

# ReportLogic: Evaluating Logical Quality in Deep Research Reports

Jujia Zhao<sup>1,2</sup>, Zhaoxin Huan<sup>2</sup>, Zihan Wang<sup>3</sup>, Xiaolu Zhang<sup>2</sup>, Jun Zhou<sup>2\*</sup>, Suzan Verberne<sup>1</sup>, Zhaochun Ren<sup>1\*</sup>

<sup>1</sup>Leiden Institute of Advanced Computer Science, Leiden University

<sup>2</sup>Ant Group, <sup>3</sup>CISPA Helmholtz Center for Information Security

{zhao.jujia.0913, zhw.cypher}@gmail.com

{zhaoxin.hzx, yueyin.zxl, jun.zhoujun}@antfin.com

{s.verberne, z.ren}@liacs.leidenuniv.nl

## Abstract

Users increasingly rely on Large Language Models (LLMs) for Deep Research, using them to synthesize diverse sources into structured reports that support understanding and action. In this context, the practical reliability of such reports hinges on *logical quality*: whether the report’s claims and arguments are explicitly supported and can be trusted as a basis for downstream use, rather than merely appearing fluent or informative. However, current evaluation frameworks largely overlook this requirement. To bridge this gap, we introduce **ReportLogic**, a benchmark that quantifies report-level logical quality through a reader-centric lens of *auditability*. Specifically, ReportLogic adopts a hierarchical taxonomy that evaluates whether readers can (1) trace an on-topic report structure with a unified analytical arc (*Macro-Logic*), (2) understand the progression with necessary context (*Expositional-Logic*), and (3) verify conclusions via explicit claim–support (*Structural-Logic*). Based on this taxonomy, we construct a human-annotated rubric-guided dataset and train an open-source LogicJudge for scalable evaluation. We further evaluate judge robustness via adversarial attacks, showing that off-the-shelf LLM judges are frequently influenced by superficial cues (e.g., verbosity), and reasoning modes can mask broken support relations. Overall, our results provide actionable guidance for building more robust logic evaluators and improving the logical reliability of LLM-generated reports.

## 1 Introduction

Users increasingly rely on Large Language Models (LLMs) to conduct Deep Research: gathering information from diverse sources, synthesizing it into structured analysis, and producing reports that support understanding and downstream

\*Corresponding author. This work was supported by Ant Group Research Intern Program.

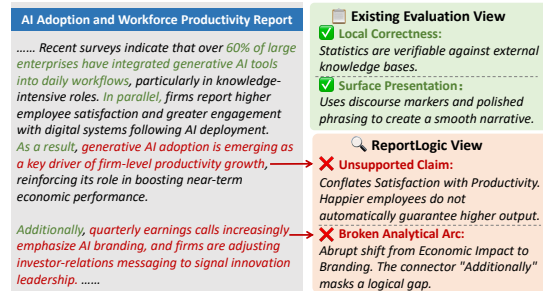


Figure 1: Comparison between existing evaluation views and ReportLogic on a Deep Research report.

action (Shi et al., 2025; Huang et al., 2025). In this context, the utility of the output report extends beyond producing correct facts or well-formed passages. For example, a report may cite accurate statistics and remain well written, yet still be unhelpful when central claims are not adequately supported, leaving readers unable to verify how the conclusion is derived (Mökander et al., 2021). Consequently, the value of Deep Research hinges on *logical quality*: the capacity to organize information into a unified report-level line of analysis, where the progression of ideas is cohesive and every claim is rigorously grounded. Without such logical quality, the narrative becomes disjointed and credibility is compromised, fundamentally undermining user trust for downstream use.

Despite its critical importance, logical quality is still under-specified and weakly measured in current evaluation frameworks for Deep Research and broader long-form generation. First, current evaluation predominantly prioritizes local correctness (atomic factual accuracy) and surface presentation (fluency, grammar, and discourse markers), focusing on fine-grained factuality checks (Min et al., 2023), verified citations (Sun et al., 2024), and transition smoothness (Du et al., 2025). However, these metrics do not assess a report’s analytical arc or support relations, i.e., whether the report develops coherently and whether major claims are explicitly supported to justify a well-founded con-

clusion (Figure 1). Second, while benchmarks like LogicBench (Parmar et al., 2024) evaluate formal logical reasoning, they focus on formal logic puzzles (e.g., symbolic inference rules) and are not designed to evaluate long-form analytical writing required in Deep Research reports.

Evaluating logical quality in Deep Research poses two fundamental challenges. (1) Logical quality is inherently multidimensional. Failures can occur at different granularities, including misalignment with user intent at the organizational level, gaps in discourse that hinder comprehension, or unsupported key claims. Because these failures operate at distinct scopes, existing holistic scoring often conflates them and provides limited diagnostic feedback. (2) Logical evaluation lacks a universal decision boundary. Unlike factuality or citation checks with relatively fixed criteria, judgments of logical sufficiency depend on the specific research question and available context, such as how much evidence is required or which reasoning steps must be explicit. Together, these challenges motivate a rubric-guided, fine-grained evaluation protocol that translates logical requirements into explicit, instance-specific inspection items.

To bridge the above gap, we introduce **ReportLogic**, a new evaluation benchmark for assessing logical quality in Deep Research reports. To enable fine-grained diagnosis across granularities, we move beyond treating logical quality as a monolithic property and operationalize it through a reader-centric lens of *auditability* (Doshi-Velez and Kim, 2017): whether a reader can efficiently trace, understand, and verify the analytical process. Accordingly, ReportLogic adopts a hierarchical taxonomy aligned with these requirements: (1) **Macro-Logic (Traceability via Organization)**. This layer assesses whether the report stays on topic and assigns clear functional roles to sections, forming a unified analytical arc rather than disjointed points. (2) **Expositional-Logic (Understandability via Flow)**. Inspired by Information Flow Theory (Halliday and Matthiessen, 2013), this layer assesses whether the narrative is easy to follow, e.g., by introducing necessary background before abstract analysis and maintaining a coherent Given→New progression. (3) **Structural-Logic (Verifiability via Argumentation)**. Drawing on the Toulmin model (Toulmin, 2003), this layer assesses whether key claims are grounded in relevant evidence and whether war-rants linking evidence to conclusions are explicit,

with appropriate qualifiers. We further instantiate this taxonomy as eight fine-grained dimensions for diagnostic evaluation. To make the decision boundary clear for each instance, we implement a context-aware rubric generation protocol that translates each dimension into instance-specific inspection items.

Based on this protocol, we construct a human-annotated benchmark dataset from Deep Research queries drawn from diverse sources, where trained annotators provide dimension-level preferences along with an overall verdict. To enable scalable evaluation, we further train an open-source **Logic-Judge** to assess logical quality in Deep Research reports and enable community-wide comparison. We also evaluate judge robustness via adversarial attacks on ReportLogic, finding that off-the-shelf LLM judges are susceptible to superficial cues (e.g., verbosity), and reasoning-optimized judges may mask broken support relations. Together, these findings provide actionable guidance for developing more resilient logic judges and improving the logical reliability of LLM-generated long-form reports.

Our contributions are threefold: (1) We introduce ReportLogic, the first human-annotated benchmark for evaluating logical quality in Deep Research, featuring fine-grained dimensions, context-aware instance-specific rubrics, and preference-based annotations. (2) We propose and train a specialized LogicJudge for scalable evaluation, achieving stronger alignment with human preferences than off-the-shelf LLM judges on ReportLogic. (3) We perform a systematic adversarial robustness analysis of LLM judges on ReportLogic, uncovering common failure modes and providing fine-grained diagnostic insights into the limitations of current logic judges.

## 2 Related Work

### 2.1 Evaluation of Long-form Report Generation

With advances in long-context modeling (Wu et al., 2024), evaluation has moved beyond sentence-level overlap metrics such as BLEU and ROUGE toward benchmarks designed for long-form report generation. In Deep Research and related long-form reporting scenarios, existing protocols primarily target three content aspects: lo-

---

Our code and data are available at <https://github.com/Polaris-JZ/ReportLogic>.

cal correctness, including statement-level factuality (Min et al., 2023) and citation verification (Gao et al., 2023); surface presentation, including transition smoothness (Cao et al., 2024); and content coverage, including key-point recall (Wei et al., 2024). While effective for correctness and readability, these criteria do not explicitly evaluate report-level logical quality, i.e., whether a report sustains a coherent line of analysis and supports its conclusions through explicit claim–support relations.

**Rubric-Guided Evaluation.** To operationalize the above aspects on open-ended outputs, most recent long-form benchmarks adopt a rubric- or checklist-guided protocol that decomposes holistic quality into a fixed set of sub-judgments. HelloBench evaluates writing quality and fluency via checklist-style LLM-as-a-judge prompts (Que et al., 2024); ReportBench grounds atomic factuality in multi-model voting with web-connected LLMs and citation-based verification (Li et al., 2025); and long-form factuality evaluation relies on atomic claim schemas for fact-checking (Wei et al., 2024; Min et al., 2023). A common property of these protocols is that they apply a *fixed global rubric*, i.e., a shared set of criteria uniformly instantiated across all instances, which offers comparability and reproducibility. For report-level logical quality, however, this fixed-schema design leaves the decision boundary underspecified on open-domain queries. In contrast, we adopt an instance-specific rubric conditioned on the query and the compared reports.

**LLM-as-Judge and Judge Training.** Regardless of how rubrics are designed, their practical utility depends on the evaluator that applies them. Using LLMs as evaluators has become a standard practice for open-ended generation (Kocmi and Federmann, 2023), but off-the-shelf LLM judges are known to exhibit biases related to length, position, and self-preference (Zheng et al., 2023), motivating specialized judge training. Representative paradigms include supervised fine-tuning on preference data (Kim et al., 2023), reinforcement learning with judge-specific rewards (Mahan et al., 2024), and critique-oriented training that elicits structured rationales before a verdict (Liu et al., 2025). Most existing trained judges target general-purpose preference or short-form evaluation, rather than long-form, rubric-conditioned logical assessment. Our LogicJudge is devel-

oped for this latter setting: a domain-specialized judge trained on ReportLogic supervision with a rubric-grounded, dimension-wise objective, producing dimension-level diagnostic signals aligned with the benchmark’s taxonomy.

## 2.2 Logical Evaluation and Argumentation

Prior work related to logic in NLP can be broadly grouped into formal logical reasoning benchmarks and argumentation mining. Benchmarks such as LogicBench (Parmar et al., 2024), FOLIO (Han et al., 2024), and JustLogic (Chen et al., 2025) evaluate puzzle-style deductive reasoning in constrained settings with explicit premises and pre-defined conclusions. While valuable, they are not designed for open-ended, report-level analytical writing, where quality depends on organization, discourse flow, and sustained claim–support structure across multiple paragraphs. Argumentation mining studies how argumentative components and relations are expressed in human-written text (Lawrence and Reed, 2019), typically focusing on extracting argument structure at the sentence or argument level (Wachsmuth et al., 2017). It does not typically operationalize logical quality as a multidimensional, report-level property, nor offer an end-to-end evaluation protocol that jointly assesses organization, information flow, and evidential support in generated reports.

## 3 The ReportLogic Benchmark

To systematically evaluate the logical quality in Deep Research reports, we propose **ReportLogic**. Specifically, we focus on **Deep Research** scenarios, where an LLM synthesizes information from diverse sources into a long-form analytical report in response to an open-domain query. Each instance comprises (i) a *report-appropriate* query  $q$  that requires multi-aspect analysis, and (ii) retrieved contexts  $\mathcal{D} = \{d_1, \dots, d_k\}$  returned by a web search pipeline. Given  $(q, \mathcal{D})$ , the model generates a report  $y$ . ReportLogic evaluates the internal logical quality of the report, rather than external factual correctness or surface writing quality.

### 3.1 Evaluation Taxonomy

As noted in Section 1, logical quality in Deep Research is difficult to evaluate because it is inherently multidimensional. To bridge this, we define logical quality in terms of *auditability* (Doshi-Velez and Kim, 2017), namely the extent to which

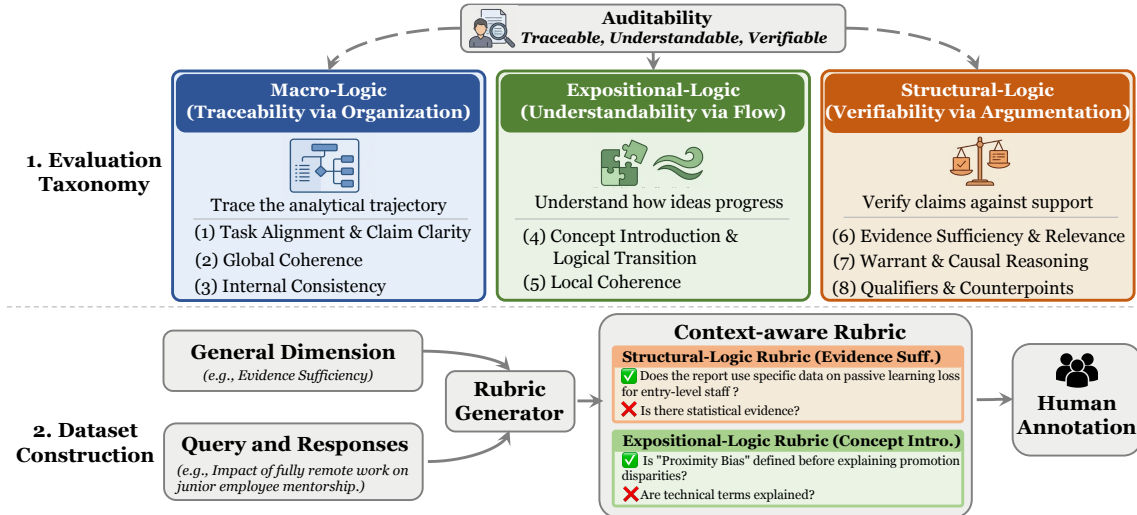


Figure 2: ReportLogic framework. We define logical quality as auditability and decompose it into a three-layer taxonomy with eight dimensions. Given a query and paired reports, a rubric generator instantiates each dimension into context-aware inspection items that guide pairwise human annotation.

a reader can trace, understand, and verify the analytical process. We adopt auditability as the organizing principle because Deep Research reports are consumed as decision-support artifacts: a logically reliable report is one whose overall analytical trajectory a reader can *trace*, whose intermediate steps a reader can *understand* without implicit gaps, and whose conclusions a reader can *verify* against explicit reasoning and evidence. These three reader actions directly induce our top-level taxonomy, providing an actionable grounding for logical evaluation that abstract logic ontologies struggle to supply at the report level. As Figure 2 shows, we operationalize auditability through a hierarchical taxonomy that decomposes report-level logic into complementary constraints at different granularities, yielding eight fine-grained dimensions.

**Macro-Logic** (Traceability via Report-Level Organization). Macro-Logic supports auditability at the document scale by providing a clear map of the report’s analytical structure. A report exhibits strong Macro-Logic if readers can readily trace its central stance and understand how sections contribute to that stance. This layer includes (1) **Task Alignment & Claim Clarity**, which assesses whether the report commits to a clear central position that directly addresses the query; (2) **Global Coherence**, which examines whether sections have well-defined functional roles that support the overall argument; and (3) **Internal Consistency**, which checks whether definitions, premises, and quantitative statements remain

compatible across the report.

**Expositional-Logic** (Understandability via Discourse-Level Explicitness). Auditability further requires that the progression of ideas remains *understandable*. A report may possess a coherent high-level organization yet fail logically if it forces the reader to speculate on missing context or bridge implicit gaps between statements. Guided by Information Flow Theory and the Given-to-New principle (Halliday and Matthiessen, 2013), Expositional-Logic evaluates discourse-level explicitness, focusing on (4) **Concept Introduction & Logical Transition**, which checks whether new concepts or assumptions are adequately motivated before use, and (5) **Local Coherence**, which examines whether adjacent units advance the argument through explicit step-by-step relations rather than rhetorical jumps.

**Structural-Logic** (Verifiability via Claim-Support Structure). Finally, the logical validity of the report depends on the verifiability of the argument itself. Even a well-structured and easy-to-follow report fails logically if its conclusions rest on weak evidence or hidden inference rules. Structural-Logic assesses whether major conclusions are grounded in stated support and whether the inferential links from that support to the conclusions are explicit enough to justify the claimed strength. We ground this layer in Toulmin’s model of practical argumentation, which decomposes an argument into Claim, Grounds, and Warrant, optionally with Qualifiers

and Rebuttals, providing a concrete schema for verifying completeness and validity of support chains in natural language (Toulmin, 2003). It includes (6) **Evidence Sufficiency & Relevance**, which checks whether major claims are supported by specific and probative evidence; (7) **Warrants & Causal Reasoning**, which examines whether inferential bridges are made explicit and causality is used appropriately; and (8) **Qualifiers & Counterpoints**, which evaluates whether conclusions are properly scoped through uncertainty, boundary conditions, or salient alternatives.

### 3.2 Dataset Construction

To instantiate our taxonomy, we construct **Report-Logic**, a human-annotated dataset for evaluating the logical quality of model-generated Deep Research reports.

**Data Collection.** To obtain sufficient coverage beyond the queries provided in existing Deep Research benchmarks (Xu et al., 2025; Du et al., 2025), we additionally source open-domain queries from *Quora* and *Zhihu*. We use GPT-5 to filter and retain only report-appropriate queries, namely prompts that require multi-aspect synthesis and structured analysis rather than short factual replies (Figure 8 in Appendix G). For each retained query, we retrieve supporting context through a web search and passage extraction pipeline, and generate analytical reports using DeepSeek-V3, Claude-4-Sonnet, GPT-4o, and Qwen-Max, with standardized generation prompts in Figure 9 in Appendix G. We choose these models as representative systems from different model families and training paradigms, enabling Report-Logic to capture diverse writing and argumentation behaviors under a unified evaluation protocol.

**Context-aware Rubric Generation.** To address the lack of a universal decision boundary described in Section 1, we instantiate our taxonomy into context-aware rubrics. As illustrated in Figure 2, our rubric generator takes as input (i) the target dimension definition, and (ii) the instance context, including the Deep Research query and the paired candidate reports to be compared. We employ Claude-4.5-Sonnet as the rubric generator to translate each abstract dimension into instance-specific inspection items (Figure 10 in

---

*Quora* and *Zhihu* are prominent community-driven Q&A platforms known for hosting high-level, open-ended discussions rather than simple factoid queries.

Appendix G). Concretely, for each instance and dimension, it produces (1) a targeted comparison question, (2) span-level cues indicating where to inspect in the reports, and (3) paired good/bad examples that clarify the intended standard in this specific context. This content-grounded instantiation reduces ambiguity and yields more accurate supervision relative to context-agnostic dimension descriptions, as validated in Section 5.4. Implementation details of rubric generator selection, annotator-side verification, and schema compliance are reported in Appendix C.2.

**Human Annotation.** To obtain reliable preference supervision, we adhere to a multi-annotator adjudication protocol where three trained experts independently evaluate reports using these generated rubrics, with final preference labels determined by majority vote. Annotator details are provided in Appendix E.

## 4 Method: LogicJudge

While ReportLogic enables fine-grained human assessment, fully manual evaluation is expensive and slows iteration. To support scalable benchmarking, model development, and a community-facing leaderboard, we train **LogicJudge**, an open-source and lightweight judge that predicts pairwise logical preference under our eight-dimensional rubric. LogicJudge is designed to (i) align with human judgments of logical quality rather than surface cues, (ii) provide actionable, dimension-level diagnostics, and (iii) remain deployable as an automated evaluation component for future research.

### 4.1 Task Formulation

We formulate logical evaluation as a rubric-guided pairwise preference prediction task. Given a query  $q$ , context-aware rubrics  $\mathcal{R}$ , and two candidate reports  $(y_A, y_B)$ , the judge outputs a structured reasoning chain followed by a final verdict. The model critiques each of the eight dimensions before predicting the overall preference  $d_{\text{all}} \in \{A > B, A < B, \text{Tie}\}$ . Our label space includes Ties to filter out noise from indistinguishable pairs. This design mitigates the Halo Effect (Kocmi and Federmann, 2023) by grounding the final verdict in specific logical evidence rather than general impressions or length bias, and provides actionable, fine-grained diagnostics rather than a scalar score.

## 4.2 Distilled Training Data

Constructing a large-scale logic training corpus is challenging due to the high cognitive load of manual annotation. To overcome this, we synthesize supervision signals via an ensemble distillation protocol using three frontier models: o3, GPT-5, and Gemini 2.5 Pro. The generation prompts are provided in Figure 11 in Appendix G. We prioritize label precision over recall through a Dual-Stage Filtering Protocol: (1) Consensus Filtering: We retain an instance if and only if all three teacher models reach a unanimous verdict. This strict consensus filters out subjective ambiguity and noise. (2) Swap-Consistency Filtering: For surviving pairs, we swap the order of reports  $(y_B, y_A)$  and re-evaluate. We discard instances where the preference does not logically invert. The final dataset consists of tuples  $(q, \mathcal{R}, y_A, y_B, y^*)$ , where the reasoning trace  $y^*$  is sampled from one of the teacher models. To facilitate parsing, we constrain the output to a strict JSON schema (schema details in Appendix A.1).

## 4.3 Training Procedure

We adopt a two-stage alignment curriculum to balance schema validity and discriminative accuracy. (1) **Supervised Fine-Tuning (SFT)**. LogicJudge is first initialized via SFT on the distilled training corpus to enforce schema alignment and rubric grounding. This stage trains the model to produce complete, machine-parseable evaluations aligned with the eight-dimensional taxonomy, rather than generic or surface-level critiques. (2) **Group Relative Policy Optimization (GRPO)**. Although SFT ensures stable formatting and rubric adherence, it may not fully capture decision boundaries for hard negative response pairs. We therefore apply GRPO (Shao et al., 2024) with a hierarchical reward to sharpen discriminative preference prediction while preserving strict output structure. To prevent schema collapse, format validity is enforced as a hard constraint, while label accuracy is treated as a soft incentive. This *Format-First, Logic-Second* principle enables exploration of reasoning paths without violating the required output schema. Full objective definitions and reward details are provided in Appendix A.2.

## 5 Experiments

We conduct a comprehensive evaluation to address three questions: **RQ1**: Does LogicJudge re-

liably align with expert judgments of logical quality? **RQ2**: How do frontier LLMs perform with ReportLogic, and which logical deficits are most prevalent? **RQ3**: How robust are LLM judges to targeted logical perturbations and surface-level bias attacks?

### 5.1 Experimental Setup

**Datasets.** We use three datasets: DeepResearch (1,204 queries derived from established deep research benchmark (Xu et al., 2025; Du et al., 2025)), Zhihu (1,262 queries collected from a professional Chinese Q&A community), and Quora (1,198 queries sampled from a large English community-driven Q&A forum). Dataset statistics are detailed in Appendix B.1.

**Models and baselines.** We evaluate LogicJudge against 16 state-of-the-art LLMs across major families, including GPT, Claude, Gemini, Qwen, and DeepSeek, covering both instruction-tuned and reasoning-enhanced variants. In addition, we include two inference-time ensemble baselines that aggregate the judgments of the same three frontier LLM judges used in our distillation setup: Ensemble Consensus, which outputs a label only when all three judges agree (otherwise abstaining), and Ensemble Vote, which applies majority voting. All judge baselines are evaluated in a one-shot setting using the same rubric definitions and prompt format as LogicJudge to ensure fairness. Full baseline configurations are provided in Appendix B.2.

**Implementation Details.** We initialize LogicJudge with Qwen-3-30B-A3B and train it via a two-stage pipeline (SFT followed by GRPO) using OpenRLHF. Hyperparameters and training dynamics are provided in Appendix B.3.

### 5.2 Judge Performance Analysis (RQ1)

We first assess the reliability of **LogicJudge** as a surrogate for human experts. To control for position bias, we report Agreement Accuracy. For a test pair  $(y_A, y_B)$  with ground truth label  $l^*$ , a prediction is counted as correct only if the model identifies the same winner as the human judge under both the original  $(y_A, y_B)$  and swapped  $(y_B, y_A)$  orders. Table 1 reports results across all

<https://www.zhihu.com/>.  
<https://www.quora.com/>.  
<https://huggingface.co/Qwen/Qwen3-30B-A3B/>.  
<https://github.com/OpenRLHF/OpenRLHF/>.

Judge Model	DeepResearch	Zhihu	Quora
gpt-o3	63.73	71.70	69.67
gpt-4.1	68.63	70.77	34.00
gpt-5	62.75	71.60	70.33
gpt-5.1	61.76	65.00	66.00
gemini-2.5-pro	68.63	<u>74.00</u>	70.00
gemini-2.5-pro-think	65.69	41.90	68.00
gemini-3-pro	54.90	68.20	62.33
gemini-3-pro-think	59.80	69.60	65.00
claude-4-sonnet	46.08	56.70	51.67
claude-4-sonnet-think	45.10	56.20	51.33
claude-4.5-sonnet	61.76	64.00	68.00
claude-4.5-sonnet-think	63.73	64.50	66.00
qwen3-max	<u>73.53</u>	71.70	69.67
qwen-max	65.69	61.62	40.67
qwen3-235b	43.14	49.20	53.67
deepseek-v3	62.75	66.41	53.33
Ensemble Consensus	49.02	61.11	56.67
Ensemble Vote	66.67	73.57	<u>71.31</u>
<b>LogicJudge (ours)</b>	<b>74.50</b>	<b>75.00</b>	<b>73.00</b>

Table 1: Judge agreement accuracy (%) on Report-Logic across three domains. The best and second-best results are highlighted in bold and underlined fonts.

three datasets, from which we draw three key findings. First, LogicJudge achieves the highest agreement with human experts across all domains, outperforming not only strong single-model judges (e.g., GPT-5 and o3) but also the ensemble baselines. Second, many baselines exhibit large cross-domain variance, with sharp drops on specific datasets, indicating sensitivity to writing style and distribution shift rather than logical quality. This highlights the need for domain-robust evaluation trained under a fixed logical rubric. Third, enabling “thinking” modes does not consistently improve judgment and can even degrade performance on some domains, suggesting that additional test-time compute introduces noise without task-specific tuning.

### 5.3 Leaderboard Analysis (RQ2)

After validating LogicJudge as a reliable surrogate evaluator, we use it to benchmark the report-level logical quality of 16 state-of-the-art LLMs for analytical report generation. Figure 3 summarizes dimension-wise win rates across three domains. We highlight three main findings: (1) Cross-domain stability varies across models. Gemini-2.5-Pro and GPT-5 exhibit relatively stable performance across DeepResearch, Zhihu, and Quora, indicating that their report-level logic generalizes beyond domain-specific writing styles. By con-

Annotation Setting	Fleiss’ Kappa ( $\kappa$ )	Pairwise Agmt.
No-Rubric	0.37	70.11%
General-Rubric	0.67	78.73%
<b>Context-Aware</b>	<b>0.71</b>	<b>81.70%</b>

Table 2: Human inter-annotator agreement across three settings.

Dimension	General-Rubric	Context-Aware
Task Align.	0.68	0.71
Global Coher.	0.67	0.72
Consistency	0.35	0.42
Concept Intro.	0.73	0.76
Local Coher.	0.64	0.68
Evidence Suff.	0.74	0.78
Warrants	0.73	0.73
Qualifiers	0.67	0.74

Table 3: Per-dimension human inter-annotator agreement (Fleiss’  $\kappa$ ) under the General-Rubric and Context-Aware settings. The No-Rubric setting does not elicit dimension-level labels and is therefore omitted.

trast, models such as Gemini-3-Pro and Claude-4.5-Sonnet shift noticeably across domains, indicating greater sensitivity to dataset-specific conventions and topics. (2) Longer inference does not reliably improve logical quality. Reasoning-style (“think”) variants do not consistently outperform their base counterparts and sometimes underperform. A plausible explanation is that longer reasoning chains can introduce redundancy, topic drift, or over-elaboration that obscures the core argumentative thread, without strengthening rubric-targeted requirements such as clear organization or explicit claim–support links. (3) Models specialize in different aspects of logical quality. Specifically, some models perform relatively stronger on *Macro-Logic* and *Expositional-Logic* but are less dominant on *Structural-Logic* (e.g., Gemini-3-Pro, Gemini-2.5-Pro), whereas others show comparatively stronger performance on *Structural-Logic* with weaker or less consistent advantages on *Expositional-Logic* (e.g., o3, GPT-5). Additional experiment details are provided in Appendix C.1.

### 5.4 Rubric Effectiveness Analysis (RQ1)

To assess whether explicit rubric design is necessary for reliable logical evaluation, we study the impact of rubric instantiation on both automated judges and human annotators. We compare three settings: (i) **No-rubric**, which requests a holistic comparison without guidance; (ii) **General-rubric**, which provides fixed, context-agnostic definitions for the eight dimensions; and




Figure 3: ReportLogic Leaderboard. Heatmap of win-rates for 16 frontier models across three domains. Darker colors indicate higher win-rates (stronger logical quality), while lighter colors indicate lower win-rates. Columns are ordered and grouped as follows: the first three dimensions correspond to *Macro-Logic*, the next two to *Expositional-Logic*, the following three to *Structural-Logic*, and the final column reports *Overall*.




Figure 4: Ablation Study on Rubric Effectiveness.

(iii) **Context-aware rubric** (ours), which instantiates each dimension into instance-specific inspection items.

We first evaluate multiple LLM judges, including LogicJudge, under all three settings by re-running the same pairwise comparisons on ReportLogic and computing agreement with the human preference labels. As shown in Figure 4 (Quora omitted for brevity), rubric instantiation yields a consistent ordering across judges and domains: context-aware rubric > general rubric > no rubric. This suggests that high-level dimension definitions alone leave the decision boundary underspecified in long-form logical assessment, while instance-specific inspection items make the comparison criteria operational and more stable.

Additionally, we observe a similar pattern for human annotators. As shown in Table 2, rubric guidance substantially increases inter-annotator agreement on the overall preference, with the context-aware rubric producing the most consistent judgments. The improvement is not confined to the overall decision: Table 3 reports per-dimension Fleiss’  $\kappa$ , showing that context-

aware instantiation yields higher or equal agreement across all eight dimensions relative to the general rubric.

Overall, context-aware rubric instantiation improves reliability by making the decision boundary explicit and enabling more accurate and consistent comparisons for both judges and humans. Additional experiment details are provided in Appendix C.2.

## 5.5 Bad Logic Case Analysis (RQ2)

While leaderboard and attack results quantify relative logical performance, they do not reveal how models fail in practice. To make these failures concrete, we conduct a qualitative analysis of frequently losing reports. Representative cases, trigger spans, and minimal fixes for each failure mode are provided in Appendix C.3. Across datasets and models, we identify three recurring failure patterns. First, models violate global consistency constraints by producing locally plausible claims that cannot jointly hold at the document level. Second, reports often exhibit smooth but weakly connected discourse, relying on rhetorical transitions rather than explicit inferential links. Third, conclusions are frequently supported by implicit or missing warrants, with asserted causes lacking specified mechanisms.

## 5.6 Attack Analysis (RQ3)

To assess whether LLM judges can distinguish genuine logical defects from superficial biases and evaluate their robustness under adversarial

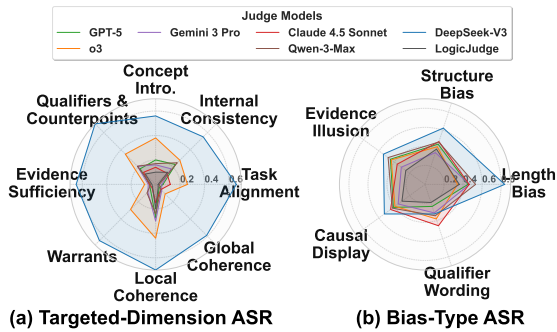


Figure 5: Attack analysis of judge robustness. (a) Targeted-dimension Attack Success Rate (ASR): fraction of cases where the judge prefers the attacked response, with lower indicating better robustness. (b) Bias-type ASR: fraction of cases where the judge prefers a logically equivalent response with surface manipulations.

perturbations, we conduct an attack analysis on ReportLogic. We apply adversarial attacks to  $N = 300$  sampled instances using two complementary attack types: (i) a **Targeted-Dimension Attack**, which injects a localized defect into a specific rubric dimension; and (ii) a **Bias-Type Attack**, which preserves the underlying logical content while manipulating surface cues such as verbosity or formatting. We report Attack Success Rate (ASR) to measure vulnerability. Results are shown in Figure 5.

First, sensitivity to targeted logical defects varies substantially across judges. As shown in Figure 5(a), DeepSeek-V3 is the most vulnerable, exhibiting high ASR across nearly all dimensions, while GPT-5, Gemini 3 Pro, Claude 4.5 Sonnet, and Qwen-3-Max show substantially lower ASR. Notably, o3 exhibits elevated ASR on both *Warrants* and *Local Coherence*, suggesting that reasoning-optimized judges tend to actively repair perceived gaps by inferring implicit warrants and smoothing over local incoherence, which can mask genuine logical defects. Second, surface-level biases can substantially distort judge preferences under the bias-type attack. Figure 5(b) shows that *Length Bias* is the most consistently effective manipulation, and *Structure Bias* and *Evidence Illusion* also produce non-trivial vulnerability. In contrast, *Qualifier Wording* exhibits the weakest effect overall. LogicJudge is relatively more robust overall, exhibiting consistently low ASR across most dimensions. Additional attack details and analyses are provided in Appendix C.4.

## 6 Conclusion

In this work, we argued that the reliability of Deep Research reports hinged not only on factual correctness or fluency, but on their *logical quality*: whether claims were coherently organized, explicitly supported, and verifiable by readers. Specifically, we introduced ReportLogic, a benchmark that operationalizes logical quality through a reader-centric notion of auditability. By decomposing logic into a hierarchical taxonomy and instantiating it via context-aware rubrics, ReportLogic enables fine-grained, diagnostic evaluation of analytical reports. Our human annotations, LogicJudge, and adversarial attack show that explicit rubric guidance is critical for reliable logical assessment, and that off-the-shelf LLM judges remain vulnerable to superficial cues and implicit reasoning shortcuts. Improving the logical quality of LLM-generated reports is essential for building user trust and enabling effective real-world use. We hope ReportLogic provides a foundation for logic-aware evaluators and principled evaluation of Deep Research reports.

## **Limitations**

First, ReportLogic is designed for Deep Research-style analytical report generation, where logical quality is operationalized from the view of auditability. Other genres such as creative writing, narrative storytelling, or purely persuasive rhetoric may follow different conventions and trade-offs for what constitutes “good logic,” and adapting our taxonomy and rubrics to those settings requires further validation and redesign. Second, while our benchmark and judge identify logical deficits in Deep Research reports, we do not directly address how to improve logical generation itself. Retrieval and planning offer promising opportunities to enforce a stronger global structure and verifiable support through evidence coverage, argumentative organization, and explicit intermediate outlines.

## **Ethical Considerations and Artifacts.**

This work uses publicly available datasets and model-generated reports, and all third-party resources are cited and used under their original research terms. We release ReportLogic and LogicJudge solely for research and diagnostic evaluation of long-form logical quality, not for deployment in downstream decision-making or moderation. We do not collect or model personally identifiable information, and we report results only in aggregate. We filter and remove instances containing offensive content during curation and annotation.

## References

- Boxi Cao, Mengjie Ren, Hongyu Lin, Xianpei Han, Feng Zhang, Junfeng Zhan, and Le Sun. 2024. Structeval: Deepen and broaden large language model assessment via structured evaluation. *arXiv preprint arXiv:2408.03281*.
- Michael K Chen, Xikun Zhang, and Dacheng Tao. 2025. Justlogic: A comprehensive benchmark for evaluating deductive reasoning in large language models. *arXiv preprint arXiv:2501.14851*.
- Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- Mingxuan Du, Benfeng Xu, Chiwei Zhu, Xiaorui Wang, and Zhendong Mao. 2025. Deepresearch bench: A comprehensive benchmark for deep research agents. *arXiv preprint arXiv:2506.11763*.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. Enabling large language models to generate text with citations. *arXiv preprint arXiv:2305.14627*.
- Michael Alexander Kirkwood Halliday and Christian MIM Matthiessen. 2013. *Halliday's introduction to functional grammar*. Routledge.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenyuan Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, and 1 others. 2024. Folio: Natural language reasoning with first-order logic. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22017–22031.
- Yuxuan Huang, Yihang Chen, Haozheng Zhang, Kang Li, Huichi Zhou, Meng Fang, Linyi Yang, Xiaoguang Li, Lifeng Shang, Songcen Xu, and 1 others. 2025. Deep research agents: A systematic examination and roadmap. *arXiv preprint arXiv:2506.18096*.
- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoon Yun, Seongjin Shin, Sungdong Kim, James Thorne, and 1 others. 2023. Prometheus: Inducing fine-grained evaluation capability in language models. In *The Twelfth International Conference on Learning Representations*.
- Tom Kocmi and Christian Federmann. 2023. Large language models are state-of-the-art evaluators of translation quality. *arXiv preprint arXiv:2302.14520*.
- John Lawrence and Chris Reed. 2019. Argument mining: A survey. *Computational linguistics*, 45(4):765–818.
- Minghao Li, Ying Zeng, Zhihao Cheng, Cong Ma, and Kai Jia. 2025. Reportbench: Evaluating deep research agents via academic survey tasks. *arXiv preprint arXiv:2508.15804*.
- Chris Yuhao Liu, Liang Zeng, Yuzhen Xiao, Jujie He, Jiakai Liu, Chaojie Wang, Rui Yan, Wei Shen, Fuxiang Zhang, Jiacheng Xu, and 1 others. 2025. Skywork-reward-v2: Scaling preference data curation via human-ai synergy. *arXiv preprint arXiv:2507.01352*.
- Dakota Mahan, Duy Van Phung, Rafael Rafailov, Chase Blagden, Nathan Lile, Louis Castricato, Jan-Philipp Fränken, Chelsea Finn, and Alon Albalak. 2024. Generative reward models. *arXiv preprint arXiv:2410.12832*.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100.
- Jakob Mökander, Jessica Morley, Mariarosaria Taddeo, and Luciano Floridi. 2021. Ethics-based auditing of automated decision-making systems: Nature, scope, and limitations. *Science and Engineering Ethics*, 27(4):44.
- Mihir Parmar, Nisarg Patel, Neeraj Varshney, Mutsumi Nakamura, Man Luo, Santosh Mashetty, Arindam Mitra, and Chitta Baral. 2024. Logicbench: Towards systematic evaluation of logical reasoning ability of large language models. *arXiv preprint arXiv:2404.15522*.
- Haoran Que, Feiyu Duan, Liqun He, Yutao Mou, Wangchunshu Zhou, Jiaheng Liu, Wenge Rong, Zekun Moore Wang, Jian Yang, Ge Zhang, and 1 others. 2024. Hellobench: Evaluating long text generation capabilities of large language models. *arXiv preprint arXiv:2409.16191*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Zhengliang Shi, Yiqun Chen, Haitao Li, Weiwei Sun, Shiyu Ni, Yougang Lyu, Run-Ze Fan, Bowen Jin, Yixuan Weng, Minjun Zhu, and 1 others. 2025. Deep research: A systematic survey. *arXiv preprint arXiv:2512.02038*.
- Hao Sun, Hengyi Cai, Bo Wang, Yingyan Hou, Xiaochi Wei, Shuaiqiang Wang, Yan Zhang, and Dawei Yin. 2024. Towards verifiable text generation with evolving memory and self-reflection. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8211–8227.
- Stephen E Toulmin. 2003. *The uses of argument*. Cambridge university press.

- Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Alberdingk Thijm, Graeme Hirst, and Benno Stein. 2017. Computational argumentation quality assessment in natural language. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 176–187.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, and 1 others. 2024. Long-form factuality in large language models. *Advances in Neural Information Processing Systems*, 37:80756–80827.
- Yuhao Wu, Ming Shan Hee, Zhiqing Hu, and Roy Ka-Wei Lee. 2024. Longgenbench: Benchmarking long-form generation in long context llms. *arXiv preprint arXiv:2409.02076*.
- Tianze Xu, Pengrui Lu, Lyumanshan Ye, Xiangkun Hu, and Pengfei Liu. 2025. Researcherbench: Evaluating deep ai research systems on the frontiers of scientific inquiry. *arXiv preprint arXiv:2507.16280*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.

## Appendix

### A LogicJudge Training

#### A.1 Output Schema for LogicJudge

To enforce comprehensive rubric coverage and enable reliable automated parsing, we constrain the output of LogicJudge to follow a strict, machine-verifiable schema. The training prompt used to elicit this format is provided in Figure 12. Each prediction consists of two components: a structured reasoning block enclosed in a `<think>` tag, followed by a single final preference token.

**Structured Reasoning Block.** The `<think>` block contains a JSON object that explicitly records the judge’s intermediate, dimension-level evaluations. This object must include the following two top-level fields:

- `aspect_evaluations`: a dictionary mapping each of the eight logical dimensions to a discrete comparison decision and a brief justification.
- `overall_explanation`: a short natural-language explanation summarizing the primary factors driving the overall preference.

Each entry in `aspect_evaluations` corresponds to one logical dimension (Section 3.1) and follows a uniform structure with two required fields:

- `decision`: a categorical label from the set `{A>B, A<B, both_good, both_bad}`, indicating the relative logical quality under that dimension.
- `justification`: a concise textual rationale grounded in the content of the two reports and the instantiated rubric.

**Final Decision Token.** After the `<think>` block, the model outputs exactly one deterministic token representing the overall preference label (e.g., `A>B`). This separation allows downstream systems to parse the final verdict efficiently while retaining access to dimension-level diagnostics when needed.

**Schema Illustration.** An example of a valid output structure is shown below:

```
<think>
{
  "aspect_evaluations": {
```

```
    "task_alignment_claim_clarity": {
      "decision": "Tie",
      "justification": "Both reports
directly answer the query by
attributing Transformer dominance
to architectural advantages over
RNN/LSTM, and both keep this thesis
stable throughout."
    },
    "global_coherence": {
      "decision": "A<B",
      "justification": "B cleanly
organizes the answer into a
criteria-driven structure
(parallelism, long-range
dependency, architectural
flexibility, multi-head attention),
making the argument easier to
follow at the document level."
    },
    "...": { }
  },
  "overall_explanation": "On the query
of why Transformers overtook
RNN/LSTMs, Report B is more
convincing because it
systematically decomposes the
advantages and backs them with
concrete benchmarks and deeper
mechanism explanations (e.g.,
constant path length mitigating
long-range issues). A is smoother
and more balanced, but B’s
structure-plus-evidence makes the
core argument stronger overall."
}
</think>
A<B
```

**Design Rationale.** This schema serves two purposes. First, by requiring explicit evaluations for all eight dimensions, it prevents the judge from collapsing into holistic or surface-level judgments. Second, the strict separation between structured reasoning and the final decision ensures robustness for both training and deployment: unparsable outputs are immediately identifiable, and valid outputs provide fine-grained diagnostic signals in addition to a scalar preference.

#### A.2 Training Details for LogicJudge

This appendix provides the full objective definitions, reward design, and optimization details for training the LogicJudge model, complementing the high-level description in Section 4.

##### A.2.1 Supervised Fine-Tuning Objective

The primary objective of SFT is to adapt the model to the specific rubric-following and schema-compliance requirements of ReportLogic. Given the dataset of curated tuples  $D = \{(x, y^*)\}$ , where

$x = (q, \mathcal{R}, y_A, y_B)$  is the input and  $y^*$  is the label, we optimize the standard autoregressive object:

$$\mathcal{L}_{SFT}(\theta) = -\mathbb{E}_{(x, y^*) \sim D} \left[ \sum_{t=1}^T \log \pi_{\theta}(y_t^* | y_{<t}^*, x) \right] \quad (1)$$

SFT serves as a critical initialization for two reasons: (1) Schema Alignment: It teaches the model to adhere to the rigid format required for automated parsing. (2) Rubric Grounding: It aligns the model’s natural language explanations with our eight-dimensional taxonomy, preventing the generation of generic or irrelevant critiques.

### A.2.2 Group Relative Policy Optimization

While SFT establishes formatting and basic reasoning, it may not fully capture the decision boundaries for hard negatives. To further optimize discriminative performance, we employ Group Relative Policy Optimization (GRPO) (Shao et al., 2024).

Unlike standard Proximal Policy Optimization (PPO), which relies on a separate value network (critic) to estimate the baseline and often introduces significant memory overhead, GRPO estimates the baseline directly from the group statistics of sampled outputs. Formally, for each input  $x$ , we sample a group of  $G$  outputs  $\{o_1, \dots, o_G\}$  from the current policy  $\pi_{\theta_{old}}$ , and let the probability ratio be  $r_i(\theta) = \frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}$ . The GRPO objective is:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{q, \{o_i\}} \left[ \frac{1}{G} \sum_{i=1}^G (\mathcal{M}_i(\theta) - \beta \mathbb{D}_{KL}) \right] \quad (2)$$

where  $\mathcal{M}_i(\theta)$  represents the clipped surrogate objective, with  $\epsilon$  denoting the clipping threshold that limits the deviation of the policy update:

$$\min \left( r_i(\theta) \hat{A}_i, \text{clip}(r_i(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_i \right) \quad (3)$$

where the advantage  $\hat{A}_i$  is computed by normalizing the rewards within the group:  $\hat{A}_i = \frac{r_i - \text{mean}(\{r_j\})}{\text{std}(\{r_j\})}$ . This group-based normalization is particularly effective for reasoning tasks, as it robustly distinguishes better reasoning paths from worse ones within the same context, independent of the absolute difficulty of the input query.

A key challenge in RL for structured generation is structure collapse, where the model optimizes for label correctness but forgets the output schema. To prevent this, we design a hierarchical reward function that treats format validity as

a hard prerequisite. For a generated output  $o$ , the total reward  $r(o)$  is computed as:

$$r(o) = \begin{cases} r_{format} & \text{if format is invalid} \\ r_{acc} & \text{if format is valid} \end{cases} \quad (4)$$

1. **Hard Format Constraint ( $r_{format}$ ):** We parse the output to verify it is a valid JSON object containing all eight required dimension keys and a final decision token. If parsing fails, the model receives a severe penalty (e.g.,  $r_{format} = -1$ ). This ensures that structural validity is prioritized above all else, as an unparsable report is useless to the pipeline.
2. **Soft Accuracy Incentive ( $r_{acc}$ ):** If and only if the format is valid, we evaluate the correctness of the final decision label against the consensus ground truth. Correct predictions yield a positive reward (e.g.,  $+1$ ), while incorrect ones yield a mild penalty (e.g.,  $-0.5$ ).

This “Format-First, Logic-Second” shaping stabilizes training, ensuring the judge remains deployable while iteratively improving its reasoning accuracy.

## B Implementation Details

### B.1 Dataset Details

We summarize the construction of the datasets used in our experiments, including query sourcing, Deep Research suitability filtering, query-level splitting, and pairwise expansion. Table 4 reports the resulting statistics for evaluation and judge training.

**Query sources and filtering.** We collect open-domain candidate queries from three sources: (i) existing Deep Research benchmarks (Xu et al., 2025; Du et al., 2025), (ii) Zhihu, a professional Chinese community-driven Q&A platform, and (iii) Quora, a large English community-driven Q&A forum. To ensure that retained queries are genuinely report-appropriate, we apply the Deep Research suitability filtering protocol described in Section 3 (Figure 8), retaining only queries that require retrieval-dependent, multi-aspect synthesis and structured long-form analysis. Queries violating hard constraints (e.g., political content or closed-form tasks) are discarded.

Dataset	#Queries	Train/Val/Test (queries)	#Pairs (human-labeled)	#Pairs (distilled)
DeepResearch	1204	1055 / 17 / 17	102	7912 / 94
Zhihu	1262	843 / 224 / 195	1170	15210 / 1424
Quora	1198	1098 / 50 / 50	300	6590 / 336
<b>Total</b>	<b>3779</b>	<b>2996 / 291 / 262</b>	<b>1572</b>	<b>29712 / 1854</b>

Table 4: Dataset statistics. All splits are performed at the query level. Human-labeled pairs appear only in the test split and are annotated by three experts with majority voting. Distilled pairs are used for LogicJudge training and validation after consensus and swap-consistency filtering.

**DeepResearch domain and synthetic expansion.** The DeepResearch Benchmark is intrinsically small, containing only 102 original queries (Xu et al., 2025; Du et al., 2025). Given the high cost of expert human annotation and the need to allocate original queries across training, validation, and testing, it is not feasible to reserve the entire set for human evaluation. To increase coverage during LogicJudge training, we therefore expand only the training split with synthetic query variants derived from the original DeepResearch training queries, which provide additional topical and linguistic diversity. Synthetic variants are produced by controlled rewriting of the original training queries and are subjected to near-duplicate filtering to avoid trivial overlap. These synthetic queries are used exclusively for training and do not define the evaluation distribution. Importantly, all human evaluation on the DeepResearch domain is conducted on held-out original (non-synthetic) queries, ensuring that test results reflect the authentic query distribution of the benchmark.

**Query-level split and pairwise expansion.** We split each dataset at the query level into train/validation/test, ensuring no query overlaps across splits. For each query, we generate multiple candidate reports using the set of report-generation models described in Section 3. We then construct pairwise comparison instances by pairing reports generated for the same query.

**Human annotation vs. distilled supervision.** Because evaluating long-form logical quality is cognitively demanding, we reserve human annotation for the test split only. Each test pair is labeled by three trained experts with majority vote. In contrast, the train/validation splits are used exclusively to train and tune LogicJudge and are constructed via the distillation protocol in Section 4, which applies consensus filtering and swap-consistency filtering to retain high-

confidence preference signals.

## B.2 Models and Baselines

We detail the models used as judge baselines in our experiments, including model families, variants, and evaluation settings.

**Model Families.** We evaluate **LogicJudge** against 16 state-of-the-art LLMs spanning both proprietary and open-weight ecosystems. The selected baselines cover five major model families:

- **GPT Series.** We include GPT-4.1, GPT-5, and GPT-5.1 as representative instruction-tuned models, as well as o3, a frontier reasoning-oriented model designed for complex multi-step inference.
- **Claude Series.** We evaluate Claude-3.5-Sonnet and Claude-4-Sonnet, together with their reasoning-enhanced variants (Sonnet-3.5-Think and Sonnet-4-Think). These variants explicitly increase test-time computation, allowing us to isolate the effect of extended reasoning on logical judgment.
- **Gemini Series.** We include Gemini-2.5-Pro and Gemini-3-Pro, along with their corresponding thinking-mode variants. Similar to Claude, these models enable inference-time scaling for deeper chain-of-thought reasoning.
- **Qwen Family.** To represent strong open-weight models, we evaluate Qwen-Max, Qwen3-Max, and Qwen3-235B. These models are competitive instruction-tuned systems widely used in open research settings.
- **DeepSeek.** We include DeepSeek-V3, a high-performing open-weight model with strong generation capabilities, which serves as a representative baseline for non-specialized judges.

**Reasoning-Enhanced Variants.** For models that provide explicit “thinking” or reasoning modes (e.g., Claude-Think, Gemini-Think), we evaluate both the standard and reasoning-enhanced versions. These variants are designed to allocate additional inference-time compute to internal reasoning processes. Evaluating both variants allows us to assess whether increased test-time reasoning alone improves logical judgment, independent of task-specific alignment.

### B.3 Training Settings and Hyperparameters

We provide the detailed training configuration for LogicJudge, complementing the high-level description in Section 4. All training runs were conducted on NVIDIA A100 80GB GPUs.

**Stage 1: Supervised Fine-Tuning (SFT).** The SFT stage aligns the model with the rubric-guided reasoning format and the structured output schema required for automated parsing. We fine-tune the model on the distilled training split for 5 epochs using a global batch size of 128 and a learning rate of  $1 \times 10^{-5}$ . To support the long contexts required by Deep Research reports, the maximum input sequence length is set to 40,960 tokens. SFT is trained using 16 A100 GPUs and takes approximately 6 hours. To prevent overfitting and to provide a stable initialization for reinforcement learning, we select the checkpoint with the lowest validation loss on the distilled validation set as the starting point for the subsequent GRPO stage.

**Stage 2: Group Relative Policy Optimization (GRPO).** After SFT convergence, we further refine LogicJudge using GRPO to sharpen discriminative preference judgments on challenging report pairs. During GRPO training, we use a large global batch size of 2,048 and a rollout batch size of 512 to ensure reliable estimation of group-relative advantages. The initial KL coefficient is set to 0.001 to constrain policy drift, and the sampling temperature is fixed at 1.0 to balance exploration and stability. The maximum generation length for the judge’s reasoning output is capped at 4,096 tokens. GRPO is trained using 32 A100 GPUs and takes approximately 16 hours.

**Evaluation Protocol.** All baseline models are evaluated strictly as judges, not generators. Each model receives the same input format consisting of: (i) the user query, (ii) the instantiated context-aware rubric, and (iii) a pair of candidate reports

$(y_A, y_B)$ .

To ensure fairness and isolate judgment capability: (1) All baselines are evaluated in a one-shot setting without any additional demonstrations. (2) The exact same prompt template and rubric definitions used for LogicJudge are applied to all baselines, as shown in Figure 12. (3) No model-specific prompt engineering or output post-processing is used.

Models are required to output a pairwise preference decision, which is evaluated using the agreement-based metrics described in Section 5.2.

**Deterministic Parsing Rule.** Some baseline models do not reliably follow the structured output schema required by LogicJudge. For such models, we extract the final pairwise preference using a deterministic parsing rule. Specifically, we identify the model’s explicit comparative judgment between the two responses (e.g., “A is better than B”, “B outperforms A”, or tie statements such as “both are equally good”), prioritizing the final explicit decision when multiple judgments are present. If no unambiguous preference can be identified, the instance is excluded from agreement-based evaluation. This procedure is fully rule-based and does not introduce additional learned components or heuristics.

## C Experiment Analysis

### C.1 Leaderboard Analysis

We provide additional details and analyses for the leaderboard in Figure 3, including the evaluation setting, how win-rates are computed.

For each domain, we use the full set of test queries in ReportLogic. For every query, we prompt each of the 16 evaluation models to generate a report using the same prompt shown in Figure 9, resulting in one report per model per query. We then construct pairwise comparisons by matching reports produced for the same query across all model pairs. For each comparison pair, we apply the same context-aware rubric generation protocol as in Section 3.2 to produce instance-specific inspection items under our taxonomy, conditioned on the query and the two compared reports. Given these rubrics, LogicJudge performs a dimension-wise pairwise judgment and selects the preferred report for each dimension. A model receives a win on a dimension whenever its report is preferred over its paired opponent for the same

query. The win-rate on a given domain and dimension is computed as the fraction of pairwise comparisons a model wins, averaged over all queries and opponents; ties are excluded. Figure 3 visualizes the resulting dimension-wise win-rates as a heatmap.

## C.2 Rubric Effectiveness

We provide full experiment setting details of the rubric ablation summarized in Section 5.4 here, together with a discussion of rubric-side quality assurance that underpins the ablation.

**Rubric Quality Assurance.** Beyond the downstream effectiveness analysis, we also ensure rubric quality at the generation stage through model selection and annotator-side verification. During development, we compared several candidate rubric generators, including Gemini-2.5-Pro, GPT-5.1, and Claude-4.5-Sonnet. For each candidate, we manually inspected sampled rubrics for schema compliance, semantic alignment with the target dimension, and usability during annotation. We selected Claude-4.5-Sonnet because it consistently produced clearer, more structured, and better dimension-aligned rubrics than the alternatives. During annotation, generated rubrics are not treated as unquestioned ground truth: each annotator first verifies that the rubric is well-formed and usable (i.e., the guiding question targets the intended dimension, the good/bad examples are consistent with that question, and the span hints point to inspectable regions) before applying it to report comparison. Ill-formed rubrics are flagged and re-generated. Across the full dataset, we observed no violations of the required JSON schema, which we attribute to constrained prompting and the strong instruction-following behavior of the selected generator.

**Ablation Protocol.** We compare three conditions that differ only in how decision boundaries are specified: No-Rubric (holistic comparison without explicit criteria), General-Rubric (fixed, context-agnostic definitions of the eight dimensions), and Context-Aware Rubric (ours; instance-specific inspection targets including a comparison question, span-level cues, and paired good/bad examples for each dimension). For this ablation, we randomly sample 200 instances from the Report-Logic test set. All other factors are controlled (same report pairs, same judge prompt template

aside from the rubric content, and the same evaluation metric).

**Judge-side experiment setting.** We evaluate multiple judge models (including LogicJudge and selected frontier LLM judges) on all three datasets. We use the same agreement-based metric as the main experiments: a prediction is counted as correct only if the judge selects the same winner when the pair is presented in the original order and when the two responses are swapped, controlling for position bias.

**Human-side experiment setting.** We conduct a controlled annotation study with three trained expert annotators under the same three rubric conditions. We report inter-annotator agreement using Fleiss’ Kappa ( $\kappa$ ) and average pairwise agreement. Specifically, for each report pair, each annotator provides a categorical preference label in  $\{A > B, A < B, \text{Tie}\}$ . Average pairwise agreement is computed as the mean, over all report pairs, of the fraction of annotator pairs that assign the same label; with three annotators, this equals the average of the three pairwise match indicators. Fleiss’ Kappa  $\kappa$  measures inter-annotator agreement beyond chance for multiple raters. It compares the observed agreement among annotators with the agreement that would be expected if annotators made decisions according to the overall label distribution at random. Formally,  $\kappa = \frac{\bar{P} - \bar{P}_e}{1 - \bar{P}_e}$ , where  $\bar{P}$  denotes the average observed agreement across report pairs, and  $\bar{P}_e$  denotes the expected agreement under random labeling based on empirical label frequencies.

## C.3 Qualitative Bad-Case Analysis

While the leaderboard provides a quantitative ranking of relative logical quality, it does not reveal the concrete mechanisms behind model failures in Deep Research-style report generation. To make these deficits visible and actionable, we conduct a qualitative bad-case analysis on frequently losing reports in Figure 6.

**(1) Internal Inconsistency.** Internal consistency requires a report to satisfy global constraints and avoid contradictions across sections, which is a prerequisite for auditable long-form argumentation. However, we observe that models can produce locally plausible details that violate global constraints. In Case 1 (Figure 6), when reporting “Chinese Social Stratification,” the model assigns

population shares to multiple strata, yet the resulting quantities cannot simultaneously hold: the implied totals (e.g., 45% for Class 4 and ~25% for Classes 6–9, in addition to other stated groups) exceed 100%. This pattern indicates that numerical statements are often generated as individually reasonable fragments without being reconciled under shared global constraints, undermining the report’s verifiability.

**(2) Local Incoherence.** We find that many reports maintain surface smoothness via generic discourse markers (e.g., “however,” “besides,” “from a career perspective”), yet the paragraph-to-paragraph progression often reflects dimension switching rather than inferential advancement. In Case 2 (resigning to pursue a fully funded Princeton PhD), the response first develops a risk-oriented block. It discusses financial pressure (cost of living, travel, and loss of income) and psychological adaptation (uncertainty, stress, and cultural adjustment). This setup introduces a set of constraints that would need to be reconciled for the argument to progress.

However, the response does not derive such a criterion. It does not summarize the risks into a decision rule (e.g., whether expected long-term gains outweigh short-term costs given one’s financial buffer and risk tolerance). Instead, it abruptly pivots to a benefit-oriented block about long-term career ceilings and academic networks.

The transition is therefore only superficial. Connective phrases like “from a career development perspective” and “therefore, in the long run” provide rhetorical flow, but they do not establish a discourse relation that links the two blocks. In particular, the text never explains how the earlier constraints enter the later argument (as assumptions, thresholds, or explicit trade-off terms), nor does it clarify which factors dominate under which scenario. As a result, the response reads like a coherent checklist of considerations rather than a step-by-step decision argument with explicit dependencies and a resolved trade-off.

**(3) Unsupported Warrants.** Analytical claims require explicit warrants because, without an inferential bridge, readers cannot verify whether the conclusion follows from the presented evidence, and fluent prose can mask unsupported causal leaps. In Case 3 (the “Cell Lysis Device Market” report), the model repeatedly treats high-level drivers (e.g., “biomedical industry expansion,” “policy support,” “precision medicine,” and “technology progress”) as sufficient evidence for the

conclusion that “market demand increases.” However, the report rarely specifies the intermediate links that would make these causal statements verifiable. For example, the claim that “biomedical industry expansion has pushed market demand” leaves the mechanism under-specified: it does not identify which downstream workflows are expanding (e.g., nucleic-acid testing, protein extraction, or drug discovery), why cell lysis becomes a bottleneck in those workflows (e.g., throughput or reproducibility constraints in sample preparation), or through which purchasing channel demand materializes (e.g., capacity expansion, automation upgrades, or replacement cycles). As a result, the statement functions as a correlation-like association rather than a justified causal pathway.

The same failure is evident for “policy support.” The report asserts that policy “creates a favorable environment” and “promotes expansion,” yet it does not articulate any concrete instrument that could plausibly translate into device procurement (e.g., research grants that expand equipment budgets, platform-building programs that trigger centralized purchasing, or standardization/regulatory changes that incentivize automated sample-processing pipelines). Instead, the text uses policy as rhetorical cause and jumps directly to market growth, leaving the reader unable to check whether the conclusion follows. Overall, the report matches the tone of professional analysis—enumerating drivers with confident language—but omits the inferential bridge that makes the argument auditable, and therefore reads as a list of plausible factors rather than a supported causal account.

**Implication.** Across cases, the most damaging failures are not surface fluency errors but missing global constraint reconciliation, under-specified discourse relations between adjacent sections, and absent warrants for key inferences. These observations motivate evaluation and training signals that explicitly target consistency, warranted inference, and discourse-level linkage in long-form analytical reports, which are precisely the failure modes emphasized by the ReportLogic taxonomy.

### Case 1: Internal Consistency (Global constraints violated)

**Query (abbr.).** Collect and summarize the real income and financial status of China's nine social strata; analyze the defining traits, size, and financial capacity of the middle class.

**Trigger span (excerpt; translated from original text).** Stratum 2: Upper-middle-class households account for 1% of all households (about 3.2 million households), with net assets around 12.12 million RMB and annual net income around 570k RMB.

Stratum 3: Ordinary middle-class households account for 10% of all households (about 32 million households), with net assets around 3.42 million RMB and annual net income 189k RMB.

Stratum 4: "Xiaokang" households account for 45% of all households (about 144 million households), with net assets around 600k RMB and annual net income 96k RMB.

Stratum 5: Subsistence households are described as "the lower half of 75% of all households" (about 80 million households), with net assets 117k RMB and annual net income 39k RMB.

Strata 6–9: Basic survival groups together account for about 25% of all households (annual income typically below 20k RMB).

(The report also assigns non-zero shares to Stratum 1 and other groups.)

**Failure diagnosis (what goes wrong).** The problem is not whether any single number is individually plausible; it is a violation of *global accounting constraints*. The excerpt mixes (i) shares that are framed as fractions of *all households* (e.g., 45%, 10%, 1%, plus 25%) with (ii) an additional statement whose *denominator and nesting are unclear* (the "lower half of 75%" phrasing). Because the report never specifies whether strata are mutually exclusive, nested, or overlapping—and never performs any reconciliation—the strata cannot be verified as a coherent partition.

**Missing bridge / Minimal fix.** Explicitly define whether the nine strata are mutually exclusive; standardize the denominator (all households vs. subset); and provide a cumulative sum-check to ensure totals equal 100% (up to rounding).

### Case 2: Local Coherence ("smooth but checklist-like" progression)

**Query (abbr.).** How should one think about resigning to pursue a fully funded Princeton PhD offer?

**Trigger span (excerpt; translated from original text; adjacent paragraphs).** *However, resigning to study abroad is not without risks. First, financial factors cannot be ignored. Even with full funding, living costs, international travel expenses, and other potential costs still require careful planning... In addition, resigning means giving up the current income source, which may create short-term financial pressure. Therefore, before making this decision, one must comprehensively assess one's financial situation and create a detailed budget plan to ensure basic living needs during the study period.*

*From a career development perspective, whether resigning to study abroad is worthwhile depends on personal career goals and the characteristics of the industry... Princeton's PhD program not only provides high-quality academic training, but also helps students build extensive academic networks, which is important for future career development... Moreover, Princeton's reputation and influence can offer more job opportunities and higher starting salaries.*

**Failure diagnosis (why coherence breaks locally).** The paragraph boundary marks a *dimension switch* (short-term risks → long-term upsides) rather than a stepwise progression. Although the second paragraph begins with "From a career development perspective", it does not state how the prior risk analysis enters the subsequent argument (as assumptions, thresholds, or explicit trade-off terms). Crucially, the text never introduces an intermediate decision criterion—e.g., "if the financial runway is at least  $X$  months and the stress tolerance is above  $Y$ , then the short-term risks are acceptable." Without such a bridge, the argument does not accumulate: each paragraph contributes a relevant factor, but the reader is not shown how the previous paragraph becomes a premise for the next. The result is "smooth but checklist-like" writing: cohesive at the surface, but locally under-connected.

**Missing bridge / Minimal fix.** Add one bridging sentence that summarizes the risks into an operational condition (financial buffer / risk tolerance threshold), and open the next paragraph by stating that long-term benefits dominate only when that condition holds.

### Case 3: Warrants & Causal Reasoning (Unsupported causal drivers)

**Query (abbr.).** Report on the global and China market status and future trends of cell lysis devices.

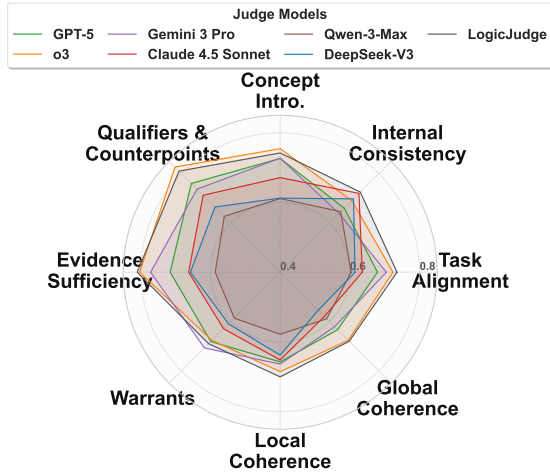
**Trigger span (excerpt; translated from original text).** *The growth of the global cell lysis device market is mainly driven by the following factors. First, continued progress in life-science research—especially rapid development in genomics, proteomics, and cell biology—raises higher requirements for efficient and reliable lysis equipment. Second, the expansion of the biomedical industry, especially increasing demand in biopharmaceuticals and molecular diagnostics, has driven market demand for cell lysis devices. In addition, technological progress is also an important driver of market growth. In recent years, emerging technologies such as automation, microfluidics, and nanotechnology applied to cell lysis devices not only improve device performance but also expand their application scope.*

**Failure diagnosis (unsupported warrants).** Warrants matter because they make causal claims *checkable*: they specify *why* a purported driver should increase demand, rather than merely asserting correlation. In this excerpt, causal language

is used repeatedly (“industry expansion” → “driven demand”, “technology progress” → “market growth”) but the mechanism is left implicit. The text does not explain which downstream workflows scale up (e.g., diagnostics sample preparation, protein extraction in biopharma), why lysis becomes a throughput bottleneck, or why that bottleneck translates into additional device procurement. Likewise, “technological progress” is stated as a driver without clarifying whether it increases throughput, reduces unit cost, enables new use cases, or changes adoption incentives. As a result, the report matches the tone of professional market analysis but omits the inferential steps needed to verify the causal chain.

**Missing bridge / Minimal fix.** Provide at least one complete mechanism chain (industry expansion → higher sample volume/throughput → more lysis operations → capacity bottleneck → increased procurement), and make explicit how “technology progress” affects demand (e.g., cost reduction / throughput gain / new use cases / regulatory or workflow constraints).

Figure 6: Bad-case visualization for Deep Research-style reports. Red highlights mark trigger spans where a rubric constraint is violated; blue highlights indicate the missing logical element required by that dimension to make the report auditable.



Targeted-Dimension Attack Isolation Ratio

Figure 7: Isolation Ratio (IR) under targeted-dimension attacks.

#### C.4 Attack Analysis Details

This appendix provides additional details for the attack-based evaluation in RQ3, including adversarial construction protocols, metrics, and extended interpretations. The main paper emphasizes Attack Success Rate (ASR) as the primary robustness indicator; here we additionally report Isolation Rate (IR) for the Targeted-Dimension Attack to analyze whether a judge can localize penalties to the attacked rubric dimension.

**Adversarial Construction.** We randomly sample  $N = 300$  ReportLogic instances, each consisting of a Deep Research query and an associated model response  $x$ . For each instance, we generate an adversarial candidate  $x_{adv}$  under one of two attack suites using Gemini-2.5-Pro (selected as a high-performing model on our logical-quality leaderboard). We then perform manual screening with two trained annotators to ensure each candidate satisfies the suite-specific constraints (targeted degradation for the Targeted-Dimension Attack, or logical equivalence with only surface manipulation for the Bias-Type Attack), resolving disagreements by discussion. The full prompt templates used for adversarial generation are released in our anonymous repository. After finalizing  $x_{adv}$ , we apply the same context-aware rubric generation pipeline as in Section 3.2 to produce instance-specific inspection items for the paired comparison, and then evaluate different judges by asking them to compare  $x$  against  $x_{adv}$  under these rubrics.

The details of two kinds of attack are as follows:

- **Targeted-Dimension Attack (Sensitivity Probe).** This suite tests whether a judge can penalize a *localized* logical defect within a designated rubric dimension. We construct  $x_{adv}$  by injecting an error that targets only the chosen dimension (e.g., removing a warrant in a causal chain or introducing a local contradiction), while keeping other aspects unchanged by construction. A reliable judge should consistently prefer the original response  $x$  over the attacked version  $x_{adv}$ .
- **Bias-Type Attack (Robustness Probe).** This suite tests whether a judge is influenced by superficial presentation cues when the adversarial response is *logically equivalent* to the original. We instantiate five bias types: (1) *Length Bias*: inflating verbosity via redundant paraphrasing without adding new information; (2) *Structure Bias*: adding outline-style scaffolding (e.g., light headers, numbering, discourse signposts, and brief transitions or summaries) while preserving the original claims and order; (3) *Qualifier Wording*: strengthening cautious hedges and generic limitation statements without introducing new conditions or counterexamples; (4) *Evidence Illusion*: labeling and listing existing statements as “evidence” (optionally with light aggregation phrasing) without introducing new supporting content; and (5) *Causal Display*: mechanically making an existing causal chain explicit (e.g., rewriting implicit links into “because–therefore” form and step-by-step presentation) without adding new intermediate claims, mechanisms, or evidence.

**Screening Protocol.** Raw adversarial candidates generated by Gemini-2.5-Pro are manually filtered by two trained annotators before being used in the evaluation. Each candidate  $x_{adv}$  is independently assessed against three criteria, and is retained only if both annotators agree it satisfies all three: (1) **Dimension Isolation.** For Targeted-Dimension attacks, the injected defect must affect only the designated rubric dimension, without introducing collateral damage to other dimensions (e.g., a *Warrants* attack must not simultaneously alter task alignment or internal consistency). For Bias-Type attacks, the manipulation must be confined to surface cues (length, structure, wording) without changing the underlying claims, evidence, or reasoning chain on any dimension. (2) **Minimal Edits.** The adversarial version must preserve

the original response as much as possible outside the attack target. For Targeted-Dimension attacks, annotators verify that unaffected content remains lexically and semantically close to the original. For Bias-Type attacks, annotators verify that no new propositions, evidence, or qualifications are introduced; only surface re-packaging is allowed. (3) **Bounded Strength**. The defect must be perceptible but not catastrophic: a Targeted-Dimension attack should degrade the targeted aspect to a clearly sub-standard level without producing an obviously incoherent report, and a Bias-Type attack should apply a realistically plausible surface manipulation rather than an exaggerated distortion. This bound ensures that judge vulnerabilities reflect subtle reasoning shortcuts rather than trivial edge cases. Candidates failing any criterion are regenerated (for Targeted-Dimension attacks where the defect is too broad, too weak, or leaks into other dimensions) or discarded (for Bias-Type attacks that introduce new content). Disagreements between the two annotators are resolved via discussion until consensus is reached.

**Representative Examples.** To make the two attack families concrete, Table 5 presents five paired excerpts showing reports before and after attack, drawn from our actual attack set. The first three illustrate *Targeted-Dimension* attacks that weaken a specific logical dimension (Qualifiers, Internal Consistency, Concept Introduction), while the last two illustrate *Bias-Type* attacks that preserve the underlying argument but inject a surface signal (report-style numbering, explicit causal connectives). In each row, modifications relative to the original are highlighted in bold, omitted passages are marked with “...”.

**Evaluation Metrics.** To quantify model performance, we employ two key metrics: (1) **Attack Success Rate (ASR)**: This measures the judge’s vulnerability to adversarial manipulation. For Targeted-Dimension attacks, ASR is the proportion of cases where the judge erroneously prefers the degraded  $x_{adv}$  or issues a tie on the targeted dimension (expecting a clear penalty on  $x_{adv}$ ). A tie is counted as a success here because the attack objective is to prevent the judge from clearly penalizing the injected defect, and a tie already means the judge fails to prefer the clean report on the attacked aspect. For Bias-Type attacks, ASR denotes the frequency with which the judge favors

the manipulated text (e.g., the longer one) over the logical original (expecting resistance). A tie is counted as non-success in this setting, because the attack objective is to actively flip the judge toward the attacked side. A tie indicates that the judge does not prefer the attacked output and is therefore not deceived by the surface manipulation. (2) **Isolation Rate (IR)**: Exclusively calculated for the Targeted-Dimension Attack, this metric assesses the selectivity of the judge’s critique. This metric verifies whether the penalty imposed by the judge is confined to the attacked dimension. A high IR indicates that the judge accurately localizes the error (e.g., penalizing only *Warrants & Causal Reasoning* for a broken chain) without reducing scores for unaffected dimensions (e.g., *Task Alignment*), thereby demonstrating precise, disentangled reasoning capabilities rather than a generalized negative halo effect.

**Extended Interpretation.** Beyond the main observations in Section 5.6, the targeted-dimension attack further reveals a clear separation between error localization and attack robustness. Figure 7 shows that o3 achieves the highest IR across most dimensions, indicating that when a localized defect is introduced, its penalties are largely confined to the attacked aspect rather than spilling over broadly. This suggests that o3 more consistently follows a dimension-conditioned evaluation procedure, producing sharper aspect attribution.

However, high isolation does not necessarily imply low attack success. As shown in Figure 5(a), o3 still exhibits relatively elevated ASR on several dimensions compared to GPT-5 and Gemini 3 Pro. Taken together, these results suggest that o3 can localize the defect once it affects its judgment, but may still be easier to flip under certain targeted degradations. A plausible explanation is that reasoning-oriented judges may attempt to preserve a coherent interpretation by implicitly reconstructing missing bridges. This tendency can reduce sensitivity to attacks that weaken explicit support relations, making the model more likely to accept a subtly degraded answer, while still assigning the resulting penalty primarily to the most relevant dimension.

In contrast, Qwen-3-Max exhibits the lowest IR overall, and DeepSeek-V3 also shows consistently low IR, indicating substantial cross-dimension spillover. This pattern is consistent with impression-driven scoring: once a response

is judged worse, multiple aspect ratings are jointly reduced to match the global preference, producing a halo effect rather than precise attribution. Overall, IR provides a complementary diagnostic signal to ASR, distinguishing judges that are vulnerable because they are easily flipped from those that are vulnerable because they cannot precisely localize the source of degradation.

**LogicJudge: Extended Analysis.** We additionally examine LogicJudge as an evaluation-aligned baseline judge trained on ReportLogic supervision. Overall, LogicJudge exhibits relatively low ASR across most targeted dimensions, and achieves strong error localization: in Figure 7, it ranks among the top judges in IR, particularly on *Warrants & Causal Reasoning* and *Concept Intro*, where it is consistently second only to the most localized judge. This suggests that once a defect is recognized, LogicJudge tends to attribute the penalty to the intended rubric aspect rather than triggering broad cross-dimension spillover. Nevertheless, two failure modes remain salient. First, LogicJudge is more affected by *Length Bias*. A plausible explanation is that our training objective is pairwise and rubric-conditioned: when two responses are logically equivalent, the decision boundary becomes narrow and the model must rely on limited comparative signals. Longer rewrites systematically increase redundancy and restatement, which can reduce apparent ambiguity and make support relations appear easier to trace, thereby shifting borderline comparisons even without changing the underlying claims. Second, LogicJudge shows higher ASR under *Local Coherence* attacks. Local coherence defects rarely manifest as a single clearly wrong sentence; instead, they arise from subtle cross-sentence mismatches (e.g., drifting conditions, inconsistent references, or quietly altered claim scopes) that only become evident when aligning information across multiple sentences. In our distillation setup, the supervision includes dimension-level pairwise judgments with brief, high-level rationales, but it may not provide fine-grained, sentence-level attributions that explicitly pinpoint the exact spans responsible for the inconsistency. As a result, when coherence defects are implicit and distributed, they may be insufficient to trigger a decisive preference flip, making it harder for the model to learn a stable boundary for detecting such fine-grained inconsistencies.

## D Artifact Documentation and Ethical Considerations

**Artifacts and attribution.** This work relies on a combination of publicly available datasets, model-generated reports, and evaluation artifacts introduced in this paper. All external datasets and models used in our experiments are properly cited in the main paper. We do not claim ownership over any third-party artifacts, and we follow the original terms under which these resources are released.

**Licensing and terms of use.** All datasets used in this study are either publicly accessible for research purposes or consist of model-generated content. Our use of these resources is strictly limited to non-commercial research and evaluation. We do not redistribute raw data that is subject to restrictive licenses; instead, we report aggregate statistics, evaluation outcomes, and derived annotations consistent with common academic practice.

**Intended use and consistency.** The artifacts created in this work—including the evaluation benchmark ReportLogic and judge model LogicJudge—are intended solely for research and diagnostic evaluation of long-form logical quality. Our use of existing datasets and model outputs is consistent with their intended research use. Derived annotations and judgments are not designed for deployment in downstream decision-making systems or real-world moderation settings.

**Personally identifiable information and offensive content.** The data analyzed in this study consist of public or model-generated text and may contain natural language references typical of open-domain content. We do not collect, analyze, or model personally identifiable information. User identifiers and metadata are removed or anonymized where applicable, and the evaluation focuses exclusively on report-level logical structure rather than individual attributes or identities. Potentially offensive content is not targeted or amplified. During dataset curation and human annotation, we filter and remove instances containing offensive content; remaining data are used solely for research evaluation under standard anonymization procedures.

## E Human Subjects Including Annotators

This study involves human annotators for constructing and validating the ReportLogic benchmark. All annotation procedures were designed to follow established ethical standards for research involving human subjects and to pose minimal risk to participants.

**Instructions to Annotators.** All annotators received comprehensive written guidelines prior to participation. For each annotation task, we first provided a plain-language overview of the study goal and the specific judgment to be made, so that annotators could build an accurate mental model of what constitutes a valid comparison. We then supplied precise definitions of the evaluation criteria and the decision rules for each dimension, together with multiple concrete examples illustrating both acceptable and unacceptable judgments for each possible outcome. The instructions also specified quality-control expectations (e.g., careful reading of both reports, avoiding reliance on superficial cues, and following the prescribed procedure before submitting a verdict). Annotators were informed that their labels would be used for research purposes. No deception was involved, and the task consisted solely of expert assessment of model-generated text.

**Annotation Workflow and Quality Assurance.** Each instance was independently annotated by three annotators following the same rubric-guided procedure. Annotators first reviewed the query and the provided context, then compared the two reports dimension by dimension, and finally selected a three-way verdict for each dimension and overall preference. We monitored inter-annotator agreement throughout the annotation process and used it as a primary quality signal.

For instances where the three annotations did not yield a majority decision (i.e., three-way disagreement) or where the disagreement indicated substantive ambiguity rather than minor differences in interpretation, we applied an adjudication step. Specifically, an expert adjudicator (a domain expert Ph.D. researcher with prior experience in logical evaluation and analytic writing) re-examined the full instance, including the rubric, both reports, and the annotators' rationales. The adjudicator either (i) selected the final verdict when one option was clearly better supported under the rubric, or (ii) initiated a short reconcilia-

tion procedure by documenting the decisive rubric criteria and resolving the disputed points to ensure the final label was consistent with the written guidelines. This process produced a single consolidated gold annotation for each instance.

**Recruitment and Compensation.** Annotators were recruited by the research team based on prior experience with analytical writing and logical evaluation tasks. In total, 25 annotators participated in the study. All annotators hold at least a master's degree in relevant fields, including linguistics, literature, philosophy, or technical writing. Annotators were compensated on a per-task basis at a rate equivalent to approximately \$28 USD per annotated instance, reflecting the substantial cognitive effort required for fine-grained, rubric-guided logical assessment of long-form reports. Compensation was determined in advance and was independent of annotation outcomes.

**Consent and Data Use.** All annotators provided informed consent prior to participation. The instructions explicitly described how the annotated data would be used, stored, and reported. Annotations were analyzed only in aggregate form, and no personally identifiable information was collected, stored, or released. Annotators were informed that participation was voluntary and that they could withdraw at any time without penalty.

**Ethics Review.** The annotation protocol and data collection procedures were reviewed and approved through the internal ethics and compliance review process of the authors' research team. The study involves minimal risk, as it consists solely of expert evaluation of text generated by language models.

**Annotator Characteristics.** Annotators are adult participants (18+), fluent in the language of the annotated content, and selected based on demonstrated expertise rather than demographic attributes. Basic professional characteristics (educational background and domain expertise) were considered for recruitment, while no sensitive personal attributes were collected or used.

## F Use of AI assistants

AI assistants were used exclusively for language polishing, such as grammar correction, clarity improvement, and minor stylistic edits. They were not used for idea generation, methodological de-

sign, experimental execution, result analysis, or interpretation. All scientific contributions and conclusions remain the responsibility of the authors.

## **G Prompt Template**

Table 5: Representative excerpts showing reports before and after attack. In each row, bold marks the spans affected by the attack—in the original column, the passages that are modified; in the attacked column, the resulting modifications. “...” denotes omitted context. All five examples are drawn from real attack outputs in our benchmark.

Attack	Original (excerpt)	Attacked (excerpt)
<i>Targeted-Dimension attacks</i>		
Qualifiers & Counterpoints	... The long-term implication of DLT is the <b>potential</b> creation of a new financial market infrastructure that is more transparent, efficient, and interoperable, <b>though significant challenges related to scalability, regulation, and energy consumption remain to be fully addressed.</b> ... AI adoption... <b>raises concerns regarding job displacement and ethical considerations; the opacity of “black box” algorithms poses challenges for transparency and accountability.</b> ...	... The long-term implication of DLT is the <b>inevitable</b> creation of a new financial market infrastructure that is more transparent, efficient, and interoperable, <b>with only minor adjustments likely to be needed over time.</b> ... AI adoption... <b>will deliver substantial benefits across the industry; modern algorithms consistently provide reliable outcomes in lending and insurance.</b> ...
Internal Consistency	Open Banking is a regulatory and technological framework that mandates financial institutions to securely share customer data with third-party providers (TPPs) through standardized APIs. ... This fosters unprecedented competition and innovation. ...	Open Banking is a regulatory and technological framework... through standardized APIs. <b>In its narrow sense, it is primarily about enabling data portability and account access for licensed payment and information service providers.</b> ... <b>In practice, this has increasingly meant that Open Banking is treated as a broad umbrella for almost any form of API-based financial integration, including services that go well beyond data access and into areas like full-service embedded lending and insurance.</b>
Concept Introduction & Logical Transition	The global financial industry is presently undergoing a transformation of unprecedented scale and velocity, <b>driven by a confluence of technological advancement, evolving regulatory frameworks, and shifting consumer expectations.</b> ... The catalysts... <b>include the lingering structural lessons from the 2008 global financial crisis, the rapid proliferation of digital connectivity, and... new consumers.</b> ... leveraging modern API protocols <b>to facilitate seamless communication between disparate banking institutions.</b> ...	The global financial industry is presently undergoing a transformation of unprecedented scale and velocity, <b>with RegTech, tokenization, and embedded finance now redefining the sector.</b> ... [background on 2008 crisis and digital connectivity moved to a later paragraph]... leveraging modern API protocols <b>for inter-bank messaging.</b> ...
<i>Bias-Type attacks</i>		
Structure Bias	<b>Chess Progression and Technical Analysis</b> In this game, Ke Jie played white... captured the half-point victory. From a technical standpoint, the reasons for Ke Jie’s win can be grouped into three points: accurate judgment of complex positions, decisive use of sacrifice tactics, and flawless endgame execution.	<b>II. Chess Progression and Technical Analysis: A Complete Arc from Opening through Middlegame to Endgame</b> In this game, Ke Jie played white... captured the half-point victory. From a technical standpoint, the reasons for Ke Jie’s win can be grouped into three points: accurate judgment of complex positions, decisive use of sacrifice tactics, and flawless endgame execution, <b>laying a clear technical foundation for the subsequent discussion of strategy and psychology.</b>
Causal Display Bias	The global financial industry is presently undergoing a transformation of unprecedented scale and velocity, driven by a confluence of technological advancement, evolving regulatory frameworks, and shifting consumer expectations. This period of intense innovation represents a fundamental departure from the incremental changes that characterized the late 20th century. ...	... driven by a confluence of technological advancement, evolving regulatory frameworks, and shifting consumer expectations. <b>Because these forces are acting together rather than in isolation,</b> this period of intense innovation represents a fundamental departure from the incremental changes that characterized the late 20th century. ...

### Query Filtering Prompt (Deep Research Suitability)

**Prompt:** You are a rigorous content suitability judge for a Retrieval-Augmented Generation (RAG) based Deep Research system. Given a query, decide whether it is suitable for Deep Research style long-form reporting. Always output exactly one JSON dictionary on a single line, and output nothing else.

**Overall Objective.** ACCEPT queries that typically (i) require consulting up-to-date or specialized external resources or systematically synthesizing substantial background knowledge, theories, and cases, (ii) support analysis across multiple dimensions rather than a single fact or definition, and (iii) are best answered as a structured long-form report, usually extendable to 1500–2000 words or more with multiple sections. Queries better suited to short answers, encyclopedic explanations, simple listings, or common-sense replies should generally be REJECTED.

**RAG / Deep Research Suitability Criteria.** Output decision “ACCEPT” only if the query clearly satisfies multiple criteria below (e.g., at least three). If it only weakly satisfies one or two, output “REJECT”.

1. **Strong retrieval dependence:** A reasonable answer requires consulting external sources such as academic papers, industry reports, empirical data, technical documentation, news, regulations, standards, or real-world cases, or systematically organizing such information.
2. **Multi-source synthesis:** Answering requires integrating information from multiple sources or perspectives rather than citing a single fact or data point, and often involves comparing viewpoints, methods, metrics, or solutions.
3. **Multi-dimensional decomposability:** The query can be naturally decomposed into multiple analytical dimensions (e.g., background, core concepts, current approaches, trade-offs, cases or data comparisons, risks, and future trends). If it essentially asks to “look up one thing”, output “REJECT”.
4. **Complexity and openness:** The query is not a simple yes/no question or a mechanical listing task. It requires reasoning, trade-off analysis, or argumentation, often involving “why”, “with what impact”, or “how to design, optimize, or balance”.
5. **Long-form report suitability:** A reasonable answer would naturally include multiple sections (e.g., motivation, foundations, existing approaches, case studies or comparisons, challenges, and recommendations). If it can be fully addressed in one or two short paragraphs or a short list of steps, output “REJECT”.

**Hard Constraints (Must Reject).** If any condition holds, output decision “REJECT” regardless of the criteria above:

- Sexual content, extreme violence, self-harm, suicide, or illegal activities.
- Real-world politics, governments, political parties, elections, diplomacy, or ideology, regardless of stance or tone.
- Attacks, denigrates, or stigmatizes a specific country, region, or city, explicitly or implicitly, including suggestive or veiled framing.
- Closed-form tasks solvable without external retrieval (e.g., pure mathematical proofs, basic programming exercises, or logic puzzles).
- Casual chat, personal life issues, emotional venting, or family or relationship disputes.
- Purely definitional or short FAQ-style queries that do not support long-form structured analysis.
- Queries that are inherently short and cannot be reasonably expanded into a multi-section analytical report.

**Output Format (JSON dictionary, two keys).** Return exactly one JSON dictionary on a single line, with exactly two keys: “decision” and “reason”. The value of “decision” must be either “ACCEPT” or “REJECT”. The value of “reason” must be a brief justification. For “REJECT”, state the primary reason (e.g., too simple, definitional, political, or not retrieval-dependent).

**Input.** Query: {query}

**Example output.** {“decision”:“ACCEPT”,“reason”:“Requires external retrieval and multi-source synthesis; supports multi-dimensional analysis and naturally fits a long, structured report.”}

Figure 8: Prompts used to filter open-domain queries for Deep Research suitability.

### Report Generation Prompt for Deep Research

**Prompt:** You are a domain expert writing a formal analytical REPORT based on your own expertise and the provided relevant documents retrieved via a Retrieval-Augmented Generation (RAG) pipeline.

**Objective.** Write a comprehensive and detailed report, not a summary or abstract. Your task is to reason deeply, synthesize evidence, and produce a coherent long-form narrative that could be published as an expert analytical report or white paper.

**Writing Requirements (STRICT).** Paragraphized prose only. Each paragraph must contain at least four complete sentences. Absolutely no bullet points, numbered lists, outlines, or tables. Maintain a formal, academic, and explanatory tone throughout. Logical flow is required: every section must connect naturally to the previous one with explicit transitions. Prioritize depth over breadth. For every major point, explain mechanisms (why and how), contextual background, and implications. Length requirement: for English or non-CJK queries, produce at least 3,000 words. For Chinese (CJK) queries, produce at least 6,000 characters excluding spaces. If uncertain, follow the CJK requirement.

**Expansion Policy (When RAG Evidence Is Insufficient).** If the retrieved RAG documents lack coverage, contain gaps, or address the topic only partially, expand the discussion using your domain knowledge while ensuring that all additional content remains reasonable, verifiable, and logically consistent with the RAG context. Never fabricate facts, numbers, or citations.

**Input.** Query: {new\_query}  
RAG Documents (with IDs): {docs\_block}

**Output.** Produce a single, continuous REPORT that strictly follows all requirements above. Begin the report immediately below this instruction without any prefatory remarks.

Figure 9: Prompt used for Deep Research style long-form report generation.

## Context-aware Rubric Generation Prompt

**Prompt:** You are a logic analysis expert. Your task is: given a QUERY and two response texts (A and B), generate a logic evaluation rubric for human reviewers to compare A and B across eight logical aspects.

**Three-layer Logic Framework.** This evaluation is grounded in classical logic and argumentation analysis, and decomposes text-level logic into three layers: (1) **Macro level:** focuses on the overall logical trajectory and topical focus, judging whether the response matches the task and remains globally consistent. (2) **Expositional level:** guided by information flow theory (Given→New), evaluates whether the exposition unfolds naturally from shared context to new information and remains easy to follow. (3) **Structural level:** inspects the completeness of the argumentation chain, including claims, evidence, reasoning, qualifiers, and consistency.

### Eight Logical Dimensions and the Focus of the Rubric Question.

#### 1. Task alignment & claim clarity.

**Definition:** Whether the response clearly addresses the core task of the QUERY in salient positions (e.g., the beginning or the end), and whether it states a clear, concentrated central claim. Evaluate whether the response stays on-topic or drifts.

**Question focus:** Compare which side more clearly responds to the question at the beginning or end, states a central stance, and maintains topic consistency.

#### 2. Global coherence.

**Definition:** Whether the overall logical structure is clear and appropriate, and whether the content unfolds in a task-appropriate order (e.g., for analysis: background→problem→analysis→conclusion; for comparison: comparison→evaluation→conclusion), forming a stable global thread.

**Question focus:** Compare which side is better organized overall, follows a more appropriate task logic, and whose conclusion better aligns with prior content.

#### 3. Internal consistency.

**Definition:** Whether numbers, definitions, and stances remain consistent throughout, without ratio conflicts, shifting definitions, or self-contradictions.

**Question focus:** Compare which side is more internally consistent and logically self-contained.

#### 4. Concept introduction & logical transition.

**Definition:** The naturalness of information unfolding, especially the smoothness of the Given→New transition, evaluated at two levels: (a) Opening level: whether the beginning provides necessary background, context, or motivation before moving to the main claim. Jumping directly into abstract discussion or using undefined terms indicates weak transitions. (b) Within-body level: when introducing new concepts, models, or terms, whether the response provides brief explanations or context so readers can understand their role and origin.

**Question focus:** Prioritize comparing which side better sets up background and naturally introduces the main topic at the opening; if both openings are reasonable, compare which side introduces new concepts more naturally and with sufficient explanation within the body.

#### 5. Local coherence.

**Definition:** Whether adjacent paragraphs or sentences are logically connected, with reasonable transitions and continuity, rather than breaks or leaps.

**Question focus:** Compare which side has more reasonable paragraph-to-paragraph connections and more coherent transitions.

#### 6. Evidence sufficiency & relevance.

**Definition:** Whether key claims are supported by sufficient and on-topic evidence (e.g., data, experiments, facts, literature, cases), and whether the relation between evidence and conclusion is clear and traceable.

**Question focus:** Compare which side provides more concrete, query-relevant evidence that better supports the main conclusion.

#### 7. Warrants & causal reasoning.

**Definition:** Whether the response explains why the evidence supports the conclusion, whether the reasoning chain is complete, whether causal direction is reasonable, and whether there are logical jumps.

**Question focus:** Compare which side provides a more complete reasoning chain, clearer causal relations, and more justified step-by-step transitions.

#### 8. Qualifiers & counterpoints.

**Definition:** Whether the response identifies and handles potential counterexamples, uncertainty, or exceptions, and whether it avoids absolute claims by using appropriate qualifiers, concessions, or balanced framing.

**Question focus:** Compare which side better identifies counterexamples and sets conditions or assumptions that make the argument more balanced.

**Output Requirements.** Output **only** a JSON array of length 8. Each object must correspond to one of Dimensions 1–8 in the same order. Each object must contain exactly five fields:

- “aspect”: the exact name of one of the eight dimensions.
- “question”: an instance-specific judgment question for the current QUERY and texts A/B. The question must be concrete and directly comparable between A and B, avoiding vague prompts such as “Is it reasonable?” or “Is it sufficient?”. It should be phrased so annotators can decide which response is better for the specified logical focus.
- “good\_example”: a positive example that matches the same question and illustrates what good performance looks like.
- “bad\_example”: a negative example that matches the same question and illustrates what poor performance looks like. The good and bad examples must be paired to the same logical point under the same question.
- “span\_hint”: a customized hint indicating linguistic or logical cues in TEXT A/B that annotators should attend to.

#### Generation Guidelines.

1. **Instance-specific.** All rubrics must be customized to the current QUERY and the content of TEXT A/B; avoid generic or template-like questions.
2. **Comparative and self-contained.** Each “question” must be understandable on its own and explicitly invite an A-versus-B comparison.
3. **Claim anchored.** Each “question” must name the concrete focal claim or topic in this instance (not vague phrases such as “key claim” or “core point”).
4. **Logic focus explicit.** Each “question” must specify the exact logical focus being compared (e.g., evidence relevance, warrant completeness, transition bridging), avoiding vague wording such as “more reasonable” or “clearer” without criteria.
5. **Good/bad paired to the same point.** “good\_example” and “bad\_example” must correspond to the same rubric question and illustrate positive versus negative performance on that exact point.
6. **Span hints tied to observable cues.** “span\_hint” must point to concrete linguistic or structural signals in TEXT A/B that annotators can inspect (e.g., explicit evidence statements, bridging sentences, qualifiers).

#### Input.

QUERY: {QUERY}  
 TEXT A: {TEXT\_A}  
 TEXT B: {TEXT\_B}

**Output Format.** The model must output a JSON array of length 8. Below is an example illustrating the required structure, keys, and style:

```
[ { “aspect”: “Evidence sufficiency & relevance”, “question”: “Comparing A and B, which response provides more specific evidence (e.g., emissions data, country cases, or mechanism explanations) to support the claim that a carbon tax reduces emissions, and explains how this evidence justifies the conclusion?”, “good_example”: “Cites concrete country-level data or empirical studies and explains how a carbon tax changes firm incentives, thereby reducing emissions.”, “bad_example”: “States that a carbon tax is effective without providing concrete evidence or explaining the supporting mechanism.”, “span_hint”: “Look for signals such as empirical data, study references, concrete cases, and explicit links between evidence and conclusions.” }, ... ]
```

Figure 10: Prompt used to generate context-aware, instance-specific rubric for eight-dimensional logical evaluation.

## LogicJudge Distillation Prompt

**Prompt:** You are a logic analysis expert tasked with comparing the logical quality of two long-form reports. Given a user **QUERY**, two candidate reports (**A** and **B**), and a set of **context-aware logic rubrics** consisting of eight dimensions, you must determine which report exhibits stronger logical reasoning.

**Objective.** Perform a rubric-guided *pairwise preference judgment* over logical quality. Your task is not to assess factual correctness or writing style, but to evaluate the *logical structure and argumentative soundness* of the two reports. You must first conduct explicit, dimension-level critiques as intermediate reasoning scaffolding, and then aggregate these judgments into an overall preference decision.

### Evaluation Procedure (STRICT).

- For each of the eight logic dimensions provided in the rubrics, determine whether report A or report B demonstrates stronger logical performance.
- Base each judgment on the dimension-specific comparison question, illustrative good/bad examples, and span-level reasoning cues provided in the rubric.
- For every dimension, produce (i) a winner label and (ii) a concise but concrete explanation grounded in logical reasoning.
- If both reports perform similarly well or similarly poorly on a dimension, this must be explicitly stated.
- After completing all dimension-level evaluations, synthesize them into a final overall preference decision.

### Evaluation Principles.

- Focus exclusively on logical reasoning quality, including claim clarity, evidence use, reasoning chains, coherence, and internal consistency.
- Do not rely on global impressions, fluency, verbosity, or stylistic features.
- Avoid vague judgments; all explanations must cite concrete logical characteristics.
- Do not quote or copy text from the reports; explanations must paraphrase logical patterns.

**Logic Dimensions (Fixed Schema).** The evaluation must cover exactly the following eight dimensions:

1. Task alignment & claim clarity
2. Global coherence
3. Internal consistency
4. Concept introduction & logical transition
5. Local coherence
6. Evidence sufficiency & relevance
7. Warrants & causal reasoning
8. Qualifiers & counterpoints

### Input.

QUERY: {QUERY}  
TEXT A: {TEXT\_A}  
TEXT B: {TEXT\_B}  
Logic Rubric (JSON): {RUBRIC\_JSON}

**Output Format (STRICT).** The model must output *exactly one valid JSON object* following the schema below. This structured format explicitly supervises both dimension-level reasoning and the final preference decision, enabling fine-grained analysis and scalable training.

```
{
  "aspect_evaluations": {
    "logic_dimension_name": {
      "winner": "A>B" | "A<B" | "Tie",
      "explanation": "Explanation of logical differences"
    },
    ...
  },
  "overall_winner": "A>B" | "A<B" | "Tie",
  "overall_explanation": "A short paragraph summarizing the overall logical
  preference"
}
```

No additional text, commentary, or formatting outside this JSON object is allowed.

Figure 11: Prompt used to distill rubric-guided, dimension-level pairwise preferences for training LogicJudge.

## LogicJudge Training Prompt

**Prompt:** You are an expert in logical analysis. Given a **QUERY**, two candidate reports (A and B), and a **context-aware rubric** with eight dimensions, you will decide which report is logically better on each dimension and provide targeted explanations. Finally, you will aggregate the eight dimensions into an overall verdict with a brief justification.

**Overall Objective.** You must produce two levels of judgments: (1) **Dimension-wise evaluation:** for each of the eight dimensions in the rubric, decide whether A or B is better and justify the decision; (2) **Overall evaluation:** aggregate the eight dimension-wise decisions into a final overall winner (`overall_winner`) and a short overall explanation.

**Procedure.** (1) Interpret the rubric signals: `question` specifies the logical focus of the dimension; `good_example` and `bad_example` illustrate desired and undesired patterns; `span_hint` highlights cues and locations to inspect for reasoning links. (2) Compare A vs. B for each dimension: choose the winner and provide a concrete explanation (do not quote the original text verbatim); if both are similarly strong or weak, state this explicitly. (3) Produce the overall judgment: `overall_explanation` must be a coherent paragraph (no bullet points), stating the overall winner and the main reason.

**Evaluation Principles.** Evaluate **logical structure and argumentation quality** only (not factual correctness or writing style). All explanations must be grounded in the rubric focus and the reports' reasoning. Avoid vague statements (e.g., "better" without justification). Do not quote the reports verbatim; describe reasoning patterns in your own words. The output must be a single JSON object with exactly the required fields and format.

**Eight Logic Dimensions (names must match exactly).**

1. Task alignment & claim clarity
2. Global coherence
3. Internal consistency
4. Concept introduction & logical transition
5. Local coherence
6. Evidence sufficiency & relevance
7. Warrants & causal reasoning
8. Qualifiers & counterpoints

**Input Format.**

QUERY: {QUERY}

TEXT A: {TEXT\_A}

TEXT B: {TEXT\_B}

Logic Rubric (JSON): {RUBRIC\_JSON}

**Output Format (must follow strictly).**

```
<think>
{
  "aspect_evaluations": {
    "Task alignment & claim clarity": {
      "winner": "A>B" | "A<B" | "both good" | "both bad",
      "explanation": "A concise, specific justification without verbatim
quotes."
    },
    "Global coherence": {...},
    "Internal consistency": {...},
    "Concept introduction & logical transition": {...},
    "Local coherence": {...},
    "Evidence sufficiency & relevance": {...},
    "Warrants & causal reasoning": {...},
    "Qualifiers & counterpoints": {...}
  },
  "overall_explanation": "A single coherent paragraph summarizing the overall
preference and main reasons (no lists)."
```

```
</think>
A>B
```

**Constraints.** (1) The content inside `<think> ... </think>` must be a JSON object containing exactly "aspect\_evaluations" and "overall\_explanation". (2) "aspect\_evaluations" must cover all eight dimensions above. For each dimension, "winner" must be one of "A>B", "A<B", "both good", or "both bad". (3) After `</think>`, output exactly one final verdict token, chosen from "A>B", "A<B", "both good", or "both bad", and output nothing else.

Figure 12: LogicJudge Training Prompt.