Scaling Environments for LLM Agents in the Era of Learning from Interaction: A Survey

Anonymous ACL submission

Abstract

LLM-based agents can autonomously accomplish complex tasks across various domains. However, to further cultivate capabilities such as adaptive behavior and long-term decision-making, training on static datasets built from human-level knowledge is insufficient. These datasets are costly to construct and lack both dynamism and realism. A growing consensus is that agents should instead interact directly with environments and learn from experience through reinforcement learning. We formalize this iterative process as the Generation–Execution–Feedback (GEF) loop, where environments generate tasks to challenge agents, return observations in response to agents' actions during task execution, and provide evaluative feedback on rollouts for subsequent learning. Under this paradigm, environments function as indispensable producers of experiential data, highlighting the need to scale them toward greater complexity, realism, and interactivity. In this survey, we systematically review representative methods for environment scaling from a pioneering environment-centric perspective and organize them along the stages of the GEF loop, namely task generation, task execution, and feedback. We further analyze benchmarks, implementation strategies, and applications, consolidating fragmented advances and outlining future research directions for agent intelligence.

1 Introduction

The rapid progress of large language models (LLMs) has catalyzed a transformative shift in artificial intelligence, precipitating a surge of research on LLM-based agents (Luo et al., 2025b; Xi et al., 2025). Such agents inherit strong reasoning and task-decomposition capabilities from their base models and, when augmented with modules for tool use and memory, can execute actions, interact with real or simulated environments, accumulate experience over time, and progressively improve their own behavior. This design has achieved remarkable progress across diverse domains, including automated coding (Qwen Team, 2025; Anthropic, 2025), interactive web navigation (OpenAI, 2025a; He et al., 2025), tool use (Zhang et al., 2025d; Anthropic, 2024), and deep research (Tongyi DeepResearch Team, 2025; OpenAI, 2025b; Google DeepMind, 2024).

However, as agent capabilities continue to evolve, it is infeasible to attain intelligence beyond the human-level merely by supervised fine-tuning (SFT) pretrained models on static datasets (Huang et al., 2025a; Su et al., 2025a; Zhao et al., 2025). Such datasets are typically manually annotated or extracted under human oversight; they are intrinsically limited to human-level knowledge, costly and labor-intensive to produce at scale, and often lack realism and adaptability. By contrast, reinforcement learning provides a more aligned training paradigm (Tao et al., 2024; Zhang et al., 2025a): agents explore and interact with the environment, accumulate experience, and acquire new knowledge and skills. We formalize this interactive process as the Generation–Execution–Feedback (GEF) loop, illustrated in Figure 1 (A). In each iteration, the environment first generates diverse tasks, then the agent executes them within the environment, producing action–observation trajectories. The environment subsequently evaluates these rollouts and retains useful experience for subsequent training. Repeated iterations progressively refine the policy and expand the agent's capabilities. Notably, unlike prior work (Gao et al., 2025), we adopt a broad view of the environment: everything external to the current agent, including the state space, the executable action space, the design of feedback for interaction and evaluation, and the activities of users and other agents, is considered part of it. In this setting, the environment is no longer a mere container for agents' activities; it has become an active producer of experiential data, underscoring

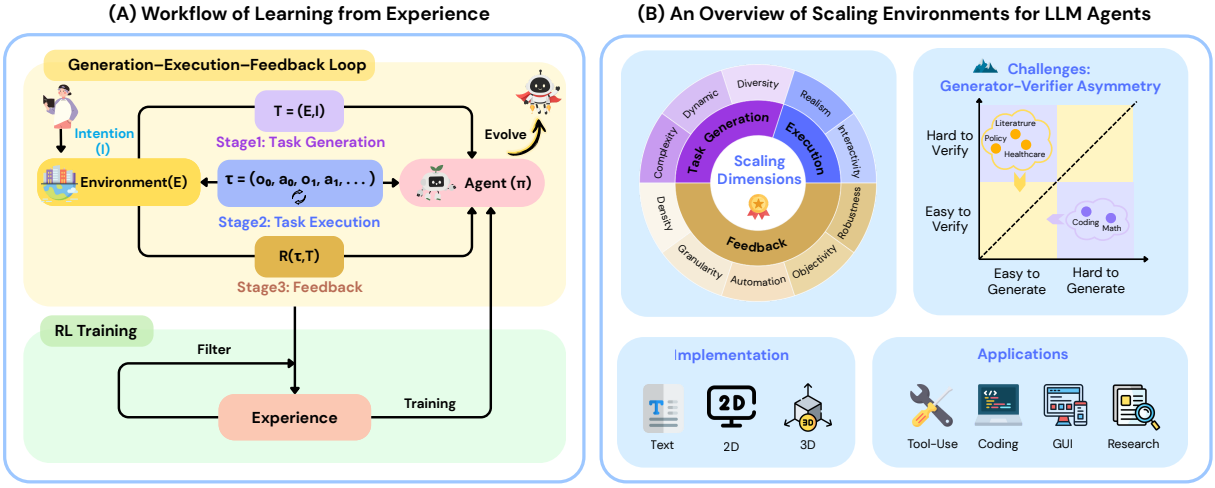

Figure 1: (A) Experience arises from the Generation–Execution–Feedback (GEF) loop, where environments generate tasks, agents execute them, and environments evaluate and filter useful experience for RL training. (B) Overview of environment scaling: a GEF-aligned taxonomy of environment-scaling methods, alongside implementation, applications, and the unique challenge of Generator–Verifier asymmetry.

the growing need for scaling environments to create a more complex, realistic, and richly interactive world (CAMEL-AI, 2025).

Recent research has embraced this trend of scaling the environment from different perspectives. For instance, systems like AgentGen (Hu et al., 2025c), AgentGym (Xi et al., 2024), and GEM (Liu et al., 2025b) devise heterogeneous environments to increase the diversity of the generated tasks. R-Zero (Huang et al., 2025a) proposes a challenger–solver framework that autonomously generates increasingly difficult tasks. Random-World (Sullivan et al., 2025) scales up the interactivity by procedural generation of diverse tools for agents to access. ARE (Andrews et al., 2025) develops an event-driven environment that supports asynchronous interactions between the environment and agents, scaling up the environmental dynamics that conform to realistic settings. However, a systematic analysis that connects these research directions remains absent.

Therefore, we first comprehensively investigate current environment scaling methods and propose a unified taxonomy aligned with the stages of the GEF loop, adopting a pioneering environment-centric perspective. In the **task generation** stage, we categorize scaling methods into *static complexity*, *dynamic complexity*, and *diversity* scaling, which together characterize an environment's ability to generate challenging, adaptive, and diverse tasks continuously. In the **task execution** stage, we highlight interactivity and realism, since these properties determine the richness and fidelity of the interaction data from which agents learn. In the **feedback** stage, we categorize the scaling of evaluative signals along five aspects: *density*, *granularity*, *automation*, *objectivity*, and *robustness*, clarifying how environments evaluate and guide agent behavior. Beyond these three stages, we also analyze the evaluation benchmarks for environments and their implementation frameworks, highlight key applications, and outline future research directions. Figure 1 (B) presents a high-level overview of the taxonomy and its supporting branches, and Figure 2 summarizes representative works across the three stages. Building on these perspectives, our survey provides an organized framework for systematically understanding environment scaling and underscores its central role in the pursuit of agent intelligence.

The survey is organized as follows. We first introduce the background and conceptual framework in §2 and §B. We then categorize representative environment-scaling methods along the three stages of the GEF loop: task generation (§3), task execution (§4), and feedback (§5). Next, we discuss evaluation benchmarks and implementation frameworks (§C and §6), as well as key applications (§D). Finally, we outline emerging research directions (§7).

2 Background

Scaling Laws for LLM Agents LLM agents exhibit scaling laws along population size, environ-

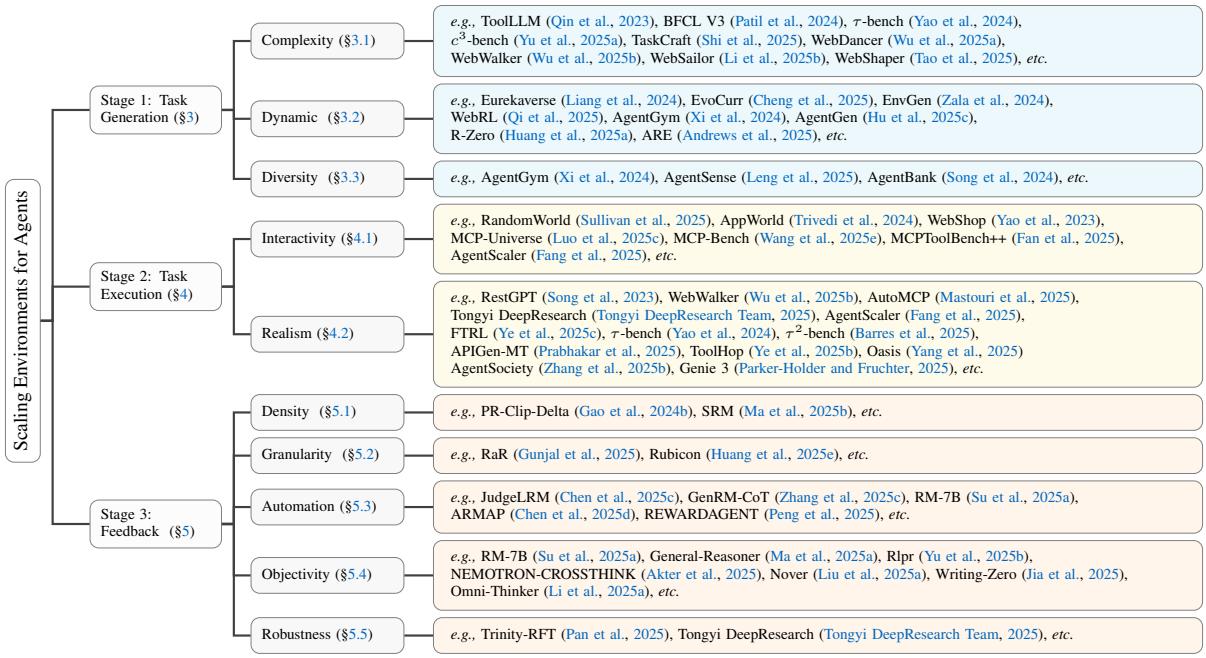

Figure 2: GEF-aligned taxonomy of environment scaling with dimensions for *Task Generation*, *Task Execution*, and *Feedback*. Representative works are illustrated as leaves on the branches.

mental complexity, and horizon length (CAMEL-AI, 2025). See § A for more background.

Generator-Verifier Asymmetry Challenge A fundamental characteristic in many real-world tasks is the inherent *Generator–Verifier Asymmetry* (Wei, 2025), namely the mismatch between the intelligence required for generator, which generates (§3) or executes (§4) tasks, and that required for verifier, which provides feedback (§5). These two kinds of intelligence naturally form two axes critical to next-generation Agentic AI, as illustrated in Figure 1(B). From this perspective, scaling up environments essentially corresponds to scaling intelligence along the x-axis and the y-axis. Current progress in RL largely exploits the regime on the easy-to-verify side of this asymmetry. These **Easy-to-Verify, Hard-to-Generate Domains** include fields such as mathematics and programming (Wei et al., 2025; Jimenez et al., 2023; Phan et al., 2025). For these domains, generating and solving a continual stream of high-quality, non-trivial tasks is challenging. In contrast, verification is objective and computationally inexpensive (e.g., via unit tests or exact match on mathematical results). This enables weak-to-strong supervision, where a simple verifier can provide accurate feedback to train a much stronger agent for solving hard tasks. On the contrary, the **Hard-to-Verify, Easy-to-Generate Domains** include areas such as creative writing, policy-making, or healthcare (Lin et al., 2024a; Arora et al., 2025). For these easy-to-propose, open-ended tasks, verification is subjective, requires substantial expert judgment, or unfolds over long horizons, making high-quality feedback scarce and expensive. This bottleneck, corresponding to the upper-left region of coordinate system, poses more difficulty in modeling the environment, and rendering environment scaling more challenging yet offering greater potential for advancing agent capabilities. Notably, the asymmetry also presents an opportunity: if the generator's stronger intelligence can be systematically leveraged to strengthen the verifier, so that it can supervise an even stronger generator, then such asymmetric property can be exploited to drive agents' self-evolution (Huang et al., 2025a; Hong et al., 2025; Lu et al., 2025c; Chen et al., 2025b; Wang et al., 2025c).

3 Stage 1: Task Generation

In the task generation stage, the environment is required to propose challenging tasks to push the agent's capabilities forward. Scaling at this stage therefore targets three facets of task supply: increasing difficulty, introducing dynamics, and expanding diversity. Accordingly, we organize representative approaches into three directions—*complexity scaling* (§ 3.1), *dynamic scaling* (§ 3.2), and *diversity scaling* (§ 3.3)—and illustrate them with a case study in Figure 3. For clarity, in *complexity scaling*,

we only consider the intrinsic difficulty of a task (i.e., static complexity). We group the temporal evolution of task difficulty (dynamic complexity) together with changes in the environment itself under the *dynamic scaling* subsection.

3.1 Complexity Scaling

Static complexity increases a task's inherent structural intricacy, moving beyond single-step commands to challenges defined by dependencies, logical flows, and hierarchical relationships. A typical example is LLM tool use, from early single-step tasks to multi-turn, multi-step scenarios, where complexity scales up as the number of turns and steps increases (Qin et al., 2023; Patil et al., 2024; Yao et al., 2024; Yu et al., 2025a). More sophisticated tasks exhibit hierarchical or compositional structure, decomposing high-level objectives into nested sub-goals and thereby testing compositional generalization, namely an agent's ability to solve novel problems by recombining known skills (Shao et al., 2023). TaskCraft (Shi et al., 2025) operationalizes this by expanding tasks both in depth (longer sequences of tool executions) and in width (multiple sub-goals per objective), enhancing hierarchical reasoning. At the highest level, conditional and graph-based tasks involve non-linear structures with branching logic, where planning must adapt dynamically to intermediate outcomes. Recent efforts in information-seeking agents extend linear sequences to complex graph-based information chains (Wu et al., 2025a; Tao et al., 2025; Li et al., 2025b; Wu et al., 2025b), and multi-agent settings further amplify this complexity, as optimal plans become contingent on the actions of other agents, producing intrinsically interdependent, graph-structured challenges.

3.2 Dynamic Scaling

Task Difficulty Dynamics Scaling dynamic complexity is critical for encouraging robust agent generalization, as it establishes non-stationary learning targets and distinct action and state spaces, which can be adjusted either according to predetermined schedules or in reaction to agent performance. A common strategy is performance-driven scheduling, where difficulty is regulated by the success rate (SR), as in Eurekaverse (Liang et al., 2024). Other approaches target newly acquired or weaker skills, as in EvoCurr (Cheng et al., 2025) and EnvGen (Zala et al., 2024). AgentGen's BI-EVAL mechanism (Hu et al., 2025c) introduces bidirectional variation, adjusting complexity upward or downward to match agent capability, in contrast to earlier methods that mostly increased difficulty (Xu et al., 2025; Luo et al., 2025a). Beyond these, WebRL (Qi et al., 2025) implements self-adjusting curricula across complex web settings, and AgentGym (Xi et al., 2024) generalizes performance-adaptive scheduling to diverse benchmarks. R-Zero (Huang et al., 2025a) formalizes a challenger–solver paradigm in which a challenger proposes near-boundary tasks based on the solver's uncertainty, and the solver improves by training on filtered task sets, yielding iterative and targeted curricula.

Environmental Dynamics The Meta Agents Research Environments (ARE) platform (Andrews et al., 2025) pushes this paradigm further by introducing a more realistic form of dynamic complexity: an asynchronous environment that evolves independently of the agent. In traditional setups, environments typically pause while the agent is reasoning; in contrast, ARE allows the world state to change continuously through random or scheduled events even while the agent is still deliberating or acting, shifting the challenge from solving discrete, static tasks to maintaining continuous, proactive engagement. As demonstrated in ARE's Mobile and Gaia2 benchmarks, agents in such dynamic environments must learn to balance cognitive depth and temporal responsiveness in order to manage latency and uncertainty as the environment continuously evolves.

3.3 Diversity Scaling

Scaling the diversity of tasks and environments is essential for developing robust and generalizable LLM agents. While regulating task-level diversity helps maintain balanced curricula and avoid redundancy (Hu et al., 2025c; Huang et al., 2025a), an equally critical dimension is at the environment level, where exposing agents to a wide spectrum of scenarios mitigates overfitting and enhances adaptability to novel situations. Frameworks such as AgentGen (Hu et al., 2025c) and AgentGym (Xi et al., 2024) synthesize a wide range of heterogeneous settings, enriching the contextual and structural variety of training signals. Beyond these, AgentSense (Leng et al., 2025) generates diverse virtual sensor data by simulating varied human personas and routines, while AgentBank (Song et al., 2024) shows that fine-tuning on tens of thousands

of heterogeneous interaction trajectories markedly improves generalization. Collectively, these approaches demonstrate that diversity across tasks and environments is fundamental to cultivating LLM agents that are adaptive and broadly capable.

4 Stage 2: Task Execution

In the task execution stage, after the agent takes an action, it receives an observation from the environment. Consequently, whether the agent can interact with the environment in real time (*interactivity* § 4.1) and whether the returned observations are consistent with real-world scenarios (*realism* § 4.2) are both critical to the quality of the resulting experience. Accordingly, we organize environment scaling in this stage into two directions: *interactivity scaling* and *realism scaling*, as shown next to the action space in Figure 3.

4.1 Interactivity Scaling

While early datasets (Liu et al., 2024b; Qin et al., 2023) primarily provided static interaction trajectories for training, recent methods (Zhang, 2025; Tang et al., 2024; Wang et al., 2024; Yao et al., 2023; Trivedi et al., 2024; Sullivan et al., 2025) allow agents to call a diverse set of real-world APIs or leverage auxiliary models as hierarchical tools via function calling or executable code generation. Standardized protocols such as the Model Context Protocol (MCP) (Anthropic, 2024; Luo et al., 2025c; Wang et al., 2025e; Fan et al., 2025) present tools, resources, and prompts through a unified interface, further improving controllability and efficiency. Complementary to MCP, another direction constructs offline environments from real-world data, sidestepping the high cost and brittleness of live API calls (Tongyi DeepResearch Team, 2025; Fang et al., 2025; Ye et al., 2025c; Yao et al., 2024; Barres et al., 2025; Prabhakar et al., 2025; Ye et al., 2025b). This approach enables more efficient interaction while preserving realism, as the database originates from the real world. Taken together, these scaling methods enable agents to interact effectively with their environments and accumulate meaningful experiences.

4.2 Realism Scaling

To ensure that large language model (LLM) agents can generalize effectively to complex, real-world scenarios, the training data collected from these environments should reflect higher levels of realism.

Tool-Use Realism In tool-use environments, earlier works have utilized fine-tuned LLMs to generate the results of tool calls (Qin et al., 2023; Lu et al., 2025a; Sun et al., 2025). However, more recent approaches have advanced by incorporating real-world APIs (Song et al., 2023; Wu et al., 2025b; Mastouri et al., 2025) or executing tasks within offline real-world databases as simulations (Tongyi DeepResearch Team, 2025; Fang et al., 2025; Ye et al., 2025c; Yao et al., 2024; Barres et al., 2025; Prabhakar et al., 2025; Ye et al., 2025b). These methods help mitigate hallucinations and enhance consistency during the collection of experiential data. Specifically, Tongyi DeepResearch Team (2025) constructs a simulated environment using an offline Wikipedia database and a custom tool suite, enabling the generation of knowledge-intensive tasks that provide rich, realistic learning signals in a cost-effective and controllable manner. Beyond traditional tool-use scenarios, which typically rely on text-based interactions, the emerging paradigm of *thinking with images* advocates for the use of visual information as a dynamic, manipulable workspace for intermediate reasoning, bridging the semantic gap between perception and symbolic thought (Su et al., 2025b). Genie 3 (Parker-Holder and Fruchter, 2025) further focuses on simulating physical environments in immersive 3D scenes. It preserves realistic physical properties, ensures long-term consistency, and supports real-time interactivity, offering a robust framework for training agents in dynamic, physically grounded settings.

Multi-Agent Realism In multi-agent settings where agents may coordinate or compete with each other (Tran et al., 2025; Qian et al., 2024; Li et al., 2024; Zhang et al., 2024; Kim et al., 2024), the scalability of realism has received growing attention. In such scenarios, each agent's behavior naturally becomes part of the environment for others, and as the number of agents increases, these interactions give rise to emergent social and economic phenomena such as information diffusion, opinion polarization, and herding effects (Yang et al., 2025; Zhang et al., 2025b). To address these challenges, Yang et al. (2025) leverage real-world social media data stored in a relational database to simulate interactive environments, enabling more realistic modeling of social dynamics. Meanwhile, Zhang et al. (2025b) move beyond earlier frameworks (Li et al., 2023; Gao et al., 2024a) that enforce execution order through message-passing SOPs by

adopting the MQTT communication protocol to enable asynchronous decision-making among autonomous agents, more closely approximating real-world processes. Similarly, ARE (Andrews et al., 2025) decouples agent and environment clocks so that the world state evolves asynchronously, and it supports simulating other agents' activities by treating them as independent events in the environment's event stream. These environments make it possible to simulate real-world societal processes and to collect more realistic experiential data.

5 Stage 3: Feedback

In the feedback stage, the environment assesses the trajectories collected during task execution and generates feedback signals for subsequent RL training. Scaling at this stage focuses on how feedback is provided, including its frequency and richness (*density* § 5.1 and *granularity* § 5.2), its level of automation (*automation* § 5.3), as well as how objectively and reliably it is delivered (*objectivity* § 5.4 and *robustness* § 5.5). Accordingly, we categorize representative approaches along five dimensions: *density scaling*, *granularity scaling*, *automation scaling*, *objectivity scaling*, and *robustness scaling*, which are illustrated in Figure 4.

5.1 Density Scaling

The first aspect is the density of feedback, which refers to how frequently the environment provides evaluative signals. Early environments offered sparse, outcome-based rewards that only indicated the final success or failure of a trajectory. While such trajectory-level signals can yield stable training dynamics in reasoning-intensive tasks such as mathematics or code generation (Jiang et al., 2025; Da et al., 2025; Qian et al., 2025), they suffer from credit assignment challenges. The overall success or failure cannot be easily attributed to intermediate actions. To mitigate this issue, recent studies have investigated step-level rewards derived from intermediate actions, providing denser supervision and detailed guidance for agents to learn. However, the effectiveness and underlying mechanisms of such step-level feedback remain insufficiently understood. Gao et al. (2024b) shows that naively combining step-level reward and outcome reward may degrade the agent's performance due to reward hacking. In this situation, trivial actions accumulate high rewards and lead to repetitive reasoning behavior. They mitigate this problem through threshold clipping and training with reward differences between adjacent steps. Similarly, Ma et al. (2025b) demonstrates that step-level reward models effectively capture the logical coherence of mathematical reasoning yet are unreliable for natural language tasks, highlighting both the promise and the limitations of denser feedback signals.

5.2 Granularity Scaling

Granularity scaling captures how detailed the evaluative feedback of task performance is. Coarse-grained evaluation only focuses on end-to-end success and emphasizes the final outcome (Jin et al., 2025; Li et al., 2025c), whereas fine-grained evaluation decomposes agents' performance into structured components and provides richer feedback for error attribution and more targeted improvement. Moving beyond a single scalar reward, fine-grained schemes are often operationalized as multi-dimensional feedback, such as a set of scores across criteria that offers more informative guidance (Huang et al., 2025c). Some work employs manually designed, rule-based rewards to improve different aspects such as correctness, format and efficiency (Qian et al., 2025; Song et al., 2025; Wang et al., 2025b; Zhang et al., 2025e). Based on this idea, the Rubrics as Rewards (RaR) framework (Gunjal et al., 2025; Huang et al., 2025e) transforms checklist-based rubrics into structured, multi-criteria reward signals, thereby improving the alignment with human preferences and ensuring more consistent performance across diverse evaluation standards.

5.3 Automation Scaling

Automation scaling refers to the transition in feedback from human-provided supervision to AI-generated evaluation. Reinforcement Learning from Human Feedback (RLHF) encompasses a broad class of frameworks aimed at aligning LLMs with human preferences (Sheng et al., 2025; Hu et al., 2025a; Christiano et al., 2017). While effective, these methods depend heavily on human annotators, making feedback slow, costly, and difficult to scale. To automate feedback and reduce reliance on labor-intensive preference labeling, a growing body of work replaces human evaluators with capable LLMs acting as automated judges, forming the paradigm of Reinforcement Learning from AI Feedback (RLAIF) (Zhang et al., 2025c; Su et al., 2025a; Chen et al., 2025c; Lee et al., 2024; Bai et al., 2022). Pioneer works such as

REWARDAGENT (Peng et al., 2025) employ a Verification Agent to assess both the factual correctness of responses and their adherence to instructions, integrating these evaluations with base human feedback to effectively guide training. In contrast, ARMAP (Chen et al., 2025d) innovatively bypasses the need for a stronger LLM-as-judge by using LLM-generated positive and negative trajectories to train a classification model, thereby creating an automated and effective reward mechanism. By embedding such automated evaluation within the environment, agents are able to acquire higher-order capabilities progressively without continuous human supervision. However, this paradigm also amplifies risks such as propagating biases from the LLM judge and reward hacking, where agents exploit the reward model in ways that diverge from intended goals, potentially leading to model collapse. This presents a critical trade-off between scalability and safety, underscoring the need for research into robust mitigation strategies.

5.4 Objectivity Scaling

Scaling objectivity is about increasing the extent to which environmental feedback is grounded in verifiable truth, moving from subjective, preference-based, or generative evaluations to strictly objective signals. Objective feedback exemplified by Reinforcement Learning with Verifiable Rewards (RLVR) has demonstrated efficacy and is extensively utilized in structured domains, such as mathematics and programming, where solutions can be objectively validated (Shao et al., 2024; DeepSeek-AI et al., 2025). Nonetheless, numerous real-world tasks, like writing and medical consultation, are open-ended and often lack objectively verifiable ground truth, making verification much harder. To mitigate this deficiency, recent studies have expanded RLVR to non-verifiable domains (Andrews et al., 2025; Ma et al., 2025a; Akter et al., 2025; Su et al., 2025a; Yu et al., 2025b; Liu et al., 2025a). For instance, Writing-Zero (Jia et al., 2025) tackles the creative writing tasks by training a pairwise generative reward model, therefore transforming subjective assessments into reliable, verifiable reward signals. Li et al. (2025a) proposes Omni-Thinker, which combines rule-based verifiable rewards with generative preference signals from an LLM-as-a-judge, forming a unified multi-task RL training loop. ARE's verifier (Andrews et al., 2025) instantiates RLVR by taking the minimal correct sequence of write operations as ground truth, combining hard checks (exact parameter matches) and soft checks (LLM-based semantic judgments) under strict causal ordering. This work demonstrates the feasibility of structured verifiable rewards in open-ended environments.

5.5 Robustness Scaling

Feedback robustness requires the environment to provide stable and reliable reward signals. However, it is common for environmental instabilities (e.g., delays, crashes, corrupted tool outputs) to undermine feedback reliability and degrade the training process. To scale robustness **at the reward level**, some approaches mitigate the effects of noise by generating soft, probabilistic rewards with generative verifiers (Lin et al., 2024b; Su et al., 2025a). For the reward-hacking problem, where agents tend to develop undesirable behaviors to maximize reward without achieving the intended goals, some studies propose frameworks such as MONA (Farquhar et al., 2025) to evaluate the future utility of actions via an overseer, constraining unstable behaviors while preserving explainability. **At the interaction level**, in order to solve challenges related to environmental failures, Trinity-RFT (Pan et al., 2025) proposes asynchronous inference and retry mechanisms, while Tongyi DeepResearch Team (2025) improves robustness by constructing a fast and deterministic interaction environment, employing caching, retrying failed calls, and switching to similar providers to prevent corrupted trajectories. Looking ahead, future reward design should prioritize both efficacy and the prevention of hacking via tricky patterns. Also, attention should also be paid to system-level robustness maintenance.

6 Implementation Frameworks

There exist diverse implementations of simulated environments that vary in modality and complexity.

Visual Environments Visual environments include both 2D and 3D settings. In 2D, grid-based designs offer symbolic or pixel-level tasks, exemplified by MiniGrid (Chevalier-Boisvert et al., 2023) and the Atari-focused ALE (Bellemare et al., 2013). Extensions such as Kinetix (Matthews et al., 2024) add physics-based interactions, while WebArena (Zhou et al., 2023), Mind2Web (Deng et al., 2023), and VisualWebArena (Koh et al., 2024b) bring agents into realistic web domains. Building on this line, ARE (Andrews et al., 2025) provides a mobile-style web environment with interconnected

apps and tools, combining time control and a structured verifier for scalable evaluation. In 3D, environments emphasize embodied interaction with realistic physics and visual perception. General-purpose simulators such as Habitat (Savva et al., 2019) and ThreeDWorld (Gan et al., 2020) support navigation and manipulation, while domain-specific worlds like MineRL (Guss et al., 2019) and MineDojo (Fan et al., 2022) exploit the flexibility of Minecraft. Household benchmarks such as ALFRED (Shridhar et al., 2020a) and EmbodiedQA (Das et al., 2018) further evaluate grounding of language in photo-realistic 3D spaces, requiring agents to follow instructions, answer questions, or perform multi-step tasks.

Text-based Environments Text-based environments instead emphasize reasoning and decision-making from natural language descriptions. TextWorld (Côté et al., 2018) provides a framework for procedurally generating text games with goal-driven challenges. Reddit-RL-simulator (He et al., 2016; Chan and King, 2018) implements an RL environment for iteratively tracking and recommending popular discussion threads on Reddit. ALFWorld (Shridhar et al., 2020b) adapts ALFRED tasks into textual form, and JerichoAgentBench (Liu et al., 2023) extends the Jericho suite (Hausknecht et al., 2020) with annotated benchmarks and concrete objectives. Such text environments stress the agent's ability to infer latent dynamics under partial observability, making them a natural testbed for language-based decision-making.

7 Future Directions

Co-Evolution via Embedded External Tools As task complexity increases, feedback mechanisms must scale in tandem to avoid sparse, noisy, or misaligned learning signals. Future environments can achieve this by embedding external tools such as verifiers, simulators, compilers, or executable systems directly into the learning loop. These tools co-evolve with task generation, providing structured, verifiable feedback and enabling agents to interact with increasingly sophisticated challenges. By integrating embedded tools with automated evaluators such as LLM-based formative signals, environments can better support multi-step, open-ended, and creative problem solving, while also helping to mitigate the Generator–Verifier Asymmetry.

Scaling Through Generator–Verifier Synergy Future environments can encourage strong generators to decompose solutions from complex problems into intermediate solutions from smaller subproblems that are tractable for weak verifiers or environments with scarce feedback. This enables scalable supervision even in domains such as creative cultural production and policy-making where holistic verification is difficult. Conversely, weak generators can provide diverse candidate solutions that are filtered, ranked, or refined by stronger Verifiers equipped with richer evaluative capabilities. By embedding such co-evolutionary dynamics into the environment design, Generator–Verifier asymmetry can shift into a catalyst for continuous self-improvement.

Open-Ended, Multi-Agent Environments Future environments can scale to support large-scale multi-agent interactions, emergent social dynamics, and economic- or organizational-level simulations, providing rich contexts for studying both collaborative and non-cooperative behaviors. In particular, scaling to massively multi-cultural multi-lingual settings requires scaling environment construction for agents to navigate the subtle semantics of concepts and values that vary across different societies, not only in the textual but also in the multi-modal domain (Koh et al., 2024a; Huang et al., 2025d). By exposing agents to these open-ended, interactive scenarios, such environments foster generalization and strategic planning, equipping LLM agents to better handle complex, real-world challenges beyond isolated task execution.

8 Conclusion

The era of experience places environments at the core of LLM agent development, positioning them as active producers of experiential data and underscoring the growing need for scaling environments to create a more complex, realistic, and richly interactive world. From a pioneering environment-centric perspective, our survey proposes a unified taxonomy that organizes representative work across the GEF loop (task generation, task execution, and feedback), together with evaluation, implementation, and applications. We also surface key challenges for advancing agent intelligence, including the asymmetry between generators and verifiers and the construction of more open, large-scale multi-agent environments, thereby shedding light on future agentic work.

Limitations

As *scaling environments* is still emerging, relatively few works explicitly adopt this framing. Thus, we take a broad view of "environment" and organize representative studies along the GEF loop (task generation–task execution–feedback). Although this wider lens brings in some adjacent lines of work, our taxonomy is comprehensive from a pioneering environment-centric perspective. Besides, given the rapid pace of agentic research, some of the latest papers may fall outside this snapshot. We will keep the survey updated as the literature evolves.

References

Syeda Nahida Akter, Shrimai Prabhumoye, Matvei Novikov, Seungju Han, Ying Lin, Evelina Bakhturina, Eric Nyberg, Yejin Choi, Mostofa Patwary, Mohammad Shoeybi, et al. 2025. Nemotron-crossthink: Scaling self-learning beyond math reasoning. *arXiv preprint arXiv:2504.13941*.

Pierre Andrews, Amine Benhalloum, Gerard Moreno-Torres Bertran, Matteo Bettini, Amar Budhiraja, Ricardo Silveira Cabral, Virginie Do, Romain Froger, Emilien Garreau, Jean-Baptiste Gaya, Hugo Laurençon, Maxime Lecanu, Kunal Malkan, Dheeraj Mekala, Pierre Ménard, Grégoire Mialon, Ulyana Piterbarg, Mikhail Plekhanov, Mathieu Rita, Andrey Rusakov, Thomas Scialom, Vladislav Vorotilov, Mengjue Wang, and Ian Yu. 2025. Are: Scaling up agent environments and evaluations. *Preprint*, arXiv:2509.17158.

Anthropic. 2024. Introducing the model context protocol. https://www.anthropic.com/news/model-context-protocol.

Anthropic. 2025. Introducing claude 4. https://www.anthropic.com/news/claude-4.

Rahul K Arora, Jason Wei, Rebecca Soskin Hicks, Preston Bowman, Joaquin Quiñonero-Candela, Foivos Tsimpourlas, Michael Sharman, Meghan Shah, Andrea Vallone, Alex Beutel, et al. 2025. Healthbench: Evaluating large language models towards improved human health. *arXiv preprint arXiv:2505.08775*.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. Constitutional ai: Harmlessness from ai feedback. *Preprint*, arXiv:2212.08073.

Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. 2025. τ^2-bench: Evaluating conversational agents in a dual-control environment. *Preprint*, arXiv:2506.07982.

Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of artificial intelligence research*, 47:253–279.

CAMEL-AI. 2025. Scaling environments for agents. https://www.camel-ai.org/blogs/scaling-environments-for-agents.

Hou Pong Chan and Irwin King. 2018. Thread popularity prediction and tracking with a permutation-invariant model. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3392–3401, Brussels, Belgium. Association for Computational Linguistics.

Delong Chen, Willy Chung, Yejin Bang, Ziwei Ji, and Pascale Fung. 2025a. Worldprediction: A benchmark for high-level world modeling and long-horizon procedural planning. *arXiv preprint arXiv:2506.04363*.

Lili Chen, Mihir Prabhudesai, Katerina Fragkiadaki, Hao Liu, and Deepak Pathak. 2025b. Self-questioning language models. *arXiv preprint arXiv:2508.03682*.

Nuo Chen, Zhiyuan Hu, Qingyun Zou, Jiaying Wu, Qian Wang, Bryan Hooi, and Bingsheng He. 2025c. Judgelrm: Large reasoning models as a judge. *Preprint*, arXiv:2504.00050.

Zhenfang Chen, Delin Chen, Rui Sun, Wenjun Liu, and Chuang Gan. 2025d. Scaling autonomous agents via automatic reward modeling and planning. *Preprint*, arXiv:2502.12130.

Yang Cheng, Zilai Wang, Weiyu Ma, Wenhui Zhu, Yue Deng, and Jian Zhao. 2025. Evocurr: Self-evolving curriculum with behavior code generation for complex decision-making. *arXiv preprint arXiv:2508.09586*.

Maxime Chevalier-Boisvert, Bolun Dai, Mark Towers, Rodrigo Perez-Vicente, Lucas Willems, Salem Lahlou, Suman Pal, Pablo Samuel Castro, and Jordan Terry. 2023. Minigrid & miniworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. *Advances in Neural Information Processing Systems*, 36:73383–73394.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.

Marc-Alexandre Côté, Akos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew Hausknecht, Layla El Asri, Mahmoud Adada, et al. 2018. Textworld: A learning environment for text-based games. In *Workshop on Computer Games*, pages 41–75. Springer.

Jeff Da, Clinton Wang, Xiang Deng, Yuntao Ma, Nikhil Barhate, and Sean Hendryx. 2025. Agent-rlvr: Training software engineering agents via guidance and environment rewards. *arXiv preprint arXiv:2506.11425*.

Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. 2018. Embodied question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–10.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *Preprint*, arXiv:2501.12948.

Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. 2023. Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, 36:28091–28114.

Haoyi Duan, Hong-Xing Yu, Sirui Chen, Li Fei-Fei, and Jiajun Wu. 2025. Worldscore: A unified evaluation benchmark for world generation. *arXiv preprint arXiv:2504.00983*.

Linxi Fan, Guanzhi Wang, Yunfan Jiang, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, and Anima Anandkumar. 2022. Minedojo: Building open-ended embodied agents with internet-scale knowledge. *Advances in Neural Information Processing Systems*, 35:18343–18362.

Shiqing Fan, Xichen Ding, Liang Zhang, and Linjian Mo. 2025. Mcptoolbench++: A large scale ai agent model context protocol mcp tool use benchmark. *arXiv preprint arXiv:2508.07575*.

Runnan Fang, Shihao Cai, Baixuan Li, Jialong Wu, Guangyu Li, Wenbiao Yin, Xinyu Wang, Xiaobin Wang, Liangcai Su, Zhen Zhang, Shibin Wu, Zhengwei Tao, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025. Towards general agentic intelligence via environment scaling. *Preprint*, arXiv:2509.13311.

Sebastian Farquhar, Vikrant Varma, David Lindner, David Elson, Caleb Biddulph, Ian Goodfellow, and Rohin Shah. 2025. Mona: Myopic optimization with non-myopic approval can mitigate multi-step reward hacking. *Preprint*, arXiv:2501.13011.

Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. 2025. Retool: Reinforcement learning for strategic tool use in llms. *Preprint*, arXiv:2504.11536.

Fosowl. 2024. AgenticSeek: A fully local and open-source AI agent for autonomous task execution. GitHub repository. Accessed: 2024-09-02.

Chuang Gan, Jeremy Schwartz, Seth Alter, Damian Mrowca, Martin Schrimpf, James Traer, Julian De Freitas, Jonas Kubilius, Abhishek Bhandwaldar, Nick Haber, et al. 2020. Threedworld: A platform for

interactive multi-modal physical simulation. *arXiv preprint arXiv:2007.04954*.

Dawei Gao, Zitao Li, Xuchen Pan, Weirui Kuang, Zhijian Ma, Bingchen Qian, Fei Wei, Wenhao Zhang, Yuexiang Xie, Daoyuan Chen, Liuyi Yao, Hongyi Peng, Zeyu Zhang, Lin Zhu, Chen Cheng, Hongzhu Shi, Yaliang Li, Bolin Ding, and Jingren Zhou. 2024a. Agentscope: A flexible yet robust multi-agent platform. *Preprint*, arXiv:2402.14034.

Huan-ang Gao, Jiayi Geng, Wenyue Hua, Mengkang Hu, Xinzhe Juan, Hongzhang Liu, Shilong Liu, Jiahao Qiu, Xuan Qi, Yiran Wu, et al. 2025. A survey of self-evolving agents: On path to artificial super intelligence. *arXiv preprint arXiv:2507.21046*.

Jiaxuan Gao, Shusheng Xu, Wenjie Ye, Weilin Liu, Chuyi He, Wei Fu, Zhiyu Mei, Guangju Wang, and Yi Wu. 2024b. On designing effective rl reward at training time for llm reasoning. *arXiv preprint arXiv:2410.15115*.

Google DeepMind. 2024. Gemini Deep Research. https://blog.google/products/gemini/google-gemini-deep-research/.

Anisha Gunjal, Anthony Wang, Elaine Lau, Vaskar Nath, Bing Liu, and Sean Hendryx. 2025. Rubrics as rewards: Reinforcement learning beyond verifiable domains. *Preprint*, arXiv:2507.17746.

William H Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. 2019. Minerl: A large-scale dataset of minecraft demonstrations. *arXiv preprint arXiv:1907.13440*.

Matthew Hausknecht, Prithviraj Ammanabrolu, Marc-Alexandre Côté, and Xingdi Yuan. 2020. Interactive fiction games: A colossal adventure. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7903–7910.

Ji He, Mari Ostendorf, Xiaodong He, Jianshu Chen, Jianfeng Gao, Lihong Li, and Li Deng. 2016. Deep reinforcement learning with a combinatorial action space for predicting popular Reddit threads. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1838–1848, Austin, Texas. Association for Computational Linguistics.

Zhitao He, Zijun Liu, Peng Li, Yi R. Fung, Ming Yan, Ji Zhang, Fei Huang, and Yang Liu. 2025. Advancing language multi-agent learning with credit re-assignment for interactive environment generalization. In *Second Conference on Language Modeling*.

Haitao Hong, Yuchen Yan, Xingyu Wu, Guiyang Hou, Wenqi Zhang, Weiming Lu, Yongliang Shen, and Jun Xiao. 2025. Cooper: Co-optimizing policy and reward models in reinforcement learning for large language models. *arXiv preprint arXiv:2508.05613*.

Jian Hu, Xibin Wu, Wei Shen, Jason Klein Liu, Zilin Zhu, Weixun Wang, Songlin Jiang, Haoran Wang, Hao Chen, Bin Chen, Weikai Fang, Xianyu, Yu Cao, Haotian Xu, and Yiming Liu. 2025a. Openrlhf: An easy-to-use, scalable and high-performance rlhf framework. *Preprint*, arXiv:2405.11143.

Mengkang Hu, Tianxing Chen, Yude Zou, Yuheng Lei, Qiguang Chen, Ming Li, Yao Mu, Hongyuan Zhang, Wenqi Shao, and Ping Luo. 2025b. Text2world: Benchmarking large language models for symbolic world model generation. *arXiv preprint arXiv:2502.13092*.

Mengkang Hu, Pu Zhao, Can Xu, Qingfeng Sun, Jian-Guang Lou, Qingwei Lin, Ping Luo, and Saravan Rajmohan. 2025c. Agentgen: Enhancing planning abilities for large language model based agent via environment and task generation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pages 496–507.

Chengsong Huang, Wenhao Yu, Xiaoyang Wang, Hongming Zhang, Zongxia Li, Ruosen Li, Jiaxin Huang, Haitao Mi, and Dong Yu. 2025a. R-zero: Self-evolving reasoning llm from zero data. *arXiv preprint arXiv:2508.05004*.

Kung-Hsiang Huang, Haoyi Qiu, Yutong Dai, Caiming Xiong, and Chien-Sheng Wu. 2025b. Gui-kv: Efficient gui agents via kv cache with spatio-temporal awareness. *Preprint*, arXiv:2510.00536.

Shijue Huang, Hongru Wang, Wanjun Zhong, Zhaochen Su, Jiazhan Feng, Bowen Cao, and Yi R Fung. 2025c. Adactrl: Towards adaptive and controllable reasoning via difficulty-aware budgeting. *arXiv preprint arXiv:2505.18822*.

Yuchen Huang, Zhiyuan Fan, Zhitao He, Sandeep Polisetty, Wenyan Li, and Yi R Fung. 2025d. Cultureclip: Empowering clip with cultural awareness through synthetic images and contextualized captions. In *Second Conference on Language Modeling*.

Zenan Huang, Yihong Zhuang, Guoshan Lu, Zeyu Qin, Haokai Xu, Tianyu Zhao, Ru Peng, Jiaqi Hu, Zhanming Shen, Xiaomeng Hu, et al. 2025e. Reinforcement learning with rubric anchors. *arXiv preprint arXiv:2508.12790*.

Ruipeng Jia, Yunyi Yang, Yongbo Gai, Kai Luo, Shihao Huang, Jianhe Lin, Xiaoxi Jiang, and Guanjun Jiang. 2025. Writing-zero: Bridge the gap between non-verifiable tasks and verifiable rewards. *Preprint*, arXiv:2506.00103.

Dongfu Jiang, Yi Lu, Zhuofeng Li, Zhiheng Lyu, Ping Nie, Haozhe Wang, Alex Su, Hui Chen, Kai Zou, Chao Du, et al. 2025. Verltool: Towards holistic agentic reinforcement learning with tool use. *arXiv preprint arXiv:2509.01055*.

Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik

Narasimhan. 2023. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*.

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*.

Kyungha Kim, Sangyun Lee, Kung-Hsiang Huang, Hou Pong Chan, Manling Li, and Heng Ji. 2024. Can llms produce faithful explanations for fact-checking? towards faithful explainable fact-checking via multi-agent debate. *Preprint*, arXiv:2402.07401.

Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Russ Salakhutdinov, and Daniel Fried. 2024a. VisualWebArena: Evaluating multimodal agents on realistic visual web tasks. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 881–905, Bangkok, Thailand. Association for Computational Linguistics.

Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Ruslan Salakhutdinov, and Daniel Fried. 2024b. Visualwebarena: Evaluating multimodal agents on realistic visual web tasks. *arXiv preprint arXiv:2401.13649*.

Hanyu Lai, Xiao Liu, Iat Long Iong, Shuntian Yao, Yuxuan Chen, Pengbo Shen, Hao Yu, Hanchen Zhang, Xiaohan Zhang, Yuxiao Dong, et al. 2024. Autowebglm: A large language model-based web navigating agent. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5295–5306.

Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, and Sushant Prakash. 2024. Rlaif vs. rlhf: Scaling reinforcement learning from human feedback with ai feedback. *Preprint*, arXiv:2309.00267.

Zikang Leng, Megha Thukral, Yaqi Liu, Hrudhai Rajasekhar, Shruthi K Hiremath, and Thomas Plötz. 2025. Agentsense: Virtual sensor data generation using llm agent in simulated home environments. *arXiv preprint arXiv:2506.11773*.

Dawei Li, Zhen Tan, Peijia Qian, Yifan Li, Kumar Satvik Chaudhary, Lijie Hu, and Jiayi Shen. 2024. Smoa: Improving multi-agent large language models with sparse mixture-of-agents. *Preprint*, arXiv:2411.03284.

Derek Li, Jiaming Zhou, Amirreza Kazemi, Qianyi Sun, Abbas Ghaddar, Mohammad Ali Alomrani, Liheng Ma, Yu Luo, Dong Li, Feng Wen, et al. 2025a. Omnithink: Scaling cross-domain generalization in llms via multi-task rl with hybrid rewards. *arXiv preprint arXiv:2507.14783*.

Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for "mind" exploration of large language model society. *Preprint*, arXiv:2303.17760.

Kuan Li, Zhongwang Zhang, Huifeng Yin, Liwen Zhang, Litu Ou, Jialong Wu, Wenbiao Yin, Baixuan Li, Zhengwei Tao, Xinyu Wang, Weizhou Shen, Junkai Zhang, Dingchu Zhang, Xixi Wu, Yong Jiang, Ming Yan, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025b. Websailor: Navigating super-human reasoning for web agent. *Preprint*, arXiv:2507.02592.

Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025c. Torl: Scaling tool-integrated rl. *arXiv preprint arXiv:2503.23383*.

William Liang, Sam Wang, Hung-Ju Wang, Osbert Bastani, Dinesh Jayaraman, and Yecheng Jason Ma. 2024. Environment curriculum generation via large language models. In *8th Annual Conference on Robot Learning*.

Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2024a. Wildbench: Benchmarking llms with challenging tasks from real users in the wild. *arXiv preprint arXiv:2406.04770*.

Muhan Lin, Shuyang Shi, Yue Guo, Behdad Chalaki, Vaishnav Tadiparthi, Ehsan Moradi Pari, Simon Stepputtis, Joseph Campbell, and Katia Sycara. 2024b. Navigating noisy feedback: Enhancing reinforcement learning with error-prone language models. *arXiv preprint arXiv:2410.17389*.

Wei Liu, Siya Qi, Xinyu Wang, Chen Qian, Yali Du, and Yulan He. 2025a. Nover: Incentive training for language models via verifier-free reinforcement learning. *arXiv preprint arXiv:2505.16022*.

Xiao Liu, Bo Qin, Dongzhu Liang, Guang Dong, Hanyu Lai, Hanchen Zhang, Hanlin Zhao, Iat Long Iong, Jiadai Sun, Jiaqi Wang, Junjie Gao, Junjun Shan, Kangning Liu, Shudan Zhang, Shuntian Yao, Siyi Cheng, Wentao Yao, Wenyi Zhao, Xinghan Liu, Xinyi Liu, Xinying Chen, Xinyue Yang, Yang Yang, Yifan Xu, Yu Yang, Yujia Wang, Yulin Xu, Zehan Qi, Yuxiao Dong, and Jie Tang. 2024a. Autoglm: Autonomous foundation agents for guis. *Preprint*, arXiv:2411.00820.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. 2023. Agentbench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*.

Zichen Liu, Anya Sims, Keyu Duan, Changyu Chen, Simon Yu, Xiangxin Zhou, Haotian Xu, Shaopan Xiong, Bo Liu, Chenmien Tan, Chuen Yang Beh, Weixun Wang, Hao Zhu, Weiyan Shi, Diyi Yang, Michael Shieh, Yee Whye Teh, Wee Sun Lee, and Min Lin. 2025b. Gem: A gym for agentic llms. *Preprint*, arXiv:2510.01051.

Zuxin Liu, Thai Hoang, Jianguo Zhang, Ming Zhu, Tian Lan, Shirley Kokane, Juntao Tan, Weiran Yao, Zhiwei Liu, Yihao Feng, Rithesh Murthy, Liangwei Yang, Silvio Savarese, Juan Carlos Niebles, Huan Wang, Shelby Heinecke, and Caiming Xiong. 2024b. Apigen: Automated pipeline for generating verifiable and diverse function-calling datasets. *Preprint*, arXiv:2406.18518.

Jiarui Lu, Thomas Holleis, Yizhe Zhang, Bernhard Aumayer, Feng Nan, Felix Bai, Shuang Ma, Shen Ma, Mengyu Li, Guoli Yin, Zirui Wang, and Ruoming Pang. 2025a. Toolsandbox: A stateful, conversational, interactive evaluation benchmark for llm tool use capabilities. *Preprint*, arXiv:2408.04682.

Quanfeng Lu, Wenqi Shao, Zitao Liu, Lingxiao Du, Fanqing Meng, Boxuan Li, Botong Chen, Siyuan Huang, Kaipeng Zhang, and Ping Luo. 2025b. Guiodyssey: A comprehensive dataset for cross-app gui navigation on mobile devices. *Preprint*, arXiv:2406.08451.

Songshuo Lu, Hua Wang, Zhi Chen, and Yaohua Tang. 2025c. Urpo: A unified reward & policy optimization framework for large language models. *arXiv preprint arXiv:2507.17515*.

Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, Yansong Tang, and Dongmei Zhang. 2025a. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *Preprint*, arXiv:2308.09583.

Junyu Luo, Weizhi Zhang, Ye Yuan, Yusheng Zhao, Junwei Yang, Yiyang Gu, Bohan Wu, Binqi Chen, Ziyue Qiao, Qingqing Long, Rongcheng Tu, Xiao Luo, Wei Ju, Zhiping Xiao, Yifan Wang, Meng Xiao, Chenwu Liu, Jingyang Yuan, Shichang Zhang, Yiqiao Jin, Fan Zhang, Xian Wu, Hanqing Zhao, Dacheng Tao, Philip S. Yu, and Ming Zhang. 2025b. Large language model agent: A survey on methodology, applications and challenges. *Preprint*, arXiv:2503.21460.

Ziyang Luo, Zhiqi Shen, Wenzhuo Yang, Zirui Zhao, Prathyusha Jwalapuram, Amrita Saha, Doyen Sahoo, Silvio Savarese, Caiming Xiong, and Junnan Li. 2025c. Mcp-universe: Benchmarking large language models with real-world model context protocol servers. *arXiv preprint arXiv:2508.14704*.

Xueguang Ma, Qian Liu, Dongfu Jiang, Ge Zhang, Zejun Ma, and Wenhu Chen. 2025a. General-reasoner: Advancing llm reasoning across all domains. *arXiv preprint arXiv:2505.14652*.

Yiran Ma, Zui Chen, Tianqiao Liu, Mi Tian, Zhuo Liu, Zitao Liu, and Weiqi Luo. 2025b. What are step-level reward models rewarding? counterintuitive findings from mcts-boosted mathematical reasoning. *Preprint*, arXiv:2412.15904.

Meriem Mastouri, Emna Ksontini, and Wael Kessentini. 2025. Making rest apis agent-ready: From openapi to mcp servers for tool-augmented llms. *Preprint*, arXiv:2507.16044.

Michael Matthews, Michael Beukman, Chris Lu, and Jakob Foerster. 2024. Kinetix: Investigating the training of general agents through open-ended physics-based control tasks. *arXiv preprint arXiv:2410.23208*.

OpenAI. 2023. Function calling. https://platform.openai.com/docs/guides/function-calling.

OpenAI. 2025. Computer-using agent: Introducing a universal interface for ai to interact with the digital world. https://openai.com/index/computer-using-agent/. Accessed: 2025-09-16.

OpenAI. 2025a. Introducing ChatGPT agent: bridging research and action. https://openai.com/index/introducing-chatgpt-agent/.

OpenAI. 2025b. Introducing Deep Research. https://openai.com/index/introducing-deep-research/.

Xuchen Pan, Yanxi Chen, Yushuo Chen, Yuchang Sun, Daoyuan Chen, Wenhao Zhang, Yuexiang Xie, Yilun Huang, Yilei Zhang, Dawei Gao, Weijie Shi, Yaliang Li, Bolin Ding, and Jingren Zhou. 2025. Trinity-rft: A general-purpose and unified framework for reinforcement fine-tuning of large language models. *Preprint*, arXiv:2505.17826.

Jack Parker-Holder and Shlomi Fruchter. 2025. Genie 3: A new frontier for world models.

Shishir G. Patil, Huanzhi Mao, Charlie Cheng-Jie Ji, Fanjia Yan, Vishnu Suresh, Ion Stoica, and Joseph E. Gonzalez. 2024. The berkeley function calling leaderboard (bfcl): From tool use to agentic evaluation of large language models. In *Advances in Neural Information Processing Systems*.

Hao Peng, Yunjia Qi, Xiaozhi Wang, Zijun Yao, Bin Xu, Lei Hou, and Juanzi Li. 2025. Agentic reward modeling: Integrating human preferences with verifiable correctness signals for reliable reward systems. *Preprint*, arXiv:2502.19328.

Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, et al. 2025. Humanity's last exam. *arXiv preprint arXiv:2501.14249*.

Akshara Prabhakar, Zuxin Liu, Ming Zhu, Jianguo Zhang, Tulika Awalgaonkar, Shiyu Wang, Zhiwei Liu, Haolin Chen, Thai Hoang, Juan Carlos Niebles, Shelby Heinecke, Weiran Yao, Huan Wang, Silvio Savarese, and Caiming Xiong. 2025. Apigen-mt: Agentic pipeline for multi-turn data generation via simulated agent-human interplay. *Preprint*, arXiv:2504.03601.

Zehan Qi, Xiao Liu, Iat Long Iong, Hanyu Lai, Xueqiao Sun, Jiadai Sun, Xinyue Yang, Yu Yang, Shuntian Yao, Wei Xu, et al. 2025. Webrl: Training llm web agents via self-evolving online curriculum reinforcement learning. In *The Thirteenth International Conference on Learning Representations*.

Chen Qian, Zihao Xie, Yifei Wang, Wei Liu, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2024. Scaling large-language-model-based multi-agent collaboration. *CoRR*.

Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. Toolrl: Reward is all tool learning needs. *arXiv preprint arXiv:2504.13958*.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. 2023. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*.

Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, Wanjun Zhong, Kuanye Li, Jiale Yang, Yu Miao, Woyu Lin, Longxiang Liu, Xu Jiang, Qianli Ma, Jingyu Li, Xiaojun Xiao, Kai Cai, Chuang Li, Yaowei Zheng, Chaolin Jin, Chen Li, Xiao Zhou, Minchao Wang, Haoli Chen, Zhaojian Li, Haihua Yang, Haifeng Liu, Feng Lin, Tao Peng, Xin Liu, and Guang Shi. 2025. Ui-tars: Pioneering automated gui interaction with native agents. *Preprint*, arXiv:2501.12326.

Qwen Team. 2024. Qwen3-coder technical blog. `https://qwenlm.github.io/blog/qwen3-coder/`. Accessed: 2024-07-23.

Qwen Team. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. 2019. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9339–9347.

Nan Shao, Zefan Cai, Hanwei Xu, Chonghua Liao, Yanan Zheng, and Zhilin Yang. 2023. Compositional task representations for large language models. In *ICLR*.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *Preprint*, arXiv:2402.03300.

Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, page 1279–1297. ACM.

Dingfeng Shi, Jingyi Cao, Qianben Chen, Weichen Sun, Weizhen Li, Hongxuan Lu, Fangchen Dong, Tianrui Qin, King Zhu, Minghao Liu, et al. 2025. Taskcraft: Automated generation of agentic tasks. *arXiv preprint arXiv:2506.10055*.

Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020a. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10740–10749.

Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. 2020b. Alfworld: Aligning text and embodied environments for interactive learning. *arXiv preprint arXiv:2010.03768*.

Significant-Gravitas. 2023. Auto-GPT: An Autonomous GPT-4 Experiment. GitHub repository. Accessed: 2023-09-02.

Huatong Song, Jinhao Jiang, Yingqian Min, Jie Chen, Zhipeng Chen, Wayne Xin Zhao, Lei Fang, and Ji-Rong Wen. 2025. R1-searcher: Incentivizing the search capability in llms via reinforcement learning. *arXiv preprint arXiv:2503.05592*.

Yifan Song, Weimin Xiong, Xiutian Zhao, Dawei Zhu, Wenhao Wu, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. Agentbank: Towards generalized llm agents via fine-tuning on 50000+ interaction trajectories. In *EMNLP (Findings)*.

Yifan Song, Weimin Xiong, Dawei Zhu, Wenhao Wu, Han Qian, Mingbo Song, Hailiang Huang, Cheng Li, Ke Wang, Rong Yao, Ye Tian, and Sujian Li. 2023. Restgpt: Connecting large language models with real-world restful apis. *Preprint*, arXiv:2306.06624.

Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang, and Dong Yu. 2025a. Crossing the reward bridge: Expanding rl with verifiable rewards across diverse domains. *Preprint*, arXiv:2503.23829.

Zhaochen Su, Peng Xia, Hangyu Guo, Zhenhua Liu, Yan Ma, Xiaoye Qu, Jiaqi Liu, Yanshu Li, Kaide Zeng, Zhengyuan Yang, et al. 2025b. Thinking with images for multimodal reasoning: Foundations, methods, and future frontiers. *arXiv preprint arXiv:2506.23918*.

Michael Sullivan, Mareike Hartmann, and Alexander Koller. 2025. Procedural environment generation for tool-use agents. *arXiv preprint arXiv:2506.11045*.

Hao Sun, Zile Qiao, Jiayan Guo, Xuanbo Fan, Yingyan Hou, Yong Jiang, Pengjun Xie, Yan Zhang, Fei Huang, and Jingren Zhou. 2025. Zerosearch: Incentivize the search capability of llms without searching. *Preprint*, arXiv:2505.04588.

Liangtai Sun, Xingyu Chen, Lu Chen, Tianle Dai, Zichen Zhu, and Kai Yu. 2022. Meta-gui: Towards multi-modal conversational agents on mobile gui. *Preprint*, arXiv:2205.11029.

Hao Tang, Darren Key, and Kevin Ellis. 2024. Worldcoder, a model-based llm agent: Building world models by writing code and interacting with the environment. *Advances in Neural Information Processing Systems*, 37:70148–70212.

Zhengwei Tao, Ting-En Lin, Xiancai Chen, Hangyu Li, Yuchuan Wu, Yongbin Li, Zhi Jin, Fei Huang, Dacheng Tao, and Jingren Zhou. 2024. A survey on self-evolution of large language models. *Preprint*, arXiv:2404.14387.

Zhengwei Tao, Jialong Wu, Wenbiao Yin, Junkai Zhang, Baixuan Li, Haiyang Shen, Kuan Li, Liwen Zhang, Xinyu Wang, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025. Webshaper: Agentically data synthesizing via information-seeking formalization. *Preprint*, arXiv:2507.15061.

Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.

Tongyi DeepResearch Team. 2025. Tongyi-deepresearch. https://github.com/Alibaba-NLP/DeepResearch.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O'Sullivan, and Hoang D Nguyen. 2025. Multi-agent collaboration mechanisms: A survey of llms. *arXiv preprint arXiv:2501.06322*.

Harsh Trivedi, Tushar Khot, Mareike Hartmann, Ruskin Manku, Vinty Dong, Edward Li, Shashank Gupta, Ashish Sabharwal, and Niranjan Balasubramanian. 2024. Appworld: A controllable world of apps and people for benchmarking interactive coding agents. *Preprint*, arXiv:2407.18901.

Haoming Wang, Haoyang Zou, Huatong Song, Jiazhan Feng, Junjie Fang, Junting Lu, Longxiang Liu, Qinyu Luo, Shihao Liang, Shijue Huang, Wanjun Zhong, Yining Ye, Yujia Qin, Yuwen Xiong, Yuxin Song, Zhiyong Wu, Aoyan Li, Bo Li, Chen Dun, Chong Liu, Daoguang Zan, Fuxing Leng, Hanbin Wang, Hao Yu, Haobin Chen, Hongyi Guo, Jing Su, Jingjia Huang, Kai Shen, Kaiyu Shi, Lin Yan, Peiyao Zhao, Pengfei Liu, Qinghao Ye, Renjie Zheng, Shulin Xin, Wayne Xin Zhao, Wen Heng, Wenhao Huang, Wenqian Wang, Xiaobo Qin, Yi Lin, Youbin Wu, Zehui Chen, Zihao Wang, Baoquan Zhong, Xinchun Zhang, Xujing Li, Yuanfan Li, Zhongkai Zhao, Chengquan Jiang, Faming Wu, Haotian Zhou, Jinlin Pang, Li Han, Qi Liu, Qianli Ma, Siyao Liu, Songhua Cai, Wenqi Fu, Xin Liu, Yaohui Wang, Zhi Zhang, Bo Zhou, Guoliang Li, Jiajun Shi, Jiale Yang, Jie Tang, Li Li, Qihua Han, Taoran Lu, Woyu Lin, Xiaokang Tong, Xinyao Li, Yichi Zhang, Yu Miao, Zhengxuan Jiang, Zili Li, Ziyuan Zhao, Chenxin Li, Dehua Ma, Feng Lin, Ge Zhang, Haihua Yang, Hangyu Guo, Hongda Zhu, Jiaheng Liu, Junda Du, Kai Cai, Kuanye Li, Lichen Yuan, Meilan Han, Minchao Wang, Shuyue Guo, Tianhao Cheng, Xiaobo Ma, Xiaojun Xiao, Xiaolong Huang, Xinjie Chen, Yidi Du, Yilin Chen, Yiwen Wang, Zhaojian Li, Zhenzhu Yang, Zhiyuan Zeng, Chaolin Jin, Chen Li, Hao Chen, Haoli Chen, Jian Chen, Qinghao Zhao, and Guang Shi. 2025a. Ui-tars-2 technical report: Advancing gui agent with multi-turn reinforcement learning. *Preprint*, arXiv:2509.02544.

Hongru Wang, Cheng Qian, Wanjun Zhong, Xiusi Chen, Jiahao Qiu, Shijue Huang, Bowen Jin, Mengdi Wang, Kam-Fai Wong, and Heng Ji. 2025b. Otc: Optimal tool calls via reinforcement learning. *arXiv e-prints*, pages arXiv–2504.

Ruoyao Wang, Graham Todd, Eric Yuan, Ziang Xiao, Marc-Alexandre Côté, and Peter Jansen. 2023. Bytesized32: A corpus and challenge task for generating task-specific world models expressed as text games. *arXiv preprint arXiv:2305.14879*.

Shaobo Wang, Zhengbo Jiao, Zifan Zhang, Yilang Peng, Xu Ze, Boyu Yang, Wei Wang, Hu Wei, and Linfeng Zhang. 2025c. Socratic-zero : Bootstrapping reasoning via data-free agent co-evolution. *Preprint*, arXiv:2509.24726.

Xingyao Wang, Yangyi Chen, Lifan Yuan, Yizhe Zhang, Yunzhu Li, Hao Peng, and Heng Ji. 2024. Executable code actions elicit better llm agents. In *Forty-first International Conference on Machine Learning*.

Xinyuan Wang, Bowen Wang, Dunjie Lu, Junlin Yang, Tianbao Xie, Junli Wang, Jiaqi Deng, Xiaole Guo, Yiheng Xu, Chen Henry Wu, Zhennan Shen, Zhuokai Li, Ryan Li, Xiaochuan Li, Junda Chen, Boyuan Zheng, Peihang Li, Fangyu Lei, Ruisheng Cao, Yeqiao Fu, Dongchan Shin, Martin Shin, Jiarui Hu, Yuyan Wang, Jixuan Chen, Yuxiao Ye, Danyang Zhang, Dikang Du, Hao Hu, Huarong Chen, Zaida Zhou, Haotian Yao, Ziwei Chen, Qizheng Gu, Yipu Wang, Heng Wang, Diyi Yang, Victor Zhong, Flood Sung, Y. Charles, Zhilin Yang, and Tao Yu. 2025d. Opencua: Open foundations for computer-use agents. *Preprint*, arXiv:2508.09123.

Zhenting Wang, Qi Chang, Hemani Patel, Shashank Biju, Cheng-En Wu, Quan Liu, Aolin Ding, Alireza Rezazadeh, Ankit Shah, Yujia Bao, et al. 2025e. Mcp-bench: Benchmarking tool-using llm agents with complex real-world tasks via mcp servers. *arXiv preprint arXiv:2508.20453*.

Jason Wei. 2025. Asymmetry of verification and verifier's rule. Blog post. https://www.jasonwei.net/blog/asymmetry-of-verification-and-verifiers-law.

Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. 2025. Browsecomp: A simple yet challenging benchmark for browsing agents. *arXiv preprint arXiv:2504.12516*.

Jialong Wu, Baixuan Li, Runnan Fang, Wenbiao Yin, Liwen Zhang, Zhengwei Tao, Dingchu Zhang, Zekun Xi, Gang Fu, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025a. Webdancer: Towards autonomous information seeking agency. *Preprint*, arXiv:2505.22648.

Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, and Fei Huang. 2025b. Webwalker: Benchmarking llms in web traversal. *Preprint*, arXiv:2501.07572.

Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2025. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2):121101.

Zhiheng Xi, Yiwen Ding, Wenxiang Chen, Boyang Hong, Honglin Guo, Junzhe Wang, Dingwen Yang, Chenyang Liao, Xin Guo, Wei He, et al. 2024. Agentgym: Evolving large language model-based agents across diverse environments. *CoRR*.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. 2024. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *Preprint*, arXiv:2404.07972.

Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. 2025. Wizardlm: Empowering large pre-trained language models to follow complex instructions. *Preprint*, arXiv:2304.12244.

Zihui Xue, Kumar Ashutosh, and Kristen Grauman. 2024. Learning object state changes in videos: An open-world perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18493–18503.

Ziyi Yang, Zaibin Zhang, Zirui Zheng, Yuxian Jiang, Ziyue Gan, Zhiyu Wang, Zijian Ling, Jinsong Chen, Martz Ma, Bowen Dong, Prateek Gupta, Shuyue Hu, Zhenfei Yin, Guohao Li, Xu Jia, Lijun Wang, Bernard Ghanem, Huchuan Lu, Chaochao Lu, Wanli Ouyang, Yu Qiao, Philip Torr, and Jing Shao. 2025. Oasis: Open agent social interaction simulations with one million agents. *Preprint*, arXiv:2411.11581.

Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2023. Webshop: Towards scalable real-world web interaction with grounded language agents. *Preprint*, arXiv:2207.01206.

Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. 2024. τ-bench: A benchmark for tool-agent-user interaction in real-world domains. *arXiv preprint arXiv:2406.12045*.

Jiabo Ye, Xi Zhang, Haiyang Xu, Haowei Liu, Junyang Wang, Zhaoqing Zhu, Ziwei Zheng, Feiyu Gao, Junjie Cao, Zhengxi Lu, Jitong Liao, Qi Zheng, Fei Huang, Jingren Zhou, and Ming Yan. 2025a. Mobile-agent-v3: Fundamental agents for gui automation. *Preprint*, arXiv:2508.15144.

Junjie Ye, Zhengyin Du, Xuesong Yao, Weijian Lin, Yufei Xu, Zehui Chen, Zaiyuan Wang, Sining Zhu, Zhiheng Xi, Siyu Yuan, Tao Gui, Qi Zhang, Xuanjing Huang, and Jiecao Chen. 2025b. Toolhop: A query-driven benchmark for evaluating large language models in multi-hop tool use. *Preprint*, arXiv:2501.02506.

Junjie Ye, Changhao Jiang, Zhengyin Du, Yufei Xu, Xuesong Yao, Zhiheng Xi, Xiaoran Fan, Qi Zhang, Tao Gui, Xuanjing Huang, and Jiecao Chen. 2025c. Feedback-driven tool-use improvements in large language models via automated build environments. *Preprint*, arXiv:2508.08791.

Asaf Yehudai, Lilach Eden, Alan Li, Guy Uziel, Yilun Zhao, Roy Bar-Haim, Arman Cohan, and Michal Shmueli-Scheuer. 2025. Survey on evaluation of llm-based agents. *arXiv preprint arXiv:2503.16416*.

Peijie Yu, Yifan Yang, Jinjian Li, Zelong Zhang, Haorui Wang, Xiao Feng, and Feng Zhang. 2025a. c^3-bench: The things real disturbing llm based agent in multi-tasking. *arXiv preprint arXiv:2505.18746*.

Tianyu Yu, Bo Ji, Shouli Wang, Shu Yao, Zefan Wang, Ganqu Cui, Lifan Yuan, Ning Ding, Yuan Yao, Zhiyuan Liu, et al. 2025b. Rlpr: Extrapolating rlvr to general domains without verifiers. *arXiv preprint arXiv:2506.18254*.

Abhay Zala, Jaemin Cho, Han Lin, Jaehong Yoon, and Mohit Bansal. 2024. Envgen: Generating and adapting environments via llms for training embodied agents. *arXiv preprint arXiv:2403.12014*.

Guibin Zhang, Hejia Geng, Xiaohang Yu, Zhenfei Yin, Zaibin Zhang, Zelin Tan, Heng Zhou, Zhongzhi Li, Xiangyuan Xue, Yijiang Li, Yifan Zhou, Yang Chen, Chen Zhang, Yutao Fan, Zihu Wang, Songtao Huang, Yue Liao, Hongru Wang, Mengyue Yang, Heng Ji, Michael Littman, Jun Wang, Shuicheng Yan, Philip Torr, and Lei Bai. 2025a. The landscape of agentic reinforcement learning for llms: A survey. *Preprint*, arXiv:2509.02547.

Jun Zhang, Yuwei Yan, Junbo Yan, Zhiheng Zheng, Jinghua Piao, Depeng Jin, and Yong Li. 2025b. A parallelized framework for simulating large-scale llm agents with realistic environments and interactions. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, pages 1339–1349.

Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. 2025c. Generative verifiers: Reward modeling as next-token prediction. *Preprint*, arXiv:2408.15240.

Wentao Zhang, Liang Zeng, Yuzhen Xiao, Yongcong Li, Ce Cui, Yilei Zhao, Rui Hu, Yang Liu, Yahui Zhou, and Bo An. 2025d. Agentorchestra: A hierarchical multi-agent framework for general-purpose task solving. *arXiv preprint arXiv:2506.12508*.

Yanfei Zhang. 2025. Agent-as-tool: A study on the hierarchical decision making with reinforcement learning. *arXiv preprint arXiv:2507.01489*.

Yusen Zhang, Ruoxi Sun, Yanfei Chen, Tomas Pfister, Rui Zhang, and Sercan Arik. 2024. Chain of agents: Large language models collaborating on long-context tasks. *Advances in Neural Information Processing Systems*, 37:132208–132237.

Zhong Zhang, Yaxi Lu, Yikun Fu, Yupeng Huo, Shenzhi Yang, Yesai Wu, Han Si, Xin Cong, Haotian Chen, Yankai Lin, et al. 2025e. Agentcpm-gui: Building mobile-use agents with reinforcement fine-tuning. *arXiv preprint arXiv:2506.01391*.

Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Yang Yue, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. 2025. Absolute zero: Reinforced self-play reasoning with zero data. *Preprint*, arXiv:2505.03335.

Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, et al. 2023. Webarena: A realistic web environment for building autonomous agents. *arXiv preprint arXiv:2307.13854*.

A Additional Related Work

Scaling Laws for LLM Agents Just as large language models exhibit predictable performance scaling with increases in the number of parameters, the volume of training data, and the compute budget, agent systems likewise display scaling regularities along three axes: (i) expanding the agent population and identifying properties that emerge as interactions increase; (ii) increasing environmental complexity and assessing how realistic, dynamic settings shape learning and adaptation; and (iii) extending the horizons of evolution and memory to study how agents generalize and improve through accumulated experience (CAMEL-AI, 2025). While most existing surveys on LLM agents adopt an agent-centric view (Luo et al., 2025b; Xi et al., 2025; Yehudai et al., 2025; Gao et al., 2025), covering topics from multi-agent interaction (Qian et al., 2024; Tran et al., 2025) to self-evolution (Gao et al., 2025; Tao et al., 2024), environment scaling remains underexplored and has not been systematically organized, despite its central role in experience formation through the Generation–Execution–Feedback (GEF) loop. In this work, we take an environment-centric perspective on scaling environments and examine how dynamic, richly interactive, high-fidelity worlds can accelerate agent development and evolution.

B Conceptual Framework

Following the formalization by Gao et al. (2025), we model the environment E as a partially observable Markov decision process (POMDP). At the beginning of each episode, the environment generates a task $T = (E, I)$, where $I \in \mathcal{I}$ represents user intention drawn from intention space $\mathcal{I}$. The agent π interacts with the environment over horizon $T \in \mathbb{N}$, producing an interleaved observation–action trajectory $\tau = (o_0, a_0, o_1, a_1, \ldots, o_T) \in (\mathcal{O} \times \mathcal{A})^{T+1}$, where $\mathcal{O}$ and $\mathcal{A}$ denote the observation and action spaces respectively. The environment then evaluates performance and provides feedback $r \in \mathbb{R}^k$, which may take the form of step-level signals $r_{\text{step}}^{(t)} = R_{\text{step}}(s_t, a_t, T) \in \mathbb{R}$ for $t \in \{0, \ldots, T\}$, where $s_t \in \mathcal{S}$ denotes the state of the environment at step t, trajectory-level signals $r_{\text{traj}} = R_{\text{traj}}(\tau, T) \in \mathbb{R}^m$, or a combination of both $r = f(r_{\text{step}}^{(0:T)}, r_{\text{traj}})$ where $f : \mathbb{R}^{T+1} \times \mathbb{R}^m \to \mathbb{R}^k$. These signals need not be limited to sparse scalar rewards but can encode structured or adaptive assessments reflecting correctness, efficiency, reasoning depth, or long-term outcomes. This Generation-Execution-Feedback (GEF) Loop $\mathcal{L} = (T_{\text{gen}}, \text{Exec}, \text{Eval}) : \mathcal{I} \to \mathcal{T} \times (\mathcal{O} \times \mathcal{A})^* \times \mathbb{R}^k$, which encompasses task generation, task execution, and feedback, defines the essential mechanics of environments. Repeated iterations drive the accumulation of experience and the progressive evolution of the agent π.

C Evaluation Benchmarks

Direct evaluation of environments remains underexplored, with most research relying on indirect metrics. For example, works such as TaskCraft (Shi et al., 2025) and AgentScaler (Fang et al., 2025) assess environments by training agents through the resulting trajectories of environment–agent interactions, where strong agent performance indirectly reflects environmental quality. While some direct evaluation has emerged, it has largely been confined to symbolic or text-based settings. Bytesized32 (Wang et al., 2023) introduces task-specific

text games and evaluates them using automated metrics on fidelity, validity, specification adherence, and winnability, bridging code generation and cognitive modeling. Text2World (Hu et al., 2025b) further benchmarks symbolic world model generation, employing structural similarity for overall evaluation while also using component-wise F1 scores to capture finer-grained aspects such as action dynamics. Recent work has begun to extend direct evaluation to video-based domains. VidOSC (Xue et al., 2024) examines open-world dynamics, WorldScore (Duan et al., 2025) proposes a unified framework for evaluating world generation, and WorldPrediction (Chen et al., 2025a) targets high-level visual reasoning, focusing on long-horizon program planning and semantic-temporal abstraction. Despite these advances, comprehensive and general-purpose evaluation protocols remain scarce, highlighting the need for broader, domain-agnostic methodologies to rigorously and directly assess environment quality beyond agent-centric performance metrics.

D Key Applications

Recent progress in agentic systems, typically built on state-of-the-art LLM families such as GPT (OpenAI, 2025a), Claude (Anthropic, 2025), Gemini (Team et al., 2024), LLaMA (Touvron et al., 2023), and Qwen (Bai et al., 2023), is increasingly driven by interactions with dynamic and multifaceted environments.

Tool-use environments Tool-use environments expose APIs and function calls as structured action spaces, and many LLMs now natively support function invocation, thereby extending reasoning with external tools (OpenAI, 2023; Anthropic, 2024; Song et al., 2023; Wu et al., 2025b; Mastouri et al., 2025; Luo et al., 2025c; Wang et al., 2025e; Fan et al., 2025).

Coding environments Coding environments leverage repositories, test frameworks, and IDE integration to support long-horizon programming. Within these settings, systems such as Qwen3-Coder (Qwen Team, 2024) and Claude 4 (Anthropic, 2025) demonstrate reliability in code editing and debugging, and ReTool (Feng et al., 2025) integrates code-interpreter execution into the reasoning loop, enabling agents to exhibit code self-correction and adaptive tool selection.

Web navigation environments Web navigation environments build on HTML/DOM structures to support tasks such as browsing, form filling, and transactions. AutoWebGLM (Lai et al., 2024) illustrates how simplified webpage representations can enhance navigation efficiency, while open-source agents such as AgenticSeek (Fosowl, 2024) and Auto-GPT (Significant-Gravitas, 2023) showcase how browser-control environments enable general-purpose task execution on the web.

GUI agent environments GUI agent environments extend these capabilities to graphical user interfaces on desktops and mobile devices (Qin et al., 2025; Wang et al., 2025d; Ye et al., 2025a; Lu et al., 2025b; Liu et al., 2024a; Sun et al., 2022; Huang et al., 2025b). Beyond these, more comprehensive platforms (Wang et al., 2025a; Xie et al., 2024; OpenAI, 2025) integrate terminals, operating systems, applications, and APIs to provide more interactive, open-ended, and real-world computer-use environments.

Deep research environments Deep research environments require both long-context reasoning and robust retrieval. Systems such as Gemini1.5Pro (Google DeepMind, 2024) and OpenAI's deep research agents (OpenAI, 2025b) show that extended context windows enable sustained analysis, while effective retrieval pipelines mitigate distraction and context dilution. Across these domains, environments transform general LLM capacities into agent skills by defining action spaces, providing feedback, and maintaining state continuity.

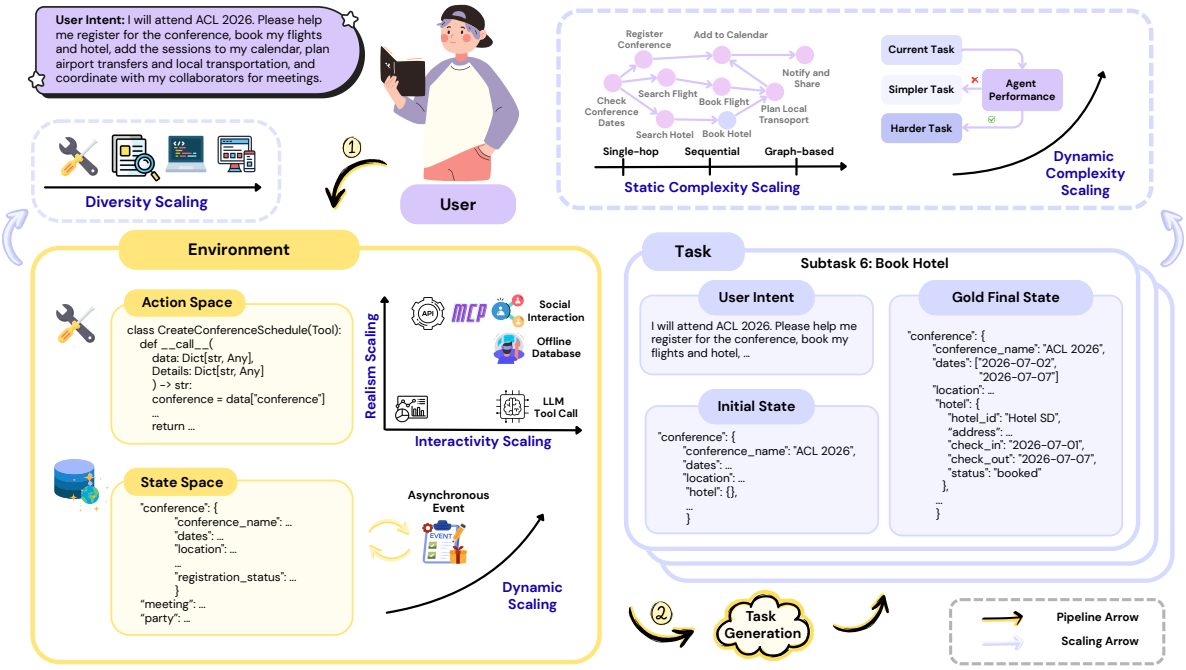

Figure 3: Illustration of environment scaling in the **task generation** and **task execution** stages, using the example of conference scheduling. Given a user intent, the environment produces a set of tasks for the agent to complete. Scaling in the task generation stage covers *complexity scaling*, *dynamic scaling*, and *diversity scaling*, while in the task execution stage scaling encompasses *interactivity scaling* and *realism scaling*.

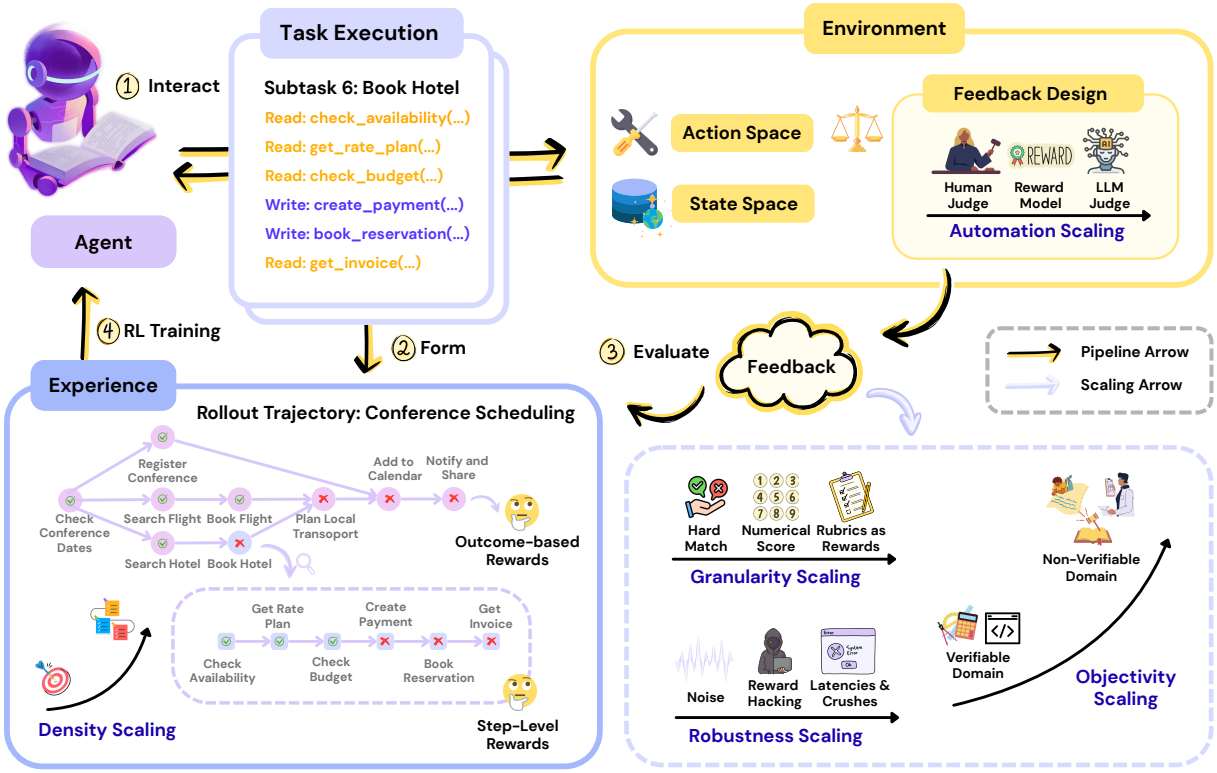

Figure 4: Illustration of environment scaling in the **feedback** stage using a conference-scheduling example. The agent first executes tasks in the environment and produces action–observation trajectories. The environment then evaluates these trajectories and returns feedback, yielding the experience used to train the agent. Scaling in the feedback stage covers *density*, *granularity*, *automation*, *objectivity*, and *robustness*.