

# VirtualEnv: A Platform for Embodied AI Research

Kabir Swain<sup>1</sup>, Sijie Han<sup>2</sup>, Ayush Raina<sup>3</sup>, Jin Zhang<sup>3</sup>, Shuang Li<sup>4</sup>, Michael Stopa<sup>3</sup>, Antonio Torralba<sup>1</sup>

<sup>1</sup>Massachusetts Institute of Technology

<sup>2</sup>University of Toronto

<sup>3</sup>Sony Interactive Entertainment

<sup>4</sup>Google Deepmind

kswain@mit.edu, hs.han@mail.utoronto.ca, ayush.raina@sony.com, jin.1.zhang@sony.com, shuangligoogle@google.com, michael.stopa@sony.com, torralba@mit.edu

## Abstract

As large language models (LLMs) continue to improve in reasoning and decision-making, there is a growing need for realistic and interactive environments where their abilities can be rigorously evaluated. We present VirtualEnv, a next-generation simulation platform built on Unreal Engine 5 that enables fine-grained benchmarking of LLMs in embodied and interactive scenarios. VirtualEnv supports rich agent–environment interactions, including object manipulation, navigation, and adaptive multi-agent collaboration, as well as game-inspired mechanics like escape rooms and procedurally generated environments. We provide a user-friendly API built on top of Unreal Engine, allowing researchers to deploy and control LLM-driven agents using natural language instructions. We integrate large-scale LLMs and vision-language models (VLMs), such as GPT-based models, to generate novel environments and structured tasks from multimodal inputs. Our experiments benchmark the performance of several popular LLMs across tasks of increasing complexity, analyzing differences in adaptability, planning, and multi-agent coordination. We also describe our methodology for procedural task generation, task validation, and real-time environment control. VirtualEnv is released as an open-source platform, we aim to advance research at the intersection of AI and gaming, enable standardized evaluation of LLMs in embodied AI settings, and pave the way for future developments in immersive simulations and interactive entertainment.

## Introduction

Simulators have become essential tools for researchers to develop, evaluate, and test AI models in controlled, reproducible, and scalable environments (Savva et al. 2019; Kolve et al. 2017; Xia et al. 2018). They offer a cost-effective way to generate large-scale data and enable the study of complex models such as deep neural networks. Simulators are widely used across computer vision (Müller et al. 2018; Shen et al. 2021; Li et al. 2024a; Ge et al. 2024; Li et al. 2024b; Rudek, Valle, and Bertolin 2024) and reinforcement learning (Ferigo et al. 2020; Dosovitskiy et al. 2017; Brockman et al. 2016; Tassa et al. 2018; Kaup et al. 2024) for tasks such as scene understanding, robotic navigation, and object

interaction. In parallel, simulation platforms have increasingly been adopted in gaming research for character control (Parberry 2017; Todorov, Erez, and Tassa 2012) and procedural gameplay mechanics (Bel and Vimont 2011; Interactive and Eidos 2016; Pérez et al. 2015).

Despite these advances, existing simulators remain limited in scale, diversity, and interactivity. Many focus exclusively on small indoor household settings (e.g., VirtualHome (Puig et al. 2018), House3D (Wu et al. 2018), AI2-THOR (Kolve et al. 2017)), with rigid environments and static object arrangements. These constraints hinder progress on tasks requiring generalization, planning, and emergent behavior. Simulators designed for gaming typically offer higher visual fidelity but often lack the modularity, programmability, and semantic richness required for embodied AI research. As research increasingly incorporates large language models (LLMs) and vision-language models (VLMs) into embodied settings, there is a growing need for flexible simulation platforms that support multimodal grounding, interactive task generation, and dynamic environment editing at scale.

To address these limitations, we introduce **VirtualEnv**, a next-generation simulation platform built on Unreal Engine 5 (Games 2024), designed to support language-driven and multimodal research in embodied AI. VirtualEnv offers a modular framework for simulating expansive, richly interactive environments that span urban settings, multi-room buildings, and outdoor spaces. The platform supports fine-grained agent-object interactions, dynamic scene editing, and procedural environment generation through integration with LLMs and VLMs (See Figure 1). Unlike prior simulators, VirtualEnv enables real-time interaction with large-scale environments and supports a wide range of tasks, including spatial reasoning, tool use, goal-conditioned planning, and multi-agent collaboration.

To demonstrate its capabilities, we introduce a suite of Escape Room–style environments (Heikkinen and Shumeyko 2016), where LLM-driven agents are tasked with solving cognitive puzzles that require multi-step reasoning, object manipulation, and sequential planning. These scenarios serve as benchmarks for evaluating language models in grounded, task-oriented environments with escalating difficulty levels. Our experiments compare several vLLMs

**System Overview** This system enables multi-agent collaboration for task solving in a simulated environment using LLMs to plan and issue actions to agents.

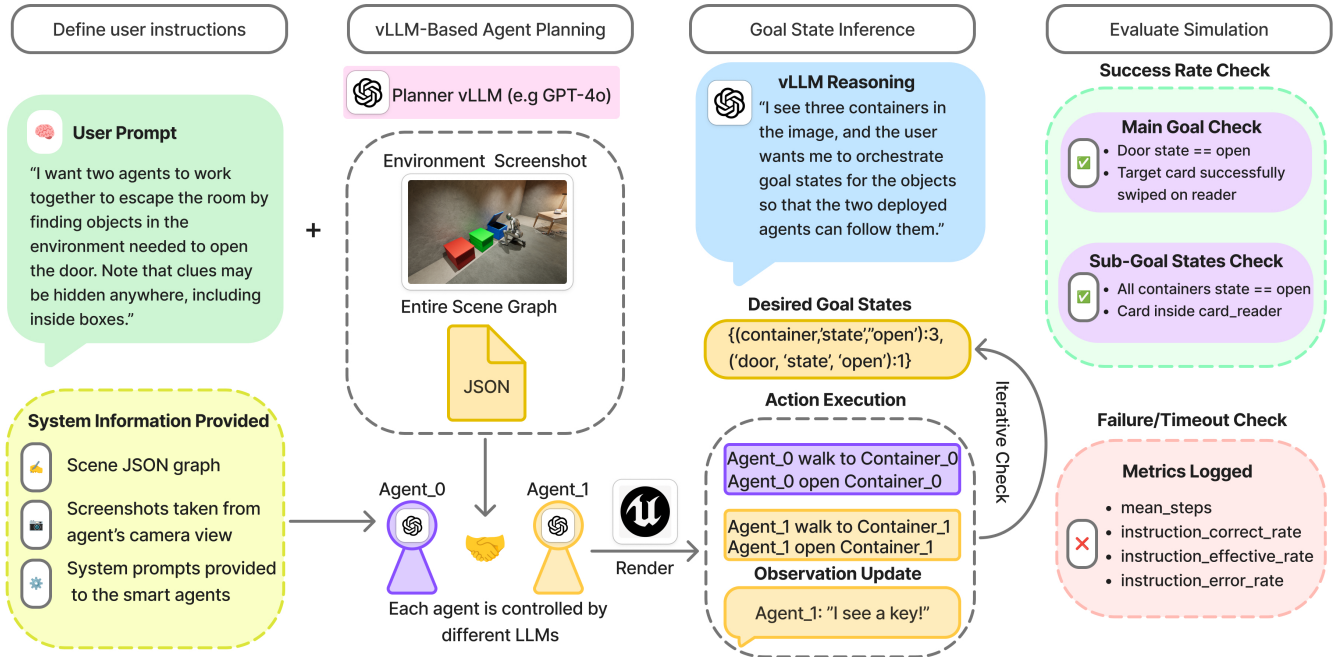


Figure 1: System overview of multi-agent planning and execution in VirtualEnv. VirtualEnv is a high-fidelity simulation environment built on Unreal Engine 5, designed for evaluating large language models (LLMs) in interactive, task-oriented settings. This figure shows the full pipeline for multi-agent task execution: users provide natural language instructions, and the system uses vLLMs (e.g., GPT-4o) to interpret goals, generate symbolic plans, and coordinate multiple agents in real time. Each agent receives environment information—including a scene graph and visual context—and executes actions via the VirtualEnv API. Performance is evaluated through goal completion checks and instruction-based success metrics.

across task success rate, instruction following, and generalization under environmental variation.

By releasing VirtualEnv as an open-source platform, we aim to provide the community with a standardized testbed for multimodal learning, embodied reasoning, and simulation-driven AI development. In particular, we position VirtualEnv as a foundation for benchmarking large language models in interactive environments—enabling consistent, reproducible comparisons across tasks, modalities, and levels of embodiment. We hope VirtualEnv accelerates progress at the intersection of AI, gaming, and simulation, offering a flexible foundation for research in language-guided agents, procedural task generation, and virtual environment control.

## Related Work

Simulators have become indispensable tools in computer vision, embodied AI, and reinforcement learning research, serving as environments for data generation, algorithm validation, and performance benchmarking. Several platforms have been developed, each tailored to specific research goals. Popular examples include VirtualHome (Puig et al. 2018), AI2-THOR (Kolve et al. 2017), OmniGibson (Li et al. 2024a), Habitat (Savva et al. 2019), and ProcTHOR (Deitke et al. 2022). While these simulators exhibit varying degrees of interactivity, realism, and complexity, they typi-

cally emphasize specific capabilities such as navigation, manipulation, or visual perception, limiting their applicability to broader AI research tasks.

**VirtualHome** (Puig et al. 2018) is designed explicitly to simulate household environments, providing structured scene graphs and scripted activities for everyday tasks such as cooking and cleaning. Although influential in facilitating high-level reasoning and task planning, VirtualHome is constrained by its exclusive focus on indoor residential settings, limiting the scope and diversity of potential research scenarios.

**AI2-THOR** (Kolve et al. 2017) expands the range of embodied AI research through interactive indoor scenarios that support visual question answering, navigation, and object manipulation. Despite its extensive adoption, AI2-THOR similarly emphasizes indoor environments, lacking the broader urban or outdoor contexts necessary for more generalizable AI agent development.

**OmniGibson** (Li et al. 2024a) extends interactive simulations by including dynamic object states, physics-driven interactions, and virtual reality interfaces to facilitate transfer to real-world settings. Its detailed indoor scenarios have proven useful for navigation and manipulation tasks. Nevertheless, its emphasis remains largely domestic, limiting ap-



Figure 2: Core capabilities of VirtualEnv. The platform supports highly realistic and interactable indoor and outdoor environments, a large curated library of over 20,000 diverse assets, and controllable humanoid agents with fine-grained motion support. These features enable the creation of complex, multimodal simulation scenarios suitable for embodied AI training, evaluation, and benchmarking.

	Environment	Multi-Agent	Language	Action Space	Task Types	Num tasks
AI2Thor (Kolve et al., 2019)	3D-S		✓	HL	CST	48,000
OmniGibson (Li et al., 2023a)	3D-M		✓	LL+HL	CST	1,000
VirtualHome (Puig et al., 2021)	3D-M	✓		HL	C	1,200
Habitat 3.0 (Puig et al., 2023)	3D-M	✓	✓	LL+HL	CSTH	100,000
<b>VirtualEnv</b>	3D-MIO	✓	✓	HL	CSTH	140,000

Table 1: Overview of Embodied AI simulation platforms and their key features. We position **VirtualEnv** alongside widely used simulators across several dimensions: scene complexity (single-room (S), multi-room (M), indoor-outdoor (IO)), support for multiple agents, language interaction, and action space (high-level (HL), low-level (LL)). Task coverage is annotated using the following categories: constraint-free (C), spatial (S), temporal (T), and heterogeneous (H). "Num tasks" reflects the approximate number of unique, predefined or generated scenarios supported by each platform.

plicability to complex urban environments and large-scale AI tasks involving diverse interactions.

**Habitat** (Savva et al. 2019) provides a high-performance engine primarily optimized for navigation and exploration, enabling scalable training and evaluation of embodied agents. However, Habitat’s primary limitation lies in its restricted interactivity and narrow task scope, making it insufficient for complex reasoning, planning, or rich object-agent interactions required in broader AI scenarios.

**ProcTHOR** (Deitke et al. 2022) introduces a procedural generation framework for AI2-THOR, enabling large-scale creation of diverse indoor scenes via compositional layouts. This approach improves generalization by increasing environmental diversity, yet remains confined to constrained household domains. Moreover, its generated scenes are primarily geometric variants rather than semantically grounded or multimodal in origin.

In summary, existing simulators including VirtualHome, AI2-THOR, OmniGibson, Habitat, and ProcTHOR have significantly advanced embodied AI research by offering specialized environments tailored to specific domains. Yet, their scope and interactivity constraints underscore the need for a more comprehensive simulator capable of addressing diverse research demands simultaneously. VirtualEnv responds to these limitations by offering a unified, highly interactive, and scalable urban simulation platform that incorporates vision-language models (VLMs) and large lan-

guage models (LLMs) to procedurally generate semantically grounded environments and tasks. This enables the study of advanced cognitive behaviors, adaptive decision-making, and systematic benchmarking of LLMs in complex, dynamic scenarios.

## VirtualEnv

In this section, we introduce the main features of VirtualEnv, built on Unreal Engine 5 to support complex agent-object interactions and procedurally generated environments. VirtualEnv includes diverse indoor and outdoor settings, multimodal sensing capabilities, and high-resolution object models. The platform is tightly integrated with large language models (LLMs) through a lightweight Python API, enabling agents to interpret instructions, plan actions, and dynamically interact with the environment using language. We describe core components such as scene representation and interactive task design in Section , and demonstrate how VirtualEnv enables escape room-style tasks for evaluating problem-solving and emergent AI behaviors.

### Main Features of the VirtualEnv Simulator

Figure 2 showcases VirtualEnv’s three core capabilities: photorealistic indoor-outdoor environments, an extensive library of 20,000+ interactive assets, and precise control over humanoid agents. These features form the foundation for our comprehensive embodied AI research platform.

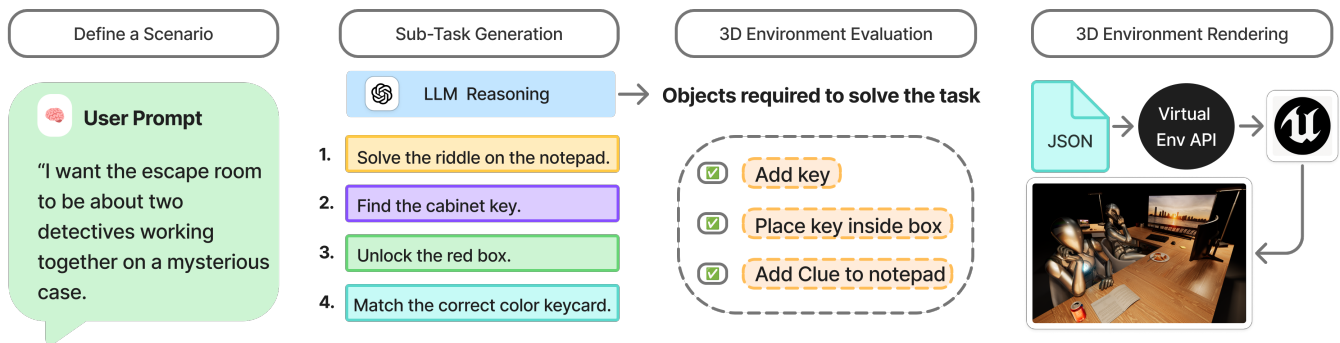


Figure 3: Language-based task and scenario generation in VirtualEnv. A user provides a natural language prompt describing a high-level scenario, which is parsed by a vLLM into sub-tasks and structured goals (e.g., solving riddles, unlocking containers). Based on the task requirements, VirtualEnv automatically evaluates which objects are needed, updates the scene graph accordingly, and renders the environment through the VirtualEnv API. This pipeline enables flexible and scalable creation of interactive, goal-driven scenarios without manual scripting.

As shown in Table 1, VirtualEnv uniquely combines multi-agent support and language interaction capabilities while offering the most comprehensive environment type (3D-MIO) and the largest task library (140,000 tasks) among existing platforms. This combination enables complex scenarios that span both indoor and outdoor settings, supporting a broader range of embodied AI research than single-room (AI2Thor) or multi-room-only (Habitat 3.0, VirtualHome) alternatives.

**High-Fidelity Engine:** VirtualEnv provides a highly dynamic and interactive AI research environment. Built on Unreal Engine 5 (Games 2024), it features diverse settings, including offices, retail venues, and urban streetscapes. At its core, VirtualEnv emphasizes realism and complex agent-environment interactions, using advanced rendering pipelines and procedural generation to enable boundless variations in physical layouts, object placements, and lighting conditions.

**Rich Object and Action Library:** With over 20,000 distinct objects, VirtualEnv supports diverse real-world scenarios such as home furnishing, household activities, urban navigation, and multi-step decision-making. Each object is embedded with affordances—e.g., openable doors, movable furniture, graspable objects, and interactive appliances—allowing agents to perform fine-grained interactions. Many objects use photogrammetry scans for high-resolution modeling, ensuring realistic physics and visual accuracy. Unreal Engine’s physics engine enables authentic object responses to interactions such as movement, deformation, and state transitions. Additionally, structured metadata allows AI agents to reason about object properties, enhancing their ability to learn real-world physical interactions.

**Multi-Modal Sensing and Observations:** VirtualEnv offers rich multimodal sensing capabilities to support perception and decision-making in dynamic environments. Agents access RGB and depth sensors for photorealistic input and spatial understanding, as well as semantic segmentation for pixel-level object recognition. Panoramic top-down views further aid spatial reasoning and large-scale navigation.

These modalities together enable agents to interpret and act within complex scenes with greater precision and adaptability.

**User-Friendly API and Language-Driven AI Agents:** VirtualEnv natively supports the integration of large language models (LLMs) and vision-language models (VLMs), allowing AI agents to interact with the world through natural language. This enables flexible task execution, dynamic decision-making, and interactive environment control based on high-level instructions. Researchers can explore how AI models interpret, respond to and act on language commands in real time, supporting advances in LLM-based robotics and embodied language understanding.

**Scene Graph Representation:** VirtualEnv organizes its environments using a scene graph, which encodes objects, agents, and spatial relationships in a hierarchical structure. This representation allows for efficient querying of environmental states, enabling agents to make informed decisions based on their surroundings. It also supports semantic reasoning, allowing agents to understand object affordances and spatial constraints, improving their ability to interact with the environment meaningfully. Additionally, the scene graph facilitates partial observations, making it possible to study agent behavior in scenarios where only limited information is available, which is crucial for research on planning under uncertainty. By providing a programmable scene graph API, VirtualEnv simplifies the process of creating and modifying scenes and tasks, offering researchers greater flexibility in designing interactive and adaptive simulation environments.

**Environment Construction and Asset Selection:** All environments in VirtualEnv were constructed using a hybrid approach that combined manual scene authoring and procedural generation within Unreal Engine 5. We curated high-resolution 3D assets from the Unreal Engine Marketplace, selecting those with rich affordances (e.g., openable, graspable, movable) and clear semantic categories. These assets were chosen based on their relevance to embodied AI tasks—such as indoor navigation, object manipulation, and

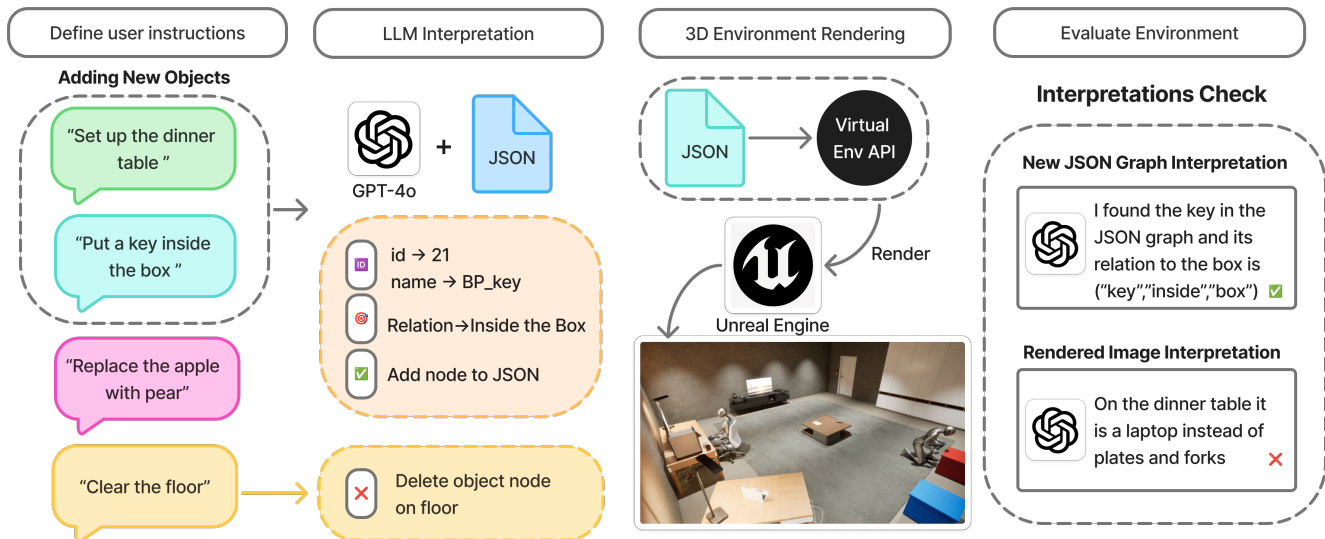


Figure 4: Interactive environment editing and interpretation validation in VirtualEnv. Users provide natural language instructions to modify the environment (e.g., adding, replacing, or removing objects). A vLLM interprets these instructions and updates the scene graph, which is then rendered using the VirtualEnv API. To ensure semantic alignment between the symbolic scene graph and the visual output, the system performs interpretation checks on both the JSON graph and the rendered image. This process enables reliable, language-driven environment manipulation and validation.

goal-directed behavior—and were organized into thematically consistent scenes (e.g., offices, kitchens, retail stores, and urban outdoor environments).

**Language Driven Task and Scenario Generation:** VirtualEnv supports dynamic scenario generation through a language based interface. Users provide natural language prompts describing desired task setups, such as escape room challenges or collaborative problem solving scenarios. An LLM interprets the prompt and decomposes it into a sequence of subgoals or puzzles (e.g., “find the key,” “solve a riddle,” “unlock the box”). Based on these subgoals, the system identifies the required environmental components including objects, clues, and spatial arrangements, and automatically updates the scene graph to instantiate them. The environment is then rendered through the VirtualEnv API, resulting in a complete interactive scene where agents can be evaluated on their ability to complete the language defined tasks. This pipeline allows for scalable and diverse environment generation grounded in multimodal reasoning and task aware scene construction. (See Figure 3)

### Escape Room Challenge Framework

To evaluate higher-level reasoning in VirtualEnv, we introduce an Escape Room Challenge Framework. Unlike simple navigation or retrieval tasks, escape rooms blend puzzle-solving, object interactions, narrative clues, and occasional multi-agent coordination. Agents must discover clues in the environment, connect information across the scene, and solve an overarching puzzle, encouraging more flexible and deliberate behavior.

**Rationale and Design.** Puzzle design in games provides a useful structure for planning and problem-solving because

each puzzle is new and cannot be solved through memorization. Many puzzles also involve abstract or fictional objects, which require agents to adapt to unfamiliar situations. Incorporating these ideas into VirtualEnv encourages exploration, strategy refinement, and responsiveness to environmental feedback. Our framework integrates cognitive puzzles, interaction mechanics, and narrative hints, following the experience-pyramid model of (Heikkinen and Shumeyko 2016). Agents must navigate, manipulate objects, and make decisions as the puzzle unfolds.

**Levels of Complexity:** We categorize our escape room challenges into four difficulty levels based on puzzle length and inter-clue dependencies, progressively increasing cognitive demands:

1. **Level 1 - One Step Problem:** A single clue leads directly to the key, requiring minimal inference. Agents parse a textual hint and execute a basic action sequence to unlock the door.
2. **Level 2 - Sequential Puzzles:** Agents must complete an intermediate task (e.g., arranging colored objects) to reveal the real clue, introducing multi-step reasoning.
3. **Level 3 - Meta Clues:** Two parallel puzzles generate separate clues, both necessary for finding the key. Agents must integrate multiple information sources, reinforcing contextual reasoning.
4. **Level 4 - Deceptive Clues:** Agents receive two clues, one accurate and one misleading. They must determine which clue is correct. These puzzles challenge critical thinking, error-checking, and contextual understanding.

**Implications for AI Research:** By embedding escape-room puzzles in VirtualEnv’s realistic, dynamic settings, we



Figure 5: Comparison of Visual Realism Rankings Across Platforms. A qualitative benchmarking study was conducted using a survey with 31 respondents. Participants ranked each platform based on visual realism, assigning a score from 5 (most realistic) to 1 (least realistic) in a label-blind test. We observe that the participants consider **VirtualEnv** to be significantly more visually realistic than all other environments.




Figure 6: Distribution of Failure Modes in Embodied AI Tasks. Analysis of failure modes reveals six primary categories.

bridge embodied AI and next-generation gaming applications. Success in these tasks signals advances in AI reasoning, contextual comprehension, and adaptive problem-solving. Additionally, our flexible puzzle design provides a framework to explore LLM fine-tuning, embodied AI integration, and generalization to novel tasks. As AI continues to evolve, these benchmarks will help assess how well AI agents adapt to unfamiliar, multi-step challenges in real-time interactive environments.

### Environment Modification with vLLM

As illustrated in Figure 4, VirtualEnv uses a vision-language model (vLLM) to modify an existing 3-D scene in response to natural-language commands. The model converts a prompt (e.g., “put a key inside the box”) into JSON-encoded edits that specify the target objects, spatial relations, and placement rules. These edits are merged into

the current observation graph and rendered in Unreal Engine 5. An interpretation check then compares the symbolic graph with the rendered view, flagging any mismatches between the intended and visualized states. This language-guided editing pipeline enables rapid, precise adjustment of complex environments without manual intervention.

## Experiments

Our experiments begin with a comprehensive evaluation of VirtualEnv’s visual fidelity through a qualitative benchmarking study, comparing it against leading simulation platforms in the field. We then assess the effectiveness of LLM-based planners within this high-fidelity environment, focusing on their ability to make structured decisions, coordinate tasks, and adapt to dynamic environments in both single-agent and multi-agent scenarios.

To achieve this, we design controlled experiments where agents operate in partially observable environments. The single-agent experiments evaluate fundamental skills such as navigation, object retrieval, and environmental manipulation, while the multi-agent experiments examine how collaborative planning improves task efficiency in scenarios requiring synchronized actions and division of labor.

To contextualize our findings, we benchmark agent performance within VirtualEnv, evaluating generalization across different task complexities. We assess each planner’s ability to adapt to new tasks, using structured comparisons to measure how well agents scale to increasing environmental challenges, including both routine household tasks and complex puzzle-solving scenarios.

Task	Reasoning LLMs				Non-Reasoning LLMs			
	Claude 3 Opus	Gemini 2.5 Pro	o3	Grok 3 Think	GPT-4o	Llama 3.1 405B	Qwen 2.5 Max	Llama 4 (MoE)
Clean Floor (S)	0.85±0.03	0.83±0.04	0.82±0.04	0.76±0.05	0.68±0.05	0.62±0.06	0.61±0.05	0.60±0.06
Watch TV (S)	0.88±0.02	0.86±0.03	0.85±0.03	0.80±0.04	0.72±0.04	0.66±0.05	0.65±0.05	0.64±0.05
Find Object (S)	0.70±0.05	0.68±0.05	0.64±0.06	0.60±0.06	0.48±0.06	0.46±0.07	0.45±0.07	0.40±0.08
Prepare Food (M)	0.92±0.03	0.90±0.03	0.88±0.04	0.84±0.04	0.75±0.04	0.70±0.05	0.69±0.05	0.68±0.05
Clean Room (M)	0.93±0.02	0.92±0.03	0.90±0.03	0.86±0.04	0.78±0.04	0.74±0.05	0.73±0.05	0.72±0.05

Table 2: Performance Comparison of Reasoning vs. Non-Reasoning LLMs on Embodied Tasks. Success rates ( $\pm 1$  Standard Deviation) across five benchmark tasks, including both single-agent (S) and multi-agent (M) scenarios. Reasoning LLMs consistently outperform their non-reasoning counterparts, with the performance gap being most pronounced in complex tasks like *Find Object* and *Prepare Food*. Multi-agent tasks generally show higher success rates, demonstrating the benefits of collaborative planning.

## Visual Realism

To evaluate the visual realism of VirtualEnv in comparison to existing simulation platforms, we conducted a qualitative benchmarking study. Participants (N=31) were asked to rank multiple platforms (VirtualEnv, OmniGibson, AI2THOR, VirtualHome, and Habitat) based on visual realism through a label-blind survey. Each platform was rated from 5 (most realistic) to 1 (least realistic). As shown in Figure 5, VirtualEnv achieved a significantly higher realism score ( $4.46 \pm 1.02$ ) compared to other platforms, clearly demonstrating its advantage in generating visually realistic environments for embodied AI tasks.

## Baseline Model Performance Analysis

**“Reasoning” vs. “Non-Reasoning” LLMs.** As shown in Table 2, our experiments compare four LLM variants with chain-of-thought capabilities against their base models across five distinct tasks. The chain-of-thought models show an average improvement of 11% in task completion rates, with particularly strong gains in complex, multi-step activities like *Find Object* and *Prepare Food*. These improvements stem from the models’ enhanced ability to break down tasks into logical steps and maintain context throughout execution. The performance is also more consistent, with standard deviations below 0.05 for routine tasks, suggesting that structured reasoning leads to both higher success rates and more reliable performance.

**Task-specific difficulty profile.** Success rates vary significantly across tasks. While *Watch TV* achieves high performance (above 0.85) for state-of-the-art models, the open-ended search required in *Find Object* reduces performance by up to 25 percentage points and nearly doubles variance ( $\sigma = 0.06 - 0.08$ ). This suggests that partial observability remains the primary challenge, even with photorealistic rendering and large language priors. To address this limitation, we propose augmenting the planner with explicit spatial memory or learned exploration heuristics.

## Collaborative Planning Analysis

Building on our single-agent findings, we observe that multi-agent collaboration consistently improves performance. For

instance, on *Prepare Food*, Claude 3 Opus’s success rate increases from 0.88 to 0.92, while GPT-4o improves from 0.68 to 0.75. Analysis of replay logs reveals that this improvement stems from effective task allocation. For example, one agent would handle utensil retrieval while another manages appliance operation, thereby reducing action horizons and minimizing occlusion-related uncertainty.

## Failure Modes

Our analysis identifies six key failure modes in embodied AI tasks, with their distribution shown in Figure 6. The most common failure (30.4%) occurs when agents get stuck in exploration loops, repeatedly visiting the same rooms without finding their target. This happens because agents lack effective strategies for systematically exploring unseen areas. The second most frequent issue (18.5%) involves agents pursuing non-existent objects, often because their planning process loses track of what’s actually available in the environment. Other significant failures include: agents incorrectly assuming object states (15.2%), multi-agent coordination problems (14.1%), physically impossible action sequences (12.0%), and confusion between similar-looking objects (9.8%). These failure modes have clear implications for improving embodied AI systems. The top three categories, exploration loops, phantom goals, and state tracking errors account for nearly two-thirds of all failures. Addressing these core issues could potentially improve overall task success rates by 7.4%, bringing the best-performing models closer to human-level performance on routine tasks.

## Conclusion

VirtualEnv is a next-generation Unreal Engine 5 simulation platform for embodied AI and language-driven interaction. It offers scalable, richly interactive environments where agents perform complex tasks such as navigation, multi-step manipulation, collaboration, and goal-directed planning. Powered by vision-language models, it supports procedural scenario generation, language-guided scene editing, and task planning in dynamic environments, providing improved interactivity, diversity, and visual fidelity over existing platforms and a strong foundation for future work at the intersection of AI and simulation.

## Acknowledgments

We would like to thank Xavier Puig for his helpful insight and advice for creating a virtual environment designed for Embodied AI.

## References

- Bel, R.; and Vimont, B. 2011. Developing the interactive dynamic natural world of "From Dust". In *ACM SIGGRAPH 2011 Talks*, 1–1. New York, NY, United States: Association for Computing Machinery.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym: A Toolkit for Developing and Comparing Reinforcement Learning Algorithms.
- Deitke, M.; Vander Bilt, E.; Herrasti, A.; Weihs, L.; Salvador, J.; Ehsani, K.; Han, W.; Kolve, E.; Farhadi, A.; Kembhavi, A.; and Mottaghi, R. 2022. ProcTHOR: large-scale embodied AI using procedural generation. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781713871088.
- Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An open urban driving simulator. In *Conference on robot learning*, 1–16. PMLR.
- Ferigo, D.; Traversaro, S.; Metta, G.; and Pucci, D. 2020. Gym-ignition: Reproducible robotic simulations for reinforcement learning. In *2020 IEEE/SICE International Symposium on System Integration (SII)*, 885–890. IEEE.
- Games, E. 2024. Unreal Engine 5. <https://www.unrealengine.com/en-US/unreal-engine-5>. Accessed: 2026-01-16.
- Ge, Y.; Tang, Y.; Xu, J.; Gokmen, C.; Li, C.; Ai, W.; Martinez, B. J.; Aydin, A.; Anvari, M.; Chakravarthy, A. K.; et al. 2024. BEHAVIOR Vision Suite: Customizable Dataset Generation via Simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22401–22412.
- Heikkinen, O. K.; and Shumeyko, J. 2016. Designing an escape room with the Experience Pyramid model.
- Interactive, I.; and Eidos, W. P. 2016. Hitman. *Computer Game], Square Enix, Tokyo*.
- Kaup, M.; Wolff, C.; Hwang, H.; Mayer, J.; and Bruni, E. 2024. A review of nine physics engines for reinforcement learning research. *arXiv preprint arXiv:2407.08590*.
- Kolve, E.; Mottaghi, R.; Gordon, D.; Zhu, Y.; Gupta, A.; and Farhadi, A. 2017. AI2-THOR: An Interactive 3D Environment for Visual AI. *CoRR*, abs/1712.05474.
- Li, C.; Zhang, R.; Wong, J.; Gokmen, C.; Srivastava, S.; Martín-Martín, R.; Wang, C.; Levine, G.; Ai, W.; Martinez, B.; Yin, H.; Lingelbach, M.; Hwang, M.; Hiranaka, A.; Garlanka, S.; Aydin, A.; Lee, S.; Sun, J.; Anvari, M.; Sharma, M.; Bansal, D.; Hunter, S.; Kim, K.-Y.; Lou, A.; Matthews, C. R.; Villa-Renteria, I.; Tang, J. H.; Tang, C.; Xia, F.; Li, Y.; Savarese, S.; Gweon, H.; Liu, C. K.; Wu, J.; and Fei-Fei, L. 2024a. BEHAVIOR-1K: A Human-Centered, Embodied AI Benchmark with 1,000 Everyday Activities and Realistic Simulation. *arXiv:2403.09227*.
- Li, X.; Park, J.; Reberg-Horton, C.; Mirsky, S.; Lobaton, E.; and Xiang, L. 2024b. Photorealistic Arm Robot Simulation for 3D Plant Reconstruction and Automatic Annotation using Unreal Engine 5. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5480–5488.
- Müller, M.; Casser, V.; Lahoud, J.; Smith, N.; and Ghanem, B. 2018. Sim4cv: A photo-realistic simulator for computer vision applications. *International Journal of Computer Vision*, 126: 902–919.
- Parberry, I. 2017. *Introduction to Game Physics with Box2D*. CRC Press.
- Pérez, L. J. F.; Calla, L. A. R.; Valente, L.; Montenegro, A. A.; and Clua, E. W. G. 2015. Dynamic game difficulty balancing in real time using evolutionary fuzzy cognitive maps. In *2015 14th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*, 24–32. IEEE.
- Puig, X.; Ra, K.; Boben, M.; Li, J.; Wang, T.; Fidler, S.; and Torralba, A. 2018. VirtualHome: Simulating Household Activities via Programs. *CoRR*, abs/1806.07011.
- Rudek, M.; Valle, A. P.; and Bertolin, R. 2024. Building Realistic Environment from Computer Vision Approach Applied to Manufacturing Simulation in the Digital Twin Context. In *International Conference on Innovative Intelligent Industrial Production and Logistics*, 223–235. Springer.
- Savva, M.; Kadian, A.; Maksymets, O.; Zhao, Y.; Wijmans, E.; Jain, B.; Straub, J.; Liu, J.; Koltun, V.; Malik, J.; Parikh, D.; and Batra, D. 2019. Habitat: A Platform for Embodied AI Research. *CoRR*, abs/1904.01201.
- Shen, B.; Xia, F.; Li, C.; Martín-Martín, R.; Fan, L.; Wang, G.; Pérez-D'Arpino, C.; Buch, S.; Srivastava, S.; Tchapmi, L.; Tchapmi, M.; Vainio, K.; Wong, J.; Fei-Fei, L.; and Savarese, S. 2021. iGibson 1.0: A Simulation Environment for Interactive Tasks in Large Realistic Scenes. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 7520–7527.
- Tassa, Y.; Doron, Y.; Muldal, A.; Erez, T.; Li, Y.; Casas, D. d. L.; Budden, D.; Abdolmaleki, A.; Merel, J.; Lefrancq, A.; et al. 2018. Deepmind control suite. *arXiv preprint arXiv:1801.00690*.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.
- Wu, Y.; Wu, Y.; Gkioxari, G.; and Tian, Y. 2018. Building generalizable agents with a realistic and rich 3D environment. *arXiv preprint arXiv:1801.02209*.
- Xia, F.; Zamir, A. R.; He, Z.; Sax, A.; Malik, J.; and Savarese, S. 2018. Gibson Env: Real-World Perception for Embodied Agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9068–9079.