WideSeek: Advancing Wide Research via Multi-Agent Scaling

Anonymous ACL submission

Abstract

Search intelligence is evolving from Deep Research to Wide Research, a paradigm essential for retrieving and synthesizing comprehensive information under complex constraints in parallel. However, progress in this field is impeded by the lack of dedicated benchmarks and optimization methodologies for search breadth. To address these challenges, we take a deep dive into Wide Research from two perspectives: **Data Pipeline** and **Agent Optimization**. First, we produce WideSeekBench, a General Broad Information Seeking (GBIS) benchmark constructed via a rigorous multi-phase data pipeline to ensure diversity across the target information volume, logical constraints, and domains. Second, we introduce WideSeek, a dynamic hierarchical multi-agent architecture that can autonomously fork parallel sub-agents based on task requirements. Furthermore, we design a unified training framework that linearizes multi-agent trajectories and optimizes the system using end-to-end RL. Experimental results demonstrate the effectiveness of WideSeek and multi-agent RL, highlighting that scaling the number of agents is promising for advancing the Wide Research paradigm.

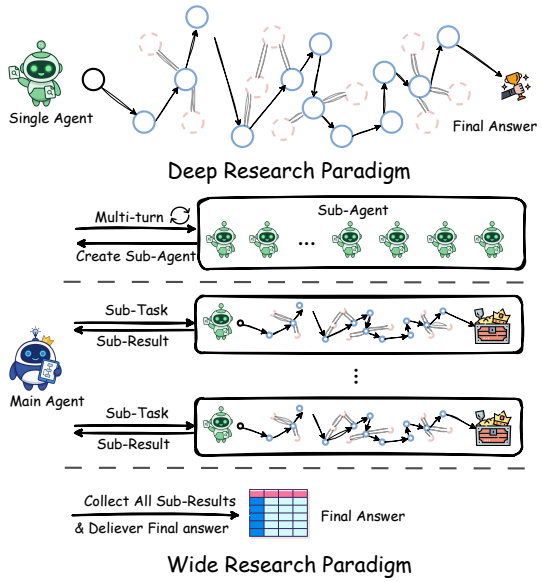

Figure 1: Deep Research paradigm vs. Wide Research paradigm.

1 Introduction

Search Intelligence constitutes the cornerstone of Agentic AI (Shi et al., 2025; Abou Ali et al., 2025), serving as an essential module for complex applications including code generation (Jimenez et al., 2024), data intelligence (Lei et al., 2025), and GUI manipulation (Xie et al., 2024). Existing research has predominantly focused on **Deep Research** (Wei et al., 2025), which employs multi-step reasoning to locate a single hard-to-find piece of information. As AI enters its *Second Half* (Yao, 2025), the community is shifting toward real-world utility, necessitating **Wide Research** (Manus, 2025) (as shown in Figure 1), a parallel orchestration paradigm that prioritizes high-breadth synthesis over sequential depth. Wide Research decomposes complex objectives into parallel sub-tasks executed by autonomous agents, as exemplified by Kimi Agent-Swarm (Team et al., 2026) and Manus (Manus, 2025). A representative application is generating competitive analysis tables that synthesize information from thousands of sources, substantially reducing human labor while enhancing productivity at scale.

Despite its promise, Wide Research faces three challenges: (1) **Benchmark Limitations:** Existing benchmarks (Wong et al., 2025; Lan et al., 2025) are human-constructed, limiting scale and diversity, and provide only test sets without training data; (2) **Data Synthesis Deficiencies:** Current methods focus on sampling graph topologies for multi-step reasoning paths (Li et al., 2025; Tao et al., 2025), optimizing search depth but lacking capacity for

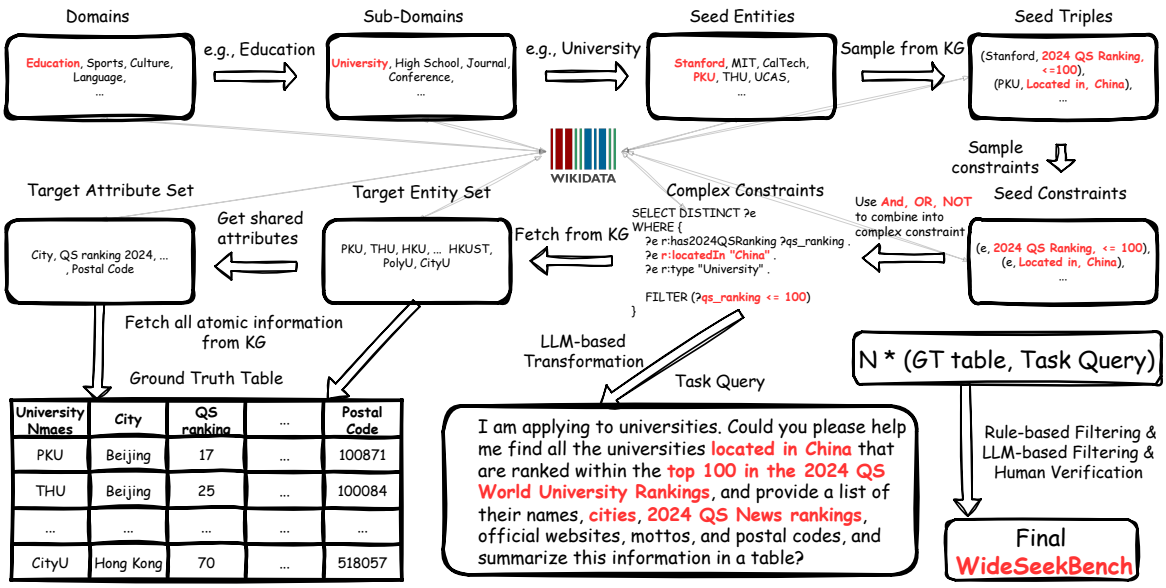

Figure 2: The data pipeline of WideSeekBench construction, which mines a set of target information under complex constraints.

synthesizing broad information under complex constraints; (3) **Optimization Gaps:** Prior work relies on closed-source models in static frameworks (Roucher et al., 2025) or enhances single-agent depth (Lu et al., 2025), with no exploration of end-to-end optimization for autonomously broadening search paths. In this work, we address these challenges from two perspectives:

Data Pipeline & Benchmark. Unlike conventional methods that construct reasoning paths toward a single answer, we leverage Knowledge Graphs (Schmelzeisen et al., 2021) to extract clusters of interconnected world knowledge and apply formal set operations to construct complex constraints. Based on a meticulous multi-phase pipeline, we introduce **WideSeekBench**, a benchmark for General Broad Information Seeking (GBIS) with 5,156 tasks spanning training and test sets, balanced across information volume, constraint complexity, and 18 domains.

Agent Optimization. We propose **WideSeek**, a dynamic multi-agent system where a planner autonomously forks any number of executor sub-agents at each step, in contrast to static frameworks with pre-defined agent roles. We optimize this system through the following pipeline: (1) trajectory distillation via SFT to bootstrap the multi-agent scaffold; (2) end-to-end RL that linearizes all agent trajectories into a unified sequence and optimizes the entire system jointly via RL, enabling the planner and executors to co-evolve toward maximizing search utility.

Our experiments demonstrate these key findings: (1) the WideSeek scaffold alone yields substantial gains over single-agent ReAct baselines without any training, validating the architectural advantage of dynamic multi-agent decomposition; (2) end-to-end multi-agent RL further unlocks coordination at scale, enabling a small open-source model to match or surpass much larger proprietary models; (3) WideSeek generalizes effectively from wide to deep research and from simulated to real-world environments, confirming that multi-agent coordination learned for breadth transfers to both depth and open-web settings; (4) fine-grained analysis across information volume, constraint type, and domain provides deeper insights into multi-agent scaling behavior.

2 Data Pipeline & Benchmark

In contrast to Deep Research, Wide Research represents an application that is more closely aligned with real-world productivity scenarios. It aims to retrieve a collection of relevant information that satisfies complex constraints. Specifically, we can compile all relevant information into a table for comparative analysis. We define this task as General Broad Information Seeking (GBIS). To systematically evaluate models' Wide Research capabilities and to further investigate how post-training can enhance these capabilities in base models, we

propose a rigorous multi-stage data pipeline and thus construct the **WideSeekBench**.

2.1 Task Definition

We define the GBIS task over a universe of entities $\mathcal{E}$ within a world knowledge space $\mathcal{W}$. A task instance is formally defined as a tuple $\mathcal{T} = (\mathcal{Q}, \mathcal{A})$, where $\mathcal{Q}$ is a task query encoding a complex semantic constraint, and $\mathcal{A} = \{a_1, a_2, \ldots, a_m\}$ is the set of required attributes.

The query $\mathcal{Q}$ maps to a latent semantic filter function $\Phi : \mathcal{E} \rightarrow \{0, 1\}$. The objective is to construct a ground truth table $\mathbf{T}^*$ corresponding to the target entity set $\mathbf{E}^* = \{e \in \mathcal{E} \mid \Phi(e) = 1\}$. Formally, $\mathbf{T}^*$ is a table of size $|\mathbf{E}^*| \times m$:

$$\mathbf{T}^* = \begin{bmatrix} v_{1,1} & \cdots & v_{1,m} \\ \vdots & \ddots & \vdots \\ v_{|\mathbf{E}^*|,1} & \cdots & v_{|\mathbf{E}^*|,m} \end{bmatrix}, v_{i,j} = \text{Value}(e_i, a_j) \quad (1)$$

GBIS requires the agent to comprehensively synthesize $\mathbf{T}^*$. This requires not only the precision of the search but also the recall.

2.2 Data Pipeline

We employ a multi-phase approach on a knowledge graph $\mathcal{K}$ to synthesize complete benchmark instances of the form $(\mathcal{Q}, \mathcal{A}, \mathbf{T}^*, \mathcal{R})$, where $\mathcal{R}$ denotes the evaluation rubrics. We provide more details in Appendix B.

Phase 1: Seed Constraint Construction. To ensure comprehensive coverage and diversity, we adopt a top-down sampling strategy. (a) *Domain Definition & Sampling:* We start with a human-defined set of high-level domains $\mathcal{D}_{domain}$ (e.g., Education, Sports). From each high-level domain, we sample specific sub-domains $\mathcal{D}_{sub}$ (e.g., University, Basketball). (b) *Seed Sampling:* Within each sub-domain, we sample seed entities e_{seed} and extract their relations (triples) $\mathcal{R}_{seed} = \{(e_{seed}, p, v)\}$ from $\mathcal{K}$. This process yields a diverse pool of atomic constraints $\mathcal{C}_{atom}^{(e_{seed})} = \{(p, v)\}$ associated with each seed entity.

Phase 2: Logical Composition & Schema Extension. We compose atomic constraints into complex constraints and extend the attribute schema. (a) *Logical Composition:* Using operators $\mathcal{O} = \{\wedge, \vee, \neg\}$, we recursively define the composite filter Φ as:

$$\Phi(e) := c(e) \mid \neg\Phi(e) \mid \Phi_1(e) \wedge \Phi_2(e) \mid \Phi_1(e) \vee \Phi_2(e) \quad (2)$$

where $c(\cdot)$ denotes a boolean predicate induced by an atomic constraint $(p, v) \in \mathcal{C}_{atom}^{(e_{seed})}$, and $c(e) = 1$ if entity e satisfies property p with value v. We execute Φ over $\mathcal{K}$ to retrieve the target entity set $\mathbf{E}^*$. (b) *Schema Extension:* Given the validated entity set $\mathbf{E}^*$, we construct a candidate attribute set $\mathcal{A}_{cand} = \bigcup_{e \in \mathbf{E}^*} \text{Attributes}(e)$, from which we select target attributes $\mathcal{A} \subset \mathcal{A}_{cand}$ by enforcing entity coverage and sufficient value diversity, and retrieve all corresponding values to populate $\mathbf{T}^*$. This phase yields approximately 30,000 candidate tasks across all sampled domains.

Phase 3: Agent Task Synthesis. This phase converts complex constraints and target attributes into user-facing tasks using LLMs. (a) *Self-Refining Query Synthesis:* We treat query generation as an iterative, self-refining process. An LLM generator $\mathcal{M}_{gen}$ converts Φ into a query $\mathcal{Q}$, while a LLM verifier $\mathcal{M}_{ver}$ extracts logic $\hat{\Phi}$ back from $\mathcal{Q}$. Discrepancies ($\hat{\Phi} \not\equiv \Phi$) trigger feedback loops for $\mathcal{M}_{gen}$ to regenerate $\mathcal{Q}$ until consistency is achieved. The consistency is also evaluated by $\mathcal{M}_{ver}$. (b) *Column-wise Rubric Generation:* For each attribute a_j, we generate a specific evaluation rubric $\mathcal{R}_j$ based on column semantics and cell values $\mathbf{T}^*_{\cdot,j}$, defining acceptance criteria for formats and tolerances. This phase yields approximately 15,000 candidate tasks.

Phase 4: Multi-Stage Filtering. To ensure high quality, we apply a three-level filtering protocol: (a) *Rule-based Filter:* We perform web searches to discard tasks where entities in $\mathbf{E}^*$ are not grounded in a web page. Moreover, we discard tasks where some cells lack natural language descriptions or $\mathbf{T}^*$ is sparse ($> 50\%$ empty cells). (b) *LLM-based Filter:* An LLM scores tasks against five dimensions: Human-Likeness, Solvability, Common Sense, Temporal Stability, and Rubric Rationality. A final task passes all these standards. (c) *Human Verification:* A final manual review removes subtle semantic irrationalities. This cascaded filtering phase yields 5,156 high-quality final tasks.

2.3 WideSeekBench

We introduce **WideSeekBench**, a comprehensive benchmark designed to evaluate Wide Research capabilities. The dataset comprises a total of 5,156 tasks, which are partitioned into a training set of 4,436 tasks $\mathcal{D}_{train}$ and a held-out test set of 720 tasks $\mathcal{D}_{test}$. The comparison of different search agent benchmarks is shown in Table 1.

To enable fine-grained evaluation, we meticu-

Table 1: Comparison of WideSeekBench with existing information-seeking benchmarks. **Task Type**: Deep (finding specific hidden info) vs. Wide (collecting broad structured info). **Auto Gen.**: automated data pipeline. **Multi-dim.**: fine-grained task classification.

Benchmark	Size	Domains	Type	Train	Auto	Multi-dim
GAIA	103	-	Deep	✗	✗	✗
BrowseComp	1,266	9	Deep	✗	✗	✗
WideSearch	200	14	Wide	✗	✗	✗
DeepWideSearch	220	15	Wide	✗	MIX	✗
WebShaper	5,000	-	Deep	✓	✓	✗
WideSeekBench	**5,156**	**18**	**Wide**	✓	✓	✓

lously controlled the distribution of the test set. This allows for a multi-dimensional task classification and detailed analysis. Specifically, the test tasks are categorized based on three distinct dimensions: (1) **Volume of Target Information:** We quantify the volume based on the total number of cells in the ground truth table. Based on this, tasks are divided into 10 distinct intervals to assess performance across varying information volume. The specific distribution is illustrated in Figure 7b. (2) **Constraint Complexity:** To evaluate how agents handle complex tasks, we classify the tasks into 7 types based on the nature of the constraints involved. The distribution of these constraint types is presented in Table 8. (3) **Domain Diversity:** We categorize the tasks into 18 distinct domains to ensure broad topical coverage. The domain-wise distribution is shown in Figure 7d.

Furthermore, we ensure that all entities in ground truth tables correspond to existing real-world web pages via search. To guarantee a fair, transparent, and reproducible evaluation, we constructed a standalone Simulated Environment. This environment includes a local document corpus and a local search engine. Detailed specifications of the simulated environment are provided in the Appendix B.7. Following WideSearch (Wong et al., 2025), we use Success Rate, Row F1, and Item F1 as the evaluation metrics. We show the details of evaluation in Appendix B.8.

3 WideSeek

Given a task $(\mathcal{Q}, \mathcal{A})$, the objective is to retrieve related information to construct a structured table $\hat{\mathbf{T}}$ containing a set of entities $\hat{\mathbf{E}} = \{e_1, e_2, \ldots, e_N\}$ and their corresponding attribute values $v(\hat{\mathbf{E}}, \mathcal{A})$, satisfying a complex semantic constraint Φ derived from $\mathcal{Q}$. To address the complexity of this task, which often exceeds the context and reasoning limits of a single serial trajectory, we propose **WideSeek**. WideSeek operates as a dynamic, hierarchical multi-agent system governed by a unified policy π_θ.

3.1 Multi-Agent Rollout

The inference process, as shown in the left of Figure 3, is modeled as a hierarchical Markov Decision Process (MDP) (Luo et al., 2025). Unlike static multi-agent architectures with fixed roles, WideSeek employs a centralized Main Agent (Planner) that dynamically forks variable instances of Sub-Agents (Executors) at any step.

Hierarchical State Transition. At the top level, the Main Agent operates at time steps t. Let s_t^{main} denote the global state, encompassing the user query $\mathcal{Q}$ and the history of high-level thoughts and sub-results. The Main Agent's policy $\pi_\theta(a_t^{main}|s_t^{main})$ selects an action a_t^{main} from a hierarchical action space $\mathbf{A} = \mathbf{A}_{\text{planning}} \cup \mathbf{A}_{\text{termination}}$.

If $a_t^{main} \in \mathbf{A}_{\text{planning}}$, the agent invokes the function `create_sub_agent`$(q_{sub}^{(1)}, \ldots, q_{sub}^{(k)})$. This action triggers the parallel instantiation of k Sub-Agents, where k is dynamically determined by the policy rather than a hyperparameter. Each Sub-Agent j ($j \in \{1, \ldots, k\}$) operates in its own local MDP defined by the sub-task $q_{sub}^{(j)}$. It generates a trajectory $\mathcal{T}_{sub}^{(j)} = (s_0^j, a_0^j, s_1^j \ldots)$[1] using the same unified policy π_θ, utilizing atomic search tools (e.g., `search`, `open_page`). Each action execution receives an observation o_t^j from the environment and updates the sub-agent state: $s_{t+1}^j \leftarrow s_t^j \cup o_t^j$. Upon completion, the sub-agent returns a textual sub-result r_j, which updates the global state: $s_{t+1}^{main} \leftarrow s_t^{main} \cup \{r_1, \ldots, r_k\}$. If $a_t^{main} \in \mathbf{A}_{\text{termination}}$, the agent synthesizes the accumulated information in s_t^{main} to produce the final answer $\mathbf{T}_{ans}$ and terminates the rollout.

This hierarchical execution generates a composite trajectory $\mathcal{T}$ that interleaves the planner's reasoning traces with the execution traces of all dynamically created sub-agents.

3.2 Cold Start

Given the complexity of the task, we distill high-quality trajectories from multiple teacher models and fine-tune the policy via SFT (Supervised Fine-Tuning). Further details are provided in the Appendix C.1.

[1] We reuse $\mathcal{T}$ to represent trajectories.

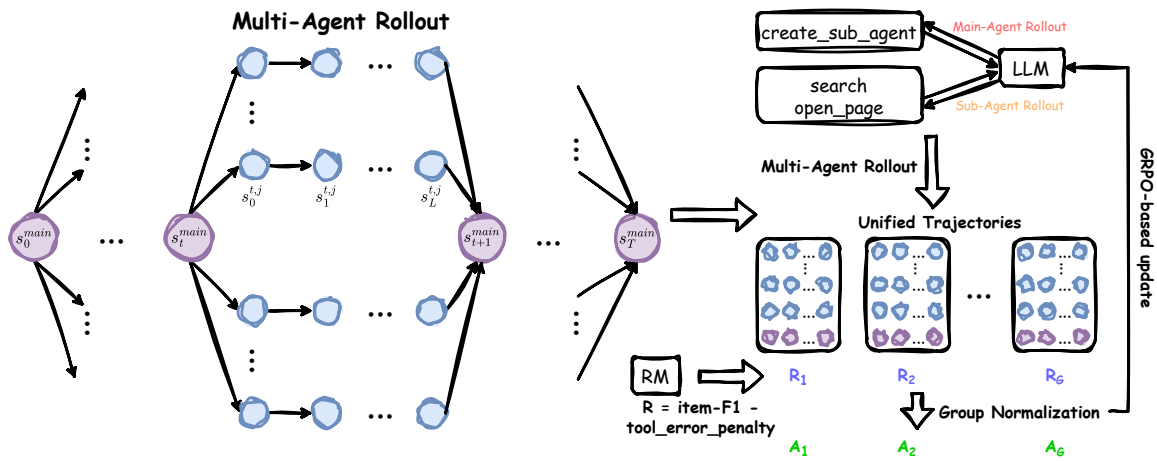

Figure 3: Overview of WideSeek's multi-agent reinforcement learning framework. Left: the planner dynamically forks sub-agents for parallel search execution. Right: all agent trajectories are linearized into a unified sequence for joint optimization via GRPO.

3.3 Multi-Agent Reinforcement Learning

Standard single-agent RL optimizes a sequential trajectory. However, WideSeek's execution graph is a dynamic tree structure. We propose a **Unified Multi-Agent RL** framework that models the entire system as a single generative process optimized via Group Relative Policy Optimization (GRPO) (Shao et al., 2024).

Unified Trajectory Modeling. We model the multi-agent interaction as a unified joint distribution. Since all agents share the same LLM checkpoint π_θ, we linearize the hierarchical execution trace into a single sequence. First, we define the trajectory of the j-th Sub-Agent forked at the Main Agent's time step t as a complete sequence of local state-action pairs:

$$\mathcal{T}_{\text{sub}}^{(t,j)} = \left[(s_0^{t,j}, a_0^{t,j}), (s_1^{t,j}, a_1^{t,j}), \ldots, (s_L^{t,j}, r_{t,j})\right] \quad (3)$$

The global unified trajectory $\boldsymbol{\mathcal{T}}$ is then constructed by interleaving each Main Agent step $(s_t^{\text{main}}, a_t^{\text{main}})$ with the set of trajectories from all K_t Sub-Agents forked at that step:

$$\boldsymbol{\mathcal{T}} = \left[(s_0^{\text{main}}, a_0^{\text{main}}), \bigcup_{j=1}^{K_0} \mathcal{T}_{\text{sub}}^{(0,j)}, \ldots, (s_t^{\text{main}}, a_t^{\text{main}}), \underbrace{\bigcup_{j=1}^{K_t} \mathcal{T}_{\text{sub}}^{(t,j)}}_{\text{Executors at step } t}, \ldots, (s_T^{\text{main}}, Y)\right] \quad (4)$$

Reward Function Design. To guide the policy toward both accurate information retrieval and robust tool usage, we define a comprehensive global reward $R(\boldsymbol{\mathcal{T}})$ that serves as the sparse training signal. The reward is composed of a correctness score based on Item-F1 and a penalty for format violations.

To discourage structural degradation, we impose a format penalty. Let n_{err} be the total count of format errors (e.g., invalid tool calls) in trajectory $\boldsymbol{\mathcal{T}}$, and N_{max} be a predefined maximum tolerance for errors. The final reward function is defined as:

$$R(\boldsymbol{\mathcal{T}}) = \text{Item-F1}(\mathbf{T}_{ans}, \mathbf{T}^*) - \lambda \cdot \underbrace{\left(\frac{n_{err}}{N_{max}}\right)}_{\text{Format Penalty}} \quad (5)$$

where λ is a balancing coefficient. This ensures that the agent is penalized proportionally to the frequency of format hallucinations relative to the tolerance threshold.

Optimization via Unified GRPO. We optimize π_θ to maximize the expected reward of the unified trajectory. For each query $\mathcal{Q}$, we sample a group of G unified trajectories $\{\boldsymbol{\mathcal{T}}_1, \ldots, \boldsymbol{\mathcal{T}}_G\}$. The Global GRPO objective is formally defined as:

$$\mathcal{J}(\theta) = \mathbb{E}_{\mathcal{Q}\sim\mathcal{D},\{\boldsymbol{\mathcal{T}}_g\}\sim\pi_{\theta_{old}}}\left[\frac{1}{G}\sum_{g=1}^{G}\frac{1}{|\boldsymbol{\mathcal{T}}_g|}\sum_{u=1}^{|\boldsymbol{\mathcal{T}}_g|}\frac{1}{|a_{u,k}|}\sum_{k=1}^{|a_{u,k}|}\min\left(\rho_{g,u,k}\hat{A}_g, \text{clip}(\rho_{g,u,k}, 1-\epsilon, 1+\epsilon)\hat{A}_g\right)\right] \quad (6)$$

Here, k indexes the action tokens generated by the model across the each step in linearized unified trajectory $\boldsymbol{\mathcal{T}}_g$, covering both Main Agent planning steps and Sub-Agent execution steps. The

Table 2: Main results on WideSeekBench (4 runs per task). We compare proprietary models, open-source baselines, and WideSeek variants across success rate, Row-F1, Item-F1, sub-agent count, and tool call count. Gains over Qwen3-8B base are shown in red.

Model	Success Rate (%)	Row F1 Score (%)		Item F1 Score (%)		# Sub-Agents	# Tool Calls
	Pass@4	Mean@4	Max@4	Mean@4	Max@4		
Proprietary Models							
GPT-5.2	0.00	**4.45**	6.75	**21.03**	**26.88**	11.21	**408.64**
GPT-5.1	0.00	4.11	6.75	20.44	27.88	6.02	121.36
DeepSeek-v3.2	0.00	4.34	**6.85**	20.51	27.09	**31.25**	326.41
Kimi-K2-Thinking	0.00	3.17	5.86	17.48	25.19	8.74	85.36
Seed-1.8	0.14	3.44	5.92	17.88	25.23	7.93	88.36
Open-Sourced Models							
Qwen3-8B-Thinking	0.00	0.53	1.51	7.37	12.71	4.18	9.50
Qwen3-30B-A3B-Thinking	0.00	1.26	3.00	10.11	16.51	7.53	17.15
WideSeek-8B-RL	0.00	1.09 (+0.56)	2.59 (+1.08)	10.86 (+3.49)	16.61 (+3.90)	9.57 (×2.29)	41.09 (×4.33)
WideSeek-8B-SFT	0.14	1.74 (+1.21)	3.66 (+2.15)	11.35 (+3.98)	18.92 (+6.21)	13.16 (×3.15)	121.98 (×12.84)
WideSeek-8B-SFT-RL	0.00	**1.95 (+1.42)**	**3.88 (+2.37)**	**12.87 (+5.50)**	**19.73 (+7.02)**	**26.60 (×6.36)**	**273.75 (×28.82)**

term $\rho_{g,u,k} = \frac{\pi_\theta(a_{u,k}|s_u, a_{u,<k})}{\pi_{\theta_{old}}(a_{u,k}|s_{u,k}, a_{u,<k})}$ represents the importance sampling ratio for the k-th token in the u-th action. The group-relative advantage $\hat{A}_g$ is computed using the global reward $R(\boldsymbol{\mathcal{T}}_g)$ as $\hat{A}_g = (R(\boldsymbol{\mathcal{T}}_g) - \mu_R)/\sigma_R$, where μ_R and σ_R are the mean and standard deviation of rewards within the sampled group, respectively.

4 Experiment

4.1 Setting

We test proprietary models and open-sourced models on WideSeekBench. We use Qwen3-8B (Yang et al., 2025) as the base for agent optimization. For more training settings, we show in Appendix C.3. To test the generalization to the Deep Research dataset, we test the agent on Browsecomp-plus (Chen et al., 2025). We also show a WideSeek trajectory example in Appendix C.5 for better understanding of the agent behavior.

Table 3: BrowseComp-Plus accuracy (%). We evaluate generalization from wide research training to deep research tasks. The WideSeek scaffold alone (no training) outperforms larger ReAct models, and RL optimization yields the best result (26.42%), surpassing Qwen3-32B ReAct by 15.70 points.

Model	Scaffold	Acc
Gemini-2.5-Pro	ReAct	**29.52**
GPT-OSS-120B-Low	ReAct	25.54
DeepSeek-R1-0528	ReAct	16.39
Search-R1-32B	ReAct	11.08
Qwen3-32B	ReAct	10.72
Qwen3-8B	ReAct	6.02
ReAct-8B-SFT	ReAct	15.18
ReAct-8B-RL	ReAct	16.63
Qwen3-30B-A3B	WideSeek	14.82
Qwen3-8B	WideSeek	14.22
WideSeek-8B-SFT	WideSeek	23.61
WideSeek-8B-SFT-RL	WideSeek	23.61
WideSeek-8B-RL	WideSeek	**26.42 (+12.20)**

4.2 Main Results

Scalability Gaps. As shown in Table 2, even the strongest proprietary model GPT-5.2 achieves only 21.03 Mean@4 Item-F1 on WideSeekBench, underscoring the inherent difficulty of information seeking at scale. Proprietary models do spontaneously fork more sub-agents (e.g., DeepSeek-v3.2: 31.25) and execute substantially more tool calls (e.g., GPT-5.2: 408), yet they fail to coordinate them effectively without specialized optimization, revealing that merely scaling agent count is insufficient for this task.

Efficacy of WideSeek Optimization. We decompose gains into three orthogonal sources. (1) *Scaffold Gain*: The WideSeek scaffold alone (no training) lifts Qwen3-8B Item-F1 from 28.44 to 39.13 (+10.69) on WideSearch and accuracy from 6.02 to 14.22 (+8.20) on BrowseComp-Plus, demonstrating the architectural advantage of dynamic multi-agent decomposition. (2) *SFT Gain*: Trajectory distillation teaches effective scaffold utilization. Specifically, WideSeek-8B-SFT scales tool usage by 12.84× and sub-agents by 3.15× on WideSeekBench (Item-F1: 7.37→11.35), with consistent gains on WideSearch (+5.72) and BrowseComp-Plus (+9.39). (3) *RL Gain*: End-to-end RL refines coordination beyond imitation. WideSeek-8B-SFT-RL achieves the best WideSeekBench (Item-F1 12.87) and WideSearch (Item-F1 45.44) performance, scaling tool calls by 28.82×. Notably, RL from scratch (WideSeek-8B-RL) ex-

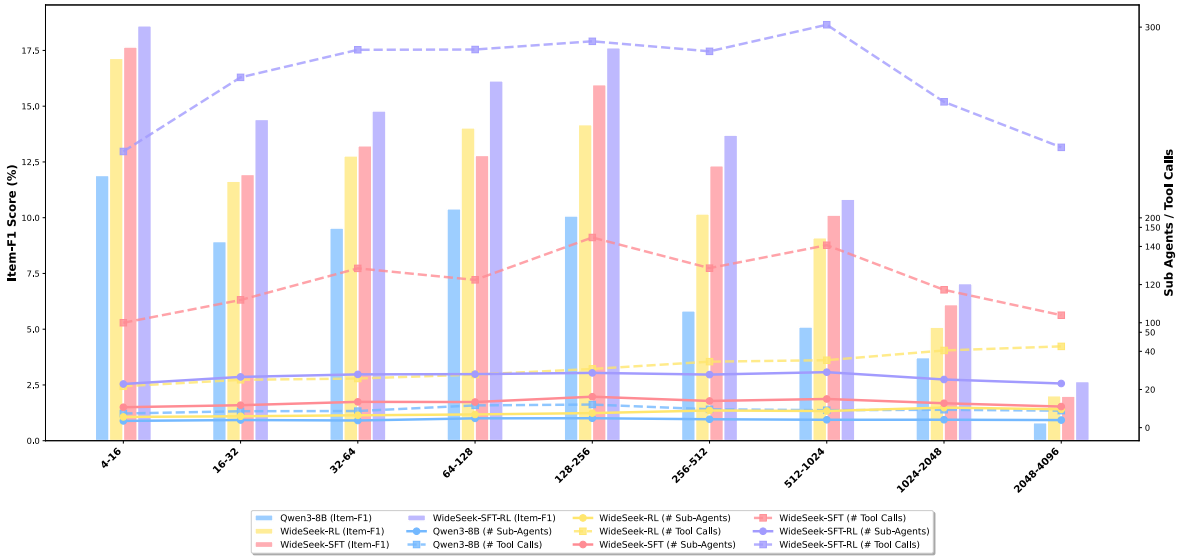

Figure 4: Performance and resource scaling with target information volume on WideSeekBench. Left: Item-F1 across volume buckets. Right: sub-agent and tool call counts for each model variant.

Table 4: Out-of-distribution evaluation on WideSearch (4 runs per task). Models are trained in our simulated environment and tested on real-world web search. WideSeek-8B-SFT-RL surpasses the much larger DeepSeek-R1-0528 on all metrics, demonstrating effective sim-to-real transfer of multi-agent coordination strategies.

Model	Scaffold	Row F1 (%)		Item F1 (%)	
		Mean@4	Max@4	Mean@4	Max@4
DeepSeek-R1-0528	ReAct	20.70	31.70	41.30	55.10
Qwen3-8B	ReAct	6.26	12.45	28.44	40.14
ReAct-8B-SFT	ReAct	10.40	20.15	34.09	48.04
ReAct-8B-RL	ReAct	11.46	20.42	37.89	50.10
Qwen3-8B	WideSeek	13.70	22.37	39.13	50.29
WideSeek-8B-SFT	WideSeek	19.41	31.18	44.85	57.13
WideSeek-8B-SFT-RL	WideSeek	**19.66** (+13.40)	**33.03** (+20.58)	**45.44** (+17.00)	**58.81** (+18.67)
WideSeek-8B-RL	WideSeek	16.15	27.30	42.03	55.39

cels on BrowseComp-Plus (26.42%), suggesting that bypassing SFT avoids teacher refusal tendencies and enables more aggressive exploration. Performance gains remain bounded by the 8B parameter size. Figure 9 confirms that the rising reward correlates with increasing tool calls, showing the model discovers broader search as better policy.

WideResearch to DeepResearch Generalization. We evaluate on BrowseComp-Plus (Table 3) to test whether wide research training transfers to deep research tasks. Without any training, the WideSeek scaffold alone enables Qwen3-8B (14.22%) to outperform the much larger Qwen3-32B ReAct (10.72%), confirming that parallel decomposition effectively mitigates context management burden even in depth-oriented settings. After optimization, WideSeek-8B-RL achieves 26.42% (+12.20 over base), demonstrating that coordination strategies learned for breadth generalize effectively to deep research.

Simulated Environment to Real-World Generalization. We evaluate on WideSearch (Wong et al., 2025), an independent benchmark operating over the live web (Table 4). Without training, the WideSeek scaffold lifts Qwen3-8B from 28.44 to 39.13 Item-F1, showing that the architectural design alone provides meaningful gains in real-world settings. After optimization, WideSeek-8B-SFT-RL achieves 45.44 Mean@4 Item-F1, surpassing both ReAct-8B-RL (37.89) and the much larger DeepSeek-R1-0528 (41.30) on all metrics. This validates that multi-agent coordination strategies learned in our simulated environment transfer effectively to the open web environment.

We further provide a fine-grained token usage analysis in Appendix C.4 to demonstrate the context management benefits.

5 Analysis

WideSeekBench enables granular evaluation through multi-dimensional task classification. Our results confirm that multi-agent RL consistently enhances performance across all dimensions.

Volume of Target Information. We categorize tasks by the total cell count in the ground truth table, from [4, 16] to [2048, 4096]. As shown in Figure 4, the hierarchy WideSeek-8B-SFT-RL > SFT > RL holds across all intervals. In the lower range ([4, 128]), gaps are minimal; beyond [128,

4096], performance degrades sharply, confirming that massive-scale retrieval remains challenging. Notably, at extreme volumes ([2048, 4096]), SFT-initialized models exhibit "early stopping"—a counter-intuitive drop in tool calls likely inherited from teacher models that assess high-volume tasks as infeasible. In contrast, WideSeek-8B-RL (trained from scratch) scales tool usage positively with volume, having learned to deploy more extensive search to maximize recall without this bias.

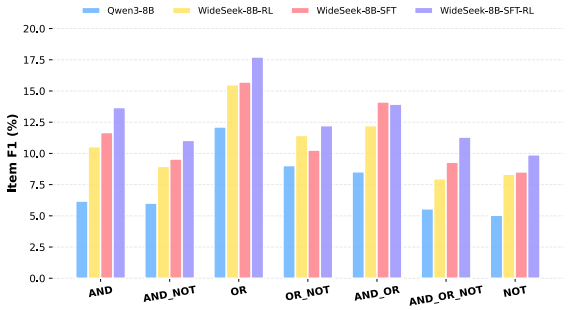

Figure 5: Item-F1 across 7 logical constraint types on WideSeekBench for all model variants.

Constraint Type. Tasks are classified into seven logical constraint types (see Appendix B.3). As shown in Figure 5, `OR` constraints yield the highest performance, as disjunctive logic aligns naturally with parallel sub-agent execution. Conversely, `NOT` yields the lowest scores, and compounding with negation (e.g., `OR_NOT`) invariably degrades performance further. This indicates that set-difference operations, which require agents to exclude specific entity sets, constitute a distinct reasoning bottleneck.

Figure 6: Item-F1 across 18 domains on WideSeekBench for all model variants.

Domain. We evaluate across 18 domains (Figure 6). The hierarchy SFT-RL > SFT > RL holds universally, validating that our optimization yields robust improvements regardless of topic. Domain sensitivity is consistent: *Infrastructure* achieves the highest scores while *Education & Academia* remains challenging.

6 Related Work

6.1 Data Synthesis for Search Agent

Synthetic data has become central to search agent training (Li et al., 2025; Tao et al., 2025; Team et al., 2025). Existing methods focus on search depth—information-driven approaches (Wu et al., 2025b,a; Li et al., 2025) simulate navigation paths for single-answer retrieval, and formalization-driven methods (Xin et al., 2024; Ren et al., 2025; Xia et al., 2025; Tao et al., 2025) enhance logical rigor but remain depth-oriented. We instead formalize search *width* via Knowledge Graph entity clusters and set operators to regulate breadth and constraint complexity.

6.2 LLM-based Multi-Agent Reinforcement Learning

Multi-agent systems have shifted from static role-based designs (Qian et al., 2024; Hong et al., 2023) toward RL-optimized coordination (Liu et al., 2025; Hu et al., 2025). M-GRPO (Hong et al., 2025) and Fold-GRPO (Sun et al., 2025) use branch-return paradigms but fork a fixed number of sub-agents per step; Kimi K2.5 (Team et al., 2026) optimizes the orchestrator while keeping sub-agents static. WideSeek performs end-to-end RL jointly optimizing the planner and dynamically forked sub-agents. Concurrent works FlashResearch (Nie et al., 2026) and Flash-Searcher (Qin et al., 2026) exploit parallelism via static tree/DAG structures but lack end-to-end multi-agent optimization.

7 Conclusion

To address the paradigm shift from Deep to Wide Research, we introduce WideSeekBench to formalize the General Broad Information Seeking (GBIS) task. We construct it via a rigorous multi-phase data pipeline that mines intersected world knowledge from KGs. We propose WideSeek, a dynamic hierarchical multi-agent architecture optimized via an end-to-end reinforcement learning framework. Our results demonstrate that WideSeek effectively leverages agent scaling to solve complex, parallel retrieval tasks, significantly advancing Wide Research capabilities.

Limitations

First, all experiments are conducted with an 8B-parameter base model due to computational constraints. While this demonstrates the effectiveness of our optimization pipeline at small scale, the scaling behavior to larger models remains unexplored. We leave it as the future work (we currently have not enough resources to conduct large-scale experiments). Second, our simulated search environment uses a static Wikipedia snapshot, which does not fully capture the dynamic nature of the live web. But our WideSearch evaluation demonstrates effective sim-to-real transfer, which is still a strong validation of our approach.

References

Mohamad Abou Ali, Fadi Dornaika, and Jinan Charafeddine. 2025. Agentic ai: a comprehensive survey of architectures, applications, and future directions. *Artificial Intelligence Review*, 59(1).

Hannah Bast and Björn Buchhold. 2017. Qlever: A query engine for efficient sparql+text search. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, CIKM '17, page 647–656, New York, NY, USA. Association for Computing Machinery.

Zijian Chen, Xueguang Ma, Shengyao Zhuang, Ping Nie, Kai Zou, Andrew Liu, Joshua Green, Kshama Patel, Ruoxi Meng, Mingyi Su, Sahel Sharifymoghaddam, Yanxi Li, Haoran Hong, Xinyu Shi, Xuye Liu, Nandan Thakur, Crystina Zhang, Luyu Gao, Wenhu Chen, and Jimmy Lin. 2025. Browsecomp-plus: A more fair and transparent evaluation benchmark of deep-research agent. *Preprint*, arXiv:2508.06600.

Haoyang Hong, Jiajun Yin, Yuan Wang, Jingnan Liu, Zhe Chen, Ailing Yu, Ji Li, Zhiling Ye, Hansong Xiao, Yefei Chen, Hualei Zhou, Yun Yue, Minghui Yang, Chunxiao Guo, Junwei Liu, Peng Wei, and Jinjie Gu. 2025. Multi-agent deep research: Training multi-agent systems with m-grpo. *Preprint*, arXiv:2511.13288.

Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, and 1 others. 2023. Metagpt: Meta programming for a multi-agent collaborative framework. In *The twelfth international conference on learning representations*.

Mengkang Hu, Yuhang Zhou, Wendong Fan, Yuzhou Nie, Bowei Xia, Tao Sun, Ziyu Ye, Zhaoxuan Jin, Yingru Li, Qiguang Chen, and 1 others. 2025. Owl: Optimized workforce learning for general multi-agent assistance in real-world task automation. *arXiv preprint arXiv:2505.23885*.

Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. 2024. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*.

Tian Lan, Bin Zhu, Qianghuai Jia, Junyang Ren, Haijun Li, Longyue Wang, Zhao Xu, Weihua Luo, and Kaifu Zhang. 2025. Deepwidesearch: Benchmarking depth and width in agentic information seeking. *Preprint*, arXiv:2510.20168.

Fangyu Lei, Jinxiang Meng, Yiming Huang, Junjie Zhao, Yitong Zhang, Jianwen Luo, Xin Zou, Ruiyi Yang, Wenbo Shi, Yan Gao, Shizhu He, Zuo Wang, Qian Liu, Yang Wang, Ke Wang, Jun Zhao, and Kang Liu. 2025. Dacomp: Benchmarking data agents across the full data intelligence lifecycle. *Preprint*, arXiv:2512.04324.

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

Shuo Liu, Tianle Chen, Zeyu Liang, Xueguang Lyu, and Christopher Amato. 2025. Llm collaboration with multi-agent reinforcement learning. *arXiv preprint arXiv:2508.04652*.

Rui Lu, Zhenyu Hou, Zihan Wang, Hanchen Zhang, Xiao Liu, Yujiang Li, Shi Feng, Jie Tang, and Yuxiao Dong. 2025. Deepdive: Advancing deep search agents with knowledge graphs and multi-turn rl. *Preprint*, arXiv:2509.10446.

Xufang Luo, Yuge Zhang, Zhiyuan He, Zilong Wang, Siyun Zhao, Dongsheng Li, Luna K. Qiu, and Yuqing Yang. 2025. Agent lightning: Train any ai agents with reinforcement learning. *Preprint*, arXiv:2508.03680.

Manus. 2025. Introducing wide research.

Lunyiu Nie, Nedim Lipka, Ryan A. Rossi, and Swarat Chaudhuri. 2026. Efficient tree-structured deep research with adaptive resource allocation. *Preprint*, arXiv:2510.05145.

Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, and 1 others. 2024. Chatdev: Communicative agents for software development. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15174–15186.

Tianrui Qin, Qianben Chen, Sinuo Wang, He Xing, King Zhu, He Zhu, Dingfeng Shi, Xinxin Liu, Ge Zhang, Jiaheng Liu, Xitong Gao, Yuchen Eleanor Jiang, and Wangchunshu Zhou. 2026. Flash-searcher: Fast and

effective web agents via DAG-based parallel execution. In *The Fourteenth International Conference on Learning Representations*.

ZZ Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanjia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, and 1 others. 2025. Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. *arXiv preprint arXiv:2504.21801*.

Aymeric Roucher, Albert Villanova del Moral, Thomas Wolf, Leandro von Werra, and Erik Kaunismäki. 2025. 'smolagents': a smol library to build great agentic systems. `https://github.com/huggingface/smolagents`.

Lukas Schmelzeisen, Corina Dima, and Steffen Staab. 2021. Wikidated 1.0: An evolving knowledge graph dataset of wikidata's revision history. *Preprint*, arXiv:2112.05003.

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

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

Zhengliang Shi, Yiqun Chen, Haitao Li, Weiwei Sun, Shiyu Ni, Yougang Lyu, Run-Ze Fan, Bowen Jin, Yixuan Weng, Minjun Zhu, Qiujie Xie, Xinyu Guo, Qu Yang, Jiayi Wu, Jujia Zhao, Xiaqiang Tang, Xinbei Ma, Cunxiang Wang, Jiaxin Mao, and 7 others. 2025. Deep research: A systematic survey. *Preprint*, arXiv:2512.02038.

Weiwei Sun, Miao Lu, Zhan Ling, Kang Liu, Xuesong Yao, Yiming Yang, and Jiecao Chen. 2025. Scaling long-horizon llm agent via context-folding. *Preprint*, arXiv:2510.11967.

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

Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, S. H. Cai, Yuan Cao, Y. Charles, H. S. Che, Cheng Chen, Guanduo Chen, Huarong Chen, Jia Chen, Jiahao Chen, Jianlong Chen, Jun Chen, Kefan Chen, Liang Chen, Ruijue Chen, Xinhao Chen, and 307 others. 2026. Kimi k2.5: Visual agentic intelligence. *Preprint*, arXiv:2602.02276.

Tongyi DeepResearch Team, Baixuan Li, Bo Zhang, Dingchu Zhang, Fei Huang, Guangyu Li, Guoxin Chen, Huifeng Yin, Jialong Wu, Jingren Zhou, and 1 others. 2025. Tongyi deepresearch technical report. *arXiv preprint arXiv:2510.24701*.

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

Ryan Wong, Jiawei Wang, Junjie Zhao, Li Chen, Yan Gao, Long Zhang, Xuan Zhou, Zuo Wang, Kai Xiang, Ge Zhang, Wenhao Huang, Yang Wang, and Ke Wang. 2025. Widesearch: Benchmarking agentic broad info-seeking. *Preprint*, arXiv:2508.07999.

Jialong Wu, Baixuan Li, Runnan Fang, Wenbiao Yin, Liwen Zhang, Zhengwei Tao, Dingchu Zhang, Zekun Xi, Gang Fu, Yong Jiang, and 1 others. 2025a. Webdancer: Towards autonomous information seeking agency. *arXiv preprint arXiv:2505.22648*.

Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, and 1 others. 2025b. Webwalker: Benchmarking llms in web traversal. *arXiv preprint arXiv:2501.07572*.

Tianle Xia, Liang Ding, Guojia Wan, Yibing Zhan, Bo Du, and Dacheng Tao. 2025. Improving complex reasoning over knowledge graph with logic-aware curriculum tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 12881–12889.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. 2024. OSWorld: Benchmarking multimodal agents for open-ended tasks in real computer environments. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Huajian Xin, Daya Guo, Zhihong Shao, Zhizhou Ren, Qihao Zhu, Bo Liu, Chong Ruan, Wenda Li, and Xiaodan Liang. 2024. Deepseek-prover: Advancing theorem proving in llms through large-scale synthetic data. *arXiv preprint arXiv:2405.14333*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Shunyu Yao. 2025. The second half.

A AI Usage Disclosure

We use LLM for paper writing to check grammar and boost the clarity.

B The Details of WideSeekBench

B.1 Knowledge Graph Source and Infrastructure

We ingest the Wikidata Truthy Dump (October 1, 2025) into a local QLever (Bast and Buchhold, 2017) SPARQL engine to support efficient, rate-limit-free execution of complex SPARQL queries over the full knowledge graph.

B.2 Seed Constraint Construction

We construct a diverse set of seed entities to serve as the semantic basis for downstream constraint construction and task synthesis.

Domain Taxonomy. We define 18 high-level domains (e.g., *Computer Science, Life Sciences, Governance*). Each domain is mapped to a set of Wikidata classes, which are treated as domain-specific sub-domains. In total, this mapping yields 200 sub-domains across all domains. These sub-domains jointly define a controlled search scope $\mathcal{S}_{sub-domain}$ (refer to Appendix B.5 for details).

Retrieval and Ranking. For each sub-domain, we identify 80 informative seed entities from the knowledge base $\mathcal{K}$ using a three-stage SPARQL-based workflow. (1) *Retrieval:* Given a sub-domain class, we retrieve a candidate entity set $\mathbf{E}_{cand}$ by recursively querying the class and all its subclasses via the transitive closure of the `wdt:P279` (subclass of) relation (Listing 1). (2) *Ranking:* Each candidate entity $e \in \mathbf{E}_{cand}$ is ranked by its information density, approximated by the number of outgoing RDF triples associated with e (Listing 2). Entities with higher information density are preferred, as they support the construction of richer constraints and attribute schemas. (3) *Filtering:* We remove non-entity artifacts and structurally uninformative entries, including entities whose labels begin with `"List of"` or `"Category:"`. The remaining entities constitute the seed entity set $\mathbf{E}_{seed}$.

Listing 1: Candidate entity retrieval via recursive subclass matching, where `wdt:P31` and `wdt:P279` represent the 'instance of' and 'subclass of' relations in Wikidata, respectively.

```
SELECT DISTINCT ?entity WHERE {
  ?entity (wdt:P31/wdt:P279*) wd:TARGET_ID .
}
```

Listing 2: Ranking entities by information density (triple count).

```
SELECT ?entity ?label (COUNT(?p) AS ?count) WHERE {
  VALUES ?entity { wd:Q_CANDIDATE_1 ... }
  ?entity ?p ?o .
  OPTIONAL { ?entity rdfs:label ?label . FILTER(LANG(?label) = "en") }
}
GROUP BY ?entity ?label
ORDER BY DESC(?count)
```

B.3 Logical Composition and Task Synthesis

We describe the procedures for composing atomic constraints into executable queries, executing and validating the resulting retrievals, and constructing bounded tables. For each seed entity, we generate up to 200 composite constraints. To control redundancy and dataset balance, each seed contributes at most 4 validated tables.

Query Formulation. Given a sampled seed entity $e_{\text{seed}} \in \mathbf{E}_{\text{seed}}$ and its associated relations $\mathcal{R}_{\text{seed}} = \{(e_{\text{seed}}, p, v)\}$, retrieved from $\mathcal{K}$ via property-seeking SPARQL queries, we define the associated atomic constraint set $\mathcal{C}_{atom}^{(e_{\text{seed}})} = \{(p, v)\}$. We then sample atomic constraints $c \in \mathcal{C}_{atom}^{(e_{\text{seed}})}$ and compose them into composite SPARQL filters using seven predefined logical patterns (Table 5), yielding a composite constraint Φ. Apart from the domain constraint, each composite constraint Φ is required to contain at least 1 and at most 8 atomic constraints.

Table 5: Logical patterns in WideSeekBench. $\mathcal{D}$ denotes the domain constraint.

Pattern	Prob.	Formulation $\Phi(e)$	SPARQL Implementation
AND	20%	$\mathcal{D}(e) \wedge (\bigwedge_i c_i(e))$	`SELECT ?item WHERE { ?item wdt:P31 wd:Q_dom . ?item wdt:P1 wd:Q1 . ?item wdt:P2 wd:Q2 . }`
OR	20%	$\mathcal{D}(e) \wedge (\bigvee_i c_i(e))$	`SELECT ?item WHERE { ?item wdt:P31 wd:Q_dom . { { ?item wdt:P1 wd:Q1 } UNION { ?item wdt:P2 wd:Q2 } } }`
NOT	15%	$\mathcal{D}(e) \wedge c_{base}(e) \wedge \neg(\bigvee_i c_{ex_i}(e))$	`SELECT ?item WHERE { ?item wdt:P31 wd:Q_dom . ?item wdt:P_base wd:Q_base . FILTER NOT EXISTS { { ?item wdt:P_ex1 wd:Q_ex1 } UNION { ?item wdt:P_ex2 wd: Q_ex2 } } }`
AND_OR	15%	$\mathcal{D}(e) \wedge [(\bigwedge_i c_i(e)) \vee (\bigwedge_j c_j(e))]$	`SELECT ?item WHERE { ?item wdt:P31 wd:Q_dom . { { ?item wdt:P1 wd:Q1 . ?item wdt:P2 wd:Q2 } UNION { ?item wdt:P3 wd:Q3 . ?item wdt:P4 wd:Q4 } } }`
AND_NOT	15%	$\mathcal{D}(e) \wedge (\bigwedge_i c_{in_i}(e)) \wedge \neg(\bigvee_j c_{ex_j}(e))$	`SELECT ?item WHERE { ?item wdt:P31 wd:Q_dom . ?item wdt:P1 wd:Q1 . ?item wdt:P2 wd:Q2 . FILTER NOT EXISTS { { ?item wdt:P_ex1 wd:Q_ex1 } UNION { ?item wdt:P_ex2 wd: Q_ex2 } } }`
OR_NOT	10%	$\mathcal{D}(e) \wedge (\bigvee_i c_{in_i}(e)) \wedge \neg(\bigvee_j c_{ex_j}(e))$	`SELECT ?item WHERE { ?item wdt:P31 wd:Q_dom . { { ?item wdt:P1 wd:Q1 } UNION { ?item wdt:P2 wd:Q2 } } FILTER NOT EXISTS { { ?item wdt:P_ex1 wd:Q_ex1 } UNION { ?item wdt:P_ex2 wd: Q_ex2 } } }`
AND_OR_NOT	5%	$\mathcal{D}(e) \wedge [(\bigwedge_i c_i(e)) \vee (\bigwedge_j c_j(e))] \wedge \neg(\bigvee_k c_{ex_k}(e))$	`SELECT ?item WHERE { ?item wdt:P31 wd:Q_dom . { { ?item wdt:P1 wd:Q1 . ?item wdt:P2 wd:Q2 } UNION { ?item wdt:P3 wd:Q3 . ?item wdt:P4 wd:Q4 } } FILTER NOT EXISTS { { ?item wdt:P_ex1 wd:Q_ex1 } UNION { ?item wdt:P_ex2 wd: Q_ex2 } } }`

Execution and Verification. Each composite filter Φ is executed against the knowledge base $\mathcal{K}$ to retrieve a candidate entity set $\mathbf{E}^*$. We restrict the cardinality of the candidate entity set $\mathbf{E}^*$ to the interval $[1, 1024]$. As shown in Listing 3, a verification step enforces this constraint prior to attribute retrieval, discarding any queries where $|\mathbf{E}^*|$ falls outside the bound.

Table Construction and Quality Control. Given the validated entity set $\mathbf{E}^*$, we first collect a candidate attribute set $\mathcal{A}_{cand} = \bigcup_{e \in \mathbf{E}^*} \text{Attributes}(e)$, and dynamically select target attributes $\mathcal{A} \subset \mathcal{A}_{cand}$ by retaining only those with at least 50% coverage across entities and sufficient value diversity (Listing 4). Next, we compute the potential table size $N_{\text{cells}} = |\mathbf{E}^*| \times |\mathcal{A}|$, and retain tasks satisfying $N_{\text{cells}} \in [8, 8192]$. Entities that fail to resolve to valid labels are removed, resulting in the cleaned entity set $\mathbf{E}_{\text{clean}}$. Finally, we perform batch SPARQL queries to retrieve all cell values (Listing 5) and populate the table $\mathbf{T}^*$. To avoid redundancy, we further deduplicate tables by discarding those with identical entity sets and attribute schemas, retaining only one representative table per equivalence class.

Listing 3: Pre-flight cardinality check.

```
SELECT (COUNT(DISTINCT ?item) AS ?count) WHERE {
  # Synthesized Logical Constraints (see Table 4)
  ?item wdt:P31 wd:Q_domain .
  ...
}
```

Listing 4: Attribute frequency analysis.

Table 6: Property ID Blacklist.

Property ID	Label
P31	Instance of
P106	Occupation
P108	Employer
P248	Stated in
P279	Subclass of
P361	Part of
P373	Commons category
P460	Said to be the same as
P527	Has part
P646	Freebase ID
P910	Topic's main category
P1001	Applies to jurisdiction
P1343	Described by source
P1709	Equivalent class
P1754	Category related to list
P1889	Different from
P2671	Google Knowledge Graph ID
P3876	Category for alumni of educational institution
P6104	Maintained by WikiProject

```
SELECT ?prop (COUNT(DISTINCT ?item) AS ?cnt) WHERE {
  VALUES ?item { wd:Q_sample1 ... }
  ?item ?prop ?value .
  FILTER(STRSTARTS(STR(?prop),
    "http://www.wikidata.org/prop/direct/"))
} GROUP BY ?prop
```

Listing 5: Batch value retrieval.

```
SELECT ?item ?prop ?value ?valueLabel WHERE {
  VALUES ?item { wd:Q_e1 ... }
  VALUES ?directProp { wdt:P1 ... }
  ?item ?directProp ?value .
  ?prop wikibase:directClaim ?directProp .
  OPTIONAL {
    ?value rdfs:label ?valueLabel .
    FILTER(LANG(?valueLabel) = "en")
  }
}
```

B.4 Agent Task Synthesis and Multi-Stage Filtering

We implement a cyclic generation-verification pipeline to transform structured logical filters Φ into diverse, human-like search tasks Q, followed by a rigorous quality assurance protocol. In this subsection, all LLM-based operations are powered by GPT-5.

Self-Refining Query Synthesis. The transformation process employs a dual-model architecture to ensure both linguistic diversity and logical fidelity. First, raw constraints are mapped into a structured f-string template (e.g., `"Find all {sub-domain} that {prop} is {val}..."`). A generator model M_{gen} then transforms this template into natural language using a style randomization protocol, sampling a syntactic mode $s \sim U(1, 10)$ from a predefined set (*Action*, *Question*, *Imperative*, *Need*, *Context*, *Interest*, *Description*, *Casual*, *Professional* and *Task*) for each task. To ensure semantic accuracy, a critic model M_{ver} extracts the logic $\hat{\Phi}$ back from the generated query Q and performs a constraint-by-constraint equivalence check $S(\Phi, \hat{\Phi})$. The verifier rigorously compares entity preservation, operator logic ($\wedge, \vee, \neg$), filtering scope, and output schema consistency. Any discrepancy triggers a feedback loop with specific error correction instructions, capped at $k = 5$ iterations.

Data-Driven Rubric Synthesis. We leverage an LLM to synthesize adaptive evaluation criteria R_j by analyzing the data distribution of each ground truth column $\mathbf{T}^*_{:,j}$. Unlike rigid string matching, the model generates semantic compliance standards tailored to the specific data type: (1) *Entities* explicitly accept

aliases and naming variations; (2) *Dates* enforce semantic exactness regardless of format; (3) *Numerics* require value equality within defined tolerances; and (4) *Sets* enforce equality independent of item order.

Quality Assurance Protocol. We apply a three-tier filtering mechanism. (1) *Rule-Based Filtering* discards tasks with sparse ground truth ($> 50\%$ empty cells) or weak web grounding where target entities lack verifiable search API hits, as determined by their English sitelink counts in Wikidata, where entities with zero English sitelinks are strictly filtered out. (2) *LLM-Based Filtering* employs a judge model to evaluate tasks on a 5-point scale across *Human-Likeness*, *Solvability*, *Common Sense*, *Temporal Stability*, and *Rubric Rationality*, a violation in any category results in immediate rejection.(3) *Human Verification* removes subtle semantic irrationalities (e.g., logical contradictions) that automated filters might overlook.

Prompt: LLM-Based Filtering

```
You are an expert Data Quality Auditor for a Search Agent Benchmark.
Your goal is to validate synthesized widesearch tasks to ensure they are high-quality, solvable, and
      fairly evaluated.

## Input Data provided for each request:
1. Task Info: The natural language user query and basic metadata.
2. Rubrics: The evaluation criteria for each column (preprocess, metric, criterion).
3. Ground Truth Stats & Sample: Statistics of the expected answer table and a preview of the data.
4. Intermediate Logic: The SPARQL query and constraints used to generate the task.

## Evaluation Criteria (Pass/Fail)

Analyze the task against the following 5 dimensions. If it fails on ANY dimension, mark as INVALID.

### 1. Human-Likeness & Phrasing
* Goal: Ensure the query sounds like a real human request.
* Reject if phrasing is unnatural, overly pedantic, or a raw SPARQL translation.
    * Fail: "Find entities where mass is exactly 70.53 kg."
    * Fail: "Find items with property P31 equal to Q5." (Leaked IDs).
    * Pass: "Find all public hospitals in Portugal founded before 1950."

### 2. Task Solvability
* Goal: Ensure the task is logically possible.
* Reject if logic contains contradictions (e.g., "born > 2000 AND died < 1999").
* Reject if it asks for private/internal info not publicly available.

### 3. Answer Common Sense
* Goal: Detect "Hallucinated" or "Incomplete" Ground Truth.
* Reject if GT clearly contradicts general world knowledge.
    * Fail: "All US States" -> GT has only 5 rows (Recall too low).
    * Fail: "Current UK Monarchs" -> GT has 10 rows (Recall too high/Wrong).
    * Fail: GT has >50% empty cells for key attributes.

### 4. Temporal Stability & Wiki-Fit
* Goal: Ensure facts are static and suitable for Wikipedia search.
* Reject if query relies on "Current/Now" without a time anchor (e.g., "Current CEO") UNLESS the GT
     explicitly handles specific dates.
* Reject if attributes are purely database metadata (e.g., "Wikidata QID", "GND ID") unlikely to be
     found in Wikipedia text.

### 5. Rubric Rationality (Criterion Quality Check)
* Goal: Ensure the criterion explicitly instructs the judge to handle the variability of web data (
     formats, aliases, precision).
* Reject if the criterion enforces rigid constraints that are unrealistic for web extraction.
    * Fail (False Strictness on Non-Strings): The criterion demands "exact string match" or "no
         formatting differences" for Dates, Coordinates, or Floating Point Numbers.
        * Bad: "Accept only exact string match." (for a Date like '2020-01-01').
        * Good: "Accept dates representing the same day, regardless of format (e.g., 'Jan 1st,
             2020')."
    * Fail (Unrealistic Precision): The criterion for Coordinates or Statistics requires matching to
          extreme precision (e.g., >4 decimal places) which rarely appears in web text.
        * Bad: "Reject if coordinate differs by 0.000001."
        * Good: "Accept if coordinates are geographically close (within ~1km)."
    * Fail (Missing Alias Logic): The criterion for Entities (e.g., Universities, Journals) fails to
          mention allowing "Aliases" or "Common Abbreviations".
        * Bad: "Accept only 'University of California, Los Angeles'."
        * Good: "Accept 'UCLA' or 'Univ. of Calif. LA' as valid matches."
    * Fail (Vague Instructions): The criterion is too generic to ensure consistency.
        * Bad: "Check if the answer is correct." (Too open-ended).
        * Good: "Accept exact matches or minor spelling variations; reject completely different
             entities."

## Output Format
Output a JSON object ONLY:
```json
{
```
```

```
  "status": "VALID" | "INVALID",
  "reasoning": "Concise explanation...",
  "flags": {
    "unnatural_phrasing": boolean,
    "logically_unsolvable": boolean,
    "gt_contradicts_commonsense": boolean,
    "temporal_or_wiki_issue": boolean,
    "bad_rubric": boolean
  },
  "suggestion": "Optional fix"
}
```

B.5 WideSeekBench Statistics

Scale of Target Information Figure 7 depicts the scale of target information across diverse top-level domains. The dataset contains 4,436 training instances and 720 test instances, covering 18 domains. Tables 7 and 8 provide detailed distributions of subdomains. High-frequency categories in the training set include *film* (252), *video game* (197), and *airport* (176). The test set preserves a similar distribution (e.g., *film* 41, *video game* 38).

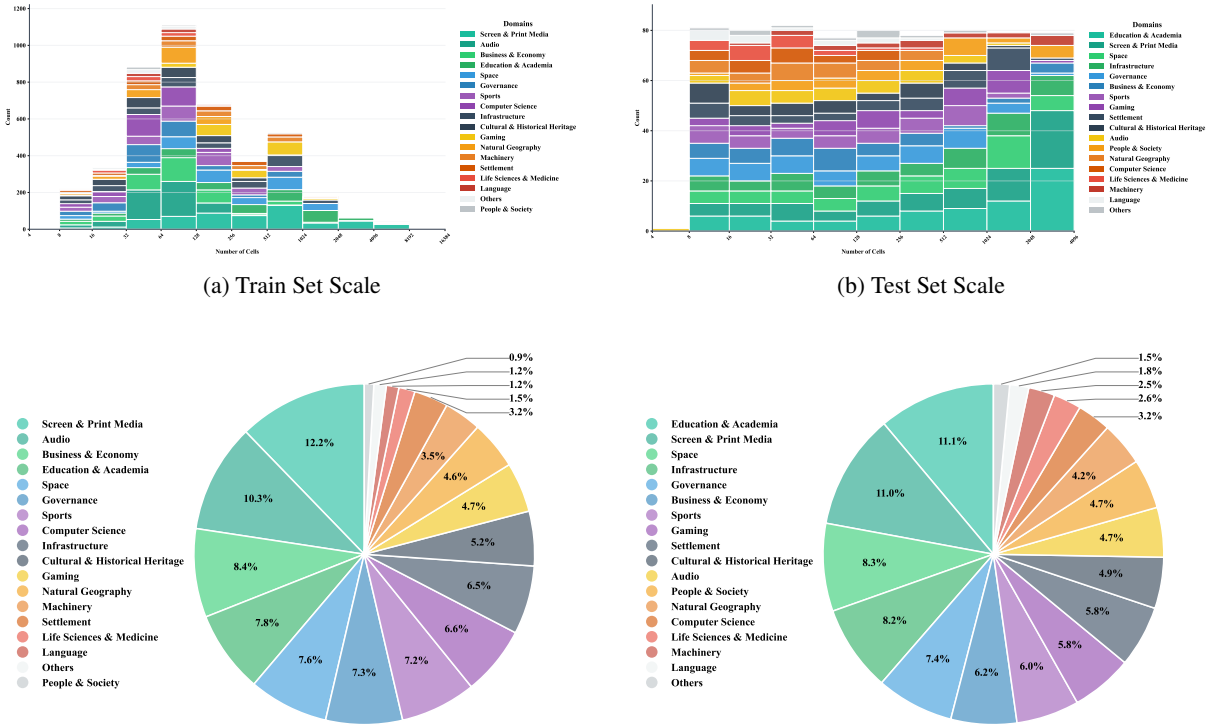

(a) Train Set Scale

(b) Test Set Scale

(c) Train Set Domain Distribution

(d) Test Set Domain Distribution

Figure 7: Comprehensive statistics of WideSeekBench. Top row (a, b) illustrates the scale of target information. Bottom row (c, d) depicts the distribution of domains across Training and Test sets.

Constraint Complexity Table 9 shows the distribution of logical patterns in the dataset, which directly reflects the distribution of constraints. The training set is dominated by single-type patterns, with pure conjunctions (`AND`) accounting for 37.8%, followed by `AND_NOT` (19.5%). The test set exhibits a more balanced distribution, with simple `AND` patterns reduced to 20.0% and complex composite patterns substantially increased. The most complex combination, `AND_OR_NOT`, constitutes 11.5% of the test set (compared to 5.1% in training), and other high-complexity patterns such as `AND_OR` and `OR_NOT` are also more evenly represented.

Domain Diversity Figure 8 shows the distribution of topics in WideSeekBench across training and test sets. Dominant domains such as *Screen & Print Media* and *Gaming* are represented by subdomains

Table 7: Training set domain and sub-domain summary.

Domain	Subdomain	Count
Screen & Print Media	film	252
	short film	109
	television series	92
	literary work	32
	television program	31
	publisher	7
	comics	6
	magazine	5
	episode	5
	periodical	2
	poem	2
Audio	single	146
	album	89
	rock band	86
	song	74
	musical group	40
	orchestra	18
	Musical Work	3
	concert	2
	rock	1
Business & Economy	bank	73
	public company	72
	goods	70
	manufactured good	56
	enterprise	18
	stock exchange	17
	business	17
	brewery	13
	brand	11
	company	9
	trademark	8
	currency	8
	farm	1
Education & Academia	university	143
	college	117
	scientific journal	25
	research institute	25
	academic journal	13
	educational institution	6
	laboratory	6
	school	6
	library	5
Space	airport	176
	space mission	47
	artificial satellite	34
	rocket launch	31
	asteroid	25
	aircraft model	10
	exoplanet	8
	variable star	3
Space	planetary nebula	1
Governance	political party	122
	charitable organization	28
	non-governmental organization	18
	polity	17
	government agency	15
	armed organization	14
	political organization	13
	battle	13
	international organization	12
	war	11
	former administrative territorial entity	10
	treaty	8
	legal case	7
	organization	6
	administrative territorial entity	5
	firearm	5
	public election	5
	executive branch	3
	conflict	3
	association	3
	legal norm	3
	crime	2
Sports	sporting event	76
	sports season	64
	competition stage	36
	association football club	28
	competition	22
	recurring sporting event edition	17
	recurring sporting event	16
	racing	12
	Olympic Games	11
	physical activity	7
	sports venue	5
	sports competition	5
	association football match	5
	tennis tournament	5
	sport	4
	nation at sport competition	4
	baseball player	1
	sports club	1
Computer Science	programming language	114
	operating system	94
	free software	35
	computer	28
	computer network protocol	10
	software	7
	database	3
Infrastructure	metro station	146
	dam	42
Infrastructure	railway station	26
	controlled-access highway	21
	lighthouse	19
	hotel	13
	road	10
	power station	5
	wind farm	3
	house	2
	building	1
	industrial building	1
Cultural & Historical Heritage	historical country	56
	tomb	40
	ceremony	15
	church building	15
	cultural heritage	14
	museum	12
	archaeological site	10
	heritage	10
	cultural property	10
	architectural heritage monument	9
	shrine	9
	heritage site	9
	temple	6
	location of worship	6
	funerary structure	4
	structure of worship	3
	chapel	1
	cemetery	1
Gaming	video game	197
	electronic game	12
	board game	1
Natural Geography	national park	38
	mountain	32
	island	27
	lake	18
	protected area	18
	canal	14
	park	11
	disaster	7
	glacier	7
	landform	6
	earthquake	6
	natural heritage	5
	hill	3
	valley	3
	forest	3
	nature reserve	3
	watercourse	2
	mineral	1
Machinery	machine	43
Machinery	vehicle	42
	vehicles and vehicle parts product	26
	tool	17
	equipment	11
	automobile model	9
	ship	5
	physical tool	4
Settlement	town	39
	municipality	32
	village	22
	city	20
	neighborhood	10
	district	9
	province	8
	human settlement	3
	region	1
Life Sciences & Medicine	taxon	17
	protein family	9
	hospital	8
	mammal	6
	Chordata	6
	fungi	6
	Vertebrata	4
	medication	3
	anatomical structure	2
	bird	2
	disease	1
	enzyme	1
	insect	1
	plant	1
	anomaly	1
Language	language	43
	languoid	9
	language variety	1
Others	visual artwork	13
	flag	11
	dish	8
	data	4
	artificial physical object	3
	physical process	2
	philosophy	2
	knowledge organization system	2
	sculpture	2
	communications media	2
	assembly	1
	chemical process	1
	disposable product	1
People & Society	human	30
	ethnic group	9
	occupation	1
Total: 4436		

Table 8: Test set domain and sub-domain summary.

Domain	Subdomain	Count
Education & Academia	college	34
	university	26
	research institute	5
	laboratory	4
	school	3
	academic journal	3
	educational institution	3
	scientific journal	1
	library	1
Screen & Print Media	film	41
	short film	16
	television series	12
	television program	6
	literary work	2
	photograph	1
	magazine	1
Space	airport	31
	space mission	12
	artificial satellite	7
	aircraft model	3
	asteroid	3
	exoplanet	2
	rocket launch	1
	astronomical object	1
Infrastructure	metro station	35
	railway station	7
	controlled-access highway	5
	hotel	4
	dam	3
	power station	2
	lighthouse	1
	wind farm	1
	road	1
Governance	political party	31
	armed organization	4
	polity	3
Governance	former administrative territorial entity	3
	charitable organization	2
	battle	2
	political organization	2
	government agency	2
	non-governmental organization	1
	conflict	1
	international organization	1
	war	1
Business & Economy	bank	9
	public company	8
	manufactured good	7
	stock exchange	5
	brewery	4
	goods	4
	enterprise	3
	company	2
	currency	1
	business	1
	brand	1
Sports	sports season	16
	association football club	6
	sporting event	6
	competition stage	3
	recurring sporting event	3
	association football match	2
	recurring sporting event edition	2
	sports competition	2
	racing	1
	Olympic Games	1
	sports venue	1
Gaming	video game	38
	board game	3
	electronic game	1
Settlement	town	19
	municipality	6
Settlement	human settlement	4
	village	3
	city	3
	province	2
	neighborhood	2
	region	2
	district	1
Cultural & Historical Heritage	historical country	11
	church building	6
	ceremony	4
	historical event	3
	architectural heritage monument	2
	cultural heritage	2
	tomb	2
	heritage site	2
	museum	1
	heritage	1
	cultural property	1
People & Society	human	31
	ethnic group	3
Audio	musical group	7
	album	6
	rock band	6
	single	6
	song	5
	orchestra	3
	musician	1
Natural Geography	national park	6
	island	5
	mountain	5
	canal	3
	hill	2
	landform	2
	protected area	2
	watercourse	1
	park	1
Natural Geography	natural heritage	1
	lake	1
	forest	1
Computer Science	operating system	11
	programming language	8
	free software	1
	computer network protocol	1
	computer	1
	archive	1
Life Sciences & Medicine	hospital	6
	protein family	4
	fungi	2
	symptom	2
	Vertebrata	1
	taxon	1
	plant	1
	mammal	1
	bird	1
Machinery	automobile model	6
	vehicle	3
	machine	3
	ship	2
	equipment	2
	vehicles and vehicle parts product	1
	tool	1
Language	language	8
	human language	4
	language variety	1
Others	visual artwork	4
	dish	2
	flag	2
	unit of measurement	1
	science	1
	artificial physical object	1
Total: 720		

Table 9: Distribution of logical patterns in WideSeekBench.

Patterns	**Training Set**		**Test Set**	
	Count	**Percentage**	**Count**	**Percentage**
AND	1,676	37.8%	144	20.0%
AND_NOT	866	19.5%	119	16.5%
AND_OR	704	15.9%	104	14.4%
OR	502	11.3%	93	12.9%
NOT	233	5.3%	94	13.1%
OR_NOT	229	5.2%	83	11.5%
AND_OR_NOT	226	5.1%	83	11.5%
Total	**4,436**	**100.0%**	**720**	**100.0%**

including *film* and *video game*. Scientific and technical sectors are also covered, notably *Space* (e.g., *airport*, *space mission*) and *Infrastructure* (e.g., *metro station*). The dataset exhibits a long-tailed distribution that includes specialized concepts ranging from *Life Sciences* (e.g., *protein family*, *enzyme*) to *Natural Geography* features (e.g., *planetary nebula*, *glacier*). The test set (Figures 8c and 8d) maintains a similar distribution across domains.

(a) Train: Domain Distribution

(b) Train: Subdomain Distribution

(c) Test: Domain Distribution

(d) Test: Subdomain Distribution

Figure 8: Word clouds illustrating the diversity of WideSeekBench. Top row (a, b) shows the Training set, and bottom row (c, d) shows the Test set. The size of each term corresponds to its frequency.

B.6 Task Cases

Example 1: AND

Domain: Education & Academia
Sub-domain: scientific journal

Query:

```
Curious about scientific journals published by Wiley that focus on geography and either use English as
     their work language or have an English name.

Please find scientific journal that meet these conditions, and provide the following attributes for
     these scientific journal: scientific journal, inception, country of origin, indexed in
     bibliographic review, Danish Bibliometric Research Indicator level

**Column Specifications:**
1. **scientific journal**
   * Precise Explanation: A periodical journal that publishes peer-reviewed scientific research focused
        on geography
   * Format Specification: Full official title of the journal in standard title case
   * Example: "Journal of Coastal Geographic Studies"

(15 lines skipped)

5. **Danish Bibliometric Research Indicator level**
   * Precise Explanation: The official ranking level assigned to the journal under the Danish
       Bibliometric Research Indicator system for academic publication channels
```

```
   * Format Specification: Single integer value (either 1 or 2)
   * Example: "1"

Try to provide all information. Only use '/' when you believe the information is truly unavailable or
     cannot be obtained.

Please organize the results in one Markdown table with the following columns: scientific journal,
     inception, country of origin, indexed in bibliographic review, Danish Bibliometric Research
     Indicator level.

Don't ask me any questions, just output the results according to the columns without omitting cells
     arbitrarily. The output format is
```markdown
{data_content}
```.
```

SPARQL Query:

```
PREFIX wd: <http://www.wikidata.org/entity/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>

SELECT DISTINCT ?item WHERE {
  ?item wdt:P31/(wdt:P279|wdt:P1647)* wd:Q5633421 .
  ?item wdt:P123/(wdt:P279|wdt:P1647)* wd:Q1479654 .
  ?item wdt:P407/(wdt:P279|wdt:P1647)* wd:Q1860 .
  ?item wdt:P921/(wdt:P279|wdt:P1647)* wd:Q1071 .
}
```

Table:

scientific journal	inception	country of origin	indexed in bibliographic review	Danish Bibliometric Research Indicator level
The Geographical Journal	1831-01-01T00:00:00Z	United Kingdom	Social Sciences Citation Index; Scopus	1
The Canadian Geographer	1950-01-01T00:00:00Z	Canada	Social Sciences Citation Index; Scopus	1
Geographical Analysis	1969-01-01T00:00:00Z	United States	Social Sciences Citation Index; Scopus	1
Journal of Biogeography	1974-01-01T00:00:00Z	United Kingdom	Science Citation Index Expanded; Scopus	2
Transactions of the Institute of British Geographers	1935-01-01T00:00:00Z	United Kingdom	Scopus	2
Antipode	1969-01-01T00:00:00Z	United Kingdom	Social Sciences Citation Index; Scopus	2
Area	1969-01-01T00:00:00Z	United Kingdom	Social Sciences Citation Index; Scopus	1
International Journal of Urban and Regional Research	1977-01-01T00:00:00Z	United Kingdom	Social Sciences Citation Index; Scopus	2
International Migration	1961-01-01T00:00:00Z	United Kingdom	Social Sciences Citation Index; Scopus	1
Natural Resources Forum	1976-01-01T00:00:00Z	Kingdom of the Netherlands	Science Citation Index Expanded; Scopus	1
(showing first 10 rows, total: 11 rows)				

Rubrics:

Column	Metric	Criterion
scientific journal	llm_judge	Accept exact journal names, common recognized abbreviations, and minor title variations (such as omitting leading articles like 'The' or adjusting capitalization) as long as they refer to the same published journal. Reject completely distinct journal titles that do not correspond to the same publication.
inception	llm_judge	Accept dates that represent the same year, month, and day regardless of string format, including ignoring timestamp suffixes like T00:00:00Z which do not alter the core date value (e.g., '1831-01-01T00:00:00Z' matches 'January 1, 1831' or '1831'). Reject dates referring to a different year, month, or day.
country of origin	llm_judge	Accept the exact official country name, common colloquial aliases, and official alternative names (e.g., 'United Kingdom' matches 'UK', 'Kingdom of the Netherlands' matches 'Netherlands' or 'Holland'). Reject country names that refer to an entirely different sovereign nation.
indexed in bibliographic review	llm_judge	Accept lists that contain the same set of indexing services regardless of order, punctuation separators (semicolon vs comma), or minor spacing differences around separators. For example, 'Social Sciences Citation Index; Scopus' matches 'Scopus, Social Sciences Citation Index'. Reject lists that have missing or additional indexing services not present in the reference value.
Danish Bibliometric Research Indicator level	llm_judge	Accept values that are numerically equivalent to the reference level, regardless of formatting (e.g., the string '2' matches the number 2 or the phrase 'Level 2') as long as they refer to the correct indicator level. Reject values that correspond to a different integer level entirely.

Example 2: AND_NOT

Domain: Natural Geography
Sub-domain: lake

Query:

```
What are all the lakes located in the Northwest Territories, with a basin country of Canada, excluding
     those that outflow into the Buffalo River?
Please find lake that meet these conditions, and provide the following attributes for these lake: lake,
      elevation above sea level, coordinate location, country

**Column Specifications:**
1. **lake**
   - Precise Explanation: The formal, titled name of a contained body of still water situated within a
       natural basin
   - Format Specification: Proper noun written in title case, may include hyphens for connected lake
       systems
   - Example: "Silvermoss Lake"

(10 lines skipped)

4. **country**
   - Precise Explanation: The sovereign state or recognized territorial entity containing the lake;
       multiple entries separated by semicolons for cross-border lakes
   - Format Specification: Proper noun country or territorial entity names, separated by semicolons
       when applicable
   - Example: "Canada; Deninu First Nation"

Try to provide all information. Only use '/' when you believe the information is truly unavailable or
     cannot be obtained.

Please organize the results in one Markdown table with the following columns: lake, elevation above sea
      level, coordinate location, country.

Don't ask me any questions, just output the results according to the columns without omitting cells
     arbitrarily. The output format is
```markdown
{data_content}
```.
```

SPARQL Query:

```
PREFIX wd: <http://www.wikidata.org/entity/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>

SELECT DISTINCT ?item WHERE {
  ?item wdt:P31/(wdt:P279|wdt:P1647)* wd:Q23397 .
  ?item wdt:P131/(wdt:P279|wdt:P1647)* wd:Q2007 .
  ?item wdt:P205/(wdt:P279|wdt:P1647)* wd:Q16 .
  FILTER NOT EXISTS { ?item wdt:P201/(wdt:P279|wdt:P1647)* wd:Q22605327 . }
}
```

Table:

lake	elevation above sea level	coordinate location	country
Clinton-Colden Lake	375.0	POINT(-107.466667 63.966667)	Canada
Colville Lake	244.0	POINT(-126.000000 67.166667)	Canada
Point Lake	375.0	POINT(-113.067962 65.250299)	Canada
Dubawnt Lake	236.0	POINT(-102.096400 63.062790); POINT(-101.533333 63.100000)	Canada
Keller Lake	255.0	POINT(-121.584942 63.950142)	Canada
Artillery Lake	365.0	POINT(-107.866667 63.166667)	Canada
Faber Lake	213.0	POINT(-117.466667 63.883333)	Canada
Lac La Martre	265.0	POINT(-117.950000 63.316667)	Tlicho Nation; Canada
Wholdaia Lake	364.0	POINT(-104.167297 60.716880)	Canada
MacKay Lake	431.0	POINT(-111.033333 63.916667)	Canada
(showing first 10 rows, total: 34 rows)			

Rubrics:

Column	Metric	Criterion
lake	llm_judge	Accept the exact lake name, minor formatting variations like hyphen or spacing changes, and established common aliases for the lake. Reject names referring to an entirely different lake.
elevation above sea level	llm_judge	Accept elevation values that are numerically equivalent or very close, allowing for reasonable rounding differences or minor measurement variations within ±1 meter or 1% difference. Ignore formatting differences like trailing .0 or string versus numeric representation.
coordinate location	llm_judge	Accept coordinate values that represent the same general locations, allowing minor positional deviations within approximately 1km. Ignore formatting differences like POINT() wrapping or separator type (comma vs semicolon), and accept sets of coordinates in any order as long as they refer to the same key locations for the lake.
country	llm_judge	Accept country or recognized indigenous territorial entity names that correctly represent the lake's location, including combined entries separated by any separator in any order. Allow common aliases for countries or associated indigenous nations, and reject entirely unrelated country or territorial names.

Example 3: AND_OR

Domain: Screen & Print Media
Sub-domain: television series

Query:

```
What are all the television series that are either science fiction programs set on an island, or
     television series that are followed by Ashes to Ashes and have Ashes to Ashes as their spin-off?
Please find television series that meet these conditions, and provide the following attributes for
     these television series: television series, original broadcaster, start time, end time, country of
     origin, number of seasons, number of episodes

**Column Specifications:**
1.  **television series**
```

```
    *   Precise Explanation: A connected collection of episodes released under one unified official
         title
    *   Format Specification: Full official title of the television series
    *   Example: "Island of Echoes"

(20 lines skipped)

7.  **number of episodes**
    *   Precise Explanation: Total count of episodes that officially aired as part of the television
         series
    *   Format Specification: Decimal numeric value
    *   Example: "40.0"

Try to provide all information. Only use '/' when you believe the information is truly unavailable or
     cannot be obtained.

Please organize the results in one Markdown table with the following columns: television series,
     original broadcaster, start time, end time, country of origin, number of seasons, number of
     episodes.

Don't ask me any questions, just output the results according to the columns without omitting cells
     arbitrarily. The output format is
```markdown
{data_content}
```.
```

SPARQL Query:

```
PREFIX wd: <http://www.wikidata.org/entity/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>

SELECT DISTINCT ?item WHERE {
  ?item wdt:P31/(wdt:P279|wdt:P1647)* wd:Q5398426 .
  {
  {
    ?item wdt:P136/(wdt:P279|wdt:P1647)* wd:Q336059 .
    ?item wdt:P8411/(wdt:P279|wdt:P1647)* wd:Q23442 .
  }
  UNION
  {
    ?item wdt:P156/(wdt:P279|wdt:P1647)* wd:Q725195 .
    ?item wdt:P2512/(wdt:P279|wdt:P1647)* wd:Q725195 .
  }
  }
}
```

Table:

television series	start time	country of origin	number of episodes
Intergalactic	2021-04-30T00:00:00Z	United Kingdom	8.0
Life on Mars	2006-01-09T00:00:00Z	United Kingdom	16.0
Department S	1969-03-09T00:00:00Z	United Kingdom	28.0
The Hoobs	2001-01-15T00:00:00Z; 2009-10-05T00:00:00Z	United Kingdom; United States	/
Survivors	1975-04-16T00:00:00Z	United Kingdom	38.0
Summer in Transylvania	2010-10-25T00:00:00Z	United Kingdom	20.0
Chocky	1984-01-01T00:00:00Z	United Kingdom	18.0
Agent Carter	2015-01-06T00:00:00Z	United States	18.0
Sense8	2015-06-05T00:00:00Z	United States	24.0
Alice	2009-12-06T00:00:00Z	United Kingdom; Canada	2.0
(showing 4 of 7 columns), (showing first 10 rows, total: 41 rows)			

Rubrics:

875

Column	Metric	Criterion
television series	llm_judge	Accept the exact show title, widely recognized alternate spellings, official abbreviations, and titles with minor qualifying additions (like "(TV series)") that clearly refer to the same show. Reject completely unrelated television series names.
original broadcaster	llm_judge	Accept exact broadcaster names, common official aliases and abbreviations (e.g., "ABC" = "American Broadcasting Company"), lists of broadcasters in any order regardless of separator punctuation, and ignore duplicate entries in the list. Confirm all key original broadcasters are present in the prediction; reject predictions missing major broadcasters or listing entirely unrelated ones.
end time	llm_judge	Accept dates representing the exact same day, month, and year regardless of string format (e.g., ISO 8601 format like "2007-04-10T00:00:00Z" matches "April 10, 2007"). For entries with multiple end times, accept lists in any order regardless of separator punctuation, and confirm all listed end dates are present. Ignore time zone suffixes if the core date is accurate. Reject predictions with significantly mismatched end dates.
number of seasons	llm_judge	Accept numeric values that are equivalent to the reference count, whether formatted as integer or decimal (e.g., 1 == 1.0). Allow for minor discrepancies only if they correspond to widely cited alternative season counts for the specific series; otherwise, reject values that differ by more than one season from the reference.
number of episodes	llm_judge	Accept numeric values that are equivalent to the reference count, whether formatted as integer or decimal (e.g.,8 ==8.0). Allow for minor discrepancies only if they correspond to widely cited alternative episode counts for the specific series; otherwise, reject values that differ by more than five episodes from the reference.
(showing 5 of 7 columns)		

Example 4: AND_OR_NOT

Domain: Business & Economy
Sub-domain: public company

Query:

```
Curious about all public companies that either operate in the retail industry with Derica W. Rice
     serving as a board member, or are based in the United States and are members of the Linux
     Foundation, excluding any companies that had the Target Corporation 2013 data breach as a
     significant event or have Bullseye as their mascot.

Please find public company that meet these conditions, and provide the following attributes for these
     public company: public company, headquarters location, stock exchange, legal form, inception

**Column Specifications:**
1.  **public company**
    *   Precise Explanation: A company that offers its securities for sale to the general public,
         allowing members of the public to invest and own shares in the business
    *   Format Specification: Properly formatted official company name, following standard business
         naming conventions
    *   Example: Northwood Retail Group

(15 lines skipped)

5.  **inception**
    *   Precise Explanation: The official date when the public company first came into existence,
         distinct from official opening or launch dates
    *   Format Specification: ISO 8601 date format (yyyy-mm-ddTHH:MM:SSZ), using semicolons to separate
          multiple inception dates if applicable
    *   Example: 2001-05-15T00:00:00Z; 2005-09-02T00:00:00Z

Try to provide all information. Only use '/' when you believe the information is truly unavailable or
     cannot be obtained.

Please organize the results in one Markdown table with the following columns: public company,
     headquarters location, stock exchange, legal form, inception.

Don't ask me any questions, just output the results according to the columns without omitting cells
     arbitrarily. The output format is
```markdown
{data_content}
```.
```

SPARQL Query:

```
PREFIX wd: <http://www.wikidata.org/entity/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>

SELECT DISTINCT ?item WHERE {
  ?item wdt:P31/(wdt:P279|wdt:P1647)* wd:Q891723 .
  {
  {
    ?item wdt:P452/(wdt:P279|wdt:P1647)* wd:Q126793 .
    ?item wdt:P3320/(wdt:P279|wdt:P1647)* wd:Q109743489 .
  }
  UNION
  {
    ?item wdt:P17/(wdt:P279|wdt:P1647)* wd:Q30 .
    ?item wdt:P463/(wdt:P279|wdt:P1647)* wd:Q858851 .
  }
  }
  FILTER NOT EXISTS { { ?item wdt:P793/(wdt:P279|wdt:P1647)* wd:Q76313694 . } UNION { ?item wdt:P822/(
      wdt:P279|wdt:P1647)* wd:Q4997048 . } }
}
```

Table:

public company	headquarters location	stock exchange	legal form	inception
CME Group	Chicago; London	Nasdaq	Delaware corporation	2007-07-12T00:00:00Z
Capital One	McLean	New York Stock Exchange	private limited company; Commercial Bank; Delaware corporation	1994-01-01T00:00:00Z; 1994-07-21T00:00:00Z
Juniper Networks	Sunnyvale	New York Stock Exchange	Delaware corporation	1996-02-06T00:00:00Z
Ciena	Hanover	New York Stock Exchange; Nasdaq	Delaware corporation	1999-01-01T00:00:00Z
Nuance Communications	Burlington	Nasdaq	Delaware corporation	1992-01-01T00:00:00Z
VMware	Palo Alto	New York Stock Exchange	corporation	1998-01-01T00:00:00Z
Adobe	San Jose	London Stock Exchange; Nasdaq	Delaware corporation	1982-02-28T00:00:00Z
Netflix, Inc.	Los Gatos	Nasdaq	Delaware corporation; public company	1997-08-29T00:00:00Z
Micron Technology	Boise	Nasdaq	Delaware corporation	1978-10-01T00:00:00Z
EPAM Systems	Newtown; Pennsylvania	New York Stock Exchange	Delaware corporation	1993-01-01T00:00:00Z
(showing first 10 rows, total: 63 rows)				

Rubrics:

Column	Metric	Criterion
public company	llm_judge	Accept the exact company name, common aliases, standard suffix variations (e.g., 'Netflix' matches 'Netflix, Inc.'), and minor punctuation differences. Reject any entry referring to a completely unrelated public company.
headquarters location	llm_judge	Accept if the set of headquarters locations matches, regardless of order or separator (comma vs semicolon). Allow common location aliases (e.g., 'PA' for 'Pennsylvania') and accept entries referring to the same general metropolitan or geographic area. Reject entries that list entirely unrelated locations.
stock exchange	llm_judge	Accept if the set of stock exchanges matches, regardless of order or separator (comma vs semicolon). Allow standard abbreviations (e.g., 'NYSE' matches 'New York Stock Exchange', 'LSE' matches 'London Stock Exchange'). Reject entries that list entirely different stock exchanges.
legal form	llm_judge	Accept semantic equivalents of the legal form(s), regardless of order or separator (comma vs semicolon). Allow abbreviations (e.g., 'Delaware corp' matches 'Delaware corporation') and equivalent terms (e.g., 'corporation' matches 'Delaware corporation' where the jurisdiction is implied). Reject entries that describe an entirely different legal structure unrelated to the reference.
inception	llm_judge	Accept dates representing the same day, month, and year, regardless of string format (e.g., '2007-07-12' matches 'July 12, 2007'). Ignore trailing timestamp suffixes like 'T00:00:00Z' as non-essential. If multiple inception dates are listed, accept any matching equivalent date from the reference set. Reject entries with completely different inception dates.

878

Example 5: NOT

Domain: Gaming
Sub-domain: video game

Query:

```
Curious about video games that have been nominated for The Game Awards - Most Anticipated Game,
     excluding any that have a GRAC rating of 15+

Please find video games that meet these conditions, and provide the following attributes for these
     video games: video game, review score, input device, platform, distributed by, publication date,
     distribution format, country of origin, developer, publisher, language of work or name, genre,
     announcement date, game mode

**Column Specifications:**
1. **video game**
   - Precise Explanation: The official title of an electronic game that features a user interface and
       provides visual feedback to players during gameplay
   - Format Specification: Properly capitalized official game title matching standard marketing
       punctuation and spelling
   - Example: "Lunar Frontier: Operation Genesis"

(60 lines skipped)

14. **game mode**
    - Precise Explanation: The official playing modes available for the video game, including single-
        player and multiplayer options
    - Format Specification: Standard industry terms for video game playing modes
    - Example: "single-player video game; competitive multiplayer"

Try to provide all information. Only use '/' when you believe the information is truly unavailable or
     cannot be obtained.

Please organize the results in one Markdown table with the following columns: video game, review score,
      input device, platform, distributed by, publication date, distribution format, country of origin,
     developer, publisher, language of work or name, genre, announcement date, game mode.

Don't ask me any questions, just output the results according to the columns without omitting cells
     arbitrarily. The output format is
```markdown
{data_content}
```.
```

SPARQL Query:

```
PREFIX wd: <http://www.wikidata.org/entity/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>

SELECT DISTINCT ?item WHERE {
  ?item wdt:P31/(wdt:P279|wdt:P1647)* wd:Q7889 .
  ?item wdt:P1411/(wdt:P279|wdt:P1647)* wd:Q68094302 .
  FILTER NOT EXISTS { ?item wdt:P9866/(wdt:P279|wdt:P1647)* wd:Q23005410 . }
}
```

Table:

video game	distributed by	publisher	game mode
Resident Evil 4	PlayStation Store; Microsoft Store; Steam; ...	Capcom	single-player video game
Final Fantasy VII Rebirth	PlayStation Store; Steam	Square Enix	single-player video game
Hades II	Nintendo eShop; Steam; Epic Games Store	Supergiant Games	single-player video game
Death Stranding 2: On the Beach	PlayStation Store	Sony Interactive Entertainment	single-player video game
Like a Dragon: Infinite Wealth	PlayStation Store; Microsoft Store; Steam	Sega	single-player video game
Star Wars Outlaws	PlayStation Store; Microsoft Store; Steam; ...	Ubisoft	single-player video game
Monster Hunter Wilds	PlayStation Store; Microsoft Store; Steam	Capcom	single-player video game
Ghost of Yōtei	PlayStation Store	Sony Interactive Entertainment	single-player video game
Quantum Break	Microsoft Store; Steam; Humble Store	Xbox Game Studios	single-player video game
No Man's Sky	PlayStation Store; Microsoft Store; GOG.com; ...	Sony Interactive Entertainment; Hello Games	co-op mode; single-player video game; multiplayer video game
(showing 4 of 14 columns), (showing first 10 rows, total: 34 rows)			

Rubrics:

Column	Metric	Criterion
video game	llm_judge	Accept exact official game titles and widely recognized common aliases (e.g., 'RE4' matches 'Resident Evil 4'). Reject entirely unrelated game titles that do not reference the same core game.
review score	llm_judge	Accept lists of scores that contain the same core numerical values regardless of formatting (percentages vs integer scores, semicolon vs comma separators). Allow minor rounding differences and ignore trivial punctuation variations. Reject scores that are significantly different in magnitude from the reference.
country of origin	llm_judge	Accept exact country names and standard common aliases (e.g., 'US' matches 'United States', 'UK' matches 'United Kingdom'). Reject countries that are not the official origin of the game's development.
announcement date	llm_judge	Accept dates that represent the exact same day, month, and year regardless of string format (e.g., '2022-06-02' matches 'June 2, 2022'). Ignore list order and separator punctuation. Reject dates that refer to a different official announcement day.
game mode	llm_judge	Accept game mode names and common shorthand equivalents (e.g., 'co-op' matches 'co-op mode', 'multiplayer' matches 'multiplayer video game'). Ignore list order and separator punctuation. Reject game modes that are not officially supported by the game.
(showing 5 of 14 columns)		

Example 6: OR

Domain: Space
Sub-domain: airport

Query:

```
Please find all airports that have Paris-Charles de Gaulle Airport as a scheduled service destination,
     or are located on Crete, or are owned by Greece, or are located in Heraklion.
Please find airport that meet these conditions, and provide the following attributes for these airport:
      airport, IATA airport code, elevation above sea level, coordinate location, place served by
     transport hub, ICAO airport code, aerodrome reference point, date of official opening, named after,
      patronage, runway, country, located in time zone

**Column Specifications:**
1.  **airport**
    * Precise Explanation: A location where aircraft take off and land with extended support facilities
        , mostly for commercial air transport
    * Format Specification: Full official airport name including any honorific or nickname enclosed in
        quotes
    * Example: "Thessaloniki Airport \"Makedonia\""

(55 lines skipped)
```

```
13. **located in time zone**
    * Precise Explanation: Official time zone the airport operates within per the tz database
    * Format Specification: Tz database time zone identifier string
    * Example: "Europe/Athens"

Try to provide all information. Only use '/' when you believe the information is truly unavailable or
     cannot be obtained.

Please organize the results in one Markdown table with the following columns: airport, IATA airport
     code, elevation above sea level, coordinate location, place served by transport hub, ICAO airport
     code, aerodrome reference point, date of official opening, named after, patronage, runway, country,
      located in time zone.

Don't ask me any questions, just output the results according to the columns without omitting cells
     arbitrarily. The output format is
```markdown
{data_content}
```.
```

SPARQL Query:

```
PREFIX wd: <http://www.wikidata.org/entity/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>

SELECT DISTINCT ?item WHERE {
  ?item wdt:P31/(wdt:P279|wdt:P1647)* wd:Q1248784 .
  {   { ?item wdt:P521/(wdt:P279|wdt:P1647)* wd:Q46280 . }
    UNION
  { ?item wdt:P706/(wdt:P279|wdt:P1647)* wd:Q34374 . }
    UNION
  { ?item wdt:P127/(wdt:P279|wdt:P1647)* wd:Q41 . }
    UNION
  { ?item wdt:P131/(wdt:P279|wdt:P1647)* wd:Q160544 . }
  }
}
```

Table:

airport	place served by transport hub	named after	located in time zone
Heraklion International Airport "Nikos Kazantzakis"	Heraklion	Heraklion; Nikos Kazantzakis	Europe/Athens
Chania International Airport	Chania	Daskalogiannis; Chania	Europe/Athens
Thessaloniki Airport "Makedonia"	Thessaloniki	Macedonia; Thessaloniki; Macedonia	Europe/Athens
Ottawa Macdonald-Cartier International Airport	Ottawa	John A. Macdonald; Ottawa; George-Étienne Cartier	Eastern Time Zone
Roland Garros Airport	Saint-Denis	Réunion; Roland Garros; Saint-Denis	Indian/Reunion
Athens International Airport "Eleftherios Venizelos"	Athens	Athens; Eleftherios Venizelos	Europe/Athens
Toronto Pearson International Airport	Toronto	Lester B. Pearson; Malton; Toronto	Eastern Time Zone
Vancouver International Airport	Vancouver	Vancouver	America/Vancouver
Guiyang Longdongbao International Airport	Guiyang	/	Asia/Shanghai
Warsaw Chopin Airport	Warsaw	Frederic Chopin; Warsaw; Okecie	Europe/Warsaw
(showing 4 of 13 columns), (showing first 10 rows, total: 21 rows)			

Rubrics:

Column	Metric	Criterion
airport	llm_judge	Accept the exact airport name, shortened variations (e.g., omitting "International"), title case or case-insensitive matches, and common aliases that clearly refer to the same airport. Reject names of entirely different airports.
IATA airport code	llm_judge	Accept the standard 3-letter IATA code regardless of case (e.g., "HER" matches "her" or "Her"). Reject any non-matching 3-letter code that refers to a different airport.
aerodrome reference point	llm_judge	Accept reference point coordinates that match the airport's official reference location, allowing minor decimal deviations (within ~100m), and any valid coordinate format regardless of string structure. Reject points that do not correspond to the airport's reference point.
country	llm_judge	Accept the exact country name or common aliases (e.g., "China" matches "People's Republic of China" or "PRC"), case variations, and abbreviations that are widely recognized. Reject names of entirely different countries.
located in time zone	llm_judge	Accept standard time zone names, common abbreviations (e.g., "EST" matches "Eastern Time Zone"), and semantic equivalents that refer to the same time zone. Reject any time zone that does not correspond to the airport's geographic location.
(showing 5 of 13 columns)		

Example 7: OR_NOT

Domain: Sports
Sub-domain: recurring sporting event edition

Query:

```
Curious about all recurring sporting event editions that either have Germany at the 2016 Summer
     Olympics as a participant, are part of the Summer Olympic Games sports season, or include Olympic
     sporting events as part of their programming, excluding any editions located in Brazil or hosted at
      Maracana Stadium.

Please find recurring sporting event edition that meet these conditions, and provide the following
     attributes for these recurring sporting event edition: recurring sporting event edition, number of
     participants, location, start time, end time, officially opened by

**Column Specifications:**
1.  recurring sporting event edition
    Precise Explanation: A specific yearly or regularly scheduled run of a recurring sporting event,
         representing one distinct iteration of the competition
    Format Specification: Formal proper noun title including the event name and associated year
    Example: 2032 Winter Paralympic Games

(20 lines skipped)

6.  officially opened by
    Precise Explanation: The public figure or official who performed the formal opening ceremony for
         the event
    Format Specification: Full formal name of the individual who conducted the opening ceremony
    Example: Charles III of the United Kingdom

Try to provide all information. Only use '/' when you believe the information is truly unavailable or
     cannot be obtained.

Please organize the results in one Markdown table with the following columns: recurring sporting event
     edition, number of participants, location, start time, end time, officially opened by.

Don't ask me any questions, just output the results according to the columns without omitting cells
     arbitrarily. The output format is
```markdown
{data_content}
```.
```

SPARQL Query:

```
PREFIX wd: <http://www.wikidata.org/entity/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>
```

```
SELECT DISTINCT ?item WHERE {
  ?item wdt:P31/(wdt:P279|wdt:P1647)* wd:Q114609228 .
  {   { ?item wdt:P710/(wdt:P279|wdt:P1647)* wd:Q17739057 . }
    UNION
  { ?item wdt:P3450/(wdt:P279|wdt:P1647)* wd:Q159821 . }
    UNION
  { ?item wdt:P2670/(wdt:P279|wdt:P1647)* wd:Q18536594 . }
  }
  FILTER NOT EXISTS { { ?item wdt:P17/(wdt:P279|wdt:P1647)* wd:Q155 . } UNION { ?item wdt:P115/(wdt:
       P279|wdt:P1647)* wd:Q155174 . } }
}
```

Table:

recurring sporting event edition	number of participants	location	start time	end time	officially opened by
2028 Summer Olympics	/	Los Angeles	2028-07-14T00:00:00Z	2028-07-30T00:00:00Z	/
2020 Summer Olympics	206.0; 11420.0	Tokyo	2021-07-24T00:00:00Z	2021-08-08T00:00:00Z	Naruhito
1896 Summer Olympics	14.0; 241.0	Athens	1896-04-06T00:00:00Z	1896-04-15T00:00:00Z	George I of Greece
1900 Summer Olympics	26.0; 1226.0	Paris	1900-05-14T00:00:00Z	1900-10-28T00:00:00Z	/
1904 Summer Olympics	12.0; 651.0	St. Louis	1904-07-01T00:00:00Z	1904-11-23T00:00:00Z	David R. Francis
1908 Summer Olympics	22.0; 2008.0	London	1908-04-27T00:00:00Z	1908-10-31T00:00:00Z	Edward VII
1912 Summer Olympics	28.0; 2407.0	Stockholm	1912-05-05T00:00:00Z	1912-07-27T00:00:00Z	Gustaf V of Sweden
1916 Summer Olympics	/	Berlin	/	/	/
1920 Summer Olympics	29.0; 2622.0	Antwerp	1920-04-20T00:00:00Z	1920-09-12T00:00:00Z	Albert I of Belgium
1924 Summer Olympics	44.0; 3088.0	Paris	1924-05-04T00:00:00Z	1924-07-27T00:00:00Z	Gaston Doumergue
(showing first 10 rows, total: 27 rows)					

Rubrics:

Column	Metric	Criterion
recurring sporting event edition	llm_judge	Accept any reference to the exact same edition of the recurring sporting event, including common aliases and formatting variations (e.g., "1896 Summer Olympics" matches "Athens 1896 Olympics" or "1st Summer Olympic Games"). Reject references to a different year, a different recurring sporting event, or an incorrect edition entirely.
number of participants	llm_judge	Accept participant counts that match the reference values within a 5% tolerance for each numerical entry, regardless of separator used (semicolon, comma, space, etc.). Allow omission of trailing decimal zeros (e.g., 206 matches 206.0) and accept the two counts presented in either order as long as both values are correct within tolerance. Reject counts that differ by more than 10% or are entirely unrelated to the event's participant numbers.
start time	llm_judge	Accept dates that represent the exact same calendar day, month, and year for the event's start, regardless of string format or time zone notation (e.g., "2028-07-14T00:00:00Z" matches "July 14, 2028" or "14/07/2028"). Reject any date that does not correspond to the correct start date of the event edition.
end time	llm_judge	Accept dates that represent the exact same calendar day, month, and year for the event's end, regardless of string format or time zone notation (e.g., "2028-07-30T00:00:00Z" matches "July 30, 2028" or "30/07/2028"). Reject any date that does not correspond to the correct end date of the event edition.
officially opened by	llm_judge	Accept the exact name of the dignitary, standard title variations, common aliases, and omissions of non-essential title details (e.g., "George I of Greece" matches "King George I" or "George I"). Ignore minor variations in name formatting like spacing or punctuation. Reject any person who did not officially open the specified event edition.
(showing 5 of 6 columns)		

B.7 Simulated Environment

To facilitate training and validation, we construct a stable and realistic simulated search engine, utilizing a snapshot of WikiPedia 2025 as the corpus. To guarantee task solvability, we verified that all entities appearing in the answer tables possess corresponding Wikipedia pages and are contained within the utilized dump. We employ Qwen3-0.6B-Embedding[2] to extract features from all text data, converting them into corresponding embeddings.This environment exposes two functions:

[2]https://huggingface.co/Qwen/Qwen3-Embedding-0.6B

- search: Computes the query embedding on the fly, retrieves the top-k nearest documents from the corpus, and returns their URLs and abstracts.
- open_page: Retrieves the full content of a specific page given its DocID or URL.

We show the schema of these tools as below:

Schema of Tools

```
SEARCH = {
    'type': 'function',
    'function': {
        "name": "search",
        "description": "Performs a web search: supply a string 'query' and optional 'topk'. The tool
             retrieves the top 'topk' results (default 10) for the query, returning their docid, url,
             and document content (may be truncated based on token limits).",
        "parameters": {
            "type": "object",
            "properties": {
                "query": {
                    "type": "string",
                    "description": "The query string for the search."
                },
                "topk": {
                    "type": "integer",
                    "description": "Return the top k pages.",
                }
            },
            "required": [
                "query"
            ]
        }
    }
}


OPEN_PAGE = {
    'type': 'function',
    'function': {
        'name': 'open_page',
        'description': (
            "Open a page by docid or URL and return the complete content. "
            "Provide either 'docid' or 'url'; if both are provided, prefer 'docid'. "
            "The docid or URL must come from prior search tool results."
        ),
        'parameters': {
            'type': 'object',
            'properties': {
                'docid': {
                    'type': 'string',
                    'description': 'Document ID from search results to resolve and fetch.',
                },
                'url': {
                    'type': 'string',
                    'description': 'Absolute URL from search results to fetch.',
                },
            },
            'required': [],
        },
    },
}
```

B.8 Evaluation

To comprehensively assess the quality of the generated tables across different granularities, we employ three evaluation metrics: Success Rate, Row F1, and Item F1. These metrics evaluate the performance at the table, row, and cell levels, respectively. Specifically, we use the LLM-based judge with column-wise rubrics to evaluate whether each generated cell is aligned with the corresponding ground truth cell. We use the GPT-4.1 as the default judge LLM.

- **Success Rate:** This is the strictest metric, operating at the table level. A sample is considered a success only if the answer table exactly matches the ground truth in terms of both content and structure, without any errors.
- **Row F1:** This metric evaluates the retrieval and generation accuracy at the row level. We calculate the precision and recall of the generated rows against the ground truth rows to compute the F1 score.

A predicted row is considered a correct match only if all the cells within that row are perfectly consistent with the corresponding ground truth row.

- **Item F1:** To provide a fine-grained assessment, Item F1 evaluates performance at the cell level. It calculates the F1 score based on the individual data items (cells) within the table. This metric focuses on the model's ability to extract or generate specific details correctly, regardless of whether the entire row is perfect.

C Experiments

C.1 Cold Start

To bootstrap the unified policy π_θ with the capability to perform complex task decomposition and robust information seeking, we employ a Cold Start phase via Supervised Fine-Tuning (SFT).

Trajectory Collection and Filtering. We utilize multiple teacher policies (e.g., DeepSeek-V3.2, Kimi-K2) to generate a diverse set of rollout trajectories on the training set $\mathcal{D}_{train}$. For each query $\mathcal{Q}_i$, we collect a set of candidate trajectories $\{\boldsymbol{\mathcal{T}}_{i,m}\}_{m=1}^{M}$. To ensure the quality of the training signal, we introduce a strict filtering mechanism based on the Item-level F1 score ($F1_{\text{item}}$) against the ground truth table $\mathbf{T}_i^*$. A trajectory is retained for the SFT dataset $\mathcal{D}_{SFT}$ if and only if its performance exceeds a threshold η:

$$\mathcal{D}_{SFT} = \{\boldsymbol{\mathcal{T}}_{i,m} \mid \text{Item-F1}(\text{Answer}(\boldsymbol{\mathcal{T}}_{i,m}), \mathbf{T}_i^*) > \eta\} \tag{7}$$

We set the η as 0.6.

SFT Optimization. The policy π_θ is initialized by minimizing the standard negative log-likelihood loss over the filtered high-quality trajectories. Let $\boldsymbol{\mathcal{T}}$ be represented as a sequence of tokens $(x_1, x_2, \ldots, x_L)$. The SFT objective is defined as:

$$\mathcal{L}_{SFT}(\theta) = -\mathbb{E}_{\boldsymbol{\mathcal{T}} \sim \mathcal{D}_{SFT}} \left[\sum_{t=1}^{|\boldsymbol{\mathcal{T}}|} \log \pi_\theta(x_t \mid x_{<t}) \right] \tag{8}$$

The loss is only computed on the tokens generated by models itself (the thoughts and actions).

C.2 Training Dynamics

C.3 Setting

We use VERL (Sheng et al., 2025) and AgentLightning (Luo et al., 2025) as the RL training framework. We use Qwen3-8B (Yang et al., 2025) as the base model. The RL hyper parameters are shown in Table 10. We use 64 H100 GPUs for RL Training. We use function calling to create sub-agent for main agent. The corresponding tool schema is as below. We use the GPT-4.1 as the default judge LLM.

Hyper Parameter	Value
Batch Size	64
Num of Rollout	6
Max Prompt Length	32768
Max Response Length	8192
Learning Rate	1e-6
Clip High	0.28
Clip Low	0.2
Training Step	80

Table 10: The hyper parameters for RL traning.

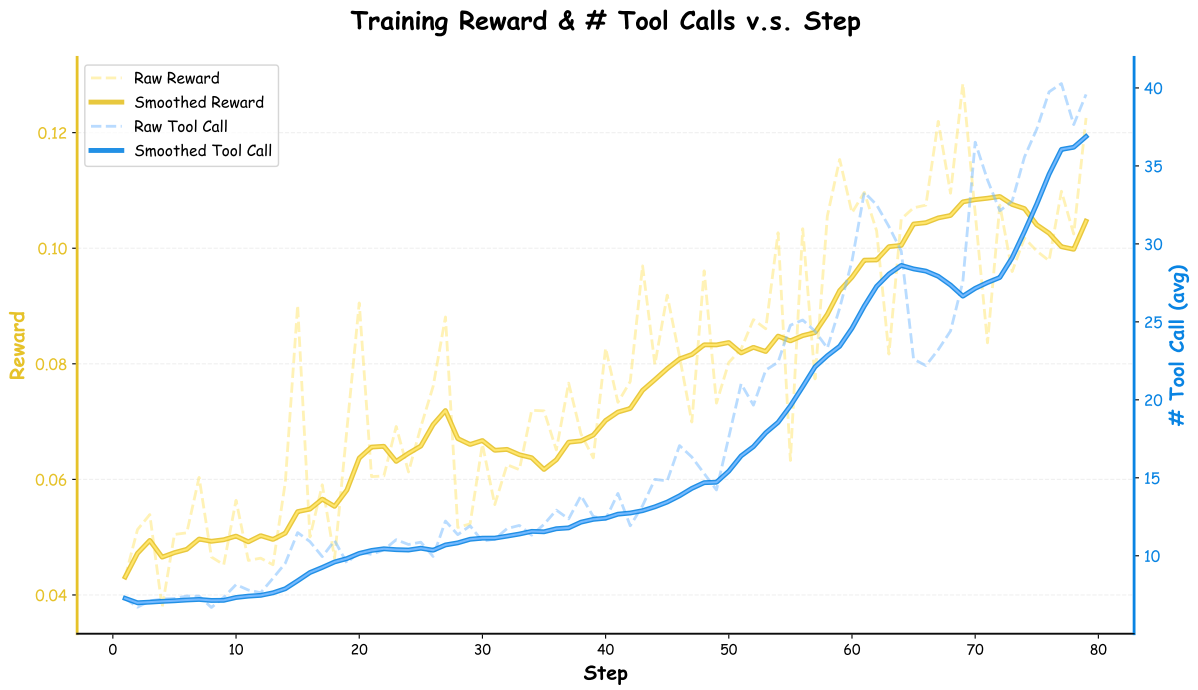

Figure 9: Training dynamics of WideSeek-8B-RL. Left: the reward increases steadily over 80 training steps. Right: the number of tool calls per trajectory grows in tandem, indicating that the agent autonomously discovers broader search as the optimal policy under GRPO optimization.

Schema of Tools

```
CREATE_SUB_AGENT = {
    "type": "function",
    "function": {
        "name": "create_sub_agent",
        "description": "Create more sub-agents to gather information in parallel.",
        "parameters": {
            "type": "object",
            "properties": {
                "tasks": {
                    "type": "array",
                    "description": "A list of tasks, each for a new sub-agent. The number of subtasks
                         created at one time can be large or small.",
                    "items": {
                        "type": "object",
                        "properties": {
                            "agent_id": {"type": "string", "description": "A unique identifier for the
                                 sub-agent, e.g., 'agent_001'."},
                            "task": {"type": "string", "description": "The specific task this sub-agent
                                  must perform."}
                        },
                        "required": ["agent_id", "task"]
                    }
                }
            },
            "required": ["tasks"]
        }
    }
}
```

C.4 Used Token Analysis

To understand the computational cost of different models and scaffolds, we report the average number of tokens consumed per task on the WideSearch benchmark in Table 11.

Several observations emerge from the token usage analysis. First, the WideSeek scaffold is more token-efficient than ReAct under the same base model: Qwen3-8B with WideSeek consumes 11,548 tokens compared to 11,942 with ReAct, indicating that parallel sub-agent decomposition avoids redundant

Table 11: Performance and average token usage per task on WideSearch. Main Agent tokens represent the planner's reasoning cost; Sub Agent tokens represent the total cost of all forked executors. ReAct models have no sub-agents.

Model	Scaffold	Item-F1 (Mean@4)	Item-F1 (Max@4)	Main Agent Tokens	Sub Agent Tokens
Qwen3-8B	ReAct	28.44	40.14	11,942	0
ReAct-8B-SFT	ReAct	34.09	48.04	24,439	0
ReAct-8B-RL	ReAct	37.89	50.10	28,441	0
Qwen3-8B	WideSeek	39.13	50.29	11,548	76,258
WideSeek-8B-SFT	WideSeek	44.85	57.13	17,070	147,077
WideSeek-8B-SFT-RL	WideSeek	45.44	58.81	21,873	258,775
WideSeek-8B-RL	WideSeek	42.03	55.39	18,356	184,805

sequential reasoning. Second, optimization increases token usage as expected, reflecting the model's learned strategy of deploying more extensive search. ReAct-8B-RL uses 28,441 tokens (2.38× the base), while WideSeek-8B-SFT-RL uses 21,873 tokens (1.89× the base), demonstrating that WideSeek achieves superior performance (Item-F1 45.44 vs. 37.89) with fewer tokens. Third, WideSeek-8B-SFT-RL consumes only 21,873 tokens yet surpasses DeepSeek-R1-0528 on all metrics, confirming that our multi-agent coordination strategy is both more effective and more computationally efficient than scaling model parameters alone.

C.5 Case Study

We illustrate the unified trajectory of the same task query produced by 4 models in Figure 10: Qwen3-30B-A3B-Thinking, WideSeek-8B-SFT-RL, WideSeek-8B-SFT, WideSeek-8B-RL. And for better understanding, we show a case trajectory of WideSeek-8B-RL as follows.

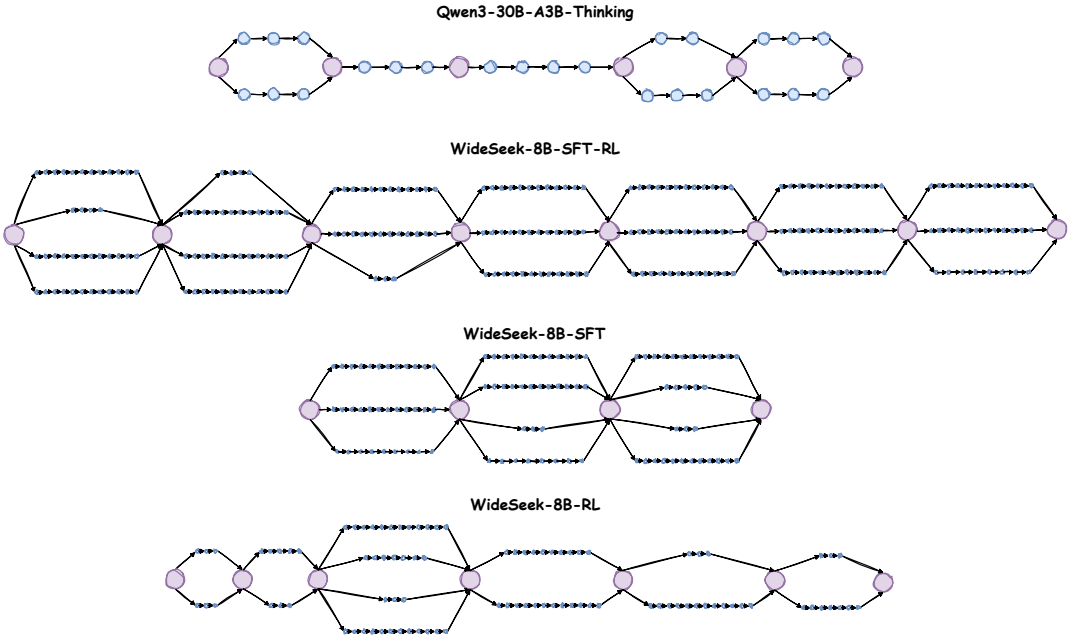

Figure 10: Multi-Agent Trajectory

Case trajectory on WideSeekBench

Question: Please find all films that were nominated for the Academy Award for Best Sound and received the National Board of Review: Top Ten Films award, excluding any films that feature Constantin Alexandrov as a cast member or have Anna Hamilton Phelan as a screenwriter.
Please find film that meet these conditions, and provide the following attributes for these film: film,

color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language of film or TV show, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating

*Column Specifications:**

1. **film**
Precise Explanation: A sequence of moving images stored on film stock, released as a feature-length motion picture.
Format Specification: Full official title of the film as a single string.
Example: "Midnight Tide"

2. **color**
Precise Explanation: The color format used to produce the film.
Format Specification: Single term indicating the color format, typically "color" or "black and white".
Example: "black and white"

3. **country of origin**
Precise Explanation: The country or countries where the film was primarily produced and created.
Format Specification: List of country names separated by semicolons if multiple countries apply.
Example: "United Kingdom; Canada"

4. **producer**
Precise Explanation: The primary person or persons responsible for overseeing the film's core production work, excluding executive or associate producers.
Format Specification: List of producer full names separated by semicolons if multiple producers are credited.
Example: "Elara Voss; Marcus Hale"

5. **review score**
Precise Explanation: Aggregate or individual critical review scores the film received from rating platforms or professional critics.
Format Specification: Multiple score values (percentages, numerical ratings) separated by semicolons.
Example: "88%; 9.2; 95"

6. **CNC film rating (France)**
Precise Explanation: Age restriction or content warning category assigned to the film by France's CNC Board of Film Classification.
Format Specification: Rating category label, with multiple categories separated by semicolons if applicable.
Example: "12 years and over; violence warning"

7. **assessment**
Precise Explanation: List of social representation or media diversity tests the film has been evaluated against.
Format Specification: Names of assessment tests separated by semicolons.
Example: "Bechdel test; Mako Mori test"

956

8. **narrative location**
Precise Explanation: The fictional or real-world settings where the film's plot and story take place.
Format Specification: List of location names separated by semicolons if multiple settings are featured.
Example: "Iceland; Reykjavík; Northern Fjords"

9. **distributed by**
Precise Explanation: The companies responsible for distributing the film to audiences via theatrical, streaming, or home video channels.
Format Specification: List of distribution company names separated by semicolons if multiple distributors are involved.
Example: "Warner Bros. Pictures; Max"

10. **original language of film or TV show**
Precise Explanation: The primary language used during the original production and recording of the film's dialogue and core audio.
Format Specification: Single language name or list separated by semicolons if multiple original languages are used.
Example: "Spanish; Catalan"

11. **composer**
Precise Explanation: The person or persons who created the original musical score for the film, separate from lyricists.
Format Specification: List of composer full names separated by semicolons if multiple composers are credited.
Example: "Lila Marquez; Theo Carter"

12. **distribution format**
Precise Explanation: The primary method used to distribute the film to viewers.
Format Specification: Single term describing the distribution format, such as theatrical release or video on demand.
Example: "theatrical release"

13. **box office**
Precise Explanation: The total ticket sales revenue earned by the film, potentially including domestic and international totals.
Format Specification: Numeric box office values separated by semicolons if multiple totals are provided.
Example: "78245100.0; 195320000.0"

14. **director**
Precise Explanation: The person or persons responsible for overseeing the film's creative direction and artistic vision throughout production.
Format Specification: List of director full names separated by semicolons if multiple directors are credited.
Example: "Jasper Reed"

15. **Kijkwijzer rating**
Precise Explanation: The age-based content rating assigned to the film under the Dutch Kijkwijzer

media rating system.
Format Specification: Numeric rating value, with multiple ratings separated by semicolons if applicable.
Example: "9; 12"

16. **production designer**
Precise Explanation: The person responsible for creating the film's overall visual aesthetic, including set design, props, and visual conceptualization.
Format Specification: List of production designer full names separated by semicolons if multiple designers are credited.
Example: "Zara Bennett"

17. **duration**
Precise Explanation: The total runtime of the film measured in minutes.
Format Specification: Numeric value representing the total runtime in minutes.
Example: "122.0"

18. **filming location**
Precise Explanation: The actual real-world locations where the film was physically shot, distinct from its narrative setting.
Format Specification: List of filming location names separated by semicolons if multiple locations were used.
Example: "Vancouver; British Columbia; Stanley Park"

19. **film editor**
Precise Explanation: The person or persons responsible for assembling raw footage into the final finished cut of the film.
Format Specification: List of film editor full names separated by semicolons if multiple editors are credited.
Example: "Noah Brooks; Sienna Torres"

20. **production company**
Precise Explanation: The companies that funded and oversaw the core production process of the film.
Format Specification: List of production company names separated by semicolons if multiple companies were involved.
Example: "A24; Reel Pictures Studio"

21. **FSK film rating**
Precise Explanation: The age restriction rating assigned to the film under the German FSK film classification system.
Format Specification: Rating in the standard format "FSK [age]".
Example: "FSK 12"

Try to provide all information. Only use '/' when you believe the information is truly unavailable or cannot be obtained.

Please organize the results in one Markdown table with the following columns: film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language of film or TV show, composer, distribution format, box

office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating.

Don't ask me any questions, just output the results according to the columns without omitting cells arbitrarily. The output format is
“‘markdown
{data_content}
“‘.

Round 1 (Main Agent + Sub Agents: agent_001, agent_002)

Main Agent

<think> Okay, let's tackle this query step by step. The user wants all films that were nominated for the Academy Award for Best Sound and received the National Board of Review: Top Ten Films award, but exclude those with Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter. Then, they need a list of attributes for each film. First, I need to find films that are in both the Academy Award Best Sound nominees and the National Board of Review Top Ten Films. Then, filter out any films that have either of the two excluded people. But the user also wants a lot of specific details about each film, which might not be straightforward to get. Since I don't have access to external databases like Wikidata, I'll need to create sub-agents to search for this information. Let me start by creating sub-agents to get the list of Academy Award Best Sound nominees and the National Board of Review Top Ten Films. Then, I can cross-reference these lists to find overlapping films. Next, check each film in the overlap for the presence of Constantin Alexandrov or Anna Hamilton Phelan, and exclude those that do. Finally, collect all the required attributes for each remaining film. But how to structure this? Let me first split the tasks. First, get the Academy Award nominees for Best Sound. Then, get the National Board of Review Top Ten Films. Then, find the intersection. Then, for each of those, check for the two excluded...
</think>
<tool_call> {"name":"create_sub_agent","arguments":{"tasks":[{"agent_id":"agent_001","task":"List all films nominated for the Academy Award for Best Sound."},{"agent_id":"agent_002","task":"List all films that received the National Board of Review: Top Ten Films award."}]}} </tool_call> <tool_response> Sub-agents for call chatcmpl-tool-7f2e53c1e9bd41d6bf46824416e0be1f completed. Reports:

— Report from agent_001 —
Here is a list of films nominated for the Academy Award for Best Sound, based on the historical data available in the provided Wikipedia content (up to the 1970s):

1930s
- **1929–1930 (3rd)**: *The Big House*, *The Case of Sergeant Grischa*, *The Love Parade*, *Raffles*, *Song of the Flame*
- **1930–1931 (4th)**: *Paramount Publix Studio Sound Department*, *Metro-Goldwyn-Mayer Studio Sound Department*, *RKO Radio Studio Sound Department*, *Samuel Goldwyn-United Artists Studio Sound Department*
- **1931–1932 (5th)**: *Paramount Publix Studio Sound Department*, *Metro-Goldwyn-Mayer Studio Sound Department*, *RKO Radio Studio Sound Department*, *Warner Bros.-First National Studio Sound Department*
- **1932–1933 (6th)**: *A Farewell to Arms*, *42nd Street*, *Gold Diggers of 1933*, *I Am a Fugitive from a Chain Gang*
- **1934 (7th)**: *One Night of Love*, *The Affairs of Cellini*, *Cleopatra*, *Flirtation Walk*, *The Gay Divorcee*, *Imitation of Life*, *Viva Villa!*, *The White Parade*

- **1935 (8th)**: *Naughty Marietta*, *$1,000 a Minute*, *Bride of Frankenstein*, *Captain Blood*, *The Dark Angel*, *I Dream Too Much*, *Lost Horizon*, *Maytime*, *Thanks a Million*
- **1936 (9th)**: *San Francisco*, *Banjo on My Knee*, *The Charge of the Light Brigade*, *Dodsworth*, *General Spanky*, *Mr. Deeds Goes to Town*, *The Texas Rangers*, *That Girl from Paris*, *You Can't Take It with You*
- **1937 (10th)**: *The Hurricane*, *The Girl Said No*, *Hitting a New High*, *In Old Chicago*, *The Life of Emile Zola*, *Lost Horizon*, *Maytime*, *Topper*, *Wells Fargo*

1940s
- **1940 (13th)**: *Strike Up the... </tool_response>

Sub Agent: agent_001

Task: List all films nominated for the Academy Award for Best Sound.

Tool Call 1:

<think> Okay, I need to list all films that were nominated for the Academy Award for Best Sound. Let me think about how to approach this. First, the user is asking... </think>
<tool_call> {"name":"search","arguments":{"query":"Academy Award for Best Sound nominations","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user is asking for a list of all films nominated for the Academy Award for Best Sound. Let me start by understanding the search results they provided.... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/99028"}} </tool_call>
<tool_response> ... </tool_response>

Response

<think> Okay, so the user wants a list of all films nominated for the Academy Award for Best Sound. Let me start by understanding the information I have from the search... </think>
<agent_response> Here is a list of films nominated for the Academy Award for Best Sound, based on the historical data available in the provided Wikipedia content (up to the 1970s):

###... </agent_response>

Sub Agent: agent_002

Task: List all films that received the National Board of Review: Top Ten Films award.

Tool Call 1:

<think> Okay, the user wants a list of all films that have won the National Board of Review's Top Ten Films award. First, I need to figure out how to get... </think>
<tool_call> {"name":"search","arguments":{"query":"National Board of Review Top Ten Films award winners","topk":20}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user wants a list of all films that received the National Board of Review: Top Ten Films award. Let me check the search results they provided. Looking... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/8644529"}} </tool_call>
<tool_response> ... </tool_response>

Response

<think> Okay, the user is asking for a list of all films that received the National Board of Review: Top Ten Films award. Let me look through the information I have.... </think>
<agent_response> Here is a list of all films that received the **National Board of Review: Top Ten Films** award, organized by year (1929–2024):

1920s
- **1929**: *Applause*, *Broadway*, *Bulldog Drummond*,... </agent_response>

Round 2 (Main Agent + Sub Agents: agent_003, agent_004)

Main Agent

<think> Okay, so the user wants a list of films that were nominated for the Academy Award for Best Sound and also received the National Board of Review: Top Ten Films award. But there's a catch: we need to exclude any films that have Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter. Then, for each qualifying film, we need to gather a bunch of attributes like color, country of origin, producer, etc. First, I need to find the intersection of the two lists. The first sub-agent gave me the list of Best Sound nominees from the 1930s to 1970s, and the second sub-agent provided the Top Ten Films from the National Board of Review from 1929 to 2024. So, I need to compare these two lists to find films that are present in both. But looking at the reports, the Best Sound list from the first agent has some entries that are studio sound departments, like "Paramount Publix Studio Sound Department." Those might not be actual films, so I should probably ignore those. The user mentioned "films," so maybe those entries are errors or not applicable. I should check each entry to see if they are actual films. Next, I need to cross-reference the two lists. Let me start by looking for common films in both lists. For example, in the 1930s, the Best Sound list includes "The Love Parade" and "Viva Villa!" which also appear in the National Board of Review list. But I need... </think>
<tool_call> {"name":"create_sub_agent","arguments":{"tasks":[{"agent_id":"agent_003","task":"Identify films that are common between the Academy Award for Best Sound nominees and the National Board of Review: Top Ten Films award list, from 1930s to 1970s."},{"agent_id":"agent_004","task":"Check if any of the common films listed in agent_003's report include Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter."}]}} </tool_call> <tool_response> Sub-agents for call chatcmpl-tool-

a9815765f0ec4795b345b11a27db9fbd completed. Reports:

— Report from agent_003 —
The films common between the Academy Award for Best Sound nominees and the National Board of Review's Top Ten Films from the 1930s to 1970s are as follows:

1930s
- **1930**:
- *The Love Parade* (National Board) – Academy Award nominee (1930).
- **1931**:
- *City Lights* (National Board) – Academy Award nominee (1931).
- **1932**:
- *I Am a Fugitive from a Chain Gang* (National Board) – Academy Award nominee (1932).
- **1933**:
- *Topaze* (National Board) – Academy Award nominee (1933).
- **1934**:
- *It Happened One Night* (National Board) – Academy Award nominee (1934).
- **1935**:
- *The Informer* (National Board) – Academy Award nominee (1935).
- **1936**:
- *Mr. Deeds Goes to Town* (National Board) – Academy Award nominee (1936).
- **1937**:
- *A Star Is Born* (National Board) – Academy Award nominee (1937).
- **1938**:
- *Snow White and the Seven Dwarfs* (National Board) – Academy Award nominee (1939).
- **1939**:
- *Stagecoach* (National Board) – Academy Award nominee (1939).

1940s
- **1940**:
- *The Grapes of Wrath* (National Board) – Academy Award nominee (1940).
- **1941**:
- *Citizen Kane* (National Board) – Academy Award nominee (1941).
- **1942**:
- *Casablanca* (National Board) – Academy Award nominee (1942).
- **1943**:
- *Casablanca* (National Board) – Academy Award nominee (1943).
- **1944**:
- *The Miracle of Morgan's Creek* (National Board) – Academy Award nominee (1944).
- **1945**:
- *The Lost Weekend* (National Board)... `</tool_response>`

962

Sub Agent: agent_003

Task: Identify films that are common between the Academy Award for Best Sound nominees and the National Board of Review: Top Ten Films award list, from 1930s to 1970s.

Tool Call 1:

<think> Okay, so I need to find films that are common between the Academy Award for Best Sound nominees and the National Board of Review's Top Ten Films list from the... </think>

<tool_call> {"name":"search","arguments":{"query":"National Board of Review Top Ten Films 1930s to 1970s","topk":10}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user wants to find films that are common between the Academy Award for Best Sound nominees and the National Board of Review's Top Ten Films list from... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/99028"}} </tool_call>

<tool_response> ... </tool_response>

Response

<think> Okay, so the user wants to find films that are common between the Academy Award for Best Sound nominees and the National Board of Review's Top Ten Films list from... </think>

<agent_response> The films common between the Academy Award for Best Sound nominees and the National Board of Review's Top Ten Films from the 1930s to 1970s are as follows:

1930s </agent_response>

Sub Agent: agent_004

Task: Check if any of the common films listed in agent_003's report include Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter.

Tool Call 1:

<think> Okay, I need to figure out if Constantin Alexandrov is a cast member in any of the common films listed in agent_003's report or if Anna Hamilton Phelan is a... </think>
<tool_call> {"name":"search","arguments":{"query":"agent_003's report common films","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user wants to check if any of the common films listed in agent_003's report include Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a... </think>
<tool_call> {"name":"search","arguments":{"query":"Constantin Alexandrov filmography Anna Hamilton Phelan screenwriting credits","topk":5}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, let's see. The user wants to check if any of the common films listed in agent_003's report include Constantin Alexandrov as a cast member or Anna Hamilton Phelan as... </think>
<tool_call> {"name":"search","arguments":{"query":"Constantin Alexandrov filmography","topk":5}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, so the task is to check if any of the common films listed in agent_003's report include Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/18716679"}} </tool_call>
<tool_response> ... </tool_response>

Response

<think> Okay, so the task is to check if any of the common films listed in agent_003's report include Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a... </think>
<agent_response> After reviewing the search results and full content of the films listed in agent_003's report, there is no evidence that **Constantin Alexandrov** appears as a cast member or **Anna Hamilton... </agent_response>

Round 3 (Main Agent + Sub Agents: agent_005, agent_006, agent_007, agent_008)

Main Agent

<think> Okay, so the user wants all films that were nominated for the Academy Award for Best Sound and also received the National Board of Review: Top Ten Films award, but exclude any that have Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter. Then, they need a bunch of attributes for each film. First, I need to find the intersection of the two lists. From the first sub-agent, agent_001 gave the Best Sound nominees, and agent_002 gave the National Board Top Ten Films. Then agent_003 found the common films between those two lists. The result was a list of films from 1930s to 1970s. But then agent_004 checked if any of those films

include Constantin Alexandrov or Anna Hamilton Phelan. The report said no, so those films are safe. But I should make sure that agent_004's conclusion is accurate. Maybe double-check some of the films mentioned. For example, *Casablanca* is in both lists. Let me verify if Constantin Alexandrov was in *Casablanca*. Wait, the report says agent_004 checked and found no mention of him. Similarly, Anna Hamilton Phelan isn't listed as a screenwriter for any of them. Now, the next step is to gather all the required attributes for each of these films. The columns are film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating. But... **</think>**
<tool_call> {"name":"create_sub_agent","arguments":{"tasks":[{"agent_id":"agent_005","task":"Retrieve detailed attributes for *It Happened One Night* (1934): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."},{"agent_id":"agent_006","task":"Retrieve detailed attributes for *Casablanca* (1942): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."},{"agent_id":"agent_007","task":"Retrieve detailed attributes for *The Grapes of Wrath* (1940): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."},{"agent_id":"agent_008","task":"Retrieve detailed attributes for *Citizen Kane* (1941): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."}]}} **</tool_call>**
<tool_response> Sub-agents for call chatcmpl-tool-b83f2882255149a29a02640892fa7924 completed. Reports:

— Report from agent_005 —
The detailed attributes for *It Happened One Night* (1934) are as follows:

Film Attributes
- **Title**: *It Happened One Night*
- **Color**: Black-and-white (typical for 1934 films).
- **Country of Origin**: United States.
- **Producer**: Frank Capra and Harry Cohn.
- **Review Score**:
- **Rotten Tomatoes**: 98% score, 9.1/10 average.
- **Metacritic**: 87/100.
- **CNC Film Rating (France)**: Not explicitly mentioned in available sources. France's CNC (Commission Nationale du Cinema) ratings may require checking French-specific databases or archives.
- **Assessment**: Widely regarded as one of the greatest films ever made, winning all five major Academy Awards (Best Picture, Best Director, Best Actor, Best Actress, Best Adapted Screenplay).
- **Narrative Location**: Florida (Ellie's escape), New York City, and cross-country bus journey.
- **Distributed By**: Columbia Pictures.

- **Original Language**: English.
- **Composer**: Howard Jackson and Louis Silvers.
- **Distribution Format**: Theatrical release (standard for 1934).
- **Box Office**: $2.5 million (worldwide rentals).
- **Director**: Frank Capra.
- **Kijkwijzer Rating (Netherlands)**: Not found in available sources. Dutch ratings would need a specialized database.
- **Production Designer**: Not explicitly listed in the Wikipedia infobox. Research into production design credits or industry archives may reveal this.
- **Duration**: 105 minutes.
- **Filming Location**: Part of the film was shot on Thousand Oaks Boulevard in Thousand Oaks, California.
- **Film Editor**: Gene Havlick.
- **Production Company**: Columbia Pictures.
- **FSK Film Rating (Germany)**: Not specified in sources. Germany's FSK ratings would require checking... `</tool_response>`

966

Sub Agent: agent_005

Task: Retrieve detailed attributes for *It Happened One Night* (1934): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating.

Tool Call 1:

<think> Okay, so the user wants detailed attributes for the 1934 film "It Happened One Night." Let me start by figuring out what information I need to retrieve. The task includes... </think>

<tool_call> {"name":"search","arguments":{"query":"It Happened One Night 1934 detailed attributes","topk":10}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user wants detailed attributes for the 1934 film "It Happened One Night." Let me look through the search results they provided. The first result is the Wikipedia... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/44696"}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, let me start by looking at the user's request. They want detailed attributes for the 1934 film "It Happened One Night." The attributes listed include film, color, country of... </think>

<tool_call> {"name":"search","arguments":{"query":"It Happened One Night 1934 FSK rating","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, let me figure out how to approach this. The user wants detailed attributes for the 1934 film "It Happened One Night." So far, I've used the search tool to... </think>

<tool_call> {"name":"search","arguments":{"query":"It Happened One Night 1934 production designer","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 5:

<think> Okay, so the user is asking for detailed attributes of the 1934 film "It Happened One Night." Let me go through what they need. The attributes listed include film, color,... </think>

<tool_call> {"name":"search","arguments":{"query":"It Happened One Night 1934 production designer","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 6:

<think> Okay, let's see. The user is asking for detailed attributes of the 1934 film "It Happened One Night." They want info on film, color, country of origin, producer, review score,... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/44696"}} </tool_call>

<tool_response> ... </tool_response>

Response

<think> Okay, let me try to figure out how to get all the detailed attributes for "It Happened One Night" (1934). The user is asking for a bunch of specific info:... </think>

<agent_response> The detailed attributes for *It Happened One Night* (1934) are as follows:

Film Attributes

Sub Agent: agent_006

Task: You have reached the token limit. Please summarize your relevant findings immediately based on existing information.

Tool Call 1:

<think> Okay, I need to retrieve detailed attributes for the 1942 film Casablanca. Let me start by breaking down the user's request. They want information on various aspects like film, color,... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 detailed attributes","topk":10}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user wants detailed attributes for the 1942 film "Casablanca." Let me look through the search results they provided. The first result is about 1942 in film, which... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/57585"}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, let's see. The user wants detailed attributes for the 1942 film Casablanca. The search results gave me a few documents, but the most relevant one is the Wikipedia page... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 CNC rating France","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, let's tackle this task step by step. The user wants detailed attributes for the 1942 film Casablanca, including things like color, country of origin, producer, review score, CNC film... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 Kijkwijzer rating","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 5:

<think> Okay, let's tackle this task step by step. The user wants detailed attributes for the 1942 film "Casablanca." So far, I've used the search and open_page tools to get information... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 CNC film rating France","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 6:

<think> Okay, let's see. The user is asking for detailed attributes of the 1942 film Casablanca. They want information on film, color, country of origin, producer, review score, CNC film rating... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 FSK film rating","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 7:

<think> Okay, let's see. The user wants detailed attributes for the 1942 film Casablanca. They listed a bunch of specific fields: film, color, country of origin, producer, review score, CNC film... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 CNC rating France","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Sub Agent: agent_006 (continued)

Tool Call 8:

<think> Okay, let's see. The user wants detailed attributes for the 1942 film Casablanca. They listed a bunch of specific pieces of information: film, color, country of origin, producer, review score,... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 production designer","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 9:

<think> Okay, let's tackle this task step by step. The user is asking for detailed attributes of the 1942 film "Casablanca." They want information on film, color, country of origin, producer,... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/57585"}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 10:

<think> Okay, let me tackle this task. The user wants detailed attributes for the 1942 film Casablanca. They listed a bunch of categories like film, color, country of origin, producer, review... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 production designer","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 11:

<think> Okay, let's see. The user wants detailed attributes for the 1942 film Casablanca. They listed a bunch of categories: film, color, country of origin, producer, review score, CNC film rating... </think>

<tool_call> {"name":"search","arguments":{"query":"Casablanca 1942 production designer","topk":1}} </tool_call>

<tool_response> reached the token limit </tool_response>

Response

<think> Okay, let's tackle this. The user wants detailed attributes for the 1942 film Casablanca. They listed a bunch of categories like film, color, country of origin, producer, etc. First, I... </think>

<agent_response> Based on the available information from the Wikipedia page for *Casablanca* (1942), here is a comprehensive summary of its detailed attributes:

Film Details
- **Title**: Casablanca
- **Type**: American... </agent_response>

Sub Agent: agent_007

Task: You have reached the token limit. Please summarize your relevant findings immediately based on existing information.

Tool Call 1:

<think> Okay, I need to find detailed attributes for the 1940 film "The Grapes of Wrath." Let me start by breaking down the task. The user is asking for a lot... </think>
<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, let's tackle this task. The user wants detailed attributes for the 1940 film "The Grapes of Wrath." First, I need to figure out which of the search results are... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/4350811"}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, let's see. The user wants detailed attributes for the 1940 film "The Grapes of Wrath." They listed a bunch of specific fields like film, color, country of origin, producer,... </think>
<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 FSK rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, so the user wants detailed attributes for the 1940 film "The Grapes of Wrath." Let me start by going through the information I have. First, from the initial search... </think>
<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 production designer","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 5:

<think> Okay, so I need to retrieve detailed attributes for the 1940 film "The Grapes of Wrath." Let me start by recalling the information I already have from the previous searches.... </think>
<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 CNC rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 6:

<think> Okay, so the user is asking for detailed attributes of the 1940 film "The Grapes of Wrath." Let me see what information I have so far. From the previous search... </think>
<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 production designer","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 7:

<think> Okay, so the user is asking for detailed attributes of the 1940 film "The Grapes of Wrath." Let me start by going through the information I have. First, from the... </think>
<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 production designer","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Sub Agent: agent_007 (continued)

Tool Call 8:

<think> Okay, let's see. The user wants detailed attributes for the 1940 film "The Grapes of Wrath." They've already done a search and opened the Wikipedia page for the film. The... </think>

<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 production designer","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 9:

<think> Okay, let's tackle this task step by step. The user wants detailed attributes for "The Grapes of Wrath" (1940), including film, color, country of origin, producer, review score, CNC film... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/4350811"}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 10:

<think> Okay, let me see. The user is asking for detailed attributes of the 1940 film "The Grapes of Wrath." They listed a bunch of categories like film, color, country of... </think>

<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 production designer","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 11:

<think> Okay, so I need to retrieve detailed attributes for "The Grapes of Wrath" (1940) film. Let me start by recalling the user's query. They want information on film, color, country... </think>

<tool_call> {"name":"search","arguments":{"query":"The Grapes of Wrath 1940 production designer","topk":1}} </tool_call>

<tool_response> reached the token limit </tool_response>

Response

<think> Okay, so the user wants detailed attributes for the 1940 film "The Grapes of Wrath." Let me start by recalling the information I have from the previous searches and the... </think>

<agent_response> ### Detailed Attributes for *The Grapes of Wrath* (1940)

1. **Film**:

- **Title**: *The Grapes of Wrath* (1940)
- **Type**: Drama / Adaptation of John Steinbeck's novel.
- **Country of... </agent_response>

Sub Agent: agent_008

Task: Retrieve detailed attributes for *Citizen Kane* (1941): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating.

Tool Call 1:

<think> Okay, I need to retrieve detailed attributes for the film *Citizen Kane* (1941). The user has listed a bunch of specific pieces of information they want, so I should start... </think>
<tool_call> {"name":"search","arguments":{"query":"Citizen Kane 1941","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, let me start by understanding the task. The user wants detailed attributes for the film "Citizen Kane" (1941). They listed a bunch of specific pieces of information: color, country... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/5224"}} </tool_call>
<tool_response> ... </tool_response>

Response

<think> Okay, let me start by figuring out what the user needs. They want detailed attributes for the film *Citizen Kane* (1941), including things like color, country of origin, producer, review... </think>
<agent_response> Here is the detailed information about *Citizen Kane* (1941):

1. **Film**: *Citizen Kane*
2. **Color**: Black and white (though the innovative 1940 trailer was in color).
3. **Country of Origin**:... </agent_response>

Round 4 (Main Agent + Sub Agents: agent_009, agent_010)

Main Agent

<think> Okay, so the user is asking for all films that were nominated for the Academy Award for Best Sound and received the National Board of Review: Top Ten Films award, excluding any films with Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter. Then, they want all the specified attributes for those films. First, I need to find the intersection of the two lists: Academy Award Best Sound nominees and National Board of Review Top Ten Films. The previous agents did that for the 1930s to 1970s and found some common films like *It Happened One Night*, *Casablanca*, *The Grapes of Wrath*, and *Citizen Kane*. Agent_004 checked if any of these have the excluded names and found none, so those four films are still in the running. Now, the user wants detailed attributes for each of these films. Agents 005, 006, 007, and 008 retrieved info for each, but some details are missing, like regional ratings (CNC, Kijkwijzer, FSK) and production designers. The user wants all information, using '/' for unavailable data. I need to compile the data from these reports into the required markdown table. Let me check each film's details: For *It Happened One Night* (1934): - Color: Black-and-white - Country of Origin: United States - Producer: Frank Capra and Harry Cohn - Review Score: 98% on Rotten Tomatoes, 9.1/10 average - CNC Film Rating: Not found - Assessment: Won 5 Academy Awards - Narrative

Location: Florida, New York City, cross-country bus journey -... </think>
<tool_call> {"name":"create_sub_agent","arguments":{"tasks":[{"agent_id":"agent_009","task":"Retrieve detailed attributes for *The Life of Emile Zola* (1937): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."},{"agent_id":"agent_010","task":"Retrieve detailed attributes for *The Best Years of Our Lives* (1946): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."}]}} </tool_call> <tool_response> Sub-agents for call chatcmpl-tool-82a2e428e1c443c29b5e1a8683cbeb7b completed. Reports:

— Report from agent_009 —
Detailed Attributes for *The Life of Emile Zola* (1937):

1. **Film**: Biographical drama based on Émile Zola's life, focusing on the Dreyfus Affair and his literary career.
2. **Color**: Black-and-white (standard for 1937 films).
3. **Country of Origin**: United States.
4. **Producer**: Henry Blanke (produced by Warner Bros.).
5. **Review Score**:
- **Rotten Tomatoes**: 92% from 72 reviews (average score: 7.5/10).
- **Critics Consensus**: Praised for its historical significance, Paul Muni's performance, and thematic relevance to 1930s political tensions.
6. **CNC Film Rating (France)**: Not explicitly mentioned in available sources. The film's connection to the Dreyfus Affair may have influenced its reception in France, but no official CNC rating is listed.
7. **Assessment**: Regarded as a landmark biographical film, winning three Oscars (Best Picture, Best Supporting Actor, Best Screenplay). Critics highlighted its moral urgency and Paul Muni's portrayal of Zola.
8. **Narrative Location**: Set in 19th-century France, spanning Paris, French Guiana (Devil's Island), and other locations tied to Zola's life and the Dreyfus Affair.
9. **Distributed By**: Warner Bros. Pictures.
10. **Original Language**: English (though the story is about a French author and historical events).
11. **Composer**: Max Steiner (score awarded to Leo F. Forbstein).
12. **Distribution Format**: Theatrical release (no mention of home media formats).
13. **Box Office**: Specific revenue figures are not provided in available sources.
14. **Director**: William Dieterle.
15. **Kijkwijzer Rating (Netherlands)**: Not mentioned in sources.
16. **Production Designer**: Not explicitly listed in the... </tool_response>

Sub Agent: agent_009

Task: Retrieve detailed attributes for *The Life of Emile Zola* (1937): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating.

Tool Call 1:

<think> Okay, I need to find detailed attributes for the 1937 film "The Life of Emile Zola." Let's start by breaking down the task. The user is asking for a lot... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 film details","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user is asking for detailed attributes of the 1937 film "The Life of Emile Zola." Let me start by looking at the search results they provided. The... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/44210"}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, let me start by going through the information provided from the open_page tool. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." The first... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 CNC rating FSK rating box office","topk":3}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, so I need to find detailed attributes for the 1937 film "The Life of Emile Zola." The user is asking for a lot of information: film, color, country of... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/44210"}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 5:

<think> Okay, let's see. The user is asking for detailed attributes of the 1937 film "The Life of Emile Zola." They want information on film, color, country of origin, producer, review... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 box office revenue CNC rating FSK rating Kijkwijzer rating","topk":3}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 6:

<think> Okay, let me tackle this task step by step. The user is asking for detailed attributes of the 1937 film "The Life of Emile Zola." They want information on film,... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 box office revenue FSK rating CNC rating Kijkwijzer rating","topk":3}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 7:

<think> Okay, let's tackle this query. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." They listed a bunch of specific fields: film, color, country, producer,... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 box office revenue FSK rating CNC rating Kijkwijzer rating","topk":3}} </tool_call>
<tool_response> ... </tool_response>

Sub Agent: agent_009 (continued)

Tool Call 8:

<think> Okay, so the user is asking for detailed attributes of the 1937 film "The Life of Emile Zola." They want information on film, color, country of origin, producer, review score,... </think>

<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 box office revenue FSK rating CNC rating Kijkwijzer rating","topk":3}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 9:

<think> Okay, so I need to find detailed attributes for the 1937 film "The Life of Emile Zola." The user is asking for a bunch of specific info: film, color, country... </think>

<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 box office revenue FSK rating CNC rating Kijkwijzer rating","topk":3}} </tool_call>

<tool_response> ... </tool_response>

Response

<think> Okay, so I need to retrieve detailed attributes for the 1937 film "The Life of Emile Zola." The user is asking for a list of specific information including film, color,... </think>

<agent_response> ### Detailed Attributes for *The Life of Emile Zola* (1937):

1. **Film**: Biographical drama based on Émile Zola's life, focusing on the Dreyfus Affair and his literary career.
2. **Color**:... </agent_response>

Sub Agent: agent_010

Task: You have reached the token limit. Please summarize your relevant findings immediately based on existing information.

Tool Call 1:

<think> Okay, I need to retrieve detailed attributes for the movie "The Best Years of Our Lives" (1946). Let me start by breaking down the task. The user is asking for... </think>
<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user is asking for detailed attributes of the film "The Best Years of Our Lives" from 1946. Let me start by looking at the search results they... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/42159"}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, so the user wants detailed attributes for "The Best Years of Our Lives" (1946). Let me start by going through the information I have from the search results and... </think>
<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, let me see. The user wants detailed attributes for "The Best Years of Our Lives" (1946). They listed a bunch of categories like film, color, country of origin, producer,... </think>
<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 5:

<think> Okay, let's tackle this query step by step. The user wants detailed attributes for "The Best Years of Our Lives" (1946). They listed a bunch of categories like film, color,... </think>
<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 6:

<think> Okay, so the user is asking for detailed attributes of the 1946 film "The Best Years of Our Lives." Let me start by going through the information I have from... </think>
<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 7:

<think> Okay, let's tackle this task step by step. The user wants detailed attributes for the film "The Best Years of Our Lives" (1946). They listed a bunch of specific pieces... </think>
<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Sub Agent: agent_010 (continued)

Tool Call 8:

<think> Okay, so the user is asking for detailed attributes of the film "The Best Years of Our Lives" from 1946. They listed a bunch of specific pieces of information they... </think>

<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 9:

<think> Okay, let's tackle this query step by step. The user wants detailed attributes for "The Best Years of Our Lives" (1946). From the previous search results, I found the Wikipedia... </think>

<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 10:

<think> Okay, so I need to retrieve detailed attributes for "The Best Years of Our Lives" (1946). Let's start by looking at the information I have from the previous searches. First,... </think>

<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 11:

<think> Okay, so I need to retrieve detailed attributes for "The Best Years of Our Lives" (1946). Let me start by going through the information I have so far. First, the... </think>

<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 CNC rating","topk":1}} </tool_call>

<tool_response> reached the token limit </tool_response>

Response

<think> Okay, let me try to figure this out. The user is asking for detailed attributes of the 1946 film "The Best Years of Our Lives." They listed a bunch of... </think>

<agent_response> ### Summary of *The Best Years of Our Lives* (1946)

*Film**: *The Best Years of Our Lives* (also known as *Glory for Me* and *Home Again*).

*Color**: Color.

*Country of... </agent_response>

Round 5 (Main Agent + Sub Agents: agent_011, agent_012)

Main Agent

<think> Okay, let me start by understanding the user's question again. They want all films that were nominated for the Academy Award for Best Sound and received the National Board of Review: Top Ten Films award, excluding any films that feature Constantin Alexandrov as a cast member or have Anna Hamilton Phelan as a screenwriter. Then, they need specific attributes for each of these films. First, I need to find the intersection of the two lists: Academy Award Best Sound nominees and National Board of Review Top Ten Films. From the previous reports, agent_003 provided a list of common films between these two categories from the 1930s to 1970s. That list

includes films like *It Happened One Night*, *Casablanca*, *The Grapes of Wrath*, *Citizen Kane*, and others. Next, I have to check if any of these films include Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter. Agent_004 checked this and found no evidence of either name in the listed films. So, the initial list seems okay, but maybe I should verify a couple of films just to be thorough, especially if there's any chance of a misspelling or alternative names. However, since the report says no matches, I can proceed. Now, the user wants detailed attributes for each film. Agents 005 to 010 have retrieved information for *It Happened One Night*, *Casablanca*, *The Grapes of Wrath*, *Citizen Kane*, and *The Life of Emile Zola* and *The Best Years of Our Lives*. Each of these reports provides... **</think>**
<tool_call> {"name":"create_sub_agent","arguments":{"tasks":[{"agent_id":"agent_011","task":"Retrieve detailed attributes for *The Best Years of Our Lives* (1946): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."},{"agent_id":"agent_012","task":"Retrieve detailed attributes for *The Life of Emile Zola* (1937): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."}]}} **</tool_call>** **<tool_response>** Sub-agents for call chatcmpl-tool-989f200e99584922b8e6826a3e5e742c completed. Reports:

— Report from agent_011 —
The Best Years of Our Lives (1946) is a 172-minute American drama film directed by William Wyler. Here are the detailed attributes based on available information:

1. **Film**: Drama
2. **Color**: Color (as cinematographer Gregg Toland used deep focus and color photography).
3. **Country of Origin**: United States
4. **Producer**: Samuel Goldwyn
5. **Review Score**: 97% "Fresh" rating on Rotten Tomatoes (8.9/10 average).
6. **CNC Film Rating (France)**: Not explicitly mentioned in available sources.
7. **Assessment**: Critically acclaimed as a masterpiece, praised for its realistic portrayal of WWII veterans' struggles and ensemble casting. Roger Ebert called it "modern, lean, and honest."
8. **Narrative Location**: Boone City (fictional, based on Cincinnati, Ohio).
9. **Distributed By**: RKO Radio Pictures
10. **Original Language**: English
11. **Composer**: Hugo Friedhofer (with Emil Newman as musical director).
12. **Distribution Format**: Theatrical release (initially limited to six shows per day due to length).
13. **Box Office**: $23.7 million (adjusted for inflation, one of the top 100 grossing films in U.S. history).
14. **Director**: William Wyler
15. **Kijkwijzer Rating (Dutch)**: Not found in available sources.
16. **Production Designer**: Not explicitly listed in the text, but Wyler emphasized realistic sets and life-size constructions.
17. **Duration**: 172 minutes
18. **Filming Locations**: Los Angeles County Arboretum, Ontario International Airport (California), Raleigh Studios, Samuel Goldwyn/Warner Hollywood Studios.
19. **Film Editor**: Daniel Mandell
20. **Production Company**: Samuel Goldwyn Productions

978

21. **FSK Film Rating (Germany)**: Not mentioned in available sources.

Notes:
- The... </tool_response>

Sub Agent: agent_011

Task: Retrieve detailed attributes for *The Best Years of Our Lives* (1946): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating.

Tool Call 1:

<think> Okay, I need to retrieve detailed attributes for the film "The Best Years of Our Lives" from 1946. Let me start by breaking down the user's query. They want information... </think>

<tool_call> {"name":"search","arguments":{"query":"The Best Years of Our Lives 1946 detailed attributes","topk":10}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user wants detailed attributes for the film "The Best Years of Our Lives" from 1946. Let me start by looking through the search results provided. The first... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/42159"}} </tool_call>

<tool_response> ... </tool_response>

Response

<think> Okay, let's start by looking at the user's query. They want detailed attributes for "The Best Years of Our Lives" (1946). The attributes listed include film, color, country of origin,... </think>

<agent_response> The Best Years of Our Lives (1946) is a 172-minute American drama film directed by William Wyler. Here are the detailed attributes based on available information:

1. **Film**: Drama
2.... </agent_response>

Sub Agent: agent_012

Task: You have reached the token limit. Please summarize your relevant findings immediately based on existing information.

Tool Call 1:

<think> Okay, I need to find detailed attributes for the 1937 film "The Life of Emile Zola." Let's start by breaking down the user's query. They want information on various aspects... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola 1937 film details","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user is asking for detailed attributes of the 1937 film "The Life of Emile Zola." Let me start by looking at the search results they provided. The... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/44210"}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, let's tackle this task. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." First, I need to go through the information provided in the... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/4970252"}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, let's see. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." So far, I've used the search tool and found the Wikipedia page for... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola CNC rating France","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 5:

<think> Okay, let's see. The user is asking for detailed attributes of the 1937 film "The Life of Emile Zola." So far, I've used the search and open_page tools to get... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola CNC rating France","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 6:

<think> Okay, let's tackle this step by step. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." So far, I've used the search and open_page tools... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola FSK rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 7:

<think> Okay, let me try to figure out how to get all the details for "The Life of Emile Zola" (1937). The user is asking for a bunch of attributes like... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola Kijkwijzer rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Sub Agent: agent_012 (continued)

Tool Call 8:

<think> Okay, let's tackle this task step by step. The user is asking for detailed attributes of the 1937 film "The Life of Emile Zola." They want information on film, color,... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola CNC rating France","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 9:

<think> Okay, let's tackle this task step by step. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." So far, I've used the search and open_page... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola box office","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 10:

<think> Okay, let's tackle this task step by step. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." So far, I've used the search and open_page... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola FSK rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 11:

<think> Okay, so I need to retrieve detailed attributes for the 1937 film "The Life of Emile Zola." Let me start by listing all the required information: film, color, country of... </think>
<tool_call> {"name":"search","arguments":{"query":"The Life of Emile Zola CNC rating France","topk":1}} </tool_call>
<tool_response> reached the token limit </tool_response>

Response

<think> Alright, let's tackle this task step by step. The user wants detailed attributes for the 1937 film "The Life of Emile Zola." So far, I've used the search and open_page... </think>
<agent_response> ### Summary of **The Life of Emile Zola** (1937) Attributes:
1. **Film**: Biographical drama based on the life of 19th-century French author Émile Zola.
2. **Color**: Black-and-white (as it was... </agent_response>

Round 6 (Main Agent + Sub Agents: agent_013, agent_014)

Main Agent

<think> Okay, so the user is asking for a list of films that were nominated for the Academy Award for Best Sound and also won the National Board of Review: Top Ten Films award, but they need to exclude any films that have Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter. Then, they want all those films with specific attributes like color, country of origin, etc. First, I need to make sure I have the correct list of films that meet both awards. From the previous reports, agents 003, 005, 006, 007, 008, 010, and 011 have provided some of these films. Let me check if there's any overlap between the Academy Award Best Sound nominees and the

National Board of Review's Top Ten Films. Looking at the reports, agent_003's list includes films like *It Happened One Night*, *Casablanca*, *The Grapes of Wrath*, *Citizen Kane*, *The Best Years of Our Lives*, and *The Life of Emile Zola*. These are the ones that are common in both lists. Now, I need to check if any of these films include Constantin Alexandrov or Anna Hamilton Phelan. Agent_004's report says there's no mention of these names in the films listed, so that's a relief. But I should still double-check each film's cast and crew to be thorough. Next, for each film, I need to gather all the specified attributes. Agents 005, 006, 007, 008, 010, and 011 have provided detailed info on these films, but some details are missing... **</think>**
<tool_call> {"name":"create_sub_agent","arguments":{"tasks":[{"agent_id":"agent_013","task":"Retrieve detailed attributes for *A Star Is Born* (1937): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."},{"agent_id":"agent_014","task":"Retrieve detailed attributes for *The Hurricane* (1937): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating."}]}} **</tool_call>** **<tool_response>** Sub-agents for call chatcmpl-tool-0a2be8050ff54e8a8bd2deb5ac4f7056 completed. Reports:

— Report from agent_013 —
Detailed Attributes for *A Star Is Born* (1937)
- **Film**: Drama/musical
- **Color**: Technicolor
- **Country of Origin**: United States
- **Producer**: David O. Selznick
- **Review Score**: 100% on Rotten Tomatoes (based on 23 reviews, average 7.7/10).
- **CNC Film Rating (France)**: **Not found** in available sources.
- **Kijkwijzer Rating (Netherlands)**: **Not found** in available sources.
- **FSK Film Rating (Germany)**: **Not found** in available sources.
- **Assessment**: Praised as a groundbreaking color film and a powerful human interest drama. Won the Academy Award for Best Original Story.
- **Narrative Location**: Hollywood, California.
- **Distributed By**: United Artists
- **Original Language**: English
- **Composer**: Max Steiner
- **Distribution Format**: Blu-ray (restored 4K version released in 2022).
- **Box Office**: $2.37 million in worldwide rentals.
- **Director**: William A. Wellman
- **Production Designer**: **Not explicitly mentioned** in the source material.
- **Duration**: 111 minutes
- **Filming Location**: Hollywood, California (implied by the plot and production details).
- **Film Editor**: James E. Newcom and Anson Stevenson (uncredited).
- **Production Company**: Selznick International Pictures
- **FSK Rating**: **Not found**; the film's content likely aligns with general theatrical standards but lacks specific classification.

Notes:
- **CNC/FSK/Kijkwijzer Ratings**: No direct information was found in accessible sources. These ratings may not have been assigned or documented for this film.

982

- **Production Designer**: The Wikipedia page does not specify this detail, and further research beyond the current tools would be required.
- **Remakes**:... `</tool_response>`

Sub Agent: agent_013

Task: Retrieve detailed attributes for *A Star Is Born* (1937): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating.

Tool Call 1:

<think> Okay, I need to find detailed attributes for the 1937 film "A Star Is Born." Let me start by breaking down the task. The user is asking for a lot... </think>
<tool_call> {"name":"search","arguments":{"query":"A Star Is Born 1937 detailed attributes","topk":10}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, let me start by understanding the user's task. They want detailed attributes for the 1937 film "A Star Is Born." The attributes listed include film, color, country of origin,... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/3450569"}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 3:

<think> Okay, let's tackle this query. The user wants detailed attributes for the 1937 film "A Star Is Born." First, I need to go through the information provided from the search... </think>
<tool_call> {"name":"search","arguments":{"query":"A Star Is Born 1937 FSK rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 4:

<think> Okay, so I need to find detailed attributes for the 1937 film "A Star Is Born." Let's start by recalling the information I already have from the previous search results.... </think>
<tool_call> {"name":"search","arguments":{"query":"A Star Is Born 1937 FSK rating","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 5:

<think> Okay, so the user is asking for detailed attributes of the 1937 film "A Star Is Born." Let me start by going through the information I have. First, from the... </think>
<tool_call> {"name":"search","arguments":{"query":"A Star Is Born 1937 production designer","topk":1}} </tool_call>
<tool_response> ... </tool_response>

Tool Call 6:

<think> Okay, so the user is asking for detailed attributes of the 1937 film "A Star Is Born." Let me start by going through the information I have. First, from the... </think>
<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/3450569"}} </tool_call>
<tool_response> ... </tool_response>

Response

<think> Okay, so the user is asking for detailed attributes of the 1937 film "A Star Is Born." Let me start by going through all the information I have gathered so... </think>
<agent_response> ### Detailed Attributes for *A Star Is Born* (1937)
- **Film**: Drama/musical
- **Color**: Technicolor
- **Country of Origin**: United States
- **Producer**: David O. Selznick
- **Review Score**: 100%... </agent_response>

Sub Agent: agent_014

Task: Retrieve detailed attributes for *The Hurricane* (1937): film, color, country of origin, producer, review score, CNC film rating (France), assessment, narrative location, distributed by, original language, composer, distribution format, box office, director, Kijkwijzer rating, production designer, duration, filming location, film editor, production company, FSK film rating.

Tool Call 1:

<think> Okay, I need to find detailed attributes for the 1937 film "The Hurricane." Let me start by breaking down the user's query. They want information on various aspects like film,... </think>

<tool_call> {"name":"search","arguments":{"query":"The Hurricane 1937 film details","topk":10}} </tool_call>

<tool_response> ... </tool_response>

Tool Call 2:

<think> Okay, so the user is asking for detailed attributes of the 1937 film "The Hurricane." Let me start by looking at the search results they provided. The first result, DocID... </think>

<tool_call> {"name":"open_page","arguments":{"docid":"enwiki/1364234"}} </tool_call>

<tool_response> ... </tool_response>

Response

<think> Okay, let's start by looking at the user's task. They want detailed attributes for the 1937 film "The Hurricane." The attributes include film, color, country of origin, producer, review score,... </think>

<agent_response> The Hurricane (1937) film details are as follows:

- **Film**: The Hurricane (1937 film)
- **Color**: Black-and-white (implied by absence of color mention in sources).
- **Country of Origin**: United... </agent_response>

Round 7 (Main Agent)

Main Agent

<think> Okay, let me see. The user wants all films that were nominated for the Academy Award for Best Sound and received the National Board of Review: Top Ten Films award, excluding any films with Constantin Alexandrov as a cast member or Anna Hamilton Phelan as a screenwriter. Then, they need detailed attributes for each of those films. First, I need to find the intersection of the two lists: Academy Award Best Sound nominees and National Board of Review Top Ten Films. From the previous reports, agent_003 provided a list of common films between the two, such as *It Happened One Night*, *Casablanca*, *The Grapes of Wrath*, *Citizen Kane*, *The Life of Emile Zola*, *The Best Years of Our Lives*, and *A Star Is Born* (1937), *The Hurricane* (1937). Next, I have to check if any of these films include Constantin Alexandrov or Anna Hamilton Phelan. Agent_004's report said there's no evidence of either name in the listed films. But just to be thorough, maybe I should cross-verify a couple of the films mentioned. For example, *Casablanca* has a stellar cast like Bogart and Bergman, but no Alexandrov. *The Grapes of Wrath* has Henry Fonda, Jane Darwell, etc., also no Alexandrov. *Citizen Kane* has Orson Welles, but again, no mention of the names. So it seems safe to proceed without excluding any films. Now, the next step is to collect all the detailed attributes for each of these films. Agents 005 to 014 have provided

information for *It Happened One Night*,... **</think>**

<answer>

Film	Color	Origin	Producer	Score	CNC	Assessment	Narrative Loc	Distributor	Lang	Composer	Format	Box Office	Director	Kijk	Designer	Dur	Filming Loc	Editor	Company	FSK
It Happened One Night	B&W	USA	F. Capra, H. Cohn	98%; 9.1; 87	/	Cultural milestone, 5 Oscars	FL, NYC	Columbia	En	H. Jackson	theatrical	$2.5m	F. Capra	/	/	105	CA	G. Havlick	Columbia	/
Casablanca	B&W	USA	Hal B. Wallis	/	/	Cultural icon, 3 Oscars	Morocco	Warner Bros.	En	M. Steiner	theatrical	$3.7m	M. Curtiz	/	/	102	CA	O. Marks	Warner Bros.	/
The Grapes of Wrath	Color*	USA	D.F. Zanuck, N. Johnson	96%; 9/10	/	Social commentary, 3 Oscars	OK, CA	20th Century	En	A. Newman	theatrical	$1.6m	J. Ford	/	/	129	OK, CA	R. Simpson	20th Century	/
Citizen Kane	B&W	USA	Orson Welles	/	/	Innovation, 1 Oscar	Xanadu, FL	RKO Radio	En	B. Herrmann	theatrical	$1.8m	O. Welles	/	P. Ferguson	119	CA	R. Wise	RKO Radio	/
The Life of Emile Zola	B&W	USA	Henry Blanke	92%; 7.5	/	Historical, 3 Oscars	France	Warner Bros.	En	M. Steiner	theatrical	/	W. Dieterle	/	A. Grot	116	USA	W. Low	Warner Bros.	/
Best Years of Our Lives	Color	USA	Samuel Goldwyn	97%; 8.9	/	Post-war, 7 Oscars	Ohio	RKO Radio	En	H. Friedhofer	theatrical	$23.7m	W. Wyler	/	/	172	CA	D. Mandell	Goldwyn Prod.	/
A Star Is Born (1937)	Technicolor	USA	David O. Selznick	100%; 7.7	/	Color film	CA	United Artists	En	M. Steiner	theatrical	$2.37m	W. Wellman	/	/	111	CA	J.E. Newcom	Selznick Int.	/
The Hurricane	B&W	USA	Samuel Goldwyn	/	/	VFX, 1 Oscar	South Seas	United Artists	En	A. Newman	theatrical	$3.2m	J. Ford	/	/	110	/	L. Nosler	Goldwyn Prod.	/

Notes:\n1. **Missing Ratings**: CNC, FSK, and Kijkwijzer ratings are not available in public sources for these films. \n2. **Production Designer**: For *It Happened One Night*, *Casablanca*, *The Grapes of Wrath*, *The Life of Emile Zola*, *The Best Years of Our Lives*, *A Star Is Born* (1937), and *The Hurricane* (1937), specific names were not explicitly listed in sources. \n3. **Box Office**: *The Grapes of Wrath* and *The Life of Emile Zola* lack precise figures, but *The Best Years of Our Lives* is noted as a top-grossing film. \n4. **Color**: *A Star Is Born* (1937) is confirmed as Technicolor; others are assumed black-and-white unless noted (e.g., *The Grapes of Wrath* is color). \n5. **Exclusions**: No films listed include Constantin Alexandrov or Anna Hamilton Phelan.

</answer>

986