

AILS-NTUA at SemEval-2026 Task 12: Graph-Based Retrieval and Reflective Prompting for Abductive Event Reasoning

Nikolas Karafyllis, Maria Lymperaïou, Giorgos Filandrianos,
Athanasios Voulodimos, Giorgos Stamou

School of Electrical and Computer Engineering, AILS Laboratory
National Technical University of Athens

nickarafyllis@gmail.com, {marialymp, geofila}@ails.ece.ntua.gr,
thanosv@mail.ntua.gr, gstam@cs.ntua.gr

Abstract

We present a winning three-stage system for SemEval 2026 Task 12: Abductive Event Reasoning that combines graph-based retrieval, LLM-driven abductive reasoning with prompt design informed by reflective prompt evolution, and post-hoc consistency enforcement; our system ranks first on the evaluation-phase leaderboard with an accuracy score of 0.95. Cross-model error analysis across 14 models (7 families) reveals three shared inductive biases: causal chain incompleteness, proximate cause preference, and salience bias, whose cross-family convergence (51% cause-count reduction) indicates systematic rather than model-specific failure modes in multi-label causal reasoning.

1 Introduction

In a rapidly evolving digital world, information travels at an unprecedented scale, resulting in a constant flood of events. Such dissemination of information does not occur in isolation; instead, it triggers interconnected event chains in which each specific occurrence can act as the root cause of subsequent events. Thus, understanding *why* an event happens within context becomes critical, requiring the discrimination of true cause-and-effect relationships from mere correlations.

Large Language Models (LLMs) have demonstrated great potential in event extraction and prediction (Chang et al., 2025; Tanev et al., 2025; Srivastava et al., 2025; Li et al., 2025). Still, this ability does not account for *abductive reasoning*, which concerns inferring the most plausible cause from incomplete information. Naturally, abduction explains why an event most possibly happened within a causal chain, offering valuable insights and interpretations for chains of events.

A large body of research focuses on LLM reasoning (Giadikiaroglou et al., 2024; Ke et al., 2025; Liu et al., 2025), yet abductive reasoning remains

particularly challenging, since it engages parallel streams of thinking and implicit causal attribution rather than applying deterministic inference rules. For this reason, it serves as a stress test for LLMs, exposing limitations that are less apparent in hypothesis-free reasoning settings (Xu et al., 2025; Floridi et al., 2025; He and Chen, 2025).

Motivated by the challenges entailed within real-world, uncertainty-rich data setups, SemEval 2026 Task 12 **Abductive Event Reasoning: Towards Real-World Event Causal Inference for Large Language Models** (Cao et al., 2026) investigates LLMs’ abilities in event-level causal reasoning: given an event and a set of retrieved documents, the LLM is tasked with identifying the most direct and plausible cause and forming the best explanation by harnessing prior knowledge, context understanding and hypothesis justification. In our work, we address this task through a three-stage pipeline that **ranks first on the evaluation leaderboard** scoring at 0.95/1.00. Overall, our contributions are:

- We evaluate 18 model configurations across seven families, characterizing the performance landscape and persistent challenges of frontier LLMs on abductive event reasoning.
- We perform a thorough error analysis across 14 models (7 families), identifying three shared inductive biases: causal chain incompleteness, proximate cause preference, and salience bias, that produce conservative cause selection as the dominant failure mode.

The code for our system is available on GitHub¹.

2 Background

Task description The task targets the selection of the most direct, plausible cause of a real-world event, as evidenced via a *single-step* reasoning process based on the given textual evidence. Each

¹<https://github.com/nickarafyllis/semEval-2026-task12>

instance comprises an *event*, *context* and four *candidate explanation* options, framing abductive reasoning within a multiple-choice question-answering setup. The *event* field contains a textual description of a real-world event; *context* refers to a set of documents that may either be related to the event or act as distractors; finally, one or more *candidate explanations* may be correct given the input context, while an option stating that “None of the others are correct causes” may be included. An LLM is evaluated on its ability to reason over the context for an event by selecting suitable candidate explanations. A full match between predicted and ground truth responses yields 1 point, a partial match earns 0.5 points, and a wrong/invalid selection receives 0 points.

Related work Abductive reasoning with LLMs has received some notable contributions in recent years. Early benchmarks, such as aNLI, established the field by requiring the selection of the best explanatory bridge between observations (Bhagavathula et al., 2020). With the rise of LLMs, generation of plausible hypotheses is feasible but often inconsistent, especially under ambiguous evidence (Pareschi, 2023). Genuine abduction may also be obscured by dataset artifacts, which can be compromised via appropriate prompting (Balepur et al., 2024). In parallel, abduction is reframed as verifiable reasoning, promoting faithful hypothesis evaluation and robustness to noise (He and Lu, 2024). Additional work evaluates abduction in more formal logical settings (e.g., syllogistic forms), revealing systematic biases and limitations that persist even in advanced LLMs (Abe et al., 2025). At the same time, works emphasizing abduction as a stress test for LLMs show that generating explanations under uncertainty is particularly challenging compared to other forms of reasoning (Xu et al., 2025), while evaluating explanations within diverse viewpoints provides insightful probing for LLM failures (He et al., 2025). In total, robust abductive inference persists as a key challenge, despite promising evidence arising from prompting and task design (Dougrez-Lewis et al., 2025).

3 System Overview

Our proposed system follows a three-stage pipeline, illustrated in Figure 1: 1) graph-based retrieval for distractor filtering, 2) an LLM-based abductive reasoner with reflectively optimized prompt design, 3) post-hoc consistency enforcement.

Component	Description
Dense embeddings	Cohere Embed v4 (1024-dim)
Sparse retrieval	BM25+ (Robertson and Zaragoza, 2009) with $3 \times$ entity boosting
Hybrid weighting	$\alpha=0.7$ dense + 0.3 sparse
Entry points	3 dense + 2 sparse (deduplicated)
Traversal	Full connected component

Table 1: Hybrid document graph retrieval configuration.

3.1 Retrieval: Distractor Filtering

Hybrid Document Graphs. For each topic, we construct a document similarity graph $G = (V, E)$ (Figure 2) where each node $v \in V$ is a context document and edges encode hybrid similarity combining dense and sparse (BM25+ with entity boosting) signals (Lewis et al., 2020; Luan et al., 2021):

$$w(d_i, d_j) = \alpha \cdot \text{sim}_{\text{sem}}(d_i, d_j) + (1 - \alpha) \cdot \text{sim}_{\text{lex}}(d_i, d_j)$$

$w(d_i, d_j)$ is the hybrid edge weight between documents d_i, d_j , sim_{sem} denotes cosine similarity over dense embeddings, sim_{lex} is BM25+ lexical similarity with entity boosting, and $\alpha=0.7$ (Table 1).

Unlike entity-centric GraphRAG (Edge et al., 2024), which extracts named entities as nodes, our graph preserves intact documents, retaining the full narrative context needed for causal reasoning. At query time, we identify *entry points*, e.g. the top-ranked documents by each signal (3 dense, 2 sparse), deduplicated so that a document retrieved by both signals counts once. From these seeds, we traverse the graph via breadth-first search (BFS) over the full connected component, collecting every reachable document. This prioritizes recall over precision: a missing document can break a multi-hop causal chain, whereas extra documents add manageable noise. The traversal enables multi-hop evidence discovery, where indirectly connected documents may provide critical causal evidence (Wang et al., 2025).

Topic-Wide Aggregation. Since questions within the same topic retrieve heavily overlapping document sets, we aggregate retrieved documents across all questions within each topic into a single *topic-wide context*. The first question per topic populates the cache, and subsequent questions reuse it (Anthropic, 2024), yielding a 91% cache hit rate and 87% cost reduction (Table 12). Retrieval selects 73% of available

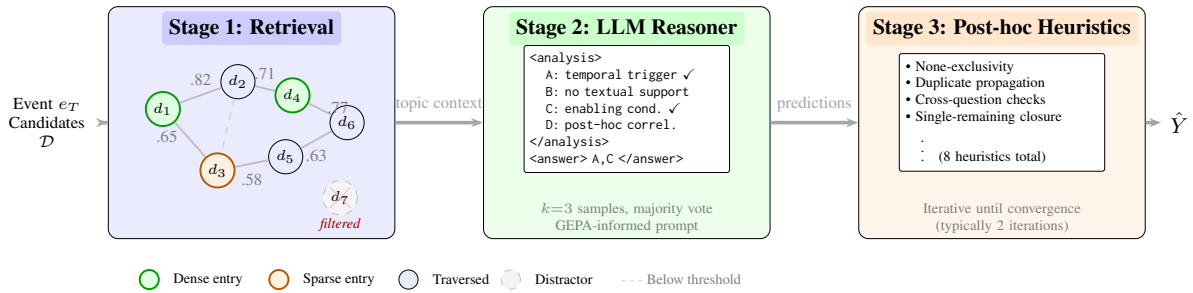


Figure 1: System pipeline. Stage 1 constructs a hybrid document graph (Figure 2), selects dense/sparse entry points, retrieves the connected component, and filters disconnected distractors. Stage 2 performs structured analysis-before-answer prompting with self-consistency. Stage 3 applies eight post-hoc consistency heuristics.




Figure 2: Hybrid document-graph retrieval in three steps. **Step 1:** Build a hybrid similarity graph ($\alpha=0.7$ dense + 0.3 sparse); disconnected documents (d_9-d_{12}) are potential distractors. **Step 2:** At query time, pick entry points from dense and sparse signals (3+2, deduplicated). **Step 3:** Retrieve the full connected component from the seeds, filter disconnected documents, and pass the selected topic context to the LLM reasoner.

documents per topic, excluding 27% as likely distractors (Table 13).

3.2 Inference

Structured Prompting. We use an XML-structured prompt with an analysis-before-answer format that forces the LLM to verbalize its reasoning before committing to a final answer. This acts as a structured scratchpad (Nye et al., 2021), eliciting chain-of-thought (CoT) reasoning (Wei et al., 2022) within a constrained output format: an `<analysis>` block containing per-option reasoning followed by an `<answer>` block restricted to valid option letters. The prompt specifies explicit reasoning criteria (requiring direct textual support and logical sufficiency), selection rules (supporting multi-label answers), and quality checks (enforcing disciplined evaluation per candidate).

Prompt Design. As a prompt optimizer, we implement GEPA (Agrawal et al., 2026) via DSPy (Khatab et al., 2024) to explore the space of effective prompt formulations through automatic reflective prompt evolution. Rather than deploying GEPA-optimized prompts directly as black-box

solutions, which raises concerns about data contamination and overfitting to the optimization set, we use the evolved prompts as a source of insight: GEPA-evolved prompts are never used at inference time. The structural heuristics surfaced by GEPA (e.g., single-step causal reasoning, explicit causal language prioritization, duplicate option handling) inform the final manually-curated prompt as transferable design principles, while the two deterministic dataset invariants it independently uncovered, duplicate option consistency and “None” mutual exclusivity, are implemented as post-hoc verification heuristics (Section 3.3) rather than prompt instructions, since logical invariants are more reliably enforced programmatically (Appendix C.2).

Self-Consistency. We sample $k=3$ responses at temperature $\tau=1.0$ and aggregate via per-option majority voting (Wang et al., 2023; Taubenfeld et al., 2025), with an inclusion threshold of 0.5. For questions containing “None of the others”, we apply conflict resolution where the higher-voted category prevails. We compare majority voting against union, intersection, and confidence-weighted aggregation (Appendix D.5).

3.3 Post-Hoc Consistency Enforcement

We apply eight de-hallucination and consistency-aware heuristics that enforce logical invariants the LLM may violate (e.g. selecting mutually exclusive options simultaneously or assigning different truth values to identical candidates) iteratively until convergence (typically 2 iterations). Details and per-heuristic impact are reported in Appendix A.9.

4 Experimental Setup

We evaluate our system on the SemEval 2026 Task 12 development set, comparing 18 model configurations under zero-shot structured prompting with optional inference enhancements.

Dataset The development set comprises 400 questions across 36 topics (775 documents), where one to four out of the four candidate explanations may be correct. Notably, 43.6% of questions have multiple correct answers. Questions are grouped by topic, with sibling questions sharing document context, a property exploited for caching and consistency enforcement (Sections 3.1, 3.3; App. A).

Evaluation Metric We use the task’s multi-label classification metric with partial credit. The final score is computed as the unweighted arithmetic mean across all questions: a model predicts a subset $\hat{y} \subset \{A, B, C, D\}$ per question, compared against the gold label set y via set-based evaluation:

- **Exact match** ($\hat{y} = y$): score = 1.0
- **Partial match** ($\hat{y} \subset y, \hat{y} \neq \emptyset$): score = 0.5
- **All other cases**: score = 0.0

Partial credit rewards models that correctly identify at least one valid cause without requiring full recall, thereby providing a finer-grained evaluation signal. The final score is the arithmetic mean across all questions, weighting each equally regardless of gold label count.

Models We evaluate models accessed through AWS Bedrock, Google’s Gemini API, and the OpenAI API. Table 2 summarizes the evaluated configurations. *Standard* LLMs use default decoding, while *reasoning* LLMs employ extended thinking (Claude’s `extended_thinking`, Gemini’s high thinking level, OpenAI’s `xhigh` reasoning effort, and DeepSeek R1’s native CoT). All models employ zero-shot structured prompting (Section 3.2), with self-consistency ($k=3, \tau=1.0$) selectively applied due to cost-related concerns.

Family	Standard	Reasoning
Claude	Haiku 4.5, Sonnet 4, Sonnet 4.5, Opus 4.0, Opus 4.5	Sonnet 4 Thinking, Sonnet 4.5 Thinking
Gemini	—	Flash 3 Preview
OpenAI	—	GPT-5.2 (xhigh)
DeepSeek	V3.1	V3.1 Think, R1
Llama	3.3-70B	—
Kimi	—	K2 Think
Qwen	3-32B	—

Table 2: Model configurations evaluated. *Reasoning* models use extended thinking or chain-of-thought.

	Model	Score
<i>Standard Models</i>	Claude Haiku 4.5	0.788
	Claude Sonnet 4	0.781
	Claude Sonnet 4.5	0.790
	Claude Opus 4.0	0.770
	Claude Opus 4.5	0.801
	DeepSeek V3.1	0.709
	Llama 3.3-70B	0.633
	Qwen3-32B	0.611
<i>Reasoning Models</i>	Claude Sonnet 4 Thinking	0.794
	Claude Sonnet 4.5 Thinking	0.828
	Gemini 3 Flash Preview	0.804
	GPT-5.2	0.810
	DeepSeek V3.1 Think	0.774
	DeepSeek R1	0.718
	Kimi K2 Think	0.776

Table 3: Zero-shot baseline results (dev set). **Bold** marks the best score within each group under the partial-credit metric (Section 4).

5 Results and Analysis

Table 3 reports zero-shot dev performance across 15 LLMs (scores ranging between 0.611-0.828), highlighting the gap between open-weight and frontier models. Extended thinking consistently helps: Sonnet 4.5 Thinking (0.828) surpasses Sonnet 4.5 by 3.8 pp. GPT-5.2 (0.810) and Gemini 3 Flash (0.804) are competitive, while DeepSeek models underperform relative to their scale and Qwen3-32B (0.611) confirms the mid-size open-weight capacity gap. Bootstrapped significance tests yield $p > 0.05$ for all pairwise frontier-model comparisons on the 400-question dev set (Appendix D.4), so fine-grained score differences among top systems should be interpreted as suggestive rather than statistically significant.

Graph-based retrieval (Section 3.1) provides marginal accuracy gains for frontier LLMs but yields a +9 pp improvement for Haiku 3.5 (Ta-

	Configuration	Base	+ Post-hoc
<i>Individual Models</i>	Claude Son. 4.5 Thinking	0.904	0.952
	GPT-5.2	0.912	0.949
	Gemini 3 Flash	0.907	0.943
<i>Self-Consistency</i>	SC: Sonnet 3× ($\theta=0.50$)	0.902	0.948
	SC: Gemini 5× ($\theta=0.50$)	0.902	0.943
<i>Ensemble</i>	Sonnet + GPT + Gemini	0.926	0.952

Table 4: Test set results. Post-hoc = deterministic consistency enforcement via heuristics. SC = self-consistency with majority voting; θ = inclusion threshold.

ble 13), consistent with findings that retrieval disproportionately benefits smaller models (Xu et al., 2024; Raptopoulos et al., 2025), while reducing cost by 87% via topic-wide caching, mitigating accuracy degradation (Liu et al., 2024; Du et al., 2025) and inference cost (Li et al., 2024) of long-context processing. Self-consistency ($k=3$, $\tau=1.0$) yields modest gains (e.g., +1.6 pp for Sonnet 4.5); majority voting across three systems reaches 0.811 (Table 17, Appendix D.5). Post-hoc heuristics (Section 3.3) produce the single largest gain: Sonnet 4.5 Thinking improves from 0.828 to 0.884 (+5.6 pp), with 85.4% of corrections being genuine improvements. We emphasize that these gains reflect the enforcement of logical invariants on model outputs (e.g., identical options must share truth values, “None” is mutually exclusive) rather than improvements to the underlying abductive inference. Cross-question checks contribute the largest individual gains (details in Appendix A.9).

Test Set Table 4 reports results on the held-out test set (612 questions, 18.3% multi-answer). Post-hoc heuristics again yield the largest gains (+4.8 pp for Claude, +3.7 pp for GPT-5.2), consistent with the dev-set pattern, while self-consistency alone has minimal impact: the gain materializes only when combined with post-hoc heuristics (0.943, +4.1 pp). The multi-model ensemble reaches 0.926 without heuristics and matches the best individual model with heuristics (Claude Sonnet 4.5 Thinking at 0.952). The higher test-set scores reflect the lower multi-answer prevalence (18.3% vs. 47.5%), consistent with the inherently greater difficulty of multi-answer questions.

5.1 Error Analysis

LLMs exhibit systematic biases in causal inference, including the *post-hoc* fallacy (Joshi et al., 2024) and degraded multi-hop reasoning (Chi et al., 2024;

System	Single	Multi	Δ
GPT-5.2	91.0	52.6	−38.4
Claude Son. 4.5 Thinking	88.6	60.5	−28.1
Gemini 3 Flash	85.2	63.7	−21.5

Table 5: Exact-match accuracy (%) by answer cardinality across three frontier systems. Multi-answer questions (190/400, 47.5%) are inherently harder due to the combinatorial requirement.

Yu et al., 2025). Cross-validation across all 14 models (7 families) yields Fleiss’ $\kappa = 0.690$ (Fleiss, 1971) (*substantial* agreement), with within-family agreement ($\kappa = 0.794$) exceeding cross-family agreement ($\kappa = 0.661$; Figure 17). Only 3 questions are truly *unbreakable* (all 14 models score zero), but 42 receive no exact match from any model, 38 of which are multi-answer under-selection failures (Figure 20). An oracle selecting the best model per question achieves 0.895, indicating 6.7 pp headroom from model complementarity (Appendix D.4).

Single vs. Multi-Answer Performance. All models show lower multi-answer accuracy (Table 5), with gaps of 25-56 pp across 14 models. Total under-selections (1,389) vastly exceed over-selections (52), confirming conservative cause selection as the dominant bias (Figure 21).

Failure Taxonomy. Of the 42 no-exact-match questions (3 truly unbreakable; Table 14, Appendix D.1), models select on average 1.2 options when gold requires 2.4, a 51% cause-count reduction, and in 83% of cases the consensus is a *single option* (Appendix D.3). Three shared inductive biases persist across all seven model families: (1) *causal chain incompleteness* (Chi et al., 2024; Yu et al., 2025), the most prevalent pattern (18/42), where models select one link of a multi-step causal chain and omit the rest; (2) *proximate cause preference* (Joshi et al., 2024) (11/42), which favors the most recent antecedent over enabling conditions; and (3) *saliency bias* (9/42), where models select dramatic, newsworthy causes over subtler contributing factors. These biases are correlated across families (error overlap ratio 2.0× above independence), suggesting shared pretraining priors rather than prompt-specific artifacts (Chi et al., 2024).

6 Conclusion

Our proposed three-stage system (graph-based retrieval, GEPA-informed prompt design, deterministic consistency enforcement via heuristics) ranks first on the evaluation-phase leaderboard (0.95), with post-hoc heuristics providing the largest dev-set gain (+5.6 pp). Our error analysis across 14 models (7 families) identifies three shared inductive biases: causal chain incompleteness (18/42), proximate cause preference (11/42), and salience bias (9/42), which together produce a single-cause default (51% cause-count reduction, 83% of failures converging on one option). With an oracle upper bound of 0.895 (best of 14 models per question), no single family dominates across all question types, suggesting that targeted ensemble or multi-agent architectures, together with multi-expert causal graph construction (Koupae et al., 2025) for explainable reasoning traces could close much of the remaining gap.

Acknowledgements

This work was carried out within the framework of the Pharos AI Factory project, funded by the European High-Performance Computing Joint Undertaking (EuroHPC JU) under Grant Agreement No. 101234269 as part of the Horizon Europe and by the Greek Public Investments Program programme.

Limitations

Generalizable vs. Benchmark-Specific Heuristics. Our post-hoc heuristics split into two groups with different transferability. The local invariants, duplicate option consistency (R2) and “None” mutual exclusivity (R1), encode general logical properties that apply to any multiple-choice benchmark containing these structures, and were surfaced automatically by GEPA without manual engineering, suggesting that reflective optimization pressure offers a principled route to discovering analogous invariants on new tasks. In contrast, the cross-question propagation rules (R4, R6, R7) exploit the sibling-question groupings of this benchmark, where multiple questions share the same target event, and would require adaptation or omission on tasks without such grouping. Separating the two clarifies that while the largest gain component is partially benchmark-specific, its mechanism (logical-invariant enforcement on model outputs) is reusable.

Model Coverage. Our evaluation focuses on frontier-scale to mid-size (32B) models; smaller open-weight models remain unexplored.

References

- Hirohiko Abe, Risako Ando, Takanobu Morishita, Kentaro Ozeki, Koji Mineshima, and Mitsuhiro Okada. 2025. [Abductive reasoning with syllogistic forms in large language models](#). In *Human and Artificial Rationalities. Advances in Cognition, Computation, and Consciousness: Third International Conference, HAR 2024, Paris, France, September 17–20, 2024, Proceedings*, page 3–17, Berlin, Heidelberg. Springer-Verlag.
- Lakshya A Agrawal, Shangyin Tan, Dilara Soyulu, Noah Ziem, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alexandros G. Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. 2026. [GEPA: Reflective prompt evolution can outperform reinforcement learning](#). In *Proceedings of the Fourteenth International Conference on Learning Representations (ICLR)*. Oral presentation.
- Anthropic. 2024. [Introducing contextual retrieval](#). <https://www.anthropic.com/news/contextual-retrieval>. Accessed: 2026-02-06.
- Nishant Balepur, Abhilasha Ravichander, and Rachel Rudinger. 2024. [Artifacts or abduction: How do LLMs answer multiple-choice questions without the question?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10308–10330, Bangkok, Thailand. Association for Computational Linguistics.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen tau Yih, and Yejin Choi. 2020. [Abductive commonsense reasoning](#). In *International Conference on Learning Representations*.
- Pengfei Cao, Yubo Chen, Mingxuan Yang, Chenlong Zhang, Mingxuan Liu, Kang Liu, and Jun Zhao. 2026. [SemEval-2026 task 12: Abductive event reasoning: Towards real-world event causal inference for large language models](#). In *Proceedings of the 20th International Workshop on Semantic Evaluation (SemEval-2026)*.
- He Chang, Chenchen Ye, Zhulin Tao, Jie Wu, Zhengmao Yang, Yunshan Ma, Xianglin Huang, and Tat-Seng Chua. 2025. [A comprehensive evaluation of large language models on temporal event forecasting](#). *Preprint*, arXiv:2407.11638.
- Haoang Chi, He Li, Wenjing Yang, Feng Liu, Long Lan, Xiaoguang Ren, Tongliang Liu, and Bo Han. 2024. [Unveiling causal reasoning in large language](#)

- models: Reality or mirage? In *Advances in Neural Information Processing Systems*, volume 37.
- John Dougrez-Lewis, Mahmud Elahi Akhter, Federico Ruggeri, Sebastian Löbbers, Yulan He, and Maria Liakata. 2025. [Assessing the reasoning capabilities of LLMs in the context of evidence-based claim verification](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 20604–20628, Vienna, Austria. Association for Computational Linguistics.
- Yufeng Du, Minyang Tian, S. Ronanki, Subendhu Rongali, S. Bodapati, A. G. Galstyan, Azton Wells, Roy Schwartz, E. A. Huerta, and Hao Peng. 2025. [Context length alone hurts LLM performance despite perfect retrieval](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitan, Robert Osazuwa Ness, and Jonathan Larson. 2024. [From local to global: A graph RAG approach to query-focused summarization](#). *arXiv preprint arXiv:2404.16130*.
- Joseph L. Fleiss. 1971. [Measuring nominal scale agreement among many raters](#). *Psychological Bulletin*, 76(5):378–382.
- Luciano Floridi, Jessica Morley, Claudio Novelli, and David Watson. 2025. [What kind of reasoning \(if any\) is an llm actually doing? on the stochastic nature and abductive appearance of large language models](#). *Preprint*, arXiv:2512.10080.
- Debanjan Ghosh, Kai North, Ekaterina Kochmar, Mamoru Komachi, and Marcos Zampieri, editors. 2026. *Proceedings of the 20th International Workshop on Semantic Evaluation (SemEval-2026)*. Association for Computational Linguistics, San Diego, United States.
- Panagiotis Giadikiaroglou, Maria Lymperaiou, Giorgos Filandrianos, and Giorgos Stamou. 2024. [Puzzle solving using reasoning of large language models: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11574–11591, Miami, Florida, USA. Association for Computational Linguistics.
- Jinwei He and Feng Lu. 2024. [Causejudger: Identifying the cause with llms for abductive logical reasoning](#). *Preprint*, arXiv:2409.05559.
- Kaiyu He and Zhiyu Chen. 2025. [From reasoning to learning: A survey on hypothesis discovery and rule learning with large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Kaiyu He, Peilin Wu, Mian Zhang, Kun Wan, Wentian Zhao, Xinya Du, and Zhiyu Chen. 2025. [Gear: A general evaluation framework for abductive reasoning](#). *Preprint*, arXiv:2509.24096.
- Nitish Joshi, Abulhair Saparov, Yixin Wang, and He He. 2024. [LLMs are prone to fallacies in causal inference](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10553–10569. Association for Computational Linguistics.
- Zixuan Ke, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li, Chengwei Qin, PeiFeng Wang, Silvio Savarese, Caiming Xiong, and Shafiq Joty. 2025. [A survey of frontiers in LLM reasoning: Inference scaling, learning to reason, and agentic systems](#). *Transactions on Machine Learning Research*. Survey Certification.
- Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2024. [DSPy: Compiling declarative language model calls into state-of-the-art pipelines](#). In *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*.
- Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. 2024. [Benchmarking cognitive biases in large language models as evaluators](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 517–545. Association for Computational Linguistics.
- Mahnaz Koupaee, Xueying Bai, Mudan Chen, Greg Durrett, Nathanael Chambers, and Niranjana Balasubramanian. 2025. [Causal graph based event reasoning using semantic relation experts](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 26169–26199. Association for Computational Linguistics.
- Klaus Krippendorff. 2011. [Computing Krippendorff’s alpha-reliability](#). *Communication Methods and Measures*, 5(1):1–12.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33 (NeurIPS)*.
- Bobo Li, Xudong Han, Jiang Liu, Yuzhe Ding, Liqiang Jing, Zhaoqi Zhang, Jinheng Li, Xinya Du, Fei Li, Meishan Zhang, Min Zhang, Aixin Sun, Philip S. Yu, and Hao Fei. 2025. [Event extraction in large language model](#). *Preprint*, arXiv:2512.19537.
- Zhuowan Li, Cheng Shi, Xinwei Gao, Shuo Chen, Yiwei Zhu, Zelong Li, Yifei Shen, and Xiang Yan. 2024. [Retrieval augmented generation or long-context LLMs? a comprehensive study and hybrid approach](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.

- Hanmeng Liu, Zhizhang Fu, Mengru Ding, Ruoxi Ning, Chaoli Zhang, Xiaozhang Liu, and Yue Zhang. 2025. [Logical reasoning in large language models: A survey](#). *Preprint*, arXiv:2502.09100.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. [Sparse, dense, and attentional representations for text retrieval](#). *Transactions of the Association for Computational Linguistics*, 9:329–345.
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. 2021. [Show your work: Scratchpads for intermediate computation with language models](#). *arXiv preprint arXiv:2112.00114*.
- Remo Pareschi. 2023. [Abductive reasoning with the gpt-4 language model: Case studies from criminal investigation, medical practice, scientific research](#). *Preprint*, arXiv:2307.10250.
- Petros Raptopoulos, Giorgos Filandrianos, Maria Lymperaiou, and Giorgos Stamou. 2025. [PAKTON: A multi-agent framework for question answering in long legal agreements](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 7959–7995, Suzhou, China. Association for Computational Linguistics.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Saurabh Srivastava, Sweta Pati, and Ziyu Yao. 2025. [Instruction-tuning LLMs for event extraction with annotation guidelines](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 13055–13071, Vienna, Austria. Association for Computational Linguistics.
- Hristo Tanev, Nicolas Stefanovitch, Tomáš Harmatha, and Diana F. Sousa. 2025. [Exploring the performance of large language models for event detection and extraction in the health domain](#). In *Proceedings of the 15th International Conference on Recent Advances in Natural Language Processing - Natural Language Processing in the Generative AI Era*, pages 1237–1247, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.
- Amir Taubenfeld, Tom Sheffer, Eran Ofek, Amir Feder, Ariel Goldstein, Zorik Gekhman, and Gal Yona. 2025. [Confidence improves self-consistency in LLMs](#). In *Findings of the Association for Computational Linguistics: ACL 2025*.
- Nengbo Wang, Xiaotian Han, Jagdip Singh, Jing Ma, and Vipin Chaudhary. 2025. [CausalRAG: Integrating causal graphs into retrieval-augmented generation](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 22680–22693, Vienna, Austria. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35 (NeurIPS)*.
- Fangzhi Xu, Qika Lin, Jiawei Han, Tianzhe Zhao, Jun Liu, and Erik Cambria. 2025. [Are large language models really good logical reasoners? a comprehensive evaluation and beyond](#). *IEEE Trans. on Knowl. and Data Eng.*, 37(4):1620–1634.
- Peng Xu, Wei Ping, Xianchao Wu, Lawrence McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. 2024. [Retrieval meets long context large language models](#). In *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*.
- Longxuan Yu, Delin Chen, Siheng Xiong, Qingyang Wu, Dawei Li, Zhikai Chen, Xiaoze Pan, and Liangming Pan. 2025. [CausalEval: Towards better causal reasoning in language models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 12512–12540. Association for Computational Linguistics.

A Exploratory Data Analysis

A.1 Dataset formulation

The questions file is represented in a jsonl object with the following format for a specific instance:

```

{
  "topic_id": 11,
  "id": "q-101",
  "target_event": "Short description
of an observed event.",
  "option_A": "Candidate explanation
A.",
  "option_B": "Candidate explanation
B.",
  "option_C": "Candidate explanation
C.",
  "option_D": "Candidate explanation
D.",
  "golden_answer": "A,B"
}

```

Example. Input instance object.

Consequently, the context comprises records in the docs file, with a jsonl instance as following:

```

{
  "topic_id": 11,
  "topic": "OpenAI releases ChatGPT",
  "docs": [
    {
      "title": "Article title",
      "id": "doc-001",
      "link": "https://example.com",
      "snippet": "Short summary of
the document.",
      "source": "News source",
      "imageUrl": "Base64imageUrl",
      "content": "Full document text"
    }
  ]
}

```

Example. Background context for an event.

A.2 Overview and Statistics

The dataset comprises 2,831 questions distributed across multiple splits as shown in Table 6. Each question presents a target event and four candidate explanations, where one or more may be correct causes.

Figure 3 illustrates the composition of each data split, showing the distribution of questions, topics, and documents.

Table 7 shows representative examples from the dataset, one for each answer cardinality. Each question presents a target event and four candidate options; the task is to identify all plausible causes.

Statistic	Train	Dev	Test
Questions	1,819	400	612
Topics	36	36	24
Documents	775	775	405
Multi-answer (%)	42.7	47.5	–
Avg. Q/Topic	50.5	11.1	25.5
Avg. Doc Length	1188	1188	899

Table 6: Dataset statistics across splits. Multi-answer percentage indicates questions with more than one correct answer.




Figure 3: Dataset composition across splits.

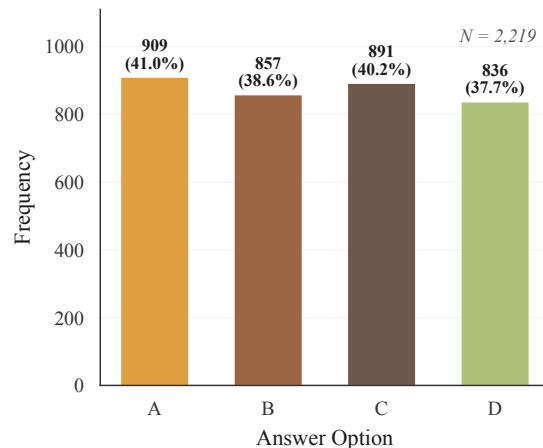


Figure 4: Answer frequency by option.

A.3 Answer Distribution

A distinctive characteristic of this dataset is its multi-label nature: 43.6% of questions have multiple correct answers. Table 8 details the answer distribution. The answer options (A, B, C, D) are roughly balanced, indicating no positional bias in the correct answers. Figure 4 shows the frequency distribution, and Figure 5 visualizes the answer cardinality breakdown. Figure 6 illustrates the most

Target Event	Candidate Options	Gold
Ontario placed a 25% tax on electricity exports to the U.S.	(A) Trump placed 25% tariffs on imports from Canada. (B) McKinley was shot by anarchist Czolgosz in 1901. (C) The 16th Amendment was ratified in 1913. (D) McKinley died on Sept. 14, 1901.	A
The Chinese government criticized Japan’s plan to release Fukushima wastewater.	(A) IAEA confirmed the release would have negligible radiological effect. (B) Japan’s cabinet approved a treated-water plan in 2021. (C) Japan announced its plan to release treated wastewater into the sea. (D) TEPCO has used water to cool nuclear fuel rods since 2011.	B,C
Facebook changed its corporate name to Meta at the annual Connect event.	(A) Zuckerberg introduced Meta at the Connect 2021 event. (B) Facebook announced Bosworth’s promotion to CTO. (C) Facebook plans to change its name to reflect metaverse focus. (D) Zuckerberg announced Facebook was betting on the metaverse.	A,C,D
Goldman Sachs estimated AI could automate 300 million jobs worldwide.	(A) OpenAI released ChatGPT Plus and later introduced GPT-4 Turbo. (B) Stack Overflow banned ChatGPT-generated answers. (C) None of the others are correct causes. (D) OpenAI launched DALL-E 2 beta in July 2022.	C

Table 7: Representative examples illustrating each answer cardinality. Correct options are shown in **bold**. Option texts are abbreviated for space.

Category	Count	Percentage
Option A selected	909	41.0%
Option B selected	857	38.6%
Option C selected	891	40.2%
Option D selected	836	37.7%
Single answer	928	41.8%
Two answers	660	29.7%
Three answers	307	13.8%
“None sufficient”	324	14.6%

Table 8: Answer distribution analysis on the combined train and dev sets.




Figure 5: Distribution of answer cardinality.




Figure 6: Most common multi-answer combinations.

Metric	Value
Label cardinality	1.574
Label density	0.394
Multi-label rate	43.6%

Table 9: Multi-label characteristics of the dataset.

A.4 Multi-Label Characteristics

The multi-label nature of the dataset poses unique challenges for evaluation. We compute standard multi-label metrics, such as label cardinality, density and multi-label rate (Table 9).

The label cardinality of 1.574 indicates that on average, each question has approximately 1.5 correct answers. This is higher than typical single-

common multi-answer combinations, revealing patterns in how answers co-occur.




Figure 7: Answer co-occurrence matrix.

label QA datasets and reflects the inherent ambiguity in causal reasoning, since multiple events can plausibly cause a given outcome.

Figure 7 shows the co-occurrence matrix of answer labels, revealing which answer pairs frequently appear together in multi-label questions. There are no significant discrepancies in label combination frequencies; combinations containing label A are more prevalent, since A is the most frequent answer option in the dataset (Figure 4).

A.5 Topic Analysis

The dataset covers 36 real-world news topics spanning significant events from 2016 to 2025, including political events (Brexit, US elections), technological developments (ChatGPT, Huawei), and natural disasters (California wildfires, earthquakes).

Figure 8 provides a comprehensive view of topic characteristics, showing the relationship between question count, complexity, and context size.

A.6 Document Corpus

Each topic is associated with a set of retrieved documents providing context for causal reasoning. Figure 9 shows the distribution of document lengths.

Document lengths vary substantially, from 11 to 44,818 words (median: 719 words), reflecting the heterogeneous, long-tailed nature of web-sourced content. This variation presents challenges for models with limited context windows. Moreover, the long-tailed nature of this curve may trigger length biases in LLMs, showing preference to the most frequently occurring length over truly consulting the documents’ context.

A.7 Linguistic Analysis

Target events are concise descriptions averaging 13.3 words (Figure 10). Candidate options have similar lengths across positions (A through D), ranging from 12.6 to 12.8 words on average. This uniform length distribution across response candidates helps reduce possible length biases in response selection.

A.8 Logical Consistency Properties

We identify two logical consistency properties inherent in the question structure that LLMs may occasionally violate in their predictions.

Mutual Exclusivity of “None of the Others.”

14.6% of questions include a “None of the others are correct causes” option. By definition, this option is mutually exclusive with all other candidates and selecting it alongside other options constitutes a logical contradiction. However, LLMs are not constrained to respect this exclusivity, and may produce predictions that include “None” together with other candidates.

Duplicate Option Consistency. 22.9% of questions contain duplicate answer options (identical text appearing in two or more positions), an artifact of the dataset construction process. Identical options necessarily share the same truth value and thus selecting one but not the other would be logically inconsistent. These cases expose instances where LLM predictions lack internal consistency.

These observations motivate a set of post-hoc consistency heuristics applied to verify or correct model predictions (Section A.9).

A.9 Post-Hoc Prediction Refinement

We apply a set of deterministic post-hoc consistency heuristics that enforce logical invariants the LLM may violate during inference. These heuristics operate entirely on the model’s own predictions, acting as restorative heuristics, requiring no gold labels or additional LLM calls.

The core heuristics address two types of inconsistency (Figure 11): (1) mutual exclusivity violations, where “None of the others” is predicted alongside other options (14.6% of questions); and (2) duplicate inconsistencies, where identical options receive different predictions (21.6% of questions). Additional heuristics enforce structural invariants across sibling questions, i.e. questions sharing the same target event (63.3% of sibling groups contain

ID	Topic Description	Questions	Docs	Multi%	Context
19	Historic winter storm paralyzes Texas, causin...	132	24	53.0	18,289
30	Killing of George Floyd by police sparks Blac...	113	25	57.5	70,271
23	Japan begins discharging treated Fukushima nu...	111	24	45.0	38,702
24	Halloween crowd crush in Seoul's Itaewon kill...	111	28	45.0	14,375
22	Iranian General Qasem Soleimani killed by US ...	110	20	36.4	12,160
32	Taliban regains control of Afghanistan after ...	109	27	55.0	24,098
21	Myanmar military stages coup, detains Aung Sa...	106	29	70.8	27,639
3	Trump supporters storm US Capitol	91	19	57.1	15,253
25	Yellow Vest movement sparks widespread protes...	77	18	51.9	17,805
13	Facebook changes company name to Meta to focu...	76	21	46.1	26,242

Table 10: Topic-level statistics (top 10 topics by question count). Full table in supplementary materials.

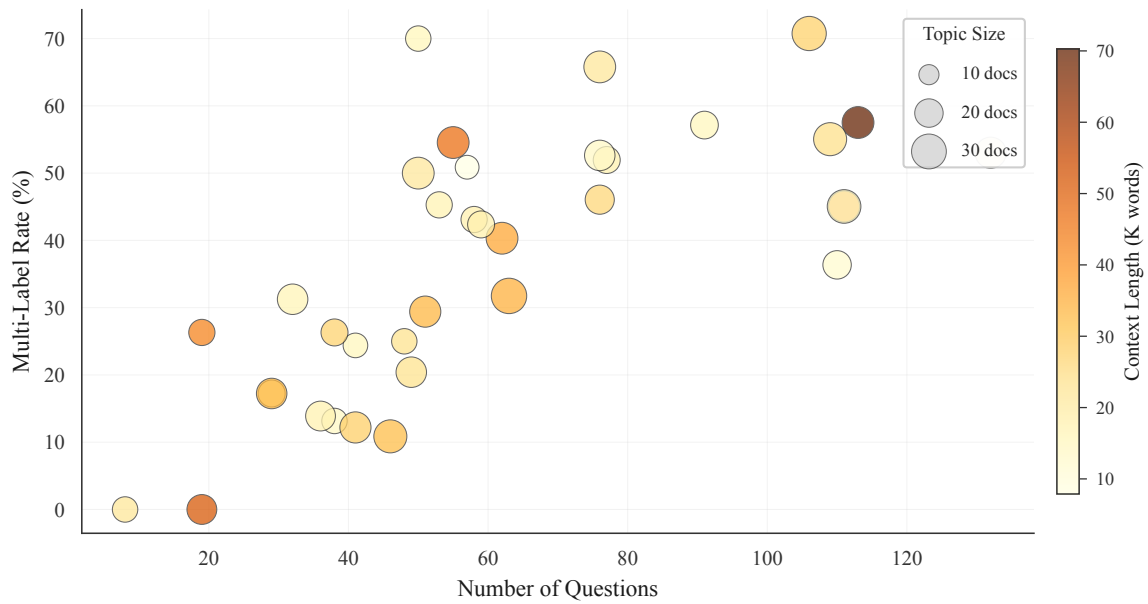


Figure 8: Topic characteristics: questions vs complexity (bubble size = document count, color = context length).

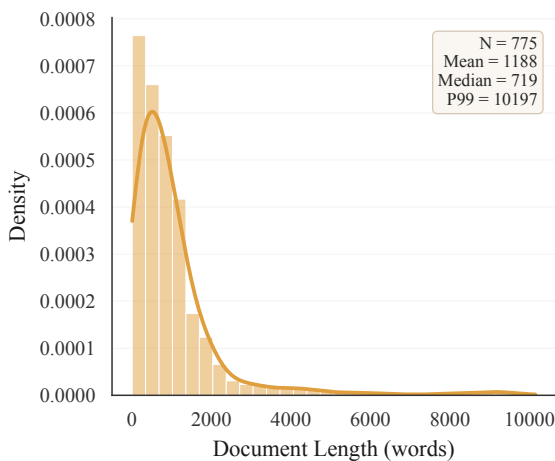


Figure 9: Distribution of document lengths (99th percentile shown).

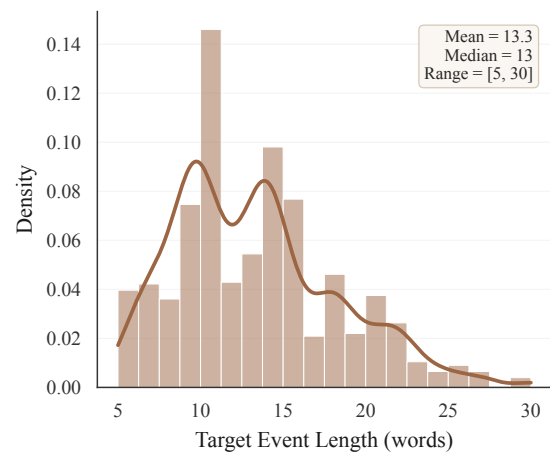


Figure 10: Target event length distribution.

2+ questions; Figure 12). Since identical option text must carry the same truth value regardless of which question it appears in, and since “None of the

others” logically excludes its co-occurring candidates across all siblings, these invariants propagate the logical consequences of the model’s within-question judgments to related questions.

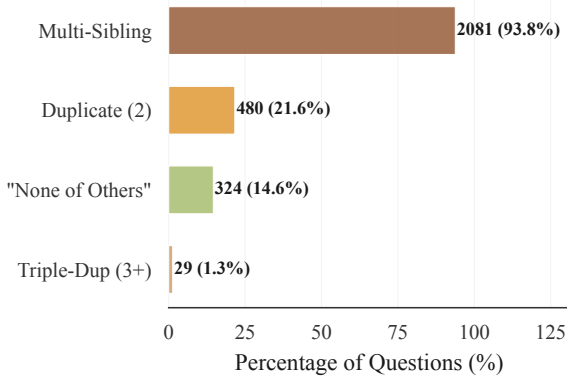


Figure 11: Frequency of structural patterns addressed by post-hoc consistency via heuristics.

Table 11 describes all eight heuristics explored.

Iterative Application. Heuristics from Table 11 are applied iteratively until convergence (fixed-point), as resolving one inconsistency may reveal another. For example, enforcing “None” exclusivity (R1) triggers cross-question exclusion (R6), which may leave a single valid option, triggering propagation (R7) or closure (R8). Similarly, enforcing duplicate consistency (R2) triggers cross-sibling propagation (R4), which may trigger R2 again in receiving questions. Convergence typically requires 2 iterations.

Figure 13 shows the cumulative score improvement as heuristics are added incrementally. Cross-question consistency heuristics (R4, R6) provide the largest gains, confirming that enforcing structural invariants across siblings is the most effective signal. On the dev set, these consistency heuristics improve accuracy from 0.8275 to 0.8838 (+6.8%), with 85.4% of corrections being improvements and 14.6% neutral (changing one incorrect prediction to another also incorrect response), with zero cases where a correct prediction was made incorrect. The absence of degradations indicates that the structural invariants enforced by these heuristics (duplicate truth-value consistency and mutual exclusivity of “None”) are well-aligned with the ground truth.

B Graph-Based Retrieval

This appendix formalizes the retrieval algorithm and provides visualization examples of the hybrid document graph retrieval system described in Section 3.1.

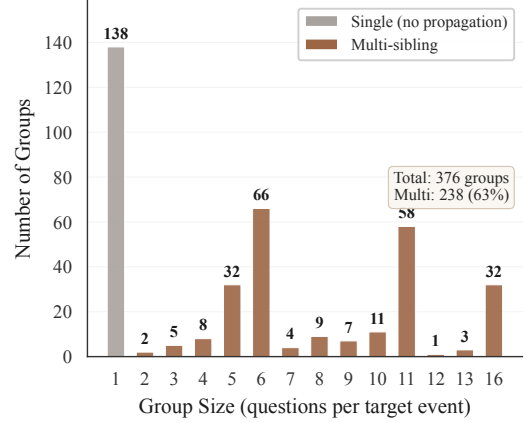


Figure 12: Sibling group size distribution. Groups of 2+ questions enable cross-question consistency enforcement.

Algorithm 1 Hybrid Graph Retrieval

Require: Document graph $G = (V, E)$ with hybrid edge weights w ; query q ; dense embedder ϕ ; BM25+ index \mathcal{B} ; parameters $k_{\text{sem}}, k_{\text{lex}}, \tau$

Ensure: Selected document set $\mathcal{S} \subseteq V$

— **Entry Point Identification** —

- 1: $\mathcal{E}_{\text{sem}} \leftarrow$ top- k_{sem} docs by $\cos(\phi(q), \phi(d))$
- 2: $\mathcal{E}_{\text{lex}} \leftarrow$ top- k_{lex} docs by $\mathcal{B}(q)$
- 3: $\mathcal{E} \leftarrow \mathcal{E}_{\text{sem}} \cup \mathcal{E}_{\text{lex}}$ ▷ Deduplicate

— **Connected-Component Traversal (BFS)** —

- 4: $\mathcal{S} \leftarrow \mathcal{E}$ ▷ Initialize with entry points
- 5: QUEUE $\leftarrow \mathcal{E}$
- 6: **while** QUEUE $\neq \emptyset$ **do**
- 7: $d \leftarrow$ QUEUE.dequeue()
- 8: **for each** neighbor d' s.t. $(d, d') \in E$ and $w(d, d') \geq \tau$ **do**
- 9: **if** $d' \notin \mathcal{S}$ **then**
- 10: $\mathcal{S} \leftarrow \mathcal{S} \cup \{d'\}$
- 11: QUEUE.enqueue(d')
- 12: **end if**
- 13: **end for**
- 14: **end while**

— **Distractor Filtering** —

- 15: **return** \mathcal{S} ▷ Disconnected docs excluded
-

B.1 Retrieval Algorithm

Algorithm 1 formalizes the hybrid graph retrieval procedure described in Section 3.1.

The algorithm prioritizes *recall over precision*: a missing document can break a multi-hop causal chain, whereas extra documents add manageable noise. Entry points are identified via both semantic (dense embedding cosine similarity) and lexical

Heuristic	Name	Scope	Mechanism	Changes
R1	None Exclusivity	Local	Enforce mutual exclusivity of “None” option	6
R2	Duplicate Consistency	Local	Ensure identical options receive same prediction	13
R3	Over-selection Guard	Local	Resolve over-selection when all four options predicted	0
R4	Cross-Q Dup. Prop.	Cross-Q	Propagate duplicate consistency to sibling questions	20
R5	Triple Exclusion	Cross-Q	Exclude unique option from triple-duplicate siblings	1
R6	None Cross-Exclusion	Cross-Q	Extend None exclusivity to sibling questions	4
R7	Single-Valid Prop.	Cross-Q	Propagate last valid option to siblings	0
R8	Closure	Local	Select only remaining valid option	1

Table 11: Post-hoc consistency heuristics applied iteratively until convergence (fixed-point). Scope indicates whether the heuristic operates within a single question (Local) or across sibling questions sharing the same target event (Cross-Q). Changes reported on Claude Sonnet 4.5 (thinking), dev set.

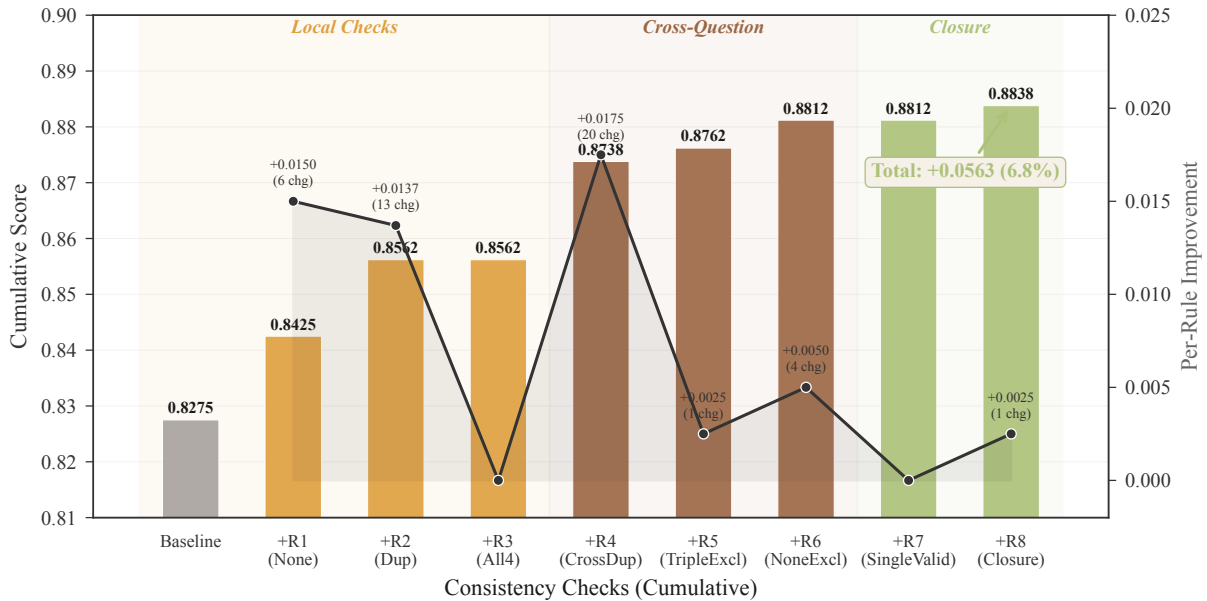


Figure 13: Cumulative accuracy improvement from post-hoc consistency heuristics on Claude Sonnet 4.5 (thinking) predictions. Background colors indicate heuristic scope: local (blue), cross-question (red), and closure (green). Black line shows per-rule improvement.

(BM25+ with $3\times$ entity boosting) signals, then deduplicated. The BFS traversal expands from these seeds through the full connected component, ensuring that indirectly related documents providing bridging causal evidence are included.

B.2 Document Graph and Query-Time Selection

For each topic, we construct a document similarity graph where nodes represent documents and edges encode hybrid similarity scores (weighted combination of dense and sparse signals). Figure 14 illustrates the complete retrieval process for Topic 7 (“South Korean President Yoon Suk-yeol faces impeachment crisis”), which contains 14 documents across 8 connected components.

Panel (a) shows the full graph structure with edges colored by their dominant similarity sig-

nal. At query time (b), dense and sparse retrieval identify entry points, then BFS traversal expands through the connected component. The retrieval selects 6 of 14 documents (43%) through three mechanisms: 3 semantic entry points (dense embedding matches), 2 lexical entry points (BM25+ matches), and 1 traversal-discovered document (reached via graph edges). The remaining 8 documents are excluded as likely distractors: they discuss related but distinct events that would add noise without providing direct causal evidence. This example demonstrates how graph traversal enables multi-hop evidence discovery: the traversal node (document 11) was not a direct match for the query but is connected to selected documents via the similarity graph, providing additional context that supports causal reasoning. Across the full dev set, retrieval selects 73% of documents on average, with a 91%

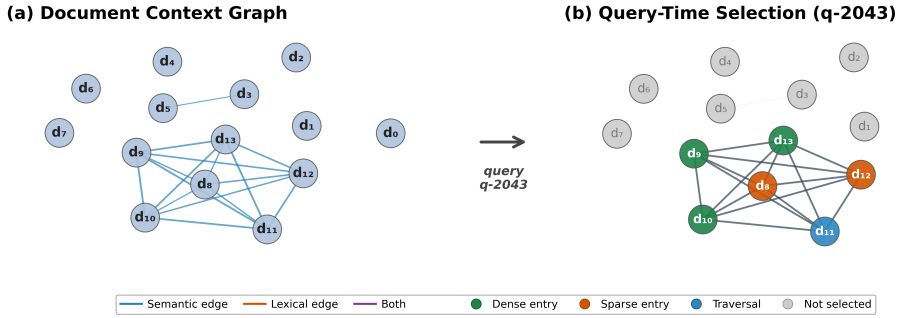


Figure 14: Document graph and query-time selection for Topic 7 (South Korean impeachment crisis). Node positions are identical across panels. **(a)** Document context graph: 14 documents connected by hybrid similarity edges (semantic-dominant, lexical-dominant, balanced). The densely connected component captures the core impeachment narrative; peripheral nodes represent tangentially related events. **(b)** Query-time selection for q-2043 (target: “Yoon’s aides offered to resign”): 6/14 documents (43%) selected via 3 dense entries, 2 sparse entries, and 1 traversal-discovered node. Grey nodes are filtered distractors.

Metric	Value
Avg. documents per topic	21.5
Avg. retrieved per question	16.6 (73%)
Avg. in topic-wide context	15.8 (73%)
Likely distractors excluded	27%
Prompt cache hit rate	91%
Overall inference cost reduction	87%

Table 12: Retrieval statistics on the dev set (400 questions, 36 topics, 775 documents).

Model	Base	+Retr.	Δ
Haiku 3.5 [†]	0.440	0.530	+0.090
Sonnet 4	0.781	0.785	+0.004
Gemini 3 Flash	0.804	0.806	+0.003
GPT-5.2	0.810	0.818	+0.008

[†]100 questions; others on full dev set (400).

Table 13: Impact of distractor filtering on accuracy. Smaller models benefit substantially from context reduction (Xu et al., 2024).

topic-wide cache hit rate (Section 3.1).

C Prompt Templates

In this section, we present the key prompt templates used in our system. All prompts leverage XML-structured formatting to enforce disciplined reasoning and constrained output (Section 3.2). Placeholder tokens in brackets are filled at inference time.

C.1 Direct Prompting

Our primary inference prompt employs a structured analysis-before-answer format that requires the model to evaluate each candidate option independently before committing to a final answer.

C.2 GEPA Optimization Insights

We employ GEPA (Agrawal et al., 2026) via DSPy (Khatab et al., 2024) to explore the prompt design space through reflective prompt optimization. Rather than deploying the resulting prompts directly, we analyze the heuristics GEPA automatically discovered to extract transferable design principles for causal reasoning that informed our final manually-designed prompt (Figure 15).

Optimization Setup. The optimization signature provides the target event, context documents, and candidate options as inputs, with reasoning and answer as outputs. We intentionally kept the baseline signature minimal (a single docstring with no engineered heuristics) to give GEPA maximum room for discovery. The metric rewards exact match (1.0) and partial match (0.5), with diagnostic feedback identifying failure types (under-selection, over-selection, wrong answer) to guide reflective evolution.

Practical Observations. Our experience with GEPA yields several practical observations for prompt optimization in general. GEPA is most effective as a *bootstrapping* tool: given no starting prompt, it can rapidly produce a functional one, making it well-suited for agent retraining when adapting to new data distributions or edge cases. However, it can only mildly optimize an already well-functioning prompt, and optimization quality depends heavily on the pre-optimization starting point. We also observe that prompt evolution plateaus at the knowledge boundary of the reflection LLM: the optimizer cannot discover reasoning strategies that the reflection model itself does not

Structured Zero-Shot Prompt

```
<role>
You are an expert in identifying the direct cause of events from textual evidence.
</role>

<task>
Given an event, context documents, and candidate explanations, analyze systematically to identify the
most plausible direct cause(s).
</task>

<input_format>
<context_documents>
<document_1>: [document content]</document_1>
<document_2>: [document content]</document_2>
[additional documents as needed...]
</context_documents>

<target_event>[event description]</target_event>

<options>
<option_a>[option A]</option_a>
<option_b>[option B]</option_b>
<option_c>[option C]</option_c>
<option_d>[option D]</option_d>
</options>
</input_format>

<instructions>

<reasoning_criteria>
- Base your reasoning ONLY on evidence from the provided context documents
- Look for direct causal relationships, not just correlations or temporal sequences
- Test logical sufficiency: Would this factor alone reasonably be enough to cause the event?
- Require both conditions: Direct textual support AND logical sufficiency to cause the event
- Use single-step reasoning: Avoid multi-step causal chains or indirect relationships
- Prioritize explicit causal language: "caused by," "resulted from," "led to," "triggered by," "due to"
</reasoning_criteria>

<selection_rules>
- Multiple options can be correct - choose ALL that apply
- Select multiple options only if each cause has strong evidence and is individually sufficient
- If options contradict each other, select the one with stronger textual evidence
- Always output ALL correct options, including duplicates: if options are duplicate/identical but
correct, include both letters
- If none seems perfectly sufficient, select the single best-supported among A-D.
- NEVER create options beyond A, B, C, D
- There is always at least one correct option from A-D
</selection_rules>

<quality_checks>
- Verify each selected option has direct quotes or paraphrases from context
- Ensure you haven't made assumptions beyond what's explicitly stated
- Confirm logical sufficiency: could this realistically cause the event by itself?
- Valid answers in the <answer> section are only A,B,C,D; never output anything else; if uncertain, pick
the best-supported among A-D and output it without explanations.
</quality_checks>

</instructions>

<output_format>

Provide your answer in EXACTLY this format (no additional text before or after):

<analysis>
Option A: [Your brief reasoning for option A - 1-2 sentences]
Option B: [Your brief reasoning for option B - 1-2 sentences]
Option C: [Your brief reasoning for option C - 1-2 sentences]
Option D: [Your brief reasoning for option D - 1-2 sentences]
</analysis>

<answer>
[Letter(s) ONLY - e.g., "B" or "B,D" or "C"]
</answer>

CRITICAL FORMATTING RULES:
- Start your response with <analysis> (no text before it)
- End your response with </answer> (no text after it)
- In <answer> tags, write ONLY letters: A, B, C, or D (comma-separated for multiple)
- DO NOT write "Option A" or "Option B" in the <answer> tags - just the letter(s)

</output_format>
```

Figure 15: Structured zero-shot prompt for direct inference (Section 3.2).

possess. Data contamination, a common concern with automatic prompt optimization, can be mitigated by adjusting the reflection prompt to discourage memorization of training examples. Finally, we note that incorporating an explicit optimization-step directive within the reflection prompt (e.g., instructing the reflector to propose targeted modifications rather than wholesale rewrites) improves stability and convergence across runs.

Discovered Heuristics. The best evolved prompt independently surfaced five causal reasoning heuristics without any manual engineering:

1. **Path length prioritization:** Strong causes are 1-2 steps from the target; distal causes are treated as background context unless they meet an exception criterion.
2. **Confidence thresholds:** High-confidence causal links are “almost always” strong; moderate links require supporting causal language; weak links are generally rejected.
3. **Directionality and sibling detection:** Effects of the target are distinguished from causes; shared-root events are classified as side effects rather than causes.
4. **Evidence-cause distinction:** Media reports about an event are separated from the event itself as causal triggers.
5. **Duplicate handling:** Identical options describing the same event must both be selected.

These heuristics are notable because they emerged purely from optimization pressure on answer correctness, without any explicit instruction about causal theory.

Dataset Bias Discovery. Beyond causal reasoning heuristics, GEPA independently surfaced two exploitable artifacts of the dataset construction process. First, 22.9% of questions contain *duplicate options* (identical text in different positions) and GEPA learned that selecting one without the other always incurs a scoring penalty, converging on the invariant that identical options must share truth values. Second, 14.6% of questions include a “None of the others are correct causes” option, and GEPA discovered that co-selecting it with other candidates always produces incorrect answers, effectively learning mutual exclusivity without being told. These patterns are significant because they emerged purely from optimization pressure: GEPA identified them as deterministic “free score” invariants before any manual dataset analysis. These

GEPA-surfaced invariants directly motivated the post-hoc consistency heuristics described in Section A.9, specifically R1 (None Exclusivity) and R2 (Duplicate Consistency), which together account for 19 corrections and seed the cross-question propagation chain that yields the full +5.6 pp improvement.

Error-Informed Refinement. After analyzing the hardest questions (Section D.1), we created a second-generation prompt that directly addresses the identified failure patterns:

- **Proximal-cause bias:** Replaced the strict proximity rule with a counterfactual necessity test (“Could the target have occurred without this event?”) allowing enabling conditions at any distance to qualify as valid causes.
- **Single-cause default** (under-selection): Added explicit multi-answer guidance requiring re-evaluation of each remaining option after identifying the strongest cause.
- **Causal chain length disagreement:** Distinguished “proximate triggers” from “enabling conditions,” both treated as valid cause types.
- **Over-selection penalty:** Added a precision warning noting that incorrect extra selections result in zero score.

From GEPA to Final Prompt. We deliberately curated which GEPA discoveries to incorporate into the final prompt (Figure 15). Structural heuristics that generalize across domains were retained: single-step reasoning preference (reasoning criterion 5), explicit causal language prioritization (criterion 6), duplicate option handling (selection rule 4) and the analysis-before-answer format that GEPA consistently preserved across all evolved prompts. Conversely, domain-specific patterns that GEPA discovered, such as “Request → Approval” and “Crisis → Response” causal templates, were deliberately excluded to avoid overfitting to the optimization set’s topical distribution. Numeric confidence thresholds were abstracted into the general logical sufficiency test (criterion 3). The dataset-specific invariants (duplicate consistency, “None” exclusivity) were implemented as post-hoc heuristics rather than prompt instructions, since they are deterministic properties better enforced programmatically than via natural-language instruction.

D Error Analysis Details

This appendix provides extended analysis supporting the findings in Section 5.1. We analyze inter-annotator agreement across 14 models from 7 families, all evaluated on the full dev set (400 questions). Table 16 summarizes global agreement metrics and Figure 17 shows family-level mean pairwise Cohen’s κ , revealing that frontier families (Claude, Gemini, GPT) form a high-agreement cluster ($\kappa > 0.78$) while mid-range models (Llama, Qwen) agree less with frontier systems ($\kappa \approx 0.55$).

D.1 Failure Taxonomy

Table 14 lists the 3 truly unbreakable questions (all 14 models score zero) and 6 near-unbreakable questions (≥ 12 models score zero). Four distinct failure patterns emerge:

Proximal-Cause Bias. Question q-2182 asks for the cause of “Taliban fighters spread across Kabul streets.” The gold answer is C (“US troops began withdrawing from Afghanistan in early July”), but 12/14 models select D (“Afghan President Ashraf Ghani fled the country”). Models apply a temporal proximity heuristic (Joshi et al., 2024), preferring the most recent antecedent, over enabling conditions.

Under-Selection + Wrong Option. Question q-2341 (gold: B,C,D) sees 11/14 models selecting only A, the sole *incorrect* option, indicating that models not only under-select but converge on the wrong cause entirely when the gold answer is a large multi-answer set.

Annotation Ambiguity. Question q-2149 (Amazon fires, gold: C, 12/14 fail) involves cases where the context supports multiple reasonable interpretations; model disagreement may reflect genuine annotation uncertainty rather than reasoning failure.

D.2 Topic-Level Difficulty

Table 15 reports the hardest topics by exact-match accuracy, alongside structural features that predict difficulty. Figure 16 shows per-topic Fleiss’ κ across all 14 models, revealing substantial cross-topic variance (range: 0.46-0.90).

D.3 Under-Selection Mechanism Analysis

Across the 42 no-exact-match questions (14 models), models select on average 1.2 options when gold requires 2.4, a 51% cause-count reduction. In

83% of cases, the consensus prediction is a single option. Three mechanisms drive under-selection (Figure 19):

Causal chain incompleteness (18/42 questions).

When the gold answer spans multiple links in a causal chain (root cause \rightarrow intermediate events \rightarrow target), models select only one link and omit the rest (Chi et al., 2024; Yu et al., 2025). For example, q-2135 (Korean martial law) has gold={A,C,D} covering the full chain (declaration \rightarrow soldiers stormed \rightarrow annulment \rightarrow impeachment), yet all 14 models select only A. The pattern recurs across the Capitol riot questions (q-2044, q-2094, q-2252, q-2299), where models consistently miss intermediate steps in the protest \rightarrow breach \rightarrow lockdown \rightarrow halt sequence.

Proximate cause preference (11/42 questions).

Models favour the most causally proximate antecedent over distal enabling conditions (Joshi et al., 2024), whether the distance is temporal (months between events) or structural (enabling process vs. its outcome). This is most severe in the wrong-answer failures: in q-2261 (Afghanistan), all 14 models choose “Ghani fled” (proximal trigger) over “US withdrawal” (enabling condition), scoring zero. The same pattern recurs across five Brexit questions (q-2090, q-2210, q-2233, q-2286, q-2322), where models pick the vote outcome (“voted to leave”) but miss the enabling institutional act (“held a referendum”), events on the same day, distinguished by causal role rather than temporal order.

Salience bias (9/42 questions).

Models prefer dramatic, newsworthy causes over subtler contributing factors (Koo et al., 2024). In q-2226 (Beirut explosion), all 14 models select “explosion killing 200 people” but miss “destroyed medical infrastructure”: both are valid causes of the PM’s resignation, but models anchor on the more dramatic event.

Multi-answer prevalence correlates positively with topic difficulty (Spearman $\rho > 0$) across all 36 topics. Topic 33 (Amazon fires, $\kappa = 0.457$) reflects cyclical causation where deforestation causes fires, fires trigger policy responses, and policy failures enable further deforestation. Topic 30 (George Floyd/BLM, $\kappa = 0.458$, 72% multi-answer) shows how social movement causation creates branching chains where many options are partially valid.

QID	Topic	Gold	Consensus	Fail	Multi?	Pattern
<i>Truly unbreakable: all 14 models score zero (all 7 families fail)</i>						
q-2261	Afghanistan	A,C	D	14/14	Yes	Under-sel. + proximal
q-2356	DeepSeek	D	C	14/14	No	Consensus wrong (7/14 C)
q-2411	DOGE/Musk	A	D	14/14	No	Consensus wrong (13/14 D)
<i>Near-unbreakable: ≥ 12 models score zero</i>						
q-2341	Afghanistan	B,C,D	A	13/14	Yes	Under-sel. + wrong opt.
q-2028	Credit Suisse	A	D	13/14	No	Distal-cause preference
q-2182	Afghanistan	C	D	12/14	No	Proximal-cause bias
q-2225	Wagner	C	A	12/14	No	Proximal-cause bias
q-2149	Amazon fires	C	B	12/14	No	Annot. ambiguity
q-2036	Myanmar	B	A	12/14	No	Distal-cause preference

Table 14: Hardest questions across 14 models (7 families). Top: 3 questions where all models score zero. Bottom: 6 near-unbreakable questions where ≥ 12 models fail. *Fail* = models scoring zero. *Under-sel.* = multi-answer under-selection; *Consensus wrong* = majority agrees on incorrect option; *Proximal-cause bias* = models prefer temporally nearest cause (Joshi et al., 2024).




Figure 16: Per-topic Fleiss' κ across 14 models (bar labels = question counts per topic); dashed line marks the per-topic mean ($\kappa = 0.697$). Topics are sorted by agreement, coloured by Landis–Koch interpretation: green = almost perfect ($\kappa \geq 0.81$), yellow = substantial (0.61–0.80), orange = moderate (0.41–0.60).




Figure 17: Family-level mean pairwise Cohen’s κ across 14 models (7 families). Frontier families (Claude, Gemini, GPT, Kimi) show high mutual agreement ($\kappa > 0.77$); DeepSeek is dragged down by R1’s low accuracy. Numbers in parentheses = model count per family. Diagonal masked (within-family comparison undefined for single-model families).

ID	Domain	N	Multi%	κ_{14}
26	California Fire	3	—	0.456
33	Amazon fires	14	43	0.457
30	George Floyd	18	72	0.458
18	DeepSeek	14	—	0.492
32	Afghanistan	23	57	0.545
12	Wagner	11	—	0.547

Table 15: Hardest topics by Fleiss’ κ across 14 models. Topics with complex causal structures and high multi-answer prevalence show the lowest agreement.

D.4 Model Agreement and Ensemble Analysis

Inter-Annotator Agreement. Table 16 reports agreement metrics across 14 models. Fleiss’ $\kappa = 0.690$ (*substantial*) reflects the inclusion of weaker models (DeepSeek R1, Llama, Qwen) that increase prediction variance. Within-family agreement ($\bar{\kappa} = 0.794$) exceeds cross-family agreement ($\bar{\kappa} = 0.661$), confirming that architecturally similar models share prediction patterns. Agreement is higher on single-answer questions ($\kappa_{14} = 0.699$) than multi-answer questions ($\kappa_{14} = 0.668$), consistent with the inherently greater difficulty of multi-

Metric	Value
Fleiss’ κ (Fleiss, 1971)	0.690
Kripp. α (Jaccard) (Krippendorff, 2011)	0.645
Kripp. α (nominal)	0.613
Unanimous agreement	10.0%
Majority agreement	100%
Within-family $\bar{\kappa}$	0.794
Cross-family $\bar{\kappa}$	0.661
Unbreakable (all score 0)	3
No exact match	42
Oracle (best per Q)	0.895

Table 16: Inter-annotator agreement metrics across 14 models (7 families: Claude, DeepSeek, Gemini, GPT, Kimi, Llama, Qwen) on 400 dev questions.

answer evaluation.

Ensemble Analysis. The oracle upper bound of 0.895 (selecting the best model per question) indicates 6.7 pp headroom from model complementarity, with unique correct predictions distributed across families: Gemini Flash 7, Claude (family




Figure 18: Dev-set difficulty distribution across 14 models: Easy 38, Medium 225, Hard 95, Very Hard 42 (based on how many models score zero). The majority of questions are medium-difficulty, with only 42 very hard questions where ≥ 10 models fail.

Configuration	k	Base	+SC	Δ
<i>Same-Model SC ($\tau=1.0$)</i>				
Claude Sonnet 4.5	3	0.790	0.806	+1.6
Gemini 3 Flash	3	0.804	0.810	+0.6
<i>Multi-Model Ensemble</i>				
Son. Think + GPT + Gemini	3	—	0.811	—

Table 17: Self-consistency and multi-model ensemble on the dev set (400 questions). k = number of predictions aggregated via majority voting ($\theta=0.5$). Same-model SC provides negligible gains; multi-model combination provides modest gains through model complementarity.

Strategy	θ	Dev	Test
<i>Sonnet 4.5 ($k=3, \tau=1.0$)</i>			
Baseline ($k=1$)	—	0.790	0.904 [†]
Majority	0.50	0.806	0.902
Union	0.33	0.788	0.869
Strict Majority	0.67	0.784	0.897
Intersection	1.00	0.784	0.897
<i>Gemini 3 Flash ($k=3$ dev, $k=5$ test, $\tau=1.0$)</i>			
Baseline ($k=1$)	—	0.804	0.907
Majority	0.50	0.810	0.902
Union	0.33	0.788	0.888
Strict Majority	0.67	0.806	0.903
Intersection	1.00	0.806	0.900

[†]Test uses Sonnet 4.5 Thinking ($k=3$).

Table 18: Self-consistency aggregation strategies on dev (400 questions) and test (612 questions). θ = per-option inclusion threshold. **Bold** = best SC strategy per section.

total) 9, Llama 3, GPT-5.2 2. Bootstrapped significance tests yield $p > 0.05$ for all pairwise model comparisons, indicating that observed differences are not statistically significant at $\alpha = 0.05$ given the 400-question evaluation set, a limitation inher-

ent to shared-task evaluations.

D.5 Self-Consistency and Ensemble Analysis

Table 17 reports the impact of self-consistency sampling and multi-model ensembling on the dev set. Same-model self-consistency ($k=3, \tau=1.0$) provides modest accuracy gains: Sonnet 4.5 improves by 1.6 pp and Gemini 3 Flash improves by 0.6 pp, suggesting that sampling variation at temperature 1.0 benefits causal reasoning moderately. In contrast, majority voting across three architecturally diverse systems (Claude Sonnet 4.5 Thinking, GPT-5.2, Gemini 3 Flash) yields 0.811, exceeding any same-model SC configuration.

Table 18 compares aggregation strategies across both models and both splits. Four per-option inclusion thresholds are evaluated: union ($\theta=0.33$, include if any sample selects), majority ($\theta=0.50$, strict majority ($\theta=0.67$), and intersection ($\theta=1.0$, all samples must agree). With $k=3$, strict majority and intersection yield identical results because no option can exceed $\frac{2}{3}$ without unanimous agreement; the difference emerges with $k=5$ (Gemini test).

Majority voting ($\theta=0.5$) is the most robust strategy, matching or marginally improving baseline across all configurations. Union ($\theta=0.33$) consistently reduces accuracy by 1-4 pp across all configurations due to over-selection, particularly on the test set where single-answer prevalence is higher. Strict thresholds ($\theta \geq 0.67$) under-select on the dev set, where 47.5% of questions require multi-label answers; with $k=3$ these thresholds reduce to unanimous agreement and yield identical results.

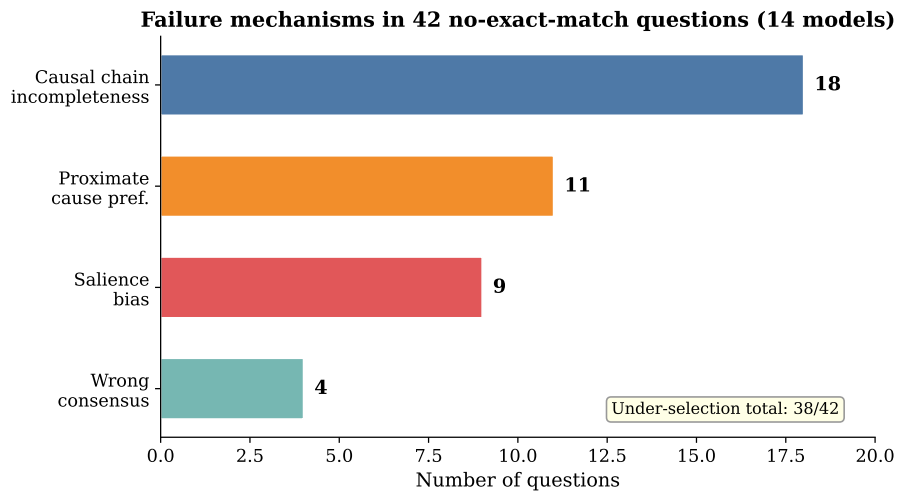


Figure 19: Failure mechanism breakdown across 42 no-exact-match questions (14 models). Causal chain incompleteness, proximate cause preference, and salience bias together account for 38/42 failures, all manifestations of the single-cause default bias.

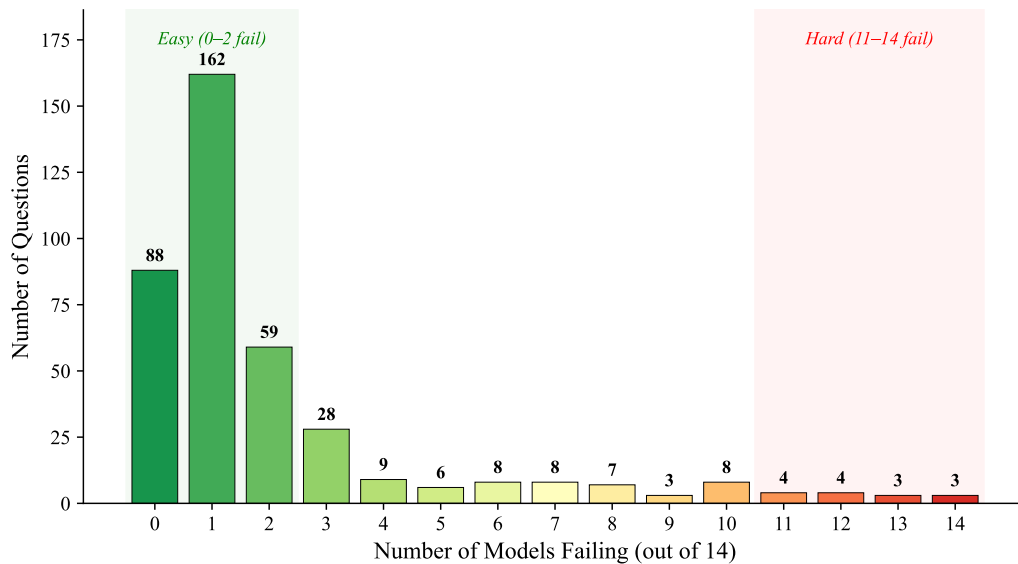


Figure 20: Failure gradient across 14 models: distribution of how many models fail each question. 88 questions (22%) are solved by all 14 models; only 3 questions (0.75%) defeat all 14. The long right tail (10-14 failures) contains 22 questions requiring systematic improvements.




Figure 21: Per-model exact-match accuracy by answer cardinality across 14 models. Every model shows a single-to-multi gap (25-56 pp); under-selections (1,389 total) vastly exceed over-selections (52), confirming conservative cause selection as the dominant shared bias.