

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  

# AN INTERACTIVE PARADIGM FOR DEEP RESEARCH

Anonymous authors

Paper under double-blind review

## ABSTRACT

Recent advances in large language models (LLMs) have enabled deep research systems that synthesize comprehensive, report-style answers to open-ended queries by combining retrieval, reasoning, and generation. Yet most frameworks rely on rigid workflows with one-shot scoping and long autonomous runs, offering little room for course correction if user intent shifts mid-process. We present **STEER**, a framework for **Steerable deEp Research** that introduces interpretable, mid-process control into long-horizon research workflows. At each decision point, **STEER** uses a cost–benefit formulation to determine whether to pause for user input or to proceed autonomously. It combines diversity-aware planning with utility signals that reward alignment, novelty, and coverage, and maintains a live persona model that evolves throughout the session. **STEER** outperforms state-of-the-art open-source and proprietary baselines by up to 22.80% on alignment, leads on quality metrics such as breadth and balance, and is preferred by human readers in 85%+ of pairwise alignment judgments. We also introduce a persona–query benchmark and data-generation pipeline. To our knowledge, this is the first work to advance deep research with an interactive, interpretable control paradigm, paving the way for controllable, user-aligned agents in long-form tasks.

## 1 INTRODUCTION

Recent advances in large language models (LLMs) have shifted information access from ranked retrieval to systems that generate comprehensive, report-style answers to complex and open-ended queries (Du et al., 2025). These *deep research* systems, spanning proprietary platforms (Google, 2024; OpenAI, 2025; xAI, 2025) and open-source frameworks (Elovic, 2025; LangChain, 2025), combine iterative retrieval with multi-step reasoning to synthesize well-supported outputs (Coelho et al., 2025). Benchmarks such as DeepResearchGym (Coelho et al., 2025) and DeepResearch Bench (Du et al., 2025) have begun to standardize this setting, providing realistic research-style questions and automated evaluation protocols for long-form, citation-heavy reports. On the system side, current deep research agents largely fall into two paradigms: multi-agent pipelines that divide planning, search, and synthesis (Huang et al., 2025; Alzubi et al., 2025; Li et al., 2025a; Zhang et al., 2025a), and RL-trained agents that learn to search and reason effectively (Zheng et al., 2025b; Jin et al., 2025; Song et al., 2025). Yet, regardless of architecture, most systems follow a rigid workflow: one-time scoping (often with a single clarification), followed by a long autonomous run. If user intent shifts mid-process, there is little room to course-correct, resulting in wasted cost and misaligned reports. This highlights the need for an alternative design where mid-process interaction is central, not optional.

Two research threads closely relate to our work. *Personalization and alignment* examine how to tailor LLM outputs to user intent, from profile-conditioned generation (Wu et al., 2025) to long-form checklists (Salemi et al., 2024; Salemi & Zamani, 2025) and interactive preference elicitation. While these works show the value of personalization, most assume fixed personas or separate preference modeling from system control, lacking a principled way to determine when to seek input. *Interactive reasoning* investigates how LLMs ask clarifying questions (Andukuri et al.; Ren et al., 2023; Wu et al., 2024), model future turns (Zhang et al., 2025b), or learn clarification policies (Chen et al., 2025). Tools like **INTERACTIVE REASONING** (Pang et al., 2025) and **REASONGRAPH** (Li et al., 2025c) enhance transparency, but focus on local clarification or static visualization. Existing approaches thus either optimize autonomous agents or isolate clarification as a narrow skill. In contrast, we aim to offer an integrated control paradigm that governs when to pause, what to explore, and how to update personalization mid-process.

054	System	Mid-process steering	Adaptive pause decision	Live persona modeling
<i>Deep Research Framework (Open-Source)</i>				
056	DeepResearcher (Zheng et al., 2025b)	✗	✗	✗
057	Search-r1 (Jin et al., 2025)	✗	✗	✗
058	ManuSearch (Huang et al., 2025)	✗	✗	✗
059	Search-o1 (Li et al., 2025a)	✗	✗	✗
060	GPT-Researcher (Elovic, 2025)	✗	✗	✗
061	Open Deep Research (LangChain, 2025)	◆	✗	✗
<i>Deep Research Framework (Proprietary Web-Based)</i>				
062	OpenAI Deep Research (OpenAI, 2025)	◆	✗	✗
063	Google Gemini Deep Research (Google, 2024)	◆	✗	✗
<i>Interactive/Persona-Aware Reasoning Frameworks</i>				
064	PersonaAgent (Zhang et al., 2025c)	✗	✗	✓
065	ReasonGraph (Li et al., 2025b)	✓	✗	✗
066	HITL CoT MCS (Cai et al., 2023)	✓	✗	✗
067	<b>STEER</b>	✓	✓	✓

069 Table 1: Comparison of deep research frameworks (open-source and proprietary web-based) and  
070 interactive/persona-aware reasoning frameworks in terms of mid-process steering via adaptive pause  
071 decisions and live persona modeling. To our knowledge, **STEER** is the first benchmarkable deep-  
072 research framework that jointly offers all.

073  
074 We introduce **STEER**, a framework for **Steerable deEp Research** that brings interactive control to  
075 long-horizon research workflows (Figure 1). The key intuition is that deep research should occa-  
076 sionally *ask*, not just *answer*: **STEER** uses a cost–benefit formulation at each decision point to  
077 determine whether to pause for user input or to proceed autonomously. To remain both user-aligned  
078 and exploratory, it combines diversity-aware planning with utility signals that reward alignment,  
079 novelty, and coverage. A live persona is continuously updated based on interactions and conditions  
080 all downstream planning, scoring, and synthesis, enabling the system to adjust as user needs evolve.

081 Table 1 positions **STEER** relative to both deep research frameworks and interactive/persona-aware  
082 reasoning systems. Open-source frameworks and proprietary web-based services overwhelmingly  
083 adopt a long autonomous-run paradigm: once a query is issued, a fixed pipeline executes to com-  
084 pletion with no interpretable control over where in the research tree to intervene or when to ask  
085 for guidance (the ◆ marks indicate only a single upfront scoping or plan-confirmation step before  
086 a fully autonomous run). Interactive/persona-aware systems address complementary aspects: they  
087 either focus on maintaining an evolving user representation, or expose the full reasoning process for  
088 inspection and manual correction, but they do not implement an adaptive pause policy that decides  
089 *where* to pause and *how* to balance autonomy and control. In contrast, **STEER** is, to our knowl-  
090 edge, the first benchmarkable deep research framework that jointly supports mid-process steering,  
091 an adaptive pause mechanism, and live persona modeling.

092 Our contributions are as follows:

093 • We propose **STEER**, an interactive deep research framework that supports interpretable, mid-  
094 process control and dynamic user alignment throughout the research loop.

095 • Extensive experiments show that **STEER** outperforms the strongest open-source and proprietary  
096 OpenAI baselines on persona-tailoredness and report quality, while offering fine-grained control to  
097 tune trade-offs between alignment and user burden, as well as between under-exploration and over-  
098 personalization. A human study further confirms its preference among readers, with significant  
099 gains in alignment, focus, and usability.

100 • We introduce a persona–query evaluation suite and a reusable data generation pipeline grounded  
101 in prior benchmarks, suitable for future evaluation and training of interactive deep research agents.

102 In summary, **STEER** consistently outperforms strong open and proprietary baselines, achieving  
103 7.83%–22.80% higher alignment and leading on general quality metrics such as breadth and bal-  
104 ance. Human readers prefer **STEER** in over 85% of alignment and 83% of focus pairwise com-  
105 parisons. To our knowledge, *this is the first work to advance deep research with an interactive,*  
106 *interpretable control paradigm*. We believe that this paradigm shift will shape future long-horizon  
107 research agents, enabling decision policies that adapt to individual users and their evolving needs,  
108 rather than relying on a single upfront clarification.

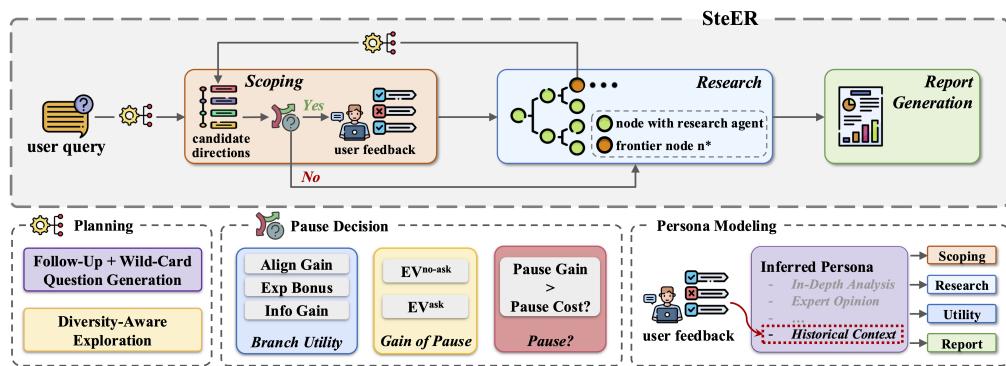
## 108 2 STEER

### 110 2.1 PROBLEM SETUP AND OBJECTIVES

112 We formulate steerable deep research as an interactive planning task. Given a user query  $Q$ , the  
 113 system incrementally constructs a research tree and produces a cited synthesis report  $R$ . The goal is  
 114 to generate a report that is both high-quality and aligned with the evolving preferences of the user,  
 115 while keeping the number of interruptions minimal and well-timed.

116 Each user is represented by a persona  $P = (p_{\text{text}}, \mathcal{A})$ , where  $p_{\text{text}}$  is a natural-language description  
 117 combining profile and personality traits (following Wu et al. (2025)), and  $\mathcal{A}$  is a set of aspects the  
 118 user expects to see addressed in the final report. We evaluate reports along two complementary  
 119 dimensions: **(1) Alignment**: the extent to which the report covers the aspects in  $\mathcal{A}$ ; and **(2) Focus**:  
 120 the proportion of content that remains on-topic with respect to  $\mathcal{A}$ .

### 122 2.2 SYSTEM OVERVIEW



131 Figure 1: Overview of **STEER**. The upper panel shows the end-to-end pipeline. The lower panels  
 132 zoom into the three core modules: *Planning*, *Pause Decision*, and *Persona Modeling*.

133 Our framework **STEER** transforms monolithic deep-research pipelines into an interactive process.  
 134 The system is structured around three core components: **diversity-aware exploration**, **pause decision**, and **persona modeling**. At a high level, the framework incrementally builds a research tree  
 135 that represents possible exploration paths and selectively engages the user at key checkpoints.

136 We denote the research tree as  $T = (N, E)$ , where each node  $n \in N$  represents a sub-problem  
 137 query with partial research results and each edge  $(n, n') \in E$  indicates a decomposition into sub-  
 138 directions. The tree expands level by level up to a maximum depth  $D$ , a hyper-parameter controlling  
 139 how many layers are explored before synthesis. At each step, the system operates at a **frontier node**  
 140  $n^*$  and performs the following actions (Figure 1):

- 141 **1. Diversity-aware exploration:** Generate candidate follow-up directions from  $n^*$  and select a  
 142 diversified subset of size  $K$  to serve as potential expansions (see the *Planning* panel in Figure 1).
- 143 **2. Pause decision and expansion:** Compute branch utilities, execution costs, and the expected  
 144 gain of asking, and compare this to the pause cost. If a pause occurs, present the diversified  
 145 subset to the user and then expand the user-selected items together with any newly suggested  
 146 directions. Otherwise, expand the system-proposed diversified subset directly. Sub-agents then  
 147 perform retrieval and reasoning at each expanded child to produce node-level reports (see the  
 148 *Pause Decision* panel in Figure 1).
- 149 **3. Persona modeling:** Update the inferred persona  $\hat{P}$  with signals from the query, initial profile,  
 150 and any user feedback gathered during pauses. The updated persona conditions planning, utility  
 151 scoring, and synthesis in subsequent steps (see the *Persona Modeling* panel in Figure 1).

152 The process terminates once all nodes at depth  $D$  have been expanded, at which point the accumulated  
 153 node reports are aggregated into the final report  $R$ . This interactive loop enables reports that  
 154 are better aligned with user goals while minimizing redundant or off-topic exploration.

162 2.3 DIVERSITY-AWARE EXPLORATION  
163164 As described in Section 2.2, at each frontier node  $n^*$ , the system generates a set of follow-up ques-  
165 tions as potential next steps. To promote exploration and reduce redundancy, we explicitly prompt  
166 for *distinct facets* and include one *wild-card* direction (see Appendix K for prompt details). From  
167 this candidate set, we select a diversified subset of size  $K$  to either present to the user (if a pause is  
168 triggered) or expand automatically.169 To select this subset, we apply a greedy Maximal Marginal Relevance (MMR) strategy (Carbonell  
170 & Goldstein, 1998; Wang et al., 2025), which balances confidence scores with dissimilarity to pre-  
171 viously chosen directions. MMR is particularly well-suited to our setting: it is simple, efficient, and  
172 interpretable, while effectively encouraging topical coverage across different aspects. In contrast,  
173 alternative diversity methods (e.g., clustering or determinantal point processes) introduce additional  
174 complexity and hyper-parameters without clear gains in this context. Appendix C provides the full  
175 algorithmic details.176 2.4 PAUSE DECISION AND EXPANSION  
177178 After the proposal stage has produced a diversified set of candidates, the system must decide whether  
179 to involve the user or continue autonomously. Asking everywhere is undesirable: user tolerance for  
180 interruptions is limited and varies widely. Some users prefer high-level guidance while trusting the  
181 system to handle details; others are more detail-oriented but want control only in specific themes.  
182 Preferences also shift across the depth of the research tree and over time. A well-calibrated system  
183 must respect these preferences while steering the exploration toward the user’s goals. Below, we  
184 present the pause decision mechanism in a top-down structure: we begin with the overall decision  
185 rule and then unpack its components, including pause cost, expected gain, and branch utility.186 **Decision rule** At each frontier node  $n^*$ , the system evaluates whether pausing to ask the user is  
187 beneficial. This decision is framed as a cost–benefit comparison:

188 
$$a(n^*) = \begin{cases} \text{PAUSEASK}, & \Delta EV(n^*) > C(n^*), \\ \text{PROCEED}, & \text{otherwise.} \end{cases}$$

189 Here,  $\Delta EV(n^*)$  denotes the expected utility gain from pausing – by allowing the user to refine  
190 or redirect the next steps – while  $C(n^*)$  denotes the cost of interruption, scaled by user-specific  
191 tolerance.192 **Pause cost** Not all users interact in the same way. To model this, we assume two things: (1) a  
193 user’s tolerance for interruptions decreases over time, and (2) users differ in how much interruption  
194 they are willing to tolerate in total, and how fast that tolerance depletes.

195 To capture this, we introduce two hyper-parameters:

200 

- 201 •  $C_0 \in [0, 1]$ : the *base pause cost*. This reflects a user’s general sensitivity to interruptions. A  
202 lower  $C_0$  implies the user is open to frequent interaction; a higher  $C_0$  indicates a preference  
203 for minimal disruption.
- 204 •  $\text{Tol} \in \mathbb{N}$ : the *tolerance budget*. This governs how quickly the pause cost increases with the  
205 number of questions asked. Intuitively, Tol represents the approximate number of clarifi-  
206 cation questions the user is comfortable answering across the entire session.

207 A user may tolerate multiple clarifications within a single topic but become frustrated by interrup-  
208 tions scattered across too many unrelated ones. To reflect this, we distribute the global tolerance  
209 budget Tol across all active *top-level directions*, defined as the root’s immediate children. While  
210 users may have different preferences across themes, we simplify by evenly dividing the tolerance  
211 budget across top-level directions. Each node  $n$  belongs to a top-level direction  $j \in K'$ , where  $K'$   
212 denotes the number of currently active directions. If the system proceeds automatically,  $K' = K$   
213 (the full diversified set). If a pause occurs,  $K'$  equals the number of user-selected plus user-added  
214 directions. The pause cost at a frontier node  $n^*$  is then computed as:

215 
$$C(n^*) = C_0 \cdot \left(1 + \frac{\text{pauses}_j}{\text{Tol}_j}\right),$$

216 where  $\text{pauses}_j$  is the number of times the system has previously paused in direction  $j$ . As the  
 217 number of pauses grows within a direction, the cost of pausing again increases proportionally.  
 218

219 **Pause gain** The *gain of pausing* should reflect two factors: the utility we forgo by pruning  
 220 branches and the execution cost we save by not pursuing them.  
 221

222 Let  $\{n_k^*\}_{k=1}^K$  be the candidate children at the frontier node, with branch utilities  $U(n_k^*)$  and normal-  
 223 ized execution costs  $C^{\text{exec}}(n_k^*)$ . If we proceed automatically, we pursue all  $K$ , and the expected  
 224 value of the frontier node without pausing is  $EV^{\text{no-ask}}(n^*) = \sum_{k=1}^K U(n_k^*) - \sum_{k=1}^K C^{\text{exec}}(n_k^*)$ .  
 225 If we pause, the user keeps a subset  $S \subseteq \{1, \dots, K\}$ , so the expected value of the frontier node  
 226 with pause is  $EV^{\text{ask}}(n^*) = \sum_{k \in S} U(n_k^*) - \sum_{k \in S} C^{\text{exec}}(n_k^*)$ . To estimate  $S$ , we retain candidates  
 227 whose upper utility bound overlaps the leader’s lower bound, capturing all options that are plausibly  
 228 optimal. Equivalently, this decision rule prunes all branches whose best-case utility still falls below  
 229 the worst-case value of the current leader. See Appendix D for bound construction and filtering.  
 230

231 A pause only changes which branches we do *not* execute. The gain of pausing at the frontier node  
 232 is the saved cost minus the lost utility of those pruned branches:  
 233

$$\Delta EV(n^*) = EV^{\text{ask}}(n^*) - EV^{\text{no-ask}}(n^*) = \sum_{k \in S^c} (-U(n_k^*) + C^{\text{exec}}(n_k^*)).$$

235 **Branch utility.** We score each candidate child  $n_k^*$  using a weighted combination of three factors:  
 236

$$U(n_k^*) = \Delta \text{Align}(n_k^*) + \lambda_{\text{exp}} \text{Explore}(n_k^*) + \lambda_{\text{info}} \text{InfoGain}(n_k^*),$$

237 where each component is scaled to  $[0, 1]$  for direct comparability with the pause cost. (See Ap-  
 238 pendix D for exact computations and normalization.)  
 239

- 240 • **Alignment gain** ( $\Delta \text{Align}$ ) computes predicted increase in persona alignment relative to the parent  
 241 under the current inferred aspects  $\hat{A}_s$ . It rewards branches that cover more of what the user actually  
 242 cares about.
- 243 • **Exploration bonus** ( $\text{Explore}$ ) adds a small reward for under-explored facets to discourage repeat-  
 244 edly selecting the same angle. We capture this “reward under-explored, penalize over-explored”  
 245 behavior using a Upper Confidence Bound (UCB) algorithm (Auer, 2002; Auer et al., 2002; Li  
 246 et al., 2010), which assigns larger bonuses to rarely used facets and decays naturally as they are  
 247 chosen more frequently.
- 248 • **Information gain** ( $\text{InfoGain}$ ) measures the content-level novelty of a candidate’s expected  
 249 evidence relative to accumulated learnings. While  $\text{Explore}$  encourages facet-level diversity,  
 250  $\text{InfoGain}$  focuses on semantic-level novelty, prioritizing branches that are more likely to yield  
 251 genuinely new information from the web.

252  $\lambda_{\text{exp}}$  and  $\lambda_{\text{info}}$  balance breadth and novelty against alignment. Both  $\text{Explore}$  and  $\text{InfoGain}$  com-  
 253 plement the diversify-aware exploration described in Section 2.3: while the latter ensures that the  
 254 *initial question set* spans distinct facets, it does not guarantee that the resulting content will be di-  
 255 verse.  $\text{Explore}$  and  $\text{InfoGain}$  help mitigate this by promoting long-term diversity at the facet and  
 256 content levels, respectively. While our process uses a minimal three-factor utility for clarity and  
 257 stability, the framework is easily extensible – additional criteria (e.g., risk, credibility) can be incor-  
 258 porated as needed.

259 **Execution cost.**  $C^{\text{exec}}(n_k^*)$  estimates remaining work if we expand  $n_k^*$ . It is also normalized to  $[0, 1]$   
 260 so it is commensurate with utilities. We approximate the cost by the tokens of a saturated subtree  
 261 beneath  $n_k^*$ , as tokens provide a consistent, model-agnostic proxy, and correlate with both latency  
 262 and spend. See Appendix D for computation details.  
 263

## 264 2.5 PERSONA MODELING

265 Beyond deciding *when* to ask (Section 2.4), the system must also know *who* it is optimizing for.  
 266 In deep research, users often do not know exactly what they want at the start. Their goals shift as  
 267 they encounter new information, and partial results may reveal new priorities. Fixing a full persona  
 268 upfront risks overfitting to stale assumptions or flooding the system with irrelevant detail. To address  
 269 this, we maintain a *live* persona that evolves dynamically as the research progresses.

At each  $n^*$ , **STEER** maintains an updated persona estimate  $\hat{P}(n^*) = (\hat{p}_{\text{text}}(n^*), \hat{\mathcal{A}}(n^*))$ , where  $\hat{p}_{\text{text}}(n^*)$  captures the user’s profile and  $\hat{\mathcal{A}}(n^*)$  represents the current inferred set of aspects the user cares about. When a pause occurs, we update  $\hat{P}(n^*)$  based on user-selected directions and any new suggestions, and implicitly incorporate recent research findings. This evolving persona conditions all downstream modules: it guides research and follow-up question generation, shapes the branch utility score via alignment to  $\hat{\mathcal{A}}(n^*)$  (decision), and steers final report synthesis. See Appendix K for full details on how the persona is inferred and updated using LLM prompts (*Persona Checklist Inference* and *Persona Modeling* prompts), and how the evolving  $\hat{P}(n^*)$  is used across the planning, research, and synthesis pipeline.

A live persona keeps the interaction tightly aligned with the user’s current interests. It prevents drift caused by outdated assumptions, reduces unnecessary questions by filtering irrelevant directions, and adapts to new priorities that emerge during exploration.

### 3 EXPERIMENTS

#### 3.1 EXPERIMENTAL SETUP

**Evaluation data** We synthesize query–persona pairs by adapting established datasets and methods, with light modifications to better suit our goals. We begin with 1k queries from the *Researchy Questions* dataset (Rosset et al., 2024), as used in *DeepResearchGym* (Coelho et al., 2025). For each query, we generate a plausible user persona  $p_{\text{text}}$  by adapting the ALOE profile–personality paradigm (Wu et al., 2025): we seed from ALOE profiles and prompt an LLM to propose new profiles that would reasonably ask the given query. To ensure diversity, we apply SBERT-based filtering (Reimers & Gurevych, 2019) and keep only distinct, plausible personas, following prior work (Wu et al., 2025; Wang et al., 2023).

Given each  $p_{\text{text}}$ , we generate 5–8 evaluation aspects  $\mathcal{A}$  using prompts inspired by Salemi & Zamani (2025), following their checklist format to ensure that the aspects are actionable, measurable, and grounded in the persona. This enables robust alignment and focus evaluation, avoiding the ambiguity of more generic rubrics.

Compared to Wu et al. (2025) and Salemi & Zamani (2025), our adaptations are minimal but tailored to deep research: (i) persona generation is query-conditioned to ensure relevance, (ii) diversity filtering is stricter to avoid near-duplicates, and (iii) aspects are framed for long-form, cited outputs. We evaluate on a held-out set of 200 queries. Full details of data generation are in Appendix E.

**User Agent simulation** To enable scalable, repeatable evaluation, we simulate user interactions with a *User Agent* conditioned on the full persona  $P = (p_{\text{text}}, \mathcal{A})$ . The agent selects directions that best align with  $\mathcal{A}$  and proposes a new follow-up when uncovered aspects remain, yielding realistic steering signals without human-in-the-loop variability. (See Appendix N for the full prompt.)

**Metrics** We evaluate persona-tailored quality using two proposed metrics: **Alignment** and **Focus**, both judged by *gpt-4.1-mini* following *DeepResearchGym*. (Prompts used to obtain these metrics are listed in Appendix M.) We present the meta-evaluation results for the LLM judge in Appendix I.

- **Alignment:** Given aspect set  $\mathcal{A}$  and report  $R$ , we compute:  $\text{Align}(R, \mathcal{A}) = \frac{1}{2|\mathcal{A}|} \sum_{a \in \mathcal{A}} \text{align}(R, a)$ ,  $\text{align}(R, a) \in \{0, 1, 2\}$ . Here, 0 means that the aspect is not addressed, 1 means that it is partially addressed (e.g., mentioned or vaguely covered), and 2 means that it is fully addressed with sufficient detail and evidence, all scored by the LLM-judge. This gives an interpretable, per-aspect measure of user alignment.
- **Focus:** We extract a set of keypoints  $\mathcal{KP}$  – short, evidence-bearing spans – from  $R$  using an LLM, and then ask the judge whether each keypoint ( $k \in \mathcal{KP}$ ) maps to at least one user aspect:  $\text{Focus}_{\mathcal{KP}}(R, \mathcal{A}) = \frac{1}{|\mathcal{KP}|} \sum_{k \in \mathcal{KP}} \mathbf{I}[\text{map}(k) \neq \emptyset]$ . While alignment is akin to *recall* over aspects, focus acts as a form of *precision*, rewarding dense, on-target content.

In addition, we report *DeepResearchGym*’s quality metrics, including clarity, depth, breadth, and insight, to evaluate general writing quality beyond persona targeting.

**Baselines** We compare **STEER** to two strong open-source frameworks: *GPT-Researcher* (Elovic, 2025) and *Open Deep Research* (LangChain, 2025), both evaluated as top-performing frameworks (Coelho et al., 2025). On the proprietary model side, we benchmark against OpenAI’s o4-mini-deep-research model.

We compare systems under a controlled setting: for **STEER** and the open-source frameworks, all agents use GPT-4o, the research tree is fixed (depth 3, breadth 3), outputs share the same token cap, and the only variable is persona information. For fairness, all baselines are run under three input settings: (1) query only, (2) query + initial persona (first sentence of  $p_{text}$ ), and (3) query + full persona. This allows us to assess how well each baseline adapts to different levels of user information. Note that **STEER** always operates with only the initial persona, and must infer preferences dynamically throughout the interaction.

### 3.2 HOW MUCH DOES STEER IMPROVE PERSONA-TAILORED QUALITY?

Metric →	Persona-Tailored		Quality				
System ↓	Align	Focus <sub>kp</sub>	Clarity	Depth	Breadth	Insight	Balance
GPT-Researcher	66.63	78.42	<b>81.80</b>	86.30	88.40	76.60	81.25
GPT-Researcher <sub>initial-persona</sub>	74.59	81.68	79.05	87.37	88.71	75.52	81.71
GPT-Researcher <sub>full-persona</sub>	<b>79.48</b>	83.83	77.93	87.09	<b>90.31</b>	79.05	82.58
OpenDeepResearch	62.74	83.72	74.90	82.40	88.85	68.39	81.25
OpenDeepResearch <sub>initial-persona</sub>	69.79	85.45	72.51	81.64	84.12	68.98	77.44
OpenDeepResearch <sub>full-persona</sub>	77.20	<b>86.10</b>	74.02	83.42	87.62	73.18	79.44
o4-mini-deep-research <sub>initial-persona</sub>	72.73	86.09	75.76	<b>89.10</b>	89.51	<b>86.74</b>	<b>82.76</b>
o4-mini-deep-research <sub>full-persona</sub>	75.72	86.02	75.54	87.19	87.36	<b>85.01</b>	82.63
<b>STEER</b>	<b>85.70</b>	<b>86.45</b>	<b>79.97</b>	<b>88.67</b>	<b>91.29</b>	83.04	<b>84.19</b>

Table 2: Performance comparison between **STEER** and baselines. For **STEER**, we report performance at  $C_0 = 0.7$  (see Section 3.3 for selection rationale).

From Table 2, we see that **STEER** achieves the strongest persona-tailored performance on both metrics across all systems (e.g. 7.83% higher alignment than the runner-up GPT-Researcher<sub>full-persona</sub>), even though some of those baselines are given the full persona, while **STEER** only receives the first sentence. This highlights the effectiveness of **STEER**’s interactive pausing and live persona modeling, which enable accurate mid-process adaptation without relying on full upfront persona input. This has practical appeal: real-world deployments often face privacy constraints, onboarding friction, or noisy user profiles. **STEER**’s ability to achieve strong alignment under minimal initial input makes it more robust in such settings.

**STEER** also leads in breadth and balance, reflecting the role of **STEER**’s diversity-aware exploration and utility components, Explore and InfoGain, in promoting semantic novelty and facet diversity. **STEER** also significantly outperforms the open-source baselines in depth and insight, though it falls slightly short of the proprietary OpenAI model on these metrics.

### 3.3 HOW DOES STEER PROVIDE INTERPRETABLE CONTROLS FOR OPTIMAL PAUSING?

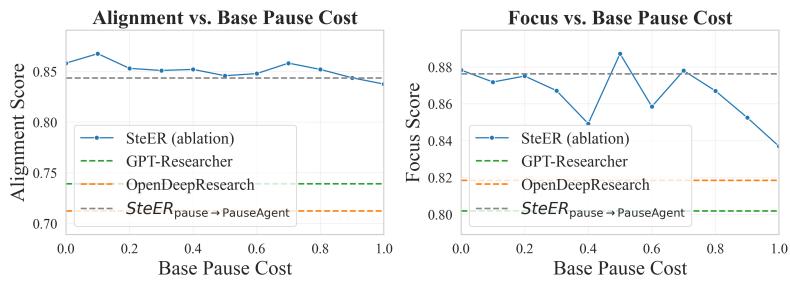


Figure 3: Effect of base pause cost on alignment (left) and focus (right). Baseline scores are shown as horizontal reference lines for comparison.

378 As introduced in Section 2.4, **STEER** offers two interpretable knobs to control pausing behavior: the base  
 379 pause cost  $C_0$ , which sets the system’s aversion to in-  
 380 terruptions, and the tolerance budget  $\text{Tol}$ , which controls  
 381 how quickly pause cost grows within a top-level direc-  
 382 tion. In this study, we vary  $C_0$  while keeping  $\text{Tol} = 3$   
 383 fixed. This is because, in shallow trees (depth 3), the  
 384 effect of  $\text{Tol}$  is limited.  $\text{Tol}$  is more impactful in long-  
 385 horizon tasks where user fatigue may accumulate across  
 386 levels. Conceptually,  $\text{Tol}$  captures a user-specific interac-  
 387 tion limit. Once set, we can tune  $C_0$  based on two criteria:  
 388 **(a)** which value yields pause counts closest to  $\text{Tol}$ , and **(b)**  
 389 which value yields the highest gain per pause.  
 390

391 To benchmark against intuitive alternatives, we introduce  
 392 a *PauseAgent* baseline that uses an LLM agent to pre-  
 393 dict pause vs. proceed at each frontier node (prompt in Appen-  
 394 dix F.2). *PauseAgent* pauses excessively, far exceeding the  $\text{Tol} = 3$  budget. In contrast, **STEER**  
 395 with  $C_0 \geq 0.4$  remains within budget, averaging fewer than 3 pauses.

396 Frequent pausing also hurts efficiency. Figure 2 shows that alignment per pause drops sharply at  
 397 low  $C_0$ , while higher  $C_0$  yields fewer but more impactful interventions. This trade-off is evident in  
 398 Figure 3: while absolute alignment declines as  $C_0$  increases, both alignment and focus reach local  
 399 maxima around  $C_0 = 0.7$ , suggesting it as a practical sweet spot.

400 In summary, **STEER** supports calibrated control of interaction.  $C_0$  adjusts interruption cost directly,  
 401 and  $\text{Tol}$  governs how that cost compounds over time. This formulation provides both interpretability  
 402 and personalization, outperforming the *PauseAgent* baseline in effectiveness and flexibility.

### 404 3.4 HOW DOES **STEER** AVOID UNDER-EXPLORATION DRIVEN BY PERSONALIZATION?

406 A potential failure mode is overfitting to per-  
 407 sonalization: when optimization focuses solely  
 408 on aspect alignment, the system quickly col-  
 409 lapses to a narrow trajectory, branch utilities  
 410 flatten as  $\Delta\text{Align}$  approaches zero, and explo-  
 411 ration stalls. To prevent this, **STEER** integrates  
 412 three complementary signals at different axes.

413 First, diversity-aware exploration ensures that  
 414 research directions span distinct facets at each  
 415 step, avoiding early myopia. As shown in Table 3, removing this component causes the largest drops  
 416 in depth, breadth, and focus, along with a significant decline in alignment, underscoring its role in  
 417 maintaining structural and semantic diversity throughout the session.

418 In addition, two utility terms guide exploration: *Explore* encourages rotation across underrepre-  
 419 sented facets, while *InfoGain* prioritizes semantic novelty. Ablating *Explore* leads to a large focus  
 420 drop and notable declination in depth and breadth, with only a small impact on alignment, show-  
 421 ing its importance in sustaining report-wide diversity. In contrast, removing *InfoGain* yields the  
 422 largest alignment drop but only relatively modest effects on other metrics. This suggests that with-  
 423 out semantic novelty, the system tends to dig deeper into already-favored lines, satisfying more  
 424 user aspects while producing redundant evidence. These complementary behaviors introduce an  
 425 interpretable trade-off: increasing  $\lambda_{\text{info}}$  prioritizes aspect satisfaction, while increasing  $\lambda_{\text{exp}}$  favors  
 426 breadth and coverage. We set both to 0.5 for balance, but these can be tuned to suit different tasks.

427 While our experiments focus on novelty and exploration, the utility function is extensible. Addi-  
 428 tional signals, such as factuality or plausibility, can be integrated into the same calibrated frame-  
 429 work. Our primary contribution lies not in these specific factors, but in the interaction paradigm that  
 430 supports modular, interpretable control over research behavior.

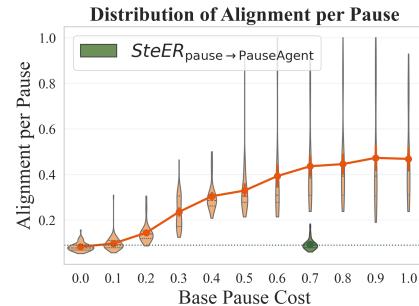
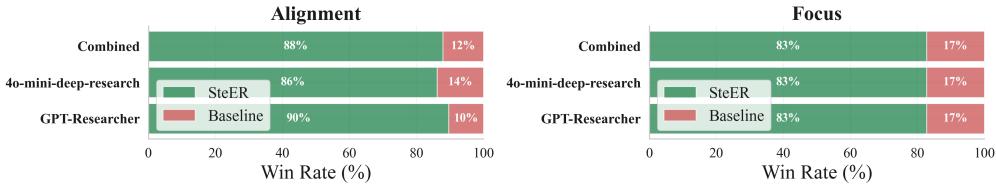


Figure 2: Alignment per pause across base pause cost values.

Method ↓	Alignment	Focus <sub>kp</sub>	Depth	Breadth
<b>STEER</b>	<b>85.82</b>	<b>87.79</b>	90.27	<b>93.15</b>
(w/o) Explore	84.98 <sub>±0.98%</sub>	85.17 <sub>±2.98%</sub>	89.86 <sub>±0.45%</sub>	92.60 <sub>±0.59%</sub>
(w/o) InfoGain	82.81 <sub>±3.51%</sub>	86.40 <sub>±1.58%</sub>	<b>90.41</b> <sub>±0.15%</sub>	92.73 <sub>±0.45%</sub>
(w/o) Div Explore	84.57 <sub>±1.46%</sub>	84.29 <sub>±3.99%</sub>	88.63 <sub>±1.82%</sub>	91.09 <sub>±2.21%</sub>

Table 3: Ablation study on novelty and explo-  
 ration components. Darker red indicates a larger  
 performance drop relative to **STEER**.

432 3.5 USER STUDY EVALUATION  
433441 Figure 4: Pairwise human preference win rates on *Alignment* and *Focus*.  
442

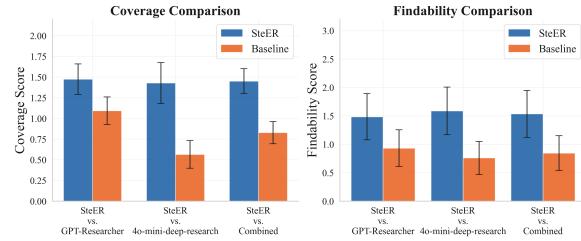
443 To complement the automated LLM-  
444 judged metrics, we conducted a user  
445 study to evaluate whether **STEER** is  
446 preferred by human users. We com-  
447 pared **STEER** with GPT-Researcher  
448 and `o4-mini-deep-research` on  
449 20 query-persona pairs. 12 annotators  
450 (all NLP/CS graduate students) viewed  
451 two reports for the same pair (one from  
452 **STEER**, one from a baseline) in  
453 randomized order on our custom annotation  
454 platform. For each comparison, annotators  
455 judged **Alignment** (better coverage of  
456 persona aspects), **Focus** (more on-topic with  
457 less redundancy), **Coverage** (aspect-level  
458 0–2, averaged), and **Findability** (report-level  
459 0–2 for ease of locating relevant information).  
460

461 This design captures both quality and usability: *Alignment* and *Focus* reflect perceived persona-fit<sup>1</sup>;  
462 *Coverage* measures how thoroughly user interests are addressed, and *Findability* assesses how easily  
463 users can locate what matters. Full platform design and annotator instructions are in Appendix G.

464 We collected 58 valid pairwise annotations. To val-  
465 idate the quality of the annotations, we computed  
466 agreement using pairwise metrics across all anno-  
467 tator pairs for each evaluation dimension. Speci-  
468 fically, we report raw pairwise agreement and Gwet’s  
469 AC1 (Gwet, 2008), a prevalence-resistant chance-  
470 corrected agreement coefficient that avoids the arti-  
471 ficial deflation often observed with Fleiss’  $\kappa$  (Fleiss,  
472 1971) when the label distribution is skewed. As  
473 summarized in Table 4, for *Alignment* and *Focus*  
474 pairwise preferences, annotators achieve raw pair-  
475 wise agreement of 82.0% and 73.8%, with AC1  
476 values of 0.639 and 0.475, respectively, indicating  
477 substantial and moderate agreement. For *Coverage*  
478 and *Findability*, we observe similar patterns of fair to moderate agreement. For *Coverage*, raw  
479 agreement is 65.9% for **STEER** and 65.2% when aggregating baselines (GPT-Researcher and  
480 `o4-mini-deep-research`). For *Findability*, raw agreement is 75.4% for **STEER** and 65.6%  
481 for the aggregated baselines. These values reflect consistent, non-trivial consensus across annotators  
482 on all dimensions, especially given the inherent subjectivity of report quality.

483 As shown in Figure 4, **STEER** is preferred in about 86–90% of cases for *Alignment* and about 83%  
484 for *Focus* across GPT-Researcher and `o4-mini-deep-research`.

485 Figure 5 shows significant gains in *Coverage* and *Findability* for **STEER**. On a 0–2 aspect-coverage  
486 scale, **STEER** improves the average by +0.623 (from 0.828 to 1.451,  $p = 3.05e - 12$ ), a relative  
487 improvement of about 75% which indicates a shift from below “somewhat covered” toward between

488 Figure 5: Human ratings on *Coverage* and *Findability*.  
489 *Left*: Average aspect-level *Coverage* scores of **STEER**  
490 and baselines. *Right*: Average *Findability* scores of  
491 **STEER** and baselines.

Metric	Raw Agreement (%)	Gwet’s AC1
<i>Alignment</i>	82.0	0.639
<i>Focus</i>	73.8	0.475
<i>Coverage</i>		
SteER	65.9	0.318
Baselines	65.2	0.303
<i>Findability</i>		
SteER	75.4	0.508
Baselines	65.6	0.311

492 Table 4: Inter-annotator agreement for user  
493 study annotations.

<sup>1</sup>Note that the *Alignment* and *Focus* metrics used in the user study are based on pairwise human preferences and are not directly comparable to the automatic metrics defined in Section 3.1.

486 “somewhat” and “fully” covered. On the 0–2 Findability scale, **STEER** improves by +0.690 (from  
 487 0.845 to 1.534,  $p = 1.64e - 11$ ), moving readers from mostly difficult-to-medium retrieval to com-  
 488 fortably above medium and closer to “easy to find.” Together these results indicate that **STEER**’s  
 489 reports both better aligned with persona needs and easier to navigate.  
 490

## 491 4 DISCUSSION 492

493 **User Agent simulation** To understand how our pause policy translates into user-facing behavior,  
 494 we analyzed the *User Agent* used in offline evaluation (Appendix F.4). The *User Agent* maintains  
 495 very high precision across base-pause costs ( $> 0.97$ ), while recall declines as  $C_0$  increases, and  
 496 report alignment closely tracks *User Agent* recall ( $r \approx 0.81$ ). This indicates that pausing affects  
 497 outcomes primarily by changing how many promising directions are retained and developed, offer-  
 498 ing a controllable alignment-effort frontier via  $C_0$ . We view the *User Agent* as a diagnostic tool  
 499 for sweeping policies and stress-testing settings, but acknowledge that real users may be noisier and  
 500 value exploration differently; future work will calibrate the *User Agent* with human logs and run  
 501 counterfactual replays to quantify gaps between simulated and actual behavior.  
 502

503 **Persona modeling** We also examine how well **STEER**’s live persona tracks report quality. A  
 504 useful takeaway from Appendix F.3 is that **STEER** not only pauses effectively but also recovers and  
 505 maintains an accurate persona during a run. Even with only the first persona sentence as input, the  
 506 inferred persona’s alignment with the ground-truth aspect set strongly tracks final report alignment  
 507 ( $r \approx 0.85$ ,  $p < 10^{-3}$ ), indicating that the learned persona is informative rather than decorative. As  
 508  $C_0$  increases, pauses become fewer, the inferred persona is less specified, and downstream alignment  
 509 declines. In practice, persona-report agreement is a useful diagnostic for selecting  $C_0$ : choosing the  
 510 smallest  $C_0$  that achieves a target agreement while balancing the alignment-effort trade-off.  
 511

512 **Broader application** Beyond our experiments, **STEER** suggests a general pattern for long-  
 513 horizon, high-stakes tasks that must balance personalization with exploration under interpretable  
 514 control. For instance, scientific-discovery agents and research stacks could benefit from pausing  
 515 and live-persona steering to curb drift while preserving exploration (Team et al., 2025; Schmidgall  
 516 & Moor, 2025; Zheng et al., 2025a). Likewise, high-stakes domains such as financial advising and  
 517 trading (Zhang et al., 2024; Yu et al., 2024) and law and policy research (Li et al., 2024; Pipitone  
 518 & Alami, 2024) are natural application areas for **STEER**’s interpretable, user-steerable control. Be-  
 519 cause of **STEER**’s modularity, domains can add factors such as factuality, citation quality, or safety  
 520 alongside novelty and exploration. We view validating these extensions as promising future work.  
 521

## 522 5 CONCLUSION AND FUTURE WORK 523

524 We have presented **STEER**, proposing a new *interactive paradigm* for deep research. **STEER** cou-  
 525 ples a cost-benefit pause policy with interpretable controls, a live persona that adapts mid-process,  
 526 and diversity-novelty utility signals that keep exploration purposeful. Our experiments show that  
 527 **STEER** improves persona-tailored quality by 7.83%–22.80% over strong open-source and propri-  
 528 etary systems, leads on generic quality metrics, and is preferred by human readers in over 85% of  
 529 alignment and 83% of focus pairwise judgments. We also release a persona-query evaluation suite  
 and data pipeline to support reproducible testing and future model development.  
 530

531 Looking ahead, several directions appear especially promising. On the system side, exploring spec-  
 532 ulative pre-execution to reduce latency, a dynamic breadth-depth planner, and policy learning for  
 533 pause and branch selection could further strengthen real-time usability. On the evaluation side, end-  
 534 to-end user studies that judge the full interaction, measuring task success, time to insight, perceived  
 535 control and trust, and cognitive load, would provide a fuller picture of real-world value.  
 536  
 537  
 538  
 539

540 REFERENCES  
541

542 Salaheddin Alzubi, Creston Brooks, Purva Chiniya, Edoardo Contente, Chiara von Gerlach, Lucas  
543 Irwin, Yihan Jiang, Arda Kaz, Windsor Nguyen, Sewoong Oh, et al. Open deep search: Democ-  
544 ratizing search with open-source reasoning agents. *arXiv preprint arXiv:2503.20201*, 2025.

545 Chinmaya Andukuri, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah Goodman. Star-gate:  
546 Teaching language models to ask clarifying questions. In *First Conference on Language Model-  
547 ing*.

548 Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of machine  
549 learning research*, 3(Nov):397–422, 2002.

550 Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit  
551 problem. *Machine learning*, 47(2):235–256, 2002.

552 Zefan Cai, Baobao Chang, and Wenjuan Han. Human-in-the-loop through chain-of-thought. *arXiv  
553 preprint arXiv:2306.07932*, 2023.

554 Jaime Carbonell and Jade Goldstein. The use of mmr, diversity-based reranking for reordering  
555 documents and producing summaries. In *Proceedings of the 21st annual international ACM  
556 SIGIR conference on Research and development in information retrieval*, pp. 335–336, 1998.

557 Maximillian Chen, Ruoxi Sun, Tomas Pfister, and Sercan Ö. Arik. Learning to clarify: Multi-turn  
558 conversations with action-based contrastive self-training. In *The Thirteenth International Confer-  
559 ence on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net,  
560 2025. URL <https://openreview.net/forum?id=SIE6VFps9x>.

561 João Coelho, Jingjie Ning, Jingyuan He, Kangrui Mao, Abhijay Paladugu, Pranav Setlur, Jiahe Jin,  
562 Jamie Callan, João Magalhães, Bruno Martins, et al. Deepresearchgym: A free, transparent, and  
563 reproducible evaluation sandbox for deep research. *arXiv preprint arXiv:2505.19253*, 2025.

564 Mingxuan Du, Benfeng Xu, Chiwei Zhu, Xiaorui Wang, and Zhendong Mao. Deepresearch bench:  
565 A comprehensive benchmark for deep research agents. *arXiv preprint arXiv:2506.11763*, 2025.

566 Assaf Elovic. Gpt researcher: Llm-based autonomous agent for deep research. 2025. URL <https://github.com/assafelovic/gpt-researcher>.

567 Joseph L Fleiss. Measuring nominal scale agreement among many raters. *Psychological bulletin*,  
568 76(5):378, 1971.

569 Ivar Frisch and Mario Giulianelli. LLM agents in interaction: Measuring personality consistency and  
570 linguistic alignment in interacting populations of large language models. In Ameet Deshpande,  
571 EunJeong Hwang, Vishvak Murahari, Joon Sung Park, Diyi Yang, Ashish Sabharwal, Karthik  
572 Narasimhan, and Ashwin Kalyan (eds.), *Proceedings of the 1st Workshop on Personalization of  
573 Generative AI Systems (PERSONALIZE 2024)*, pp. 102–111, St. Julians, Malta, March 2024.  
574 Association for Computational Linguistics. URL <https://aclanthology.org/2024.personalize-1.9/>.

575 Google. Try deep research and our new experimental model in gemini, your  
576 ai assistant. 2024. URL <https://blog.google/products/gemini/google-gemini-deep-research/>.

577 Kilem Li Gwet. Computing inter-rater reliability and its variance in the presence of high agreement.  
578 *British Journal of Mathematical and Statistical Psychology*, 61(1):29–48, 2008.

579 Lisheng Huang, Yichen Liu, Jinhao Jiang, Rongxiang Zhang, Jiahao Yan, Junyi Li, and Wayne Xin  
580 Zhao. Manusearch: Democratizing deep search in large language models with a transparent and  
581 open multi-agent framework. *arXiv preprint arXiv:2505.18105*, 2025.

582 Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil’ucb: An optimal ex-  
583 ploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pp. 423–439.  
584 PMLR, 2014.

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

597

598 LangChain. Open deep research. 2025. URL [https://github.com/langchain-ai/open\\_deep\\_research?tab=readme-ov-file](https://github.com/langchain-ai/open_deep_research?tab=readme-ov-file).

599

600 Haitao Li, Junjie Chen, Jingli Yang, Qingyao Ai, Wei Jia, Youfeng Liu, Kai Lin, Yueyue Wu, Guozhi  
 601 Yuan, Yiran Hu, et al. Legalagentbench: Evaluating llm agents in legal domain. *arXiv preprint*  
 602 *arXiv:2412.17259*, 2024.

603

604 Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to  
 605 personalized news article recommendation. In *Proceedings of the 19th international conference*  
 606 *on World wide web*, pp. 661–670, 2010.

607

608 Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and  
 609 Zhicheng Dou. Search-o1: Agentic search-enhanced large reasoning models. *arXiv preprint*  
 610 *arXiv:2501.05366*, 2025a.

611

612 Zongqian Li, Ehsan Shareghi, and Nigel Collier. ReasonGraph: Visualization of reasoning methods  
 613 and extended inference paths. In Pushkar Mishra, Smaranda Muresan, and Tao Yu (eds.), *Pro-  
 614 ceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume  
 615 3: System Demonstrations)*, pp. 140–147, Vienna, Austria, July 2025b. Association for Com-  
 616 putational Linguistics. ISBN 979-8-89176-253-4. doi: 10.18653/v1/2025.acl-demo.14. URL  
 617 <https://aclanthology.org/2025.acl-demo.14/>.

618

619 Zongqian Li, Ehsan Shareghi, and Nigel Collier. Reasongraph: Visualisation of reasoning paths.  
 620 *arXiv preprint arXiv:2503.03979*, 2025c.

621

622 Rock Yuren Pang, KJ Feng, Shangbin Feng, Chu Li, Weijia Shi, Yulia Tsvetkov, Jeffrey Heer, and  
 623 Katharina Reinecke. Interactive reasoning: Visualizing and controlling chain-of-thought reason-  
 624 ing in large language models. *arXiv preprint arXiv:2506.23678*, 2025.

625

626 Nicholas Pipitone and Ghita Houir Alami. Legalbench-rag: A benchmark for retrieval-augmented  
 627 generation in the legal domain. *arXiv preprint arXiv:2408.10343*, 2024.

628

629 Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-  
 630 networks. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of*  
 631 *the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th Inter-  
 632 national Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3982–  
 633 3992, Hong Kong, China, November 2019. Association for Computational Linguistics. doi:  
 634 10.18653/v1/D19-1410. URL <https://aclanthology.org/D19-1410/>.

635

636 Allen Z Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Pen Xu,  
 637 Leila Takayama Takayama, Fei Xia, Jake Varley, et al. Robots that ask for help: Uncertainty align-  
 638 ment for large language model planners. In *Conference on Robot Learning (CoRL)*. Proceedings  
 639 of the Conference on Robot Learning (CoRL), 2023.

640

641 Corby Rosset, Ho-Lam Chung, Guanghui Qin, Ethan C Chau, Zhuo Feng, Ahmed Awadallah, Jen-  
 642 nifer Neville, and Nikhil Rao. Researchy questions: A dataset of multi-perspective, decomposi-  
 643 tional questions for llm web agents. *arXiv preprint arXiv:2402.17896*, 2024.

644

645 Alireza Salemi and Hamed Zamani. Lamp-qa: A benchmark for personalized long-form question  
 646 answering. *arXiv preprint arXiv:2506.00137*, 2025.

647

648 Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. LaMP: When large  
 649 language models meet personalization. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar  
 650 (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Lin-  
 651 guistics (Volume 1: Long Papers)*, pp. 7370–7392, Bangkok, Thailand, August 2024. Asso-  
 652 ciation for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.399. URL  
 653 <https://aclanthology.org/2024.acl-long.399/>.

648 Samuel Schmidgall and Michael Moor. Agentrxiv: Towards collaborative autonomous research.  
 649 *arXiv preprint arXiv:2503.18102*, 2025.  
 650

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

654 NovelSeek Team, Bo Zhang, Shiyang Feng, Xiangchao Yan, Jiakang Yuan, Zhiyin Yu, Xiaohan He,  
 655 Songtao Huang, Shaowei Hou, Zheng Nie, et al. Novelseek: When agent becomes the scientist–  
 656 building closed-loop system from hypothesis to verification. *arXiv preprint arXiv:2505.16938*,  
 657 2025.

658 Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and  
 659 Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions.  
 660 In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual  
 661 Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13484–  
 662 13508, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/  
 663 v1/2023.acl-long.754. URL <https://aclanthology.org/2023.acl-long.754/>.  
 664

665 Zhichao Wang, Bin Bi, Yanqi Luo, Sitaram Asur, and Claire Na Cheng. Diversity enhances an llm’s  
 666 performance in rag and long-context task. *arXiv preprint arXiv:2502.09017*, 2025.

667 Cheng-Kuang Wu, Zhi Rui Tam, Chao-Chung Wu, Chieh-Yen Lin, Hung-yi Lee, and Yun-Nung  
 668 Chen. I need help! evaluating LLM’s ability to ask for users’ support: A case study on  
 669 text-to-SQL generation. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Pro-  
 670 ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp.  
 671 2191–2199, Miami, Florida, USA, November 2024. Association for Computational Linguis-  
 672 tics. doi: 10.18653/v1/2024.emnlp-main.131. URL <https://aclanthology.org/2024.emnlp-main.131/>.  
 673

674 Shujin Wu, Yi R. Fung, Cheng Qian, Jeonghwan Kim, Dilek Hakkani-Tur, and Heng Ji. Align-  
 675 ing LLMs with individual preferences via interaction. In Owen Rambow, Leo Wanner, Mar-  
 676 rianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (eds.), *Pro-  
 677 ceedings of the 31st International Conference on Computational Linguistics*, pp. 7648–7662,  
 678 Abu Dhabi, UAE, January 2025. Association for Computational Linguistics. URL <https://aclanthology.org/2025.coling-main.511/>.  
 679

680 xAI. Grok 3. 2025. URL <https://x.ai/news/grok-3>.

681

682 Zhouhang Xie, Junda Wu, Yiran Shen, Yu Xia, Xintong Li, Aaron Chang, Ryan Rossi, Sachin Ku-  
 683 mar, Bodhisattwa Prasad Majumder, Jingbo Shang, et al. A survey on personalized and pluralistic  
 684 preference alignment in large language models. *arXiv preprint arXiv:2504.07070*, 2025.

685

686 Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yuechen Jiang, Yupeng Cao, Zhi Chen,  
 687 Jordan Suchow, Zhenyu Cui, Rong Liu, et al. Fincon: A synthesized llm multi-agent system with  
 688 conceptual verbal reinforcement for enhanced financial decision making. *Advances in Neural  
 689 Information Processing Systems*, 37:137010–137045, 2024.

690

691 Dingchu Zhang, Yida Zhao, Jialong Wu, Baixuan Li, Wenbiao Yin, Liwen Zhang, Yong Jiang,  
 692 Yufeng Li, Kewei Tu, Pengjun Xie, et al. Evolvesearch: An iterative self-evolving search agent.  
 693 *arXiv preprint arXiv:2505.22501*, 2025a.

694

695 Michael Zhang, W. Bradley Knox, and Eunsol Choi. Modeling future conversation turns  
 696 to teach llms to ask clarifying questions. In Y. Yue, A. Garg, N. Peng, F. Sha, and  
 697 R. Yu (eds.), *International Conference on Representation Learning*, volume 2025, pp. 60722–  
 698 60742, 2025b. URL [https://proceedings.iclr.cc/paper\\_files/paper/2025/file/97e2df4bb8b2f1913657344a693166a2-Paper-Conference.pdf](https://proceedings.iclr.cc/paper_files/paper/2025/file/97e2df4bb8b2f1913657344a693166a2-Paper-Conference.pdf).

699

700 Weizhi Zhang, Xinyang Zhang, Chenwei Zhang, Liangwei Yang, Jingbo Shang, Zhepei Wei,  
 701 Henry Peng Zou, Zijie Huang, Zhengyang Wang, Yifan Gao, et al. Personaagent: When large lan-  
 guage model agents meet personalization at test time. *arXiv preprint arXiv:2506.06254*, 2025c.

702 Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiaze Sun, Molei Qin, Xinyi Li, Yuqing  
 703 Zhao, Yilei Zhao, Xinyu Cai, et al. A multimodal foundation agent for financial trading: Tool-  
 704 augmented, diversified, and generalist. In *Proceedings of the 30th ACM SIGKDD conference on*  
 705 *knowledge discovery and data mining*, pp. 4314–4325, 2024.

706 Tianshi Zheng, Zheyen Deng, Hong Ting Tsang, Weiqi Wang, Jiaxin Bai, Zihao Wang, and Yangqiu  
 707 Song. From automation to autonomy: A survey on large language models in scientific discovery.  
 708 *arXiv preprint arXiv:2505.13259*, 2025a.

709 Yuxiang Zheng, Dayuan Fu, Xiangkun Hu, Xiaojie Cai, Lyumanshan Ye, Pengrui Lu, and Pengfei  
 710 Liu. Deepresearcher: Scaling deep research via reinforcement learning in real-world environ-  
 711 ments. *arXiv preprint arXiv:2504.03160*, 2025b.

## 714 A RELATED WORK

715 **Deep research** LLM-based research agents combine retrieval and multi-step reasoning to pro-  
 716 duce long-form answers (Coelho et al., 2025). Among the open-source frameworks, two dominant  
 717 paradigms are multi-agent pipelines that split planning, browsing, and reading across roles (Huang  
 718 et al., 2025; Alzubi et al., 2025; Li et al., 2025a; Zhang et al., 2025a) and RL-trained agents that  
 719 learn to search and reason (Zheng et al., 2025b; Jin et al., 2025; Song et al., 2025). On the eval-  
 720 uation side, benchmarks such as DeepResearchGym (Coelho et al., 2025) and DeepResearch Bench  
 721 (Du et al., 2025) have begun to standardize this setting, providing realistic research-style questions  
 722 and automated evaluation protocols for long-form, citation-heavy reports. Despite progress, most  
 723 systems still follow one-shot scoping with at most a single clarification, then a long autonomous run  
 724 and a monolithic report, offering little mid-process control when user needs evolve.

725 **Personalization and alignment** A growing line of work pursues personalization for LLM agents,  
 726 moving from static profile–personality conditioning and long-form checklists (Wu et al., 2025;  
 727 Salemi et al., 2024; Salemi & Zamani, 2025) toward interactive, test-time adaptation and multi-  
 728 stakeholder alignment (Xie et al., 2025). Recent trends probe persona behavior in interaction (e.g.,  
 729 consistency and drift under dialogue) (Frisch & Giulianelli, 2024) and build agent mechanisms that  
 730 adapt actions to user preferences at inference time (Zhang et al., 2025c). While this work establishes  
 731 that preferences should be updated during use, most approaches still lack *interpretable, end-to-end*  
 732 *controls* for deciding *when* to seek input and *how* to steer long-horizon generation as goals evolve.

733 **Interactive reasoning and control** Another closely related line of work equips LLMs with in-  
 734 teractive reasoning via clarification. Prior studies train models to ask when information is missing  
 735 (Andukuri et al.; Ren et al., 2023; Wu et al., 2024), model future turns to decide ask vs. answer  
 736 (Zhang et al., 2025b), and learn clarification policies with contrastive objectives (Chen et al., 2025).  
 737 Visualization tools improve transparency and user steering over chains of thought (Pang et al., 2025;  
 738 Li et al., 2025c). However, these efforts mostly address local interactions or static control, rather  
 739 than providing interpretable, end-to-end controls for *when* to pause, *what* to explore, and *how* to  
 740 adapt personalization mid-process in long-horizon research.

741 Table 1 positions **STEER** against both deep research frameworks and interactive/persona-aware rea-  
 742 soning systems, clarifying which capabilities each class actually offers. Open-source deep research  
 743 frameworks and proprietary web-based services are all built around long autonomous run paradigm:  
 744 once the user issues a query, a fixed pipeline executes to completion with no exposed, interpretable  
 745 control over where in the research tree to intervene or when to ask for guidance. The  marks in the  
 746 mid-process steering column indicate the limited behavior: these systems sometimes allow a single  
 747 upfront scoping or plan-confirmation step before the full autonomous run, but provide no further  
 748 steering within the trajectory; beyond that point they are fully autonomous. In addition, for propri-  
 749 etary systems, public documentation suggests that fixed personas or user memory may be used, but  
 750 there is no evidence of live persona modeling, and their control policies and internal mechanisms  
 751 are not accessible or benchmarkable; we therefore restrict ourselves to conceptual comparison.

752 Meanwhile, interactive/persona-aware frameworks occupy the complementary side of the space:  
 753 PersonaAgent (Zhang et al., 2025c) maintains a live user representation and adapts over time but has  
 754 no steering capability, while ReasonGraph (Li et al., 2025b) and HITL CoT MCS (Cai et al., 2023)

756 expose mid-process reasoning for visualization, inspection, or human correction but implement no  
 757 adaptive pause mechanism and do not perform live persona modeling. Taken together, the table  
 758 shows that existing systems provide at most one of the three capabilities we target – mid-process  
 759 steering, adaptive pause decisions, or live persona modeling – and never all three in a benchmarkable  
 760 deep research setting. **STEER** is the only system that supports all three simultaneously, coupling a  
 761 cost–benefit pause policy with a live persona that conditions planning, branch utility, and synthesis  
 762 within a single research tree.

763

## 764 B USE OF LLMs FOR WRITING ASSISTANCE

765

766 We used ChatGPT-4o *only* for language-level editing. Concretely:

767

- 768 • Polishing prose, tightening sentences, fixing grammar and LaTeX wording, reordering or shortening  
 769 paragraphs, and suggesting alternative titles or section headers.
- 770 • No ideas, methods, claims, proofs, experiments, numbers, figures, tables, code, prompts, or citations  
 771 were produced by the model. All technical content, analyses, and results were authored and  
 772 verified by the authors.
- 773 • We supplied already written passages or outlines and requested editing (for example, “polish word-  
 774 ing, keep all technical details unchanged”).
- 775 • The model was not given proprietary data, code, or unpublished results beyond the text to be  
 776 edited. All outputs were reviewed by the authors for accuracy and tone.

777

## 778 C DIVERSIFIED SUBSET SELECTION

779

780 For completeness, we include the pseudocode of the greedy MMR selection used in our framework.  
 781 Given a candidate set of follow-up questions  $\mathcal{C} = \{q_1, \dots, q_M\}$  with confidence scores  $\text{conf}(q_i)$   
 782 and embeddings  $\mathbf{e}_i$ , the algorithm selects a diversified subset  $\mathcal{C}'$  of size  $K$ :

783

---

784 **Require:** Candidate list  $\mathcal{C} = \{q_1, \dots, q_M\}$  with confidences  $\text{conf}(q_i)$ , embeddings  $\mathbf{e}_i$ ; desired  
 785 subset size  $K$

786 **Ensure:** Diversified subset  $\mathcal{C}'$

787 1:  $\mathcal{C} \leftarrow \text{sort } \mathcal{C} \text{ in non-increasing order of } \text{conf}(q_i)$

788 2:  $\mathcal{C}' \leftarrow \emptyset, I_{\mathcal{C}'} \leftarrow \emptyset$

789 3: **while**  $|\mathcal{C}'| < K$  **do**

790 4:      $\mathcal{C} \leftarrow \{i \mid i \notin I_{\mathcal{C}'}\}$

791 5:     **if**  $I_{\mathcal{C}'} = \emptyset$  **then**

792 6:          $i^* \leftarrow \min \mathcal{C}$  ▷ top-confidence question

793 7:     **else**

794 8:         **for**  $i \in \mathcal{C}$  **do**

795 9:              $d_i \leftarrow \max_{j \in I_{\mathcal{C}'}} \text{sim}(\mathbf{e}_i, \mathbf{e}_j) + \varepsilon$

796 10:         **end for**

797 11:          $i^* \leftarrow \arg \min_{i \in \mathcal{C}} d_i$  ▷ least similar to current set (MMR criterion)

798 12:         **end if**

799 13:          $\mathcal{C}' \leftarrow \mathcal{C}' \cup \{q_{i^*}\}, I_{\mathcal{C}'} \leftarrow I_{\mathcal{C}'} \cup \{i^*\}$

800 14: **end while**

801 15: **return**  $\mathcal{C}'$

---

802

803

## 804 D DETAILS FOR GAIN OF PAUSING IMPLEMENTATION

805

806 **Alignment gain** Let  $r(n)$  denote the chunk report at node  $n$ , formed by concatenating the learn-  
 807 ings  $\{\ell_i\}_{i=1}^m$  (if there are  $m$  learnings at the node), and let  $\hat{A}_n$  be the inferred aspect set at that node.  
 808 For the  $k$ -th child node of a frontier node  $n_k^*$ ,

809

$$\Delta \text{Align}(n_k^*) = \text{Align}(r(n_k^*), \hat{A}_{n_k^*}) - \text{Align}(r(n^*), \hat{A}_{n^*}).$$

810    **Exploration bonus** For each chunk report, we prompt an LLM to assign short tags (see Appendix K, *Search Result Processing* prompt). We maintain the global tag set  $\mathcal{T}$  and a cumulative usage count  $\text{count}(T)$  for each tag  $T$  up to the current step. With a small constant  $\epsilon > 0$ , the exploration bonus is

$$814 \quad \text{Explore}(n_k^*) = \frac{1}{|\mathcal{T}|} \sum_{T \in \mathcal{T}} \frac{\epsilon}{1 + \sqrt{\text{count}(T)}}.$$

816    This UCB-style term grants larger bonus to under-tried tags and decays as a tag is reused.

818    **Information gain** To reward novelty relative to what has already been learned, we compare a candidate’s node embedding to the centroid of accumulated learnings. Let  $\mathbf{e}_\ell$  be the embedding 819    of a learning  $\ell$ . For node  $n$  with number of learnings  $L(n) = \{\ell_i\}_{i=1}^{m(n)}$ , define its embedding 820     $\mathbf{e}_n = \frac{1}{m(n)} \sum_{i=1}^{m(n)} \mathbf{e}_{\ell_i}$  (when  $m(n) > 0$ ). Let  $\mathcal{L}$  be the set of all learnings gathered so far,  $M = |\mathcal{L}|$ , 821    and  $\mu = \frac{1}{M} \sum_{\ell \in \mathcal{L}} \mathbf{e}_\ell$ . Then

$$825 \quad \text{InfoGain}(n_k^*) = \begin{cases} 1 - \text{sim}(\mathbf{e}_{n_k^*}, \mu), & m(n_k^*) > 0 \text{ and } M > 0, \\ 0, & m(n_k^*) = 0, \\ 1, & \text{otherwise.} \end{cases}$$

828    **Execution cost** Let  $D$  be the max depth,  $d(n)$  the depth of node  $n$ , and  $K$  the branching factor. 829    For child  $n_k^*$ , the remaining depth is  $d_{\text{rem}} = D - d(n_k^*)$ . The number of nodes in a saturated  $K$ -ary 830    subtree is

$$832 \quad N_{\text{rem}} = \begin{cases} \frac{K^{d_{\text{rem}}+1} - 1}{K - 1}, & K > 1, \\ d_{\text{rem}} + 1, & K = 1. \end{cases}$$

835    With a running average token cost  $\text{Tok}_{\text{avg}}$  per node, the estimated tokens are  $T_k^{\text{est}} = \text{Tok}_{\text{avg}} N_{\text{rem}}$ , 836    and the normalized execution cost is

$$837 \quad C^{\text{exec}}(n_k^*) = \frac{T_k^{\text{est}}}{T_k^{\text{est}} + \text{Tok}_{\text{avg}}} = \frac{N_{\text{rem}}}{N_{\text{rem}} + 1}.$$

840    **Filtering candidates when pausing** Let  $U_k = U(n_k^*)$  and  $\text{conf}_k \in [0, 1]$  be a confidence score 841    generated by the LLM (see Appendix K, *Search Result Processing* prompt). Define the uncertainty 842    radius

$$843 \quad r_k = (1 - \text{conf}_k) \left( \max_{i \in K} U_i - \min_{i \in K} U_i \right), \quad U_k^{\text{upper}} = U_k + r_k, \quad U_k^{\text{lower}} = U_k - r_k.$$

845    The *could-be-the-best* set is

$$846 \quad S = \{k \mid U_k^{\text{upper}} \geq \max_{i \in K} U_i^{\text{lower}}\}.$$

848    This mirrors upper and lower confidence bounds for best-arm filtering Jamieson et al. (2014).

## 850    E DATA CONSTRUCTION DETAILS

852    To evaluate our method, we need a dataset with  
853    deep research worthy questions paired with realistic  
854    personas, where personas are, as defined in Section 2.1,  $(p_{\text{text}}, \mathcal{A})$ , where  $p_{\text{text}}$  is a string, combining  
855    the user’s profile and personality, and  $\mathcal{A}$  is a set of  
856    *aspects* that the user is interested to see in a high-  
857    quality, well-aligned final report. We construct our  
858    dataset on top of the subset of 1,000 queries from  
859    Researchy Questions dataset Rosset et al. (2024)  
860    used in DeepResearchGym Coelho et al. (2025).

862    For each query, we first generate one or more  $(p_{\text{text}}$   
863    that would be reasonable to ask the query. For this, we adopt a two-step approach. In the first step,  
inspired by Wu et al. (2025) and Wang et al. (2023), we use an iterative self-generation and filtering

Data Split →	All	Eval
Total Queries	1000	200
Total Query-Persona Pairs	1381	286
Queries with 1 Persona	646 <sub>(64.6%)</sub>	125 <sub>(62.5%)</sub>
Queries with 2 Personas	327 <sub>(32.7%)</sub>	64 <sub>(32.0%)</sub>
Queries with 3 Personas	27 <sub>(2.7%)</sub>	11 <sub>(5.5%)</sub>

Table 5: Data Statistics

864 pipeline. In each round, 3 profiles are randomly selected from the profiles in the ALOE dataset  
 865 Wu et al. (2025) and used as input to an off-the-shelf LLM (GPT-4o) to generate 3 new profiles  
 866 that would be reasonable to ask the query per iteration. Then we introduce an automatic filtering  
 867 process based on semantic similarity to ensure the distinctiveness and diversity of the generated  
 868 profiles. Same as Wu et al. (2025), we use Sentence Transformers Reimers & Gurevych (2019) to  
 869 compute embedding of the generated profiles and measure the cosine similarity among the generated  
 870 new profiles. For each new profile, if the highest similarity score compared to the other profiles  
 871 exceeds 0.65, the profile is considered too similar to at least one of the other profiles and discarded.  
 872 Otherwise, it will be accepted as a successful new profile to pair with the query. We repeat the  
 873 process until 3 new accepted profiles are generated. In step 2, for each accepted profile, we generate  
 874 a reasonable personality with GPT-4o to pair with it. For this, we randomly sample personalities  
 875 from personality pool of the ALOE dataset as sample personalities fed into the LLM for generation.  
 876

877 Once we have generated one or more  $p_{text}$  for each query, we then generate the set of aspects  $\mathcal{A}$   
 878 for each  $p_{text}$ . We adopt the same approach as in Salemi & Zamani (2025) to generate 5-8 specific  
 879 aspects that a user (described by  $p_{text}$ ) would expect to see in a comprehensive and helpful report to  
 the query, along with an evidence and a reasoning for each aspect, attributed from  $p_{text}$ .

880 Table 5 details the statistics of our generated dataset. All prompts for persona/profile/aspect genera-  
 881 tion are provided in Appendix L.

## 883 F ADDITIONAL EXPERIMENT DETAILS

886 System	Focus <sub>st</sub> $\uparrow$	Relevance <sub>sup</sub> $\uparrow$	Relevance <sub>con</sub> $\downarrow$
GPT-Researcher	67.07	61.39	1.02
GPT-Researcher <sub>initial-persona</sub>	69.18	60.82	1.11
GPT-Researcher <sub>full-persona</sub>	70.78	59.94	1.04
OpenDeepResearch	73.15	60.36	<b>0.69</b>
OpenDeepResearch <sub>initial-persona</sub>	75.61	57.07	<u>0.81</u>
OpenDeepResearch <sub>full-persona</sub>	78.84	57.06	<u>0.81</u>
o4-mini-deep-research <sub>initial-persona</sub>	78.41	<b>67.36</b>	1.74
o4-mini-deep-research <sub>full-persona</sub>	80.60	<u>66.45</u>	1.94
<b>STEER</b> <sub>[<math>C_0=0.7</math>]</sub>	78.51	60.47	1.13
<b>STEER</b> <sub>[<math>C_0=0.1</math>]</sub>	<b>80.67</b>	60.19	1.10

895 Table 6: Performance comparison between **STEER** and baseline frameworks on sentence-level fo-  
 896 cusing score and report relevance scores.

### 899 F.1 ADDITIONAL METRICS

901 Table 6 reports sentence-level focus and DeepResearchGym relevance (support  $\uparrow$  and contradiction  
 902  $\downarrow$ ). We do not use sentence-level focus as a primary metric because it is length sensitive: the score  
 903 Focus<sub>st</sub> is the fraction of sentences mapped to any aspect, so longer reports with a few connective or  
 904 background sentences are penalized, whereas terse styles can inflate the ratio. Still, **STEER** achieves  
 905 competitive values (e.g., 80.67 at  $C_0=0.1$ ), on par with the proprietary model and higher than the  
 906 open-source baselines, indicating that personalization does not come at the cost of sentence-level  
 907 topicality.

908 DeepResearchGym relevance compares a report to a pre-extracted, task-generic keypoint list; be-  
 909 cause **STEER** steers into personalized directions, it is expected to score lower on relevance<sub>sup</sub> than  
 910 a system optimized for the generic keypoints (e.g., o4-mini-deep-research), while main-  
 911 taining moderate relevance<sub>con</sub>. In our results, **STEER**'s relevance<sub>sup</sub> is similar to GPT-Researcher  
 912 and OpenDeepResearch, with contradiction around 1.10–1.13; the proprietary model attains higher  
 913 support but also substantially higher contradiction, whereas OpenDeepResearch shows low contra-  
 914 diction but lower support. Taken together, these metrics are complementary diagnostics: sentence-  
 915 level focus confirms topicality at the sentence granularity, and DeepResearchGym relevance reflects  
 916 overlap with generic keypoints rather than user-specific goals.

### 917 F.2 BASE PAUSE COST VS. PAUSE BEHAVIOR

To better understand system behavior, Figure 6 shows the distribution of number of pauses per run. As expected, lowering base pause cost increases the number of pauses, with median pauses dropping from around 10 ( $C_0 = 0.0$ ) to fewer than 2 ( $C_0 \geq 0.8$ ). Compared to an LLM-based Pause-Agent baseline, which issues many more questions, **STEER**'s cost-sensitive mechanism achieves tighter control over the frequency of interruptions. This suggests that base pause cost provides a direct and interpretable knob for regulating user burden.

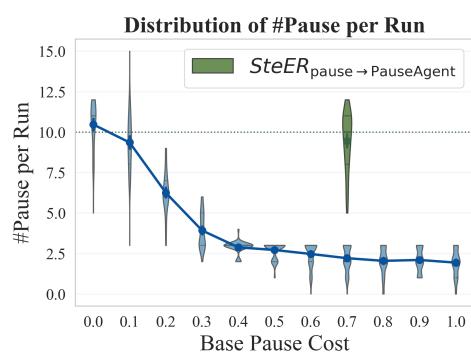


Figure 6: Distribution of number of pauses per run across base pause values.

### F.3 STEER'S PERSONA MODELING ANALYSIS

To assess the effectiveness of **STEER**'s dynamic persona modeling, we examine how well the inferred persona aligns with the system's final report over the course of interaction. Specifically, we track the alignment score between the generated report and the inferred persona at different base pause cost ( $C_0$ ) settings, alongside the alignment between the report and the ground-truth persona provided at the start.

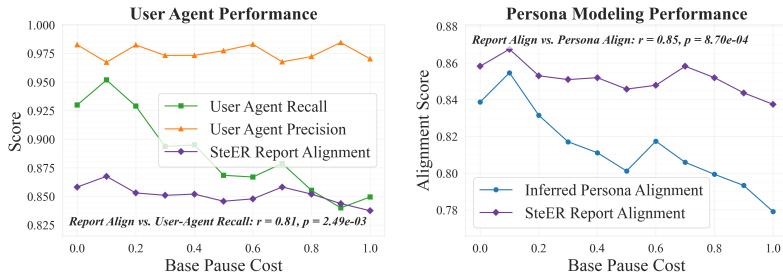


Figure 7: Analysis of User Agent and Persona Modeling Performance across Base Pause Cost ( $C_0$ ). *Left:* User Agent precision, recall, and **STEER** report alignment scores plotted across varying base pause cost values. *Right:* Alignment scores of **STEER**'s inferred persona and final report, both evaluated against the ground-truth aspect set  $\mathcal{A}$ , plotted across varying base pause cost values.

**Report Alignment Tracks Persona Alignment** As shown in the right panel of Figure 7, there is a strong positive correlation between **STEER**'s report alignment and the alignment of its inferred persona to the ground-truth aspect set. The Pearson correlation is  $r = 0.85$  ( $p = 8.7 \times 10^{-4}$ ), indicating that improvements in inferred persona accuracy are tightly coupled with improvements in report quality. This supports the intuition that **STEER**'s performance stems not only from architectural advances like mid-process pausing, but also from its ability to incrementally build an accurate model of user goals.

**Impact of Base Pause Cost** We observe a general downward trend in both inferred persona alignment and report alignment as base pause cost increases. This confirms that higher interruption costs reduce the frequency of clarifying interactions, resulting in less accurate persona estimates and, consequently, less aligned outputs. In contrast, low  $C_0$  values allow **STEER** to query the user more frequently, leading to refined persona inference and stronger downstream alignment.

These results highlight the central role of interactive refinement in personalized research workflows. Rather than relying solely on upfront persona injection, **STEER** learns about the user incrementally – and this process is empirically shown to improve alignment. The correlation between inferred and actual persona alignment validates the design of our live persona model and its integration into the decision-making process.

972 F.4 USER AGENT PERFORMANCE ANALYSIS  
973

974 To enable scalable, automatic evaluation of **STEER**, we employ a User Agent that simulates a real  
975 user interacting with the system. This User Agent is responsible for selecting preferred research  
976 directions based on a target persona and proposing new follow-up questions when relevant aspects  
977 remain uncovered. Its effectiveness directly impacts the utility of our offline evaluation framework.

978 As shown in the left panel of Figure 7, the User Agent maintains consistently high precision across  
979 a wide range of  $C_0$  values, with scores above 0.97. This suggests that when the agent chooses to  
980 retain a direction, it is highly likely to align with the user’s intended aspects. In contrast, recall is  
981 more sensitive to the pausing configuration. At lower  $C_0$  (e.g., 0.1), the User Agent achieves peak  
982 recall near 0.95, but recall steadily declines as  $C_0$  increases, falling to approximately 0.85 by  $C_0 =$   
983 1.0. This reflects the agent’s conservative behavior under higher interruption costs, where it refrains  
984 from selecting additional directions that could be beneficial.

985 We also observe that the alignment score of the final report generated by **STEER** (in purple) closely  
986 tracks the recall curve of the User Agent. The Pearson correlation between the two is strong and  
987 statistically significant ( $r = 0.81, p = 2.49 \times 10^{-3}$ ), as annotated in the plot. This indicates that  
988 the breadth of information the agent retains during interaction is highly predictive of the alignment  
989 quality of the final report. The stronger the agent’s coverage of relevant aspects (recall), the more  
990 aligned the report tends to be with the user’s needs.

991 These results confirm that the simulated User Agent is not only a faithful proxy for real user be-  
992 havior but also a critical driver of **STEER**’s alignment performance. Its high precision ensures  
993 quality, while its recall effectively governs how much of the user’s goals are ultimately realized in  
994 the research output.

995  
996 G USER STUDY DETAILS  
997

998 To complement automated evaluation, we conducted a human annotation study to directly assess  
999 how well **STEER** reports align with user personas compared to baseline systems. We developed a  
1000 custom web-based annotation platform (Figure 8) that guides annotators through a structured eval-  
1001 uation procedure with clear instructions and embedded report viewers.

**Step 1: Aspect Coverage Evaluation**

Step 1 asks the annotator to rate how well each report covers the persona's aspects. The persona's aspects are listed as: Impact of Geopolitical Events, Supply and Demand Analysis, Environmental Impact, Implications for Transportation Costs, Global Economic Considerations, Physical and Behavioral Insights, Future Projections and Trends, and Risk profile.

**Report A** and **Report B** are shown side-by-side. Report A is titled "Research Report: Understanding the Rise in Oil Prices" and Report B is titled "Why Are Oil Prices Increasing?". Both reports contain sections on oil price increases, geopolitical events, and economic factors.

**Step 2: Information Findability**

Step 2 asks the annotator to rate how easy it was to find answers to aspects in Report A and Report B. The aspects are the same as in Step 1. The rating scale is 0 (Difficult), 1 (Medium), and 2 (Easy).

**Step 3: Compare Reports**

Step 3 compares Report A and Report B based on the annotator's ratings. The table shows the average rating for each aspect across both reports.

Aspect	Report A	Report B
Impact of Geopolitical Events	0.9 0.1 0.1	0.9 0.1 0.2
Supply and Demand Analysis	0.9 0.1 0.1	0.9 0.1 0.2
Environmental Impact	0.9 0.1 0.1	0.9 0.1 0.2
Implications for Transportation Costs	0.9 0.1 0.2	0.9 0.1 0.2
Global Economic Considerations	0.9 0.1 0.2	0.9 0.1 0.2
Physical and Behavioral Insights	0.9 0.1 0.2	0.9 0.1 0.2
Future Projections and Trends	0.9 0.1 0.2	0.9 0.1 0.2

1024 Figure 8: User study interface.  
1025

1026  
1027

## G.1 SETUP

1028  
1029  
1030  
1031  
1032  
1033

Annotators were provided with a **persona card** containing (i) the query, (ii) a short persona description, and (iii) the persona’s **interested aspects**—the specific information needs that the final report should cover. These interested aspects formed the primary basis of evaluation. Annotators then evaluated two reports for the same query–persona pair: one generated by **STEER** and one by a baseline system (either GPT-Researcher or Open Deep Research). Report order was randomized to reduce bias.

1034

1035  
1036

## G.2 EVALUATION PROCEDURE

1037  
1038  
1039  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
1049

- **Step 1: Aspect Coverage.** Annotators skimmed both reports and rated, for each aspect, how well the report addressed it on a 3-point scale: **0 = not covered**, **1 = somewhat covered**, **2 = fully covered**. When assigning a score of 1 or 2, annotators were instructed to copy-paste a short supporting quote (1–2 sentences) from the report to ground their judgment. This ensured ratings were evidence-backed rather than impressionistic.
- **Step 2: Findability.** Annotators rated how easy it was to locate content relevant to each aspect in the report on a 3-point scale: **0 = difficult**, **1 = medium**, **2 = easy**. This step captured not only whether the aspect was present, but also whether it was readily discoverable by a reader.
- **Step 3: Report Comparison.** Based on their coverage and findability assessments, annotators selected a winner between the two reports along two dimensions: **Alignment** (which report better served the persona’s aspects) and **Focus** (which report stayed more on-topic versus digressing into irrelevant content).

1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057

**Interface Design.** The interface (Figure 8) displayed both reports side by side in embedded PDF viewers, alongside the persona’s aspects in a draggable panel for quick reference. Each evaluation step was clearly separated into dedicated panels, with concise instructions and tips (e.g., “**You don’t need to read every word**—scan section titles and opening sentences for relevant content”). Progress indicators guided annotators through the sequence, ensuring consistency. Importantly, the platform emphasized that judgments should be made **from the persona’s perspective**, not based on annotators’ personal preferences.

1058  
1059  
1060

**Instructions and Quality Control.** The study followed a three-step protocol:

As displayed in Figure 9, annotators were instructed to:

1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069

1. Read the persona aspects carefully, treating them as the ground truth for evaluation.
2. Provide evidence quotes for all non-zero aspect coverage ratings.
3. Complete all steps in sequence (coverage → findability → comparison).
4. Judge strictly by persona relevance, not by report verbosity, formatting, or personal opinion.

These safeguards helped ensure high-quality, reproducible annotations grounded in persona-aligned judgments.

1070  
1071  
1072

## H STEER WORKING PROTOTYPE

1073  
1074  
1075  
1076  
1077  
1078  
1079

To illustrate the functionality of **STEER**, we build an interactive web-based prototype (Figure 10) that visualizes the **STEER** framework in action. The interface consists of three synchronized panels: (i) a conversation pane for clarification prompts and user feedback, (ii) a dynamically expanding research tree that reflects the research status and partial research results, and (iii) a live persona tracker that displays the evolving inferred persona  $\hat{P}$  and monitors the updating alignment between cumulative research results and the inferred user aspects  $\hat{A}$ . This prototype supports interactive research sessions, allowing users to guide the exploration by selecting preferred subtopics or introducing new follow-up questions mid-process.

1080

**What You'll Do in Each Item**

1081

**Step 0 – Read the persona card**

1082

You'll see:

1083

The **query**.

1084

A short **persona description**.

1085

The persona's **interested aspects** — the specific key information they expect in a high-quality report for this query.

1086

1087

**Tip:** The **interested aspects** are the most important and concise part. You don't need to read every word of the full persona text, but **do** read the aspects list carefully.

1088

1089

1090

1091

**Step 1 – Rate aspect coverage for each report**

1092

Skim both full reports (embedded PDFs) and for each **aspect**, rate how well the report covers it:

1093

**0** = not covered

1094

**1** = somewhat covered

1095

**2** = fully covered

1096

If you choose **1** or **2**, copy-paste a short quote from that report (1–2 sentences or a short paragraph) that supports your rating.

1097

1098

1099

**Tip:** You don't need to read every word. To find evidence quickly, you can:

1100

1. Scan headings for relevant sections.

1101

2. If you'd like, open the report in a new tab and use **Ctrl/Cmd-F** to search keywords from the aspect.

1102

1103

1104

1105

1106

1107

1108

**Step 2 – Findability**

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

**Step 3 – Compare two reports (A vs B)**

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

**Tip:** You don't need to read every word. A quick scan of section titles plus the first 1–2 sentences of each section is usually enough to judge. Some reports have a table of contents that can help you navigate, but **don't judge a report only by whether it has one**.**I LLM-AS-JUDGE EVALUATION**

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

To validate the effectiveness of the LLM judge used throughout evaluation, we conduct a small-scale meta-evaluation of the LLM-as-judge. Specifically, we take the alignment score per aspect produced by the LLM judge (*gpt-4.1-mini*) and the Coverage score per aspect produced by human annotators in the user study (both in the scale of 0 - 2), and compute the Pearson correlation  $r$  between LLM-assigned and human-assigned scores over all overlapping aspects where both annotations are available. We obtain  $r = 0.34$ ,  $p$ -value  $< 0.0001$ , indicating a statistically significant, moderate positive correlation between the LLM-as-judge and human evaluations. At the annotator

Figure 9: User study instructions.

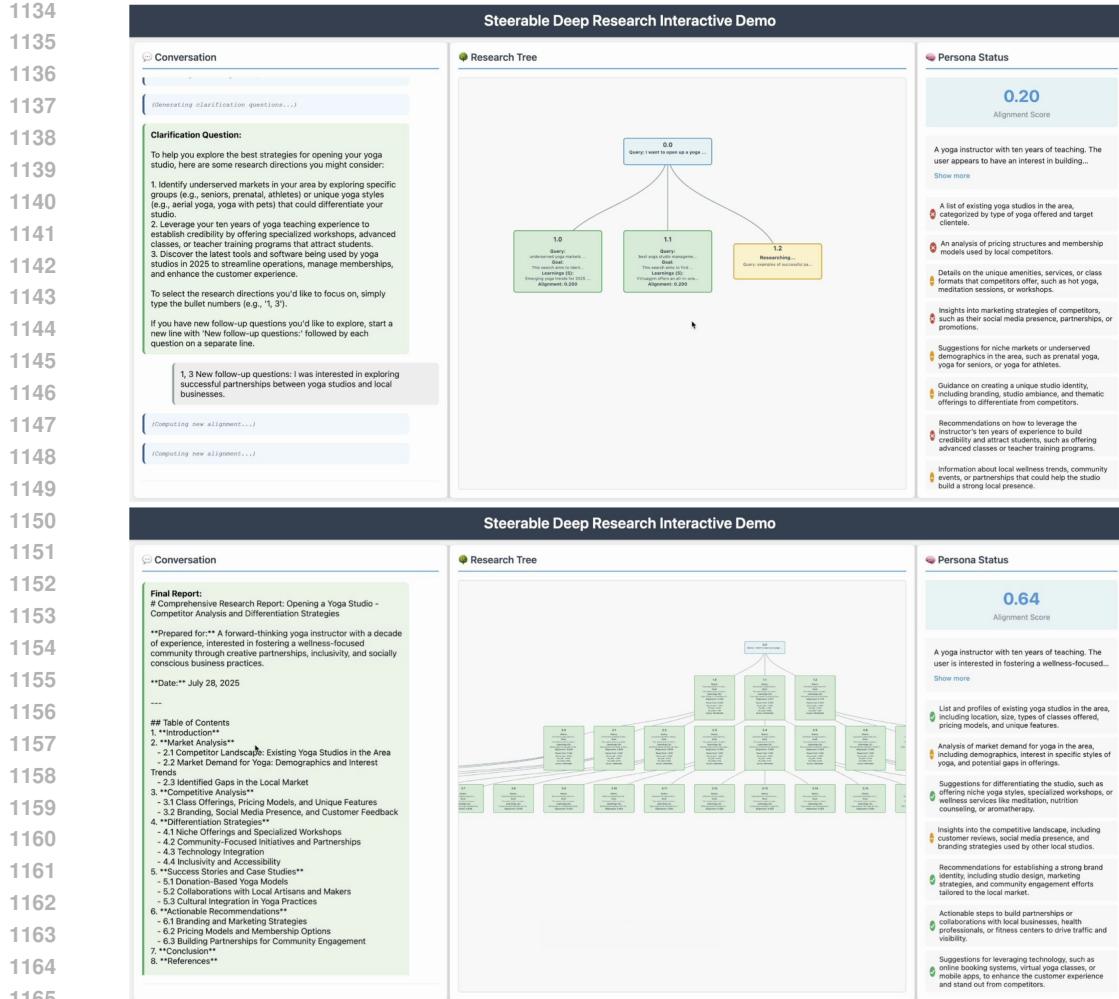


Figure 10: Interface of STEER web application.

level, all annotators exhibit positive correlations with the LLM judge, ranging from 0.19 to 0.45 with small variability (standard deviation = 0.094), suggesting that the *gpt-4.1-mini*'s scoring is consistently aligned with different human raters rather than being driven by any single annotator. While imperfect, these results indicate that the LLM-judge is directionally consistent with human judgments at the aspect level and is suitable as a scalable proxy for our large-scale evaluation, especially when interpreted alongside the human study that directly validates our main claims.

## J LIMITATIONS AND FUTURE WORK

**Simulated personas and queries vs. users' own queries.** Our user study asks annotators to *mimic* a target persona and query rather than using their own information needs. This is intentional: our goal is to evaluate system-level behavior under controlled, persona-conditioned information needs and to compare STEER against baselines on *exactly* the same persona–query pairs. This requires (*i*) a fixed set of personas and queries that all systems answer and (*ii*) a shared reference for annotators, so that cross-system differences in alignment and coverage can be attributed to the system rather than to heterogeneous user goals. If each participant were to choose their own query and implicit persona, different systems would be evaluated on tasks of varying difficulty, domain, and specificity, making it difficult to perform clean system-level comparisons and to interpret differences in outcomes as stemming from the model rather than from the task. Moreover, obtaining enough repeated measurements per system–persona–task condition under fully free-form user queries would require

1188 a substantially larger number of participants and interactions, constituting a much larger user-study  
 1189 effort in terms of human resources, annotation time, and cost. We acknowledge that this controlled  
 1190 design does not fully capture all aspects of real-world usage, particularly long-term adaptation to  
 1191 an individual's genuine information needs. As outlined in our limitations and future work, a natural  
 1192 next step is to conduct more comprehensive, end-to-end user studies in which participants bring their  
 1193 own personas and tasks, interact with the system over longer sessions, and are evaluated on richer  
 1194 metrics such as task success, time to insight, perceived control and trust, and cognitive load.  
 1195

1196 **Runtime and latency analysis.** A key limitation of this work is that we do not provide a sys-  
 1197 tematic runtime or latency analysis of **STEER** relative to baseline frameworks. Because **STEER** is  
 1198 explicitly designed as an interactive deep-research system, defining and comparing “runtime” in a  
 1199 meaningful way is non-trivial: a naive wall-clock measure would conflate *(i)* time spent on user inter-  
 1200 action (or User Agent responses in our simulations) and *(ii)* the highly variable size of the research  
 1201 tree induced by different pause policies and user choices. Simply averaging end-to-end runtimes  
 1202 could therefore be misleading – **STEER** may spend more time within a single run while reducing  
 1203 wasted effort by correcting misalignment earlier, whereas non-steerable baselines may require stack-  
 1204 ing multiple full deep-research calls to reach a comparable result. Moreover, even in a controlled  
 1205 environment like DeepResearchGym, end-to-end latency is heavily influenced by external factors  
 1206 such as search endpoint stability, network conditions, and rate limits, especially in multi-agent, tool-  
 1207 using pipelines. As a result, raw wall-clock differences are hard to attribute cleanly to the pause  
 1208 policy rather than infrastructure noise. We view a careful efficiency study – using user-centric mea-  
 1209 sures such as time- or turns-to-satisfactory-answer, and stratifying by tree size, number of pauses,  
 1210 and backend conditions – as an important direction for future work, and plan to complement our  
 1211 current alignment and quality evaluations with such analyses in follow-up studies.  
 1212

## K PROMPT TEMPLATES IN STEER

1214 We include here the core prompt templates used in our **STEER** framework, organized by function-  
 1215 ality. Each block shows the **system prompt** and the corresponding **user prompt**. Placeholders such  
 1216 as `{query}`, `{persona_text}`, and `{checklist_items}` are substituted at runtime.  
 1217

### Research Planning

#### System Prompt

1220 You are an expert researcher working with a specific user persona.  
 1221 Your task is to analyze the original query and search results, then  
 1222 generate targeted questions that explore different directions and time  
 1223 periods of the topic, specifically tailored to the user's interests  
 1224 and checklist items.  
 1225

#### User Prompt

1226 Original query: `{query}`  
 1227

1228 Current time: `{current_time}`  
 1229

1230 User persona: `{persona_text}`  
 1231

1232 User checklist (aspects they care about):  
 1233 `{checklist_items}`  
 1234

1235 Search results:  
 1236 `{search_results}`  
 1237

1238 Based on these results, the original query, the user's persona, and  
 1239 their checklist, generate 5-8 unique follow-up questions that:  
 1240

1. Explore different directions relevant to this query
2. Cover a good wide range of topics and aspects of the query
3. Consider recent developments up to `{current_time}`
4. Are somewhat tailored to the user's background and needs, but not  
 1241 constrained by the user's persona and interests

```

1242
1243     5. Each follow-up question should cover a distinct thematic facet - do
1244     not repeat other questions
1245
1246     For each question, provide a confidence score between 0.0 and 1.0
1247     indicating:
1248     - Relevance of the question to the main research query
1249     - Insightfulness of the question that would be useful for the final
1250     report generation
1251     - How likely this question is to lead to valuable information for this
1252     user
1253
1254     Return your response as a JSON object with the following structure:
1255     {
1256         "follow_up_questions": [
1257             {
1258                 "follow_up_question": "follow-up question text",
1259                 "confidence": 0.0-1.0,
1260                 "reasoning": "why this is a good follow-up question"
1261             }
1262         ]
1263     }

```

## 1262 Search Result Processing

### 1263 System Prompt

1264 You are an expert researcher analyzing search results for a specific  
1265 user persona. Focus on extracting learnings and follow-up questions  
1266 that are most relevant to the user's interests and checklist items.

### 1267 User Prompt

1268 Given the following research results for the query '{query}', extract  
1269 key learnings and suggest 5-8 follow-up questions that are  
1270 specifically relevant to the user's persona and interests.

1271 User persona: {persona\_text}

1273 User checklist (aspects they care about):  
1274 {checklist\_items}

1276 Previously seen tags: {seen\_tags}

1277 Focus on:

- 1278 1. Learnings that address the user's checklist items
- 1279 2. Information relevant to their background and interests
- 1280 3. Follow-up questions that would help address their specific needs
- 1281 4. Each follow-up question should cover a distinct thematic facet - do  
1282 not repeat other questions

1283 For each follow-up question, provide a confidence score between 0.0  
1284 and 1.0 indicating:

- 1285 - How likely this question is to lead to valuable information for this  
1286 user
- 1287 - Alignment with user's persona and checklist items
- 1288 - Relevance to the original research query

1289 Additionally, ALWAYS generate one "wild-card" question in the separate  
1290 wild\_card\_question field that goes outside the inferred persona but is  
1291 plausibly useful for broader understanding of the topic.

1292 Additionally, assign tags to categorize what aspects this research  
1293 content covers. Tags should be short phrases (2-4 words) that describe  
1294 the key topics, themes, or domains covered by the query, context, and  
1295 learnings. Be very cautious about adding new tags:

```

1296     - FIRST, check if any of the previously seen tags are relevant to this
1297     content
1298     - REUSE existing tags whenever possible
1299     - ONLY add new unseen tags if the content covers aspects not captured
1300     by existing tags
1301     - Keep tags concise and descriptive
1302
1303     Return your response as a JSON object with the following structure:
1304     {
1305         "learnings": [
1306             {
1307                 "insight": "key insight or finding relevant to the user",
1308                 "source_url": "URL of the source (if available)",
1309                 "relevance_to_user": "how this learning relates to the user's
1310                 interests"
1311             },
1312             "follow_up_questions": [
1313                 {
1314                     "follow_up_question": "follow-up question text",
1315                     "confidence": 0.0-1.0,
1316                     "reasoning": "why this is a good follow-up question"
1317                 },
1318                 "wild_card_question": {
1319                     "question": "wild-card question that goes outside the persona but
1320                     is plausibly useful",
1321                     "confidence": 0.0-1.0,
1322                     "reasoning": "why this wild-card question could be valuable"
1323                 },
1324                 "tags": ["tag1", "tag2", "tag3"]
1325             }
1326
1327             Research query: {query}
1328
1329             Search results:
1330             {context}
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

```

## Follow-up Questions to Search Queries

### System Prompt

You are an expert search query optimizer. Your task is to convert follow-up research questions into effective search queries that will yield relevant search results.

### User Prompt

Convert the following follow-up question into an optimized search query that will yield relevant search results.

Original research query: {original\_query}

User persona: {persona\_text}

User checklist (aspects they care about):  
{checklist\_items}

For each of the following follow-up question, create a search query that:

1. Effectively searches for information to answer the follow-up question
2. Is optimized for search engines
3. Maintains connection to the original research query
4. Considers the user's persona and interests

```

1350
1351 For each search query, also provide a clear research goal that
1352 describes:
1353 - What specific information or insights this search aims to discover
1354 - How it relates to the original research question
1355 - What direction of the topic it will explore
1356 =====
1357
1358 Follow-up questions:
1359 {followup_questions}
1360 =====
1361
1362 Return your response as a JSON object with the following structure:
1363 {
1364     "search_queries": [
1365         {
1366             "follow_up_question": "input follow-up question text",
1367             "search_query": "optimized search query",
1368             "research_goal": "clear description of what this search aims to
1369             discover and how it relates to the original research question"
1370         }
1371     ]
1372 }
1373
1374

```

### Persona Checklist Inference

```

1375 System Prompt
1376 You are an expert at understanding user personas and inferring what
1377 aspects they would care about in research responses. Your task is to
1378 analyze a user's persona and generate specific checklist items they
1379 would expect to see addressed.
1380
1381 User Prompt
1382 Given the following user persona and their research query, infer a
1383 checklist of specific aspects that this user would expect to see
1384 addressed in a comprehensive research response.
1385
1386 User persona: {persona_text}
1387
1388 Research query: {query}
1389
1390 Based on this persona and query, generate 5-8 specific checklist items
1391 that this user would expect to see in a helpful response. Each item
1392 should be:
1393 1. Specific to this user's background and interests
1394 2. Relevant to the research query
1395 3. Actionable and measurable
1396 4. Distinct from other items
1397
1398 Return your response as a JSON object with the following structure:
1399 {
1400     "checklist_items": [
1401         "specific aspect this user would expect to see addressed",
1402             "another specific aspect relevant to their interests"
1403     ]
1404 }

```

### Persona Modeling

1404  
1405     **System Prompt**  
1406     You are an expert at understanding user personas and updating them  
1407     based on user interactions. Your task is to analyze a user's response  
1408     and infer additional information about their persona and interests.  
1409  
1410     **User Prompt**  
1411     Given the following current persona and a user's response to a  
1412     research proposal, infer additional information about this user's  
1413     persona and interests.  
1414  
1415     Current persona: {current\_persona}  
1416  
1417     Current checklist items they already care about:  
1418     {current\_checklist}  
1419  
1420     User's response: {user\_response}  
1421  
1422     Based on this response, identify additional information about the  
1423     user's:  
1424        1. Background and interests  
1425        2. Specific preferences and priorities  
1426        3. Communication style and concerns  
1427        4. Any new aspects they care about  
1428  
1429     IMPORTANT: Do NOT output repetitive information:  
1430        - Only include NEW persona information that isn't already covered in  
1431        the current persona  
1432        - Only include NEW checklist items that aren't already in the current  
1433        checklist  
1434        - If nothing new can be inferred, return empty strings and empty  
1435        arrays  
1436  
1437     Return your response as a JSON object with the following structure:  
1438     {  
1439        "additional\_persona\_info": "new information to append to the persona  
1440        (empty if nothing new)",  
1441        "new\_checklist\_items": [  
1442            "new aspect they mentioned or implied they care about (only if not  
1443            already in checklist)",  
1444            "another new aspect if applicable"  
1445        ]  
1446     }  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457

## Clarification Question Generation

1443     **System Prompt**  
1444     You are an expert research assistant. Your task is to generate clear,  
1445     helpful clarification questions that present concise summaries of  
1446     research directions to users for selection.  
1447  
1448     **User Prompt**  
1449     Based on the following research context, generate a structured  
1450     clarification question that presents concise summaries of the  
1451     available research directions to the user for selection.  
1452  
1453     Research query: {query}  
1454  
1455     Available research directions (search queries):  
1456     {research\_directions}  
1457  
1458     User persona: {persona\_text}  
1459  
1460     Create a structured question that:

```

1458
1459     1. Starts with a natural introduction
1460     2. Lists each research direction as numbered bullet points (1., 2.,
1461        3., etc.)
1462     3. For each direction, provide a concise summary (1 sentence) that
1463        captures the essence of what that search query would explore, rather
1464        than showing the raw search query
1465     4. Provides clear selection instructions:
1466         - To select directions: just type the bullet numbers (e.g., "1, 3")
1467         - To suggest new follow-up questions: start a new line with "New
1468            follow-up questions:" followed by each new follow-up question on
1469            separate lines
1470     5. Matches the user's communication style
1471
1472     Return your response as a JSON object with the following structure:
1473     {
1474         "clarification_question": "your structured question to the user with
1475            concise summaries"
1476     }

```

## 1477 Report Generation

### 1478 System Prompt

1479 You are a professional research report writer specializing in  
1480 persona-aware reports. Create comprehensive, well-structured reports  
1481 based on research data with proper citations, tailored to the specific  
1482 user's background and interests.

### 1483 User Prompt

1484 Using the following hierarchically researched information and  
1485 citations:

1486 "context"

1487 Write a comprehensive research report answering the query: "question"

1488 User persona: {persona\_text}

1489 User interests (checklist items they care about):  
1490 {checklist\_items}

1491 The report should:

- 1492 1. Synthesize information from multiple levels of research depth
- 1493 2. Integrate findings from various research branches
- 1494 3. Present a coherent narrative that builds from foundational to  
1495 advanced insights
- 1496 4. Maintain proper citation of sources throughout
- 1497 5. Be well-structured with clear sections and subsections
- 1498 6. Have a minimum length of {total\_words} words
- 1499 7. Follow {report\_format} format with markdown syntax
- 1500 8. Use markdown tables, lists and other formatting features when  
1501 presenting comparative data, statistics, or structured information
- 1502 9. Be tailored to the user's persona and interests

1503 Additional requirements:

- 1504 - Prioritize insights that emerged from deeper levels of research
- 1505 - Highlight connections between different research branches
- 1506 - Include relevant statistics, data, and concrete examples
- 1507 - Focus on directions that align with the user's interests and  
1508 checklist
- 1509 - Use language and explanations appropriate for the user's background
- 1510 - Address the user's specific concerns and priorities

1512 - You MUST determine your own concrete and valid opinion based on the  
 1513 given information. Do NOT defer to general and meaningless  
 1514 conclusions.  
 1515 - You MUST prioritize the relevance, reliability, and significance of  
 1516 the sources you use. Choose trusted sources over less reliable ones.  
 1517 - You must also prioritize new articles over older articles if the  
 1518 source can be trusted.  
 1519 - Use in-text citation references in {report\_format} format and make it  
 1520 with markdown hyperlink placed at the end of the sentence or paragraph  
 1521 that references them like this: ([in-text citation] (url)).  
 1522 - Write in {language}  
 1523  
 1524 Citation requirements:  
 1525 - You MUST write all used source URLs at the end of the report as  
 1526 references  
 1527 - Make sure to not add duplicated sources, but only one reference for  
 1528 each  
 1529 - Every URL should be hyperlinked: [url website] (url)  
 1530 - Include hyperlinks to the relevant URLs wherever they are referenced  
 1531 in the report  
 1532 - Format example: Author, A. A. (Year, Month Date). Title of web page.  
 1533 Website Name. [url website] (url)  
 1534  
 1535 Please write a thorough, well-researched report that synthesizes all  
 1536 the gathered information into a cohesive whole, tailored specifically  
 1537 to this user's persona and interests.  
 1538 Assume the current date is {current\_date}.

### 1537 Persona Alignment Evaluation

1538  
**1539 System Prompt**  
 1540 You are an expert evaluator specializing in assessing how well  
 1541 research content aligns with user personas and interests. Your task is  
 1542 to analyze content and determine how well it addresses specific  
 1543 directions important to the user.  
 1544  
**1543 User Prompt**  
 1545 You are evaluating how well research content aligns with a user's  
 1546 persona and interests.  
 1547  
 1548 # User Persona: {persona\_text}  
 1549  
 1550 # Research Content:  
 1551 {content}  
 1552  
 1553 # Key Learnings:  
 1554 {learnings}  
 1555  
 1556 # Checklist Items to Evaluate:  
 1557 {checklist\_items}  
 1558  
 1559 For each checklist item, evaluate how well the research content and  
 1560 learnings address it.  
 1561 Provide a score from 0-2 for each item:  
 1562 - 0: Not addressed or covered  
 1563 - 1: Somewhat addressed or partially covered  
 1564 - 2: Well addressed or thoroughly covered  
 1565  
 1566 Return your evaluation as a JSON object with the following structure:  
 1567 {  
 1568 "evaluations": [  
 1569 {  
 1570 "item": "checklist item text",

```

1566
1567     "score": 0-2,
1568     "reasoning": "brief explanation of the score"
1569   }
1570 }
1571
1572
1573
1574
```

## L PROMPT TEMPLATES FOR DATA GENERATION

1575  
1576 We include here the prompt templates used in data generation.  
1577

### 1578 Profile Generation Prompt

#### 1580 User Prompt

1581 Generate a user profile for someone who would logically and reasonably  
1582 ask the following question: "{query}"

1583 The profile should include demographic and background information such  
1584 as age range, occupation, hobbies, family structure, education  
1585 background, or any other relevant facts. Note that you don't need to  
1586 include all of these details for each persona. You can use any kinds  
1587 of combinations and please think about other aspects other than  
1588 these. You should include something that can be elicited from daily and  
1589 natural conversations. You should not include too much information  
1590 about this person's work content and you should not give any  
1591 description about the user's personality traits. Focus on objective  
1592 facts about the person.

1593 Here are some example profiles for reference:  
1594 {profile\_examples}

1595 Generate a single user profile that contains around 8-10 distinct  
1596 facts about the person. The profile should logically connect to why  
1597 this person would ask the given question. You should only output the  
1598 profile in plain text format.

1599 IMPORTANT: Try to be creative and comprehensive. Make sure the profile  
1600 makes it realistic for this person to ask the specific question.

### 1601 Personality Generation Prompt

#### 1602 User Prompt

1603 Generate personality traits for a person with the following profile  
1604 who would ask this question: "{query}"

1605 Profile:

1606 {generated\_profile}

1607 Based on this profile and the question they would ask, generate  
1608 appropriate personality traits. You should include something that can  
1609 be elicited from daily and natural conversations. Each description  
1610 should contain around 8-10 personality traits about the person.

1611 Here are some example personality descriptions for reference:  
1612 {personality\_examples}

1613 Generate personality traits that are consistent with the profile and  
1614 make it logical for this person to ask the given question. You should  
1615 only output the personality description in plain text format.

1620  
 1621  
 1622     IMPORTANT: You should not include any other content that is beyond  
 1623     personality traits, such as occupation or demographic information  
 1624     (those are already in the profile). Focus only on personality  
 1625     characteristics, behavioral patterns, and psychological traits. Be  
 1626     creative and make sure the personality aligns with both the profile  
 1627     and the research question.

## 1629     Aspect Generation Prompt

### 1630     User Prompt

1631     Given a user's persona and their query, generate a list of specific  
 1632     aspects that this user would expect to see addressed in a high-quality  
 1633     response to their query. These aspects will serve as evaluation  
 1634     criteria to assess how well a response meets this specific user's  
 1635     needs and expectations.

1636     Query: "{query}"

1638     User Persona:

1639     {persona}

1640     Based on this persona and query, generate 5-8 specific aspects that  
 1641     this user would expect to see in a comprehensive and helpful response.  
 1642     Each aspect should be:

1. Specific to this user's background, needs, and context
2. Actionable and measurable (can be used to evaluate a response)
3. Relevant to the query and persona
4. Distinct from other aspects (no overlap)

1647     Format your response in JSON format where each aspect is a clear,  
 1648     specific expectation that can be used to evaluate whether a response  
 1649     adequately addresses this user's needs and provide a clear explanation  
 1650     of why each aspect is significant for the user and what specific  
 1651     details they would expect to see in the response. Focus on what  
 1652     content, depth, style, or approach would be most valuable for this  
 1653     specific user.

1654     Each aspect should have the following fields:

- aspect: a string that is the name of the aspect that is important to  
 be present in the response
- evidence: a string that points to specific details from the user's  
 persona that indicate this aspect is important
- reason: a string that explains why the aspect is important for the  
 user

1660     Use the following JSON structure:

```
{
  "aspects": [
    {
      "aspect": "Name of the aspect",
      "evidence": "Specific details from the user's persona that
                  indicate this aspect is important",
      "reason": "Explanation of why this aspect is important for the
                 user"
    }
  ]
}
```

1671     IMPORTANT: Make the aspects specific to this user's unique situation  
 1672     and needs, not generic aspects that would apply to any user asking  
 1673     this question.

1674 **M PROMPT TEMPLATES FOR EVALUATION**  
16751676 We include here the prompt templates used in evaluation scripts.  
16771678 **Alignment Evaluation Prompt**  
16791680 **User Prompt**1681 You are a fair and insightful judge with exceptional reasoning and  
1682 analytical abilities. Your task is to evaluate a user's question, a  
1683 generated response to that question, and multiple aspects that are  
1684 important to the user. Based on this information, assess how well each  
1685 aspect is addressed in the generated response. Provide a clear and  
1686 accurate assessment for each aspect.

1687 # Your input:

1688 - question: the question asked by the user  
1689 - persona: the user's persona (profile and personality) that the  
1690 aspects are based on  
1691 - response: a generated response to the user's question  
1692 - aspects: a list of aspects that are important to the user, each  
1693 consisting of:  
1694   - aspect: the title for the aspect  
1695   - reason: the reason that this aspect is important for the user  
1696   - evidence: the evidence from the user persona that the aspect was  
1697   extracted from

1698 # Your output:

1699 Your output should be a valid JSON object in ```json ``` format  
1700 containing the following fields:1701 - evaluations: A list of evaluations for each aspect, where each  
1702 evaluation contains:  
1703   - aspect: the aspect title  
1704   - match\_score: A score between 0 to 2 that indicates how well the  
1705 generated response addresses this aspect, where:  
1706     \* 0 means the response does not cover this aspect  
1707     \* 1 means the response somewhat covers this aspect  
1708     \* 2 means the response covers this aspect very well  
1709   - reasoning: A detailed explanation of why this score was assigned,  
1710     including specific examples from the response

1711 # Question: {question}

1712 # Persona: {persona}

1713 # Response: {response}

1714 # Aspects:

1715 {aspects\_formatted}

1716 Output:

1717  
1718 **Sentence Focus Evaluation Prompt**1719 **User Prompt**1720 You are an expert judge evaluating whether sentences in a report cover  
1721 specific user aspects. For each sentence, determine which aspects (if  
1722 any) it addresses.

1723 # Your input:

1724 - sentences: numbered sentences from a report  
1725 - aspects: user aspects with IDs, titles, and reasons1726  
1727 # Your task:

```

1728 For each sentence, identify whether it covers any of the user aspects.
1729 **BE EXTREMELY STRICT** in your evaluation.
1730
1731 A sentence covers an aspect ONLY if it:
1732 1. Directly addresses the specific concern or interest described in
1733    the aspect
1734 2. Provides substantive, detailed information that would be valuable
1735    to someone with that specific aspect
1736 3. Goes beyond mere keyword mentions or general background information
1737
1738 A sentence does NOT cover an aspect if it:
1739 - Only provides general background or introductory information
1740 - Mentions keywords related to the topic but doesn't address the
1741    specific concern
1742 - Gives broad overviews without targeting the particular interest
1743 - Describes general principles without connecting to the specific
1744    aspect
1745 - Is just factual information that doesn't serve the user's particular
1746    need
1747
1748 **Default to NOT covering aspects unless there is clear, direct,
1749 substantial relevance to the specific user concern.**
1750
1751 # JSON Schema for output:
1752 {
1753     "type": "object",
1754     "patternProperties": {
1755         "^\$": {
1756             "type": "object",
1757             "properties": {
1758                 "cover_aspects": {
1759                     "description": "A list of aspect IDs that the sentence
1760                     covers. If the sentence does not cover any of the aspects,
1761                     the list should be empty.",
1762                     "type": "array",
1763                     "items": {"type": "integer"}
1764                 },
1765                 "reasoning": {"type": "string"}
1766             },
1767             "required": ["cover_aspects", "reasoning"]
1768         }
1769     }
1770 }
1771
1772 # Sentences:
1773 {sentences_formatted}
1774
1775 # Aspects:
1776 {aspects_formatted}
1777
1778 Output valid JSON:
1779
1780
1781

```

## Key Point Extract Prompt

```

1782
1783 User Prompt
1784 Based on the report provided, identify key points in the report that
1785 directly help in responding to the query. The key points are not
1786 simply some key content of the text, but rather the key points that
1787 are important for **answering the query**. IMPORTANT: Ensure each
1788 point is helpful in responding to the query. Keep the point using the
1789 original language and do not add explanations. IMPORTANT: Each span
1790 must be a single consecutive verbatim span from the corresponding
1791 passages. Copy verbatim the spans, don't modify any word! Your
1792 response should state the point number, followed by its content, and
1793 spans in the text that entail the key point. Respond strictly in JSON
1794 format:
1795 {
1796   "points": [ {
1797     "point_content": point_content,
1798     "spans": [span1, span2, ...]
1799   }, ... ]
1800 }
1801
1802 Remember:
1803 - Key points can be abstracted or summarized, but the span must be a
1804 copy of the original text. The content of the key point does NOT need
1805 to be the same as that of the span.
1806 - These keypoints must be helpful in responding to the query.
1807 - If there are multiple spans for a point, add all of them in the spans
1808 list.
1809
1810 Report: {report}
1811
1812 Query: {query}
1813
1814 Output:

```

## Key Point Focus Evaluation Prompt

```

1815 User Prompt
1816 You are an expert judge evaluating whether key points of a report
1817 cover specific user aspects to answer a query. For each key point,
1818 determine which aspects (if any) it addresses. **BE EXTREMELY STRICT**
1819 in your evaluation.
1820
1821 A key point covers an aspect ONLY if it:
1822 1. Directly addresses the specific concern or interest described in
1823 the aspect
1824 2. Provides substantive, detailed information that would be valuable
1825 to someone with that specific aspect
1826 3. Goes beyond mere keyword mentions or general background information
1827
1828 A key point does NOT cover an aspect if it only provides introductory
1829 information or broad overviews
1830 **Default to NOT covering aspects unless there is clear, direct,
1831 substantial relevance to the specific user concern.**
1832
1833 Response strictly in JSON format:
1834 {
1835   "point_number": {
1836     "cover_aspects": [aspect1, aspect2, ...],
1837     "reasoning": reasoning
1838   },
1839 }
1840
1841 # Query:
1842 {query}

```

```

1836
1837
1838     # Report Key Points:
1839     {key_points_formatted}
1840
1841     # UserAspects:
1842     {aspects_formatted}
1843
1844     Output:
1845
1846 User Agent Alignment Evaluation Prompt
1847
1848 User Prompt
1849 You are a fair and insightful judge with exceptional reasoning and
1850 analytical abilities. Your task is to evaluate a user's follow-up
1851 questions in regard to a query, and multiple aspects that are
1852 important to the user. Based on this information, assess how well the
1853 follow-up questions trying to cover the user's interested aspects. An
1854 aspect is considered covered if there are follow-up questions are
1855 trying to initiate research directions that are related to the aspect.
1856 Provide a clear and accurate assessment for each aspect.
1857
1858 # Your input:
1859 - query: the query asked by the user
1860 - persona: the user's persona (profile and personality) that the
1861 aspects are based on
1862 - follow-up questions: a list of follow-up questions that the user
1863 asked
1864 - aspects: a list of aspects that are important to the user, each
1865 consisting of:
1866     - aspect: the title for the aspect
1867     - reason: the reason that this aspect is important for the user
1868     - evidence: the evidence from the user persona that the aspect was
1869     extracted from
1870
1871 # Your output:
1872 Your output should strictly be a valid JSON object:
1873 {
1874     "evaluations": [ {
1875         "aspect": aspect,
1876         "match_score": match_score,
1877         "reasoning": A detailed explanation of why this score was
1878             assigned, including specific examples from the follow-up questions
1879     }, ... ]
1880 }
1881
1882 "match_score" is a score between 0 to 2 that indicates how well the
1883 follow-up questions addresses this aspect, where:
1884     * 0 means the follow-up questions does not cover this aspect
1885     * 1 means the follow-up questions somewhat covers this aspect
1886     * 2 means the follow-up questions covers this aspect very well
1887
1888 # Query: {query}
1889
1890 # Persona: {persona}
1891
1892 # Follow-up Questions:
1893 {follow_up_questions_formatted}
1894
1895 # Aspects:
1896 {aspects_formatted}
1897
1898 Output:

```

1890

**User Response Precision Evaluation Prompt**

1891

**User Prompt**

1892

You are an expert judge evaluating whether a user's follow-up questions or responses are truly targeted to specific user aspects for answering a query. For each follow-up, determine which aspects (if any) it substantively targets. BE EXTREMELY STRICT.

1896

A follow-up COVERS an aspect ONLY if it:

1897

- 1) Clearly aims to gather information directly relevant to the specific concern described by the aspect; AND
- 2) Goes beyond surface keywords or generic curiosity.

1900

A follow-up does NOT cover an aspect if it:

1902

- Is a broad/background question without tailoring to that aspect; OR
- Only mentions related keywords but lacks a targeted objective tied to the aspect; OR
- Is unrelated to the user's stated concerns.

1904

Respond strictly in JSON format:

1905

```
{
  "response_number": {
    "cover_aspects": [aspect_id_1, aspect_id_2, ...],
    "reasoning": reasoning
  },
  ...
}

# Query:
{query}

# User Responses (indexed from 0):
{user_responses_formatted}

# User Aspects (IDs start at 0):
{aspects_formatted}
```

1921

Output:

1922

**Final Persona State Evaluation Prompt**

1924

**User Prompt**

1925

You are a fair and insightful judge with exceptional reasoning and analytical abilities. Your task is to evaluate how well items from a final persona state checklist cover user aspects. Given the user's query, the original persona, a list of checklist items, and the user aspects, assess for each aspect how well the checklist covers it.

1930

# Your input:

1931

- query: the query asked by the user
- persona: the user's original persona text
- checklist: a list of items inferred that might be important for the user to answer the query
- aspects: a list of aspects that are indeed important to the user as ground truth, each consisting of aspect, reason, and evidence

1936

# Your output:

1938

Return strictly valid JSON of the form:

1939

```
{
  "evaluations": [
    "aspect": aspect_title,
    "match_score": 0|1|2,
    "reasoning":
      detailed_reasoning_referencing_specific_checklist_items
```

```
1944
1945     }, ... ]
1946 }
1947 Interpret match_score as:
1948 - 0: the checklist does not cover this aspect
1949 - 1: the checklist somewhat covers this aspect
1950 - 2: the checklist covers this aspect very well
1951
1952 # Query: {query}
1953
1954 # Persona: {persona}
1955
1956 # Checklist Items:
1957 {checklist_formatted}
1958
1959 # Aspects:
1960 {aspects_formatted}
1961
1962 Output:
```

## N ADDITIONAL PROMPT TEMPLATES

We include here the prompt templates used for User Agent and Pause Agent.

## User Agent

## System Prompt

You are simulating a real user with a specific persona and interests. Your task is to respond to SteER's research proposals by selecting relevant directions and suggesting new directions based on your persona and research interests.

### User Prompt

You are acting as a user with the following persona:

```
# User Persona:  
{persona_text}
```

```
# Aspects and directions You Care About:  
{aspects_text}
```

```
# History of your previous asked questions:  
{questions_history_text}
```

```
# Research Query:  
{query}
```

```
# SteER's Proposal:  
{steer_proposal}
```

SteER is presenting research directions as numbered bullet points. Based on your persona and interests, respond as this user would by:

1. Selecting ONLY the most relevant direction numbers that have the highest priority for this research
2. Suggesting new follow-up questions ONLY if you feel there's a very important direction missing from the proposal
3. Providing natural commentary as this user would speak

\*\*IMPORTANT CONSTRAINTS:\*\*

\*\*IMPORTANT CONSTRAINTS:\*\*  
- \*\*DO NOT select directions or suggest questions that are outside your persona and aspects/interests\*\*

```

1998
1999   - **DO NOT suggest questions you have already asked before or that are
2000   similar to the questions you have already asked (check your history
2001   above)**
2002   - Only focus on areas that align with your specific expertise,
2003   interests, and concerns as described in your persona
2004   - If all current directions seem unrelated to your interests, it's
2005   better to select none and suggest relevant alternatives
2006
2007   Focus on quality over quantity - select only the directions that truly
2008   matter most to you and align with your expertise. You should refrain
2009   from suggesting new follow-up questions unless something critical is
2010   missing and directly relates to your interests.
2011
2012   You should at most suggest 1 new follow-up question.
2013
2014   The probability of you suggesting a new follow-up question is 50%.
2015
2016   Your response should reflect how this person would actually
2017   communicate when discussing their research preferences.
2018
2019   Return your response as a JSON object with the following structure:
2020   {
2021     "selected_directions": [
2022       {
2023         "number": 1,
2024         "direction": "direction name from the proposal",
2025         "reasoning": "why this direction is most important to you and
2026         aligns with your interests"
2027       }
2028     ],
2029     "new_follow_up_questions": [
2030       {
2031         "follow_up_question": "suggested new follow-up question. Most of
2032         the time you should not suggest new follow-up questions. But
2033         only if you feel there's a very important direction missing from
2034         the proposal, suggest one new follow-up question at most",
2035         "reasoning": "why this follow-up question is important, missing,
2036         and relevant to your interests"
2037       }
2038     ],
2039     "user_response": "natural response as this user would speak (in the
2040       format: selected numbers with reasoning in parentheses, then 'New
2041       follow-up questions:' if any)",
2042     "additional_context": "any additional preferences or clarifications
2043       related to your expertise"
2044   }
2045
2046
2047
2048
2049
2050
2051

```

## Pause Agent

### System Prompt

You are an expert research assistant specialized in making optimal pause decisions during deep research. Your task is to analyze the current research state and decide whether it's a good time to pause and ask for user guidance on which research branches to pursue, or to proceed with the current research plan.

### User Prompt

You need to decide whether to pause and ask for user guidance or proceed with the current research plan.

```

# Context:
**Original Query:** {original_query}

```

```

2052
2053    **Current Research Goal:** {current_research_goal}
2054
2055    **User Persona:** {persona_text}
2056
2057    **User Interests (Checklist):** {checklist_items}
2058
2059    **Current Search Depth:** {current_depth} / {max_depth}
2060
2061    # Available Research Branches:
2062    branch_summaries
2063
2064    # Decision Criteria:
2065    Consider pausing (PAUSEASK) when:
2066    - User input would help prioritize which direction to pursue
2067    - There's uncertainty about which direction aligns best with the
2068    user's specific interests
2069    Analyze the situation and make your decision. Your reasoning should be
2070    specific to the current research context, user persona, and branch
2071    characteristics.
2072
2073    Respond in the following JSON format:
2074    {
2075        "type": "object",
2076        "properties": {
2077            "action": {
2078                "type": "string",
2079                "enum": ["PROCEED", "PAUSEASK"],
2080                "description": "Decision to proceed with research or pause to
2081                ask user for guidance"
2082            },
2083            "reasoning": {
2084                "type": "string",
2085                "description": "Detailed explanation of the decision based on
2086                research context and user persona"
2087            }
2088        },
2089        "required": ["action", "reasoning"]
2090    }
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105

```