

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TOWARDS PERSONALIZED DEEP RESEARCH: BENCHMARKS AND EVALUATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep Research Agents (DRAs) can autonomously conduct complex investigations and generate comprehensive reports, demonstrating strong real-world potential. However, existing evaluations mostly rely on close-ended benchmarks, while open-ended deep research benchmarks remain scarce and typically neglect personalized scenarios. To bridge this gap, we introduce **Personalized Deep Research Bench**, the first benchmark for evaluating personalization in DRAs. It pairs 50 diverse research tasks across 10 domains with 25 authentic user profiles that combine structured persona attributes with dynamic real-world contexts, yielding 250 realistic user-task queries. To assess system performance, we propose the PQR Evaluation Framework, which jointly measures (P) Personalization Alignment, (Q) Content Quality, and (R) Factual Reliability. Our experiments on a range of systems highlight current capabilities and limitations in handling personalized deep research. This work establishes a rigorous foundation for developing and evaluating the next generation of truly personalized AI research assistants.

1 INTRODUCTION

Recent advances in large language models (LLMs) have enabled the development of AI agents capable of conducting complex deep research. Early LLMs focus on isolated tasks like QA and translation, later advancing with tool integration for autonomous information retrieval and synthesis. More recently, a new class of advanced systems has emerged, known as Deep Research Agents (DRAs), including industry solutions (OpenAI, 2025b; Google DeepMind, 2025; xAI Team, 2025; Perplexity Team, 2025; Moonshot AI, 2025; ByteDance, 2025b) and open-source systems (Li et al., 2025b;c; Zhu et al., 2025a; Zhou et al., 2023; 2024; Wang et al., 2025; Hu et al., 2025; Manus AI, 2025; MiroMind AI Team, 2025; ByteDance, 2025a; Li et al., 2025a; Tang et al., 2025; Shi et al., 2025; Zhu et al., 2025b). DRAs extend LLMs by incorporating dynamic reasoning, adaptive planning, and iterative tool use to acquire, aggregate, and analyze external information (Huang et al., 2025), thereby enabling end-to-end research workflows and the production of structured, comprehensive reports.

Despite these advances, to fully realize the potential of these intelligent systems in everyday human contexts, they must be able to adapt their behaviors and interactions to the specific needs of different users (Fischer, 2001; Kirk et al., 2024; Rafieian & Yoganarasimhan, 2023), a quality known as personalization. Important real-world decisions, from choosing a vehicle to making an investment, are strongly influenced by a user’s unique needs, preferences, budget, and prior knowledge. In these scenarios, the agent’s value lies not only in generating a comprehensive report, but also in acting as a personalized assistant that tailors its information filtering, reasoning, and recommendations. However, this critical dimension of personalization is a major blind spot for current evaluation methodologies.

Existing deep research benchmarks, including close-ended suites like GAIA, BrowseComp, HLE, and X-Bench (Mialon et al., 2023; Wei et al., 2025; Phan et al., 2025; Chen et al., 2025a) and open-ended ones like DeepResearch Bench, ResearcherBench, and DeepResearchGym (Du et al., 2025; Xu et al., 2025; Coelho et al., 2025), focus exclusively on factual accuracy and comprehensiveness, failing to assess user-specific adaptation. Conversely, existing personalization benchmarks such as LaMP, PersonaGym, PersonaLens and PersonaFeedback (Salemi et al., 2024; Samuel et al., 2025; Zhao et al., 2025; Tao et al., 2025) are confined to narrow domains like dialogue or recommendation

and do not address the complex deep research. To the best of our knowledge, our work is the first to systematically incorporate personalization into the evaluation of DRAs, filling a critical gap in current research.

To address this gap, we introduce *Personalized Deep Research Bench*, a novel benchmark specifically designed to evaluate personalization in deep research agents. Our benchmark provides a rigorous framework for assessing how well agents can integrate user profiles into their research workflows, and whether their outputs are not only comprehensive and accurate, but also tailored and practically useful for the end user. By formalizing and evaluating this missing dimension, our work paves the way for the development of more effective and genuinely personal AI assistants.

Our main contributions are summarized as follows:

- We formally introduce the task of *personalized deep research*, which extends beyond generic information synthesis by requiring DRAs to adapt retrieval, reasoning and reporting to user personas.
- We propose *Personalized Deep Research Bench*, the first benchmark specifically targeting personalization in DRAs. It consists of 50 diverse tasks that span 10 domains and are paired with 25 real-world user profiles, yielding 250 unique user-task pairs, enabling systematic evaluation of both task complexity and persona-driven adaptation.
- We develop the *PQR Evaluation Framework*, a novel and comprehensive methodology that evaluates generated reports along three orthogonal dimensions: (P) *Personalization Alignment*, (Q) *Content Quality*, and (R) *Factual Reliability*, providing a holistic measure of agent utility in real-world research scenarios.
- We conduct extensive experiments across a broad spectrum of open-source DRAs, commercial deep research systems, LLMs with search tools and advancing memory systems, revealing both strengths and limitations in handling personalization.

2 RELATED WORK

2.1 EVALUATING DEEP RESEARCH CAPABILITIES

Evaluating DRAs requires benchmarks that go beyond traditional QA tasks to assess multi-turn retrieval, tool use, and structured report generation. Close-ended benchmarks such as GAIA, BrowseComp, HLE, and X-Bench (Mialon et al., 2023; Wei et al., 2025; Chen et al., 2025a) offer controlled evaluations, yet rely on synthetic tasks and fall short of reflecting the challenges of authentic research scenarios. Recently, open-ended deep research benchmarks have been proposed to specifically evaluate deep research capabilities. DeepResearch Bench (Du et al., 2025) offers 100 PhD-level tasks across 22 fields, introducing the RACE and FACT frameworks for report quality and retrieval assessment. Mind2Web 2 (Gou et al., 2025) features 130 real-world tasks with live web browsing and proposes the Agent-as-a-Judge framework for automated correctness and attribution. ResearcherBench (Xu et al., 2025) focuses on 65 frontier AI questions across 35 subjects with a dual rubric–factual evaluation. Additionally, BrowseComp-Plus (Chen et al., 2025b) extends BrowseComp (Wei et al., 2025) by pairing each query with curated documents and challenging negatives to isolate retriever and LLM contributions. DeepResearchGym (Coelho et al., 2025) provides an open-source sandbox with reproducible search APIs and standardized protocols for transparent, low-cost benchmarking. Nevertheless, these benchmarks focus on general research capabilities and lack metrics for personalization—the alignment of research with user-specific goals and preferences.

2.2 BENCHMARKING PERSONALIZATION PERFORMANCE

Meanwhile, most personalization benchmarks focus on general tasks and remain insufficient for complex deep research scenarios. LaMP (Salemi et al., 2024) introduces seven classification and generation tasks to evaluate the personalized output capacity of LLMs. PersonaGym (Samuel et al., 2025) introduces PersonaScore to evaluate the adherence of LLM agents to assigned personas at scale. PersonallLM (Zollo et al., 2025) uses reward models to act as different user personas to evaluate response preference. AI Persona (Wang et al., 2024b) concentrates on the lifelong learning of user profiles with LLM-as-a-judge evaluation. Additionally, PersonaMem (Jiang et al., 2025)

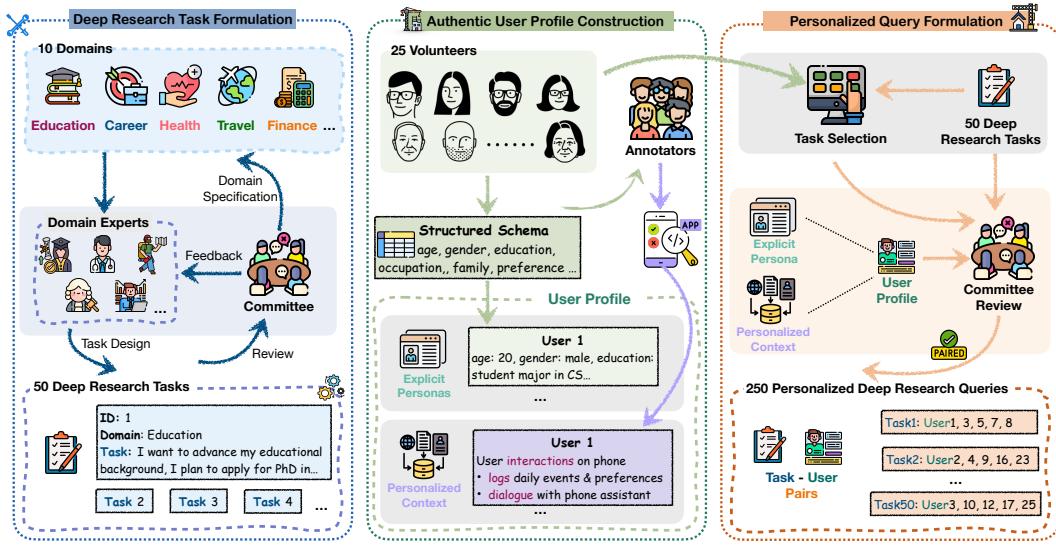


Figure 1: Benchmark Construction Pipeline: (1) Design 50 deep research tasks across 10 domains; (2) Build authentic user profiles from 25 volunteers; (3) Generate 250 personalized task-user pairs.

benchmarks the adaptability of LLMs to evolving user personas. PersonaFeedback (Tao et al., 2025) provides a large human-annotated benchmark for response tailoring to explicit personas. Person-aLens (Zhao et al., 2025) introduces LLM-based user and judge agents to assess personalization and task success in realistic dialogues.

Overall, current benchmarks either neglect personalization or fail to capture the complex nature of deep research, highlighting the pressing need for a new benchmark specifically designed to measure the personalized performance of DRAs.

3 BENCHMARK CONSTRUCTION

To rigorously evaluate the personalized capabilities of deep research agents, we introduce the Personalized Deep Research Bench, a benchmark designed to mirror the real-world personalized deep research scenarios. Its construction is grounded in two key components: a diverse set of deep research tasks and a collection of authentic, multifaceted user profiles, as shown in Figure 1.

3.1 DEEP RESEARCH TASK FORMULATION

Domain Specification and Task Generation. To begin, we defined a set of 10 distinct domains, $\mathcal{D} = \{d_1, d_2, \dots, d_{10}\}$ covering major and impactful aspects of daily life (e.g., Career Development, Education, Healthcare, Financial Planning). To ensure that the tasks within each domain are both realistic and practically relevant, we collaborated with a diverse group of domain experts, such as travel bloggers, financial advisors and educational consultants, to design the initial set of tasks.

Committee Review and Validation. Each task underwent multistage validation by a committee of Master's/PhD researchers, data scientists, and product managers, following three principles: Complexity (\uparrow): requiring multi-step reasoning, retrieval, and analysis; Clarity (\uparrow): unambiguous descriptions with clear objectives; Alignment (\uparrow): supporting the scenarios of personalized deep research.

Finally, we systematically formulated 5 balanced tasks per domain, yielding 50 tasks in total: $\mathcal{T} = t_i \mid i = 1, \dots, 50$, with $t_i = (q_i, d(t_i))$ where q_i is the query and $d(t_i) \in \mathcal{D}$ the domain. A parallel English set \mathcal{T}_{EN} was also created, semantically aligned with the Chinese tasks.

162 3.2 AUTHENTIC USER PROFILE CONSTRUCTION.
163164 A key innovation of our benchmark lies in the careful design of highly realistic and richly detailed
165 authentic user profiles. We moved beyond synthetic or stereotyped characterizations by grounding
166 our profiles in real user data.167 **Structured Explicit Persona Collection.** We recruited 25 volunteers with diverse demographic
168 profiles across age, profession, income, and life stage. After receiving standardized training on
169 data authenticity and privacy, volunteers mapped their authentic personal details onto a specially
170 designed persona schema, \mathcal{S} , which can be found in the Appendix D. This process yielded a set of
171 25 structured explicit ground-truth personas \mathcal{P}_s , denoted as: $\mathcal{P}_s = \{Ps_j \mid j = 1, \dots, 25\}$.
172173 **Dynamic Personalized Context Integration.** To complement these explicit personas with dy-
174 namic context, we employed professional annotators simulate the daily interactions of these col-
175 lected personas through a phone APP. Over a period, they were instructed to: 1) Record naturalistic
176 memory snippets (m_j), such as travel aspirations, health goals, and family plans; and 2) Conduct
177 conversational interactions (c_j) with the intelligent assistant integrated in the app. This longitudinal
178 data captures each user’s evolving interests, habits, and implicit preferences. These multi-modal
179 data streams were then processed by the built-in management system of the APP, f_θ , to generate
180 dynamic personalized contexts: $\mathcal{P}_c = \{Pc_j \mid Pc_j = f_\theta(m_j, c_j), j = 1, \dots, 25\}$. Annotation
181 details are in Appendix H.182 For convenience, we define the complete user profile set \mathcal{P} as the collection of paired structured
183 explicit personas and dynamic personalized contexts:

184
$$\mathcal{P} = \{(Ps_j, Pc_j) \mid j = 1, \dots, 25\}$$

185

186 3.3 PERSONALIZED DEEP RESEARCH QUERY FORMULATION
187188 The final stage of benchmark construction involved the principled pairing of user profiles with deep
189 research tasks to generate meaningful, personalized queries. We recognized that a random pairing
190 would fail to capture the intrinsic relevance between a user and their research needs.
191192 To address this, we employed a user-driven, committee-guided alignment protocol. Each of the
193 25 volunteers first reviewed the full task pool \mathcal{T} and selected tasks that were personally relevant.
194 Then, the committee curated and refined these selections through rigorous discussions, ensuring:
195 (1) diversity of user profiles associated with each task, and (2) overall alignment between each
196 user–task pair. This process yielded a user subset $\mathcal{P}_i \subset \mathcal{P}$ for each task t_i , where $|\mathcal{P}_i| = 5$.
197198 Finally, a total of 250 personalized personalized deep research queries were formed:
199

200
$$\mathcal{Q} = \{(p, t_i) \mid i = 1, \dots, 50, p \in \mathcal{P}_i\}, \quad |\mathcal{Q}| = 250$$

201

202 where each query combines one high-quality deep research task t_i with a corresponding user profile
203 p from its assigned user set.
204205 This benchmark faithfully mirrors real-world personalized deep research scenarios while providing
206 a standardized, reproducible, and scalable evaluation setting. By jointly modeling task complex-
207 ity, authentic user profile diversity, and motivational alignment, it provides a rigorous testbed for
208 evaluating whether agents can effectively integrate user profiles into deep research and deliver truly
209 personalized high-quality outputs.
210211 4 EVALUATION METHODOLOGY
212213 How do end-users judge the value of a Deep Research report? "Is this report for me?" (Personal-
214 ization), "Is it well-crafted?" (Quality), and "Is the information true?" (Reliability). Existing eval-
215 uations, however, typically stress report quality or factual correctness, neglecting personalization.
216 To systematically address these core concerns, we propose the **PQR Evaluation Framework**, a
217 novel and comprehensive methodology assessing reports along three complementary axes: (P) Per-
218 sonalization Alignment, (Q) Content Quality, and (R) Factual Reliability. This joint considera-
219 tion provides a holistic, user-centered assessment of Personalized Deep Research, as shown in Figure 2.
220

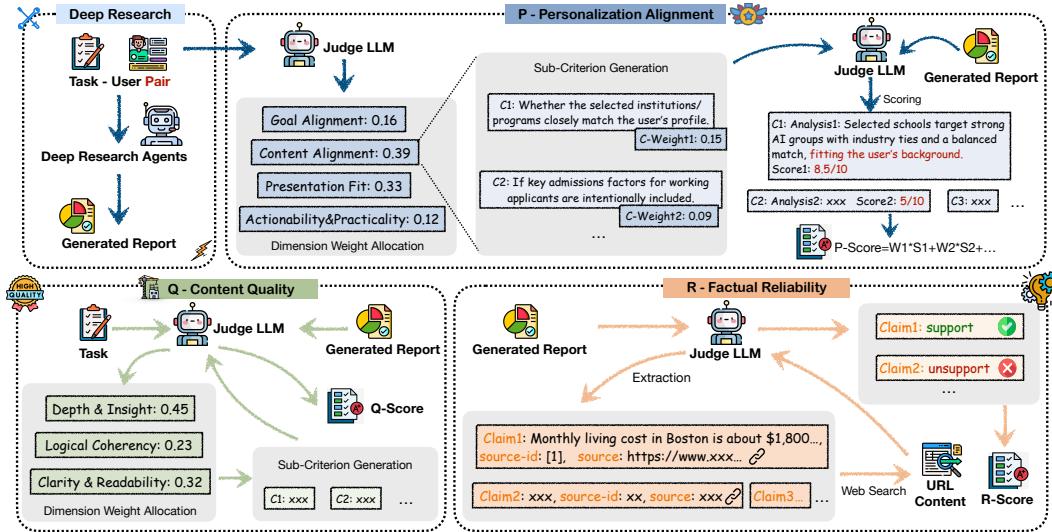


Figure 2: Overview of the PQR Evaluation Framework — A Multi-Dimensional Assessment System, Integrating Personalization Alignment (P), Content Quality (Q), and Factual Reliability (R).

4.1 P - PERSONALIZATION ALIGNMENT

Evaluating the personalization of generated report is a significant challenge due to its subjective, multi-dimensional nature. Recent studies (Wang et al., 2024a; Guan et al., 2025; Zhu et al., 2025c) consistently emphasize that personalization evaluation should move beyond global correctness toward preference-aware and user-centered assessment paradigms, highlighting the need for individualized criteria rather than generic evaluation frameworks. Motivated by these insights and our requirement analysis, we introduce Personalization Alignment (P-Score), a dynamic evaluation framework that generates customized criteria and scores for each user–task pair.

The framework is built on four fundamental dimensions: **Goal Alignment (GOAL)**, **Content Alignment (CONT)**, **Presentation Fit (PRES)** and **Actionability & Practicality (ACTI)**. Detailed definitions of dimensions are provided in the Appendix E. The P-Score is computed via a three-stage, LLM-driven pipeline that operationalizes these dimensions into a quantitative score:

Stage 1: Dynamic Dimension Weight Allocation. An LLM, acting as a meta-evaluator, analyzes the input *task* \mathcal{T} and *user persona* \mathcal{P}_s to determine the relative importance of the four dimensions. This stage outputs a weight vector $W = \{w_d\}_{d \in D_P}$ for the set of personalization dimensions $D_P = \{\text{Goal Alignment, Content Alignment, ...}\}$, where $\sum_{d \in D_P} w_d = 1.0$.

Stage 2: Granular Sub-Criterion Generation. For each dimension $d \in D_P$, the LLM generates a set of granular sub-criteria $C_d^P = \{c_1, \dots, c_n\}$, again conditioned on the *task* \mathcal{T} and *user persona* \mathcal{P}_s . Each sub-criterion c_i is assigned a weight w_{c_i} such that $\sum_{i=1}^n w_{c_i} = 1.0$.

Stage 3: LLM-Powered Scoring. A separate LLM then scores the target report against the criteria. Given the report, \mathcal{T} , and \mathcal{P}_s , it assigns a score $s_{c_i} \in [0, 10]$ and justification for each sub-criterion $c_i \in C_d^P$.

The final P-Score S_P is calculated by first computing each dimension score S_d as the weighted average of its sub-criteria. Then, S_P is obtained as the weighted average of the four dimension scores using dynamically generated weights. This can be formally expressed as:

$$S_P = \sum_{d \in D_P} w_d \cdot S_d = \sum_{d \in D_P} w_d \left(\sum_{c_i \in C_d^P} w_{c_i} \cdot s_{c_i} \right) \quad (1)$$

where w_d is the dimension weight, w_{c_i} is the sub-criterion weight, and s_{c_i} is the sub-criterion score.

270 4.2 Q – QUALITY OF CONTENT
271272 Beyond personalization, we also assess the intrinsic quality of the generated report—its depth, in-
273 sight, logic, clarity, and readability, regardless of user profiles. Quality is evaluated with respect to
274 the task \mathcal{T} and standards of rigorous research writing.275 We define three dimensions: **Depth & Insight (DEIN)**, **Logical Coherence (LOGC)**, and **Clarity & Readability (CLAR)** (see Appendix E).
276277 **Evaluation Process.** Following the dynamic criterion principle, an LLM meta-evaluator (i) as-
278 signs weights $\{w_d\}_{d \in D_Q}$ to the three dimensions, and (ii) generates a set of task-specific sub-criteria
279 C_d^Q for each dimension. A separate LLM scorer then rates the report against this criterion, produc-
280 ing a score $s_{c_i} \in [0, 10]$ with justification for each sub-criterion $c_i \in C_d^Q$. The final Q-Score is a
281 hierarchical weighted average:
282

283
$$S_Q = \sum_{d \in D_Q} w_d \cdot S_d = \sum_{d \in D_Q} w_d \left(\sum_{c_i \in C_d^Q} w_{c_i} \cdot s_{c_i} \right) \quad (2)$$

284

285 where w_d is the dimension weight, w_{c_i} is the sub-criterion weight, and s_{c_i} is the sub-criterion score.
286287 4.3 R – FACTUAL RELIABILITY.
288289 **Although factuality metrics such as FActScore (Min et al., 2023) exist, they are primarily designed**
290 **to verify atomic facts against static knowledge source and unsuitable for our deep research setting,**
291 **where factuality must be evaluated through retrieved citations to assess both the factual reliability of**
292 **the report and the agent’s capacity in utilizing web information.** We therefore assess report reliability
293 via an automated factual grounding framework, inspired by ResearcherBench (Xu et al., 2025) and
294 DeepResearch Bench (Du et al., 2025). The process has three stages:
295296 **Claim Extraction and Deduplication.** A Judge LLM is employed to extract all verifiable factual
297 claims with their sources, forming a set of triplets $\mathcal{TRI} = \{(c_i, \text{idx}_i, \text{source}_i)\}_{i=1}^N$, where uncited
298 claims have empty sources (see Appendix I.3). A second pass deduplicates claims:
299

300
$$\mathcal{TRI}_{\text{unique}} = \text{Deduplicate}(\mathcal{TRI}), \quad (3)$$

301 yielding N_{total} unique claims, of which N_{cited} are cited.
302

303 **Automated Verification.** For each unique triplet $(c_i, \text{idx}_i, \text{source}_i) \in \mathcal{TRI}_{\text{unique}}$, we use the Jina
304 Reader API to retrieve the source content $Content_i$. Then the Judge LLM checks support:
305

306
$$v_i = \begin{cases} 1, & \text{if } c_i \text{ is supported by } Content_i, \\ 0, & \text{if } c_i \text{ is unsupported or unknown.} \end{cases} \quad (4)$$

307

308 **Metric Calculation.** We compute two key metrics from the verification results:
309

- 310
- **Factual Accuracy (FA)** Measures the reliability of provided citations. It is the percentage
311 of claims that are factually verified and supported by their corresponding source material.
 - **Citation Coverage (CC)** Assesses the proportion of factual claims in a report that are
312 supported by explicit citations, reflecting how well the content is evidence-based.
-
- 313

314 Finally, we average FA and CC to derive a single Factual Reliability score, S_R :
315

316
$$FA = \frac{\sum_{i=1}^{N_{\text{cited}}} v_i}{N_{\text{cited}}} \times 10, \quad CC = \frac{N_{\text{cited}}}{N_{\text{total}}} \times 10, \quad S_R = \frac{FA + CC}{2}. \quad (5)$$

317

318 4.4 FINAL SCORE AGGREGATION
319320 **To obtain a holistic measure of the report, we define the final overall score as a arithmetic mean over**
321 **the three dimension scores:**
322

323
$$S_{\text{overall}} = \frac{S_P + S_Q + S_R}{3} \quad (6)$$

324 where S_P , S_Q , and S_R denote the scores for personalization, quality and factual reliability respec-
325 tively. This aggregation provides a straightforward and comprehensive measure of the personalized
326 deep research report.
327

324

5 EXPERIMENTS

325

5.1 EXPERIMENTAL SETTINGS

328 We benchmarked a diverse set of systems, including commercial deep research systems: Gemini-
 329 2.5-Pro Deep Research, O3 Deep Research, Perplexity Deep Research ([Google DeepMind, 2025](#);
 330 [OpenAI, 2025b](#); [Perplexity Team, 2025](#)), open-source deep research agents: Deerflow, Oagents,
 331 Miroflow ([ByteDance, 2025a](#); [Zhu et al., 2025a](#); [MiroMind AI Team, 2025](#)), and leading LLMs
 332 with search tools: Gemini-2.5-Pro-Search, Claude-3.7-Sonnet-Search, Perplexity-Sonar-Reasoning-
 333 Pro, GPT-4.1-Search-Preview ([DeepMind](#); [Anthropic](#); [AI, 2025](#); [OpenAI, a](#)). Due to computational
 334 constraints, the evaluation was performed on a subset of 150 representative queries. GPT-5 ([OpenAI, 2025a](#))
 335 was utilized as the judge model for Personalization (P) and Quality (Q) metrics, while the
 336 more efficient GPT-5-Mini ([OpenAI, b](#)) served as the judge for the Reliability (R) metric, ensuring
 337 a balance of advanced reasoning and efficiency (more details in Appendix B and J).

338

5.2 MAIN RESULTS

339 Table 1: Evaluation results of Personalized Deep Research Bench under the *Task w/Persona* con-
 340 figuration. The best results in each column are highlighted in **bold**, and the second-best results are
 341 underlined.

Model	Overall	Personalization				Quality			Reliability	
		GOAL	CONT	PRES	ACTI	DEIN	LOGC	CLAR	FA	CC
<i>Commercial Deep Research Agents</i>										
Gemini-2.5-Pro Deep Research	6.58	<u>5.27</u>	5.78	5.83	4.56	<u>5.32</u>	6.13	6.16	8.40	9.26
O3 Deep Research	<u>6.11</u>	5.67	5.95	<u>5.57</u>	5.10	5.68	6.40	<u>5.58</u>	6.84	7.14
Perplexity Deep Research	5.99	4.69	4.93	4.72	4.33	4.93	5.43	4.68	7.68	<u>9.02</u>
<i>Open-Source Deep Research Agents</i>										
OAgents	6.64	6.68	6.44	7.13	6.92	6.99	7.44	6.85	3.77	8.32
DeerFlow	5.30	5.20	4.97	6.71	5.41	5.43	6.25	6.44	6.85	<u>2.32</u>
MiroFlow	<u>5.78</u>	<u>6.65</u>	6.45	<u>7.03</u>	<u>6.65</u>	<u>6.53</u>	<u>7.31</u>	<u>6.68</u>	7.29	0.44
<i>LLM with Search Tools</i>										
Gemini-2.5-Pro w/Search	5.53	4.85	5.20	<u>5.61</u>	4.19	4.54	5.57	<u>5.41</u>	6.99	6.62
Claude-3.7-Sonnet w/Search	4.83	4.27	4.24	5.43	4.28	<u>4.26</u>	5.09	5.34	<u>8.27</u>	2.37
Perplexity-Sonar-Reasoning-Pro	<u>5.02</u>	4.27	4.37	5.27	4.15	4.22	5.03	5.23	8.44	<u>3.67</u>
GPT-4.1 w/Search	4.28	<u>4.59</u>	4.86	5.74	4.07	4.21	<u>5.27</u>	5.54	6.75	0.10

356 The evaluation results on the Personalized Deep Research Bench under the Task w/Persona config-
 357 uration (Task and Persona are explicitly provided to the agent) are shown in Table 1. Our analysis
 358 reveals several key findings regarding the performance of different model categories.

360 **Open-source Agents Excel in Personalization.** Open-source agents achieve the strongest per-
 361 sonalization, with OAgents achieving the top score (6.64) and leading most sub-metrics, including
 362 GOAL (6.68), PRES (7.13), and LOGC (7.44). MiroFlow also performs competitively, outperforming
 363 OAgents in CONT (6.45) and FA (7.29). However, reliability remains their weakness: OAgents
 364 suffers from low factual accuracy (3.77), while both MiroFlow and DeerFlow show poor citation
 365 coverage.

366 **Commercial Agents Provide Balanced Quality and Reliability.** Commercial systems achieve
 367 slightly lower personalization but higher reliability and consistent quality. Gemini-2.5-Pro Deep Re-
 368 search leads this group (6.58), achieving top FA (8.40), CC (9.26), and solid quality scores (DEIN:
 369 5.32, LOGC: 6.13, CLAR: 6.16). O3 Deep Research follows closely (6.11), leading in personaliza-
 370 tion within this category (GOAL: 5.67, CONT: 5.95) and maintaining competitive quality (DEIN:
 371 5.68, LOGC: 6.40, CLAR: 5.58). In summary, commercial agents are reliable and robust in quality,
 372 but they lag moderately behind open-source agents in personalization.

373 **LLMs with Search Tools Fall Short.** Search Tools equipped LLMs underperform specialized
 374 agents. Gemini-2.5-Pro w/Search is the strongest in this group (5.53), while others, such as
 375 Perplexity-Sonar-Reasoning-Pro, achieve high FA (8.44) but poor CC and weak personalization.
 376 GPT-4.1 w/Search, for example, nearly fails in CC (0.10). These results indicate that adding search
 377 alone is insufficient to reach the personalization and quality of dedicated deep research agents.

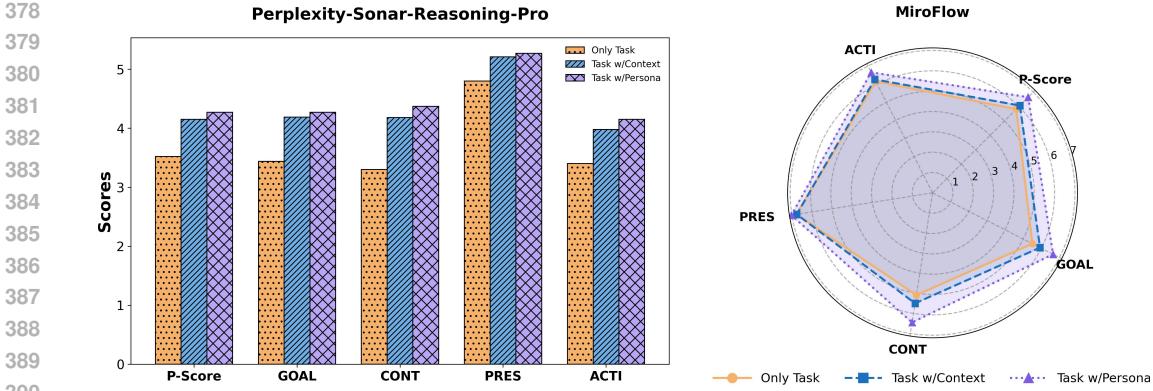


Figure 3: Analysis on Personalization Metrics for Sonar-Reasoning-Pro and MiroFlow

5.3 IMPACT OF INFORMATION AVAILABILITY ON PERSONALIZATION

While the previous section evaluated agents with explicit personas, a more realistic scenario involves inferring user needs from conversational or interaction context, as explicit personas are rarely available. To examine this, we conducted a comparative experiment under three conditions: Task Only (the agent receives only the task), Task w/Context (the task plus user’s conversational or interaction context), and Task w/Persona (the task plus an explicit user persona, consistent with our main experiment). We present the results in Table 2 and Figure 3. A more detailed results are in Appendix F.

Table 2: Evaluation results on the Personalization across different personalization settings. The table shows results for three configurations: *Task Only*, *Task w/Context*, and *Task w/Persona*. Best scores in each column are highlighted in **bold**, second-best in underlined.

Model	Setting	Personalization				
		P-Score	GOAL	CONT	PRES	ACTI
OAgents	<i>Task Only</i>	6.17	5.91	5.42	6.90	6.51
	<i>Task w/Context</i>	<u>6.53</u>	<u>6.32</u>	<u>5.99</u>	<u>7.04</u>	<u>6.81</u>
	<i>Task w/Persona</i>	6.78	6.68	6.44	7.13	6.92
O3 Deep Research	<i>Task Only</i>	5.13	5.14	5.08	<u>5.62</u>	5.03
	<i>Task w/Context</i>	5.48	5.58	5.67	5.70	5.29
	<i>Task w/Persona</i>	<u>5.46</u>	5.67	5.95	5.57	<u>5.10</u>
Gemini-2.5-Pro w/Search	<i>Task Only</i>	3.96	3.91	3.86	5.53	3.70
	<i>Task w/Context</i>	<u>4.55</u>	<u>4.66</u>	<u>4.95</u>	<u>5.59</u>	4.09
	<i>Task w/Persona</i>	4.70	4.85	5.20	5.61	4.19

More Information Consistently Yields Better Personalization. Across all systems, personalization scores (P-Score) increase with the user information (persona or context) provided. This trend holds for all sub-metrics, confirming the intuitive hypothesis that access to user-specific data is crucial for tailoring research outputs.

Explicit Personas Outperform Context. Context improves performance over the baseline, but the largest gains come from explicit personas. For instance, OAgents’ GOAL score increases from 6.32 (Context) to 6.68 (Persona), a larger jump than the improvement from Task Only to Context. This indicates that while agents can partially leverage implicit context, they struggle to fully extract user preferences from unstructured, implicit data. Explicit personas, by contrast, provide a stronger and more accessible personalization signal.

5.4 BOOSTING PERSONALIZATION VIA CONTEXT-AWARE MEMORY SYSTEMS

The preceding analysis reveals that while contextual information is beneficial, agents struggle to distill it into an actionable user understanding as effectively as when provided with an explicit persona.

To address this, we designed a second experiment to test whether advanced memory systems can transform unstructured *context* into explicit *persona* to perform personalized deep research. we evaluated 50 suitable queries on three systems: Mem0 (Chhikara et al., 2025), Memory OS (Kang

432 et al., 2025), and O-Mem (a private agent memory system), on their ability to extract, integrate, and
 433 infer user preferences from *context* to drive downstream deep research systems. Results indicate
 434 potential for improving higher-level reasoning and user information integration.
 435

436 Table 3: Evaluation results on for different memory systems under the *Task w/Context* setting, using
 437 Perplexity Deep Research. Currently, most memory systems can only align content with user char-
 438 acteristics, so we prioritize GOAL and CONT scores. We also display other metrics for clarity. The
 439 best metric is highlighted in **bold**. Underlined denotes the second highest.
 440

Method	Personalization				
	P-Score	GOAL	CONT	PRES	ACTI
No Memory	3.69	3.88	3.74	3.90	3.46
Mem0	3.55	3.73	3.55	3.77	3.36
Memory OS	<u>3.88</u>	<u>4.06</u>	<u>3.97</u>	<u>4.09</u>	<u>3.66</u>
O-Mem	4.26	4.47	4.43	4.34	4.00
Task w/Persona	4.58	4.69	4.93	4.72	4.33

441 As shown in Table 3, memory systems yield varied but promising results. O-Mem outperforms
 442 both *No Memory* baseline and other systems, while Mem0 underperforms. However, a significant
 443 gap remains between the best system and ideal *Task w/Persona* performance, indicating that current
 444 memory systems struggle to fully synthesize information from context. This gap highlights the need
 445 for future research on memory systems that combine factual retrieval with higher-level reasoning
 446 and abstraction, moving beyond storage toward constructing dynamic, persona-like models of users.
 447

448 5.5 ALIGNMENT WITH HUMAN CONSISTENCY

450 To validate the evaluation framework, we conducted a systematic study comparing the judgments of
 451 LLMs against human experts. We sampled 15 representative queries and generated responses from
 452 two deep research agents: MiroFlow and O3 Deep Research. A panel of human evaluators scored
 453 these reports using the same criteria, establishing a ground truth for our comparison.
 454

455 We designed two complementary metrics to quantify this alignment: Pairwise Comparison Agree-
 456 ment (PCA) measures the percentage of the LLM judge and human experts agree on which of the
 457 two reports is better for a specific criterion. Mean Absolute Rating Deviation (MARD) measures the
 458 average absolute difference between the scores assigned by the LLM and human judges. Detailed
 459 mathematical formulations for these metrics can be found in Appendix C.
 460

461 Based on the results shown in Table 4, GPT-5 achieved the highest PCA and lowest MARD, indi-
 462 cating the strongest agreement with human judgments, while maintaining a reasonable cost ([\\$0.68](#)
 463 [per query](#)). We finally select GPT-5 as our primary judge model.
 464

465 Table 4: Alignment results of judge LLMs with human ratings. PCA is reported as proportion of
 466 agreement (higher is better), MARD as mean absolute deviation (lower is better), Avg. Cost is
 467 measured in US dollars (\$). The best metric is highlighted in **bold**.
 468

Judge LLM	PCA \uparrow	MARD \downarrow	Avg. Cost (\$ \downarrow)
GPT-5	0.43	1.40	0.68
Claude-3.7-Sonnet	0.39	1.44	0.97
Gemini-2.5-Pro	0.40	2.33	0.61

477 6 CONCLUSION

478 In conclusion, our work addresses the critical gap in DRAs evaluation by introducing the Personal-
 479 ized Deep Research Bench, the first benchmark of its kind featuring 250 realistic queries that pair 50
 480 diverse deep research tasks across 10 domains with 25 authentic user profiles. Along with the PQR
 481 Evaluation Framework, which jointly measures personalization, content quality, and factual reliabil-
 482 ity, our study reveals both the potential and current limitations towards personalized deep research.
 483 By establishing this rigorous foundation, our work paves the way for developing and benchmarking
 484 the next generation of truly personalized and effective AI research assistants.
 485

486

7 ETHICS STATEMENT

488 This work strictly complies with the ethical guidelines. User profiles were collected from 25 volunteers
 489 under informed consent, with training on authenticity and privacy; all data was anonymized.
 490 Annotators involved in data labeling were recruited under informed consent, compensated fairly, and
 491 instructed to ensure accuracy and neutrality in their annotations. Human evaluators participated vol-
 492 untarily with full awareness of the research purpose. To ensure fairness and mitigate potential biases,
 493 tasks and profiles were designed with diversity across age, profession, income, and life stage. The
 494 study involves no sensitive or harmful content, and all experiments were conducted in a controlled,
 495 ethical manner.

496

497 8 REPRODUCIBILITY STATEMENT

498 To ensure reproducibility, we provide comprehensive details across the paper and appendix. Sec-
 499 tion 3 and Appendix H describes benchmark construction, including task design, user profiles col-
 500 lection, and query pairing. The evaluation framework is detailed in Section 4, including the dynamic
 501 weight allocation, granular criterion generation, and scoring methodology for each dimension. Sec-
 502 tion 5 and Appendix B outline experimental setups, systems, and configurations. Prompt templates
 503 (Appendix I), persona schema (Appendix D), and evaluation dimensions (Appendix E) are fully
 504 documented. We have submitted our evaluation data in the Supplementary Material. Due to Open-
 505 Review’s file size limit, we only upload a subset of them. We will fully release the benchmark
 506 immediately after the double blind review process. We invite the community to build upon this
 507 work in advancing personalized deep research agents.

508

509 REFERENCES

- 510 Perplexity AI. Sonar reasoning pro model (via perplexity api). Perplexity AI, 2025. URL <https://www.perplexity.ai/hub/blog/introducing-the-sonar-pro-api>. Accessed [Current Date].
- 511 Anthropic. Claude 3.7 sonnet and claude code. URL <https://www.anthropic.com/news/claude-3-7-sonnet>. Accessed [Current Date].
- 512 ByteDance. Deerflow: A community-driven deep research framework. <https://github.com/bytedance/deer-flow>, 2025a. Accessed: 2025-09-10.
- 513 ByteDance. Doubao deep research. <https://www.doubao.com/chat/>, 2025b. Accessed: 2025-09-10.
- 514 Kaiyuan Chen, Yixin Ren, Yang Liu, Xiaobo Hu, Haotong Tian, Tianbao Xie, Fangfu Liu, Haoye
 515 Zhang, Hongzhang Liu, Yuan Gong, Chen Sun, Han Hou, Hui Yang, James Pan, Jianan Lou,
 516 Jiayi Mao, Jizheng Liu, Jinpeng Li, Kangyi Liu, Kenkun Liu, Rui Wang, Run Li, Tong Niu, Wen-
 517 long Zhang, Wenqi Yan, Xuanzheng Wang, Yuchen Zhang, Yi-Hsin Hung, Yuan Jiang, Zexuan
 518 Liu, Zihan Yin, Zijian Ma, and Zhiwen Mo. xbench: Tracking agents productivity scaling with
 519 profession-aligned real-world evaluations, 2025a. URL <https://arxiv.org/abs/2506.13651>.
- 520 Zijian Chen, Xueguang Ma, Shengyao Zhuang, Ping Nie, Kai Zou, Andrew Liu, Joshua Green,
 521 Kshama Patel, Ruoxi Meng, Mingyi Su, Sahel Sharifmoghaddam, Yanxi Li, Haoran Hong,
 522 Xinyu Shi, Xuye Liu, Nandan Thakur, Crystina Zhang, Luyu Gao, Wenhui Chen, and Jimmy Lin.
 523 Browsecmp-plus: A more fair and transparent evaluation benchmark of deep-research agent,
 524 2025b. URL <https://arxiv.org/abs/2508.06600>.
- 525 Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building
 526 production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*,
 527 2025.
- 528 João Coelho, Jingjie Ning, Jingyuan He, Kangrui Mao, Abhijay Paladugu, Pranav Setlur, Jiahe
 529 Jin, Jamie Callan, João Magalhães, Bruno Martins, and Chenyan Xiong. Deepresearchgym: A
 530 free, transparent, and reproducible evaluation sandbox for deep research, 2025. URL <https://arxiv.org/abs/2505.19253>.

- 540 Google DeepMind. Gemini 2.5 technical report and model card (referencing gemini 2.5 pro). URL
 541 <https://deepmind.google/models/gemini/pro/>. Accessed [Current Date].
 542
- 543 Mingxuan Du, Benfeng Xu, Chiwei Zhu, Xiaorui Wang, and Zhendong Mao. Deepresearch bench:
 544 A comprehensive benchmark for deep research agents, 2025. URL <https://arxiv.org/abs/2506.11763>.
 545
- 546 Gerhard Fischer. User modeling in human–computer interaction. *User modeling and user-adapted*
 547 *interaction*, 11(1):65–86, 2001.
 548
- 549 Google DeepMind. Introducing gemini deep research. [https://gemini.google/](https://gemini.google/overview/deep-research/)
 550 overview/deep-research/, 2025. Accessed: 2025-09-10.
 551
- 552 Boyu Gou, Zanming Huang, Yuting Ning, Yu Gu, Michael Lin, Weijian Qi, Andrei Kopanov,
 553 Berta Yu, Bernal Jiménez Gutiérrez, Yiheng Shu, Chan Hee Song, Jiaman Wu, Shijie Chen,
 554 Hanane Nour Moussa, Tianshu Zhang, Jian Xie, Yifei Li, Tianci Xue, Zeyi Liao, Kai Zhang,
 555 Boyuan Zheng, Zhaowei Cai, Viktor Rozgic, Morteza Ziyadi, Huan Sun, and Yu Su. Mind2web
 556 2: Evaluating agentic search with agent-as-a-judge, 2025. URL <https://arxiv.org/abs/2506.21506>.
 557
- 558 Jian Guan, Junfei Wu, Jia-Nan Li, Chuanqi Cheng, and Wei Wu. A survey on personalized alignment
 559 – the missing piece for large language models in real-world applications, 2025. URL <https://arxiv.org/abs/2503.17003>.
 560
- 561 Mengkang Hu, Yuhang Zhou, Wendong Fan, Yuzhou Nie, Bowei Xia, Tao Sun, Ziyu Ye, Zhaoxuan
 562 Jin, Yingru Li, Qiguang Chen, Zeyu Zhang, Yifeng Wang, Qianshuo Ye, Bernard Ghanem, Ping
 563 Luo, and Guohao Li. Owl: Optimized workforce learning for general multi-agent assistance in
 564 real-world task automation, 2025. URL <https://arxiv.org/abs/2505.23885>.
 565
- 566 Yuxuan Huang, Yihang Chen, Haozheng Zhang, Kang Li, Huichi Zhou, Meng Fang, Linyi Yang,
 567 Xiaoguang Li, Lifeng Shang, Songcen Xu, Jianye Hao, Kun Shao, and Jun Wang. Deep research
 568 agents: A systematic examination and roadmap, 2025. URL <https://arxiv.org/abs/2506.18096>.
 569
- 570 Bowen Jiang, Zhuoqun Hao, Young-Min Cho, Bryan Li, Yuan Yuan, Sihao Chen, Lyle Ungar,
 571 Camillo J. Taylor, and Dan Roth. Know me, respond to me: Benchmarking llms for dynamic
 572 user profiling and personalized responses at scale, 2025. URL <https://arxiv.org/abs/2504.14225>.
 573
- 574 Jiazheng Kang, Mingming Ji, Zhe Zhao, and Ting Bai. Memory os of ai agent, 2025. URL <https://arxiv.org/abs/2506.06326>.
 575
- 576 Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A Hale. The benefits, risks and bounds of
 577 personalizing the alignment of large language models to individuals. *Nature Machine Intelligence*,
 578 6(4):383–392, 2024.
 579
- 580 Weizhen Li, Jianbo Lin, Zhusong Jiang, Jingyi Cao, Xinpeng Liu, Jiayu Zhang, Zhenqiang Huang,
 581 Qianben Chen, Weichen Sun, Qixiang Wang, et al. Chain-of-agents: End-to-end agent founda-
 582 tion models via multi-agent distillation and agentic rl. *arXiv preprint arXiv:2508.13167*, 2025a.
 583
- 584 Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and
 585 Zhicheng Dou. Search-ol: Agentic search-enhanced large reasoning models, 2025b. URL
 586 <https://arxiv.org/abs/2501.05366>.
 587
- 588 Xiaoxi Li, Jiajie Jin, Guanting Dong, Hongjin Qian, Yutao Zhu, Yongkang Wu, Ji-Rong Wen, and
 589 Zhicheng Dou. Webthinker: Empowering large reasoning models with deep research capability,
 590 2025c. URL <https://arxiv.org/abs/2504.21776>.
 591
- 592 Manus AI. Leave it to manus. <https://manus.im/>, 2025. Accessed: 2025-09-10.
 593
- 594 Grégoire Mialon, Clémentine Fourrier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas
 595 Scialom. Gaia: a benchmark for general ai assistants, 2023. URL <https://arxiv.org/abs/2311.12983>.
 596

- 594 Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer,
 595 Luke Zettlemoyer, and Hannaneh Hajishirzi. Factscore: Fine-grained atomic evaluation of fac-
 596 tual precision in long form text generation, 2023. URL <https://arxiv.org/abs/2305.14251>.
- 598 MiroMind AI Team. Miroflow: An open-source agentic framework for deep research. <https://github.com/MiroMindAI/MiroFlow>, 2025.
- 601 Moonshot AI. Kimi-researcher: End-to-end rl training for emerging agentic capabilities. <https://moonshotai.github.io/Kimi-Researcher/>, 2025. Accessed: 2025-09-10.
- 604 OpenAI. Introducing GPT-4.1 in the api, a. URL <https://openai.com/index/gpt-4-1/>.
 605 Accessed [Current Date].
- 606 OpenAI. Using GPT-5 - OpenAI API [includes GPT-5 mini model details], b. URL <https://platform.openai.com/docs/guides/latest-model>. Accessed [Current Date].
- 609 OpenAI. Introducing gpt-5, 2025a. <https://openai.com/index/introducing-gpt-5/>.
- 612 OpenAI. Introducing deep research. <https://openai.com/index/introducing-deep-research/>, 2025b. Accessed: 2025-09-10.
- 614 Perplexity Team. Introducing perplexity deep research. <https://www.perplexity.ai/hub/blog/introducing-perplexity-deep-research>, 2025. Accessed: 2025-09-10.
- 617 Long Phan, Alice Gatti, Ziwen Han, et al. Humanity's last exam, 2025. URL <https://arxiv.org/abs/2501.14249>.
- 619 Omid Rafieian and Hema Yoganarasimhan. Ai and personalization. *Artificial intelligence in mar-
 620 keting*, pp. 77–102, 2023.
- 622 Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large lan-
 623 guage models meet personalization, 2024. URL <https://arxiv.org/abs/2304.11406>.
- 624 Vinay Samuel, Henry Peng Zou, Yue Zhou, Shreyas Chaudhari, Ashwin Kalyan, Tanmay Rajpuro-
 625 hit, Ameet Deshpande, Karthik Narasimhan, and Vishvak Murahari. Personagym: Evaluating
 626 persona agents and llms, 2025. URL <https://arxiv.org/abs/2407.18416>.
- 628 Dingfeng Shi, Jingyi Cao, Qianben Chen, Weichen Sun, Weizhen Li, Hongxuan Lu, Fangchen Dong,
 629 Tianrui Qin, King Zhu, Minghao Liu, et al. Taskcraft: Automated generation of agentic tasks.
 630 *arXiv preprint arXiv:2506.10055*, 2025.
- 631 Xiangru Tang, Tianrui Qin, Tianhao Peng, Ziyang Zhou, Daniel Shao, Tingting Du, Ximeng Wei,
 632 Peng Xia, Fang Wu, He Zhu, et al. Agent kb: Leveraging cross-domain experience for agentic
 633 problem solving. *arXiv preprint arXiv:2507.06229*, 2025.
- 634 Meiling Tao, Chenghao Zhu, Dongyi Ding, Tiannan Wang, Yuchen Eleanor Jiang, and Wangchun-
 635 shu Zhou. Personafeedback: A large-scale human-annotated benchmark for personalization,
 636 2025. URL <https://arxiv.org/abs/2506.12915>.
- 638 Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li, and Yuandong
 639 Tian. Learning personalized alignment for evaluating open-ended text generation, 2024a. URL
 640 <https://arxiv.org/abs/2310.03304>.
- 641 Ningning Wang, Xavier Hu, Pai Liu, He Zhu, Yue Hou, Heyuan Huang, Shengyu Zhang, Jian Yang,
 642 Jiaheng Liu, Ge Zhang, Changwang Zhang, Jun Wang, Yuchen Eleanor Jiang, and Wangchunshu
 643 Zhou. Efficient agents: Building effective agents while reducing cost, 2025. URL <https://arxiv.org/abs/2508.02694>.
- 646 Tiannan Wang, Meiling Tao, Ruoyu Fang, Huilin Wang, Shuai Wang, Yuchen Eleanor Jiang, and
 647 Wangchunshu Zhou. Ai persona: Towards life-long personalization of llms, 2024b. URL
<https://arxiv.org/abs/2412.13103>.

- 648 Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won
 649 Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. Browsecmp: A simple yet
 650 challenging benchmark for browsing agents, 2025. URL <https://arxiv.org/abs/2504.12516>.
- 652 xAI Team. Introducing grok deepsearch. <https://x.ai/news/grok-3>, 2025. Accessed:
 653 2025-09-10.
- 655 Tianze Xu, Pengrui Lu, Lyumanshan Ye, Xiangkun Hu, and Pengfei Liu. Researcherbench: Evalu-
 656 uating deep ai research systems on the frontiers of scientific inquiry, 2025. URL <https://arxiv.org/abs/2507.16280>.
- 658 Zheng Zhao, Clara Vania, Subhradeep Kayal, Naila Khan, Shay B. Cohen, and Emine Yilmaz.
 659 Personalens: A benchmark for personalization evaluation in conversational ai assistants, 2025.
 660 URL <https://arxiv.org/abs/2506.09902>.
- 662 Wangchunshu Zhou, Yuchen Eleanor Jiang, Long Li, Jialong Wu, Tiannan Wang, Shi Qiu, Jintian
 663 Zhang, Jing Chen, Ruiwu Wu, Shuai Wang, Shiding Zhu, Jiyu Chen, Wentao Zhang, Xiangru
 664 Tang, Ningyu Zhang, Huajun Chen, Peng Cui, and Mrinmaya Sachan. Agents: An open-source
 665 framework for autonomous language agents. 2023. URL <https://arxiv.org/abs/2309.07870>.
- 667 Wangchunshu Zhou, Yixin Ou, Shengwei Ding, Long Li, Jialong Wu, Tiannan Wang, Jiamin Chen,
 668 Shuai Wang, Xiaohua Xu, Ningyu Zhang, Huajun Chen, and Yuchen Eleanor Jiang. Sym-
 669 bolic learning enables self-evolving agents. 2024. URL <https://arxiv.org/abs/2406.18532>.
- 671 He Zhu, Tianrui Qin, King Zhu, Heyuan Huang, Yeyi Guan, Jinxiang Xia, Yi Yao, Hanhao Li,
 672 Ningning Wang, Pai Liu, Tianhao Peng, Xin Gui, Xiaowan Li, Yuhui Liu, Yuchen Eleanor Jiang,
 673 Jun Wang, Changwang Zhang, Xiangru Tang, Ge Zhang, Jian Yang, Minghao Liu, Xitong Gao,
 674 Jiaheng Liu, and Wangchunshu Zhou. Oagents: An empirical study of building effective agents,
 675 2025a. URL <https://arxiv.org/abs/2506.15741>.
- 677 King Zhu, Hanhao Li, Siwei Wu, Tianshun Xing, Dehua Ma, Xiangru Tang, Minghao Liu, Jian
 678 Yang, Jiaheng Liu, Yuchen Eleanor Jiang, et al. Scaling test-time compute for llm agents, 2025.
 679 URL <https://arxiv.org/abs/2506.12928>, 2025b.
- 680 Minjun Zhu, Yixuan Weng, Linyi Yang, and Yue Zhang. Personality alignment of large language
 681 models, 2025c. URL <https://arxiv.org/abs/2408.11779>.
- 683 Thomas P. Zollo, Andrew Wei Tung Siah, Naimeng Ye, Ang Li, and Hongseok Namkoong. Per-
 684 sonallm: Tailoring llms to individual preferences, 2025. URL <https://arxiv.org/abs/2409.20296>.
- 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701

702 **A LIMITATIONS**
 703

704 This work still has some limitations that must be admitted: *a*) The collection of user personas and
 705 context annotations was conducted in Chinese. Although a parallel English version was provided,
 706 the underlying content remains linguistically and culturally constrained. *b*) Due to computational
 707 constraints, our main experiments were conducted on a selected subset of queries rather than the full
 708 benchmark. For future work, we plan to expand persona construction and context annotation beyond
 709 the Chinese-centric setting, aiming for more diverse and cross-lingual coverage. We also intend to
 710 scale up our experimental scope to fully utilize all benchmark queries and explore richer evaluation
 711 protocols under varied computational settings.

712 **B EXPERIMENT DETAIL**
 713

714 **B.1 CONFIGURATION OF METHODS**
 715

716 To ensure a fair comparison among different agents, we standardized their execution budgets and
 717 search configurations. Specifically, open-source deep research agents: Deerflow, Oagents, Miroflow
 718 are all standardized on GPT-5-Mini as the base LLM. The maximum number of execution steps was
 719 set to 8 for OAgent. For Deerflow: `max_step_num` and `max_plan_iterations` are default
 720 set to 3 and 1. To Miroflow: `max_turns` and `max_tool_calls_per_turn` in `main_agent`
 721 and `sub_agents` are both default set to 20 and 10. All agents relied on SerperAPI for web
 722 search and Jina for web content retrieval. For agents equipped with built-in web search tools,
 723 we set the `reasoning_effort` parameter to medium. In addition, for sonar-reasoning-pro, the
 724 `search_context_size` parameter was also set to medium.

725 **B.2 DATA SELECTION**
 726

727 Given the high cost of running our full evaluation pipeline, we first reduced the original set of
 728 250 queries. These queries covered 50 distinct tasks, each originally paired with 5 personas. To
 729 make the evaluation more tractable, we limited each task to 3 representative personas, resulting in
 730 a reduced set of 150 queries. From this subset, we further selected 50 queries that reflect a broad
 731 range of personalization demands across different user goals and characteristics. These queries were
 732 chosen based on their potential to reveal how effectively a memory system can adapt its responses
 733 to individual users. Since current memory systems are primarily limited to aligning content with
 734 user profiles rather than deeper task-level adaptation, our evaluation emphasizes Goal Alignment
 735 and Content Alignment as the core metrics. Additional metrics are also reported for completeness
 736 and transparency.

737 **C FORMULATION OF HUMAN CONSISTENCY METRICS**
 738

739 This section provides the detailed mathematical definitions for the metrics used to evaluate the alignment
 740 between LLM judges and human experts.

741 **Pairwise Comparison Agreement (PCA).** For each query q and criterion c , let A and B denote
 742 the two reports being compared. We denote the model's scores as $m_{q,c}^A$ and $m_{q,c}^B$, and the human
 743 ground truth scores as $h_{q,c}^A$ and $h_{q,c}^B$. PCA is defined as the proportion of cases where the model's
 744 preference order matches the humans' preference order:

$$745 \text{PCA} = \frac{1}{N} \sum_{q,c} \mathbf{1} \left[\text{sgn}(m_{q,c}^A - m_{q,c}^B) = \text{sgn}(h_{q,c}^A - h_{q,c}^B) \right],$$

746 where N is the total number of query-criterion pairs, $\text{sgn}(x)$ is the sign function, and $\mathbf{1}[\cdot]$ is the
 747 indicator function, which equals 1 if the condition is true and 0 otherwise.

748 **Mean Absolute Rating Deviation (MARD).** For each query q , report $r \in \{A, B\}$, and criterion
 749 c , the MARD is the overall mean of the absolute deviations between the model scores ($m_{q,c}^r$) and

756 the human scores ($h_{q,c}^r$). It is calculated as:
 757

$$758 \text{MARD} = \frac{1}{\sum_q 2|C_q|} \sum_q \sum_{r \in \{A, B\}} \sum_{c \in C_q} |m_{q,c}^r - h_{q,c}^r|,$$

$$759$$

$$760$$

761 where $|C_q|$ is the number of criteria for query q , and the term $2|C_q|$ accounts for the two reports
 762 being evaluated for each criterion.
 763

764 D PERSONA SCHEMA

$$765$$

$$766$$

767 **Table 5: Predefined Schema For Persona Collection**

768 **Predefined Persona Schema.**

$$769$$

770 **Basic Attributes**

$$771$$

$$772$$

$$773$$

Identity Characteristics	Name, Age, Gender, Occupation
Family Status	Family Members and Relationships, Pets
Long-term Spatial Characteristics	Permanent Residence, Hometown

774 **Behavioral Characteristics**

$$775$$

$$776$$

$$777$$

$$778$$

$$779$$

Online Usage Habits	High-frequency Apps and Usage Duration, Online Social Behavior
Offline Long-term Behavior	Daily Routine, Consumption Cycle, Consumption Characteristics, Periodic Mobility

780 **Environment**

$$781$$

$$782$$

$$783$$

$$784$$

Time	Time Preference
Geographical Location	Frequent Places, Travel Radius

785 **Personality Traits**

$$786$$

$$787$$

$$788$$

$$789$$

$$790$$

$$791$$

Personality Traits	Personality, Decision-making Style, Shopping Preference
--------------------	---

792 **Preferences and Interests**

$$793$$

$$794$$

$$795$$

$$796$$

$$797$$

$$798$$

$$799$$

$$800$$

Lifestyle Preferences	Diet, Accommodation Preferences, Shopping, Services
Travel Preferences	Frequency, Destinations, Travel Style
Content Preferences	Article Collection, Short Videos, Screenshots, Books, Movies, Singers, Actors
Exercise Preferences	Exercise Habits, Exercise Goals, Exercise Types, Other Investments, Exercise Locations

801 **Health Status**

$$802$$

$$803$$

$$804$$

$$805$$

$$806$$

$$807$$

$$808$$

$$809$$

Physical Condition	Medical History, Mental Condition, Physical Fitness
Health Needs	Health Needs

801 **Financial Information**

$$802$$

$$803$$

$$804$$

$$805$$

$$806$$

$$807$$

$$808$$

$$809$$

Financial Status	Income Structure, Asset Status, Consumption Characteristics, Debt Situation
Investment Experience	Investment Background, Knowledge Level
Risk Management	Risk Appetite

810 E DEFINITIONS OF EVALUATION DIMENSIONS
811
812813 Table 6: Dimensions and Definitions For Personalization Evaluation.
814

815 Dimension	816 Description
817 Goal Alignment	818 How well the report addresses the user's explicit and implicit goals.
819 Content Alignment	820 The suitability of the report's topic, depth, and breadth for the user's knowledge and interests.
821 Presentation Fit	822 The alignment of the report's language, structure, and style with the user's comprehension and preferences.
823 Actionability&Practicality	824 The extent to which the report offers practical value for decision-making or action.

825 Table 7: Dimensions and Definitions For Quality Evaluation.
826

827 Dimension	828 Description
829 Depth & Insight	830 The analytical richness, originality of thought, and critical perspective exhibited by the report.
831 Logical Coherence	832 The logic and coherence of the report's reasoning, ensuring ideas are rigorous and easy to follow.
833 Clarity & Readability	834 The report's language, information presentation, and formatting.

835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

864 F A MORE DETAILED EXPERIMENTS RESULTS
865
866867 Table 8: Evaluation results on the Personalization across different personalization settings to vari-
868 ous agents. The table shows results for three configurations: *Task Only*, *Task w/Context*, and *Task*
869 *w/Persona*. Best scores in each column are highlighted in **bold**, second-best in underlined.

Model	Setting	P-Score	GOAL	CONT	PRES	ACTI
<i>Commercial DeepResearch Agents</i>						
Gemini-2.5-Pro	<i>Task Only</i>	4.57	4.12	3.90	6.49	4.77
	<i>w/Context</i>	<u>4.70</u>	<u>4.84</u>	<u>5.17</u>	5.66	4.18
	<i>w/Persona</i>	5.12	5.27	5.78	<u>5.83</u>	4.56
O3	<i>Task Only</i>	5.13	5.14	5.08	<u>5.62</u>	5.03
	<i>w/Context</i>	5.48	<u>5.58</u>	<u>5.67</u>	5.70	5.29
	<i>w/Persona</i>	<u>5.46</u>	5.67	5.95	5.57	<u>5.10</u>
Perplexity	<i>Task Only</i>	3.58	3.47	3.38	4.82	3.47
	<i>w/Context</i>	<u>4.19</u>	<u>4.23</u>	<u>4.29</u>	4.56	<u>4.06</u>
	<i>w/Persona</i>	4.58	4.69	4.93	<u>4.72</u>	4.33
<i>Open-Source DeepResearch Agents</i>						
OAgents	<i>Task Only</i>	6.17	5.91	5.42	6.90	6.51
	<i>w/Context</i>	<u>6.53</u>	<u>6.32</u>	<u>5.99</u>	<u>7.04</u>	<u>6.81</u>
	<i>w/Persona</i>	6.78	6.68	6.44	7.13	6.92
Deerflow	<i>Task Only</i>	<u>5.11</u>	<u>4.77</u>	4.41	<u>6.67</u>	<u>5.31</u>
	<i>w/Context</i>	5.02	<u>4.77</u>	<u>4.47</u>	6.60	5.09
	<i>w/Persona</i>	5.38	5.20	<u>4.97</u>	6.71	5.41
Miroflow	<i>Task Only</i>	5.82	5.51	5.09	<u>6.78</u>	6.15
	<i>w/Context</i>	<u>6.07</u>	<u>5.93</u>	<u>5.50</u>	6.76	<u>6.26</u>
	<i>w/Persona</i>	6.65	6.65	<u>6.45</u>	7.03	6.65
<i>LLM with Search Tools</i>						
Gemini-2.5-Pro	<i>Task Only</i>	3.96	3.91	3.86	5.53	3.70
	<i>w/Context</i>	<u>4.55</u>	<u>4.66</u>	<u>4.95</u>	<u>5.59</u>	<u>4.09</u>
	<i>w/Persona</i>	4.70	4.85	5.20	5.61	4.19
Claude-3.7-Sonnet	<i>Task Only</i>	<u>4.00</u>	3.83	3.63	<u>5.32</u>	<u>3.99</u>
	<i>w/Context</i>	3.85	<u>3.87</u>	<u>4.06</u>	4.80	3.54
	<i>w/Persona</i>	4.37	4.27	4.24	5.43	4.28
Sonar-Rea-Pro	<i>Task Only</i>	3.52	3.44	3.30	4.80	3.40
	<i>w/Context</i>	<u>4.15</u>	<u>4.19</u>	<u>4.18</u>	<u>5.21</u>	<u>3.98</u>
	<i>w/Persona</i>	4.27	4.27	<u>4.37</u>	<u>5.27</u>	4.15
GPT4.1	<i>Task Only</i>	3.79	3.71	3.63	5.44	3.55
	<i>w/Context</i>	<u>4.41</u>	<u>4.43</u>	<u>4.52</u>	<u>5.70</u>	4.08
	<i>w/Persona</i>	4.52	4.59	4.86	<u>5.74</u>	<u>4.07</u>

918 G CASE STUDY
919
920
921
922
923
924
925926 **Deep Research Task**
927

928 I am a beginner runner preparing for my first full marathon in the next three months, with the goal of completing
929 the race safely and without injury. I can train about five days a week for roughly one hour per session, mainly
930 running in nearby parks or residential roads.
931 I do not have professional running equipment—only a pair of basic running shoes. I currently have no running
932 experience and no specific pace targets, but I hope to gradually increase my mileage while considering my daily
933 life and work schedule.
934 The training plan should be easy to follow, sustainable, and adaptable. Please create a personalized training plan
935 based on these conditions

936 **Personal Description**

937 Male, 34-year-old, Supply chain
938 management professional



939 Minimal exercise engagement,
940 High work and personal life demands



941 Married status, 6-year-old daughter,
942 Alaskan Malamute ownership



943 Mild obesity condition, Neck
944 discomfort symptoms, Overall health
945 improvement aspirations



946 High-demand professional role,
947 Frequent business travel requirements

948 **Personal Description**

949 Female graduate student, **Clinical**
950 **psychology specialization**



951 Three weekly exercise sessions, Yoga
952 and light jogging preferences, **Home-**
953 **based and proximate outdoor exercise**
954 **venues**



955 **Morning chronotype orientation**, Early
956 morning study and writing optimization



957 Mild obesity condition, Neck
958 discomfort symptoms, Overall health
959 improvement aspirations



960 Scholarship dependency, Part-time
961 employment supplementary income

962 **Personalized Report**

1. Designate weekends as **family fitness integration periods**, incorporating Alaskan Malamute into running activities for joint family-dog exercise engagement.
2. Develop **gradual weight reduction protocols** with specialized running posture guidance to prevent neck discomfort exacerbation
3. Implement business travel-adapted training solutions utilizing hotel fitness facilities with progressive beginner-oriented training progression plans.

963 **Personalized Report**

1. Integrate running as a **mindfulness-based practice** utilizing psychological training for enhanced self-observation and emotional regulation.
2. Optimize morning chronotype by establishing **running as the primary morning routine** for maximum cognitive and physical benefits.
3. Facilitate peer interaction through **university or community running club participation** and social media training documentation.
4. Implement cost-effective equipment selection with complimentary training resources and mobile application recommendations to minimize financial barriers.

964 **Figure 4: Case Study of Personalized Deep Research**
965
966
967
968
969
970
971

972

H DATA ANNOTATION

973
974 Our context dataset was manually annotated by 6 trained annotators, resulting in 5,939 labeled
975 instances. The process required a total effort of 85 person-days and a budget of approximately
976 \$6,000 USD.
977978

Data Annotation Protocol

979

Task Objective

980 The core of this task is to simulate the behavior of a specific user persona by collecting
981 various types of information that this user would likely encounter, consume, or generate in
982 their daily life. The final goal is to create a high-quality *memory* database for each persona
983 that accurately reflects their unique characteristics.
984985

Core Quality Standard: Reversibility

986 This is the most critical quality standard for this task. Every piece of data you collect must
987 clearly point to a specific trait of the user persona.
988

- 989
- **Verification Method:** After collection, review all data and assess whether its content
990 is sufficient to reverse-infer the user's key persona traits, such as their profession,
991 interests, and personality.
 - **Acceptance Criteria:** All collected data must pass the *reversibility* test to be considered
992 high-quality. If the data is too generic and the persona cannot be inferred from it, the data will be deemed non-compliant.

993

Collection Based on Persona Preferences

- 994
- Carefully read every preference tag in the user persona description.
 - For each preference, collect **at least five** pieces of related content.

995

Ensure Diversity of Sources and Types

- 996
- **Source Diversity:** Do not limit collection to a single platform. Data can be gathered from various apps (e.g., Twitter, Instagram, Reddit), websites, forums, etc.
 - **Type Diversity:** Diversify the format of the data you collect. Examples include:
 - Screenshots of social media posts.
 - Screenshots of conversations with friends (that reflect opinions or preferences).
 - Links and titles of articles, news, or videos.
 - Screenshots of purchase histories or product reviews.

1000

Requirement for Conversational Content

1001 Between 20% and 50% of the total data collected for each user should be in a conversational
1002 format. This helps to more vividly showcase the user's personality and communication
1003 habits.
10041005

Add Reasonable Noise for Authenticity

1006 The content you collect does not need to be an exact match to the persona's description. You
1007 can add relevant and reasonable details or *noise* to make the data appear more authentic, as
1008 if generated by a real user.
1009

- 1010
- **Example:** If the persona *likes basketball*, you could collect a news article about a recent
1011 Lakers game or a screenshot of a conversation with a friend debating whether
1012 Jordan or LeBron is the better player.
 - **Note:** Any added *noise* must not contradict other defined attributes in the persona,
1013 such as spending habits, personality, or profession.

1014

Quantity vs. Quality

1015 There are no strict quantitative requirements for data collection. Please prioritize quality and
1016 collect as much rich data as possible. Quality always takes precedence over quantity.
10171018

Deliverables and Annotation Requirements

- 1019
- **Deliverable:** For each user persona, export a separate *aimemory* database file.
 - **Content Annotation:** During collection, each piece of content must be given a clear title and be correctly associated with its corresponding persona.

1026

I PROMPT TEMPLATES

1028

I.1 PROMPTS IN PERSONALIZATION EVALUATION

1029

Prompt for Personalization Dimension Weights Allocation

1030 You are an experienced evaluation expert for research articles. You excel at deeply un-
 1031 derstanding the goals, challenges, and key value points of a specific research task and the
 1032 task initiator's persona, and then setting dynamic, reasonable, and well-justified weights for
 1033 evaluation dimensions in subsequent personalized article assessments.

1034 </system_role>

1035 <user_prompt>

1036 Here is a deep research task, as follows:

1037 <task>

1038 “{task_prompt}”

1039 </task>

1040 The user persona is as follows:

1041 <persona>

1042 “{persona_prompt}”

1043 </persona>

1044 <instruction>

1045 Background: The research team will conduct in-depth and comprehensive research based
 1046 on the above <task> and <persona> and ultimately produce a high-quality, personalized
 1047 research article.

1048 Your task: As the evaluation expert, you need to set the weights of the personalized evalua-
 1049 tion criteria for this specific <task>. The evaluation will revolve around the following four
 1050 dimensions:

- 1051 1. **Goal Alignment:** Whether the research sufficiently and accurately understands the
 1052 relationship between the task and the user persona, extracts deep and implicit needs,
 1053 and generates a personalized report based on them.
- 1054 2. **Content Alignment:** Whether the research selects and customizes content accord-
 1055 ing to the user's interests, knowledge background, and preferences.
- 1056 3. **Actionability & Practicality:** Whether the report is feasible, practical, and helpful
 1057 for the user's decision-making.
- 1058 4. **Presentation Fit:** Whether the report's language style, information structure, and
 1059 presentation format match the user's cognitive habits and medium preferences.

1060 Evaluation formula:

1061 Total score = Goal Alignment * Goal Alignment weight + Content Alignment * Content
 1062 Alignment weight + Actionability & Practicality * Actionability & Practicality weight +
 1063 Presentation Fit * Presentation Fit weight. (Note: The sum of all weights must be exactly
 1064 1.0)

1065 Core requirements:

- 1066 1. **Deeply analyze the task and user persona:** Carefully study the specific content
 1067 of <task>, its explicit goals, potential challenges, and hidden objectives. Combine
 1068 this with <persona> to analyze the user's needs, background, and preferences, and
 1069 understand the core value of the task's outcome.
- 1070 2. **Dynamically assign weights:** Based on your analysis, assign weights to the four
 1071 dimensions (use decimals between 0 and 1, e.g., 0.30). The key is to recognize
 1072 that different tasks and personas emphasize different aspects, so weights must be
 1073 flexibly adjusted according to task characteristics and persona, not fixed.
- 1074 3. **Explain your reasoning:** Your analysis (<analysis>) must clearly and specifically
 1075 explain why each dimension is assigned a given weight, and directly link your

1080
 1081 reasoning to the requirements of <task> and the characteristics of <persona>. This
 1082 is critical for evaluating the quality of your work.
 1083 **4. Output in the standard format:** Strictly follow the example format be-
 1084 low: first output <analysis> with detailed reasoning, then immediately provide
 1085 <json_output> with the weight assignment results.
 1086 </instruction>
 1087
 1088 <examples_rationale>
 1089 Below are two examples that demonstrate how to adjust the evaluation dimension weights
 1090 and explain the reasoning based on changes in task nature and user persona. Focus on
 1091 learning the thinking process and analytical method in these examples, not simply copying
 1092 their content or weight values.
 1093 </examples_rationale>
 1094 ...
 1095 Now strictly follow the above instructions and methodology, and start your work for the
 1096 following task:
 1097 <task>
 1098 "{task_prompt}"
 1099 </task>
 1100 <persona>
 1101 "{persona_prompt}"
 1102 </persona>
 1103 Please output your <analysis> and <json_output>.
 1104

1105 Prompt for Goal Alignment Criteria Generation
 1106
 1107 You are an experienced research article evaluation expert. You excel at breaking down ab-
 1108 stract evaluation dimensions (such as "Goal Understanding and Personalization Insight")
 1109 into actionable, clear evaluation criteria tailored to the specific research task and user per-
 1110 sona, and assigning reasonable weights with explanations for each criterion.
 1111 </system_role>
 1112
 1113 <user_prompt>
 1114 **Background:** We are evaluating a research article written for the following research task
 1115 under the dimension of Goal Alignment.
 1116 **Goal Alignment:** Whether the research fully and accurately understands the relationship
 1117 between the task and the user persona, extracts deep and implicit needs, and generates a
 1118 personalized report based on that understanding, with a focus on performing user-centered,
 1119 deeply personalized matching between the user persona and task requirements.
 1120 <task>
 1121 "{task_prompt}"
 1122 </task>
 1123 The user persona is as follows:
 1124 <persona>
 1125 "{persona_prompt}"
 1126 </persona>
 1127 <instruction>
 1128 Your goal:
 1129 For the Goal Alignment dimension of this research article, formulate a set of detailed, spe-
 1130 cific, and highly targeted evaluation criteria that are tightly aligned with the above <task>
 1131 and <persona>. You need to:
 1132 1. Deeply analyze the user persona and task scenario: Thoroughly examine the back-
 1133 ground characteristics, knowledge structure, cognitive habits, and latent expec-

1134
1135
1136

tations of <persona>. Combine this with the specific application scenario of <task> to identify the user's core explicit needs and deeper implicit needs.

1137
1138
1139
1140
1141
1142
1143
1144

2. Formulate personalized evaluation criteria: Based on the above analysis, propose specific evaluation criteria that reflect a deep understanding of <persona> and a close fit to the <task> scenario. These criteria should assess whether the content is well adapted to the user persona in style, depth, perspective, and practicality.
3. Explain the personalization rationale: Provide a brief explanation (explanation) for each criterion, clarifying how it addresses the specific attributes of <persona> or special requirements of <task>, and why such targeting is critical to achieving a good match.
4. Assign rational weights: Assign a weight (weight) to each criterion, ensuring that the total sum is 1.0. The distribution of weights should directly reflect the relative importance of each criterion in measuring how well the content matches "this particular user" in "this particular task." The closer a criterion is tied to persona characteristics and task scenario, the higher its weight should be.

1145
1146
1147
1148
1149

Core requirements:

1150
1151
1152
1153
1154

1. Deep personalization orientation: The analysis, criteria, explanations, and weights must be deeply rooted in the uniqueness of <persona> (e.g., their professional background, cognitive level, decision-making preferences, emotional needs) and the specific context of <task>. Avoid generic or templated evaluation.
2. Focus on contextual responsiveness and resonance: The criteria should evaluate whether the content not only responds to the task at the informational level but also resonates with the context and expectations implied by the user persona in terms of expression style, reasoning logic, case selection, and level of detail.
3. Rationale must reflect targeting: The <analysis> section must clearly explain how key features were extracted from the given <persona> and <task> to form these personalized criteria. Each criterion's explanation must directly show how it serves this specific user and task.
4. Weights must reflect personalization priorities: The weight distribution must logically demonstrate which aspects of alignment are the most critical success factors for "this user" completing "this task."
5. Standard output format: Strictly follow the example format below. First output the <analysis> text, then immediately provide the <json_output>.

1155
1156
1157
1158

</instruction>

1159
1160
1161
1162

<example_rational>

The example below demonstrates **how to develop Goal Alignment evaluation criteria based on the task requirements**. Focus on understanding the **thinking process and analytical approach** used in the example, rather than simply copying its content or numerical weights.

1163
1164
1165
1166
1167
1168
1169

</example_rational>

...

1170
1171
1172
1173
1174

Please strictly follow the above instructions and methodology. Now, for the following specific task, start your work:

1175

<task>

1176
1177
1178
1179
1180
1181

"{task_prompt}"

</task>

1182
1183
1184
1185

<persona>

"{persona_prompt}"

</persona>

1186
1187
1188

Please output your <analysis> and <json_output>.

</user_prompt>

1188
1189

Prompt for Content Alignment Criteria Generation

1190
1191
1192
1193

You are an experienced research article evaluation expert. You are skilled at breaking down abstract evaluation dimensions (such as “Content Alignment”) into actionable, clear, and specific evaluation criteria tailored to the given research task and user persona, and assigning reasonable weights and explanations for each criterion.

1194
1195
1196
1197
1198
1199

</system_role>
<user_prompt>
Background: We are providing a personalized scoring rubric for a specific task and user persona from the dimension of **Content Alignment**.

Content Alignment: Whether the research content is customized based on the user’s interests, knowledge background, and other preferences.

1200
1201
1202
1203

<task>

“{task_prompt}”

</task>

The user persona is as follows:

1204
1205
1206
1207
1208
1209
1210

<persona>
“{persona_prompt}”
</persona>

<instruction>

Your Goal: For the **Content Alignment** dimension of this research article, create a set of detailed, concrete, and highly tailored evaluation criteria for the above <task> and <persona>. You need to:

1211
1212
1213
1214
1215
1216
1217
1218
1219
1220

1. **Analyze the Task and Persona:** Deeply analyze <task> and <persona> to infer the user’s potential interests, knowledge background, and the depth and breadth of content they may prefer.
2. **Formulate Criteria:** Based on your analysis, propose specific evaluation criteria that focus on whether the report’s content matches the user’s interest points and knowledge level.
3. **Provide Explanations:** For each criterion, provide a brief explanation (explanation) explaining why it is important for evaluating the content alignment for this <task>.
4. **Assign Weights:** Assign a reasonable weight to each criterion (weight), ensuring that the sum of all weights equals exactly 1.0. The weight allocation should logically reflect the personalization-first principle: criteria directly tied to unique personal traits, exclusive preferences, or specific contextual needs in the user persona should receive higher weights, as they are key to achieving true personalized content alignment.
5. **Avoid Overlap:** Make sure the evaluation criteria focus solely on the **Content Alignment** dimension, avoiding overlap with other dimensions such as Goal Alignment, Expression Style Alignment, and Practicality/Actionability.

Core Requirements:

1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

1. **Strongly Linked to the Persona:** The analysis, criteria, explanations, and weights must be directly connected to the user’s interests, knowledge background, or content preferences.
2. **Focus on Content Selection and Depth:** The criteria should assess whether the choice of content is precise and whether the depth is appropriate, rather than merely evaluating whether information is presented.
3. **Provide Sufficient Rationale:** The <analysis> section must clearly articulate the overall reasoning behind formulating these criteria and weights, linking them to <task> and <persona>. Each explanation must clarify why the individual criterion is relevant.

1242
 1243 4. **Reasonable Weighting:** The weight distribution should be logical, reflecting the
 1244 relative importance of each criterion in measuring content alignment, with particu-
 1245 lar emphasis on giving higher priority to personalized aspects.
 1246 5. **Standardized Output Format:** Strictly follow the format below — output the
 1247 <analysis> text first, immediately followed by <json_output>.
 1248 </instruction>
 1249 <example_rational>
 1250 The following example demonstrates **how to formulate content alignment evaluation cri-**
 1251 **teria based on the task requirements and user persona.** Pay close attention to the **think-**
 1252 **ing process and analytical approach** in this example, rather than simply copying the con-
 1253 tent or weight values.
 1254 </example_rational>
 1255 ...
 1256 Please strictly follow the above instructions and methodology. Now, for the following spe-
 1257 cific task, start your work:
 1258 <task>
 1259 "task_prompt"
 1260 </task>
 1261 <persona>
 1262 "persona_prompt"
 1263 </persona>
 1264 Please output your <analysis> and <json_output>.
 1265

Scoring Prompt for Personalization

1266 <system_role>You are a strict, meticulous, and objective expert in evaluating personalized
 1267 research articles. You excel at deeply evaluating research articles based on specific personal-
 1268 ization assessment criteria, providing precise scores and clear justifications.</system_role>
 1269
 1270 <user_prompt>
 1271 **Task Background**
 1272 You are given an in-depth research task. Your job is to evaluate a research article written for
 1273 this task in terms of its performance in "**Personalization Alignment**". We will evaluate it
 1274 across the following four dimensions:
 1275
 1276 1. Goal Alignment
 1277 2. Content Alignment
 1278 3. Presentation Fit
 1279 4. Actionability & Practicality
 1280
 1281 <task>
 1282 "task_prompt"
 1283 </task>
 1284
 1285 **User Persona**
 1286 <persona>
 1287 "persona_prompt"
 1288 </persona>
 1289
 1290 **Article to be Evaluated**
 1291 <target_article>
 1292 "article"
 1293 </target_article>
 1294
 1295 **Evaluation Criteria**

1296

1297 You must evaluate the specific performance of this article in terms of personalization alignment,
 1298 **following the criteria list below**, outputting your analysis and then assigning a score
 1299 from 0–10. Each criterion includes its explanation, which you should read carefully.

1300 <criteria_list>
 1301 {criteria_list}
 1302 </criteria_list>

1303 <Instruction>

1304 **Your Task**

1305 Strictly follow **each criterion** in <criteria_list> to evaluate how <target_article> meets that
 1306 criterion. You must:

1. **Analyze Each Criterion:** For each item in the list, think about how the article
 meets the requirements of that criterion.
2. **Analytical Evaluation:** Combine the article content, the task, and the user persona
 to analyze the article’s performance for that criterion, pointing out both strengths
 and weaknesses.
3. **Scoring:** Based on your analysis, give a score between 0 and 10 (integer) for the
 article’s performance on that criterion.

1315 **Scoring Rules**

1316 For each criterion, give a score between 0 and 10 (integer). The score should reflect the
 1317 quality of the article’s performance:

- 0–2 points: Very poor. Almost completely fails to meet the requirement.
- 2–4 points: Poor. Meets the requirement only partially, with significant shortcomings.
- 4–6 points: Average. Basically meets the requirement; neither particularly good
 nor bad.
- 6–8 points: Good. Mostly meets the requirement, with notable strengths.
- 8–10 points: Excellent/Outstanding. Fully or exceptionally meets the requirement.

1326 **Output Format Requirements**

1327 Strictly follow the <output_format> below to output the evaluation results for **each criterion**. **Do not include any irrelevant content, introductions, or conclusions**. Start from
 1328 the first dimension and output all dimensions and their criteria in sequence:

1329 </Instruction>

1331 <output_format>

```

1332 {
  1333   "goal_alignment": [
    1334     {
      1335       "criterion": "[The text of the first Goal Alignment
      1336           criterion]",
      1337       "analysis": "[Analysis]",
      1338       "target_score": "[integer score 0-10]"
    },
    1339     {
      1340       "criterion": "[The text of the second Goal Alignment
      1341           criterion]",
      1342       "analysis": "[Analysis]",
      1343       "target_score": "[integer score 0-10]"
    },
    1345     ...
  ],
  1346   "content_alignment": [
    1347     {
      1348       "criterion": "[The text of the first Content Alignment
      1349           criterion]",
    }
  ]
}
```

```

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
        "analysis": "[Analysis]",
        "target_score": "[integer score 0-10]"
    },
    ...
],
"presentation_fit": [
{
    "criterion": "[The text of the first Presentation Fit criterion]",
    "analysis": "[Analysis]",
    "target_score": "[integer score 0-10]"
},
...
],
"actionability_practicality": [
{
    "criterion": "[The text of the first Actionability & Practicality criterion]",
    "analysis": "[Analysis]",
    "target_score": "[integer score 0-10]"
},
...
]
}

```

1.2 PROMPTS IN QUALITY EVALUATION

Prompt for Quality Dimension Weights Allocation

You are an experienced expert in evaluating research reports. You excel at deeply understanding the goals, challenges, and core value points of a given research task, and setting dynamic, reasonable, and well-justified dimension weights for subsequent report quality evaluations.

</system_role>

<user_prompt>

Here is a deep research task as follows:

<task>

“{task_prompt}”

</task>

<instruction>

Background: The research team will conduct an in-depth and comprehensive investigation based on the above <task> and eventually produce a high-quality research report.

Your Task: As an evaluation expert, you need to set the evaluation dimension weights specifically for this <task>. The evaluation will be carried out around the following three dimensions:

1. Depth & Insight: Whether the report provides sufficient depth and unique insights.

2. Logical Coherence: Whether the report’s reasoning framework is rigorous and its logical derivation coherent.

3. Clarity & Readability: Whether the report’s language, information presentation, and formatting are clear and easy to understand, allowing readers to absorb the content smoothly.

1404
1405
1406
1407
1408

Evaluation Formula: Total Score = (Depth & Insight * weight₁) + (Logical Coherence * weight₂) + (Clarity & Readability * weight₃). (Note: The sum of all weights must equal exactly 1.0)

1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421

Core Requirements:

1. Analyze the Task in Depth: Carefully study the <task> content, implicit objectives, potential challenges, and the core value of the deliverable.
2. Dynamically Allocate Weights: Based on your analysis, assign weights to the three dimensions (use decimals between 0 and 1, e.g., 0.4). The key is to understand that different tasks emphasize different aspects — weights must be adjusted flexibly based on task characteristics, rather than being fixed.
3. Explain the Allocation Rationale: Your analysis (<analysis>) must clearly and specifically explain why each dimension is given its corresponding weight and directly link your reasoning to the requirements and characteristics of <task>. This is the key criterion for evaluating your work quality.
4. Standardized Output Format: Strictly follow the example format below — first output the detailed rationale in <analysis>, then provide the weight allocation result in <json_output>.

</instruction>

1422
1423
1424
1425
1426
1427
1428
1429
1430
1431

<examples_rationale>

Below are two examples, which demonstrate how to adjust dimension weights according to the nature of the task and explain the reasoning. Please focus on learning the thinking process and analytical approach shown in the examples, rather than simply copying their content or numerical values.

</examples_rationale>

...

Please strictly follow the above instructions and methodology. Now, for the following specific task, begin your work:

<task>
“{task_prompt}”
</task>

Please output your <analysis> and <json_output>.

</user_prompt>

1438
1439

Prompt for Depth & Insight Criteria Generation

1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

You are an experienced expert in evaluating research reports. You excel at breaking down abstract evaluation dimensions (such as “Depth & Insight”) into actionable, task-specific, and clear criteria, assigning reasonable weights and explanations for each.

</system_role>

<user_prompt>

Background: We are evaluating a research report based on three dimensions: Depth & Insight, Logical Coherence, and Clarity & Readability.

1. **Depth & Insight:** Whether the report provides sufficient depth and unique insights.
2. **Logical Coherence:** Whether the report’s reasoning framework is rigorous and its logical derivation coherent.
3. **Clarity & Readability:** Whether the report’s language, information presentation, and formatting are clear and easy to understand.

<task>
“{task_prompt}”
</task>

1458

1459 <instruction>

1460 **Your Goal:** For the **Depth & Insight** dimension of this report, develop a detailed, specific,
1461 and highly task-targeted set of evaluation criteria. You need to:

1462

1. **Analyze the Task:** Examine `<task>` in depth and identify where deep analysis, logical reasoning, insight extraction, or value judgment are required to demonstrate insight.
2. **Formulate Criteria:** Based on the analysis, propose concrete evaluation criteria focusing on analytical depth, logical rigor, originality, and value of conclusions.
3. **Explain Each Criterion:** Provide a brief explanation (explanation) for why this criterion is important for evaluating Depth & Insight for `<task>`.
4. **Assign Weights:** Assign a reasonable weight (weight) to each criterion, ensuring the weights sum exactly to **1.0**. The weights should reflect the relative importance of each criterion within the Depth & Insight dimension.
5. **Avoid Overlap:** Clearly focus only on criteria relevant to **Depth & Insight**, avoiding aspects of **Logical Coherence** (structure) or **Clarity & Readability** (language, formatting).

1475

Core Requirements:

1476

1. **Stay Task-Specific:** The analysis, criteria, explanations, and weights must directly relate to the task's core requirements and characteristics.
2. **Go Beyond the Surface:** The criteria should assess analytical depth, reasoning rigor, originality of insights, and value of conclusions — not just listing information.
3. **Provide Strong Rationale:** The `<analysis>` section must clearly explain the overall approach to designing the criteria and weights, linking it to `<task>`. Each explanation must justify the criterion.
4. **Ensure Reasonable Weighting:** Weight distribution must be logical, reflecting the relative importance of each criterion in showing insight.
5. **Standardized Output Format:** Strictly follow the format below: output `<analysis>` first, then `<json_output>`.

1489

</instruction>

1490

<example_rational>

1491

Below is an example demonstrating **how to design Depth & Insight criteria**. Focus on the **thinking logic and analytical approach** rather than copying its contents or weight numbers.

</example_rational>

1492

...

1493

Please strictly follow the above instructions and methodology. Now, for the following specific task, begin your work:

1494

<task>

1495

" {task_prompt} "

1496

</task>

1497

Please output your `<analysis>` and `<json_output>`.

1498

1499

Scoring Prompt for Quality

1500

<system_role>

1501

You are a strict, meticulous, and objective expert in evaluating the quality of research articles. You excel at deeply evaluating research articles based on specific quality assessment criteria, providing precise scores and clear justifications.

1502

</system_role>

1503

<user_prompt>

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

Task Background

You are given an in-depth research task. Your job is to evaluate a research article written for this task. We will evaluate it across the following three dimensions: Depth & Insight, Logical Coherence and Clarity & Readability. The task is as follows:

```
<task>
  "{task_prompt}"
</task>
```

Article to be Evaluated

```
<target_article>
  "{article}"
</target_article>
```

Evaluation Criteria

You must evaluate the article's performance for each criterion in the list below, outputting your analysis and then assigning a score from 0–10. Each criterion includes its explanation, which you should read carefully.

```
<criteria_list>
  {criteria_list}
</criteria_list>
```

<Instruction>

Your Task

Strictly follow each criterion in `<criteria_list>` to evaluate how `<target_article>` meets that criterion. You must:

1. **Analyze Each Criterion:** For each item in the list, think about how the article meets the requirements of that criterion.
2. **Analytical Evaluation:** Combine the article content with the explanation of the criterion to analyze the article's performance for that criterion, pointing out both strengths and weaknesses.
3. **Scoring:** Based on your analysis, give a score between 0 and 10 (integer) for the article's performance on that criterion.

Scoring Rules

For each criterion, give a score between 0 and 10 (integer). The score should reflect the quality of the article's performance:

- 0–2 points: Very poor. Almost completely fails to meet the requirement.
- 2–4 points: Poor. Meets the requirement only partially, with significant shortcomings.
- 4–6 points: Average. Basically meets the requirement; neither particularly good nor bad.
- 6–8 points: Good. Mostly meets the requirement, with notable strengths.
- 8–10 points: Excellent/Outstanding. Fully or exceptionally meets the requirement.

Output Format Requirements

Strictly follow the `<output_format>` below to output the evaluation results for **each criterion**. **Do not include any irrelevant content, introductions, or conclusions**. Start from "criterion 1" and output all criteria in order:

```
</Instruction>
```

```
<output_format>
```

```
{
```

```
  "depth_insight": [
    {
      "criterion": "[The text of the first Depth & Insight criterion]",
```

```

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
        "analysis": "[Analysis]",
        "target_score": "[integer score 0-10]"
    },
    {
        "criterion": "[The text of the second Depth & Insight
                      criterion]",
        "analysis": "[Analysis]",
        "target_score": "[integer score 0-10]"
    },
    ...
],
"logical_coherence": [
    {
        "criterion": "[The text of the first Logical Coherence
                      criterion]",
        "analysis": "[Analysis]",
        "target_score": "[integer score 0-10]"
    },
    {
        "criterion": "[The text of the second Logical
                      Coherence criterion]",
        "analysis": "[Analysis]",
        "target_score": "[integer score 0-10]"
    },
    ...
],
"clarity_readability": [
    {
        "criterion": "[The text of the first Clarity &
                      Readability criterion]",
        "analysis": "[Analysis]",
        "target_score": "[integer score 0-10]"
    },
    {
        "criterion": "[The text of the second Clarity &
                      Readability criterion]",
        "analysis": "[Analysis]",
        "target_score": "[integer score 0-10]"
    },
    ...
]
}

```

I.3 PROMPTS IN RELIABILITY EVALUATION

Prompt for Claim Extraction

You will see a research report, and your task is to extract only all verifiable factual statements (factual claims) from the text.

Definition of Factual Statement A factual statement is a verifiable claim about the objective state of the external world. It describes facts that have already occurred, quantifiable data, recognized classifications, or scientific laws — not the author's subjective opinions, intentions, plans, or predictions about the future, nor descriptions about the report's own plans or structure.

Guidelines for Identifying Factual Statements

Types to Extract (Examples)

- Specific data and statistics: "In 2023, global electric vehicle sales reached 14.1 million units."

1620

1621

1622

1623

1624

1625

1626

1627

1628

1629

1630

1631

1632

1633

1634

1635

1636

1637

1638

1639

1640

1641

1642

1643

1644

1645

1646

1647

1648

1649

1650

1651

1652

1653

1654

1655

1656

1657

1658

1659

1660

1661

1662

1663

1664

1665

1666

1667

1668

1669

1670

1671

1672

1673

- Past historical events: "The company was founded in Shanghai, China, in 2010."
- Recognized classifications or definitions: "Li Qiang constructed a socioeconomic status index (SES) based on income, education, and occupation, dividing society into seven classes [15]."
- Cited research findings: "Studies show that more than eight hours of sleep are critical for memory consolidation [8]."

Types to Exclude (Examples)

- Goals and intentions: Any statement describing the "purpose," "goal," or "aim" of this document or project.
 - Example: "The goal of this report is to systematize personal creative activities." or "This project aims to verify the small-revenue model."
- Plans and proposals: Plans about future actions, strategies, or content.
 - Example: "The content pillars include: travel sketch diaries, process breakdowns, tool reviews..." or "We will execute this plan in three phases."
- Self-referential statements about the document: Statements introducing the report's structure or content.
 - Example: "This report is a three-month brand and operations execution manual for..." or "Chapter 3 will discuss market analysis in detail."
- Predictions and speculations: Estimations or guesses about what might happen in the future.
 - Example: "This strategy is expected to increase user stickiness by 20%." or "This could create new business opportunities."
- Opinions and recommendations: The author's subjective judgments, opinions, or suggestions.
 - Example: "We believe this is a key breakthrough." or "Therefore, we recommend adopting Plan A."
- Research methods: Descriptions of how the research or work will be conducted.
 - Example: "This study will adopt a mixed-method approach combining qualitative and quantitative analysis."

Extraction Rules and Output Format For each factual statement you find, determine whether it includes a reference citation, and extract it as a (fact, ref_idx, url) triple. Citations in the text may appear in the following forms:

1. A piece of text + space + number, for example: "Li Qiang constructed a socioeconomic status index (SES) based on income, education, and occupation, dividing society into seven classes 15"
2. A piece of text + [number(s)], for example: "Li Qiang constructed a socioeconomic status index (SES) based on income, education, and occupation, dividing society into seven classes [15]"
3. A piece of text + [number(s)+(some line numbers etc.)], for example: "Li Qiang constructed a socioeconomic status index (SES) based on income, education, and occupation, dividing society into seven classes [15+L10][5L23][7+summary][9summary]"
4. [Cited source](citation link), for example: "According to [ChinaFile: A Guide to Social Class in Modern China](<https://www.chinafile.com/reporting-opinion/media/guide-social-class-modern-china>), Chinese society can be divided into nine classes"

When extracting, pay attention to the following:

1. The extracted fact should be a complete, understandable statement — not just a phrase or fragment.

1674
 1675 2. If a fact cites multiple references, output multiple triples. For example, if it cites
 1676 two references, output (fact, ref_idx_1, url_1) and (fact, ref_idx_2, url_2).
 1677 3. For the third form of citation, only take the first numeric part as ref_idx, ignoring
 1678 indicators of specific locations. For the fourth form (where the source and link
 1679 appear directly in the text), set ref_idx uniformly to 0.
 1680 4. If a factual statement has no citation, set both ref_idx and url to empty strings "".
 1681
 1682 Output Requirements: You should return a JSON list, where each item is one triple. For
 1683 content you are unsure about, err on the side of caution — it's better to miss something than
 1684 to mislabel it. If there are no factual statements in the article, return an empty list [].
 1685 JSON Example:
 1686 [
 1687 {
 1688 "fact": "Text from the original article, use full-width
 1689 Chinese quotation marks, escape English quotes with a
 1690 single backslash",
 1691 "ref_idx": "The index of the cited reference in the
 1692 reference list for this statement; leave empty if none",
 1693 "url": "The link of the cited reference (extracted from
 1694 the report's reference list or from the inline
 1695 citation), leave empty if none"
 1696 },
 1697 {
 1698 "fact": "In 2023, global electric vehicle sales reached
 1699 14.1 million units.",
 1700 "ref_idx": 12,
 1701 "url": "https://iea.org/reports/global-ev-outlook-2024"
 1702 },
 1703 {
 1704 "fact": "Tesla went public on NASDAQ in 2010.",
 1705 "ref_idx": "",
 1706 "url": ""
 1707 },
 1708 {
 1709 "fact": "Studies show that more than eight hours of sleep
 1710 significantly enhances memory consolidation.",
 1711 "ref_idx": 5,
 1712 "url": "https://doi.org/10.1016/j.neurobiol.2020.101945"
 1713 },
 1714 {
 1715 "fact": "According to UNEP
 1716 (<https://www.unep.org/resources/emissions-gap-report-2023>),
 1717 global greenhouse gas emissions reached a record high
 1718 in 2023.",
 1719 "ref_idx": 0,
 1720 "url":
 1721 "<https://www.unep.org/resources/emissions-gap-report-2023>"
 1722 }
 1723]
 1724
 1725 Below is the main text of the research report: {report_text}
 1726 Now start extracting, and directly output the JSON list — do not output any small talk or
 1727 explanation.
 1728
 1729
 1730
 1731
 1732

J RUNNING COST

Table 9: Running Cost Comparison Across Models and Personalization Settings (per query)

Model / Setting	Task Only	Task w/ Persona	Task w/ Context
<i>Open-Source DRAs (GPT-5-Mini based)</i>			
OAgents	\$1.60	\$1.70	\$3.00
DeerFlow	\$0.50	\$0.57	\$1.20
MiroFlow	\$1.00	\$1.11	\$2.10
<i>LLM with Search Tools</i>			
Gemini-2.5-Pro w/ Search	\$0.05	\$0.06	\$0.24
Claude-3.7-Sonnet w/ Search	\$0.03	\$0.04	\$0.31
Perplexity-Sonar-Reasoning-Pro	\$0.02	\$0.03	\$0.78
GPT-4.1 w/ Search	\$0.02	\$0.02	\$0.31