

DeepSynth-Eval: Objectively Evaluating Information Consolidation in Deep Survey Writing

Hongzhi Zhang^{†*}, Yuanze Hu^{‡*}, Tinghai Zhang^{†*}, Jia Fu^{†*}, Tao Wang^{†*},
Junwei Jing[†], Zhaoxin Fan[‡], Wei Bi[†], Ruiming Tang^{†✉}, Han Li[†], Guorui Zhou[†], Kun Gai[†]

[†]Kuaishou Technology

[‡]Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing

Abstract

The evolution of Large Language Models (LLMs) towards autonomous agents has catalyzed progress in Deep Research. While retrieval capabilities are well-benchmarked, the post-retrieval synthesis stage—where agents must digest massive amounts of context and consolidate fragmented evidence into coherent, long-form reports—remains under-evaluated due to the subjectivity of open-ended writing.

To bridge this gap, we introduce DeepSynth-Eval, a benchmark designed to objectively evaluate information consolidation capabilities. We leverage high-quality survey papers as gold standards, reverse-engineering research requests and constructing "Oracle Contexts" from their bibliographies to isolate synthesis from retrieval noise. We propose a fine-grained evaluation protocol using General Checklists (for factual coverage) and Constraint Checklists (for structural organization), transforming subjective judgment into verifiable metrics. Experiments across 96 tasks reveal that synthesizing information from hundreds of references remains a significant challenge. Our results demonstrate that agentic "plan-then-write" workflows significantly outperform single-turn generation, effectively reducing hallucinations and improving adherence to complex structural constraints.

1 Introduction

The evolution of Large Language Models (LLMs) (DeepSeek-AI et al., 2025) marks a paradigm shift from fundamental token completion to sophisticated task completion via tool augmentation (Yao et al., 2023). This transition has notably accelerated progress in Deep Research (OpenAI, 2025; Google, 2024; Shao et al., 2024), where agents are tasked with not only retrieving information but also

synthesizing it into comprehensive, survey-style reports. While Information Retrieval (IR) has robust benchmarks such as GAIA (Mialon et al., 2023), HLE (Phan et al., 2025) and others (Wei et al., 2025; Wu et al., 2025), the latter stage—Post-Retrieval Synthesis—remains under-evaluated. In this stage, models must digest massive amounts of retrieved contexts, consolidate fragmented evidence, and organize them into a coherent narrative. Evaluating this capability is challenging. Unlike short-form QA tasks like GAIA and HLE where answers can be checked against ground-truth facts, deep survey reports are long, lack a unique reference, and vary widely in content selection, structure, and writing style. Consequently, current practices rely heavily on subjective human judgments or LLM-as-a-judge scores (Zheng et al., 2023), which are costly, unstable, and prone to reward hacking (Wang et al., 2024a). Furthermore, end-to-end evaluations often conflate retrieval noise with synthesis failures, obscuring the model's true ability to consolidate information.

To objectively evaluate this information consolidation capability, we propose to assess deep synthesis performance via survey papers. High-quality survey papers can be viewed as expert-written research reports that serve as "gold standards" of synthesis. For each survey, we first "reverse-engineer" a corresponding research request and construct the input context from its bibliography. We then manually extract a set of checklists that specify what a good report must cover under this request, including representative methods, key design choices, important datasets, and other domain-specific points. These checklists transform open-ended report evaluation into verifying whether the model's output accurately covers the required items, enabling a more objective measurement of completeness and factual correctness.

Beyond factual coverage, a hallmark of deep synthesis is structural organization—for example,

* Equal contribution.

✉ Corresponding author: tangruiming@kuaishou.com

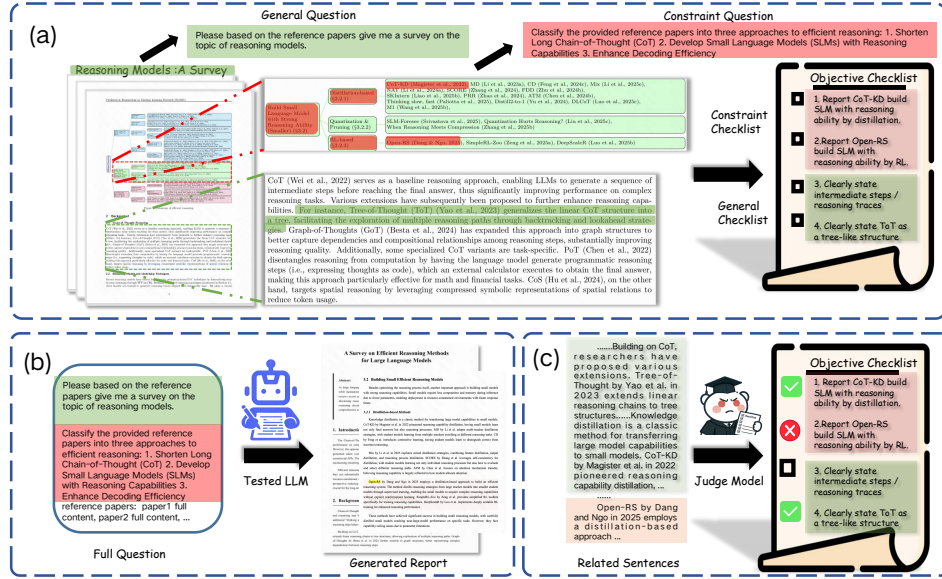


Figure 1: An illustration of DeepSynth-Eval. (a) From a reference survey, we derive a general prompt and constraint questions, and construct corresponding general/constraint checklists for evaluation. (b) The full question (prompt + constraints + reference papers) is fed to a tested LLM to generate a report. (c) A judge model verifies each checklist item on the generated report, converting subjective synthesis evaluation into objective item-wise verification.

consolidating methods into comparative tables or grouping them into meaningful taxonomies. However, such organization is inherently subjective and flexible: there are diverse valid ways to categorize works or select attributes for comparison. Consequently, expecting a model to spontaneously replicate the exact structure of a reference survey is unreasonable. To resolve this ambiguity and make organizational abilities objectively comparable, we extract a canonical set of structural constraints from each survey (e.g., specific taxonomies to adopt or attributes to include in a table). By explicitly imposing these constraints in the synthesis prompt, we reduce editorial freedom, which allows us to compare the organization of content in corresponding sections. We refer to these requirements as Constraint Questions and the corresponding checklists as Constraint Checklists.

Building on these ideas, we construct **DeepSynth-Eval (DSE)**, a benchmark designed for the objective evaluation of **deep synthesis** capabilities. DSE comprises 96 queries derived from high-quality surveys across diverse domains. For each query, the benchmark provides the reconstructed research request, a comprehensive **Oracle Context** (constructed from the bibliography), and a set of fine-grained general and constraint checklists. Given a model-generated report, our pipeline automatically aligns the

content with checklist items to compute coverage, precision, and structural constraint satisfaction scores. Compared to directly judging unstructured free-form text, this key-point verification reduces the knowledge burden on judge models, supports fine-grained and discriminative scoring, and makes the evaluation more robust and reproducible—providing a reliable foundation for training and improving deep synthesis systems.

Our main contributions are:

- **Objective Checklist Metric.** We address the challenge of evaluating long-form generation by introducing a checklist-based metric derived from gold-standard surveys. This method decomposes synthesis quality into verifiable factual coverage and structural constraint adherence, providing a robust standard for measuring deep research capabilities.
- **DeepSynth-Eval (DSE) Benchmark.** We introduce DSE, a benchmark comprising 96 tasks equipped with Oracle Contexts. This design isolates post-retrieval synthesis from retrieval noise, ensuring reproducibility and preventing ground-truth leakage.
- **Evaluation and Open Source.** We establish a rigorous baseline by benchmarking end-to-end and agentic workflows. We find that while agentic methods dominate, synthesizing context from over 100 references remains a

formidable bottleneck. We open-source the entire framework to accelerate research in this direction.

2 Related Work

The paradigm of LLMs is shifting from closed-ended tasks—exemplified by benchmarks like MMLU (Hendrycks et al., 2021), GPQA (Rein et al., 2023), and Humanity’s Last Exam (Wang et al., 2024b)—to complex **Deep Research**. Unlike these static evaluations, Deep Research requires integrating reading, searching, and writing. Existing works fall into three categories.

Long-Context Understanding Serving as the cognitive foundation for deep research, Long-Context Understanding has garnered significant attention; however, current evaluations are often confined to information extraction rather than logical restructuring. Existing benchmarks, such as LongBench (Bai et al., 2024), InfiniteBench (Zhang et al., 2024), and L-Eval (An et al., 2023), primarily assess the model’s capacity for information retention within extended context windows. However, the task formulations in these benchmarks are largely restricted to “Needle-in-a-Haystack” scenarios or local extractive QA. In such settings, models can typically succeed by locating and reciting specific segments within the context, without the necessity for global association or cross-document reasoning to synthesize unstructured evidence dispersed across numerous documents.

Agentic Retrieval and Fact-Finding Transitioning from passive information processing to active information acquisition, Deep Search benchmarks (e.g., GAIA (Mialon et al., 2023), BrowseComp (Wei et al., 2025)) focus on evaluating an agent’s ability to utilize tools for solving complex problems. However, these evaluations typically involve *low information throughput*, where retrieving a few key pages suffices. Consequently, they fall short of measuring the comprehensive ability to screen and synthesize massive information into coherent long-form reports. To address scale, WideSearch (Wong et al., 2025) targets breadth-oriented acquisition. Yet, despite the increased throughput, the task remains limited to **Structured Information Extraction** (e.g., compiling tables). This primarily tests recall in repetitive tasks, lacking the depth required for cross-document comparative analysis and the integration of heterogeneous

information needed in deep research.

End-to-End Research System Evaluations

Deep Research tasks inherently require agents to complete the full workflow from open-web searching to long-form report generation (Search-Read-Write). To align with this real-world application scenario, recent works have adopted an end-to-end evaluation paradigm; however, this setting introduces confounding variables that are difficult to isolate. For instance, PDR-Bench (Liang et al., 2025) and DeepResearchBench (Du et al., 2025) focus on personalization alignment and LLM-based subjective scoring for complex tasks, respectively, whereas ReportBench (Li et al., 2025) assesses the Search-Read-Write capabilities of end-to-end systems by reverse-engineering arXiv survey papers.

Despite these significant explorations in system-level assessment, this open-web evaluation methodology faces two major methodological challenges. The first challenge is the **Coupling of Retrieval and Generation**. The *dynamic drift of rankings* causes context variations over time, precipitating a *reproducibility crisis* where it is indistinguishable whether gains stem from model enhancements or better retrieval luck. Furthermore, retrieval noise (e.g., ads, broken links) prevents accurate measurement of the model’s *intrinsic cognitive capability* in processing complex information. The second is the **Subjectivity of Evaluation**. Existing methods predominantly rely on LLMs to assign holistic scores (score-based) to reports, lacking objective verification of specific factual point coverage and remaining susceptible to biases inherent in the judge models.

3 Benchmark Construction

We propose a benchmark construction framework that uses high-quality survey papers as reference answers. Formally, given a survey report R , we first reverse-engineer a query q_{gen} for which R serves as an answer; we term this the **General Question** (q_{gen}) (e.g., “Summarize recent advances in reasoning models”). From each survey, we extract two types of evaluation checklists. The first is the **General Checklist**, which captures the core content that a comprehensive report should cover under q_{gen} . Each item in this checklist corresponds to a specific, verifiable fact derived from the reference survey. The second is the **Constraint Checklist**, which captures structural or organizational requirements imposed on the report, such as a particular

taxonomy, grouping principle, or presentation format.

Intuitively, one might evaluate a model by directly comparing the generated report $\hat{R} = \mathcal{M}(q_{\text{gen}})$ with the original survey R . In practice, however, directly comparing two long, unstructured reports is both computationally expensive and metrically unreliable. Checklist-based evaluation provides a more tractable alternative by decomposing the reference survey into verifiable evaluation points. The distinction between the two checklist types is important because survey writing is inherently subjective. Different authors may adopt equally valid but different strategies to organize the same body of literature. For example, benchmarks for Large Language Models (LLMs) may be categorized by *capability domain* (e.g., reading comprehension, mathematics, coding) or by *task format* (e.g., multiple-choice versus open-ended generation). As a result, structural organization is often not directly comparable unless the intended organizing principle is explicitly specified.

To address this issue, we extract structural requirements from the reference survey and formulate them as **Constraint Questions** (q_{con}), which are appended to the general prompt. For example, we may require the model to organize efficient reasoning methods into three specific categories: (1) *Shortening Chain-of-Thought*, (2) *Developing Small LLMs with reasoning capabilities*, and (3) *Enhancing decoding efficiency*. The associated **Constraint Checklists** are used to evaluate whether the generated report follows this prescribed organization. Importantly, Constraint Checklists are valid only when the corresponding constraint question q_{con} is provided. Without q_{con} , alternative organizational strategies may still be reasonable and should not be penalized. By imposing q_{con} , we reduce the editorial freedom of the synthesis task and make structural aspects of the output objectively comparable. This design also enables a more rigorous evaluation of the model’s instruction-following ability in complex writing tasks. Overall, given a composite input consisting of the general question q_{gen} and a set of constraint questions Q_{con} , the model synthesizes a report based on the reference bibliography of the original survey. An evaluation system then verifies the generated report \hat{R} against the full checklist set $C = C_{\text{gen}} \cup C_{\text{con}}$ to determine whether each requirement is satisfied.

3.1 Automated Construction Pipeline

The construction of this writing benchmark itself could be very labor-consuming. In this paper, we utilize LLMs to generate candidates and human labor to review and cross-validate the dataset.

To construct the benchmark efficiently while maintaining high quality, we employ a multi-stage generation pipeline utilizing LLMs.

General Question Generation (q_{gen}). This step is straightforward. We treat the survey title and abstract as the core intent and employ an LLM to directly reverse-engineer the corresponding research request. This produces the general question q_{gen} (e.g., “Provide a comprehensive survey on ...”), which serves as the foundation of the task.

General Checklist Generation (C_{gen}). Directly generating a comprehensive set of fine-grained checklists from a lengthy document is challenging for current models, often resulting in low recall or hallucinations. To mitigate this, we adopt a hierarchical **Topic Anchor** strategy. First, we prompt the LLM to synthesize several distinct *topics* based on the survey’s section headers and content (e.g., “Core Algorithms”, “Evaluation Metrics”, “Future Directions”). Then, using these topics as anchors, we invoke the LLM to generate specific checklist items for each topic individually. This divide-and-conquer approach yields K groups of general checklists, denoted as $C_{\text{gen}} = \bigcup_{k=1}^K C_{\text{gen},k}$.

Constraint Pair Generation ($Q_{\text{con}}, C_{\text{con}}$). For structural constraints, we employ a **Structure Anchor** approach to generate coupled pairs of constraint questions and checklists.

1. **Anchor Extraction:** The model first scans the survey to identify key organizational elements, such as specific taxonomies (classification schemes) or summary tables.
2. **Guidance Generation** ($q_{\text{con}}^{(m)}$): Based on an extracted anchor (e.g., a table summarizing reasoning datasets), the model generates a **Constraint Question** $q_{\text{con}}^{(m)}$ that acts as explicit guidance. For instance, if the anchor is a dataset table, $q_{\text{con}}^{(m)}$ might be: “List a table summarizing recent reasoning datasets, specifying their data scale, question source, and answer source.”
3. **Checklist Extraction:** Simultaneously, the model extracts the corresponding cell values or classification criteria from the anchor to form a group of **Constraint Checklists** $C_{\text{con},m}$.

This process yields M groups of constraint checklists, denoted as $C_{\text{con}} = \bigcup_{m=1}^M C_{\text{con},m}$.

3.2 Human Verification

We construct DeepSynth-Eval using a hybrid strategy that combines expert curation and assisted expansion, where each survey paper defines exactly one evaluation task. We first asked computer science experts to identify 60 representative survey papers, then expanded this seed set through an automated arXiv-based retrieval pipeline, yielding 322 candidate tasks.

For each survey, we collected its bibliography and constructed the Oracle Context using structured metadata from the Semantic Scholar API, supplemented by a local index from the HuggingFace arXiv dataset; long-tail references were recovered by parsing PDFs with Qwen3-235B and matching titles through the arXiv API.

We then conducted a two-stage human verification process with six computer science experts, retaining 129 candidates after the first round and 96 final benchmark instances after cross-validation and refinement. During screening, we identified and discarded problematic instances characterized by two common issues: (1) over-specification, where constraints made checklist items trivially inferable; and (2) contextual ambiguity, especially in table-based extractions lacking sufficient textual context for verification.

3.3 Task Formalization

Formally, we define a deep synthesis evaluation task in DeepSynth-Eval as a tuple $T = (P, D, C)$, where:

- P denotes the **Unified Input Prompt**. It is constructed by concatenating the general question q_{gen} with a set of M constraint questions $\{q_{\text{con}}^{(1)}, \dots, q_{\text{con}}^{(M)}\}$. Formally:

$$P = q_{\text{gen}} \oplus q_{\text{con}}^{(1)} \oplus \dots \oplus q_{\text{con}}^{(M)}$$

where \oplus represents the string concatenation operation.

- D represents the **Oracle Context**, constructed from the reference bibliography associated with the original survey. This serves as the closed-world knowledge boundary for the task.
- C is the **Evaluation Standard**, composed of the union of general checklist groups and constraint checklist groups: $C = C_{\text{gen}} \cup C_{\text{con}}$.

3.4 Evaluation Metric

During inference, the model \mathcal{M} receives the unified prompt P and the reference context D to synthesize a report $\hat{R} = \mathcal{M}(P, D)$. Here, P consists of the general question q_{gen} and, when applicable, the corresponding constraint questions, and the full checklist set is defined as $C = C_{\text{gen}} \cup C_{\text{con}}$. The quality of \hat{R} is quantified by evaluating its satisfaction of C using a judge model.

For each checklist item $c \in C$, the judge assigns one of three verification outcomes:

$$r(c) = \begin{cases} 1, & \text{if mentioned correctly (Correct),} \\ 0, & \text{if not mentioned (Omission),} \\ -1, & \text{if mentioned incorrectly (Incorrect).} \end{cases} \quad (1)$$

Each checklist item has unit weight, i.e., every item contributes equally within its checklist group.

Fault-Tolerant Group Scoring. Checklist items are organized into semantic groups (e.g., representative methods under a subtopic). Let G_k denote the set of items in group k , and let $N_k = |G_k|$ be the number of items in that group. The cumulative reward of group k is defined as

$$S_k = \sum_{c \in G_k} r(c). \quad (2)$$

A naive normalization by N_k would over-penalize long groups and allow checklist groups with many items to dominate the final score. This is undesirable because survey writing is open-ended: covering the majority of representative works is often sufficient for a high-quality synthesis, and extremely long enumerations should not disproportionately affect evaluation.

To address this issue, we introduce a **saturation threshold** $\theta_k \leq N_k$ for each group and define

$$\text{Score}_k = \max\left(0, \min\left(1, \frac{S_k}{\theta_k}\right)\right). \quad (3)$$

Once the cumulative reward reaches the threshold, i.e., $S_k \geq \theta_k$, the group is treated as fully covered. For example, if a group contains $N_k = 15$ items and we set $\theta_k = 10$, then reaching a cumulative reward of 10 already yields full credit for that group; the remaining five items do not further increase the score. This design prevents unusually long checklist groups from dominating the overall metric and makes scores across groups more comparable.

Group-Level Aggregation. We report three metrics: **General Score** (S_{gen}), **Constraint Score**

(S_{con}), and **Overall Score** (S_{all}). To ensure fairness across checklist groups, we aggregate at the group level rather than the item level, assigning equal weight to each group. Formally, if there are K general groups and M constraint groups, then

$$S_{\text{gen}} = \frac{1}{K} \sum_{k=1}^K \text{Score}_{\text{gen},k} \times 100\%, \quad (4)$$

$$S_{\text{con}} = \frac{1}{M} \sum_{m=1}^M \text{Score}_{\text{con},m} \times 100\%, \quad (5)$$

and

$$S_{\text{all}} = \frac{\sum_{k=1}^K \text{Score}_{\text{gen},k} + \sum_{m=1}^M \text{Score}_{\text{con},m}}{K + M} \times 100\%. \quad (6)$$

This aggregation scheme ensures that no single long checklist group can dominate the final score simply because it contains more items.

Precision. In addition to coverage-oriented scores, we also report precision over mentioned checklist items:

$$\text{Precision} = \frac{N_{\text{Correct}}}{N_{\text{Correct}} + N_{\text{Incorrect}}}. \quad (7)$$

Here, N_{Correct} and $N_{\text{Incorrect}}$ denote the numbers of checklist items judged as Correct and Incorrect, respectively.

3.5 Statistics

We present benchmark statistics based on the checklist items and citations of each task.

Distribution of Checklist Items per Task: As illustrated in Figure 2, the dataset maintains a high standard of evaluation granularity: on average, each task entails **199.7** checklist items, comprising **128.2** general items and **71.5** constraint items. This dense distribution underscores the complexity of the tasks, as models are required to satisfy a vast array of atomic facts and constraints to achieve a high score, effectively precluding the success of generic responses. By transforming the holistic evaluation of lengthy reports into fine-grained checklist matching, we convert an inherently subjective task into a quantifiable metric, significantly enhancing the objectivity of the report assessment.

Distribution of Total Reference Lengths: Figure 3 illustrates the distribution of total reference lengths per task, measured in words. Notably, the cumulative text length for a significant portion of tasks in DeepSynth-Eval exceeds 1 million words. This magnitude underscores the rigorous nature of the benchmark, placing substantial demands on

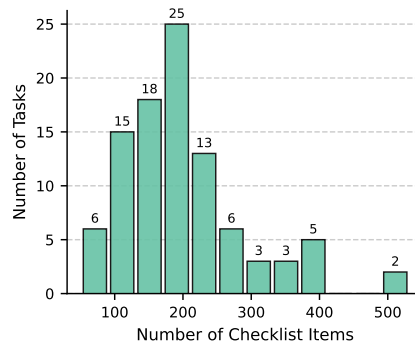


Figure 2: Distribution of checklist item counts




Figure 3: Distribution of total reference word counts

the information processing, long-context retention, and summarization capabilities of the evaluated models.

4 Experiments and Evaluation

4.1 Evaluation Settings

We aim to evaluate the deep synthesis capability of models, formalized as $\hat{R} = \mathcal{M}(P, D)$. As observed in Figure 3, the total length of the reference materials underlying the Oracle Context D exceeds 1 million words for a significant portion of tasks. Since the context window of most mainstream LLMs falls below 256k tokens—with only a few extending to 1 million—directly feeding the full context often results in context overflow.

Nevertheless, we believe that direct reading of such massive contexts is a challenging task that holds significant value for benchmarking emerging long-context capabilities and memory mechanisms. Unlike standard reading comprehension, which often relies on retrieving specific “needles in a haystack” (Liu et al., 2024; Hsieh et al., 2024), our report generation task requires high information throughput, making it a more challenging testbed for global context understanding.

To evaluate the deep synthesis ability of main-

Workflow	Model	Evaluation Metrics(%)			
		Overall	General	Constraint	Precision
-	Reference Survey	96.1	95.5	98.9	99.6
E2E Single-turn	Qwen3-30B-A3B-Thinking-2507	6.3	4.8	11.4	69.5
	Qwen3-235B-A22B-Thinking-2507	24.8	24.7	23.3	79.9
	DeepSeek-V3.2	23.4	22.2	27.0	79.8
	GPT-5.2	28.3	26.4	36.1	85.5
Agentic Multi-turn	Qwen3-30B-A3B-Thinking-2507	17.3	17.7	16.0	90.1
	Qwen3-235B-A22B-Thinking-2507	35.5	37.5	27.5	92.9
	DeepSeek-V3.2	30.4	31.4	27.0	91.4
	GPT-5.2	33.3	34.8	26.2	95.3
	GPT-5.2 (59/96 tasks subset)	37.0	39.0	28.4	95.0

Table 1: Main results under our two controlled synthesis workflows on DeepSynth-Eval.

stream models within current constraints, we designed two representative scaffolds for report writing.

E2E Single-turn workflow. Instead of reading the full text of all reference materials in D , we utilize a pre-processing stage to summarize each document relevant to the writing request. We then assemble the Oracle Context by concatenating these summaries, each paired with a stable numeric citation ID. Finally, we append a unified instruction P —constructed by merging the general request with the extracted constraint questions—and the model generates the full survey report in a single pass, without an explicit intermediate planning stage or iterative refinement.

Agentic Multi-turn workflow. The Agentic Multi-turn workflow utilizes the same summarized Oracle Context but decomposes generation into coordinated stages—*Planning*, *Iterative Chapter Writing*, and *Global Polishing*—to strengthen organization and grounding. First, the model produces an *intellectual skeleton*—a high-level plan specifying section taxonomy, comparison axes, and content allocation—which is subsequently parsed into structured sub-tasks for each section. Second, for each sub-task, the model executes a ‘read-then-write’ loop starting with *selective deep reading*: it identifies critical papers based on task requirements and generates detailed reading notes by consulting their full text from the underlying reference collection, ensuring fine-grained details are captured without context overflow. Subsequently, the model writes each section conditioned on the skeleton, detailed notes, and previous sections, adhering to rules that encourage grouped comparisons and tabular consolidation. Finally, the model aggregates




Figure 4: Checklist requirement status breakdown across four E2E models.

and polishes the sections to ensure global coherence. Importantly, the key innovation lies in the dynamic, task-driven paper selection, which significantly improves grounding compared to single-turn generation.

Detailed textual descriptions of the implementation steps for both workflows are provided in Appendix A. In the evaluation, DeepSeek v3.2 is employed as the underlying model for checklist evaluation and summary reading in both workflows.

4.2 Results Analysis

Table 1 reports the main results under two controlled synthesis workflows—E2E single-turn and agentic multi-turn—with all models given the same Oracle Context. Our goal is to figure out how well a model can *synthesize* the Oracle Context into a structured survey that satisfies both content coverage and structural constraints. We focus on two aspects: (i) model-to-model differences under the single-turn baseline, which isolates single-pass consolidation over the provided context, and (ii) the effect of introducing multi-turn planning under the same model, which tests whether staged generation improves synthesis quality.

Judge Validity. We first validate the judge system using the Reference Survey (Row 1). It




Figure 5: Score distribution of E2E versus Agentic workflows on Qwen3-235B-A22B.




Figure 6: Checklist requirement status breakdown for E2E versus Agentic workflows on Qwen3-235B-A22B.

achieves a 96.1% Overall score, with the minor gap (<5%) primarily attributable to textual ambiguity and judge error. Notably, the near-perfect Precision (99.6%) confirms the judge’s high fidelity, indicating that correct content is almost never erroneously penalized. These results establish a reliable upper bound for our evaluation.

Model comparison under the E2E single-turn workflow. Under the E2E single-turn setting, we observe large performance variance across models, revealing substantial differences in their ability to transform the same Oracle Context into a well-structured synthesis. GPT-5.2 achieves the strongest Overall score of 28.3 and leads by a clear margin on Constraint (36.1), while also obtaining the highest Precision (85.5). This pattern suggests that GPT-5.2 is comparatively better at *executing explicit structural instructions* and *grounding mentioned content* within a single pass. We further examine the checklist requirement status breakdown to understand where these gains come from. As shown in Figure 4, GPT-5.2 attains the highest fraction of correctly satisfied requirements and the lowest incorrect rate among the E2E models, indicating fewer unsupported completions when mapping Oracle evidence to checklist items. Qwen3-235B and DeepSeek-V3.2 form a middle tier with similar Overall scores (24.8 and 23.4), but without a corresponding advantage on Constraint (23.3 and 27.0), indicating that their main bottleneck lies in

organizing and covering required content rather than in factuality once content is produced. In contrast, Qwen3-30B largely collapses in the single-turn regime (Overall 6.3; General 4.8), consistent with Figure 4, where most requirements fall into the not-mentioned category, highlighting a strong capacity dependence for long-context consolidation.

Workflow effect: agentic multi-turn improves synthesis quality. Switching from single-turn generation to the agentic multi-turn workflow consistently improves Overall and General across all evaluated models: Qwen3-30B increases from 6.3→17.3, Qwen3-235B from 24.8→35.5, and GPT-5.2 from 28.3→33.3. We note that GPT-5.2 often fails to rewrite the full text during polishing, instead referring to the previous draft (Appendix E); excluding these invalid responses reveals a higher Overall score of 37.0 (Table 1).

We then perform a finer-grained distributional analysis over tasks to go beyond mean scores. As shown in Figure 5, the agentic workflow shifts a larger portion of tasks into higher score bins than E2E, indicating more consistent improvements across tasks rather than gains driven by a small subset. To further inspect subtle workflow differences, we analyze the checklist requirement status composition under the same model. Figure 6 shows that the agentic workflow increases the fraction of correctly satisfied requirements while reducing incorrect ones, suggesting that staging the process as *plan*→*reading*→*write* improves grounding and reduces unsupported completions. Overall, these results suggest that the agentic multi-turn workflow is generally beneficial for long-form synthesis: it helps models better allocate attention over full reference passages, produce more faithful content, and consequently achieve higher checklist pass rates in aggregate.

Concluding Remarks. Our analysis highlights three key insights. First, intrinsic model capacity is foundational, with stronger models consistently prevailing under identical workflows. Second, scaffolding design is pivotal; the substantial gains from the agentic workflow demonstrate that *how* we structure generation is as critical as *which* model is used. Finally, with SOTA scores still below 40%, synthesizing grounded reports from over 100 documents remains a formidable open challenge.

4.3 Case Analysis

To complement checklist pass rates, we conduct a qualitative case analysis to identify recurring failure modes and to understand how workflow design affects post-retrieval synthesis behaviors.

Failure modes under the E2E single-turn workflow. Across low-scoring models, a common pattern is *under-elaboration*: instead of consolidating as many relevant points as possible from the Oracle Context, the model produces a high-level but sparse narrative that omits many checklist items. This behavior disproportionately harms the General score, as required coverage is missed due to insufficient detail and weak synthesis effort. Moreover, under constrained requirements, these models not only omit required organizational elements (e.g., missing categories/attributes) but also more frequently introduce *unsupported statements* when attempting to satisfy constraints, leading to apparent “hallucination”-like outputs (i.e., claims or comparisons not grounded in the provided materials).

5 Conclusion

We introduce **DeepSynth-Eval** to address the lack of objective benchmarks for long-form report generation. Instead of relying on subjective holistic scoring, we ground our evaluation in atomic, verifiable checklists derived from expert surveys. On average, each evaluation task involves 199.7 checklist items, providing a granular lens to quantify the quality of open-ended generation. Beyond the dataset, we establish a foundational evaluation framework by implementing and testing diverse baselines, ranging from standard end-to-end workflows to sophisticated agentic systems. Experiments demonstrate that even with agentic workflows, state-of-the-art (SOTA) models achieve scores of only around 37%, highlighting the challenging nature of this task and the significant room for improvement. Furthermore, this verifiable checklist mechanism paves the way for designing precise reward signals in long-form generation, offering a promising avenue for future reinforcement learning optimization.

Limitations

Currently, our dataset is predominantly composed of computer science surveys, a choice driven by data accessibility and annotation feasibility. However, the proposed checklist-based evaluation

methodology is domain-agnostic and holds the potential to be extended to broader scientific disciplines.

Acknowledgments

We acknowledge the use of Gemini 3 to assist with linguistic polishing and refinement of the manuscript. The original content and ideas were authored by the researchers. All AI-suggested modifications were manually reviewed and verified by the authors to ensure strict consistency with the factual information and original intent.

References

- Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2023. *L-eval: Instituting standardized evaluation for long context language models*. *Preprint*, arXiv:2307.11088.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. *Longbench: A bilingual, multitask benchmark for long context understanding*. *Preprint*, arXiv:2308.14508.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. *Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning*. *Preprint*, arXiv:2501.12948.
- Mingxuan Du, Benfeng Xu, Chiwei Zhu, Xiaorui Wang, and Zhendong Mao. 2025. *Deepresearch bench: A comprehensive benchmark for deep research agents*. *Preprint*, arXiv:2506.11763.
- Google. 2024. Try deep research and our new experimental model in Gemini, your AI assistant. <https://blog.google/products/gemini/google-gemini-deep-research/>. Accessed: 2026-01-01.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. *Measuring massive multitask language understanding*. *Preprint*, arXiv:2009.03300.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekeshe, Fei Jia, and Boris Ginsburg. 2024. *RULER: What’s the real context size of your long-context language models?* In *First Conference on Language Modeling*.
- Minghao Li, Ying Zeng, Zhihao Cheng, Cong Ma, and Kai Jia. 2025. *Reportbench: Evaluating deep research agents via academic survey tasks*. *Preprint*, arXiv:2508.15804.

- Yuan Liang, Jiaxian Li, Yuqing Wang, Piaohong Wang, Motong Tian, Pai Liu, Shuofei Qiao, Runnan Fang, He Zhu, Ge Zhang, Minghao Liu, Yuchen Eleanor Jiang, Ningyu Zhang, and Wangchunshu Zhou. 2025. [Towards personalized deep research: Benchmarks and evaluations](#). *Preprint*, arXiv:2509.25106.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Grégoire Mialon, Clémentine Fourrier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2023. [Gaia: a benchmark for general ai assistants](#). *Preprint*, arXiv:2311.12983.
- OpenAI. 2025. [Introducing deep research](https://openai.com/index/introducing-deep-research/). <https://openai.com/index/introducing-deep-research/>. Accessed: 2026-01-01.
- Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, Michael Choi, Anish Agrawal, Arnav Chopra, Adam Khoja, Ryan Kim, Richard Ren, Jason Hausenloy, Oliver Zhang, Mantas Mazeika, and 1093 others. 2025. [Humanity’s last exam](#). *Preprint*, arXiv:2501.14249.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Driani, Julian Michael, and Samuel R. Bowman. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *Preprint*, arXiv:2311.12022.
- Yijia Shao, Yucheng Jiang, Theodore Kanell, Peter Xu, Omar Khattab, and Monica Lam. 2024. [Assisting in writing Wikipedia-like articles from scratch with large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6252–6278, Mexico City, Mexico. Association for Computational Linguistics.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024a. [Large language models are not fair evaluators](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9440–9450, Bangkok, Thailand. Association for Computational Linguistics.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024b. [Mmlu-pro: A more robust and challenging multi-task language understanding benchmark](#). *Preprint*, arXiv:2406.01574.
- Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. 2025. [Browsecomp: A simple yet challenging benchmark for browsing agents](#). *Preprint*, arXiv:2504.12516.
- Ryan Wong, Jiawei Wang, Junjie Zhao, Li Chen, Yan Gao, Long Zhang, Xuan Zhou, Zuo Wang, Kai Xi-ang, Ge Zhang, Wenhao Huang, Yang Wang, and Ke Wang. 2025. [Widesearch: Benchmarking agentic broad info-seeking](#). *Preprint*, arXiv:2508.07999.
- Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, and Fei Huang. 2025. [Web-Walker: Benchmarking LLMs in web traversal](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10290–10305, Vienna, Austria. Association for Computational Linguistics.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. [ReAct: Synergizing reasoning and acting in language models](#). In *International Conference on Learning Representations (ICLR)*.
- Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Hao, Xu Han, Zhen Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun. 2024. [∞Bench: Extending long context evaluation beyond 100K tokens](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15262–15277, Bangkok, Thailand. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.

A Detailed Workflow Implementation

In this section, we provide a step-by-step textual description of the two synthesis workflows presented in Section 4.1, as illustrated in Figure 7.

A.1 Monolithic E2E Single-turn Workflow

The E2E workflow (Figure 7a) is designed to test the model’s ability to perform long-context consolidation in a single pass. The process consists of two main steps:

1. **Oracle Context Construction:** First, we parallelize the processing of all bibliography entries. Each reference paper d_i is compressed into a dense prose summary s_i using a dedicated summarization prompt. These summaries are then concatenated to form the *Oracle Context*. Crucially, each summary is prefixed with a stable numeric citation identifier (e.g., “[1]”, “[2]”) to allow the model to ground its claims explicitly.
2. **One-Shot Generation:** We construct a unified prompt that combines the reconstructed research request, the extracted constraint questions, and the *Oracle Context*. The model is then instructed to generate the full survey report in a single inference pass. No intermediate reasoning steps or external retrieval actions are permitted; the model must rely entirely on its internal attention mechanism to organize and synthesize the provided summaries.

A.2 Agentic Multi-turn Workflow

The Agentic workflow (Figure 7b) decomposes the synthesis task to mimic the *plan-read-write* process of a human researcher. It operates in three phases to ensure better organization and grounding:

Phase 1: Global Planning (Intellectual Skeleton). Instead of writing immediately, the model acts as an “Architect.” It analyzes the *Oracle Context* (containing summaries) to construct an *Intellectual Skeleton*. This skeleton is a high-level plan that defines the section taxonomy, identifies key comparison axes, and allocates specific content to each section. This plan is then parsed into a structured list of sub-tasks (e.g., JSON objects), where each task corresponds to writing one section of the report.

Phase 2: Serial Iterative Writing Loop. The workflow executes a loop for each sub-task derived

from the skeleton. For each section k , the process involves three specialized steps:

1. **Relevance Selection:** The model evaluates the requirements of the current section against the reference bibliography to identify a subset of critical papers (represented by IDs) that are necessary for drafting this specific part.
2. **Selective Deep Reading:** For the identified critical papers, the agent retrieves their full text from the underlying reference collection. It then generates detailed reading notes focused specifically on the current section’s topic. This step allows the model to access fine-grained details that might have been lost in the initial summaries, without overflowing the context window with all papers at once.
3. **Context-Aware Writing:** The model writes the section. The generation is conditioned on three inputs: (1) the *Intellectual Skeleton* (for structure), (2) the *Deep Reading Notes* (for details), and (3) the *Writing History* (the draft of previous sections). This ensures the new section is logically coherent with the preceding text.

Phase 3: Final Polish. Once all sections are drafted, they are concatenated. An “Editor” agent performs a final pass to ensure global consistency, unify terminology, and format the document before outputting the final report.

B External System Evaluation

The controlled workflows in Table 1 isolate post-retrieval synthesis by holding the Oracle Context fixed. However, practical Deep Research deployments are typically end-to-end: systems actively search, browse, and write in a coupled loop. To contextualize our controlled findings and to understand how checklist-based synthesis evaluation manifests in real deployments, we additionally evaluate representative professional Deep Research systems and an external agentic pipeline. We emphasize that these results serve as complementary reference points rather than the primary benchmark ranking.

Fairness controls. Evaluating external systems introduces two confounds: uncontrolled retrieval may obscure synthesis failures, and the gold survey itself (or derivative pages) may be retrieved, leading to leakage and inflated scores. To mitigate these issues, we provide an explicit allowlist of reference URLs derived from the bibliography and

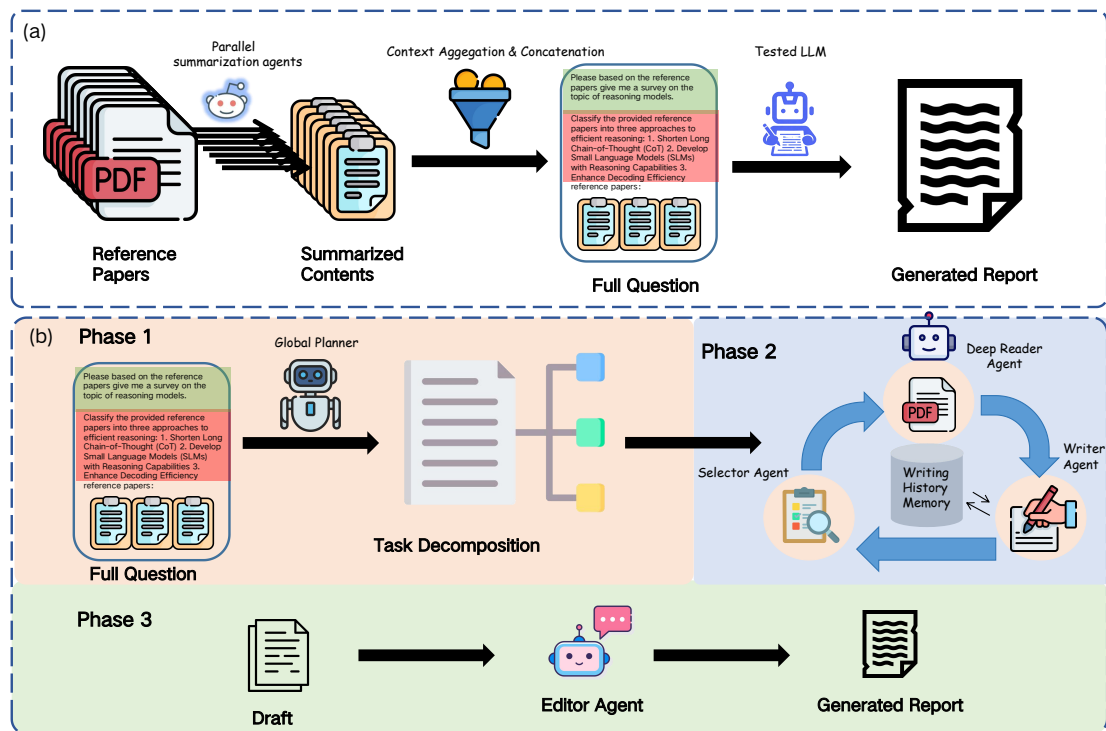


Figure 7: **Schematic overview of the controlled synthesis workflows.** (a) **E2E Single-turn workflow:** A streamlined pipeline that first summarizes all references into an *Oracle Context* with stable citation IDs, then concatenates the request and constraint questions to generate the full report in a single pass via long-context inference. (b) **Agentic multi-turn workflow:** A staged generation process. **Phase 1** generates an *Intellectual Skeleton* to plan taxonomy and content allocation. **Phase 2** executes a loop for each section: the model performs *selective deep reading* by identifying critical papers and consulting their full text from the underlying reference collection to generate specific notes, followed by writing the section conditioned on the skeleton, notes, and previous writing history to ensure grounded organization.

instruct systems to use only these sources for evidence and citations, thereby preventing the original gold survey from being used as a reference. We further impose a strict temporal cutoff: only papers published before the survey date are allowed, and any later materials must be ignored even if surfaced by search.

Systems and evaluation. We evaluate two professional Deep Research systems accessed through their official interfaces (*Doubao DeepResearch*, *OpenAI DeepResearch*) and one external agentic pipeline (*OpenAI Agentic*). We interpret the results as contextual comparisons rather than direct, full-set rankings against our controlled workflows.

Results and discussion. Table 2 summarizes the external-system performance on the evaluation subset. Among the tested systems, *OpenAI DeepResearch* achieves the highest Overall score (28.7) and particularly strong constraint-checklist performance (Constraint: 41.8), with *Doubao DeepRe-*

System	Overall	General	Constraint	Prec. (%)
Doubao DeepResearch	26.4	23.3	38.9	87.4
OpenAI Agentic	25.6	24.0	28.3	83.3
OpenAI DeepResearch	28.7	25.1	41.8	85.4

Table 2: Performance of external systems on a subset of DSE tasks. Scores are checklist pass rates (%). Due to interface and execution constraints, these systems are evaluated on 39 tasks; results are reported separately from the main table and should be interpreted as contextual comparisons.

search showing a similar pattern (Overall: 26.4; Constraint: 38.9). In contrast, the external agentic pipeline (*OpenAI Agentic*) attains a comparable Overall score (25.6) but notably lower Constraint (28.3), suggesting that generic agentic tool use does not automatically translate into stronger constraint-driven organization. We observe that the evaluated end-to-end Deep Research systems do not consistently outperform our controlled workflows, despite having access to browsing and tool

orchestration.

Why external Deep Research systems may underperform our controlled workflows. In our manual inspection, we observe that once the retrieval scope is constrained by the URL allowlist and the temporal cutoff, external Deep Research systems do not consistently utilize the full set of provided references: reports often cite only a subset of the allowlisted papers, and some key references required by the checklists are never retrieved or never appear in the final citations. Correspondingly, the generated reports exhibit more frequent coverage gaps (missing required methods/dataset/attributes), which directly lowers checklist satisfaction.

We hypothesize that this underperformance is driven by two practical factors induced by the fairness constraints. First, restricted retrieval can become *incomplete*: even with an allowlist, the system may fail to open, read, or surface all relevant papers (e.g., due to interface limits, ranking/truncation, or budgeted browsing), causing evidence coverage gaps that propagate to the final synthesis. Second, to prevent leakage and ensure protocol compliance, our prompts impose additional instructions (e.g., strict source restrictions and citation rules), which may introduce an instruction-following burden that reduces flexibility and consumes generation capacity, leading to more conservative or less detailed synthesis.

Together, these observations highlight a key advantage of DSE’s controlled Oracle-Context evaluation: it isolates consolidation and organization ability from retrieval completeness and interface-induced instruction overhead, enabling a more faithful measurement of post-retrieval synthesis.

C Notation Summary

Table 3 summarizes the main notations used throughout benchmark construction, synthesis workflows, and evaluation.

D Prompt Templates

D.1 Unified Prompt P

Figure 8 shows the unified prompt template used to formulate the survey-writing request for each task.

D.2 E2E Single-turn Workflow Prompt

Figure 9 shows the prompt template used in the end-to-end (E2E) single-turn workflow. The model

is provided with summaries of all reference papers together with the unified task prompt P , and generates the survey in a single final call.

D.3 Agentic Workflow Prompts

For brevity, we group the agentic workflow into three stages: planning, section writing, and final polishing. Internally, the planning stage includes both survey skeleton construction and task parsing, while the section-writing stage includes evidence selection, detailed reading, and section drafting.

D.3.1 Planning

Figure 10 shows the prompt template used in the planning stage. Given summaries of all reference papers, the model first constructs the intellectual skeleton of the survey and then parses it into section-level writing tasks.

D.3.2 Section Writing

Figure 11 shows the prompt template used in the section-writing stage. For each section, the model identifies the most relevant references, gathers section-specific reading notes, and writes the section grounded in the selected evidence.

D.3.3 Final Polishing

Figure 12 shows the prompt template used in the final polishing stage. After all sections are drafted, the model refines the full survey into a coherent final report.

E Case study of GPT 5.2

Notation	Type	Description
$T = (P, D, C)$	Task	A benchmark task in DeepSynth-Eval, consisting of the unified prompt, Oracle Context, and evaluation checklist set.
R	Survey	The reference survey paper used as the gold-standard synthesis target.
q_{gen}	Prompt component	The general question reconstructed from the reference survey.
$q_{\text{con}}^{(m)}$	Prompt component	The m -th constraint question that specifies an organizational or structural requirement.
Q_{con}	= Prompt set	The set of all constraint questions associated with a task.
$\{q_{\text{con}}^{(1)}, \dots, q_{\text{con}}^{(M)}\}$		
P	Prompt	The Unified Input Prompt formed by concatenating q_{gen} with all constraint questions in Q_{con} .
d_i	Paper	The i -th reference paper in the bibliography.
s_i	Summary	The summary generated for reference paper d_i .
D	Context	The Oracle Context constructed from the reference bibliography.
M	Model	The tested model used for report generation.
$\hat{R} = M(P, D)$	Output	The report generated by the tested model.
$C_{\text{gen},k}$	Checklist group	The k -th general checklist group.
$C_{\text{con},m}$	Checklist group	The m -th constraint checklist group.
C_{gen}	= Checklist set	The set of all general checklist items.
$\bigcup_{k=1}^K C_{\text{gen},k}$		
C_{con}	= Checklist set	The set of all constraint checklist items.
$\bigcup_{m=1}^M C_{\text{con},m}$		
$C = C_{\text{gen}} \cup C_{\text{con}}$	Checklist set	The complete evaluation checklist set.
$c \in C$	Checklist item	A single checklist item to be verified.
$r(c)$	Verification label	The verification result for checklist item c , where 1 denotes Correct, 0 denotes Omission, and -1 denotes Incorrect.
G_k	Group	A generic checklist group in the scoring formulation.
$N_k = G_k $	Count	The number of checklist items in group G_k .
$S_k = \sum_{c \in G_k} r(c)$	Group reward	The cumulative reward of checklist group G_k .
θ_k	Threshold	The saturation threshold used in fault-tolerant group scoring.
Score_k	Group score	The normalized score of checklist group G_k .
S_{gen}	Metric	The average score over all general checklist groups.
S_{con}	Metric	The average score over all constraint checklist groups.
S_{all}	Metric	The overall score across all checklist groups.
N_{Correct}	Count	The number of checklist items judged as correct.
$N_{\text{Incorrect}}$	Count	The number of checklist items judged as incorrect.
Precision	Metric	The precision over mentioned checklist items, defined as $\frac{N_{\text{Correct}}}{N_{\text{Correct}} + N_{\text{Incorrect}}}$.

Table 3: Summary of the main notations used in DeepSynth-Eval.




Figure 8: Unified prompt template P

=== System Prompt ===

You are a rigorous research assistant and deep analysis expert, skilled in reading academic papers, writing literature reviews, and designing research plans. Your task is to help the user analyze multiple academic papers and produce a systematic and comprehensive survey article. All of your reasoning must be in-depth and detailed, and strictly based on the materials provided by the user; you must not fabricate information or over-interpret the content.

Let us start with the first step: based on the given reference materials and survey requirements, systematically analyze the existing work and think deeply about how this survey should be organized and how to present these works, thereby constructing the “intellectual skeleton” of the entire article. Pay attention to choosing an appropriate level of granularity for the intellectual skeleton: it should be sufficiently instructive to provide guidance, but it does not need to cover every detail.

And then write a survey article in analytical prose. In each paragraph, avoid using bullet points or lists; instead, use continuous text to ensure fluidity. - Roughly follow the organizational structure discussed earlier, while allowing flexible adjustments in detail according to the actual content. - Don't list models and results one by one, but explicitly group them into clear categories, highlight their common patterns and key differences, and state what these comparisons imply for strengths, weaknesses, or design trade-offs. - Do not list the method performance in the paragraph. Provide result in table for comparison if it is necessary. - When referring to specific papers or groups of papers, use bracketed numerical citations such as [1], [2], or [1,3], and strictly follow the numbering or identifiers provided by the user. - Do not introduce any conclusions or facts that are not supported by the given materials.

=== User Prompt ===

<The start of 1 Paper's summary.>

Cite it as [1]

{summary_1}

<The End of 1 Paper's summary.>

<The start of 2 Paper's summary.>

Cite it as [2]

{summary_2}

<The End of 2 Paper's summary.>

...

The above are summaries of all the reference articles.

Please help me generate an in-depth survey report following the instruct:

Main Question: {general_prompt}

Detailed Constraints:

1. {constraint_1}

2. {constraint_2}

...

M. {constraint_M}

Figure 9: Prompt template for the E2E single-turn workflow

=== System Prompt ===

You are a rigorous research assistant and deep analysis expert, skilled in reading academic papers, writing literature reviews, and designing research plans. Your task is to help the user analyze multiple academic papers and produce a systematic and comprehensive survey article. The designed sections would be written by each writer or sub-agent. Papers in latest years should be paid more attention.

=== User Prompt ===

{all_reference_summaries}

The above are summaries of all the reference articles. Firstly, systematically analyze the existing work and think deeply about how this survey should be organized and how to present these works, thereby constructing the “intellectual skeleton” of the entire article.

Then parse the resulting survey skeleton into a JSON list of section-level writing tasks. Each task should include:

- background: the section context and writing goal
- task: the concrete writing instruction for that section

Figure 10: Prompt template for the planning stage in the agentic workflow

=== User Prompt ===

Current Phase: Writing Section {section_id}

[Task Background]
{section_background}

[Writing Instruction]
{section_task}

[Reference Selection]
Identify the papers that are critical for this section. These selected papers will be read in more detail and used as the primary evidence for writing.

[Gathered Information]
{section_specific_notes}

Now write a section in analytical prose. In each paragraph, avoid using bullet points or lists; instead, use continuous text to ensure fluidity.

- Roughly follow the planned organizational structure, while allowing flexible adjustments according to the actual content.
- Do not list models and results one by one; instead, group them into clear categories, highlight common patterns and key differences, and explain the resulting strengths, weaknesses, or trade-offs.
- When referring to papers, use bracketed numerical citations such as [1], [2], or [1,3].
- Do not introduce any conclusions or facts that are not supported by the provided materials.

[Requirements]

1. Output only the content for this specific section.
2. Use Markdown formatting (e.g., ## for headers).
3. Ensure strict academic tone and precise citation.
4. Maintain logical flow with previously written sections.

Figure 11: Prompt template for the section-writing stage in the agentic workflow

=== User Prompt ===

You are the Editor-in-Chief of a prestigious AI journal. You have been provided with a full draft of a survey paper generated section by section.

[Start of full draft]

{full_draft}

[End of full draft]

Your task is to transform this draft into a polished, high-quality survey report.

Please perform the following operations:

1. Write a concise abstract that captures the essence of the survey.
2. Ensure global consistency in terminology and style.
3. Refine the Markdown structure so that headers are logical and professional.
4. If a title is missing, generate a professional and descriptive academic title.

Format the output as a complete, standalone academic report. Keep all citations [n] intact.

Figure 12: Prompt template for the final polishing stage in the agentic workflow

Retrieval-Augmented Generation: Foundations, Methods, Multimodal Extensions, Systems, and Evaluation

1. Introduction and Formalization

Retrieval-Augmented Generation (RAG) provides a principled way to couple the parametric knowledge stored inside large language models with non-parametric, dynamically updatable external knowledge through retrieval and context curation. We formalize a generic RAG pipeline as $y = G(x, C(R(x, K)))$, where x denotes the user query, R is a retriever that selects evidence from a knowledge base K , C is a context curation function that filters, orders, compresses, or structures the retrieved content for the generator, and G is the generator that produces the final output conditioned on both x and curated context. This formulation encompasses common operational regimes, including open-book and closed-book usage; single-hop and multi-hop question answering; and iterative procedures in which retrieval and generation interleave or recur. Treating retrieval outputs as latent variables that are marginalized at decoding time connects RAG to probabilistic formulations introduced in early hybrid models, while instantiations like Fusion-in-Decoder characterize specific architectural choices for integrating multiple retrieved passages into sequence-to-sequence generation [133, 125].

...

To guide readers, we include three artifacts. A taxonomy figure maps methods across the axes of retrieval type, integration strategy, adaptivity, efficiency, robustness, modality, structure, and domain, clarifying how designs cluster and where gaps remain [225, 283]. A pipeline schematic contrasts retrieve-then-generate, generate-then-retrieve, and iterative synergy workflows, highlighting where retriever training, query refinement, and context curation intervene [133, 84, 211]. A timeline from 2020 to 2024 situates key milestones—dense retrieval and REALM, probabilistic RAG and Fusion-in-Decoder, black-box and in-context augmentation, few-shot Atlas, and neural indexing—underscoring the rapid maturation and recent emphasis on adaptivity, risk certification, and multimodal expansion [165, 198, 133, 125, 77, 37, 205, 98, 20, 80]. Within this scope, the survey contributes a unified formalization, a structured taxonomy, and a synthesis of design trade-offs and evaluation constructs, with emphasis on recent advances that define the current frontier.

2. Retrieval Foundations

[Content unchanged from draft Section II.]

3. Integration Strategies and Generation Architectures

[Content unchanged from draft Section III.]

4. Robustness, Hallucination Mitigation, and Risk

[Content unchanged from draft Section IV.]

5. Efficiency, Latency, and System Co-Design

[Content unchanged from draft Section V.]

6. Multimodal RAG (Images, Audio, Video)

[Content unchanged from draft Section VI.]

7. RAG Meets Diffusion Models and Generative Modeling

Diffusion and score-based generative models provide the backbone on which retrieval-augmented conditioning can act, offering a spectrum of formulations that trade sampling speed, likelihood, and controllability. The core DDPM framework establishes forward and reverse Markov chains with a simplified noise-prediction objective that yields high-fidelity images and interpretable denoising trajectories [26]. Subsequent refinements improve both training and sampling: learned reverse-process variances with a hybrid variational objective and cosine noise schedules enable near-optimal sampling in tens of steps and competitive likelihoods, markedly improving practical efficiency [5]; the SDE-based unification generalizes DDPMs and denoising score matching...