WEAGENT-MMSEARCH: NATIVE TEXT-VISION INTERACTION FOR MULTIMODAL SEARCH AGENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Multimodal search agents extend parametric knowledge with newly emerging and long-tail evidence from the open web. Yet many existing agentic search environments often expose retrieved evidence only as text, omitting tool-returned images and reducing visually grounded search to text-only reasoning. Long-horizon interaction also compounds tool-call, response-length, timeout, and budget failures, which can discard salvageable trajectories, waste rollout computation, and disturb policy updates. To address these issues, we introduce WEAGENT-HARNESS, a multimodal agentic harness that supports native text–vision interaction and runtime recovery. Retrieved images receive persistent disk references, allowing the model to inspect, process, and cite them throughout the trajectory. Based on this harness, we develop **WeAgent-MMSearch**, an integrated system spanning data construction, agentic post-training, and multimodal rollout. For data construction, a strong MLLM uses WEAGENT-HARNESS to discover, synthesize, and verify MMSearch-style tasks and collect expert trajectories. During post-training, our Failure-Aware GSPO (FA-GSPO) recovers salvageable abnormal rollouts and filters invalid ones to improve bounded multimodal planning and search. We also introduce VISTARGET-BENCH, a 150-task human-verified benchmark that pairs each question with a held-out target image, distinguishing image-retrieval failures from visual-perception failures. Evaluation on VISTARGET-BENCH and seven public benchmarks shows our agentic post-training yields a 19.22-point average improvement, outperforming similarly-sized open-source models and rivaling those ten times larger.

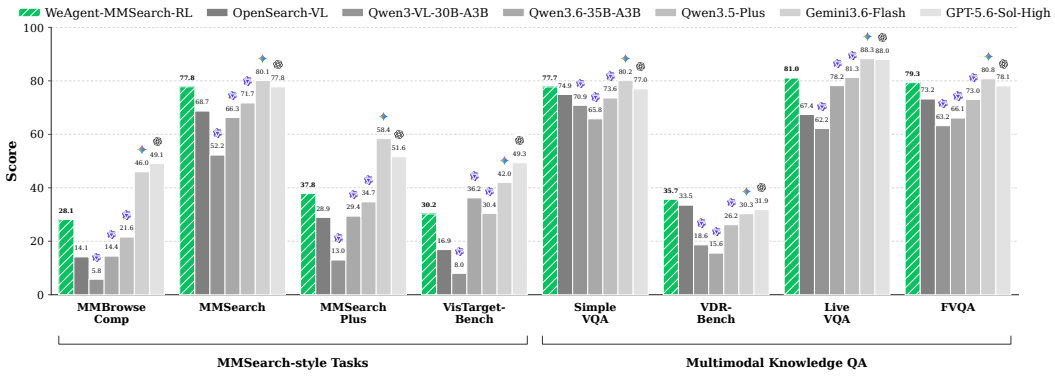

Figure 1: Benchmark performance of WeAgent-MMSearch.

1 INTRODUCTION

Multimodal large language models (MLLMs) acquire broad visual and textual knowledge from static corpora(OpenAI, 2024; Gemini Team et al., 2023; Bai et al., 2025), yet finite coverage and parametric memory leave long-tail entities, recent events, and source-specific facts unsupported or stale. Search, webpages, and image tools extend this boundary. Multimodal search (MMSearch)

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

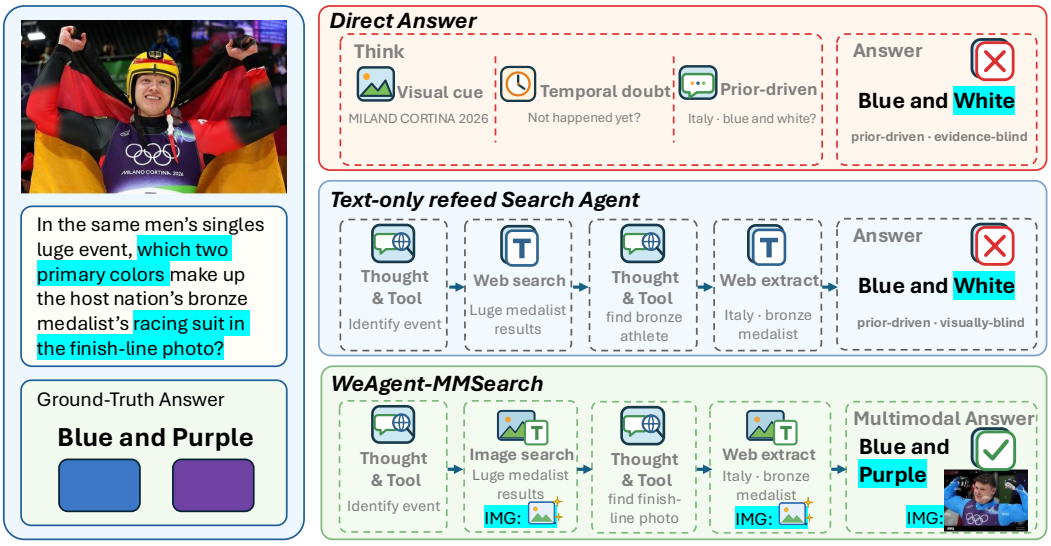

Figure 2: A time-sensitive multi-hop search example. An MLLM fails when answering directly (*top*). The same MLLM (e.g., Kimi K2.6) also fails with a text-only harness that omits retrieved images (*middle*). WeAgent-MMSearch preserves and re-feeds retrieved images across turns and answers correctly (*bottom*).

therefore requires models to identify information gaps, invoke the right tools, and integrate external multimodal evidence (Jin et al., 2025; Li et al., 2025a; Hu et al., 2023; Zhang et al., 2024; Jiang et al., 2025).

Many search environments return retrieved evidence as text and omit tool-returned images from later context. Visual facts that cannot be verbalized then disappear from the model-visible state, and multimodal trajectories collapse into text-only reasoning. Figure 2 shows that direct answering and a text-only harness miss the answer, while WeAgent-MMSearch re-feeds the images and succeeds.

Long-horizon interaction also amplifies tool-name or parameter errors, oversized outputs, timeouts, and budget exhaustion. These failures interrupt reasoning, block valid answers, waste rollout compute, and form a long tail of expensive episodes that lowers throughput. Kept as ordinary samples, they contaminate optimization with failures unrelated to search policy quality.

High-quality supervision is scarce. Unlike standard visual question answering or text-only multi-hop retrieval, MMSearch-style tasks must stay visually dependent across steps and executable under the tool contract. Expert trajectories need live search, webpage access, image acquisition, and code execution, so task synthesis, validity checking, and rollout collection limit open MMSearch-style supervision (Chu et al., 2026; Chen et al., 2026).

To address these issues, we introduce WEAGENT-HARNESS, a multimodal agentic harness for native text–vision interaction and runtime recovery. It integrates search, webpage, image, reverse-image, and code tools, and stores retrieved images on disk under persistent references with provenance, so the model can inspect, process, and cite them throughout the trajectory. It also caches repeated requests, repairs only unambiguous call errors, returns structured errors for self-correction, and intervenes before local failures or budget exhaustion invalidate a rollout. This improves answer completion, reduces wasted inference compute, salvages valid training samples, and keeps irrecoverable artifacts out of policy updates.

On this harness we develop **WeAgent-MMSearch**, spanning data construction, agentic post-training, and multimodal rollout. During data construction, a strong MLLM uses WEAGENT-HARNESS under the same tool and state contract as training and inference. It discovers seeds from Wikipedia (Wikimedia Foundation, 2026) and time-sensitive web events, expands them backward into multi-hop chains, verifies chain validity, composes MMSearch-style tasks, and collects expert trajectories. The four stages yield 3.3K in-house SFT trajectories and 4.9K RL prompts. We further filter and

normalize 56.2K open-source SFT trajectories and 8.3K open-source RL prompts, for final mixtures of 59.5K and 13.3K.

On these mixtures we perform agentic post-training. SFT teaches a ReAct-style search policy (Yao et al., 2022) from successful demonstrations, with tool observations and harness interventions masked from the loss. For RL, Failure-Aware GSPO (FA-GSPO) extends Group Sequence Policy Optimization (GSPO) (Zheng et al., 2025): salvageable abnormal rollouts are retained and invalid ones are filtered, improving bounded multimodal planning and search. We also introduce VISTARGET-BENCH, 150 human-verified tasks, each paired with a held-out target image so image-retrieval failures can be separated from visual-perception failures.

Overall, our contributions are WEAGENT-HARNESS for native text–vision interaction and runtime recovery, a four-stage pipeline for executable MMSearch-style tasks and expert trajectories, and agentic SFT with FA-GSPO. On VISTARGET-BENCH and seven public benchmarks, the harness alone raises Qwen3-VL-30B-A3B from 19.13% to 36.75% and Kimi K2.6 from 21.92% to 58.42%, without parameter updates. Post-training then reaches 51.04% (WeAgent-MMSearch-SFT) and 55.97% (WeAgent-MMSearch-RL), a 19.22-point gain over the harness-only base. WeAgent-MMSearch-RL outperforms similarly sized open-source models and rivals models with roughly ten times its parameter count.

2 RELATED WORK

Agentic search. Early visual information-seeking systems are modular: AVIS coordinates a planner and reasoner with web, image-search, and perception tools, while Vision Search Assistant couples a vision-language model with a web agent (Hu et al., 2023; Zhang et al., 2024). Later MMSearch agents adapt text-search RL (Jin et al., 2025; Li et al., 2025a) so that one vision-language model interleaves reasoning and search (Wu et al., 2025; Narayan et al., 2025; Huang et al., 2026; Geng et al., 2025). OpenSearch-VL and Skywork-R1V4 further add active image manipulation, including OCR, cropping, and enhancement (Chen et al., 2026; Zhang et al., 2025). Across these systems, images in retrieved pages are often absent from model-visible observations, and retaining tool-returned images across turns remains uneven. WEAGENT-HARNESS stores those images under stable references for later reasoning and tool calls.

Data and optimization for long-horizon search. Long-horizon search learning is limited by scarce tasks and the cost of tool-interleaved trajectories. REDSearcher sets synthetic difficulty through reasoning-graph topology and evidence dispersion, and co-designs mid-training with post-training to cut interaction cost (Chu et al., 2026). OpenSearch-VL samples Wikipedia paths, applies fuzzy entity rewriting and source-anchored visual grounding, and pairs the resulting data with a diverse tool environment and failure-aware agentic RL (Chen et al., 2026). Both mainly synthesize multi-hop tasks from existing entities, so discovering new seeds—especially time-sensitive events—falls outside their scope. Our pipeline instead discovers Wikipedia entities and time-sensitive web events, then builds visually dependent backward multi-hop tasks under the same execution contract used for training and inference.

Multimodal search benchmarks. MMSearch evaluates component-level and end-to-end search (Jiang et al., 2025); MMSearch-Plus emphasizes provenance-aware search and evidence attribution (Tao et al., 2026); and MM-BrowseComp uses expert-crafted multi-hop questions grounded in multimodal web content (Li et al., 2025b). These benchmarks broaden multimodal information seeking, but they are not built around the target-image diagnostic studied here. In VISTARGET-BENCH, every question is paired with a held-out target image, and the agent returns both its answer and an index of the supporting image. Comparing that selection with the target separates final-answer accuracy, image retrieval, and answer correctness after a match.

3 WEAGENT-HARNESS: MULTIMODAL AGENTIC HARNESS

WEAGENT-HARNESS is the shared interaction substrate for task construction, trajectory collection, post-training, and inference. It keeps newly acquired visual evidence usable throughout a long interaction, and keeps open-web execution reliable and bounded enough to produce train-

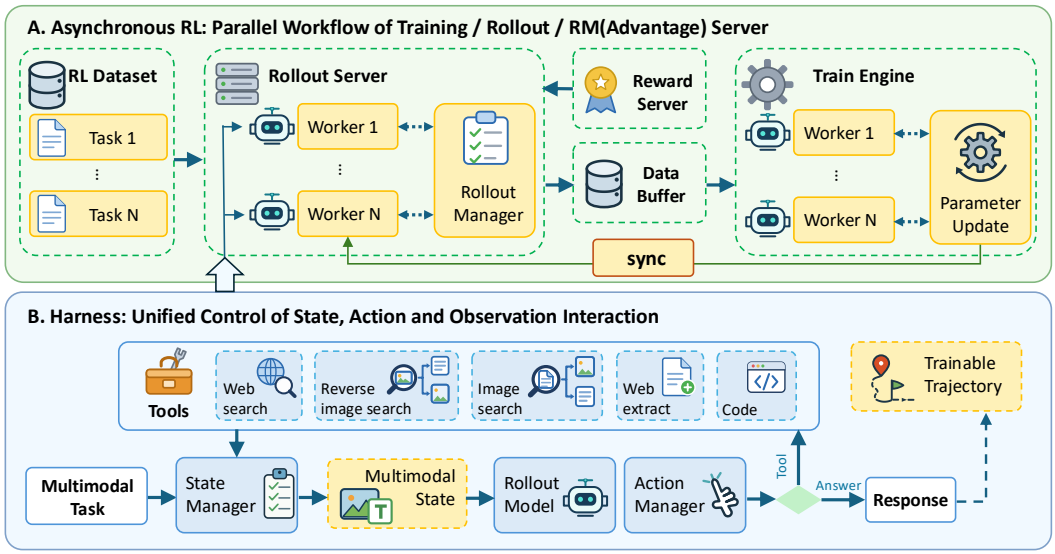

Figure 3: Asynchronous RL infrastructure and WEAGENT-HARNESS. **Top:** the SGLang rollout server, SGLang/API reward server, and Megatron training engine (Zheng et al., 2024; Shoeybi et al., 2019) advance concurrently. **Bottom:** the harness owns the state, action, and text–image cycle, routes tool calls, and serializes each result as a trainable trajectory.

able trajectories. It therefore exposes `web_search`, `web_extract`, `web_search_for_image`, `web_search_based_on_image`, and `execute_code` through one multimodal state and evidence contract. Figure 3 situates this contract within the asynchronous RL infrastructure.

3.1 PERSISTENT MULTIMODAL STATE

WEAGENT-HARNESS keeps a persistent evidence workspace across turns. Textual history and registered images are serialized together into each later policy context, so newly retrieved pixels directly condition subsequent reasoning and tool calls. Outputs of `web_extract`, `web_search_for_image`, `web_search_based_on_image`, and `execute_code` become multimodal observations of text and any returned images. Each image is bound to a model-visible reference and stable identity, such as `<user_image_1>` or `<search_image_3>`, together with its source metadata, so later turns can revisit, reverse-search, or compare it.

3.2 RELIABLE AND BOUNDED EXECUTION

Cache-backed search balances freshness and latency. Open-web latency is heavy-tailed and can stall a rollout group on a few slow searches. A shared exact-match cache over normalized requests stores successful data-collection and training responses with their content and provenance. Workers read a stable snapshot: a hit returns immediately through the same tool schema, and a miss calls the live backend. Repeated requests trade a little freshness for predictable latency; unseen or time-sensitive requests still reach the live web, and periodic refreshes keep the pool from becoming a fixed offline corpus.

Recovery and budgets shape trainable rollouts. WEAGENT-HARNESS repairs a call only when the fix is unambiguous, correcting tool names, parameter names, and serialization. Unclear intent instead returns a structured error observation for self-correction, and repeated invalid calls are suppressed. A response past `max-response-len` is rejected and regenerated once, so one anomalous turn does not discard the trajectory. Turn, tool, context, and wall-clock budgets bound every episode; once the regular turn budget is spent, tools are disabled and one final evidence-synthesis turn is allowed. Trajectories that exceed the context-length or wall-clock limit are terminated and excluded from training.

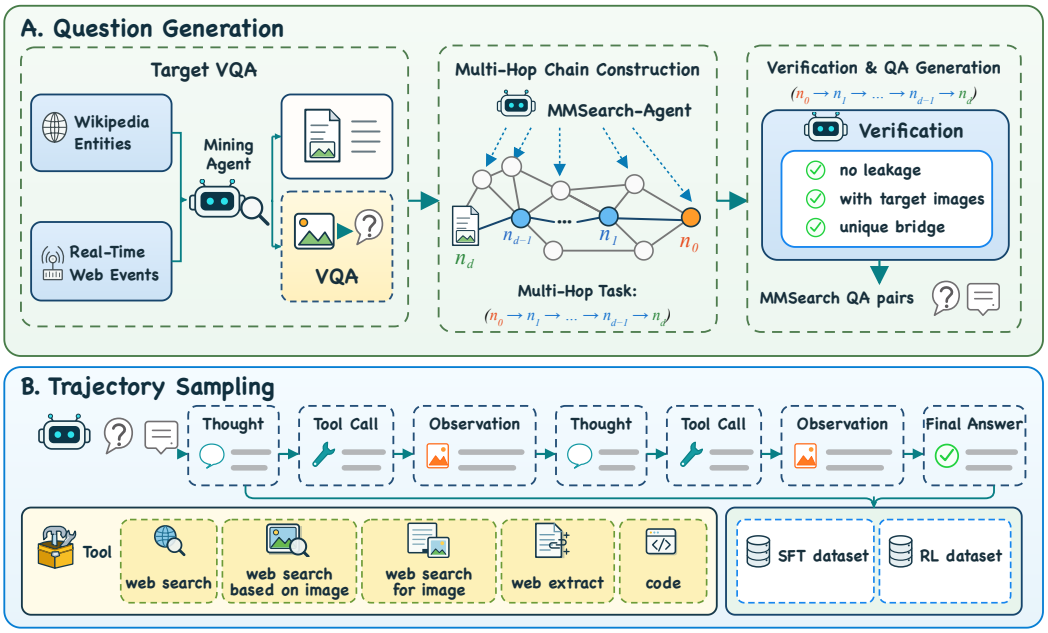

Figure 4: Overview of the WEAGENT-HARNESS-native data pipeline. **A. Task construction:** Wikipedia entities and time-sensitive web events provide target-image seeds, which are expanded backward into multi-hop chains and validated for image dependence, bridge uniqueness, and answer leakage. **B. Trajectory sampling:** the same harness executes accepted tasks; successful trajectories form SFT data, while queries with mixed rollout outcomes are retained for RL.

Persistent evidence determines what the policy can observe, while cache lookup, recovery, and budgets determine which interactions become finite, replayable training episodes. The following sections use this same contract first as a data factory and then as the environment for SFT, RL, and inference.

4 HARNESS-NATIVE MMSEARCH DATA PIPELINE

We instantiate Kimi K2.6 on WEAGENT-HARNESS as the expert MMSearch agent (Moonshot AI, 2026). It uses the same tools, observations, and evidence registry later available during post-training and inference. Four in-house stages transform visually grounded seeds into executable trajectories: seed generation, backward multi-hop expansion, task validation, and trajectory sampling.

Figure 4 summarizes the pipeline. Stage 1 builds visually grounded seeds from time-sensitive web events and Wikipedia (Wikimedia Foundation, 2026), and keeps a seed only when its target image is reachable through WEAGENT-HARNESS, its answer depends on visible pixels rather than a title, URL, caption, or page metadata, and its provenance can be audited. Stage 2 expands each verified seed backward into the chain $(n_0 \longrightarrow \cdots \longrightarrow n_d)$ with $d \in \{2, 3, 4\}$: given only the current node, the expert writes one predecessor question whose answer is that node, and repeats this step d times. Solving the task therefore requires following the bridge relations, retrieving target image C, and inspecting its pixels to produce D. Stage 3 keeps a chain only when it is internally consistent, the target image is present, every bridge is unique, and the answer is not exposed by the URL, metadata, or question text, then synthesizes the user-facing query. Stage 4 executes accepted tasks under the same WEAGENT-HARNESS contract: correct trajectories form the SFT set, and a stratified mix of model-solved and model-unsolved tasks forms the RL query set, with benchmark candidates held out at the seed and task levels. Stage procedures are detailed in Appendix A. These stages yield 3.3K in-house SFT trajectories and 4.9K RL prompts. Open-source trajectories are normalized into the same contract and mixed in; see Appendix B. The final mixtures contain 59.5K SFT trajectories and 13.3K RL prompts.

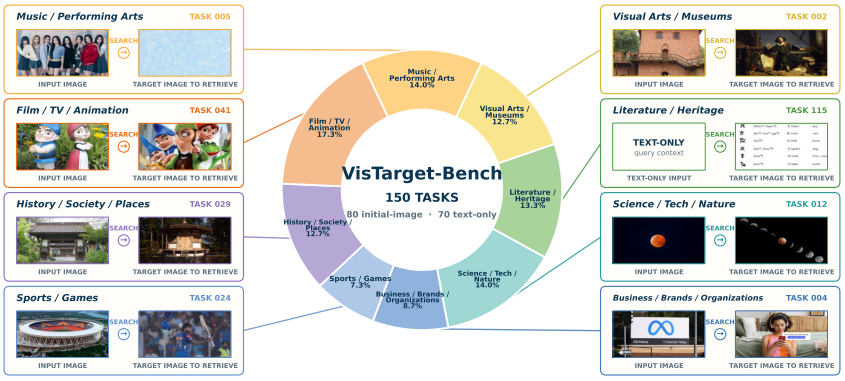

Figure 5: Composition of VISTARGET-BENCH. The donut is the terminal-evidence domain mix. Each pair is an input (image or text) and the held-out target image to retrieve.

5 VISTARGET-BENCH: DIAGNOSING VISUAL EVIDENCE USE

Final-answer accuracy alone cannot reveal whether an agent retrieved the intended visual evidence or used it correctly after retrieval. VISTARGET-BENCH pairs every task with a held-out target image and asks the agent to index its supporting image, separating final-answer accuracy from target-image retrieval and from answer correctness after a match. Tasks come from seeds held out before SFT and RL splitting and follow the harness-native pipeline in Section 4; the additional stage is human verification. Reviewers correct factual and relational ambiguities, remove answer shortcuts, and confirm that each task requires retrieving and reading the target image. The benchmark contains 150 tasks: 80 (53.3%) include an input image and 70 (46.7%) begin with text only. Figure 5 summarizes the domain mix and representative tasks.

The model returns a final answer and the index of its supporting image; accuracy uses only the final answer. For task i, a VLM judge compares the selected image $\widehat{C}_i$ with the held-out target C_i. Let $I_i \in \{0,1\}$ indicate that the response cites a tool-retrieved image, M_i denote target correspondence ($M_i = 0$ if none is cited), and $A_i \in \{0,1\}$ denote answer correctness. Image Cite, Target Match, and Match+Acc are $\text{Cite} = \sum_{i=1}^{N} I_i/N$, $\text{Match} = \sum_{i=1}^{N} M_i / \sum_{i=1}^{N} I_i$, and $\text{Match+Acc} = \sum_{i=1}^{N} A_i M_i / \sum_{i=1}^{N} M_i$. A conditional rate is zero when its conditioning set is empty.

6 AGENTIC POST-TRAINING

Using the final mixtures from Section 4, SFT imitates successful demonstrations to learn when to search and which tool to call, while RL optimizes the policy on outcome rewards beyond those demonstrations. Configurations are in Appendix D.

6.1 SUPERVISED FINE-TUNING

We apply next-token prediction to correct, complete trajectories. Tool-returned observations and the intervention prompts that WEAGENT-HARNESS inserts as user-role messages are masked, so the loss covers only model-generated tokens. This teaches a ReAct-style retrieval policy (Yao et al., 2022): the model reasons over the available evidence, searches when external information is needed, and revises its plan after a new visual observation.

6.2 AGENTIC REINFORCEMENT LEARNING

SFT learns successful interaction patterns but does not directly optimize behavioral policies. We therefore use Failure-Aware GSPO (FA-GSPO), an extension of Group Sequence Policy Optimization (GSPO) (Zheng et al., 2025). It separates runtime recoverability from optimization eligibility: candidates are sampled online, and only trajectories that remain useful after recovery enter the update.

Failure-aware recovery. Vanilla GSPO treats every group member as a valid sample. A timeout or an overlong response can instead truncate a trajectory for reasons unrelated to MMSearch; Appendix C.2 shows that such cases are not rare. FA-GSPO recovers first and then filters. An overlong turn is regenerated once, a malformed tool call is returned as a structured observation for self-correction, and an exhausted budget triggers one final tool-free answer. Trajectories that remain overlong, or that exceed the wall-clock or context budget, are excluded from optimization.

Eligible-only group normalization. For each prompt q_g in a batch $\mathcal{B} = \{q_g\}_{g=1}^{B} \sim \mathcal{D}$, we draw $\{o_{gi}\}_{i=1}^{G}$ online and keep the eligible set $\mathcal{E}_g = \{i \in \{1,\dots,G\} \mid o_{gi} \text{ remains trainable after runtime checks}\}$, with $n_g = |\mathcal{E}_g|$. Only these trajectories enter group normalization: $\bar{R}_g = \frac{1}{\max(n_g,1)} \sum_{i\in\mathcal{E}_g} R_{gi}$ and $\sigma_g = \sqrt{\frac{1}{\max(n_g-1,1)} \sum_{i\in\mathcal{E}_g} (R_{gi} - \bar{R}_g)^2}$. The reward comes from the terminal answer, $R_{gi} = 0.9\, C_{gi} + 0.1\, F_{gi}$ with $C_{gi}, F_{gi} \in \{0,1\}$, where C_{gi} is answer correctness and F_{gi} indicates that the final turn obeys the required format. The advantage is $\widehat{A}_{gi} = (R_{gi} - \bar{R}_g)/(\sigma_g + \delta)$ for $i \in \mathcal{E}_g$. Invalid trajectories contribute neither to these statistics nor to the policy gradient.

Failure-Aware GSPO objective. We plug these advantages into the clipped, sequence-level GSPO objective. For $\mathcal{B} = \{q_g\}_{g=1}^{B} \sim \mathcal{D}$, let $\mathbf{o}_g = \{o_{gi}\}_{i=1}^{G} \sim \pi_{\theta_{\text{old}}}(\cdot \mid q_g)$ and $\mathbf{o} = \{\mathbf{o}_g\}_{g=1}^{B}$:

$$
\begin{aligned}
\mathcal{J}_{\text{FA-GSPO}}(\theta) &= \mathbb{E}_{\mathcal{B},\mathbf{o}} \left[\frac{1}{\max(|\mathcal{G}_{\text{train}}|, 1)} \sum_{g\in\mathcal{G}_{\text{train}}} \frac{1}{n_g} \sum_{i\in\mathcal{E}_g} S_{gi}(\theta) \right], \\
S_{gi}(\theta) &= \min\left(\rho_{gi}(\theta)\widehat{A}_{gi},\ \text{clip}(\rho_{gi}(\theta), 1-\epsilon_{\text{lo}}, 1+\epsilon_{\text{hi}})\, \widehat{A}_{gi} \right), \\
\rho_{gi}(\theta) &= \exp\left[\frac{1}{T_{gi}} \sum_{t=1}^{T_{gi}} \log \frac{\pi_\theta(a_{gi,t} \mid h_{gi,t})}{\pi_{\theta_{\text{old}}}(a_{gi,t} \mid h_{gi,t})} \right],
\end{aligned}
\tag{1}
$$

where $\mathcal{G}_{\text{train}} = \{g \in \{1,\dots,B\} \mid n_g \geq 2\}$, $a_{gi,t}$ is the t-th model-generated action token in o_{gi}, T_{gi} is the number of such tokens, and $h_{gi,t}$ contains the text, images, and tool observations available before token t. ρ_{gi} is the geometric mean of token-level importance ratios.

7 Experiments

7.1 Experimental setup

We evaluate MMBrowseComp (Li et al., 2025b), MMSearch (Jiang et al., 2025), MMSearch-Plus (Tao et al., 2026), SimpleVQA (Cheng et al., 2025), VDR-Bench (Zeng et al., 2026), LiveVQA (Fu et al., 2025), FVQA (Wang et al., 2018), and VisTarget-Bench, scored by Kimi K2.6. Unless noted, each configuration is run three times and we report the mean and standard deviation. Gemini3.6-Flash, Opus-5-High, and GPT-5.6-Sol-High are single runs because of cost; published open-source results are point estimates. Direct answering uses the released Qwen3-VL model (Bai et al., 2025).

All frontier models use WeAgent-Harness: Qwen3VL-30B-A3B (Bai et al., 2025), Qwen3.6-35B-A3B (Qwen Team, 2026b), Qwen3.8-27B (Qwen Team, 2026c), Kimi K2.6 (Moonshot AI, 2026), Qwen3.5-Plus (Qwen Team, 2026a), Gemini3.6-Flash (Google DeepMind, 2026), Opus-5-High (Anthropic, 2026), and GPT-5.6-Sol-High (OpenAI, 2026). Opus-5-High, GPT-5.6-Sol-High, Gemini3.6-Flash, and Qwen3.5-Plus are called through official APIs in thinking mode; Kimi K2.6, Qwen3VL-30B-A3B, Qwen3.6-35B-A3B, and Qwen3.8-27B are self-deployed in non-thinking mode. Lightweight agents are OpenSearch-VL-30B-A3B-Instruct (Chen et al., 2026), REDSearcher-MM-30B-A3B-Thinking (Chu et al., 2026), and HyperEyes-30B-A3B-Instruct (Li et al., 2026); we use published scores when available. Because OpenSearch-VL releases its model and code, results missing from its report are reproduced over three runs under our protocol and scored by Kimi K2.6.

7.2 Main Results

Table 1 compares direct answering, frontier models with WeAgent-Harness, lightweight search agents, and our variants. The SFT and RL gains show that both the curated data and the post-training

Model	Params	MMSearch-style Tasks: MMBrowse Comp	MMSearch	MMSearch Plus	VisTarget-Bench	Multimodal Knowledge QA: Simple VQA	VDR Bench	Live VQA	FVQA	Avg.
Direct Answer										
Qwen3-VL	30B (3B)	7.66 ±3.19	16.96 ±2.68	2.36 ±1.88	0.00 ±0.00	35.23 ±0.29	2.02 ±0.18	37.50 ±0.35	51.34 ±0.83	**19.13** ±0.47
Kimi K2.6	1T (32B)	2.53 ±2.29	32.36 ±5.27	4.72 ±1.65	1.99 ±0.68	42.09 ±0.72	4.00 ±0.23	50.15 ±0.19	37.50 ±0.83	**21.92** ±0.75
Frontier Models with WeAgent-Harness										
Gemini3.6-Flash‡	–	45.98	80.12	58.39	42.00	80.20	30.30	88.30	80.78	**63.26**
GPT-5.6-Sol-High‡	–	49.11	77.78	51.61	49.33	76.99	31.85	88.00	78.11	**62.85**
Opus-5-High‡	–	40.54	80.70	61.44	40.67	–	–	–	–	–
Kimi K2.6	1T (32B)	33.63 ±1.80	81.09 ±2.05	49.08 ±0.68	33.78 ±3.15	75.09 ±0.54	30.43 ±0.50	85.00 ±0.68	79.27 ±0.58	**58.42** ±0.66
Qwen3.5-Plus	397B (17B)	21.58 ±1.36	71.73 ±1.35	34.73 ±1.34	30.44 ±0.77	73.59 ±0.44	26.19 ±0.25	81.29 ±0.44	73.05 ±0.14	**51.58** ±0.41
Qwen3.6-35B-A3B	35B (3B)	14.44 ±0.68	66.28 ±0.89	29.40 ±0.17	36.22 ±0.38	65.79 ±0.62	15.57 ±0.15	78.20 ±0.38	66.09 ±0.16	**46.50** ±0.30
Qwen3.8-27B	27B	28.42 ±0.25	69.59 ±0.59	25.11 ±0.60	31.45 ±1.07	73.17 ±0.33	19.15 ±0.62	85.18 ±0.26	70.43 ±0.60	**50.31** ±0.41
Lightweight Search Agents										
OpenSearch-VL-30B	30B (3B)	14.14 ±0.93*	68.70	28.94 ±0.85*	16.89 ±1.02*	74.90	33.50	67.40	73.20	**47.21**†
REDSearcher-MM-30B	30B (3B)	23.50	72.90	26.60	–	–	–	79.30	–	–
HyperEyes-30B	30B (3B)	–	86.90	31.50	–	–	–	81.60	79.30	–
Ours										
WeAgent-MMSearch-Base	30B (3B)	5.80 ±2.79	52.24 ±3.33	12.98 ±0.83	8.00 ±3.53	70.86 ±0.84	18.65 ±0.80	62.21 ±0.34	63.22 ±0.83	**36.75** ±1.01
WeAgent-MMSearch-SFT	30B (3B)	21.87 ±2.04	74.85 ±1.55	31.00 ±0.68	24.00 ±1.76	74.40 ±0.24	28.93 ±0.40	80.61 ±0.28	72.63 ±0.50	**51.04** ±0.28
WeAgent-MMSearch-RL	30B (3B)	28.13 ±0.73	77.78 ±0.48	37.82 ±0.26	30.22 ±0.32	77.73 ±0.18	35.68 ±0.25	81.05 ±0.06	79.33 ±0.27	**55.97** ±0.12

Table 1: Main comparison on seven public benchmarks and VisTarget-Bench. Entries are means ± standard deviations over three runs unless marked ‡ (single run). A dash is unavailable; * is our three-run reproduction; † is an unweighted average of the reported means. Published baselines are point estimates. For MoE models, Params lists total parameters, with activated parameters in parentheses when disclosed.

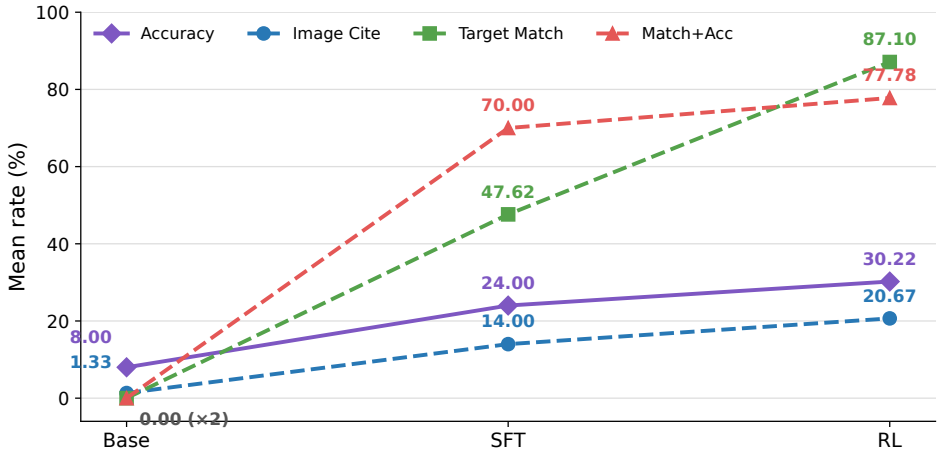

Figure 6: Final-answer and stage-wise visual evidence on VisTarget-Bench, all with WeAgent-Harness. Accuracy and Image Cite cover every response; Target Match is conditioned on a citation, and Match+Acc on a target match. Points are means over three runs.

method help. WeAgent-MMSearch-RL matches or exceeds similarly sized open-source search agents. Some frontier models still lead overall, but post-training narrows the gap on several benchmarks.

VisTarget-Bench results. The no-tool diagnostic suggests that VisTarget-Bench substantially reduces the probability of answering correctly from parametric knowledge alone: Qwen3-VL obtains 0% and Kimi K2.6 obtains $1.99 \pm 0.68\%$. Figure 6 compares final-answer accuracy, image citation, target-image acquisition, and answer correctness after a match.

Post-training improves all four metrics in Figure 6: final-answer accuracy rises from 8.00% to 24.00% and 30.22% (Base, SFT, RL), and image citation from 1.33% to 14.00% and 20.67%. Given a citation, target matching reaches 47.62% (SFT) and 87.10% (RL); given a match, answer accuracy reaches 70.00% and 77.78%. RL therefore improves both acquisition of the target image and its use in the final answer.

Configuration	MMSearch-style Tasks				Multimodal Knowledge QA				Avg.
	MMBrowse Comp	MMSearch	MMSearch Plus	VISTARGET-BENCH	Simple VQA	VDR Bench	Live VQA	FVQA	
WeAgent-MMSearch-RL without image re-feed	13.69 ±0.76	71.54 ±0.73	23.26 ±0.66	10.44 ±1.13	73.04 ±0.53	29.97 ±0.27	79.43 ±0.21	73.74 ±0.23	**46.89** ±0.27
WeAgent-MMSearch-RL with image re-feed	28.13 ±0.73	77.78 ±0.48	37.82 ±0.26	30.22 ±0.32	77.73 ±0.18	35.68 ±0.25	81.05 ±0.06	79.33 ±0.27	**55.97** ±0.12

Table 2: Image re-feed ablation for WeAgent-MMSearch-RL (mean ± std. over three runs). "No image re-feed" keeps the WEAGENT-HARNESS interface but drops returned images from later turns.

7.3 ANALYSIS EXPERIMENTS

Image re-feed. Removing re-fed images under the same WEAGENT-HARNESS interface drops WeAgent-MMSearch-RL from 55.97% to 46.89% (Table 2): benchmarks that depend on pixels returned during search lose evidence that text search cannot recover. The same mechanism explains the harness comparison. Against Hermes (Nous Research, 2026), WEAGENT-HARNESS raises the eight-benchmark average from 55.09% to 58.42% for Kimi K2.6, from 60.95% to 63.26% for Gemini3.6-Flash, and from 49.63% to 51.58% for Qwen3.5-Plus, with the largest Kimi gains on MMBrowseComp and VISTARGET-BENCH. The advantage comes from richer multimodal tool observations and from re-feeding retrieved images to later turns. Per-benchmark scores are in Appendix C.1.

Configuration	MMSearch-style Tasks				Multimodal Knowledge QA				Avg.
	MMBrowse Comp	MMSearch	MMSearch Plus	VISTARGET-BENCH	Simple VQA	VDR Bench	Live VQA	FVQA	
SFT Data Ablation									
Raw Open-Source SFT	18.00 ±1.36	67.25 ±1.75	25.94 ±1.13	9.56 ±2.69	66.93 ±0.32	26.12 ±0.33	71.61 ±0.12	67.15 ±0.47	**44.07** ±1.00
Normalized Open-Source SFT	20.83 ±1.12	71.73 ±1.47	28.31 ±1.15	14.45 ±4.34	73.12 ±0.55	28.63 ±0.33	74.77 ±0.06	72.67 ±0.11	**48.06** ±1.11
Full SFT	21.87 ±2.04	74.85 ±1.55	31.00 ±0.68	24.00 ±1.76	74.40 ±0.24	28.93 ±0.40	80.61 ±0.28	72.63 ±0.50	**51.04** ±0.28
RL Optimization Ablation									
Vanilla GSPO	21.58 ±1.57	74.85 ±2.68	36.81 ±1.13	9.78 ±5.00	75.43 ±0.52	32.35 ±0.18	80.31 ±0.22	77.67 ±0.39	**51.10** ±1.41
Failure-Aware GSPO	28.13 ±0.73	77.78 ±0.48	37.82 ±0.26	30.22 ±0.32	77.73 ±0.18	35.68 ±0.25	81.05 ±0.06	79.33 ±0.27	**55.97** ±0.12

Table 3: Data and optimization ablations (mean ± std. over three runs). Raw Open-Source SFT keeps original formats; Normalized Open-Source SFT uses the unified format; Full SFT also adds harness-native data. Vanilla GSPO omits trajectory self-recovery and invalid-trajectory filtering.

Data and optimization ablations. On the same open-source trajectories, format normalization raises the average from 44.07% to 48.06%, and adding harness-native data further raises it to 51.04%, with VISTARGET-BENCH moving from 14.45% to 24.00%. A consistent trajectory grammar and demonstrations from the same tool distribution as WEAGENT-HARNESS both matter for agentic SFT. FA-GSPO, which adds self-recovery and invalid-trajectory filtering to vanilla GSPO, improves the average from 51.10% to 55.97% and VISTARGET-BENCH from 9.78% to 30.22%. Trigger and retention rates are in Appendix C.2.

8 CONCLUSION

WeAgent-MMSearch jointly addresses persistent visual evidence, runtime recovery, environment-aligned data construction, and agentic post-training. WEAGENT-HARNESS keeps tool-retrieved images as addressable state, re-feeds them across turns, and recovers local failures before they invalidate a trajectory. Under this contract we discover seeds, build backward multi-hop tasks, collect expert trajectories, and post-train with agentic SFT and FA-GSPO. VISTARGET-BENCH pairs each question with a held-out target image, separating image retrieval from visual interpretation. WeAgent-MMSearch-RL reaches an eight-benchmark average of 55.97, a 19.22-point gain over WeAgent-MMSearch-Base, outperforming similarly sized open-source models and remaining competitive with models of roughly ten times its parameter count.

9 LIMITATIONS

Live results may time out or change, so rollout answers can differ and retain some randomness. Feeding every image can overflow the context; future work will filter them in advance.

AI USE STATEMENT

We used Kimi K2.6 to synthesize training tasks and expert trajectories and to audit RL labels (Section 4), and as the evaluation judge. VISTARGET-BENCH tasks were produced by the same pipeline and then verified by human reviewers. Generative AI was otherwise used only for minor language polishing. The authors reviewed all AI-assisted work and take responsibility for the final content.

REFERENCES

Anthropic. Claude opus 5 system card. https://www.anthropic.com/research/claude-opus-5, 2026.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025.

Shuang Chen, Kaituo Feng, Hangting Chen, Wenxuan Huang, Dasen Dai, Quanxin Shou, Yunlong Lin, Xiangyu Yue, Shenghua Gao, and Tianyu Pang. Opensearch-vl: An open recipe for frontier multimodal search agents. *arXiv preprint arXiv:2605.05185*, 2026.

Xianfu Cheng, Wei Zhang, Shiwei Zhang, Jian Yang, Xiangyuan Guan, Xianjie Wu, Xiang Li, Ge Zhang, Jiaheng Liu, Yuying Mai, et al. Simplevqa: Multimodal factuality evaluation for multimodal large language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4637–4646, 2025.

Zheng Chu, Xiao Wang, Jack Hong, Huiming Fan, Yuqi Huang, Yue Yang, Guohai Xu, Chenxiao Zhao, Cheng Xiang, Shengchao Hu, Dongdong Kuang, Ming Liu, Bing Qin, and Xing Yu. Redsearcher: A scalable and cost-efficient framework for long-horizon search agents. *arXiv preprint arXiv:2602.14234*, 2026.

Mingyang Fu, Yuyang Peng, Benlin Liu, Yao Wan, and Dongping Chen. Livevqa: Live visual knowledge seeking. *arXiv preprint arXiv:2504.05288*, 2025.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, et al. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

Xinyu Geng, Peng Xia, Zhen Zhang, Xinyu Wang, Qiuchen Wang, Ruixue Ding, Chenxi Wang, Jialong Wu, Yida Zhao, Kuan Li, et al. Webwatcher: Breaking new frontiers of vision-language deep research agent. *arXiv preprint arXiv:2508.05748*, 2025.

Google DeepMind. Gemini 3.6 flash model card. https://deepmind.google/models/model-cards/gemini-3-6-flash/, 2026.

Ziniu Hu, Ahmet Iscen, Chen Sun, Kai-Wei Chang, Yizhou Sun, David A. Ross, Cordelia Schmid, and Alireza Fathi. Avis: Autonomous visual information seeking with large language model agent. In *Advances in Neural Information Processing Systems*, 2023.

Wenxuan Huang, Yu Zeng, Qiuchen Wang, Zhen Fang, Shaosheng Cao, Zheng Chu, Qingyu Yin, Shuang Chen, Zhenfei Yin, Lin Chen, et al. Vision-deepresearch: Incentivizing deepresearch capability in multimodal large language models. *arXiv preprint arXiv:2601.22060*, 2026.

Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanmin Wu, Jiayi Lei, Pengshuo Qiu, Pan Lu, Zehui Chen, Guanglu Song, Peng Gao, et al. Mmsearch: Unveiling the potential of large models as multi-modal search engines. In *International Conference on Learning Representations*, 2025.

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

Guankai Li, Jiabin Chen, Yi Xu, Xichen Zhang, and Yuan Lu. Hypereyes: Dual-grained efficiency-aware reinforcement learning for parallel multimodal search agents. *arXiv preprint arXiv:2605.07177*, 2026.

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. Websailor: Navigating super-human reasoning for web agent. *arXiv preprint arXiv:2507.02592*, 2025a.

Shilong Li, Xingyuan Bu, Wenjie Wang, Jiaheng Liu, Jun Dong, et al. Mm-browsecomp: A comprehensive benchmark for multimodal browsing agents. *arXiv preprint arXiv:2508.13186*, 2025b.

Moonshot AI. Kimi k2.6 model card. `https://huggingface.co/moonshotai/Kimi-K2.6`, 2026.

Kartik Narayan, Yang Xu, Tian Cao, Kavya Nerella, Vishal M. Patel, Navid Shiee, Peter Grasch, Chao Jia, Yinfei Yang, and Zhe Gan. Deepmmsearch-r1: Empowering multimodal llms in multimodal web search. *arXiv preprint arXiv:2510.12801*, 2025.

Nous Research. Hermes agent. `https://github.com/NousResearch/hermes-agent`, 2026.

OpenAI. Gpt-4o system card, 2024.

OpenAI. Gpt-5.6 sol model card. `https://developers.openai.com/api/docs/models/gpt-5.6-sol`, 2026.

Qwen Team. Qwen3.5-plus model card. `https://huggingface.co/Qwen/Qwen3.5-397B-A17B`, 2026a.

Qwen Team. Qwen3.6-35b-a3b model card. `https://huggingface.co/Qwen/Qwen3.6-35B-A3B`, 2026b.

Qwen Team. Qwen3.8-27b model card. `https://huggingface.co/Qwen/Qwen3.8-27B`, 2026c.

Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.

Xijia Tao, Yihua Teng, Xinxing Su, Xinyu Fu, Jihao Wu, Chaofan Tao, Ziru Liu, Haoli Bai, Rui Liu, and Lingpeng Kong. Mmsearch-plus: Benchmarking provenance-aware search for multimodal browsing agents. In *International Conference on Learning Representations*, 2026.

Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Fvqa: Fact-based visual question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40 (10):2413–2427, 2018.

Wikimedia Foundation. Wikipedia. `https://www.wikipedia.org`, 2026.

Jinming Wu, Zihao Deng, Wei Li, Yiding Liu, Bo You, Bo Li, Zejun Ma, and Ziwei Liu. Mmsearch-r1: Incentivizing lmms to search. *arXiv preprint arXiv:2506.20670*, 2025.

Shilin Yan, Jintao Tong, Hongwei Xue, Xiaojun Tang, Yangyang Wang, Kunyu Shi, Guannan Zhang, Ruixuan Li, and Yixiong Zou. Act wisely: Cultivating meta-cognitive tool use in agentic multimodal models. *arXiv preprint arXiv:2604.08545*, 2026.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2022.

Yu Zeng, Wenxuan Huang, Zhen Fang, Shuang Chen, Yufan Shen, Yishuo Cai, Xiaoman Wang, Zhenfei Yin, Lin Chen, Zehui Chen, et al. Vision-deepresearch benchmark: Rethinking visual and textual search for multimodal large language models. *arXiv preprint arXiv:2602.02185*, 2026.

Yifan Zhang, Liang Hu, Haofeng Sun, Peiyu Wang, Yichen Wei, Shukang Yin, Jiangbo Pei, Wei Shen, Peng Xia, Yi Peng, Tianyidan Xie, Eric Li, Yang Liu, Xuchen Song, and Yahui Zhou. Skywork-r1v4: Toward agentic multimodal intelligence through interleaved thinking with images and deepresearch. *arXiv preprint arXiv:2512.02395*, 2025.

Zhixin Zhang, Yiyuan Zhang, Xiaohan Ding, and Xiangyu Yue. Vision search assistant: Empower vision-language models as multimodal search engines. *arXiv preprint arXiv:2410.21220*, 2024.

Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. Group sequence policy optimization. *arXiv preprint arXiv:2507.18071*, 2025.

Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. Sglang: Efficient execution of structured language model programs. *arXiv preprint arXiv:2312.07104*, 2024.

A Data Construction Details

This appendix records the stage procedures summarized in the main paper. The terminal visual question-answering (VQA) item is derived from a Stage 1 visually grounded seed and supplies a terminal image and pixel-dependent answer rather than the final task itself. Earlier hops are generated backward from this endpoint, and only chains that remain visually dependent and executable through WeAgent-Harness are retained. The accepted tasks are then sampled as alternating reasoning, tool-call, and multimodal-observation turns under the same evidence contract used for learning.

A.1 Stage 1: Visually Grounded Seed Generation

We first construct a seed dataset of visually grounded facts. The expert MMSearch agent searches time-sensitive topics across sports/games, science/technology/nature, finance, business/brands/organizations, among others. For each topic, construction requests advance through date windows at five-day intervals between January 1 and July 30, 2026 (e.g., January 1–5 and January 5–10), which limits repeated discovery of the same events. For each retained event, the collector stores an authoritative source page, event date, "as-of" timestamp, localized target image, source provenance, and a concise fact that must be read from that image. For the complementary stable-fact branch, we download the publicly available Wikipedia data (Wikimedia Foundation, 2026), mine each record's concept name, definition, and category, and insert only these fields into the construction prompt. The live-event branch allows the seed pool to reflect information unavailable in a fixed corpus.

A seed is accepted only if its target image is reachable through WeAgent-Harness, its answer depends on visible pixels rather than a title, URL, caption, or page metadata, and its provenance can be audited. The search results collected in this stage also populate the shared cache pool described in Section 3.2. Consequently, the seed dataset contains both the terminal visual evidence needed for later construction and an executable route by which an agent can recover it.

A.2 Stage 2: Backward Multi-Hop Expansion

Starting from each verified seed, the model expands the terminal visual fact into a complex search problem. The resulting multi-hop reasoning process follows the latent dependency chain $(n_0 \longrightarrow \cdots \longrightarrow n_{d-1} \longrightarrow n_d)$, where $d \in \{2, 3, 4\}$ is the number of retrieval hops, n_d is the verified VQA endpoint from Stage 1, and n_0 is the visual or textual anchor shown to the user. The endpoint contains target image C and a short answer D that must be read from its pixels. For every intermediate node n_i, where $1 \leq i < d$, the answer identifies the entity required to reach node n_{i+1}.

We construct this chain backward. Each agent invocation receives only the current node n_i: the expert MMSearch agent searches the open web for related evidence and writes one QA pair whose answer resolves to that entity, producing the immediate predecessor n_{i-1}. A controller repeats this single-step procedure d times to obtain the complete multi-hop retrieval chain. Solving the composed task therefore requires following the bridge relations, retrieving target image C, and inspecting its pixels to produce D.

A.3 Stage 3: Task Validation

After multi-hop expansion, the expert MMSearch agent performs semantic verification of each candidate. Structurally, we require complete anchors, bridge relations, endpoints, and provenance; at the retrieval level, we filter out records with a missing target image, a non-unique bridge, or an answer exposed by the URL, metadata, or question text. The agent further checks that the reasoning chain is internally consistent and that the answer genuinely depends on target image C, then synthesizes the final user-facing query from the validated chain. Each accepted record retains query q, answer D, the latent reasoning chain, target image C, and audit evidence.

A.4 Stage 4: Trajectory Sampling

The expert MMSearch agent next solves the accepted tasks under the same WeAgent-Harness contract. Since retrieved images are re-fed into subsequent turns, a sampled execution realizes the

visual dependencies imposed during construction. We represent a solver trajectory as

$$\tau = \big(q, (a_0, o_0), \ldots, (a_T, o_T), y\big), \tag{2}$$

where q is the constructed query, T indexes the final interaction turn, a_t and o_t are the action and observation at turn t, and y is the terminal answer. Observation o_t may contain text and newly retrieved pixels. The corresponding episode bundle preserves tool calls, observations, visual evidence, and terminal outputs. Thus successful demonstrations are executable records of the target environment.

The verified endpoint answer determines whether y is correct. Correct and complete trajectories form the SFT set. Failed or incomplete executions are excluded from behavior cloning, but the underlying tasks remain useful for RL. We construct the RL query set by stratifying model-solved and model-unsolved tasks, balancing demonstrated feasibility against headroom for policy improvement. Benchmark candidates are held out at the seed and task levels before either split is used for optimization.

B Open-Source Normalization and Final Mixtures

The four stages are followed by a separate mixture-construction step. The resulting trajectories and prompts are complemented by open-source SFT and RL data from Metis (Yan et al., 2026), OpenSearch-VL (Chen et al., 2026), and REDSearcher (Chu et al., 2026). These sources expose heterogeneous tool names, action schemas, and answer conventions. Directly merging them would make the same serialized action carry source-dependent semantics. We therefore convert all three sources into a common contract before mixing them with the in-house data.

SFT normalization. We first remove records containing corrupted or garbled text, trajectories in which any individual model response exceeds 4,096 output tokens, and trajectories whose total serialized length exceeds 48K tokens. We then place every tool in a source-specific namespace. For example, tools imported from REDSearcher receive the prefix `RED_`; analogous prefixes are assigned to the other sources. This avoids collisions when two datasets use the same tool name for different interfaces and gives the policy an unambiguous action vocabulary. Finally, every assistant turn is rewritten into one of two canonical forms: an intermediate action contains `<think>...</think><tool_call>...</tool_call>`, whereas a terminal turn contains `<think>...</think><response>...</response>`. Consequently, source identity changes the available tool namespace but not the trajectory grammar learned by the SFT policy.

RL query filtering. For every candidate query, we sample an 8-rollout group from WeAgent-MMSearch-SFT and discard queries for which all sampled answers are correct or all are incorrect. Uniformly successful queries provide little headroom, while uniformly failed queries provide no within-group reward contrast; both are uninformative for group-relative policy optimization. We additionally use Kimi K2.6 (Moonshot AI, 2026) to audit the reference labels and remove ambiguous or non-verifiable targets. This step filters, for example, some open-source records whose nominal label is a refusal placeholder such as "I cannot find the final answer." The retained RL prompts therefore combine non-trivial outcome variance with a concrete, judgeable target.

Dataset scale and composition. After normalization and filtering, the SFT mixture contains 59.5K trajectories: 3.3K in-house trajectories (5.5%) and 56.2K open-source trajectories (94.5%). The RL mixture contains 13.3K prompts: 4.9K in-house prompts (37.2%) and 8.3K open-source prompts (62.8%).

C Additional Ablations

C.1 Harness Ablation

Across the compared models, WeAgent-Harness is stronger than Hermes (Nous Research, 2026). For Kimi K2.6 the eight-benchmark average rises from 55.09% to 58.42%, with the largest gains on MMBrowseComp and VisTarget-Bench; for Gemini3.6-Flash it rises from 60.95% to 63.26%, and for Qwen3.5-Plus it rises from 49.63% to 51.58%. The advantage comes from richer multimodal tool observations and from re-feeding retrieved images to later turns.

Configuration	MMSearch-style Tasks				Multimodal Knowledge QA				Avg.
	MMBrowse Comp	MMSearch	MMSearch Plus	VISTARGET-BENCH	Simple VQA	VDR Bench	Live VQA	FVQA	
Kimi K2.6 + Hermes	23.39 ±1.16	76.22 ±4.11	50.96 ±1.54	20.01 ±3.52	76.43 ±0.40	29.83 ±0.55	84.47 ±0.55	79.43 ±0.42	**55.09** ±0.18
Kimi K2.6 + WEAGENT-HARNESS	33.63 ±1.80	81.09 ±2.05	49.08 ±0.68	33.78 ±3.15	75.09 ±0.54	30.43 ±0.50	85.00 ±0.68	79.27 ±0.58	**58.42** ±0.66
Gemini3.6-Flash‡ + Hermes	28.57	82.46	59.68	38.00	80.05	29.90	87.67	81.28	**60.95**
Gemini3.6-Flash‡ + WEAGENT-HARNESS	45.98	80.12	58.39	42.00	80.20	30.30	88.30	80.78	**63.26**
Qwen3.5-Plus + Hermes	22.02 ±0.68	63.35 ±0.89	37.34 ±0.59	21.55 ±0.39	72.63 ±0.47	26.47 ±0.65	82.43 ±0.15	71.25 ±0.23	**49.63** ±0.26
Qwen3.5-Plus + WEAGENT-HARNESS	21.58 ±1.36	71.73 ±1.35	34.73 ±1.34	30.44 ±0.77	73.59 ±0.44	26.19 ±0.25	81.29 ±0.44	73.05 ±0.14	**51.58** ±0.41

Table 4: Harness-choice ablation for frontier models. Values are means ± standard deviations over three runs unless marked with ‡, which denotes a single run owing to high inference cost.

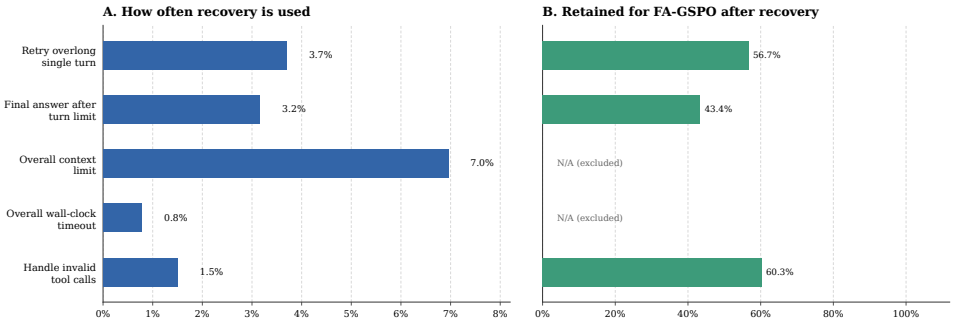

Figure 7: Harness recovery on one RL run. **Left:** how often each mechanism is triggered. **Right:** conditional on a trigger, the share retained for FA-GSPO after recovery and runtime filtering. Overall context-limit and wall-clock-timeout trajectories are N/A (excluded), because they are filtered before this outcome is evaluated.

C.2 TRAJECTORY RECOVERY

Figure 7 reports recovery percentages from one RL training run. For each mechanism, the trigger rate is its fraction of all rollout trajectories. Conditional on a trigger, the retention rate is the share of trajectories that remain in the FA-GSPO training set after recovery and runtime filtering. Retrying an overlong single-turn output, forcing a final answer after the regular turn limit, and returning invalid tool-call errors are counted separately. We additionally report overall context-limit and wall-clock-timeout triggers. These two trajectory-level violations are directly excluded and are therefore reported as N/A (excluded). A trajectory may activate more than one local mechanism, so trigger rates are not additive.

The recovery mechanisms retain a portion of locally abnormal executions for optimization, avoiding the waste of otherwise usable rollout computation. Directly excluding overall context-limit and wall-clock-timeout trajectories prevents their length- and timeout-related artifacts from introducing noise into policy optimization.

D IMPLEMENTATION DETAILS

WeAgent-MMSearch-SFT updates all language-model and visual modules (vision tower and projector) with a global batch size of 32 trajectories and learning rates of 1.0×10^{-5} for the language model and 1.0×10^{-6} for the vision tower. WeAgent-MMSearch-RL initializes from this SFT checkpoint, freezes the vision tower and projector, and trains only the language-model policy; it uses a prompt batch of 64 with 8 samples per prompt (global trajectory batch 512), a constant learning rate of 5.0×10^{-7}, and rollout temperature 1.0.

E TOOL INTERFACES AND RUNTIME CONFIGURATION

The reported system exposes five native tools. A tool-returned image is saved under a stable relative filename, registered in the episode evidence state, and attached as visual input to the next model turn.

Table 5 defines the model-visible interfaces, and Table 6 gives the episode budgets, recovery rules, and persistence policy.

Tool	Arguments	Operation and returned evidence
`web_search`	`query`: string or up to 5 strings; `num`: 1–5 (default 3)	General web search. Returns titles, URLs, snippets, and publication times; query language selects the search locale automatically.
`web_search_for_image`	`query`: string or up to 3 strings; `num`: 1–10 (default 5)	Text-to-image search. Localizes successful thumbnails, returns their titles/domains/relative paths, and directly re-feeds up to four images to the VLM.
`web_search_based_on_image`	`image`: a relative filename registered in the current rollout, or a list of up to 5 such filenames; `num`: 1–8 (default 5)	Reverse-image search. Returns related pages and an optional knowledge-graph match; result thumbnails are localized and up to five are re-fed at no more than 512 vision tokens each.
`web_extract`	`urls`: up to 5 URLs; optional `goal`	Fetches HTML or PDF content and returns a goal-conditioned Markdown summary per URL. Pages are summarized by a LLM to at most 5,000 characters; up to five localized page images are also re-fed when available.
`execute_code`	`code`: Python source	Runs Python in a persistent episode sandbox. Stdout/stderr are returned; newly generated raster images are registered and the first five are re-fed.

Table 5: Model-visible tool interfaces. Limits are the active values in the reported RL recipe.

F RL PROMPTS

Table 7 gives the system and user prompts used for all RL rollouts and evaluations. At runtime, `CURRENT_TIME` is instantiated at hour granularity, while `TOOL_LIST` is populated with the function schemas.

Category	Setting	Value / Behavior
Episode	Turn / executed-call limits	20 regular turns (plus at most one no-tool synthesis turn) / 32 calls
	Total context / wall-clock limit	32,768 tokens / 900 seconds; trajectories crossing either limit are excluded from training
	Early final synthesis	Tool use is disabled once at most 4,096 context tokens remain; one no-tool answer turn is requested
Recovery	Argument repair	Deterministically repairs only unambiguous aliases, scalar-to-list conversions, and ignored extra fields; it never infers missing information
	Duplicate suppression	A call with the same tool name and canonicalized arguments as the immediately preceding executed call is not executed and instead receives a corrective observation
	Repeated format errors	After two consecutive identical invalid calls, a loop-guard observation instructs the model to change the call or answer with available evidence
	Overlong response	A single response exceeding the output-length limit is discarded and one replacement response is regenerated
	Max-turn finalization	After 20 regular turns, tools are disabled and one no-tool synthesis turn is allowed; it is retained only if it contains a complete terminal response
	Excluded trajectories	A trajectory hitting the context-length or wall-clock limit is excluded from training, irrespective of any subsequently generated answer
Persistence	General result threshold	Results above 16,384 characters are persisted with a 1,500-character in-context preview
	Web extraction	Length is controlled independently per URL; at most five 5,000-character summaries are returned

Table 6: Harness budgets, recovery behavior, and system-level tool settings.

Prompt Type	Prompt

System Prompt	You are a helpful, knowledgeable assistant with strong multimodal understanding. Use the available tools to solve the user's task. Be targeted and efficient.
	
	`[Current date and time (Beijing Time, UTC+8): CURRENT_TIME]`
	`[Available tools and function schemas: TOOL_LIST]`
	
	Rules.
	1. Always use relative file paths in tool arguments. Do not hardcode absolute paths.
	2. Each turn: start with a short `<think>...</think>` in the SAME language as the user's prompt, then output EXACTLY ONE of the following: (i) a single tool call via the function-calling protocol—do NOT write `<tool_call>` JSON inside the text content, and call at most one tool per turn; or (ii) a final answer wrapped in `<response>...</response>`. Never emit a tool call, either textual or structured, after `<response>`.
	3. Uploaded images are shown to you directly AND are available as files (`user_image_1.png`, `user_image_2.png`, ...); images returned by tools (e.g., search thumbnails labeled `[Image N]`) are likewise already visible. Do NOT call `execute_code` just to open or display an image you can already see—only run code on an image to crop or zoom a small region or to measure or compute something you cannot judge by eye. Pass relative paths to tools.
	4. If a tool call fails or returns nothing useful, switch tactics by using different arguments or a different tool. Retry the same kind of action at most twice, then answer with the best available evidence.
	5. Prefer a final reply that mixes text with supporting images when image evidence exists. Inside `<response>...</response>`, if a claim is backed by a saved image, embed that image inline with `<IMG: filename>` (e.g., `<IMG: search_*.jpg>`). Place each image next to the claim it supports. Cite only the images that actually corroborate the answer—do not dump every saved image.

User Prompt	`[task instruction]`
	
	When input images are present, append their registered relative paths:
	`Registered image paths: - user_image_1: [filename_1], ..., - user_image_N: [filename_N].`

Table 7: System and user prompts for RL rollouts and evaluations. Bracketed fields are replaced by the corresponding runtime time, tool schemas, task, image identifier, or filename.

G FINAL-RESPONSE COMPARISON

This section compares the *final reply* of WeAgent-MMSearch-RL and Kimi K2.6 with Hermes on three samples. On these samples, WeAgent-MMSearch-RL is correct, and Kimi K2.6 with Hermes is not. Images cited in a reply are shown in that column.

Case 1 Membrane color in a Class II cytokine-receptor schematic

Query. This scholar has long taught immunology at the Catholic University of Leuven in Belgium, is also affiliated with an international cancer research institution, and is known for research on cytokines and cancer immunology. In an article the scholar wrote for a leading immunology review journal in 2003, a schematic systematically depicting multiple members of receptor families involved in antiviral and inflammatory regulation and the structures of their transmembrane complexes shows the portions of the receptor complexes embedded in the cell membrane mainly in what color?

Target image(not provided to the model at query time)

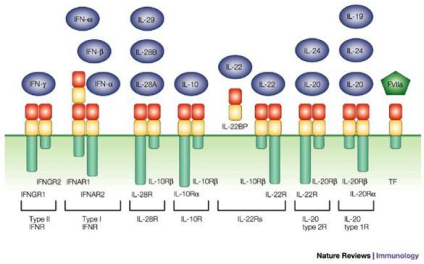

Kimi K2.6 + Hermes — Incorrect

The membrane-embedded portions are shown in **gray**.

WeAgent-MMSearch-RL — Correct

Based on my research, the scholar is **Jean-Christophe Renauld**, a professor at the Catholic University of Leuven (KU Leuven) in Belgium and a researcher at the Ludwig Institute for Cancer Research. In 2003, he published a review article titled "Class II cytokine receptors and their ligands: key antiviral and inflammatory modulators" in *Nature Reviews Immunology*.

The schematic systematically depicts multiple members of Class II cytokine receptor families involved in antiviral and inflammatory regulation, together with the structures of their transmembrane complexes. In this figure, the portions of the receptor complexes embedded in the cell membrane are shown in **green** (or a greenish/teal color).

The figure includes receptors such as IFNGR2/IFNAR1, IFNAR2, IL-28RA, IL-10Rβ, IL-10Rα, IL-22R, IL-20Rβ, IL-20Rα, and TF. Extracellular domains are drawn in blue/purple, whereas the transmembrane segments that cross the membrane bar are colored green.

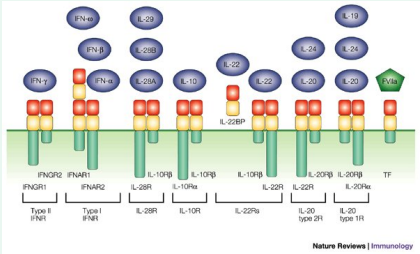

Retrieved receptor schematic

Case 2 Next Chinese equivalent after *ma* in a Shui-script chart

Query. In the Wanhu Shuizhai community on the outskirts of the county seat of China's only Shui autonomous county, there is a museum that is the country's first specialized exhibition space devoted to the culture of this ethnic minority. On its first floor, a migration-route map and classical texts show the development of an ancient script. A reference page for this script on the Omniglot website contains a comparison chart. What is the next character after *ma* ('horse') in the chart's third column?

Target image(not provided to the model at query time)

Sui character	Pronunciation	Chinese equivalent(s)m	Meaning(s)
开	[ʔbən^{13}; thjen^{13}]	天 (tiān)	sky
𠂤	[mi^{55}; fən^{13}; jaŋ^{31}]	未 (wèi)	not
𠂤	[ma^{52}]	马 (mǎ)	horse
戈	[m̥a^{13}; thon^{33}]	犬 (quǎn)	dog
𠂤	[mai^{52}]	木 (mù)	tree, wood
𠂤	[mit^{32}]	刀 (dāo)	knife
𠂤	[ma^{52}]	牛 (niú)	cow, beef

Kimi K2.6 + Hermes — Incorrect

The next character is ***yang*** ('sheep').

WeAgent-MMSearch-RL — Correct

Based on my investigation, the museum is the **Guizhou Shui Culture Museum**, located in Wanhu Shuizhai community, Fengyu Subdistrict, Sandu Shui Autonomous County, Qiannan Buyei and Miao Autonomous Prefecture, Guizhou Province. It is China's first museum devoted specifically to Shui culture.

The first floor displays a Shui migration-route map and classical texts that present the history of the ancient Shui script (*Sui* script).

The Omniglot reference page contains a comparison chart. In the third column (Chinese equivalents) the sequence is *tian* (sky), *wei* (not), *ma* (horse), ***quan*** **(dog)**, *mu* (tree/wood), *dao* (knife), *niu* (cow).

Therefore, the next character after *ma* is ***quan*** **('dog')**.

Sui character	Pronunciation	Chinese equivalent(s)m	Meaning(s)
开	[ʔbən^{13}; thjen^{13}]	天 (tiān)	sky
𠂤	[mi^{55}; fən^{13}; jaŋ^{31}]	未 (wèi)	not
𠂤	[ma^{52}]	马 (mǎ)	horse
戈	[m̥a^{13}; thon^{33}]	犬 (quǎn)	dog
𠂤	[mai^{52}]	木 (mù)	tree, wood
𠂤	[mit^{32}]	刀 (dāo)	knife
𠂤	[ma^{52}]	牛 (niú)	cow, beef

Omniglot comparison chart

Case 3 Vertical text and rate letter on a 2015 Canada Picture Postage

Query. A North American country and an East Asian country jointly issued a se-tenant postal item to commemorate the 35th anniversary of their diplomatic relations; it featured two kinds of big cat drawn separately by designers from the two countries. Exactly ten years after this se-tenant postal item featuring the two big cats was issued, the animal on the right was again selected by that country's postal service as a postage design, this time shown as an actual photograph on personalized postage. What fixed text is printed vertically along the left edge of the postage? What letter appears in the postage symbol in the lower right?

Target image(not provided to the model at query time)



Kimi K2.6 + Hermes — Incorrect

No final answer was produced. After a long search, the model hypothesized that a 2015 Canada Post Picture Postage of a cougar might use a frame with vertical "CANADA" and a Permanent "P" mark, then issued another web-search call about frame designs and stopped without boxing a response.

WeAgent-MMSearch-RL — Correct

The se-tenant item is the 2005 Canada–China joint issue commemorating the 35th anniversary of diplomatic relations, featuring a Canadian cougar (on the right) and a Chinese leopard. Exactly ten years later, in 2015, Canada Post issued a personalized Picture Postage stamp with a photograph of a cougar.

On this personalized postage stamp, the fixed text printed vertically along the left edge is **CANADA**, and the postage symbol in the lower-right corner is a maple leaf containing the letter **P** (Permanent domestic rate).



2015 Canada Picture Postage

H QUALITATIVE RL RESPONSE CASES

This section shows three WeAgent-MMSearch-RL rollouts. Tool returns are shown as a short preview: at most three text hits, and at most five images. Source URLs keep only the host and redact the path. The full payload is omitted. The target image is held out at query time.

Case 1: Cross-event athlete retrieval and suit-color reading

Query. In the same event in which the athlete shown in this photograph competed, what were the two main colors of the racing suit worn by the host nation's best-performing athlete in a news photograph showing that athlete celebrating after crossing the finish line?

Reference answer: Light blue and dark purple (blue and purple).

Input image
(provided with the query)



Target image
(not provided to the model at query time)

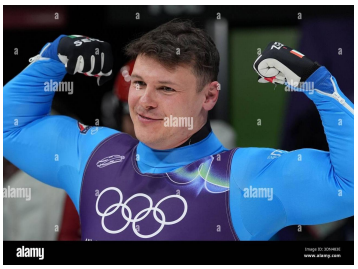

Think | Round 1

We need to answer: In the same event in which the athlete shown in this photograph competed, what were the two main colors of the racing suit worn by the host nation's best-performing athlete in a news photograph showing that athlete celebrating after crossing the finish line?
We have an image: user_image_1.png. We need to identify the athlete in the image, the event, then find the host nation's best-performing athlete in that event, and then find a news photograph of that athlete celebrating after crossing the finish line, and determine the two main colors of the racing suit worn by that athlete ...

Tool call | web_search_based_on_image | Round 1

```
{"image": "user_image_1.png", "num": 5}
```

Tool response | Round 1

Summary preview

Langenhan races to luge gold
https://www.picture-alliance.com/***

Olympia 2026: Die Winterspiele von Mailand und Cortina d'Ampezzo: Kuhne-Hellmessen, Ulrich, Ve ...
https://www.amazon.ca/***

Lieber Herr Merz - es tut mir super leid
https://www.stern.de/***





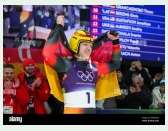

search_5785fff5.jpg search_af8aa6c4.jpg search_993575f1.jpg search_408aa347.jpg search_1a9c7c46.jpg
▷ Full tool return (2,743 characters)

Think | Round 2

The reverse image search results identify the athlete as Max Langenhan, a German luger who won gold at the Milano Cortina 2026 Winter Olympics in the men's singles luge event on February 8, 2026.
The event is the Milano Cortina 2026 Winter Olympics (host nation: Italy). The host nation's best-performing athlete in this event? Wait, the event is luge men's singles. The host nation is Italy. What was Italy's best result in men's singles luge at Milano Cortina 2026? We need to find the host nation's best-performing athlete in the same event (men's singles luge). The host nation is Italy. We need to ...

Tool call | web_search | Round 2

{"query": ["Milano Cortina 2026 luge men's singles Italy best result finish line celebration", "Italian luger Milano Cortina 2026 luge gold medal", "campioni olimpici Milano Cortina 2026 luge singolo maschile", "luge Milano Cortina 2026 Italia finish line celebrare"]}

Tool response | Round 2

Summary preview

An Olympic champion is crowned in men's singles luge ...
https://www.facebook.com/***
Olympic luge at Milan Cortina 2026: Germany dominates ...
https://www.nbcolympics.com/***
Focus on Cortina: Olympic stage for Luge in ...
https://www.fil-luge.org/***
▷ Full tool return (5,910 characters)

Think | Round 3

We need to find the host nation (Italy) best-performing athlete in the same event (men's singles luge). The search results mention Germany dominated, Italy got gold in doubles, but what about men's singles? Let's check the Olympics.com results page for luge results. Also, we need a news photograph showing that athlete celebrating after crossing the finish line. Let's extract the Olympics.com results page and the NBC article.

Tool call | web_extract | Round 3

{"urls": ["https://www.olympics.com/***", "https://www.nbcolympics.com/***"], "goal": "Find the results for men's singles luge at Milano Cortina 2026, spec ...

Tool response | Round 3

Summary preview

Milano-Cortina 2026 Luge - Olympic Results by Discipline
https://www.olympics.com/***
Olympic luge at Milan Cortina 2026: Germany dominates, Farquharson gives U.S. its medal moment
https://www.nbcolympics.com/***







extract_73f08767.jpg extract_0182f6e2.jpg extract_2a0e719b.jpg extract_420cb9c5.jpg extract_077c4bf3.jpg
▷ Full tool return (3,102 characters)

Think | Round 4

The NBC article reveals that Italy's Dominik Fischnaller won bronze in men's singles luge at Milano Cortina 2026. He was the best-performing Italian in men's singles luge. We need a news photograph showing Dominik Fischnaller celebrating after crossing the finish line, and we need to determine the two main colors of the racing suit worn by him in that photo.
Let's search for "Dominik Fischnaller Milano Cortina 2026 finish line celebration" or similar.

Tool call | web_search | Round 4

```
{"query": ["Dominik Fischnaller Milano Cortina 2026 finish
line celebration", "Dominik Fischnaller luge 2026 Italy finish
line", "Dominik Fischnaller finish line photo 2026", "Luge Dominik
Fischnaller Milano Cortina 2026"]}
```

Tool response | Round 4

Summary preview

Dominik Fischnaller races to bronze at the Winter ...
https://www.facebook.com/***
Bravo, Dominik! Our luging star from Meransen ...
https://www.instagram.com/***
Dominik Fischnaller and the long road to Olympic dreams
https://www.fil-luge.org/***
▷ Full tool return (5,641 characters)

Think | Round 5

We have a Facebook post from FILuge titled "Dominik Fischnaller races to bronze at the Winter Olympics in Milano Cortina 2026!" which might contain a photo of him celebrating after crossing the finish line. Let's extract that page. Also, we might need to search for images of Dominik Fischnaller at the finish line. Let's extract the FILuge Facebook post and maybe the Olympics.com athlete page. Also, check Getty Images or other sources. Let's start with the Facebook post.

Tool call | web_extract | Round 5

```
{"urls": ["https://www.facebook.com/***", "https://www.olympics.com/***",
"https://en.wikipedia.org/***"], "goal": "Find a news photograph
showing Dom ...
```

Tool response | Round 5

Summary preview

International Luge Federation
https://www.facebook.com/***
Dominik FISCHNALLER
https://www.olympics.com/***
Dominik Fischnaller
https://en.wikipedia.org/***



extract_2713a7e2.jpg
▷ Full tool return (3,590 characters)

Think | Round 6

The Facebook post images didn't show the actual photograph of Fischnaller celebrating at the finish line. The extracted images from the page were mostly emojis and a graphic. We need to locate the actual news photograph. Let's search for images of Dominik Fischnaller at the finish line. We can use web_search_for_image with queries like "Dominik Fischnaller finish line Milano Cortina 2026" or "Dominik Fischnaller Milano Cortina 2026 luge". Let's try.

Tool call | web_search_for_image | Round 6


```
{"query": ["Dominik Fischnaller finish line Milano Cortina 2026",
"Dominik Fischnaller luge Milano Cortina 2026 finish line", "Dominik
Fischnaller bronze medal Milano Cortina 2026 finish line"], "num":
10}
```


Tool response | Round 6

Summary preview
Tool returned a structured payload; only the preview is shown.

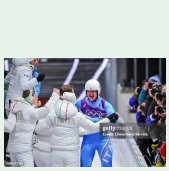





search_245c38a3.jpg search_a0b48fcc.jpg search_8be8ed1e.jpg search_d707ce80.jpg search_9c529dfd.jpg

▷ Full tool return (10,847 characters)

Think | Round 7

We have many images of Dominik Fischnaller. We need the one showing him celebrating after crossing the finish line. Let's look at the images carefully.
Image 1 (Getty Images): Shows Fischnaller in a blue and purple suit being hugged by teammates in white "ITALIA" jackets. This is likely after winning bronze, but is it at the finish line? He is being hugged by teammates who are in white jackets, not racing suits. The caption says "Italy's Dominik Fischnaller celebrates winning the bronze medal as he arrives at the finish during a men's single luge run at the 2026 Winter Olympics, in Cortina". Wait ...

Tool call | web_search | Round 7


```
{"query": ["Dominik Fischnaller suit color Milano Cortina 2026 luge",
"Dominik Fischnaller racing suit blue purple", "Italy luge suit Milano
Cortina 2026 colors"]}
```


Tool response | Round 7

Summary preview
Dominik Fischnaller (ITA) - Men's Singles Bronze - Luge
https://www.olympics.com/***
Winter Opening! #fisiintour #suit #night #milano ...
https://www.instagram.com/***
Dominik Fischnaller races to bronze at the Winter ...
https://www.facebook.com/***
▷ Full tool return (4,355 characters)

Think | Round 8

The search results don't specifically mention the suit colors beyond what we can see in the images. The images clearly show blue and purple. Let's verify by looking at the image more closely. We can use execute_code to load and inspect the image if needed, but we can see it clearly.

However, to be thorough, let's examine the Alamy image (search_a0b48fcc.jpg) or the Getty image (search_245c38a3.jpg) using code to confirm colors. But we can see them directly. The suit is predominantly blue (light blue) and purple (dark purple). The two main colors are blue and purple. Wait, the question says "two m ...

Tool call | execute_code | Round 8

```
{"code": "from PIL import Image\nimport matplotlib.pyplot as plt\n\n# Load the Getty image showing Fischnaller celebrating at finish\nimg1 = Image.open('search_245c38a3.jpg')\n# Load the Alamy image\nimg2 = Image.open('search_a0b48fcc.jpg')\n\nfig, axes = plt.subplots(1, 2, figsize=(12, 6))\naxes[0].imsho ...
```

Tool response | Round 8

Summary preview
Tool returned a structured payload; only the preview is shown.

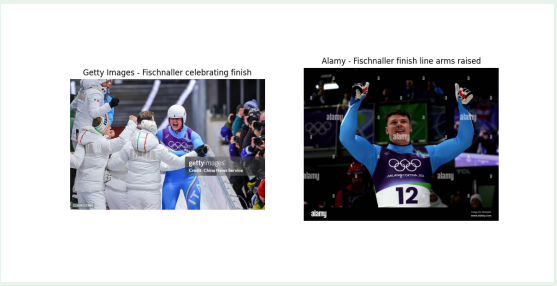

fig_001.png

▷ Full tool return (171 characters)

Think | Round 9

The images confirm the racing suit colors: blue and purple. The suit has light blue sleeves and shoulders, and a dark purple torso with Olympic rings. The two main colors are blue and purple. Now we can formulate the final answer.
But let's also verify that the host nation's best-performing athlete is indeed Fischnaller. We'll quickly extract the Olympics.com results page for men's singles luge to confirm his placement.

Tool call | web_extract | Round 9

```
{"urls": ["https://www.olympics.com/***"], "goal": "Find the medalists for men's singles luge at Milano Cortina 2026 and confirm that Italy's best performer is Dominik Fischnaller with bronze."}
```

Tool response | Round 9

Summary preview
Milano-Cortina 2026 Luge - Olympic Results by Discipline
https://www.olympics.com/***


extract_c82923a6.jpg


extract_e10f1b00.jpg


extract_009f4cf4.jpg


extract_79482084.jpg

▷ Full tool return (1,473 characters)

Think | Final

Use the verifiable evidence gathered in the trajectory and answer the question as requested.

Final response

The athlete in the photograph is **Max Langenhan** of Germany, celebrating his gold medal victory in men's singles luge at the **Milano Cortina 2026 Winter Olympics**.
The host nation of that event was Italy. Italy's best-performing athlete in men's singles luge was **Dominik Fischnaller**, who won the bronze medal after four strong runs.
In a news photograph showing him celebrating after crossing the finish line (e.g., the Getty Images photo titled "Fischnaller celebrates finish" or the Alamy photo "Cortina, Italia. 08th Feb, 2026. Italy's Dominik Fischnaller celebrates winning the bronze medal as he arrives at the finish..."), Fischnaller is wearing the standard Italian national luge racing suit.
The **two main colors of his racing suit** are:
Blue and purple – the suit features light blue sleeves and shoulders with a dark purple torso adorned with the Olympic rings and Milano Cortina 2026 branding.



search_245c38a3.jpg



search_a0b48fcc.jpg

Case 2: Scholar, article, and receptor-schematic retrieval

Query. This scholar has long taught immunology at the Catholic University of Leuven in Belgium, is also affiliated with an international cancer research institution, and is known for research on cytokines and cancer immunology. In an article the scholar wrote for a leading immunology review journal in 2003, a schematic systematically depicting multiple members of receptor families involved in antiviral and inflammatory regulation and the structures of their transmembrane complexes shows the portions of the receptor complexes embedded in the cell membrane mainly in what color?

Reference answer: green.

Target image

(not provided to the model at query time)

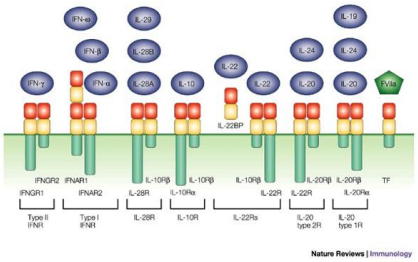

Think | Round 1

The user is asking about a specific scholar who taught immunology at the Catholic University of Leuven in Belgium, is affiliated with an international cancer research institution, and is known for research on cytokines and cancer immunology. They wrote an article for a leading immunology review journal in 2003, and in that article there's a schematic showing receptor families involved in antiviral and inflammatory regulation. The question asks what color the portions of receptor complexes embedded in the cell membrane are depicted in.
I need to identify this scholar first. Let me search for immun ...

Tool call | web_search | Round 1

```
{"query": ["Catholic University of Leuven immunology cytokines cancer immuno", "KU Leuven immunology cytokines cancer immunology scholar 2003 review article", "Leuven immunology scholar cytokines cancer international cancer research institution", "2003 immunology review journal cytokine receptor schematic"]}
```

Tool response | Round 1

Summary preview

The role of cytokines in shaping the future of Cancer ...
https://pubmed.ncbi.nlm.nih.gov/***

The "double-edged sword" effect of cytokines in cancer
https://www.frontiersin.org/***

Research Portal - Laboratory of Cell Stress & Immunity
https://research.kuleuven.be/***

▷ Full tool return (5,747 characters)

Think | Round 2

The search results didn't directly identify the scholar. Let me try a more targeted search. The scholar might be affiliated with Ludwig Cancer Research (an international cancer research institution). Let me search for KU Leuven and Ludwig Cancer Research together.

Tool call | web_search | Round 2

{"query": ["KU Leuven Ludwig Cancer Research immunology cytokines", "Leuven immunology scholar Ludwig Cancer Institute", "Ludwig Cancer Research KU Leuven cytokine receptor", "2003 Annual Review of Immunology cytokine receptor schematic"]}

Tool response | Round 2

Summary preview

The cytokine-serum amyloid A-chemokine network - PubMed
https://pubmed.ncbi.nlm.nih.gov/***

Ludwig Lausanne's
https://www.ludwigcancerresearch.org/***

Aminereactive OVA multimers for autovaccination against ...
https://jlb.onlinelibrary.wiley.com/***

▷ Full tool return (5,555 characters)

Think | Round 3

I need to identify the specific scholar. Let me search for leading immunology review journals in 2003 that might have had an article about cytokine receptors. The Annual Review of Immunology was published in 2003. Let me search for that specific issue.

Tool call | web_search | Round 3

{"query": ["Annual Review of Immunology 2003 cytokine receptors schematic", "Annual Review of Immunology 2003 receptor families antiviral inflammatory regulation", "2003 Immunology review cytokine receptor complex membrane embedded color"]}

Tool response | Round 3

Summary preview

Cytokine Receptor - an overview
https://www.sciencedirect.com/***

Mechanisms of Signal Transduction from Receptors ...
https://www.tandfonline.com/***

Cytokine Signaling Modules in Inflammatory Responses
https://www.cell.com/***

▷ Full tool return (4,272 characters)

Think | Round 4

Let me search more specifically for the 2003 Annual Review of Immunology issue and look for an article about cytokine receptor families, antiviral, and inflammatory regulation. I need to identify the author.

Tool call | web_search | Round 4

{"query": ["2003 Annual Review of Immunology volume 21 cytokine receptors", "Annual Review of Immunology 2003 \"interferon\" interleukin receptor families", "2003 Immunology review journal \"schematic\" \"receptor families\" antiviral"]}

Tool response | Round 4

Summary preview

Annual Review of Immunology - Volume 21, 2003
https://www.annualreviews.org/***

Annual review of immunology
https://www.scielo.br/***
(PDF) Annual review of immunology
https://www.researchgate.net/***
▷ Full tool return (4,330 characters)

Think | Round 5

I need to be more systematic. Let me search for the specific article in the 2003 Annual Review of Immunology about cytokine receptors and their families involved in antiviral and inflammatory regulation. The description mentions "multiple members of receptor families involved in antiviral and inflammatory regulation" and a "schematic systematically depicting" them. This sounds like it could be about the interferon receptor family or interleukin receptor family.
Let me search for 2003 Annual Review of Immunology articles specifically about cytokine receptors and receptor families.

Tool call | web_search | Round 5

```
{"query": ["2003 Annual Review of Immunology cytokine receptor
families interferon interleukin", "Annual Review of Immunology 2003
\"cytokine\" \"receptor\" family schematic", "2003 \"Annual Review of
Immunology\" \"interferon receptor\" family", "2003 \"Annual Review of
Immunology\" \"interleukin receptor\" family"]}
```

Tool response | Round 5

Summary preview
Review The interleukin-10 family of cytokines
https://www.sciencedirect.com/***
Insights into Cytokine–Receptor Interactions from ... - PMC - NIH
https://pmc.ncbi.nlm.nih.gov/***
A Life with Cytokines, Accompanied by Friends and Music
https://www.annualreviews.org/***
▷ Full tool return (5,979 characters)

Think | Round 6

The search results mention that the 2002 article by H Fickenscher in Annual Review of Immunology discusses the CRF2 (cytokine receptor family type 2) which comprises interferons. But I need the 2003 article. Let me think about this differently.
The description says: "a schematic systematically depicting multiple members of receptor families involved in antiviral and inflammatory regulation and the structures of their transmembrane complexes shows the portions of the receptor complexes embedded in the cell membrane mainly in what color?"
This sounds like a specific schematic from a review article ...

Tool call | web_search | Round 6

```
{"query": ["KU Leuven Ludwig Cancer Research immunology professor",
"Catholic University of Leuven Ludwig cancer research senior
scientist", "KU Leuven immunology Ludwig Institute for Cancer
Research"]}
```

Tool response | Round 6

Summary preview
Benoit Van den Eynde
https://www.ludwigcancerresearch.org/***
Ludwig Lausanne's Mikael Pittet
https://www.ludwigcancerresearch.org/***

Laboratory for Tumor Immunology and Immunotherapy
https://research.kuleuven.be/***
▷ Full tool return (4,098 characters)

Think | Round 7

I see Benoit Van den Eynde is mentioned as a Member of the Ludwig Institute for Cancer Research and co-director of the de Duve Institute. He's a tumor immunologist. Let me search for him specifically, as he might have published in Annual Review of Immunology in 2003.

Tool call | web_search | Round 7

{"query": ["Benoit Van den Eynde Annual Review of Immunology 2003", "Benoit Van den Eynde cytokine receptor schematic", "Benoit Van den Eynde 2003 immunology review article"]}

Tool response | Round 7

Summary preview
Tumor Antigens Recognized by T Lymphocytes
https://www.annualreviews.org/***
The tumor defense dismantler
https://www.ludwigcancerresearch.org/***
Tumour immunology
https://pubmed.ncbi.nlm.nih.gov/***
▷ Full tool return (4,133 characters)

Think | Round 8

Benoit Van den Eynde is a tumor immunologist at the Ludwig Institute, but he's more associated with tumor immunology and antigens rather than cytokines and receptor families specifically. Let me search for other KU Leuven scholars affiliated with Ludwig Cancer Research who focus on cytokines and receptor families. Maybe it's a different person.
Let me search for "KU Leuven cytokine receptor family" directly.

Tool call | web_search | Round 8

{"query": ["KU Leuven cytokine receptor family", "Catholic University of Leuven cytokine receptor", "KU Leuven immunology cytokine receptor family professor", "KU Leuven cancer immunology cytokine"]}

Tool response | Round 8

Summary preview
Structural Biology of Shared Cytokine Receptors - PMC
https://pmc.ncbi.nlm.nih.gov/***
Laboratory of Molecular Immunology
https://rega.kuleuven.be/***
Paul Proost KU Leuven
https://www.researchgate.net/***
▷ Full tool return (5,485 characters)

Think | Round 9

I found Ghislain OPDENAKKER listed as Professor of Immunology, Head of Immunobiology, focusing on cytokines, chemokines, adhesion molecules and other membrane receptors at KU Leuven. He also has an MD from Catholic University of Leuven. Let me search for him specifically.

Tool call | web_search | Round 9

{"query": ["Ghislain Opdenakker Annual Review of Immunology 2003", "Ghislain Opdenakker cytokine receptor schematic", "Opdenakker 2003 cytokine receptor review"]}

Tool response | Round 9

Summary preview

Interferons and other cytokines, genetics and beyond in ... - PMC
https://pmc.ncbi.nlm.nih.gov/***

Ghislain Opdenakker - Loop
https://loop.frontiersin.org/***

Ghislain OPDENAKKER — Professor of Immunology, Head ...
https://www.researchgate.net/***

▷ Full tool return (4,229 characters)

Think | Round 10

Ghislain Opdenakker is a professor at KU Leuven focusing on cytokines, but I'm not sure if he published in Annual Review of Immunology in 2003. Let me search more systematically for 2003 Annual Review of Immunology articles about cytokine receptors and their families. The description mentions "multiple members of receptor families involved in antiviral and inflammatory regulation" - this sounds like it could be about the interferon receptor family or interleukin receptor family.
Let me try to find the 2003 Annual Review of Immunology table of contents or specific chapter about cytokine receptors.

Tool call | web_search | Round 10

{"query": ["2003 Annual Review of Immunology table of contents cytokine receptors", "Annual Review of Immunology 2003 \"cytokine receptors\" \"receptor family\"", "2003 Annual Review of Immunology \"interferon\" \"receptor\" schematic", "2003 Annual Review of Immunology \"interleukin\" \"receptor\""]}

Tool response | Round 10

Summary preview

Structural Biology of Shared Cytokine Receptors
https://www.annualreviews.org/***

Interleukin-10 and Related Cytokines and Receptors
https://www.annualreviews.org/***

Cytokine-mediated communications: a quantitative appraisal ...
https://pmc.ncbi.nlm.nih.gov/***

▷ Full tool return (6,021 characters)

Think | Round 11

Let me try to find the specific 2003 Annual Review of Immunology article about multiple receptor families involved in antiviral and inflammatory regulation. The description mentions "schematic systematically depicting multiple members of receptor families" - this sounds like it could be a figure showing the classification of cytokine receptors.
Let me search for "2003 Annual Review of Immunology cytokine receptor families schematic figure" or try to access the actual PDF of the 2003 issue.

Tool call | web_search | Round 11

{"query": ["2003 Annual Review of Immunology cytokine receptor families figure schematic", "Annual Review of Immunology 2003 cytokine receptor classification", "2003 Annual Reviews immunology cytokine

receptor family diagram color", "2003 Annual Review of Immunology \"interferon\" \"interleukin\" receptor family"]}

Tool response | Round 11

Summary preview

The Common Cytokine Receptor Chain Family of ... - PMC
https://pmc.ncbi.nlm.nih.gov/***

Cytokine Receptor - an overview
https://www.sciencedirect.com/***

The c Family of Cytokines: Basic Biology to Therapeutic ...
https://www.cell.com/***

▷ Full tool return (5,968 characters)

Think | Round 12

Let me try a different approach. The scholar might be someone like Willem J. (WIJ) J. (maybe a professor at KU Leuven). Let me search for "KU Leuven immunology professor Ludwig Institute cytokine" more broadly. Also, maybe the scholar is a Belgian immunologist who has published in Annual Review of Immunology about cytokine receptors. Let me search for "Annual Review of Immunology 2003" with more specific terms about receptor families.
Actually, the description mentions "multiple members of receptor families involved in antiviral and inflammatory regulation" - this sounds like it could be about th ...

Tool call | web_search | Round 12

{"query": ["2003 Annual Review of Immunology cytokine receptors review article", "2003 Annual Review of Immunology \"cytokine\" \"receptor\" \"families\"", "2003 Annual Review of Immunology cytokine receptor family review", "Annual Review of Immunology 2003 \"receptor\" \"cytokine\" \"inflammatory\""]}

Tool response | Round 12

Summary preview

Structural Biology of Shared Cytokine Receptors
https://www.annualreviews.org/***

Interleukin-10 and Related Cytokines and Receptors
https://www.annualreviews.org/***

Historical Review of Cytokines - PMC - NIH
https://pmc.ncbi.nlm.nih.gov/***

▷ Full tool return (5,922 characters)

Think | Round 13

The search results are not giving me the specific article. Let me try to find the 2003 Annual Review of Immunology table of contents or specific chapter. The description mentions "multiple members of receptor families involved in antiviral and inflammatory regulation" - this could be referring to a figure showing the classification of cytokine receptors across different families. Let me search for "Annual Review of Immunology 2003 cytokine receptors schematic figure color" or try to find the actual PDF of the 2003 issue.
Actually, let me think about this differently. The scholar taught at Catholi ...

Tool call | web_search | Round 13

{"query": ["Ludwig Institute for Cancer Research KU Leuven professor cytokine", "Ludwig cancer research KU Leuven immunology professor", "KU Leuven Ludwig Institute for Cancer Research immunologist"]}

Tool response | Round 13

Summary preview
Ludwig Cancer Research and Oxford University welcome ...
https://www.ludwigcancerresearch.org/***
Stefan Constantinescu
https://www.deduveinstitute.be/***
Jean-Christophe Renauld - Professor at the Universite ...
https://be.linkedin.com/***
▷ Full tool return (4,112 characters)

Think | Round 14

I found Ghislain OPDENAKKER listed as Professor of Immunology and Head of Immunobiology at KU Leuven. He might be the one. But I need to confirm if he published in Annual Review of Immunology in 2003. Let me search more specifically for him and Annual Review of Immunology 2003. Actually, the 2003 Annual Review of Immunology is Volume 21. Let me search for Volume 21 of Annual Review of Immunology and look for a chapter about cytokine receptors.
Actually, let me reconsider the scholar. The description says "long taught immunology at the Catholic University of Leuven in Belgium" - this could be some ...

Tool call | web_search | Round 14

```
{"query": ["Jean-Christophe Renauld Ludwig Institute cancer immunology", "Jean-Christophe Renauld Annual Review of Immunology 2003", "Jean-Christophe Renauld cytokine receptor"]}
```

Tool response | Round 14

Summary preview
Jean-Christophe Renauld - Professor at the Universite ...
https://be.linkedin.com/***
Jean-Christophe RENAULD — Professor — MD, PhD
https://www.researchgate.net/***
Masters of the JAKs
https://www.ludwigcancerresearch.org/***
▷ Full tool return (4,270 characters)

Think | Round 15

Great! I found Jean-Christophe Renauld. The PubMed result shows: "Class II cytokine receptors and their ligands: key antiviral and inflammatory modulators" by JC Renauld, 2003, Rev Immunol. 2003 Aug;3(8):667-76. But wait, the user said "leading immunology review journal" and "2003". This is "Reviews in Immunology" which might be a review journal. But the user mentioned "leading immunology review journal" - could be Annual Review of Immunology, Nature Immunology, Immunity, or Reviews in Immunology.
Wait, let me look at the search result more carefully. It says "by JC Renauld 2003 Cited by 395 ...

Tool call | web_extract | Round 15

```
{"urls": ["https://pubmed.ncbi.nlm.nih.gov/***"], "goal": "Find details about the article by Jean-Christophe Renauld from 2003, including journal name, title, and any descriptions of figures showing receptor families and schematic diagrams. Also check if this is the article described in the user's question."}
```

Tool response | Round 15

Summary preview
Class II cytokine receptors and their ligands: key antiviral and inflammatory modulators - PubMed

https://pubmed.ncbi.nlm.nih.gov/***

▷ Full tool return (1,236 characters)

Think | Round 16

The article is "Class II cytokine receptors and their ligands: key antiviral and inflammatory modulators" published in Nature Reviews Immunology in August 2003. This matches the description: Jean-Christophe Renauld is a professor at the Universite catholique de Louvain (KU Leuven, Catholic University of Leuven) and is affiliated with the Ludwig Institute for Cancer Research. He has published reviews on cytokine receptors.
Now I need to find the specific figure from this article that shows "multiple members of receptor families involved in antiviral and inflammatory regulation" and "structures of ...

Tool call | web_search | Round 16

```
{"query": ["Class II cytokine receptors and their ligands: key
antiviral and inflammatory modulators figure schematic", "Renauld
Nature Reviews Immunology 2003 receptor families schematic", "Nature
Reviews Immunology 2003 \"Class II cytokine receptors\" figure",
"Renauld 2003 Nature Reviews Immunology cytok ...
```

Tool response | Round 16

Summary preview

Class II cytokine receptors and their ligands: key antiviral ...
https://pubmed.ncbi.nlm.nih.gov/***

(PDF) Class II cytokine receptors and their ligands
https://www.researchgate.net/***

Class II cytokine receptors and their ligands: Key antiviral ...
https://www.scilit.com/***

▷ Full tool return (5,916 characters)

Think | Round 17

I need to find the actual figures from this article. Let me try to access the PDF or the full text through ResearchGate or other sources. I should look for the figure showing receptor families and transmembrane complexes. Let me search for images or the specific figure from this paper.

Tool call | web_search | Round 17

```
{"query": ["site:researchgate.net \"Class II cytokine receptors\"
figure", "Renauld 2003 Nature Reviews Immunology schematic receptor
families", "Nature Reviews Immunology 2003 figure cytokine receptor
complex membrane color"]}
```

Tool response | Round 17

Summary preview

(PDF) Class II cytokine receptors and their ligands
https://www.researchgate.net/***

Phylogenetic tree of class II cytokines and receptors.The...
https://www.researchgate.net/***

Expression pattern for the classII helical cytokine receptor ...
https://www.researchgate.net/***

▷ Full tool return (4,586 characters)

Think | Round 18

Let me try to access the ResearchGate publication directly to see the figures. I need to find the specific figure that shows receptor families and transmembrane complexes with colored membrane-embedded portions.

Tool call | web_extract | Round 18

{"urls": ["https://www.researchgate.net/***"], "goal": "Find any figures from this article, especially a schematic showing multiple receptor families and transmembrane complexes with colored portions of ...

Tool response | Round 18

Summary preview

Just a moment...

https://www.researchgate.net/***

▷ Full tool return (626 characters)

Think | Round 19

The ResearchGate page is blocked. Let me try to find the figure through other sources. I can search for the specific figure or try to access the PDF through the Springer link or other repositories. Let me also search for images from this paper.

Tool call | web_search_for_image | Round 19

{"query": ["Class II cytokine receptors and their ligands Nature Reviews Immunology 2003 figure", "Renauld 2003 cytokine receptor schematic", "Nature Reviews Immunology 2003 receptor families diagram"]}

Tool response | Round 19

Summary preview

Tool returned a structured payload; only the preview is shown.

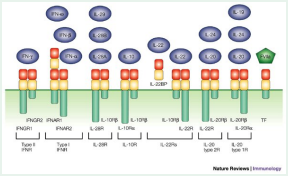

search_21d0f33d.jpg

▷ Full tool return (7,498 characters)

Think | Final

I found the key figure! The second image (search_21d0f33d.jpg) shows exactly what the user described: a schematic systematically depicting multiple members of receptor families involved in antiviral and inflammatory regulation, and the structures of their transmembrane complexes. Looking at this figure, the portions of the receptor complexes embedded in the cell membrane are shown in **green** color ...

Final response

Based on my research, the scholar is **Jean-Christophe Renauld**, a professor at the Catholic University of Leuven (KU Leuven) in Belgium and a researcher at the Ludwig Institute for Cancer Research. In 2003, he published a review article titled ”Class II cytokine receptors and their ligands: key antiviral and inflammatory modulators” in *Nature Reviews Immunology*.

The schematic you're referring to is a figure from this article that systematically depicts multiple members of Class II cytokine receptor families (involved in antiviral and inflammatory regulation) and the structures of their transmembrane complexes. In this figure, the portions of the receptor complexes embedded in the cell membrane are shown in **green** (or a greenish/teal color).
The figure shows various Class II cytokine receptors (such as IFNGR2/IFNAR1, IFNAR2, IL-28RA, IL-10R, IL-10R, IL-22R, IL-20R, IL-20R, and TF) with their characteristic transmembrane domains passing through the cell membrane, which are depicted in green. The extracellular receptor domains are shown in blue/purple, while the transmembrane portions are colored green.

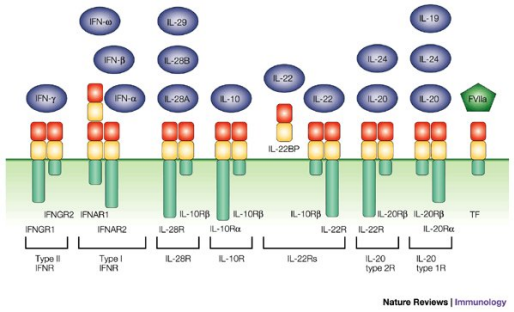

search_21d0f33d.jpg

Case 3: Game-lineage, award, and screenshot reasoning

Query. This two-dimensional puzzle game, released in 2019, uses impossible architecture as its core visual style; its game world creates the illusion of infinite extension by looping and folding through three-dimensional space. The development team later released a new VR title that prioritizes a room-scale experience. What award did the new title receive in February 2026? In addition, in an official screenshot from the new title, a large, fluid-like bright-red object stands in the center of a hall; what shape is the light-blue decoration to its left?

Reference answer: In February 2026, it received: the 29th Annual D.I.C.E. Award for Immersive Reality Technical Achievement; the shape of the light-blue decoration on the left is: a water-drop shape.

Target image

(not provided to the model at query time)



Think | Round 1

The user is asking about a specific puzzle game released in 2019 with impossible architecture as its core visual style, which creates the illusion of infinite extension by looping and folding through three-dimensional space. The development team later released a new VR title that prioritizes a room-scale experience. I need to identify: 1. The award the new title received in February 2026 2. In an official screenshot from the new title, a large, fluid-like bright-red object stands in the center of a hall; what shape is the light-blue decoration to its left?
First, I need to identify the game. Let ...

Tool call | web_search | Round 1

```
{"query": "puzzle game 2019 impossible architecture infinite
extension looping folding VR sequel room-scale", "num": 5}
```

Tool response | Round 1

Summary preview

Infinite Hallways in Roomscale VR — Hotel Infinity Devlog #2
https://www.youtube.com/***
Roomscale VR Has NEVER Been This Wild! Hotel Infinity
https://www.youtube.com/***
Building an Impossible VR Game — Hotel Infinity Devlog #1
https://www.youtube.com/***

▷ Full tool return (2,168 characters)

Think | Round 2

The search results point to ”Hotel Infinity”. Let me search for more details about Hotel Infinity, including its 2019 release and the VR title that prioritizes room-scale experience, as well as awards in February 2026.

Tool call | web_search | Round 2

```
{"query": "Hotel Infinity VR game 2019 impossible architecture award
February 2026", "num": 5}
```

Tool response | Round 2

Summary preview

Building an Impossible VR Game — Hotel Infinity Devlog #1
https://www.youtube.com/***
Hotel Infinity We Hope Your Stay is Eternal Studio Chyr
https://hotelinfinityvr.com/***
Hotel Infinity review: mind-bending corridors, hollow ...
https://www.thevrcritic.com/***
▷ Full tool return (2,256 characters)

Think | Round 3

Hotel Infinity is definitely the game. Now I need to find: 1. What award did the new title receive in February 2026? 2. In an official screenshot from the new title, a large, fluid-like bright-red object stands in the center of a hall; what shape is the light-blue decoration to its left?
Let me search for more specific information about awards in February 2026 and also look for official screenshots.

Tool call | web_search | Round 3

```
{"query": "Hotel Infinity award February 2026", "num": 5}
```

Tool response | Round 3

Summary preview

42nd Annual ICP Infinity Awards
https://www.icp.org/***
2026 Infinity Awards
https://givebutter.com/***
Receiving the 2026 ICP Infinity Award for Documentary and ...
https://hsakag.substack.com/***
▷ Full tool return (2,085 characters)

Think | Round 4

The search didn't find the specific award. Let me try different queries - maybe it's an award at a specific event like Gamescom, The Game Awards, or something like that. Also the user mentions "new VR title that prioritizes a room-scale experience" - this could be Hotel Infinity itself, or maybe a sequel? Let me search more specifically.

Tool call | web_search | Round 4

```
{"query": ["Hotel Infinity VR award 2026", "Hotel Infinity Steam
Awards", "Hotel Infinity accolades 2026"], "num": 5}
```

Tool response | Round 4

Summary preview

HOTEL INFINITY by William Chyr - XRGuide
https://xrguide.app/***
William Chyr's Post
https://www.linkedin.com/***
Hotel Infinity review: mind-bending corridors, hollow ...
https://www.thevrcritic.com/***
▷ Full tool return (5,723 characters)

Think | Round 5

I found some relevant information: 1. Hotel Infinity won the ”Immersive Reality Technical Achievement” at the DICE Awards (29th DICE Awards) 2. It won the ”Auggie Award Winner — Best Indie Creator(s)” at AWE USA 2026
But the user specifically asked about an award in February 2026. Let me search more specifically for February 2026 awards and also look for official screenshots with the described scene.

Tool call | web_search | Round 5

```
{"query": ["Hotel Infinity February 2026 award", "Hotel Infinity DICE Awards 2026 February", "Hotel Infinity Steam Awards February 2026"], "num": 5}
```

Tool response | Round 5

Summary preview

Only a few days are left to cast your votes in this year’s @ ...
https://www.instagram.com/***
AIDAC INFINITY AWARD 2026 Hospitality Architecture
https://www.instagram.com/***
Celebrity Infinity Member Review
https://www.cruisecritic.com/***

▷ Full tool return (4,803 characters)

Think | Round 6

Great! I found that Hotel Infinity won the ”Immersive Reality Technical Achievement” at the 29th D.I.C.E. Awards on February 12, 2026. That’s the award in February 2026.
Now I need to find the official screenshot with the described scene: a large, fluid-like bright-red object in the center of a hall, and a light-blue decoration to its left. Let me search for official screenshots of Hotel Infinity.

Tool call | web_search_for_image | Round 6

```
{"query": "Hotel Infinity official screenshot hall red object blue decoration", "num": 10}
```

Tool response | Round 6

Summary preview
Tool returned a structured payload; only the preview is shown.



search_b5af4025.jpg



search_1a3eb628.jpg

▷ Full tool return (2,460 characters)

Think | Round 7

I found some images. Let me examine them more carefully. Image 1 shows a large light-blue (teal/cyan) teardrop-shaped object on the left side of what appears to be a lobby or atrium, with a red fluid-like object in the center. Image 2 shows a different room. Image 3 shows a bedroom. Image 5 shows a reception desk. Image 9 shows another lobby scene.
Image 1 and Image 9 both show a large red fluid-like object in the center of a hall. In Image 1, to the left of the red object, there’s a large light-blue teardrop-shaped object. In Image 9, there’s a cyan/blue spiral-shaped decoration to the right of ...

Tool call | web_extract | Round 7

{"urls": ["https://hotelinfinityvr.com/***"], "goal": "Find official screenshots of Hotel Infinity, especially any showing a large red fluid-like object in a hall with a light-blue decoration"}

Tool response | Round 7

Summary preview

Hotel Infinity

https://hotelinfinityvr.com/***







extract_458eb29c.jpg extract_6d71ed0b.jpg extract_5eba61bd.jpg extract_a807c31b.jpg extract_6475e49c.jpg

▷ Full tool return (1,544 characters)

Think | Round 8

Excellent! I found the exact screenshot from the official website. Image extract_807c31b.jpg shows the atrium scene perfectly. Let me examine it carefully:
- The large, fluid-like bright-red object is in the center of the hall - To its left is a large light-blue decoration - The shape of this light-blue decoration is clearly a teardrop (or water drop) shape - it's a large, smooth, teardrop-shaped sculpture
Let me also verify the award information. I found that Hotel Infinity won the "Immersive Reality Technical Achievement" at the 29th D.I.C.E. Awards on February 12, 2026. Let me extract that in ...

Tool call | web_extract | Round 8

{"urls": ["http://www.shacknews.com/***"], "goal": "Confirm that Hotel Infinity won an award at the 29th D.I.C.E. Awards in February 2026 and identify the specific award"}

Tool response | Round 8

Summary preview

The D.I.C.E. Awards 2026 winners & finalists

http://www.shacknews.com/***

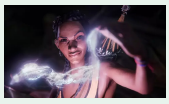

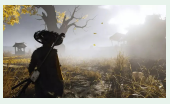


extract_cf26c5f2.jpg extract_21a27db5.jpg extract_9fcd2c5c.jpg extract_6166642b.jpg

▷ Full tool return (1,693 characters)

Think | Final

Use the verifiable evidence gathered in the trajectory and answer the question as requested.

Final response

Based on my research:
1. **The game**: The puzzle game released in 2019 with impossible architecture is **Manifold Garden**, developed by William Chyr. The development team later released **Hotel Infinity** (2025), a VR title designed for room-scale experience.
2. **The award**: In February 2026, Hotel Infinity won the **Immersive Reality Technical Achievement** award at the **29th D.I.C.E. Awards**, which took place on February 12, 2026 at the Aria Resort & Casino in Las Vegas.

3. **The screenshot** (from the official Hotel Infinity website): In the atrium/lounge hall screenshot, there is a large, fluid-like bright-red object (glowing red ooze-like liquid) in the center. To its left is a large light-blue decoration shaped like a **teardrop** (or water drop).



extract_a807c31b.jpg