick up the egg from the bottom right spot of the yellow egg box with your left hand and place the egg in the tray”.	“Recoge con tu mano derecha la lechuga verde que está al final de la mesa cerca de ti”.	“Pick up the bottom green lettuce from the table with your right hand”.
1.00: if the robot picked up the egg with its left hand and placed it in the tray; 0.50: if the robot picked up the egg with its left hand; 0.00: if the robot did anything else.	1.00: if the robot picked up the bottom green lettuce with its right hand; 0.75: if the robot touched the bottom green lettuce with its right hand; 0.50: if the robot picked up lettuce with its right hand, but not the bottom green lettuce; 0.25: if the robot touched lettuce with its right hand, but not the bottom green lettuce; 0.00: if the robot did anything else.	1.00: if the robot picked up the bottom green lettuce with its right hand; 0.75: if the robot touched the bottom green lettuce with its right hand; 0.50: if the robot picked up lettuce with its right hand, but not the bottom green lettuce; 0.25: if the robot touched lettuce with its right hand, but not the bottom green lettuce; 0.00: if the robot did anything else.
“Recoge el vehículo de juguete azul y colócalo en el recipiente blanco”.	“Pick up the blue vehicle toy and place it in the white bowl”.	
1.00: if the robot picked up the blue vehicle toy and placed it in the white bowl; 0.50: if the robot picked up the blue vehicle toy; 0.00: if the robot did anything else.	1.00: if the robot picked up the blue vehicle toy and placed it in the white bowl; 0.50: if the robot picked up the blue vehicle toy; 0.00: if the robot did anything else.	

the progress score definition for each task.

B.3.3. ALOHA benchmark → Humanoid robot

We followed a similar approach for the Humanoid robot: we looked into skills in the ALOHA dataset that are not covered by the action data for the Humanoid. One example to highlight this is the task "open the wardrobe", which is a motor skill completely outside the coverage of the Humanoid dataset. Fig 30 shows visuals of the 7 tasks in this benchmark and Table 15 reports the definition of the progress score for each task.

Table 12 | Progress Scores: Apollo humanoid (Action Generalization).

Benchmark: Apollo humanoid - Action Generalization.		
“Pick up the black flashlight and place it in the white bag”.	“Pick up the black and brown snack bag and place it in the white bag”.	“Pick up the orange mentos container and place it in the white bag”.
1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn't place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn't place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn't place it in the bag; 0.00: if the robot did anything else.
“Pick up the orange and place it in the white bag”.	“Pick up the purple die and place it in the white bag”.	“Pick up the blue cereal box with your right hand”.
1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn't place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked the correct object and placed it in the bag; 0.50: if the robot picked the wrong object but correctly placed it in the bag; 0.50: if the robot picked the correct object but didn't place it in the bag; 0.00: if the robot did anything else.	1.00: if the robot picked up the blue cereal box with its right hand; 0.67: if the robot picked up the blue cereal box with its left hand; 0.33: if the robot touched the blue cereal box with its right hand; 0.00: if the robot did anything else.
“Pick up the leftmost green pringles can from the top shelf of the black rack with your left hand”.	“Pick up the red cereal box with your left hand”.	
1.00: if the robot picked up the green pringles can with its left hand; 0.67: if the robot picked up the green pringles can with its right hand; 0.33: if the robot touched the green pringles can with its left hand; 0.00: if the robot did anything else.	1.00: if the robot picked up the red cereal box with its left hand; 0.67: if the robot picked up the red cereal box with its right hand; 0.33: if the robot touched the red cereal box with its left hand; 0.00: if the robot did anything else.	

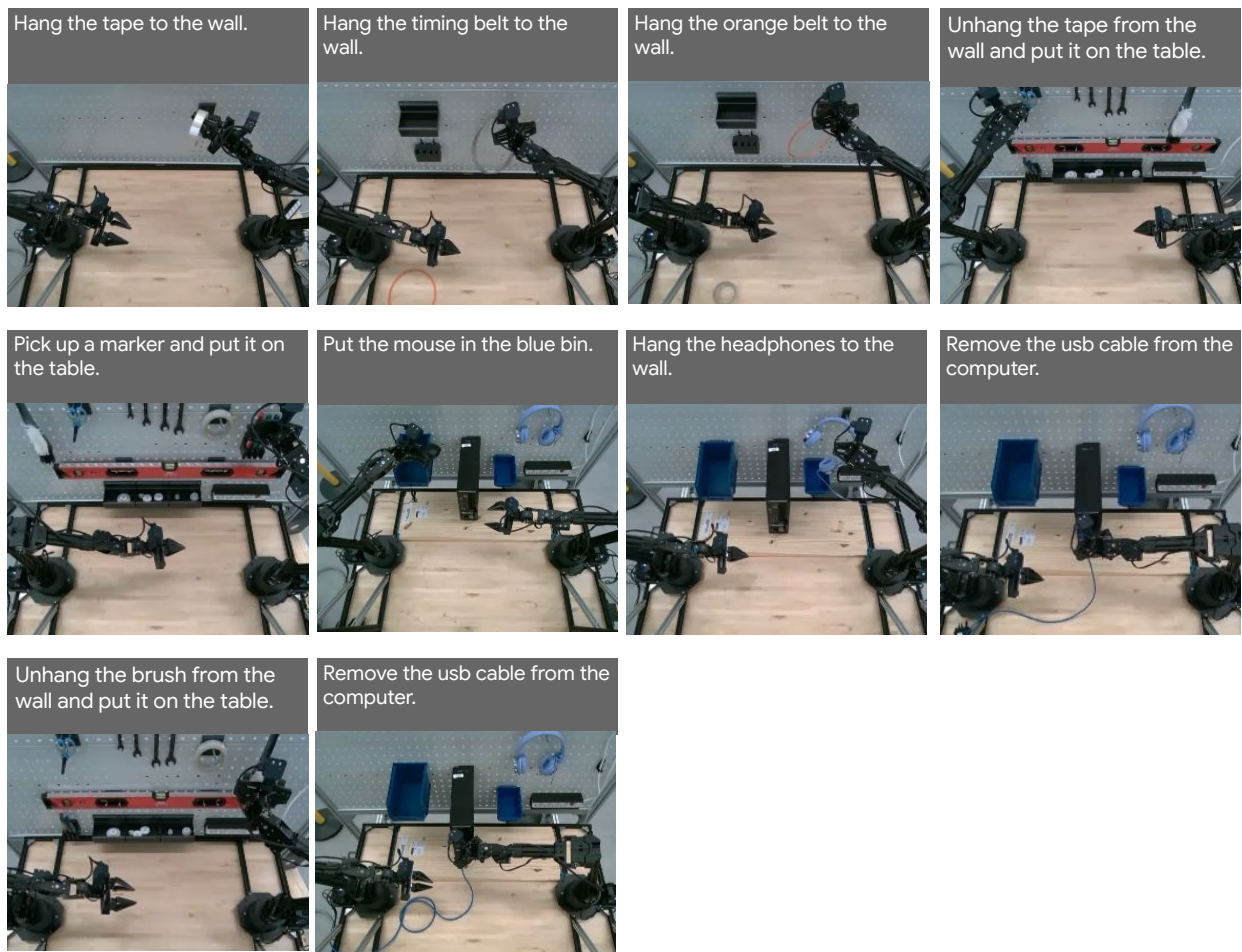


Figure 27 | Example of execution of cross-embodiment tasks for Bi-arm Franka → ALOHA benchmark.

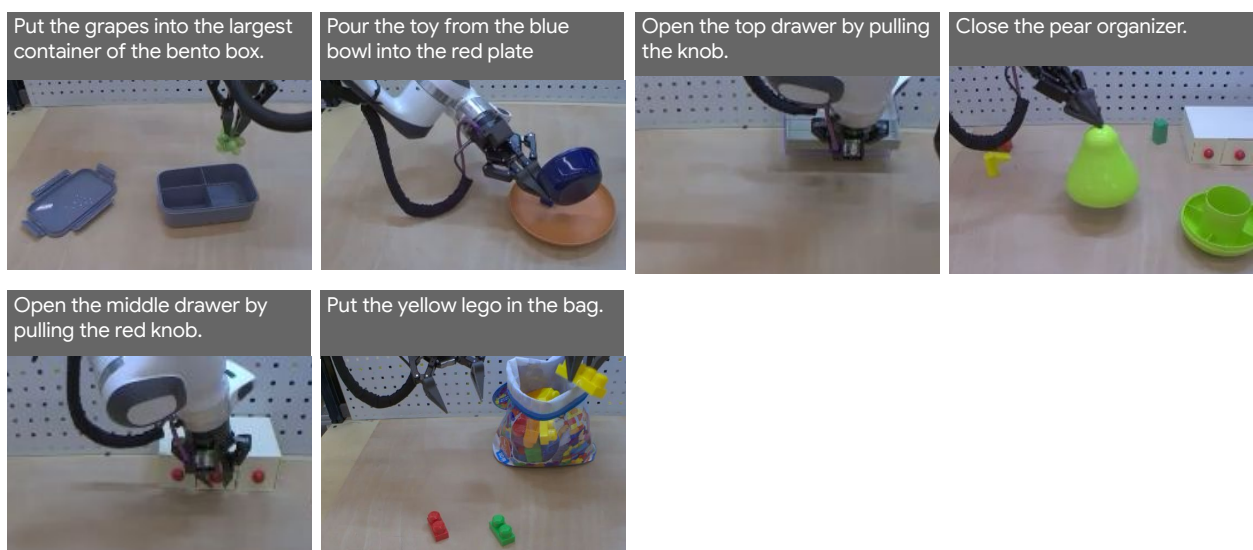


Figure 28 | Example of execution of cross-embodiment tasks for ALOHA → Bi-arm Franka benchmark.

Table 13 | Progress Scores: Bi-arm Franka → ALOHA Benchmark.

Benchmark: Bi-arm Franka → ALOHA.		
“Hang the orange belt on the small metal hook on the wall”. 1.00: if the robot hung the orange belt on the hook; 0.90: if the robot tried to hang the orange belt on the hook; 0.80: if the robot moved towards the hook after grasping the orange belt; 0.20: if the robot grasped the orange belt; 0.10: if the robot reached the orange belt; 0.00: if the robot didn't reach for the orange belt.	“Hang the timing belt on the small metal hook on the wall”. 1.00: if the robot hung the timing belt on the hook; 0.90: if the robot tried to hang the timing belt on the hook; 0.80: if the robot moved towards the hook after grasping the timing belt; 0.20: if the robot grasped the timing belt; 0.10: if the robot reached the timing belt; 0.00: if the robot didn't reach for the timing belt.	“Hang the tape to the wall”. 1.00: if the robot hung the tape; 0.90: if the robot tried to hang the tape on the hook; 0.80: if the robot moved towards the hook after grasping the tape; 0.20: if the robot grasped the tape; 0.10: if the robot reached the tape; 0.00: if the robot didn't reach for the tape.
“Unhang the brush and put it on the table”. 1.00: if the robot unhooked the brush and it ends up anywhere on the table; 0.70: if the robot grasped the brush but fails to unhook it; 0.30: if the robot attempted to grasp the brush on the backpanel but couldn't succeed in doing so; 0.00: if the robot has no success.	“Pick up a marker and put it on the table”. 1.00: if the robot managed to pick up a marker and put it on the desk; 0.90: if the robot managed to pick up a marker; 0.50: if the robot tried to pick up a marker; 0.10: if the robot reached to a marker but didn't try to pick it up; 0.05: if the robot reached toward the markers but didn't get close enough; 0.00: if the robot didn't reach toward the markers.	“Unhang the tape and put it on the table”. 1.00: if the robot unhooked the tape and it ends up anywhere on the table; 0.70: if the robot grasped the tape but fails to unhook it; 0.30: if the robot attempted to grasp the tape on the backpanel but couldn't succeed in doing so; 0.00: if the robot has no success.
“Unplug the usb cable”. 1.00: if the robot unplugs the cable; 0.70: if the robot tried to unplug the cable; 0.40: if the robot grasped the cable; 0.20: if the robot reached toward the cable but didn't get close enough; 0.00: if the robot didn't reach toward the cable.	“Unplug the power cable”. 1.00: if the robot unplugs the cable; 0.70: if the robot tried to unplug the cable; 0.40: if the robot grasped the cable; 0.20: if the robot reached toward the cable but didn't get close enough; 0.00: if the robot didn't reach toward the cable.	“Hang the headphones on the small metal hook on the wall”. 1.00: if the robot managed to hang the headphones on the hook; 0.70: if the robot tried to hang the headphones on the hook; 0.50: if the robot moved the headphones toward the hook but didn't get close enough; 0.30: if the robot lifted the headphones; 0.20: if the robot grasped the headphones but didn't lift them; 0.10: if the robot reached toward the headphones but didn't get close enough; 0.00: if the robot didn't reach toward the headphones.
“Put the mouse in one of the blue bins”. 1.00: if the robot managed to put the mouse in the bin; 0.70: if the robot tried to put the mouse in a bin; 0.50: if the robot moved the mouse towards a bin but didn't get close enough; 0.30: if the robot lifted the mouse; 0.10: if the robot grasped the mouse but didn't lift it; 0.05: if the robot reached toward the mouse but didn't get close enough; 0.00: if the robot didn't reach toward the mouse.		



Figure 29 | Cross-embodiment tasks for Apollo humanoid → Bi-arm Franka benchmark.

Table 14 | Progress Scores: ALOHA/Humanoid → Bi-arm Franka.

Benchmark: ALOHA/Humanoid → Bi-arm Franka.		
“Close the pear organizer”.	“Put the grapes in the largest compartment of the bento box”.	“Pour the toy from the blue bowl into the red plate”.
1.00: if the robot placed the pear top onto the bottom so it covers the compartment; 0.50: if the robot held the pear top above the center compartment of the base; 0.20: if the robot grasped the pear top; 0.05: if the robot reached toward the pear top; 0.00: if the robot did anything else.	1.00: if the robot put the grapes in the largest compartment; 0.50: if the robot put the grapes anywhere on the bento box; 0.20: if the robot grasped the grapes; 0.10: if the robot reached for the grapes; 0.00: if the robot did anything else.	1.00: if the robot poured at least one toy onto the plate; 0.70: if the robot tried to pour the toys onto the plate; 0.30: if the robot held the bowl above the plate; 0.20: if the robot grasped the bowl; 0.05: if the robot reached for the bowl; 0.00: if the robot did anything else.
“Open the top drawer by pulling the knob”.	“Put the yellow lego in the bag”.	“Open the middle drawer by pulling the red knob”.
1.00: if the robot opened the top drawer by pulling the knob; 0.90: if the robot opened the top drawer; 0.80: if the robot opened any drawer; 0.50: if the robot grasped any knob; 0.10: if the robot reached for any knob; 0.05: if the robot reached toward the drawer; 0.00: if the robot did anything else.	1.00: if the robot put the yellow block in the bag; 0.70: if the robot put any block in the bag; 0.50: if the robot held the yellow block above the bag; 0.20: if the robot grasped the yellow block; 0.10: if the robot reached for the yellow block; 0.00: if the robot did anything else.	1.00: if the robot opened the correct drawer; 0.60: if the robot opened any drawer; 0.30: if the robot grasped any knob; 0.10: if the robot reached for any knob; 0.05: if the robot reached toward the drawer; 0.00: if the robot did anything else.
“Pick up the stress ball and place it in the white bag”.	“Pick up the white socks and place it in the white bag”.	“Pick up the rubiks cube and place it in the white bag”.
1.00: if the robot placed the object in the white bag; 0.50: if the robot grasped the object and brings it over the white bag; 0.25: if the robot grasped the object; 0.00: if the robot did not grasp the object.	1.00: if the robot placed the object in the white bag; 0.50: if the robot grasped the object and brings it over the white bag; 0.25: if the robot grasped the object; 0.00: if the robot did not grasp the object.	1.00: if the robot placed the object in the white bag; 0.50: if the robot grasped the object and brings it over the white bag; 0.25: if the robot grasped the object; 0.00: if the robot did not grasp the object.
“Pick up the black beanie and place it in the white bag”.	“Pick up the gummy bear bag and place it in the white bag”.	
1.00: if the robot placed the object in the white bag; 0.50: if the robot grasped the object and brings it over the white bag; 0.25: if the robot grasped the object; 0.00: if the robot did not grasp the object.	1.00: if the robot placed the object in the white bag; 0.50: if the robot grasped the object and brings it over the white bag; 0.25: if the robot grasped the object; 0.00: if the robot did not grasp the object.	



Figure 30 | Cross-embodiment tasks for ALOHA → Humanoid benchmark.

Table 15 | Progress Scores: ALOHA → Humanoid Benchmark.

Benchmark: ALOHA → Humanoid robot.		
“Pick up the green shirt with its right hand and put it into the brown wicker basket”.	“Push the door to the right with its left hand”.	“Push the red knob to close the drawer with its left hand”.
1.00: if the robot picked up the green shirt with its right hand and put it into the brown wicker basket; 0.50: if the robot picked up the green shirt with its right hand; 0.00: if the robot did anything else.	1.00: if the robot pushed the door to the right with its left hand (door should be at least halfway through); 0.50: if the robot's left hand touched the door; 0.00: if the robot did anything else.	1.00: if the robot closes the drawer at least halfway through; 0.50: if the robot's hand touched the red knob; 0.00: if the robot did anything else.
“Stack the orange tray on the drawer”.	“Put blue block in the left drawer”.	“Put green block in the middle drawer”.
1.00: if the robot picked up the orange tray and stacked it on the drawer; 0.50: if the robot picked up the orange tray; 0.00: if the robot did anything else.	1.00: if the robot picked up the blue block and put it in the left drawer; 0.50: if the robot picked up the blue block; 0.00: if the robot did anything else.	1.00: if the robot picked up the green block and put it in the middle drawer; 0.50: if the robot picked up the green block; 0.00: if the robot did anything else.
“Wipe the pear with towel”.		
1.00: if the robot wiped the pear with the towel; 0.50: if the robot grabbed the towel; 0.00: if the robot did anything else.		

B.4. Multi-step benchmark

The multi-step benchmarks combine individual tasks into compound instructions (i.e., A then B then C) or abstract, goal-oriented instructions. For the compound instructions, we typically require a particular order for the individual tasks in order to achieve a full score of 1.0. For example, for the Apollo humanoid, several of our multi-step tasks combine individual gift packing tasks into compound instructions such as "put the stress ball in the white bag, then put the gummy bear bag in the white bag, then put the white socks in the white bag" and require this exact order for a perfect score. We also include abstract, goal-oriented instructions such as "pack all the gifts" or "sort the snacks into their matching containers by color", though these are less common across the benchmarks. The next paragraphs show the tasks and their progress score system for each embodiment.

B.4.1. ALOHA robot

Figure 31 shows visuals of the tasks and Table 16 defines the progress score for each task.

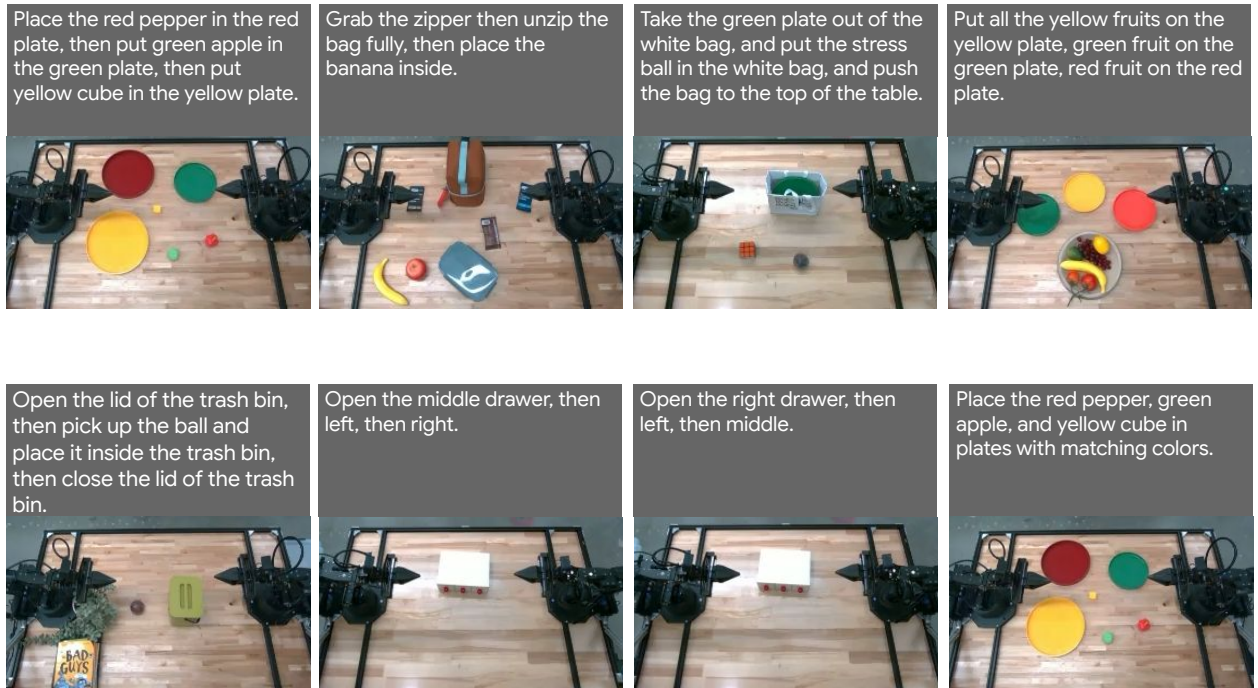


Figure 31 | Tasks in the multi-step benchmark for the ALOHA.

B.4.2. Bi-arm Franka robot

Figure 31 shows visuals of the tasks and Table 17 defines the progress score for each task.

B.4.3. Humanoid robot

Figure 33 shows visuals of the tasks and Table 18 defines the progress score for each task.

Table 16 | Progress Scores: ALOHA Robot (Multi-step benchmark).

Benchmark: ALOHA Robot - Multi-step.		
“Unzip the lunchbag and then place the banana in the lunchbag”.	“Place the red pepper, green apple, and yellow cube in plates with matching colors”.	“Place the red pepper in the red plate, then put green apple in the green plate, then put yellow cube in the yellow plate”.
1.00: if the robot fully unzipped the lunch bag and placed the banana in the lunchbag; 0.50: if the robot fully unzipped the lunch bag; 0.00: if the robot did anything else.	1.00: if the robot placed all three objects correctly in the matching plates; 0.70: if the robot placed two objects on the correct matching plates, but failed for third object; 0.40: if the robot picked one object and successfully placed it on the correct plate, but failed to do that for second object; 0.20: if the robot picked one object but failed to put it on the right plate; 0.00: if the robot didn't approach any of the objects or approached plates directly.	1.00: if the robot placed all three objects correctly in the matching plates; 0.70: if the robot placed two objects on the correct matching plates, but failed for third object; 0.40: if the robot picked one object and successfully placed it on the correct plate, but failed to do that for second object; 0.20: if the robot picked one object but failed to put it on the right plate; 0.00: if the robot didn't approach any of the objects or approached plates directly.
“Open the middle drawer, then left, then right”.	“Open the right drawer, then left, then middle”.	“Open the lid of the trash bin, then pick up the ball and place it inside the trash bin, then close the lid of the trash bin”.
1.00: if the robot opened all three drawers in correct order; 0.60: if the robot opened the middle then left drawer but failed to open the right drawer; 0.30: if the robot opened the middle drawer but failed to open the left drawer next; 0.00: if the robot failed to open the middle drawer first.	1.00: if the robot opened all three drawers in correct order; 0.60: if the robot opened the right then left drawer but failed to open the middle drawer; 0.30: if the robot opened the right drawer but failed to open the left drawer next; 0.00: if the robot failed to open the right drawer first.	1.00: if the robot successfully opened the lid, put the ball in the bin, then closes the lid; 0.75: if the robot failed to close the lid after putting the ball in the trash bin; 0.50: if the robot opened the trash bin lid, picked up the ball, but failed to put the ball in the trash bin; 0.25: if the robot opened the trash bin lid but failed to pick up the ball next; 0.00: if the robot didn't open the trash bin lid.
“Take the green plate out of the white bag, and put the stress ball in the white bag, and push the bag to the top of the table”.	“Put all the yellow fruits on the yellow plate, green fruit on the green plate, red fruit on the red plate”.	
1.00: if the robot completes the entire task successfully; 0.70: if the robot took the green plate out of the bag, put the stress ball in the bag, but failed to push the bag to the top of the table; 0.30: if the robot took the green plate out of the bag but failed to put the stress ball into the bag; 0.00: if the robot failed to take the green plate out of the bag.	1.00: if the robot put all fruits in the correct plate; 0.80: if the robot put 4 fruits in the correct plate; 0.60: if the robot put 3 fruits in the correct plate; 0.20: if the robot put 2 fruits in the correct plate; 0.20: if the robot put 1 fruit in the correct plate; 0.00: if the robot put none of the fruits in the correct plate.	

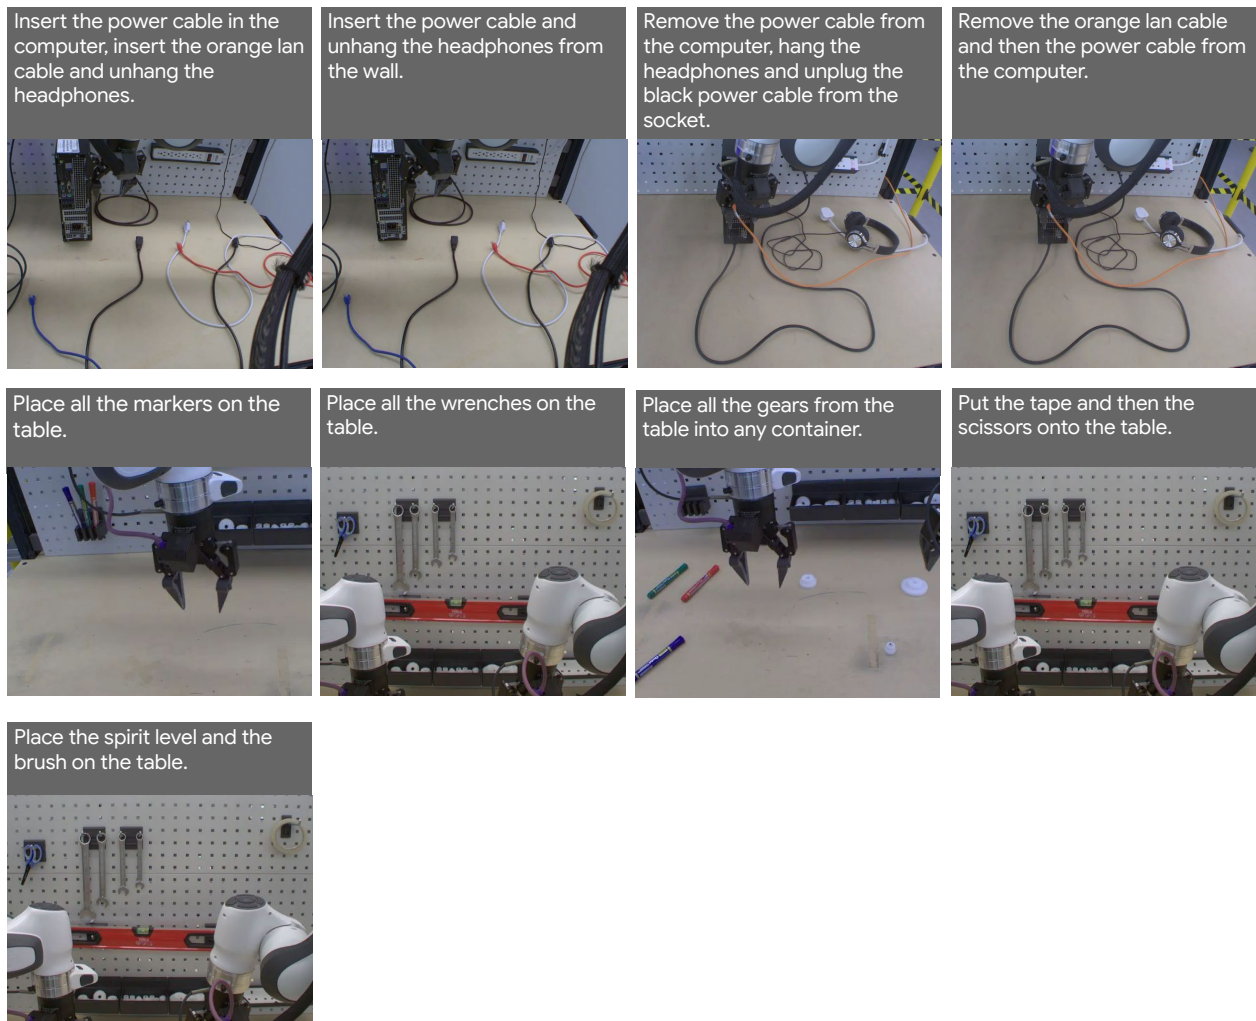


Figure 32 | Tasks in the multi-step benchmark for the Bi-arm Franka.

Table 17 | Progress Scores: Bi-arm Franka (Multi-step benchmark).

Benchmark: Bi-arm Franka - Multi-step.		
“Insert the power cable in the computer, insert the orange lan cable and unhang the headphones”.	“Insert the power cable and unhang the headphones from the wall”.	“Remove the power cable from the computer, hang the headphones and unplug the black power cable from the socket”.
1.00: if the robot successfully completed the third task; 0.75: if the robot successfully completed the remaining task; 0.60: if the robot successfully completed the second task; 0.45: if the robot grasped and attempted one of the remaining two tasks; 0.30: if the robot successfully inserted the power cable or the lan cable or unhang the headphones; 0.15: if the robot grasped and attempted to insert the power cable or the lan cable or unhang the headphones; 0.00: if the robot didn't attempt to insert the power cable or insert the lan cable or unhang the headphones.	1.00: if the robot successfully completed both tasks; 0.75: if the robot attempted to complete the remaining task; 0.50: if the robot successfully inserted the cable or unhang the headphones; 0.25: if the robot attempted to insert the cable or unhang the headphones; 0.00: if the robot didn't do anything.	1.00: if the robot successfully completed the remaining task; 0.75: if the robot grasped and attempted the remaining task; 0.60: if the robot is successful in completing one of the remaining two tasks; 0.45: if the robot grasped and attempted one of the remaining two tasks; 0.30: if the robot is successful in removing the power cable, unplugging the plug from the socket or hanging the headphones; 0.15: if the robot grasped and attempted to either remove the power cable, unplug the plug from the socket or to hang the headphones; 0.00: if the robot didn't attempt to remove the power cable, unplug the plug from the socket or to hang the headphones.
“Remove the orange lan cable and then the power cable from the computer”.	“Place all the markers on the table”.	“Place all the wrenches on the table”.
1.00: if the robot successfully removed both cables; 0.75: if the robot grasped and attempted to remove the remaining cable; 0.50: if the robot successfully removed either the lan cable or the power cable from the computer; 0.25: if the robot grasped and attempted to remove either the lan cable or the power cable from the computer; 0.00: if the robot didn't attempt to remove either cable.	1.00: if the robot placed the remaining marker on the table; 0.80: if the robot grasped the remaining marker; 0.70: if the robot reached for the remaining marker of any colour; 0.67: if the robot placed the second marker on the table; 0.50: if the robot grasped the second marker of any colour; 0.40: if the robot reached for the second marker of any colour; 0.33: if the robot placed the first marker on the table; 0.20: if the robot grasped the first marker of any colour; 0.10: if the robot reached for a first marker of any colour; 0.00: if the robot placed no markers on the table.	1.00: if the robot successfully unhung and placed the fourth wrench on the table; 0.90: if the robot grasped the fourth wrench; 0.85: if the robot reached for a fourth wrench; 0.75: if the robot successfully unhung and placed the third wrench on the table; 0.65: if the robot grasped the third wrench; 0.60: if the robot reached for a third wrench; 0.50: if the robot successfully unhung and placed the second wrench on the table; 0.40: if the robot grasped the second wrench; 0.35: if the robot reached for a second wrench; 0.25: if the robot successfully unhung and placed the first wrench on the table; 0.15: if the robot grasped the first wrench; 0.10: if the robot reached for a first wrench; 0.00: if the robot placed no wrenches on the table.
“Place all the gears from the table into any container”.	“Put the tape and then the scissors onto the table”.	“Place the spirit level and the brush on the table”.
1.00: if the robot placed the remaining gear in any container; 0.80: if the robot grasped the remaining third gear on the table; 0.70: if the robot reached for the remaining third gear on the table; 0.66: if the robot placed the second gear in any container; 0.50: if the robot grasped the second gear; 0.40: if the robot reached for a second gear on the table; 0.33: if the robot placed the first gear in any container; 0.20: if the robot grasped the first gear; 0.10: if the robot reached for a first gear on the table; 0.00: if the robot didn't place or attempt to put back any gears back in the container.	1.00: if the robot placed the scissors on the table; 0.75: if the robot grasped the scissors; 0.60: if the robot reached for the scissors; 0.50: if the robot successfully placed the tape on the table; 0.25: if the robot grasped the tape; 0.10: if the robot reached for the tape; 0.00: if the robot didn't place any item on the table.	1.00: if the robot placed the remaining item on the table; 0.75: if the robot grasped the remaining item; 0.60: if the robot reached for the remaining item; 0.50: if the robot successfully placed the spirit level or the brush on the table; 0.25: if the robot grasped the spirit level or the brush; 0.10: if the robot reached for either the spirit level or the brush; 0.00: if the robot didn't place any item on the table.

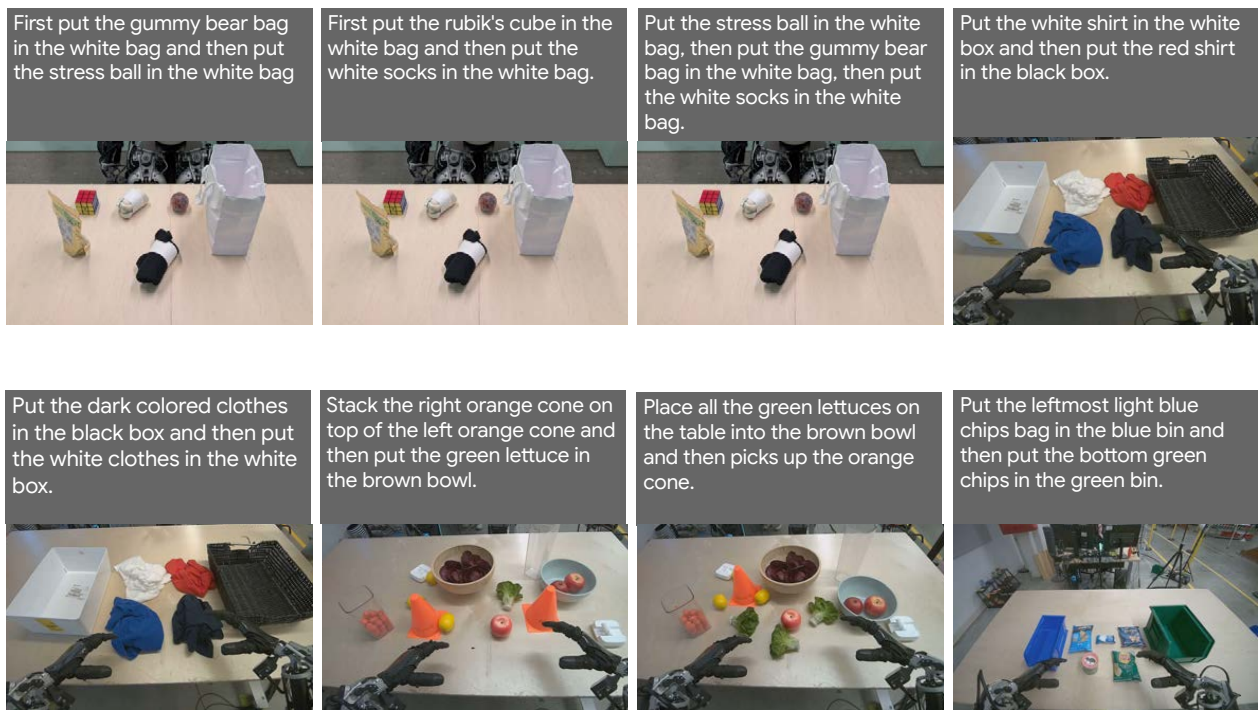


Figure 33 | Tasks in the multi-step benchmark for the Apollo humanoid.

Table 18 | Progress Scores: Humanoid Robot (Multi-step benchmark).

Benchmark: Humanoid Robot - Multi-step.		
“First put the gummy bear bag in the white bag and then put the stress ball in the white bag”.	“First put the rubik’s cube in the white bag and then put the white socks in the white bag”.	“Put the stress ball in the white bag, then put the gummy bear bag in the white bag, then put the white socks in the white bag”.
1.00: if the robot successfully put gummy bear bag into the white bag and then put stress ball into the white bag and only those 2 items; 0.90: if the robot put the gummy bear bag in the white bag and then put the stress ball in the white bag; 0.75: if the robot put the gummy bear bag in the white bag and then picked up the stress ball; 0.60: if the robot put the gummy bear bag in the white bag; 0.45: if the robot picked up the gummy bear bag and did not place it in the white bag; 0.30: if the robot put the stress ball and the gummy bear bag in the white bag but in the wrong order; 0.20: if the robot put the stress ball in the white bag; 0.10: if the robot put something in the white bag but it’s not the stress ball or the gummy bear bag; 0.00: if the robot did anything else.	1.00: if the robot successfully put the rubik’s cube in the white bag and then put the white socks into the white bag and only those 2 items; 0.90: if the robot put the rubik’s cube in the white bag and then put the white socks in the white bag – order matters; 0.75: if the robot put the rubik’s cube in the white bag and then picked up the white socks – order matters; 0.60: if the robot put the rubik’s cube in the white bag; 0.45: if the robot picked up the rubik’s cube and did not place it in the white bag; 0.30: if the robot put the white socks in the bag and then put the rubik’s cube in the white bag; 0.20: if the robot put the white socks in the white bag; 0.10: if the robot put something in the white bag but it’s *not* the rubik’s cube or white socks; 0.00: if the robot did anything else.	1.00: if the robot put the stress ball in the white bag and then put the gummy bears in the white bag and then put the white socks in the white bag; 0.90: if the robot put the stress ball in the white bag and then put the gummy bears in the white bag and then put the white socks in the white bag; 0.75: if the robot put the stress ball in the white bag and then put the gummy bears in the white bag and then picked up the white socks; 0.60: if the robot put the stress ball in the white bag and then put the gummy bears in the white bag; 0.45: if the robot put the stress ball in the white bag and then picked up the gummy bears; 0.30: if the robot put the stress ball in the white bag; 0.15: if the robot picked up the stress ball; 0.00: if the robot did anything else.
“Stack the right orange cone on top of the left orange cone and then put the green lettuce in the brown bowl”.	“Put the dark colored clothes in the black box and then put the white clothes in the white box”.	“Put the white shirt in the white box and then put the red shirt in the black box”.
1.00: if the robot stacked the right orange cone on the left orange cone and then put the green lettuce in the brown bowl; 0.75: if the robot stacked the right orange cone on the left orange cone and then touched the green lettuce; 0.50: if the robot stacked the right orange cone on the left orange cone; 0.25: if the robot touched the right orange cone; 0.00: if the robot did anything else.	1.00: if the robot put 3 colored shirts in the black box, *then* put the white shirt in the white box; 0.75: if the robot put 3 colored shirts in the black box; 0.50: if the robot put 2 colored shirts in the black box; 0.25: if the robot put 1 colored shirt in the black box; 0.00: if the robot did anything else.	1.00: if the robot put the white shirt in the white box, then put the red shirt in the black box; 0.75: if the robot put the white shirt in the white box, then touched the red shirt; 0.50: if the robot put the white shirt in the white box; 0.25: if the robot touched the white shirt; 0.00: if the robot did anything else.
“Place all the green lettuces on the table into the brown bowl and then picked up the orange cone”.	“Put the leftmost light blue chips bag in the blue bin and then put the bottom green chips in the green bin”.	“Put the white clothes in the black box and then put the dark colored clothes in the white box”.
1.00: if the robot put 3 lettuces in the brown bowl, then picked up the orange cone; 0.75: if the robot put 3 lettuce in the brown bowl; 0.50: if the robot put 2 lettuce in the brown bowl; 0.25: if the robot put 1 lettuce in the brown bowl; 0.00: if the robot did anything else.	1.00: if the robot put the leftmost blue chips bag in the blue bin and then put the bottom green chips in the green bin; 0.75: if the robot put the leftmost blue chips bag in the blue bin and then touched the bottom green chips; 0.50: if the robot put the leftmost blue chips bag in the blue bin; 0.25: if the robot put any blue snack in the blue bin; 0.00: if the robot did anything else.	1.00: if the robot put the white shirt in the black box, then put 3 colored shirts in the white box; 0.75: if the robot put the white shirt in the black box, then put 2 colored shirts in the white box; 0.50: if the robot put the white shirt in the black box, then put 1 colored shirt in the white box; 0.25: if the robot put the white shirt in the black box; 0.00: if the robot did anything else.
“First put all the blue snacks in the blue bin and then put all the green snacks in the green bin”.	“Sort the snacks into their matching containers by color”.	“Pack all the gifts”.
1.00: if the robot put 3 blue snacks in the blue bin, then put 2 green snacks in the green bin; 0.80: if the robot put 3 blue snacks in the blue bin, then put 1 green snack in the green bin; 0.60: if the robot put 3 blue snacks in the blue bin; 0.40: if the robot put 2 blue snacks in the blue bin; 0.20: if the robot put 1 blue snack in the blue bin; 0.00: if the robot did anything else.	1.00: if the robot put 5 snacks into their matching color bins; 0.80: if the robot put 4 snacks into their matching color bins; 0.60: if the robot put 3 snacks into their matching color bins; 0.40: if the robot put 2 snacks into their matching color bins; 0.20: if the robot put 1 snack into its matching color bin; 0.00: if the robot did anything else.	1.00: if the robot packed all 5 gifts into the white bag; 0.80: if the robot packed 4 gifts into the white bag; 0.60: if the robot packed 3 gifts into the white bag; 0.40: if the robot packed 2 gifts into the white bag; 0.20: if the robot packed 1 gift into the white bag; 0.00: if the robot did anything else.
“Place the stress ball and the gummy bears in the white bag”.	“First put the stress ball in the white bag and then put the gummy bear bag in the white bag”.	
1.00: if the robot successfully put stress ball and gummy bear bag into the white bag and only those 2 items; 0.80: if the robot successfully put stress ball and gummy bear bag into the white bag but also put other items in the bag; 0.60: if the robot put either the stress ball or the gummy bear bag in the white bag, but not both, and successfully picked up the other item but did not put it in the white bag; 0.40: if the robot put either the stress ball or the gummy bear bag in the white bag, but not both; 0.20: if the robot put something in the white bag but it’s *not* the stress ball or the gummy bear bag; 0.00: if the robot did anything else.	1.00: if the robot successfully put stress ball into the white bag, then put gummy bear bag into the white bag and only those 2 items – order matters; 0.90: if the robot put the stress ball in the white bag and then put the gummy bear bag in the white bag; 0.75: if the robot put the stress ball in the white bag and then picked up the gummy bear bag; 0.60: if the robot put the stress ball in the white bag; 0.45: if the robot picked up the stress ball and did not place it in the white bag; 0.30: if the robot put the gummy bear bag and the stress ball in the white bag but in the wrong order; 0.20: if the robot put the gummy bear bag in the white bag; 0.10: if the robot put something in the white bag but it’s not the stress ball or the gummy bear bag; 0.00: if the robot did anything else.	

B.5. Success rate

In this section we report the *success rate* for all our results in Section 3.

B.5.1. Success rate for generalization performance




Figure 34 | Breakdown of GR 1.5 generalization capabilities across our robots. GR 1.5 consistently outperforms the baselines and handles all four types of variations more effectively.

B.5.2. Success rate for data and model ablation




Figure 35 | Ablation on datasets and training recipes on our robots: GR 1.5 consistently outperforms our baselines: GR 1.5 trained on single or multi-robot data without the MT recipe.

B.5.3. Success rate for thinking ablation




Figure 36 | Ablation of thinking: success rate in the multi-step benchmark with and without enabling thinking during inference.

C. Gemini Robotics-ER 1.5 is a generalist embodied reasoning model

C.1. Evaluation Details: Generality

To assess an overall approximation of model embodied reasoning performance, we evaluate Gemini Robotics-ER 1.5 and other multimodal models on a mix of 15 academic benchmarks. The aggregated results are reported in Fig. 8, and the individual benchmark performance results are shown in Table 19 and Table 20. For text-based VQA evaluation benchmarks, we used Gemini 2.5 Flash to grade response accuracy for both multiple-choice and freeform question formats.

The Gemini 2.5 and GPT-5 models were accessed in between September 1, 2025 and September 20, 2025, using default thinking budgets and without tool use.

Model	GR-ER 1.5	GR-ER 1.5	GR-ER	Gemini 2.5 Pro	Gemini 2.5 Flash	GPT-5	GPT-5-mini
Thinking	Yes	No	No	Yes	Yes	Yes	Yes
Point-Bench	71.6	73.3	75.7	62.7	61.7	43.6	39.5
RefSpatial	48.5	41.8	49.3	33.6	41.2	23.5	23.0
RoboSpatial-Pointing	31.1	25.3	30.3	8.3	7.9	19.0	12.5
Where2Place	59.0	48.0	41.0	37.0	48.0	37.0	33.5
Spatial average	52.6	47.1	49.1	35.4	39.7	30.8	27.1
BLINK	57.8	65.2	60.1	69.2	46.1	71.3	66.4
CV-Bench	84.3	83.6	83.2	85.9	85.5	86.1	85.9
ERQA	54.8	47.0	45.3	56.0	47.5	59.0	57.3
EmbSpatial	78.4	73.4	56.4	78.0	76.2	81.5	78.8
MindCube	54.7	47.7	47.4	59.2	55.4	58.0	55.6
RoboSpatial-VQA	79.3	57.7	66.2	71.3	73.4	69.3	70.7
SAT	76.7	62.0	64.7	74.7	73.3	86.7	81.3
Cosmos-Reason1	72.2	68.3	62.0	73.8	72.1	79.4	76.3
Min Video Pairs	72.5	67.1	59.5	72.8	69.2	77.0	73.0
OpenEQA	55.0	50.5	38.3	55.7	45.3	64.4	59.2
VSI-Bench	45.8	39.9	34.1	51.1	45.3	52.9	46.2
QA average	66.5	60.2	56.1	68.0	62.7	71.4	68.2
ER Score	59.6	53.7	52.6	51.7	51.2	51.1	47.7
Overall Average	62.8	56.7	54.2	59.3	56.5	60.6	57.3

Table 19 | Model performance on a mix of 15 academic embodied reasoning benchmarks. GPT-5 and GPT-5-mini results obtained via API in September 2025.

C.2. Evaluation Details: Pointing

Table 21 shows detailed breakdown of the evaluation for complex pointing.

C.3. Additional Examples

Fig. 37 illustrates sampled thoughts from GR-ER 1.5 on embodied reasoning tasks.

Model	GR-ER 1.5	GR-ER 1.5	GR-ER	Pro 2.5	Flash 2.5	GPT-5	GPT-5-mini
Thinking	Yes	No	No	Yes	Yes	Yes	Yes
MMMU	80.7	79.3	67.0	82.0	79.7	82.0	78.0
GPQA	83.3	81.3	59.6	86.4	82.8	88.4	78.3
Aider Polyglot	57.3	44.4	16.0	82.2	56.7	81.3	66.7
Average	73.8	68.3	47.5	83.5	73.1	83.9	74.3

Table 20 | Model performance on MMMU, GPQA and Aider Polyglot benchmarks. Results for GPT-5 and GPT-5-mini obtained via API with default thinking settings and no tool use in September 2025.

Table 21 | Model performance on complex pointing benchmarks, broken down by subtask.

Model	GR-ER 1.5	GR-ER 1.5	GR-ER	Gemini 2.5 Pro	Gemini 2.5 Flash	GPT-5	GPT-5-mini
Thinking	Yes	No	No	Yes	Yes	Yes	Yes
<i>Standard Pointing</i>							
Point-Bench-Affordance	70.9	76.5	87.9	65.3	67.8	58.1	50.0
Point-Bench-Counting	86.8	86.8	88.4	77.5	73.1	53.7	56.8
Point-Bench-Reasoning	61.7	69.0	64.8	55.4	49.4	33.0	28.3
<i>Steerable Pointing</i>							
Point-Bench-Steerable	67.8	61.8	65.8	53.4	61.3	38.0	32.0
<i>Spatial Pointing</i>							
Point-Bench-Spatial	71.0	72.6	71.9	61.9	57.2	35.4	30.3
RefSpatial	48.5	41.8	49.2	33.6	41.1	23.5	23.0
RoboSpatial	31.1	25.3	30.3	8.3	7.9	19.0	12.5
Where2Place	59.0	48.0	41.0	22.0	48.0	37.0	33.5
<i>Point-to-Count</i>							
PixMo Count	80.0	65.0	60.0	76.0	64.0	73.0	77.0
Average	52.6	47.1	49.1	35.4	39.7	30.8	27.1



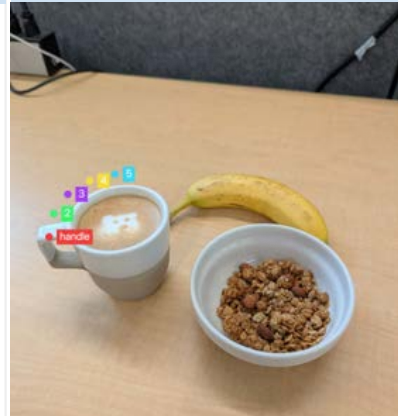

<p>What does the gauge read?</p> 	<p>Select model thoughts</p> <p>Analyze the gauge's features:</p> <ul style="list-style-type: none"> * Scale: The gauge measures temperature in degrees Fahrenheit (°F). This is indicated by "°F" written on the lower part of the dial. * Numbers: The major markings on the scale are 30, 50, 80, 110, 140, 170, 200, 230, 250 <p>Read the pointer position:</p> <ul style="list-style-type: none"> * The pointer is between the numbers 50 and 80. * If the interval between 50 and 80 is 30 degrees, and there are two tick marks dividing it into three equal sections, each section would be 10 degrees. The first tick mark after 50 would represent 60. The second tick mark would represent 70. * The pointer is pointing to the unnumbered tick mark that is immediately after the 50 mark. * Therefore, the value is $50 + 10 = 60$ degrees Fahrenheit. 	<p>Model response</p> <p>Based on the image, the gauge reads approximately 60°F.</p> <p>Here's how to interpret the reading:</p> <ol style="list-style-type: none"> 1. The large numbers on the dial are in increments, such as 30, 50, 80, 110, etc. The unit is marked as °F (degrees Fahrenheit). 2. The pointer is located between 50°F and 80°F. 3. Between 50°F and 80°F, there are two unlabeled tick marks. This divides the 30-degree interval into three sections, with each section representing 10 degrees ($50 + 10 = 60$; $60 + 10 = 70$; $70 + 10 = 80$). 4. The pointer is pointing directly at the first tick mark after 50°F, which represents 60°F.
<p>Match all of the socks by pointing to them and labeling them based on their pattern. Make sure matching socks have the same label.</p> 	<p>Select model thoughts</p> <p>Identify individual socks and their patterns:</p> <ul style="list-style-type: none"> * Sock 3 (center-top): Light blue sock with small white dots. It's partially obscured. Pattern: small white dots. Let's look for a match. * Sock 4 (bottom-center): Light blue sock with small white dots. This matches Sock 3. <p>Create pairs based on matching patterns:</p> <ul style="list-style-type: none"> * Pair 1: Cat face socks. Sock 1 ([y: 250, x: 380]) and Sock 2 ([y: 350, x: 700]). Label: "cat face". * Pair 2: Small dot socks. Sock 3 ([y: 350, x: 580]) and Sock 4 ([y: 800, x: 400]). Label: "small dots". 	<p>Model response</p> 
<p>Point to the handle of the mug. Plot a trajectory that traces the position of the handle as it is rotated in place by 90 degrees.</p> 	<p>Select model thoughts</p> <p>Determine the center of rotation: For "rotated in place," the center of rotation is likely the center of the mug itself. The mug's center is roughly around [500, 250]. The handle starts at approximately [558, 100].</p> <p>Visualize the rotation: The handle is currently at the 9 o'clock position relative to the mug's center. A 90-degree rotation clockwise would move it to the 6 o'clock position (bottom of the mug).</p> <p>Refine trajectory points: The trajectory should be an arc. Let's make it more visually representative of a rotation around the mug's center.</p> <ul style="list-style-type: none"> * Start: [558, 100] (initial handle location) * Point 1: [500, 120] (moving upwards and inwards) * Point 2: [450, 180] (approaching the top-left corner) * Point 3: [400, 250] (arriving at the top of the mug) 	<p>Model response</p> 
<p>I am a robot with a payload of 10lbs. Point to all the objects in the image I am physically able to pick up</p> 	<p>Select model thoughts</p> <p>Dumbbells: There are two purple dumbbells. The label on the dumbbells clearly reads "Prosource Fit 5LBS". This means each dumbbell weighs 5 pounds.</p> <p>Resistance Bands/Loop Bands: There are several resistance bands on the floor. Resistance bands typically weigh very little, usually a few ounces to less than a pound each. A 10 lb payload limit is more than sufficient to lift a single resistance band, or even all visible bands together.</p> <p>Furniture/Chair: There is a wooden piece of furniture (desk or cabinet) on the left and a chair base on the right. These objects are clearly much heavier than 10lbs. A robot with a 10lb payload cannot lift them.</p> <p>Dumbbells: Each dumbbell weighs 5lbs. The robot can lift one dumbbell ($5\text{lbs} \leq 10\text{lbs}$). Can it lift both? $5\text{lbs} + 5\text{lbs} = 10\text{lbs}$. The payload limit is exactly 10lbs. The question asks "pick up", which could mean one or both.</p>	<p>Model response</p> 

Figure 37 | Sample thoughts from GR-ER 1.5 performing embodied reasoning tasks.

D. Gemini Robotics 1.5: A Physical Agent

D.1. Long-horizon benchmarks

Our long-horizon benchmarks evaluate the combination of the GR-ER 1.5 model with the VLA as an autonomous agent. Fig. 38 shows visuals of the 4 tasks in the ALOHA long-horizon benchmark. The progress is scored as the sum of points scored along each subtask (Table 22). Fig. 39 shows visuals of the 4 tasks on the Bi-arm Franka long-horizon benchmark. The progress is scored as the sum of points scored along each subtask (Table 23).

Table 22 | Progress Scores: ALOHA Robot (Long-horizon Benchmark).

Benchmark: ALOHA Robot - Long-horizon.			
Trash Sorting: “Put the compostables into the green bin, the recyclables into the blue bin, and the waste into the black bin”.	Desk Organization: “What is the state of the objects in the table? Return them to their original locations”.	Packing Suitcase: “Put the hat and socks into the suitcase then pack the colorful shirt that’s on the hanger”.	Blocks in drawer: “Open each drawer, and put one block in each drawer”.
0.2 is added per item in the correct bin: <ul style="list-style-type: none"> – red grapes in the green bin; – lettuce leaf in the green bin; – aluminum can in the blue bin; – plastic cup in the blue bin; – energy bar wrapper in the black bin. 	0.2 is added per item in the correct state: <ul style="list-style-type: none"> – red pen in the pen holder; – blue pen in the pen holder; – green marker in the cork tray; – glasses case closed; – laptop closed. 	0.25 is added per: <ul style="list-style-type: none"> – white beanie in the suitcase; – blue socks in the suitcase; – shirt taken off the hanger; – shirt in the suitcase. 	0.11 is added per: <ul style="list-style-type: none"> – left drawer was opened; – any block in left drawer; – left drawer closed; – middle drawer was opened; – any block in middle drawer; – middle drawer closed; – right drawer was opened; – any block in right drawer; – right drawer closed.

Trash Sorting



Put the compostables into the green bin, the recyclables into the blue bin, and the waste into the black bin.

Desk Organization



What is the state of the objects on the table?

Return the objects to their original locations.

Packing Suitcase



Put the hat and socks into the suitcase, then pack the colorful shirt that’s on the hanger

Blocks in drawers



Open each drawer, and put one block in each drawer

Figure 38 | ALOHA long-horizon benchmark.

Table 23 | Progress Scores: Bi-arm Franka (Long-horizon Benchmark).

Benchmark: Bi-arm Franka - Long-horizon.			
Swap: "Swap the sardines and the yellow bottle".	Top shelf to the table: "Put all the objects from the top right shelf onto the table".	Mushroom risotto: "Pack all ingredients for a mushroom risotto into the basket".	Vegetarian with nut allergy: "I am vegetarian and allergic to nuts. Can you put all the food I can't eat into the basket".
0.33 is added per each subtask: <ul style="list-style-type: none"> lemon juice is in the correct location; can of sardines is in the correct location; no unrelated task done. 	0.25 per each subtask: <ul style="list-style-type: none"> rice is on the table; corn is on the table; lemon juice is on table; no unrelated task done. 	0.25 per each subtask: <ul style="list-style-type: none"> mushrooms are in the basket; rice is in the basket; stock cubes are in the basket; no unrelated task done. 	0.33 per each subtask: <ul style="list-style-type: none"> the can of sardines is into the basket; the granola is into the basket; no unrelated task done.

Vegetarian with nut allergy



Swap



Mushroom risotto



Top shelf to the table




Figure 39 | Bi-arm Franka long-horizon benchmark.

D.2. Success rate

Fig. 40 shows the *success rate* for the results in Section 5.

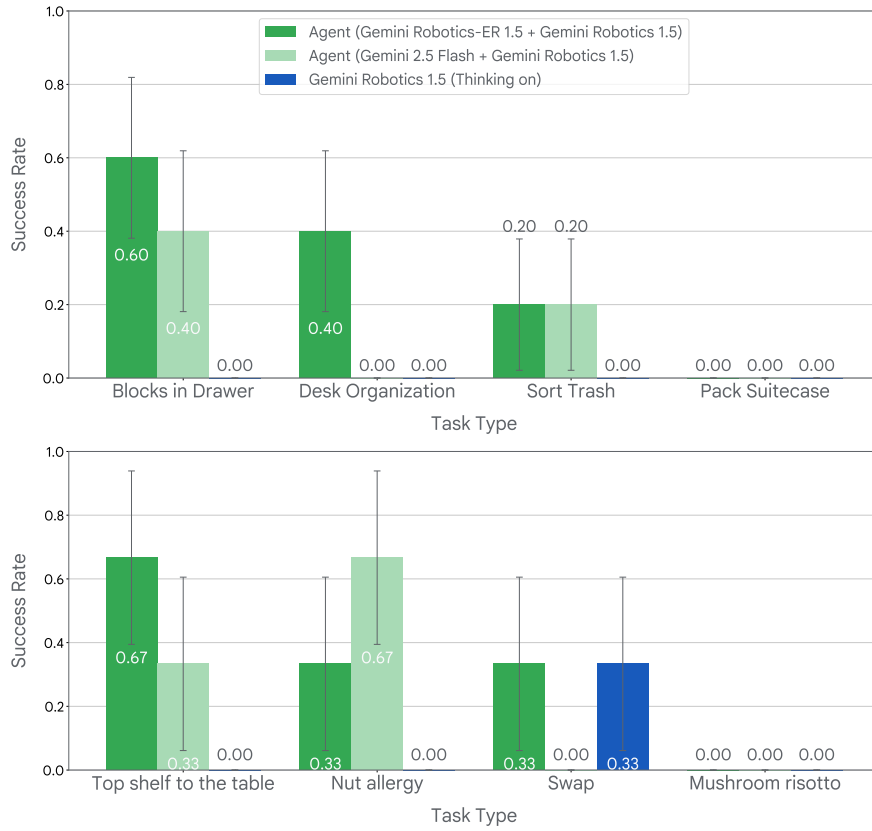


Figure 40 | Long-horizon evaluations for the GR 1.5 Agent on ALOHA (top) and Bi-arm Franka (bottom), consisting of tasks that require advanced real-world understanding, tool use, long-horizon task planning, execution and error recovery to successfully complete the complex long-horizon tasks.