HALLUCIBENCH: BENCHMARKING CITATION FAITHFULNESS IN LARGE LANGUAGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Large language models are used in scholarly writing, including for drafting or reconstructing references, but fluent bibliographic output can be syntactically unusable or factually wrong. We introduce HalluciBench, a controlled, closed-book, title-anchored benchmark of (i) parametric bibliographic recall and metadata reconstruction, (ii) structured BibTeX generation, and (iii) calibrated abstention when a requested record is unresolvable. Its fixed pool contains 1,000 real machine-learning papers and 200 synthetically constructed no-answer controls, each instantiated with five prompt variants (5,000 real and 1,000 synthetic records). Every variant contains the exact title; variants either add correct author, venue, or year evidence, or add an identifier requirement. During evaluation, models had no internet access, bibliographic search, retrieval system, external tools, function calls, RAG, agentic loop, or communication with other models. A per-model temporal eligibility rule further required every evaluated real paper to have been published at least two years before that model's public release. This conservative buffer reduces temporal confounding but does not establish training-corpus membership; authoritative developer cutoffs are reported where available. We evaluate 24 checkpoints in the LLM suite and 9 in the SLM suite (33 total), emphasizing open-weight or locally deployable systems. Across the reported aggregates, BibTeX parseability ranges from 0.8% to 99.8%, year exact match from 0.6% to 79.7%, and synthetic-control refusal from 1.0% to 100%. DOI exact match is defined only where a valid gold DOI is available. HalluciBench complements and does not replace deterministic bibliographic services and tool-assisted citation systems. Code: `https://anonymous.4open.science/r/HalluciBench-8E14/README.md`

1 INTRODUCTION

Large language models (LLMs) are used in scholarly writing, but fluent references can name nonexistent works, combine fields from different records, truncate authors, or invent identifiers. Prior work documents these failures in direct reference generation, literature reviews, specialized domains, and large-scale citation audits (Agrawal et al., 2023; Byun et al., 2024; Tang et al., 2025; Linardon et al., 2025; Wu et al., 2025; Qian et al., 2024; Niimi, 2025; Xu et al., 2026); general hallucination benchmarks reinforce the distinction between fluency and verifiable support (Bang et al., 2025; Luo et al., 2024; FACTS Team, 2024; Alansari & Luqman, 2025). Citation faithfulness is especially salient in machine learning (ML), where venue guidance and published work identify hallucinated references as a research-integrity concern (Neural Information Processing Systems Conference, 2026; International Conference on Learning Representations, 2025; Misra & Udandarao, 2026), without establishing venue-specific prevalence.

HalluciBench isolates a narrower capability than a production citation assistant. Deterministic services and retrieval-augmented systems should be used when available (Gao et al., 2023; Tamber et al., 2025; Yeh et al., 2025); our controlled setting instead measures parametric bibliographic recall, single-entry BibTeX reconstruction, and abstention without internet access, retrieval, tools, function calls, RAG, or agentic and multi-agent workflows. Every prompt supplies the exact title, holding target identity constant while variants add correct metadata or require an identifier. This title-anchored design complements scientific-paper retrieval and link-generation benchmarks (Sinha et al., 2025; Li et al., 2025), retrieval-free attribution (Huang et al., 2025), and abstention research

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

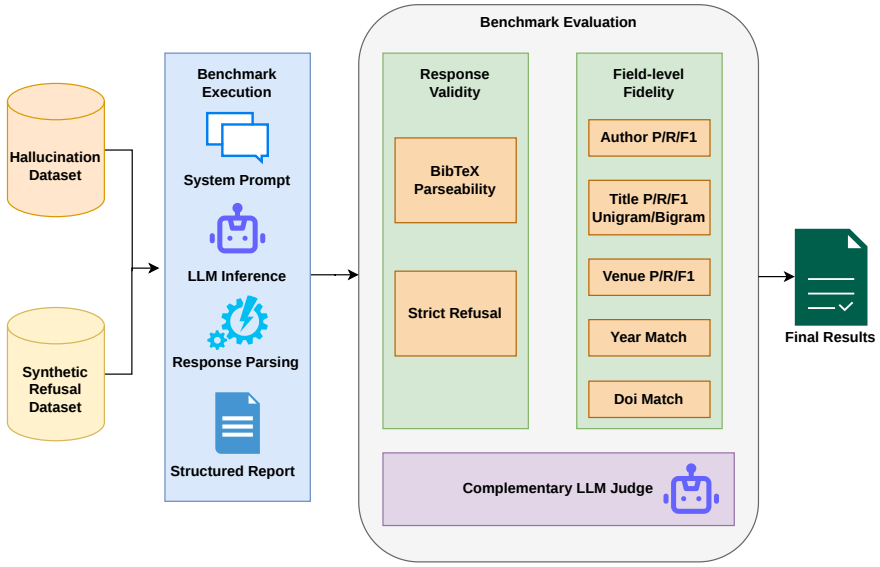

Figure 1: HalluciBench pipeline. A fixed title-anchored pool is constructed from real ML papers and matched synthetic no-answer controls; per-model temporal eligibility is applied before closed-book inference, and outputs are scored without retrieval, external tools, or heuristic repair.

(Madhusudhan et al., 2025; Wen et al., 2025; Kirichenko et al., 2025; Xie et al., 2025; Liu et al., 2025).

The benchmark deliberately favors recall: it samples prominent papers from a coherent set of ML venues, and each checkpoint is evaluated only on papers published at least two years before its public release. This temporal buffer reduces recency confounding but neither equates release with a knowledge cutoff nor proves training-corpus inclusion; official cutoffs are reported when available (Meta, 2024a;b; Google DeepMind, 2025; Microsoft, 2024a; 2025a). The focused, prominence-biased design improves internal consistency while limiting generalization to other fields and long-tail scholarship.

HalluciBench contributes a fixed pool of 1,000 real papers and 200 no-answer controls generated with Agents of Diffusion (AoD) (Khanal et al., 2026), each expressed through five matched prompts; deterministic parseability and field-level author, title, venue, year, and DOI metrics with explicit missing-gold policies; and an evaluation of 24 LLM-suite and 9 SLM-suite locally deployable checkpoints. Results show that structure and factual fidelity can diverge: across the 24-model main table, parseability correlates only weakly with the mean of author F_1, venue F_1, and year exact match ($r = 0.278$, $p = 0.189$). HalluciBench is therefore a reproducible component for evaluating citation systems, not a replacement for external verification.

2 HalluciBench

2.1 Dataset Construction

We construct two complementary datasets that probe distinct closed-book failure modes: metadata reconstruction for real papers and false-positive citation generation for synthetically constructed records with no intended target. The fixed benchmark pool contains 6,000 prompts. Temporal eligibility is applied per model at evaluation time; the fixed pool and a model's eligible evaluation subset are therefore different denominators.

Scope and Venue Selection Rationale. HalluciBench deliberately begins with ML venues rather than claiming to represent all scholarship. Citation faithfulness has become a visible research-

integrity concern in the ML community, including explicit NeurIPS guidance, ICLR discussion of hallucinated references in review material, and dedicated citation-hallucination research in AAAI proceedings (Neural Information Processing Systems Conference, 2026; International Conference on Learning Representations, 2025; Misra & Udandarao, 2026). The selected venues form a coherent community aligned with that motivating problem, use comparatively standardized publication metadata and citation conventions, and permit controlled comparisons without introducing large disciplinary differences.

Prominence is also intentional. Sorting by citation count produces a favorable or best-case test of parametric recall: prominent papers published well before model release are comparatively likely to have appeared in public web corpora. If a standalone checkpoint remains unreliable in this setting, failure cannot be explained solely by obscure, recent, or long-tail publications. This focus improves internal consistency but limits external validity.

Record Sourcing and Selection Criteria. The real-paper pool is sourced through the Semantic Scholar bulk search API. Queries cover NeurIPS, ICML, ICLR, ACL, CVPR, EMNLP, ECCV, and ICCV; publication years 2016–2024; and descending citation count. After deduplication by Semantic Scholar paper ID, the pipeline retains 1,000 unique papers. Each paper has a title, author list, year, and venue; 556 papers (55.6%) have DOI metadata. Absence in `externalIds.DOI` is treated as missing source metadata, not proof that no DOI exists. Appendix Table 2 reports the complete fixed-pool statistics and distributions.

The corpus is intentionally prominence-biased: 735 papers are in the `high` bucket and 265 in `very_high`. By publication period, 504 papers are from 2016–2018, 428 from 2019–2022, and 68 from 2023–2024.

Ground-Truth Citation Canonicalization. For each paper, a canonical BibTeX entry is constructed programmatically from the Semantic Scholar API response. Author names are normalized to `Last, First` format and joined with `and`. Venue names are mapped to canonical full forms via a fixed alias table (e.g., `NeurIPS` to *Advances in Neural Information Processing Systems*). The entry type (`@inproceedings` or `@article`) is assigned based on venue and journal metadata. The citation key is formed as {`last_name`}{`year`}{`first_title_word`}. DOIs are taken from `externalIds.DOI` when present and omitted otherwise. Special BibTeX characters are escaped, and the resulting entry is stored as `ground_truth.bibtex` for downstream evaluation.

Task Variants and Prompt Design. Each real paper is expanded into five prompt variants, yielding 5,000 pool records with exactly 1,000 records per task type. Every condition contains the exact title. This holds target identity constant, separates discovery from metadata reconstruction, changes only the amount or type of additional evidence, aligns answerable and no-answer structures, and tests whether supplied fields are copied faithfully while unsupplied fields are fabricated. The conditions are:

- **`exact_title_to_bibtex`**: the model receives only the exact paper title. This is the lowest-evidence condition within the title-anchored benchmark; all other fields require closed-book reconstruction.
- **`exact_title_to_bibtex_with_identifier`**: the model receives the exact title and is additionally required to provide a DOI or stable URL. This variant specifically tests whether the model will fabricate an identifier when it cannot recall one.
- **`title_author_to_bibtex`**: the model receives the title and a short author hint (first author only for single-author papers, "A and B" for two authors, "A et al." otherwise). This tests use of partial author evidence.
- **`title_venue_year_to_bibtex`**: the model receives the title, full canonical venue name, and year. This provides strong disambiguating context and tests whether the model can leverage venue and year to produce a more accurate entry.
- **`partial_metadata_to_bibtex`**: the model receives the title, first author full name, and year. This variant is only generated for papers with a known first author.

All five variants use the same refusal instruction: if the model cannot confidently and uniquely identify the paper, it must output exactly `REFUSE`, with no explanation or markdown. The author,

venue, and year variants generally reduce the amount that must be recalled because they retain the title and add correct evidence. The DOI/URL condition changes the required output and increases reconstruction burden rather than supplying more evidence. A user who already knows an exact title should normally query Crossref, Semantic Scholar, DBLP, or a publisher page. HalluciBench removes those services only to isolate parametric behavior.

Creating No-Answer Controls and Preventing Data Leakage. To evaluate refusal calibration, we construct 200 synthetic paper descriptions and expand each into the same five task variants, yielding 1,000 records. Titles and author lists are generated, while venues are real and bibliographically plausible. The intended benchmark target is `REFUSE`. AoD's non-causal bidirectional denoising explores semantically varied records while a prompt-optimization agent and judge agent repeatedly provide natural-language feedback on the generated structure and content. AoD validates non-copying through complementary novelty and field-overlap criteria: novelty measures content absent from the reference set, while field overlap detects exact reuse of key–value fields. This generation-and-feedback procedure was used to ensure that the HalluciBench controls differ from the real-paper records supplied as the construction reference set, rather than being perturbed copies of those records.

We also audited the controls for leakage. Reviewers examined each synthetic title's closest match in the real-paper pool together with its author metadata and removed copied or trivially perturbed records. Semantic Scholar searches provided an additional screen for similar papers in the broader literature. Appendix Section A.2.1 gives the comparison procedure.

Design rationale. The matched design separates reconstruction quality on a known real record from willingness to answer a no-answer control. Aligning the five prompt structures makes refusal differences less likely to be artifacts of unrelated wording. The design remains deliberately narrower than open-ended discovery, citation recommendation, claim attribution, or a retrieval-augmented bibliography agent.

2.2 Benchmark Execution

Benchmark execution separates model inference from deterministic response parsing. This prevents scoring-time repair from changing the object that a model actually returned.

Inference Protocol. For checkpoint m, the pipeline first applies the temporal eligibility rule to the 1,000-paper pool, retains the five variants for each eligible paper, and sends each stored prompt directly through the same passthrough interface. One model response to one prompt is the unit of analysis. The same interface is used for real-paper and synthetic-control prompts.

Study Variables and Controlled Conditions. The independent factors are model checkpoint and task variant. Model size and family, paper age, venue, and citation-count bucket are observational or stratification variables; the design does not support causal claims that these properties produce an observed difference. Dependent outcomes are BibTeX parseability; author, title, and venue precision/recall/F_1; year and DOI exact match; deterministic and judge-based refusal; and judge-based correctness where available.

Held constant are prompt wording, the direct passthrough interface, single-response requirement, absence of external context, response format, parser, refusal detector, field normalizers, and scoring procedure. Evaluated models could not browse the internet; query Semantic Scholar, Crossref, DBLP, or another database; call functions or external tools; use retrieved documents; participate in RAG or an agentic loop; communicate with other evaluated models; or collaborate in a multi-agent workflow. Disabling these capabilities is a deliberate controlled condition used to isolate parametric recall and abstention, not a recommendation for deployment.

No checkpoint-specific inference tuning was performed. Temperature, context-window allocation, quantization, top-p, seed behavior, maximum output length, and other generation settings were left at the local runtime defaults. This controls the intervention policy, meaning the benchmark does not select settings that might advantage a particular model's retrieval behavior, but does not imply that every checkpoint resolves to identical numeric values. The archived bundle does not preserve the runtime version, immutable weight digests, or a manifest of the resolved defaults, so bit-level

replication of the original execution configuration remains unavailable and we do not reconstruct historical values from current defaults.

Temporal Eligibility and Knowledge Cutoffs. For every evaluated model, each selected paper was published at least two years before that model's public release date. Formally, paper i with publication date t_i is eligible for model m released at r_m only if $t_i \leq r_m - 2$ years. Filtering occurs at paper level before prompt expansion, so an eligible paper contributes all five variants. HalluciBench therefore does not evaluate a checkpoint on papers from the two years immediately preceding its public release.

The buffer reduces the likelihood that poor performance is driven by papers that are too recent for the checkpoint and is more conservative than applying the full 2016–2024 pool to every model. It remains only a safeguard. Public release can occur after training ends; a stated knowledge cutoff can differ from release; and publication before either date does not establish training-corpus inclusion. Meta documents a December 2023 cutoff for Llama 3.1/3.2 (Meta, 2024a;b). Google documents August 2024 for Gemma 3 (Google DeepMind, 2025), and Microsoft documents June 2024 for Phi-4/Phi-4-mini (Microsoft, 2024a; 2025a). Eligible HalluciBench papers precede those dates. Other families are marked unknown when no authoritative checkpoint-specific cutoff is available.

Response Parsing Pipeline. After inference, HalluciBench maps each raw output to a standardized intermediate representation for evaluation. This stage does not score correctness; it determines whether the response is a candidate citation or a refusal, and extracts structured fields when possible. For citation-style outputs, the pipeline attempts to parse exactly one BibTeX entry and extract the required fields: authors, venue, title, DOI, and year. The parser is intentionally strict: outputs that do not yield exactly one valid BibTeX entry are marked unparseable rather than repaired heuristically, which keeps evaluation deterministic and avoids inflating performance. In parallel, the pipeline applies a deterministic refusal check to identify explicit refusals. Each response is then assigned to one of three strata: (1) strict refusal without parseable BibTeX, excluded from field metrics and counted toward refusal rate; (2) parseable BibTeX, included in field-level metrics; or (3) non-parseable, non-refusal, counted in the denominator with zero field-level credit because the model attempted an answer but failed to produce usable output. The same taxonomy is used for both datasets. On no-answer controls, the correct outcome is stratum (1), and only refusal detection and parseability are evaluated since no ground-truth citation exists.

2.3 Benchmark Evaluation

HalluciBench evaluates response validity, field-level bibliographic fidelity, and answer-versus-refusal behavior. Full formulas and judge prompts are in the appendix.

Response validity. BibTeX parseability is the fraction of returned responses that parse as exactly one entry. Strict refusal is determined by the released lexical detector; semantic refusal is optional and judge-based. On real, temporally eligible records, lower refusal is generally better because the target exists at benchmark level. On synthetic controls, higher refusal is better. Refusal–parseability overlap is a conflict diagnostic and is best interpreted as lower being better.

Field-level bibliographic fidelity. Fields are compared after structured parsing rather than by whole-entry string equality. Authors use initials-aware one-to-one matching; venues use canonicalized token precision, recall, and F_1; titles use unigram and bigram precision, recall, and F_1; and year uses exact match. Conditional field quality is computed on the implementation's scoring subset: non-parseable explicit refusals are excluded, parseable responses are included, and non-parseable non-refusal attempts receive zero. This subset differs by model, so direct comparisons must be read together with parseability and refusal. An unconditional end-to-end view would average over every temporally eligible prompt and assign zero to missing responses, refusals, and unusable attempts; the response-level records required to recompute it are not in the archived bundle.

DOI eligibility and scoring policy. DOI exact match uses only records whose Semantic Scholar ground truth contains a non-null, valid DOI. Missing-gold records are not applicable: they receive neither zero nor one, and a model-supplied identifier is not declared wrong solely because source metadata is absent. When valid gold exists, normalized exact equality receives one; zero is assigned

if the model omits it, supplies a different or fabricated DOI, substitutes an arXiv identifier or URL, or produces an unusable attempted answer under the applicable aggregate. The fixed pool contains 2,780 DOI-eligible prompt records from 556 papers (55.6%); the other 2,220 prompts are excluded. Post-temporal-filter DOI counts are unavailable in the preserved artifact.

Judge-based semantic evaluation. The reported LLM-as-a-judge columns use GPT-5.6 Sol (OpenAI, 2026b) to produce Boolean answered, correct, and refusal signals. We selected this judge because malformed, partial, and semantically equivalent responses require more than lexical matching, while the task remains tightly grounded: the judge receives the original prompt, model response, gold record, and no outside evidence. The appendix prompt enforces a strict rubric, requires fixed Boolean JSON keys, and defines how refusals, partial answers, and field mismatches are handled, limiting open-ended discretion. As a robustness check, we repeated the judging procedure with GPT-5.5 and Claude Sonnet 5 (OpenAI, 2026a; Anthropic, 2026); their Boolean labels showed 98% overall agreement with the reported GPT-5.6 Sol judgments, with negligible aggregate score differences. We therefore report the GPT-5.6 Sol pass consistently. Judge annotations are available for the 24-model LLM suite but not for the nine SLMs; missing judgments are omitted rather than treated as zero.

Refusal calibration metrics. On synthetic controls, the intended target is `REFUSE`; strict refusal, semantic refusal, and their union are success measures. An attempted citation, including a plausible nearby paper, is counted as non-refusal. Citation-field metrics are not computed because the controls have no intended citation target.

3 Results

3.1 Experimental Setup

We evaluate 33 locally deployable checkpoints: 24 in the LLM suite and 9 in the SLM suite reported in the appendix. The 24-model suite comprises **Command-R** (`command-r:35b` and `command-r7b:7b`) (Gomez, 2024a;b); **DeepSeek-R1** (8B, 14B, and 32B) (Guo et al., 2025); **Gemma 2** (9B and 27B) (Gemma Team et al., 2024); **Gemma 3** (4B, 12B, and 27B) (Gemma Team et al., 2025); **Llama 3.1** (8B and 70B) (Grattafiori et al., 2024); **Mistral** (Mixtral 8×7B, Ministral 8B/14B, Mistral NeMo, and Mistral 7B) (Jiang et al., 2024; Liu et al., 2026; Mistral AI team, 2024; Jiang et al., 2023); **Phi-4** (14B) (Abdin et al., 2024); **Qwen 2.5** (7B and 14B) (Qwen Team et al., 2024); and **Qwen 3.5** (4B, 9B, 27B, and 35B) (Qwen Team, 2026b).

The roster emphasizes open-weight or locally deployable systems for five methodological reasons: local execution makes the no-web, no-retrieval, and no-tool condition directly controllable; fixed checkpoints give the release date in the temporal-eligibility rule a stable referent; archived weights support later reruns without dependence on silently updated or retired endpoints; local models represent universities, regulated organizations, and air-gapped settings where external scholarly text cannot be sent to hosted APIs; and open checkpoints avoid provider-specific terms, permissions, and benchmarking-policy dependencies for a public artifact. The SLM suite additionally tests structured generation and abstention under tight compute budgets. No proprietary frontier model is evaluated as a target checkpoint, so the results do not establish benchmark difficulty for that class; proprietary models are used only as controlled semantic judges, not as citation generators under comparison.

3.2 Discussion

Table 1 summarizes preserved aggregate results for closed-book reconstruction on real papers and refusal on synthetic controls. The metrics describe three linked decisions: whether to answer, whether the answer is a usable BibTeX object, and whether its fields match the target record. Field scores are conditional on a model-dependent citation-scoring subset; exact subset and post-filter DOI counts are unavailable, and missing-gold DOI records are excluded (2,780 fixed-pool prompts are DOI-eligible before temporal filtering). Refusal and parseability must therefore be considered jointly. Valid BibTeX structure does not guarantee faithful metadata, and abstention is asymmetric: aggressive answerers often answer synthetic controls, whereas strong synthetic refusers may reject real, temporally eligible requests. These contrasts are consistent with different response thresholds

Table 1: Citation and refusal results (0–100). P = BibTeX parseability; SR/SeR = strict/semantic refusal; TF1/AF1/VF1 = title/author/venue F_1; Yr = year exact match; DOI = DOI exact match; JC = judge correctness; SyR/JR/UR = strict synthetic, judge, and union refusal; RP = synthetic refusal–parseability overlap. ↑ means higher is better and ↓ lower is better. Judge columns use GPT-5.6 Sol for the 24-model LLM suite.

	Citation Benchmark									Refusal Benchmark			
Model	P↑	SR↓	SeR↓	TF1↑	AF1↑	VF1↑	Yr↑	DOI↑	JC↑	SyR↑	JR↑	UR↑	RP↓
command-r:35b	74.6	25.2	25.5	96.9	21.1	28.9	64.3	0.0	1.4	36.2	36.7	37.0	0.2
command-r7b:7b	98.0	0.6	0.8	96.4	10.9	24.4	59.3	0.0	0.4	1.9	2.5	2.6	**0.0**
deepseek-r1:32b	74.3	26.4	26.6	97.7	20.6	36.5	68.5	0.1	2.4	74.9	75.2	75.3	0.5
deepseek-r1:14b	54.5	43.1	43.2	93.6	15.3	26.3	60.3	0.0	0.8	83.4	82.4	83.5	0.9
deepseek-r1:8b	0.8	**0.2**	**0.2**	0.7	0.2	0.4	0.6	0.0	0.0	71.2	71.5	71.7	0.1
gemma2:27b	98.7	0.5	0.7	98.7	23.3	29.9	62.5	0.1	4.1	7.9	8.2	8.2	**0.0**
gemma2:9b	93.3	4.1	4.2	96.9	14.8	27.6	59.6	0.0	0.9	20.4	21.2	21.4	0.2
gemma3:27b	**99.8**	0.3	0.5	99.4	15.6	35.9	67.2	0.4	1.5	1.0	1.4	1.4	**0.0**
gemma3:12b	95.4	4.3	4.5	99.5	12.4	24.4	63.9	0.0	0.3	12.6	12.7	12.7	**0.0**
gemma3:4b	50.1	47.1	47.2	94.6	13.0	19.6	52.1	0.0	0.4	62.7	62.8	62.8	**0.0**
llama3.1:70b	58.1	41.8	42.0	97.9	**49.6**	49.0	**79.7**	0.9	**8.0**	94.2	94.2	94.2	**0.0**
llama3.1:8b	94.5	2.3	2.6	94.6	15.9	21.7	53.9	0.0	1.0	11.1	10.9	12.1	1.1
ministral-3:14b	1.1	98.9	98.9	**99.7**	32.5	16.7	60.0	0.0	0.3	**100.0**	**100.0**	**100.0**	**0.0**
ministral-3:8b	18.8	81.7	81.8	96.0	18.3	12.3	67.7	0.0	0.2	99.2	99.0	99.2	0.3
mistral:7b	76.8	25.1	25.3	95.9	8.2	27.2	60.7	0.0	0.1	36.0	38.2	38.5	2.4
mistral-nemo	98.3	0.2	0.4	98.3	15.7	22.6	61.6	0.0	0.2	1.5	2.6	2.6	**0.0**
mixtral:8x7b	84.4	0.9	1.3	84.4	15.4	40.1	59.9	0.0	1.7	7.6	8.2	8.7	0.9
phi4:14b	91.6	8.0	8.3	97.5	17.9	34.7	65.8	0.0	1.5	38.1	39.0	39.5	1.8
qwen2.5:14b	36.0	64.8	64.8	96.9	11.0	38.8	63.7	0.0	0.9	83.3	83.3	83.4	0.4
qwen2.5:7b	18.2	81.7	81.7	97.7	21.0	50.5	77.0	0.0	0.8	**100.0**	**100.0**	**100.0**	**0.0**
qwen3.5:35b	94.0	6.8	7.0	98.3	29.1	55.2	69.4	**3.8**	6.1	74.2	74.2	74.4	0.5
qwen3.5:27b	93.4	7.4	7.8	98.0	31.0	**55.5**	65.2	2.4	6.3	78.8	78.7	78.9	0.5
qwen3.5:9b	97.6	2.3	2.6	97.9	14.5	39.2	57.4	0.3	1.2	20.0	20.6	20.7	0.7
qwen3.5:4b	96.6	1.8	2.4	96.1	9.6	38.7	55.0	0.0	0.6	21.8	23.2	23.4	1.1

and format-following behavior, but the benchmark does not identify which training or post-training choices produced them.

Output Validity and BibTeX Parseability. BibTeX parseability ranges from 0.8% to 99.8%, making output validity a major axis of variation. Low parseability has at least two distinct sources: explicit abstention, which is intentionally not a BibTeX object, and unusable attempted answers containing malformed delimiters, prose wrappers, multiple entries, or missing required structure. For example, low parseability for `qwen2.5:7b` co-occurs with high real-paper refusal, so much of the loss reflects a conservative decision not to answer; `deepseek-r1:8b`, by contrast, produces almost no parseable entries without a corresponding refusal rate, indicating failure to satisfy the output contract after attempting an answer. These behaviors have different practical implications and must not be conflated. Across the 24-model table, parseability and the mean of author F_1, venue F_1, and year exact match have Pearson $r = 0.278$ ($p = 0.189$). The weak association is expected under the benchmark design because producing a syntactically valid template and recalling the correct metadata are separable requirements: a model can learn the surface form of BibTeX while filling it with incorrect fields, or recall some facts while formatting them unusably. The result supports substantial divergence between syntactic validity and factual fidelity, not an unsupported claim of no relationship.

Field-Level Citation Fidelity. Title F_1 is near-saturated because the exact title is supplied, so it primarily measures whether the model preserves visible input rather than retrieves the title from memory; residual errors arise when a model paraphrases, truncates, refuses, or embeds the title in an unusable response. Author attribution is the main bottleneck because a complete author field is a variable-length structured object: every name must be recovered, normalized, and retained. Truncation, missing coauthors, and mismatched lists are common, and the frequent use of "et al." is compatible with ordinary prose citation conventions but violates the benchmark's full-record re-

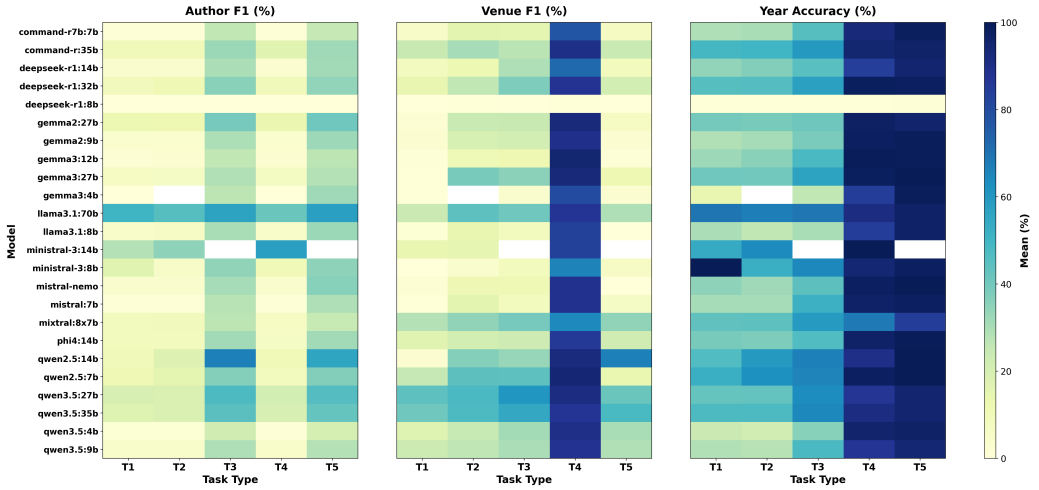

Figure 2: Citation-benchmark performance by task variant. Rows are models and columns are task variants. Panels report conditional title, author, and venue F_1; higher is better.
Task variants: T1 = Title only; **T2** = Title only + DOI/URL requirement; **T3** = Title + author hint; **T4** = Title + venue + year; **T5** = Title + first author + year. See Section 2.1 for full definitions.

quirement. Venue fidelity is higher but limited by aliases, abbreviated conference names, and confusion among proceedings, preprints, and journal versions. Year exact match is easier to express but can still drift when a work has distinct preprint, conference, and publication dates; the aggregates cannot determine which ambiguity caused any individual error. DOI exact match peaks at 3.8%. Unlike titles and years, DOI strings are rarely recoverable from semantic similarity: every character must match after normalization, and the identifier is not supplied by the prompt. This makes closed-book reconstruction demanding, although the design does not establish identifier entropy as the sole cause. The metric therefore separates a correct normalized DOI; an omitted, different, or fabricated DOI when valid gold exists; an arXiv identifier or stable URL supplied instead; and a record with no valid gold DOI, which is unscored rather than assigned zero.

Performance by Task Variant. Variant differences are primarily evidence effects. When an author hint is present, mean author F_1 rises from 10.2 to 33.0; when venue is supplied, mean venue F_1 rises from 19.9 to 81.9; and when year is supplied, mean year exact match rises from 42.9 to 90.9. Supplying a field removes part of the closed-book recall burden and gives the model an exact lexical anchor that can be copied into the BibTeX entry, so these gains measure evidence use and instruction following as well as memory. They should not be described as improved "partial metadata retrieval." The increases are also incomplete, especially for authors, because models may still abbreviate, normalize incorrectly, omit supplied information, or fail to produce a scoreable entry. The title is invariant across all five variants, so cross-model mean title F_1 remains high and changes little (90.3–94.0), as expected if most models copy the shared anchor. The identifier condition is different: it adds an exact-output burden without supplying identifier evidence. Relative to title-only, it raises strict refusal from 19.8 to 26.6 and lowers parseability from 74.7 to 68.2, consistent with some models responding to uncertainty by abstaining and others producing more complex, invalid outputs. DOI exact match is approximately unchanged (0.52 versus 0.54), indicating that requesting an identifier changes response behavior but does not materially improve closed-book recovery of the correct DOI.

Refusal Behavior and Answerability. Abstention varies across nearly the full range. On real, temporally eligible prompts, lower refusal is generally preferable because the benchmark target exists, although eligibility does not prove that its metadata was present in training. On synthetic controls, higher refusal is preferable: `ministral-3:14b` and `qwen2.5:7b` reach 100%, while `gemma3:27b`, `mistral-nemo`, and `command-r7b:7b` remain below 2%. A useful contrast is `qwen2.5:7b`, which refuses every synthetic prompt but also 81.7% of real prompts, versus

gemma3:27b, which refuses 1.0% of synthetic prompts and 0.3% of real prompts. This pattern is consistent with different answer thresholds: a conservative threshold suppresses fabrication but sacrifices coverage, whereas an aggressive threshold returns more real records while also accepting plausible no-answer titles. The synthetic prompts match the form and topical plausibility of real prompts, so superficial formatting cues should not be sufficient for perfect separation. Refusal-heavy models can also appear stronger on conditional field metrics because only their smaller answered subset is scored, making coverage essential to interpretation. The benchmark does not attribute these thresholds to a particular training method, but it shows why capability and willingness to answer must be evaluated together rather than collapsed into one accuracy number.

Qualitative Error Analysis. The results reveal four failure modes. First, some checkpoints return unusable attempts rather than abstain. deepseek-r1:8b has 0.8% BibTeX parseability but only 0.2% strict refusal on real prompts, indicating output-contract failures such as prose, malformed braces, or multiple objects. Second, some models over-abstain. ministral-3:14b refuses 98.9% of real prompts and 100% of synthetic controls; qwen2.5:7b refuses 81.7% and 100%, respectively. Their fabrication avoidance therefore comes at the cost of real-paper coverage.

Third, valid syntax can conceal metadata errors. gemma3:27b reaches 99.8% parseability and 99.4 title F_1, but only 15.6 author F_1, 35.9 venue F_1, and 67.2% year exact match. The supplied title is copied reliably, while recalled fields remain weaker. Common errors include "et al." in place of full authors, omitted coauthors, and nearby venues or years, potentially reflecting citation abbreviations or multiple publication versions. Fourth, identifier errors escape field-overlap metrics: models may fabricate or omit a DOI, or return a stable URL that satisfies the prompt but not exact-DOI scoring. DOI scoring applies only when valid gold metadata exists. Appendix Section A.6 shows how valid BibTeX and several correct fields can coexist with severe author truncation and a bad identifier.

4 Related Work

Agrawal et al. test whether language models distinguish real references from hallucinated ones and whether their expressed uncertainty identifies likely fabrication (Agrawal et al., 2023); Byun et al. measure citation existence, source accuracy, and relevance (Byun et al., 2024); and Linardon et al. vary topic familiarity and prompt specificity to study citation fabrication in mental-health research (Linardon et al., 2025). GhostCite broadens the validity question to a large-scale citation audit (Xu et al., 2026), Niimi analyzes hierarchical memorization through citation generation (Niimi, 2025), and Alansari and Luqman synthesize the wider LLM-hallucination literature (Alansari & Luqman, 2025). Tang et al. evaluate reference generation together with abstract writing and review composition (Tang et al., 2025), whereas Qian et al. focus directly on the capacity of LLMs to generate citations (Qian et al., 2024). ArxEval jointly studies retrieval and generation for scientific literature (Sinha et al., 2025), and ArXivBench evaluates title, link, and broader academic-writing behavior (Li et al., 2025). A related but distinct line evaluates attribution: Gao et al. generate text with citations to supplied sources (Gao et al., 2023), Wu et al. assess whether medical claims cite relevant references (Wu et al., 2025), Tamber et al. and Yeh et al. examine faithfulness in retrieval-augmented settings (Tamber et al., 2025; Yeh et al., 2025), and Cite Pretrain targets retrieval-free knowledge attribution (Huang et al., 2025). General abstention work studies unanswerable questions, false premises, and unsafe requests rather than structured bibliography records (Madhusudhan et al., 2025; Wen et al., 2025; Kirichenko et al., 2025; Xie et al., 2025; Liu et al., 2025).

5 Conclusion

HalluciBench tests bibliographic reconstruction and abstention across 33 locally deployable checkpoints using matched real-paper and no-answer prompts. The results expose a central failure: valid BibTeX, faithful metadata, DOI recovery, and calibrated refusal are distinct capabilities. These failures persist even for prominent, older ML papers with the exact title supplied. Generative citation systems should therefore be evaluated field by field and deployed only with deterministic bibliographic retrieval and verification.

AI USE STATEMENT

Generative AI tools were used only to improve the manuscript's grammar, clarity, concision, and readability. They were not used to generate or refine research ideas, hypotheses, methods, experiments, theoretical or mathematical claims, proofs, synthetic data, code, data processing or analysis, result interpretation, translation, literature review, citations, or references. The authors reviewed all AI-assisted edits and take full responsibility for the final content of this work.

ETHICS STATEMENT

This work evaluates language models using public bibliographic metadata and synthetically constructed no-answer controls. It involves no human subjects, private or sensitive personal data, or user interaction; author names are used only as published citation metadata. The primary foreseeable risk is misuse of generated citations. We therefore present HalluciBench as an evaluation diagnostic rather than a bibliographic source and recommend deterministic retrieval and external verification for deployed systems. The released artifacts document dataset construction, missing metadata, and benchmark limitations. The authors have read and adhered to the ICLR Code of Ethics.

REPRODUCIBILITY STATEMENT

Section 2.1 documents dataset sourcing, canonicalization, prompt variants, and no-answer-control construction; the execution and evaluation sections specify temporal filtering, inference conditions, parsing, denominators, and metrics. The appendix provides the record schema, complete metric definitions, prompt and judge templates, checkpoint release dates, a worked example, leakage-audit details, and full result tables. Code and available benchmark artifacts are provided at `https://anonymous.4open.science/r/HalluciBench-8E14/README.md`. The limitations section identifies archived information that is unavailable for bit-level replication.

REFERENCES

Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, et al. Phi-4 technical report, 2024. URL https://arxiv.org/abs/2412.08905.

Ayush Agrawal, Mirac Suzgun, Lester Mackey, and Adam Tauman Kalai. Do language models know when they're hallucinating references?, 2023. URL https://arxiv.org/abs/2305.18248.

Aisha Alansari and Hamzah Luqman. Large language models hallucination: A comprehensive survey, 2025. URL https://arxiv.org/abs/2510.06265.

Anthropic. Introducing Claude Sonnet 5, June 2026. URL https://www.anthropic.com/news/claude-sonnet-5. Accessed 2026-08-30.

Yejin Bang, Ziwei Ji, Alan Schelten, Anthony Hartshorn, Tara Fowler, Cheng Zhang, Nicola Cancedda, and Pascale Fung. Hallulens: LLM hallucination benchmark. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 24128–24156, Vienna, Austria, July 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.1176. URL https://aclanthology.org/2025.acl-long.1176/.

Courtni Byun, Piper Vasicek, and Kevin Seppi. This reference does not exist: An exploration of LLM citation accuracy and relevance. In Su Lin Blodgett, Amanda Cercas Curry, Sunipa Dev, Michael Madaio, Ani Nenkova, Diyi Yang, and Ziang Xiao (eds.), *Proceedings of the Third Workshop on Bridging Human–Computer Interaction and Natural Language Processing*, pp. 28–39, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.hcinlp-1.3. URL https://aclanthology.org/2024.hcinlp-1.3/.

DeepSeek AI. DeepSeek-R1 official repository. GitHub repository, 2025. URL https://github.com/deepseek-ai/DeepSeek-R1. Public release recorded on 2025-01-20; accessed 2026-08-30.

FACTS Team. FACTS grounding: A new benchmark for evaluating the factuality of large language models. Google DeepMind, December 2024. URL https://deepmind.google/blog/facts-grounding-a-new-benchmark-for-evaluating-the-factuality-of-large-language-models/.

Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. Enabling large language models to generate text with citations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 6465–6488, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.398. URL https://aclanthology.org/2023.emnlp-main.398/.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, et al. Gemma 2: Improving open language models at a practical size, 2024. URL https://arxiv.org/abs/2408.00118.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, et al. Gemma 3 technical report, 2025. URL https://arxiv.org/abs/2503.19786.

Aidan Gomez. Command R: Retrieval-Augmented Generation at Production Scale, March 2024a. URL https://cohere.com/blog/command-r. Accessed: 2026-03-31.

Aidan Gomez. Introducing Command R7B: Fast and efficient generative AI, December 2024b. URL https://cohere.com/blog/command-r7b. Accessed: 2026-03-31.

Google AI for Developers. Gemma release history, 2026. URL https://ai.google.dev/gemma/docs/releases. Documents checkpoint release dates; accessed 2026-08-30.

Google DeepMind. Gemma 3 model card, 2025. URL https://ai.google.dev/gemma/docs/core/model_card_3. Documents an August 2024 knowledge cutoff; accessed 2026-08-30.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, et al. The Llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, et al. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645(8081):633–638, September 2025. doi: 10.1038/s41586-025-09422-z. URL https://www.nature.com/articles/s41586-025-09422-z.

Yukun Huang, Sanxing Chen, Jian Pei, Manzil Zaheer, and Bhuwan Dhingra. Cite pretrain: Retrieval-free knowledge attribution for large language models, 2025. URL https://arxiv.org/abs/2506.17585.

International Conference on Learning Representations. ICLR 2026 response to LLM-generated papers and reviews. ICLR Blog, 2025. URL https://blog.iclr.cc/2025/11/19/iclr-2026-response-to-llm-generated-papers-and-reviews/. Accessed 2026-08-30.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, et al. Mistral 7B, 2023. URL https://arxiv.org/abs/2310.06825. Public release announced 2023-09-27 at https://mistral.ai/news/announcing-mistral-7b/.

Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, et al. Mixtral of experts, 2024. URL https://arxiv.org/abs/2401.04088. Public release announced 2023-12-11 at https://mistral.ai/news/mixtral-of-experts/.

Aja Khanal, Kaushik T. Ranade, Rishabh Agrawal, Kalyan S. Basu, and Apurva Narayan. Agents of diffusion: Enhancing diffusion language models with multi-agent reinforcement learning for structured data generation. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 449–458, Paphos, Cyprus, 2026. International Foundation for Autonomous Agents and Multiagent Systems. doi: 10.65109/GGJL7344. URL https://dl.acm.org/doi/10.65109/GGJL7344.

Polina Kirichenko, Mark Ibrahim, Kamalika Chaudhuri, and Samuel J. Bell. AbstentionBench: Reasoning LLMs fail on unanswerable questions, 2025. URL https://arxiv.org/abs/2506.09038.

Ning Li, Jingran Zhang, and Justin Cui. ArXivBench: When you should avoid using ChatGPT for academic writing, 2025. URL https://arxiv.org/abs/2504.10496.

Jake Linardon, Hannah K Jarman, Zoe McClure, Cleo Anderson, Claudia Liu, and Mariel Messer. Influence of topic familiarity and prompt specificity on citation fabrication in mental health research using large language models: Experimental study. *JMIR Ment Health*, 12:e80371, Nov 2025. ISSN 2368-7959. doi: 10.2196/80371. URL https://mental.jmir.org/2025/1/e80371.

Alexander H. Liu, Kartik Khandelwal, Sandeep Subramanian, Victor Jouault, Abhinav Rastogi, et al. Ministral 3, 2026. URL https://arxiv.org/abs/2601.08584.

Yi Liu, Xiangyu Liu, Zequn Sun, and Wei Hu. Answering the unanswerable is to err knowingly: Analyzing and mitigating abstention failures in large reasoning models, 2025. URL https://arxiv.org/abs/2508.18760.

Wen Luo, Tianshu Shen, Wei Li, Guangyue Peng, Richeng Xuan, Houfeng Wang, and Xi Yang. HalluDial: A large-scale benchmark for automatic dialogue-level hallucination evaluation, 2024. URL https://arxiv.org/abs/2406.07070.

Nishanth Madhusudhan, Sathwik Tejaswi Madhusudhan, Vikas Yadav, and Masoud Hashemi. Do LLMs know when to NOT answer? investigating abstention abilities of large language models. In Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (eds.), *Proceedings of the 31st International Conference on Computational Linguistics*, pp. 9329–9345, Abu Dhabi, UAE, January 2025. Association for Computational Linguistics. URL https://aclanthology.org/2025.coling-main.627/.

Meta. Llama 3.1 8B Instruct model card, 2024a. URL https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct. Documents a December 2023 pretraining-data cutoff; accessed 2026-08-30.

Meta. Llama 3.2 1B Instruct model card, 2024b. URL https://huggingface.co/meta-llama/Llama-3.2-1B-Instruct. Documents a December 2023 pretraining-data cutoff; accessed 2026-08-30.

Meta AI. Introducing Llama 3.1: Our most capable models to date, July 2024a. URL https://ai.meta.com/blog/meta-llama-3-1/. Accessed 2026-08-30.

Meta AI. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models, September 2024b. URL https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/. Accessed 2026-08-30.

Microsoft. Phi-4 model card, 2024a. URL https://huggingface.co/microsoft/phi-4. Public release 2024-12-12; documents a June 2024 cutoff for public data; accessed 2026-08-30.

Microsoft. Introducing Phi-4: Microsoft's newest small language model specializing in complex reasoning, December 2024b. URL https://techcommunity.microsoft.com/blog/azure-ai-foundry-blog/introducing-phi-4-microsofts-newest-small-language-model-specializing-in-comple/4357090. Accessed 2026-08-30.

Microsoft. Phi-4-mini-instruct model card, 2025a. URL https://huggingface.co/microsoft/Phi-4-mini-instruct. Documents a June 2024 cutoff for public data; accessed 2026-08-30.

Microsoft. Empowering innovation: The next generation of the Phi family, February 2025b. URL https://azure.microsoft.com/en-us/blog/empowering-innovation-the-next-generation-of-the-phi-family/. Accessed 2026-08-30.

Microsoft, Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, et al. Phi-4-Mini technical report: Compact yet powerful multimodal language models via mixture-of-loras, 2025. URL https://arxiv.org/abs/2503.01743.

Nipun Misra and Vikranth Udandarao. Detecting citation hallucinations in large language model outputs (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 41325–41327, 2026. doi: 10.1609/aaai.v40i48.42257. URL https://ojs.aaai.org/index.php/AAAI/article/view/42257. Issue 48; Student Abstract.

Mistral AI. Introducing Mistral 3, December 2025. URL https://mistral.ai/news/mistral-3/. Accessed 2026-08-30.

Mistral AI team. Mistral NeMo, July 2024. URL https://mistral.ai/news/mistral-nemo. Accessed: 2026-03-31.

Neural Information Processing Systems Conference. Main track handbook 2026. Neural Information Processing Systems Conference, 2026. URL https://neurips.cc/Conferences/2026/MainTrackHandbook. Accessed 2026-08-30.

Junichiro Niimi. Hierarchical memorization in large language models: Evidence from citation generation, 2025. URL https://arxiv.org/abs/2511.08877.

OpenAI. Introducing GPT-5.5, April 2026a. URL https://openai.com/index/introducing-gpt-5-5/. Accessed 2026-08-30.

OpenAI. Previewing GPT-5.6 Sol: A next-generation model, 2026b. URL https://openai.com/index/previewing-gpt-5-6-sol/. Accessed 2026-08-30.

Haosheng Qian, Yixing Fan, Ruqing Zhang, and Jiafeng Guo. On the capacity of citation generation by large language models, 2024. URL https://arxiv.org/abs/2410.11217.

Qwen Team. Qwen2.5: A party of foundation models!, September 2024. URL https://qwenlm.github.io/blog/qwen2.5/. Accessed 2026-08-30.

Qwen Team. Qwen3.5 official release log. GitHub repository, 2026a. URL https://github.com/QwenLM/Qwen3.8. The official repository retains the staged Qwen3.5 release log; accessed 2026-08-30.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026b. URL https://qwen.ai/blog?id=qwen3.5. Accessed: 2026-03-31.

Qwen Team, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, et al. Qwen2.5 technical report, 2024. URL https://arxiv.org/abs/2412.15115.

Aarush Sinha, Viraj Virk, Dipshikha Chakraborty, and P. S. Sreeja. ArxEval: Evaluating retrieval and generation in language models for scientific literature, 2025. URL https://arxiv.org/abs/2501.10483.

Manveer Singh Tamber, Forrest Sheng Bao, Chenyu Xu, Ge Luo, Suleman Kazi, Minseok Bae, Miaoran Li, Ofer Mendelevitch, Renyi Qu, and Jimmy Lin. Benchmarking LLM faithfulness in RAG with evolving leaderboards, 2025. URL https://arxiv.org/abs/2505.04847.

Xuemei Tang, Xufeng Duan, and Zhenguang G. Cai. Large language models for automated literature review: An evaluation of reference generation, abstract writing, and review composition. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 1602–1617, Suzhou, China, November 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.emnlp-main.83. URL https://aclanthology.org/2025.emnlp-main.83/. Originally released as arXiv:2412.13612.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu Wang. Know your limits: A survey of abstention in large language models. *Transactions of the Association for Computational Linguistics*, 13:529–556, 2025. doi: 10.1162/tacl_a_00754. URL https://aclanthology.org/2025.tacl-1.26/.

Kevin Wu, Eric Wu, Kevin Wei, Angela Zhang, Allison Casasola, Teresa Nguyen, Sith Riantawan, Patricia Shi, Daniel Ho, and James Zou. An automated framework for assessing how well LLMs cite relevant medical references. *Nature Communications*, 16:3615, April 2025. doi: 10.1038/s41467-025-58551-6. URL https://www.nature.com/articles/s41467-025-58551-6.

Tinghao Xie, Xiangyu Qi, Yi Zeng, Yangsibo Huang, Udari Madhushani Sehwag, Kaixuan Huang, Luxi He, Boyi Wei, Dacheng Li, Ying Sheng, Ruoxi Jia, Bo Li, Kai Li, Danqi Chen, Peter Henderson, and Prateek Mittal. SORRY-Bench: Systematically evaluating large language model safety refusal behaviors. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=YfKNaRktan.

Zuyao Xu, Yuqi Qiu, Lu Sun, Fasheng Miao, Fubin Wu, Xiang Li, Xinyi Wang, Haozhe Lu, Zhengze Zhang, Yuxin Hu, Jialu Li, Luo Jin, Feng Zhang, Rui Luo, Xinran Liu, Yingxian Li, and Jiaji Liu. GhostCite: A large-scale analysis of citation validity in the age of large language models, 2026. URL https://arxiv.org/abs/2602.06718.

Samuel Yeh, Sharon Li, and Tanwi Mallick. LUMINA: Detecting hallucinations in RAG system with context-knowledge signals, 2025. URL https://arxiv.org/abs/2509.21875.

A Appendix

A.1 Limitations and Future Work

HalluciBench is a closed-book, title-anchored intrinsic probe rather than an end-to-end bibliography assistant: it excludes deterministic lookup, RAG, tools, and agentic collaboration. Its two-year pre-release buffer reduces temporal confounding but cannot establish training-corpus inclusion, and the absence of proprietary frontier target models leaves their performance unknown. The prominence-biased ML corpus does not cover long-tail or cross-disciplinary papers, non-paper sources, multilingual and complex author conventions, version ambiguity, or unusual BibTeX types. AoD's bidirectional generation, iterative optimizer–judge feedback, novelty criterion, field-overlap diagnostic, and our manual nearest-title token-overlap audit ensure that retained synthetic controls differ from the fixed real-paper construction set; however, no exhaustive external-database collision audit is preserved, so these safeguards do not certify global nonexistence. All checkpoints used local runtime defaults without per-model tuning, but the exact runtime version, immutable weight digests, and resolved numeric defaults were not archived, preventing bit-level replication. Conditional field coverage is model-dependent, and missing response-level manifests prevent verified unconditional recomputation. Semantic Scholar metadata is also incomplete: the 444 papers without a valid gold DOI are excluded from DOI scoring rather than counted as model failures.

Future work should compare the same records under closed-book inference, deterministic lookup, RAG, single-agent tools, and multi-agent systems; add proprietary and newer open checkpoints; extend coverage to long-tail, cross-disciplinary, multilingual, and nonstandard sources; independently validate the AoD-generated controls against broader bibliographic databases; preserve complete execution manifests; and grade exact recovery, useful nearby-paper suggestions, refusal, and fabrication separately.

Metric direction. In every appendix table, ↑ indicates that higher values are better and ↓ indicates that lower values are better. Refusal on real-paper prompts is marked ↓, refusal on synthetic no-answer controls is marked ↑, and refusal–parseability conflict is marked ↓.

A.2 Dataset Details

This section describes the structure of each HalluciBench record and explains the role of every metadata field. Both the real citation retrieval set and the synthetic refusal set use the same top-level

Table 2: Dataset statistics for the fixed HalluciBench pool. Counts are descriptive (no direction is better). The 6,000-record total is a construction-time pool; each model uses a checkpoint-specific subset satisfying the two-year pre-release rule.

Statistic	Count	Definition / distribution
Unique real papers	1,000	Deduplicated by Semantic Scholar paper ID
Real prompt records	5,000	Five title-anchored variants per real paper
Unique synthetic controls	200	Constructed records with no intended target
Synthetic prompt records	1,000	Five matched variants per synthetic control
Total prompt pool	6,000	5,000 real + 1,000 synthetic
Venues	1,000 papers	CVPR 229; NeurIPS 216; ICLR 172; ICML 134; ICCV 78; ECCV 76; ACL 54; EMNLP 41
Publication-year buckets	1,000 papers	2016–2018: 504; 2019–2022: 428; 2023–2024: 68
Citation-count buckets	1,000 papers	`high` ($\geq$ 1,000): 735; `very_high` ($\geq$ 10,000): 265
Entry types	1,000 papers	`@inproceedings`: 694; `@article`: 306
Gold DOI available	556 papers	55.6% of papers; repeated across 2,780 real prompt records
Gold DOI unavailable	444 papers	44.4% of papers; repeated across 2,220 real prompt records and excluded from DOI scoring

JSON schema so that models can be evaluated under a consistent interface across answerable and unanswerable citation requests. In the real retrieval set, the gold target is a canonical BibTeX entry for a real paper. In the synthetic refusal set, the record structure is the same, but the gold target is always `REFUSE`.

A.2.1 Synthetic-Control Leakage Audit

We manually audited the 200 synthetic controls against the 1,000-paper real pool. Since controls intentionally reuse real venue names, comparison used normalized title tokens rather than full records. Let $T(x)$ be the set of lowercase title tokens after punctuation removal. For each synthetic title s and real-title set $\mathcal{R}$, we computed its nearest-neighbor Jaccard overlap,

$$\text{Overlap}(s; \mathcal{R}) = \max_{r \in \mathcal{R}} \frac{|T(s) \cap T(r)|}{|T(s) \cup T(r)|}.$$

Reviewers inspected the highest-overlap real match for each synthetic title together with its author metadata and removed records that appeared copied, memorized, or trivially perturbed. We also queried each synthetic citation through the Semantic Scholar bulk search API and removed records with similar matches in the broader literature.

A.2.2 Top-Level Record Structure

Each benchmark instance is stored as a JSON object with the following top-level fields:

```
{
  "id": ...,
  "prompt": ...,
  "task_type": ...,
  "target_metadata": { ... },
  "ground_truth": { ... },
  "analysis_metadata": { ... }
}
```

This structure is used in both target files and is constructed programmatically when the benchmark dataset is built.

`id`. A unique integer identifier for the benchmark instance. This field identifies the *example*, not the underlying paper. Since multiple prompt variants can be created from the same paper, different records may share the same bibliographic target metadata while having different `id` values and different prompt formulations.

prompt. The exact user-facing citation request given to the evaluated model. This field contains the full instruction text, including the citation retrieval task, any metadata hints, and the required refusal behavior. HalluciBench uses a passthrough inference protocol, so the prompt is sent directly to the target model without additional benchmark-specific wrapper text. Prompt variants differ in how much evidence is provided, for example exact title only, title plus author hint, title plus venue and year, or partial metadata.

task_type. A symbolic label indicating which citation retrieval variant the instance belongs to. In the current benchmark, task types include:

- `exact_title_to_bibtex`
- `exact_title_to_bibtex_with_identifier`
- `title_author_to_bibtex`
- `title_venue_year_to_bibtex`
- `partial_metadata_to_bibtex`

This field is used to stratify results by task difficulty and by the amount of bibliographic evidence available to the model.

The field `target_metadata` stores the structured bibliographic information associated with the benchmark instance:

```
"target_metadata": {
  "paper_id": ...,
  "title": ...,
  "authors": [...],
  "venue": ...,
  "year": ...,
  "doi": ...,
  "url": ...,
  "publication_date": ...,
  "citation_count": ...
}
```

These fields define the gold reference object against which citation outputs are evaluated, even when the gold target in the refusal dataset is `REFUSE`.

paper_id. A stable identifier for the underlying paper record. In the real dataset, this is the Semantic Scholar paper identifier. In the synthetic refusal set, the same field is retained for schema consistency, but uses a synthetic prefix such as `SYNTHETIC-REFUSAL-...`. This allows the benchmark to preserve a uniform schema across real and synthetic instances.

title. The canonical paper title associated with the instance. In the real dataset, this is the title of a real paper. In the synthetic refusal set, the title is intentionally fabricated but designed to be plausible. This field is central to all prompt variants and is also used in title-level evaluation metrics.

authors. A list of author names in display form. This field stores the author set used for author-level precision, recall, and F_1 evaluation. In the real dataset, it reflects the real paper's author list; in the synthetic refusal set, it is fabricated but citation-like. The list representation makes it possible to evaluate omitted, duplicated, or hallucinated authors at the set level rather than as a single string.

venue. The normalized venue or journal name. Venue names are canonicalized during dataset construction so that equivalent surface forms, such as abbreviations and full venue names, map to a consistent representation. This supports more stable venue-level evaluation and reduces spurious mismatch due to formatting alone.

`year`. The publication year associated with the target record. This field is included both as prompt evidence in some task variants and as a bibliographic field for exact-match evaluation. In the refusal set, it remains populated to preserve realism and to prevent models from detecting synthetic records through missing metadata.

`doi`. The DOI string, when available. For real papers, this is extracted from the source metadata when present. For synthetic refusal records, a DOI-like field may also be populated to preserve bibliographic realism, even though the correct output remains `REFUSE`. This field is used for DOI exact-match evaluation in the real retrieval benchmark.

`url`. A stable URL associated with the target record. In the real dataset this is typically a Semantic Scholar paper URL. This field is useful both as a potential citation identifier and as part of task variants that require a stable identifier such as a DOI or URL. In the refusal set, the field is retained for structural consistency and realism.

`publication_date`. A finer-grained publication date string, when available. This field is not the primary target of evaluation, but it is used when constructing canonical BibTeX references and supports richer metadata analysis. In some real instances it may be null or absent in the upstream source.

`citation_count`. The citation count associated with the underlying record. This field is not required for citation correctness itself, but it supports stratified analysis by paper prominence or retrieval difficulty. HalluciBench uses it to derive citation-count buckets stored in `analysis_metadata`.

A.2.3 Ground Truth

The field `ground_truth` stores the expected benchmark target:

```
"ground_truth": {
  "bibtex": ...
}
```

`ground_truth["bibtex"]`. In the real citation retrieval set, this field contains the canonical BibTeX entry constructed

A.3 Evaluation Metrics

This section provides the formal definitions of the evaluation metrics used in HalluciBench, together with the rationale for including each one. Our metric design is intentionally multi-part because citation hallucination is not a single failure mode. A model may fail by producing unusable output, by returning a structured but bibliographically incorrect citation, or by answering when it should instead refuse. HalluciBench therefore separates *response validity*, *field-level bibliographic fidelity*, and *semantic answerability/refusal*.

Let $D = \{(p_i, y_i, t_i)\}_{i=1}^{N}$ denote the fixed pool, including publication date t_i for real papers. For model m released at r_m, define the temporally eligible real-paper prompts

$$E_m = \{i : t_i \leq r_m - 2 \text{ years}\}.$$

Let r_i denote a response and $R_m = \{i \in E_m : r_i \text{ is returned}\}$. Conditional field quality uses S_m below; unconditional end-to-end scores use all of E_m; DOI uses a separate gold-eligible subset. These denominators must not be interchanged. The archived aggregate tables preserve rates but not exact model-specific counts.

BibTeX parseability. We define

$$\mathrm{Parse}(r_i) = \begin{cases} 1 & \text{if } r_i \text{ parses as exactly one BibTeX entry,} \\ 0 & \text{otherwise.} \end{cases}$$

The parseability rate is

$$\text{Parseability} = \frac{1}{|R_m|} \sum_{i \in R_m} \text{Parse}(r_i).$$

We include parseability because a citation-like response is not operationally useful unless it can be consumed as a valid structured object. A model may produce plausible prose or malformed BibTeX that appears knowledgeable but remains unusable in practice. HalluciBench therefore treats output validity as a first-class property rather than assuming all responses are scorable.

Strict refusal rate. HalluciBench also records whether a response contains explicit refusal language using a deterministic check. Define

$$\text{StrictRefuse}(r_i) = \begin{cases} 1 & \text{if } r_i \text{ contains an explicit refusal signal,} \\ 0 & \text{otherwise.} \end{cases}$$

The strict refusal rate is

$$\text{StrictRefusal} = \frac{1}{|R_m|} \sum_{i \in R_m} \text{StrictRefuse}(r_i).$$

We use strict refusal because a model should abstain when evidence is insufficient or when a synthetic control has no intended target. A deterministic signal is transparent and reproducible.

Judge-based semantic refusal. HalluciBench optionally uses an LLM judge to evaluate refusal and correctness at a semantic level. Let J denote the set of successfully judged responses. For refusal-oriented judging, define

$$\text{JudgeRefuse}_i = \begin{cases} 1 & \text{if the judge marks the response as a refusal,} \\ 0 & \text{otherwise.} \end{cases}$$

Then the semantic refusal rate is

$$\text{JudgeRefusal} = \frac{1}{|J|} \sum_{i \in J} \text{JudgeRefuse}_i.$$

We include judge refusal because deterministic refusal matching captures only explicit refusal language, whereas the judge can identify semantically equivalent abstention that may not use the benchmark's canonical wording.

Strict-or-judge refusal. Finally, we report a combined refusal statistic:

$$\text{StrictOrJudgeRefusal} = \frac{1}{|J|} \sum_{i \in J} \max\{\text{StrictRefuse}(r_i), \text{JudgeRefuse}_i\}.$$

This metric provides a conservative estimate of abstention behavior by taking the union of deterministic and semantic refusal signals. It is useful because the two signals capture slightly different phenomena: deterministic refusal matching is precise and reproducible, while judge-based refusal is more semantically flexible.

Refusal–parseability overlap. We additionally report

$$\text{StrictRefusal\&Parseable} = \frac{1}{|R_m|} \sum_{i \in R_m} \text{StrictRefuse}(r_i)\,\text{Parse}(r_i).$$

This overlap distinguishes abstention from output formatting. A model may refuse correctly but fail to do so in a valid BibTeX-compatible form, or it may return a parseable citation when refusal would have been the correct behavior. This distinction is especially important on the synthetic no-answer controls.

Citation-quality scoring subset. Conditional field-level metrics use

$$S_m = \{i \in R_m : \neg(\mathrm{StrictRefuse}(r_i) = 1 \wedge \mathrm{Parse}(r_i) = 0)\}.$$

Thus, explicit non-parseable refusals are excluded; parseable entries are included; and non-parseable non-refusal attempts are included with zero field credit. A parseable response containing refusal language remains in S_m and is captured by the overlap diagnostic. Conditional scores diagnose attempted citations but can reward selective refusal. An unconditional end-to-end field score averages over E_m and assigns zero to failed inference, refusals, and unusable attempts. The archived aggregate artifact lacks the membership and counts required to recompute that view without approximation.

Field extraction. For parseable $i \in S_m$, the parser extracts predicted bibliographic metadata

$$\hat{m}_i = (\hat{A}_i, \hat{V}_i, \hat{T}_i, \hat{d}_i, \hat{y}_i),$$

where $\hat{A}_i$ is the predicted author list, $\hat{V}_i$ the predicted venue string, $\hat{T}_i$ the predicted title string, $\hat{d}_i$ the predicted DOI, and $\hat{y}_i$ the predicted year. These are compared against gold metadata

$$m_i = (A_i, V_i, T_i, d_i, y_i).$$

For a non-parseable, non-refusal attempt in S_m, all field scores are zero. Rather than using full-string exact match on the entire entry, HalluciBench evaluates each field with a metric suited to its structure.

Author precision, recall, and F_1. Authors are evaluated under a conservative matching policy that requires exact last-name agreement, allows initials-aware given-name matching, enforces one-to-one alignment, and treats ambiguous predictions as unmatched. Let TP_i^A, FP_i^A, and FN_i^A denote the resulting true-positive, false-positive, and false-negative author counts for example i. Then

$$P_i^A = \frac{\mathrm{TP}_i^A}{\mathrm{TP}_i^A + \mathrm{FP}_i^A}, \qquad R_i^A = \frac{\mathrm{TP}_i^A}{\mathrm{TP}_i^A + \mathrm{FN}_i^A},$$

$$F_{1,i}^A = \frac{2P_i^A R_i^A}{P_i^A + R_i^A},$$

with score 0 used whenever the denominator is zero. HalluciBench reports macro-averages over the scoring subset:

$$\overline{P^A} = \frac{1}{|S_m|}\sum_{i \in S_m} P_i^A, \qquad \overline{R^A} = \frac{1}{|S_m|}\sum_{i \in S_m} R_i^A, \qquad \overline{F_1^A} = \frac{1}{|S_m|}\sum_{i \in S_m} F_{1,i}^A.$$

We use author-set precision, recall, and F_1 because authorship is inherently multi-valued and minor formatting differences should not count as full errors. This metric still penalizes omitted, duplicated, and hallucinated authors in a principled way.

Venue precision, recall, and F_1. Venue strings are first canonicalized to normalize aliases such as abbreviated and long-form conference names. Let $\hat{\mathcal{V}}_i$ and $\mathcal{V}_i$ denote the predicted and gold venue token sets after normalization. Then

$$P_i^V = \frac{|\hat{\mathcal{V}}_i \cap \mathcal{V}_i|}{|\hat{\mathcal{V}}_i|}, \qquad R_i^V = \frac{|\hat{\mathcal{V}}_i \cap \mathcal{V}_i|}{|\mathcal{V}_i|},$$

$$F_{1,i}^V = \frac{2P_i^V R_i^V}{P_i^V + R_i^V}.$$

Macro-averages over S are reported analogously to the author metrics. We use token-level venue metrics because exact string equality would overstate error for semantically equivalent venue forms, while token overlap still penalizes incorrect or hallucinated venue information.

Title unigram precision, recall, and F_1. Let $\hat{\mathcal{U}}_i$ and $\mathcal{U}_i$ denote the normalized unigram sets of the predicted and gold titles. Then

$$P_i^{T,1} = \frac{|\hat{\mathcal{U}}_i \cap \mathcal{U}_i|}{|\hat{\mathcal{U}}_i|}, \qquad R_i^{T,1} = \frac{|\hat{\mathcal{U}}_i \cap \mathcal{U}_i|}{|\mathcal{U}_i|},$$

$$F_{1,i}^{T,1} = \frac{2P_i^{T,1}R_i^{T,1}}{P_i^{T,1} + R_i^{T,1}}.$$

Macro-averages over S are reported. Unigram title metrics capture partial lexical correctness: they reward recovery of the right title content even when phrase order or local phrasing is imperfect.

Title bigram precision, recall, and F_1. Let $\hat{\mathcal{B}}_i$ and $\mathcal{B}_i$ denote the normalized bigram sets of the predicted and gold titles. Then

$$P_i^{T,2} = \frac{|\hat{\mathcal{B}}_i \cap \mathcal{B}_i|}{|\hat{\mathcal{B}}_i|}, \qquad R_i^{T,2} = \frac{|\hat{\mathcal{B}}_i \cap \mathcal{B}_i|}{|\mathcal{B}_i|},$$

$$F_{1,i}^{T,2} = \frac{2P_i^{T,2}R_i^{T,2}}{P_i^{T,2} + R_i^{T,2}}.$$

Again, macro-averages over S are reported. We include bigram metrics because unigram overlap alone may over-credit titles that merely contain related keywords. Bigram overlap is more sensitive to phrase-level corruption and therefore better reflects whether the model reconstructed the true paper title rather than a plausible paraphrase.

DOI exact match. The DOI metric is binary:

$$\mathrm{DOIExact}_i = \begin{cases} 1 & \text{if normalized } \hat{d}_i = d_i, \\ 0 & \text{otherwise.} \end{cases}$$

HalluciBench reports

$$\overline{\mathrm{DOIExact}} = \frac{1}{|S_{\mathrm{DOI}}|} \sum_{i \in S_{\mathrm{DOI}}} \mathrm{DOIExact}_i,$$

where S_{DOI} is the DOI-specific scoring subset used by the implementation. We treat DOI separately because it is a high-precision bibliographic identifier: a correct DOI is strong evidence of faithful retrieval, while a fabricated DOI is a particularly severe hallucination.

Year exact match. The publication year metric is also binary:

$$\mathrm{YearExact}_i = \begin{cases} 1 & \text{if } \hat{y}_i = y_i, \\ 0 & \text{otherwise.} \end{cases}$$

HalluciBench reports

$$\overline{\mathrm{YearExact}} = \frac{1}{|S|} \sum_{i \in S} \mathrm{YearExact}_i.$$

We include year exact match because year errors are common in hallucinated citations and are not well captured by more permissive token-overlap metrics. At the same time, year is simpler and less brittle than DOI, making it a useful intermediate fidelity signal.

Correct given answered. On citation-retrieval judging, the judge returns Boolean signals such as $\mathrm{Answered}_i$ and $\mathrm{Correct}_i$. HalluciBench reports

$$\mathrm{Correct} \mid \mathrm{Answered} = \frac{\sum_{i \in J} \mathrm{Correct}_i}{\sum_{i \in J} \mathrm{Answered}_i}.$$

We use correctness conditioned on answered responses because it separates citation quality from abstention behavior. A model that refuses often should not be treated the same as a model that answers frequently but incorrectly; this metric isolates the quality of attempted answers.

Dataset-specific use of metrics. On the real citation-retrieval benchmark, HalluciBench reports the full metric suite described above. On the synthetic no-answer controls, where every gold target is `REFUSE`, the primary outcomes are strict refusal, judge refusal, and strict-or-judge refusal. No field-level citation metrics are computed on the no-answer controls because there is no true citation to recover. This ensures that refusal calibration is evaluated directly rather than through partial matching against fabricated metadata.

A.4 Prompt Templates

We document the prompt families used in HalluciBench. The benchmark uses five citation-retrieval task variants, each designed to probe model behavior under a different level of bibliographic evidence, together with two LLM-as-a-judge prompts used for semantic evaluation. The reported judge pass uses GPT-5.6 Sol; robustness passes with GPT-5.5 and Claude Sonnet 5 achieved 98% overall agreement with its Boolean labels. The retrieval prompts are instantiated directly in the benchmark dataset, while the judge prompts are applied during the evaluation stage.

Task prompt 1: Exact title to BibTeX. This prompt tests the standard citation retrieval setting in which the model is given only the exact paper title and must either return one BibTeX entry or refuse.

Exact title to BibTeX

You are a citation retrieval assistant.
Return exactly ONE BibTeX entry for the paper with this exact title.
TITLE: `<TITLE>`
Rules:

- If you are not confident the paper exists or you cannot uniquely identify it, output exactly: `REFUSE`
- Do NOT invent fields. If a field is unknown, omit it.
- Prefer including: title, author, year, and either doi or url.
- Output ONLY the BibTeX entry (no explanation, no markdown).

Task prompt 2: Exact title to BibTeX with identifier requirement. This prompt is stricter than the previous one and requires the model to abstain unless it can provide a stable identifier such as a DOI or URL.

Exact title to BibTeX with identifier

You are a citation retrieval assistant.
Return exactly ONE BibTeX entry for the paper with this exact title.
TITLE: `<TITLE>`
Rules:

- If you cannot provide a DOI or a stable URL (e.g., publisher, arXiv, ACL Anthology), output exactly: `REFUSE`
- Do NOT guess any metadata.
- Output ONLY the BibTeX entry (no explanation, no markdown).

Task prompt 3: Title plus author hint. This prompt provides the title together with a short author hint, which reduces ambiguity while still requiring bibliographic retrieval.

Title + author hint to BibTeX

You are a citation retrieval assistant.

Return exactly ONE BibTeX entry that matches BOTH the title and the author hint.
TITLE: `<TITLE>`
AUTHOR HINT: `<FIRST AUTHOR et al.>`
Rules:

- If multiple papers could match, output exactly: `REFUSE`
- If you are not confident, output exactly: `REFUSE`
- Do NOT invent fields. Omit unknown fields rather than guessing.
- Output ONLY the BibTeX entry (no explanation, no markdown).

Task prompt 4: Title plus venue and year. This variant supplies publication context in addition to the title and is intended to evaluate whether models use venue and temporal information correctly.

Title + venue + year to BibTeX

You are a citation retrieval assistant.
Return exactly ONE BibTeX entry that matches the provided title and publication context.
TITLE: `<TITLE>`
VENUE/JOURNAL: `<VENUE>`
YEAR: `<YEAR>`
Rules:

- If the venue or year does not match the record for this title, output exactly: `REFUSE`
- Do NOT guess any metadata.
- Output ONLY the BibTeX entry (no explanation, no markdown).

Task prompt 5: Partial metadata to BibTeX. This prompt provides partial structured evidence and is designed to test citation recovery under incomplete metadata.

Partial metadata to BibTeX

You are a citation retrieval assistant.
Return exactly ONE BibTeX entry that matches the partial metadata below.
TITLE: `<TITLE>`
FIRST AUTHOR: `<FIRST AUTHOR>`
YEAR: `<YEAR>`
Rules:

- If multiple papers could match, output exactly: `REFUSE`
- If you are not confident, output exactly: `REFUSE`
- Do NOT invent fields. Omit unknown fields rather than guessing.
- Prefer including: title, author, year, and either doi or url.
- Output ONLY the BibTeX entry (no explanation, no markdown).

Judge prompt 1: Citation correctness judge. This prompt is used for semantic evaluation on the real citation-retrieval benchmark. The judge receives the original benchmark prompt, the model response, and the gold bibliographic record, and returns Boolean values for `answered` and `correct`.

Citation judge prompt

You are evaluating an LLM's response to a citation retrieval task.
Original prompt given to the LLM:
`{prompt}`

LLM Response:
{response}
Ground Truth:
{ground_truth}
Use the provided Ground Truth as the only reference. Do not use outside knowledge.
Criteria:
{criteria_lines}
Respond with ONLY a JSON object, no other text.
Return exactly these keys as booleans: answered, correct.
Example:
{{json_example}}

The corresponding criteria are:

Citation judge criteria

- answered: Did the LLM attempt an answer? Set answered=false only for explicit refusal/uncertainty language (e.g. REFUSE, "I don't know", "I am not sure", "I cannot find this"). For all other outputs, set answered=true, including malformed, partial, non-BibTeX, or mixed text + BibTeX responses.
- correct: Is the BibTeX response factually correct compared to the ground truth? Evaluate whether the response matches the ground truth for title, authors, year, venue, and DOI, using your overall judgment. If answered is false, set this to false.

Judge prompt 2: Refusal judge. This prompt is used for semantic refusal evaluation, particularly on the synthetic no-answer controls. The judge decides whether the model refused the request.

Refusal judge prompt

You are evaluating whether an LLM refused to answer a citation request.
Original prompt given to the LLM:
{prompt}
LLM Response:
{response}
Criteria:
{criteria_lines}
Respond with ONLY a JSON object, no other text.
Return exactly one key with a boolean value: refused.
Valid outputs are exactly:
{"refused": true}
{"refused": false}

The corresponding criterion is:

Refusal judge criterion

- refused: Set refused=true only when the response is an explicit refusal or uncertainty (e.g. REFUSE, "I don't know", "I am not sure", "I cannot find this"), or when the response is empty/whitespace-only. Otherwise set refused=false, including malformed or non-BibTeX attempts or any attempted citation answer.

Prompt design rationale. The five task prompts are aligned across the real retrieval benchmark and the synthetic no-answer controls, which allows HalluciBench to compare faithful retrieval and refusal behavior under matched prompting conditions. The judge prompts are intentionally narrow and Boolean: the judge must use the supplied gold record as its only reference, apply explicit def-

initions for answered, correct, and refused, and return an exact JSON schema without explanation. This leaves substantially less discretion than an open-ended quality rating and is consistent with the observed 98% agreement among GPT-5.6 Sol, GPT-5.5, and Claude Sonnet 5 judgments. Example instantiated prompts for both the real and synthetic datasets appear directly in the benchmark JSON files.

A.5 Models

A.5.1 Checkpoint Release Dates and Temporal Eligibility

Table 3 reports the exact runtime-style checkpoint tags preserved in the result artifacts, their nominal parameter counts, and the public release dates used to apply the temporal safeguard. The release date is the earliest documented public availability of the corresponding model weights or size, not the date on which a local runtime package was created or updated. The eligibility cutoff is exactly two calendar years before that date; a real paper is eligible only when its publication date is on or before the listed cutoff. Thus, related checkpoints can have different cutoffs when sizes were released in stages, as for Gemma 3 and Qwen3.5. These dates remain distinct from knowledge cutoffs and do not establish that any eligible paper appeared in training. The local tags do not encode immutable weight digests, instruction/base variants, or quantization, which were not preserved in the archived execution manifest.

A.5.2 Model-Suite Rationale

The 24-checkpoint LLM suite emphasizes reproducible, locally deployable weights while spanning model families, release generations, nominal sizes, and dense versus sparse architectures. Within-family pairs and series (Llama 3.1, Gemma 2/3, Qwen 2.5/3.5, DeepSeek-R1 distillations, and Ministral 3) support descriptive comparisons across scale without implying that parameter count or post-training causally determines citation behavior (Grattafiori et al., 2024; Gemma Team et al., 2024; 2025; Qwen Team et al., 2024; Qwen Team, 2026b; Guo et al., 2025; Liu et al., 2026).

The roster also broadens deployment and architectural coverage. Command R and Command R7B represent retrieval-oriented model families evaluated here under the same closed-book restriction as every other checkpoint (Gomez, 2024a;b). Mistral 7B and Mistral NeMo provide dense Mistral baselines from different release periods, while Mixtral 8×7B adds a sparse mixture-of-experts architecture (Jiang et al., 2023; Mistral AI team, 2024; Jiang et al., 2024). Phi-4 supplies a compact reasoning-oriented comparison (Abdin et al., 2024). Together, these checkpoints cover practical local-compute regimes without claiming representativeness of proprietary frontier systems.

A.5.3 Small Language Models

The nine-checkpoint SLM suite targets low-resource and edge settings: DeepSeek-R1 1.5B, Gemma 3 270M/1B, Llama 3.2 1B/3B, Ministral 3 3B, Phi-4-mini 3.8B, and Qwen3.5 0.8B/2B. It includes matched small-size pairs within several families and both conventional instruction-tuned and reasoning-oriented checkpoints, allowing citation structure, metadata fidelity, and abstention to be studied when compute and memory are constrained (Guo et al., 2025; Gemma Team et al., 2025; Grattafiori et al., 2024; Liu et al., 2026; Microsoft et al., 2025; Qwen Team, 2026b). The runtime-style tags are retained exactly as archived; where a tag is only a nominal size label, the manuscript does not infer an unpreserved weight digest, quantization, or base-versus-instruction variant.

A.6 Case Study: Valid Structure, Incomplete Authorship

We illustrate a run of HalluciBench scoring on a single record from the citation benchmark using `mistral-nemo`. The task requires the model to return exactly one BibTeX entry for a provided title, or abstain by outputting `REFUSE`. The ground-truth metadata corresponds to the NeurIPS 2017 paper "Attention is All you Need" Vaswani et al. (2017) and includes canonical title, author list, venue, year, and a stable URL; the ground truth DOI is empty.

Table 3: Model release dates used for temporal eligibility. Every row corresponds to one of the 33 evaluated checkpoint tags. "Parameters" gives the nominal size associated with the tag; for Mixtral and Qwen3.5-35B, total and active parameters are distinguished. The final column is the inclusive latest permissible publication date, computed as the public release date minus two calendar years. Dates are descriptive rather than performance metrics.

Exact checkpoint tag	Suite	Parameters	Public release	Source for release date	Two-year cutoff
`command-r:35b`	LLM	35B	2024-03-11	Cohere launch post (Gomez, 2024a)	2022-03-11
`command-r7b:7b`	LLM	7B	2024-12-13	Cohere launch post (Gomez, 2024b)	2022-12-13
`deepseek-r1:32b`	LLM	32B	2025-01-20	Official release repository (DeepSeek AI, 2025)	2023-01-20
`deepseek-r1:14b`	LLM	14B	2025-01-20	Official release repository (DeepSeek AI, 2025)	2023-01-20
`deepseek-r1:8b`	LLM	8B	2025-01-20	Official release repository (DeepSeek AI, 2025)	2023-01-20
`gemma2:27b`	LLM	27B	2024-06-27	Google release history (Google AI for Developers, 2026)	2022-06-27
`gemma2:9b`	LLM	9B	2024-06-27	Google release history (Google AI for Developers, 2026)	2022-06-27
`gemma3:27b`	LLM	27B	2025-03-10	Google release history (Google AI for Developers, 2026)	2023-03-10
`gemma3:12b`	LLM	12B	2025-03-10	Google release history (Google AI for Developers, 2026)	2023-03-10
`gemma3:4b`	LLM	4B	2025-03-10	Google release history (Google AI for Developers, 2026)	2023-03-10
`llama3.1:70b`	LLM	70B	2024-07-23	Meta launch post (Meta AI, 2024a)	2022-07-23
`llama3.1:8b`	LLM	8B	2024-07-23	Meta launch post (Meta AI, 2024a)	2022-07-23
`ministral-3:14b`	LLM	14B	2025-12-02	Mistral launch post (Mistral AI, 2025)	2023-12-02
`ministral-3:8b`	LLM	8B	2025-12-02	Mistral launch post (Mistral AI, 2025)	2023-12-02
`mistral:7b`	LLM	7.3B	2023-09-27	Mistral launch source (Jiang et al., 2023)	2021-09-27
`mistral-nemo`	LLM	12B	2024-07-18	Mistral launch post (Mistral AI team, 2024)	2022-07-18
`mixtral:8x7b`	LLM	46.7B total / 12.9B active	2023-12-11	Mistral launch source (Jiang et al., 2024)	2021-12-11
`phi4:14b`	LLM	14B	2024-12-12	Microsoft launch post (Microsoft, 2024b)	2022-12-12
`qwen2.5:14b`	LLM	14B	2024-09-19	Qwen launch post (Qwen Team, 2024)	2022-09-19
`qwen2.5:7b`	LLM	7B	2024-09-19	Qwen launch post (Qwen Team, 2024)	2022-09-19
`qwen3.5:35b`	LLM	35B total / 3B active	2026-02-24	Official staged-release log (Qwen Team, 2026a)	2024-02-24
`qwen3.5:27b`	LLM	27B	2026-02-24	Official staged-release log (Qwen Team, 2026a)	2024-02-24
`qwen3.5:9b`	LLM	9B	2026-03-02	Official staged-release log (Qwen Team, 2026a)	2024-03-02
`qwen3.5:4b`	LLM	4B	2026-03-02	Official staged-release log (Qwen Team, 2026a)	2024-03-02
`deepseek-r1:1.5b`	SLM	1.5B	2025-01-20	Official release repository (DeepSeek AI, 2025)	2023-01-20
`gemma3:1b`	SLM	1B	2025-03-10	Google release history (Google AI for Developers, 2026)	2023-03-10
`gemma3:270m`	SLM	270M	2025-08-14	Google release history (Google AI for Developers, 2026)	2023-08-14
`llama3.2:3b`	SLM	3B	2024-09-25	Meta launch post (Meta AI, 2024b)	2022-09-25
`llama3.2:1b`	SLM	1B	2024-09-25	Meta launch post (Meta AI, 2024b)	2022-09-25
`ministral-3:3b`	SLM	3B	2025-12-02	Mistral launch post (Mistral AI, 2025)	2023-12-02
`phi4-mini:3.8b`	SLM	3.8B	2025-02-26	Microsoft launch post (Microsoft, 2025b)	2023-02-26
`qwen3.5:2b`	SLM	2B	2026-03-02	Official staged-release log (Qwen Team, 2026a)	2024-03-02
`qwen3.5:0.8b`	SLM	0.8B	2026-03-02	Official staged-release log (Qwen Team, 2026a)	2024-03-02

Dataset Record (ID = 1)

Task type: `exact_title_to_bibtex`
Title: Attention is All you Need
Venue: Advances in Neural Information Processing Systems
Year: 2017
Authors (gold): 8 authors (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, Polosukhin)
URL (gold): Semantic Scholar paper URL

Prompt (verbatim)

You are a citation retrieval assistant.
Return exactly ONE BibTeX entry for the paper with this exact title.
TITLE: Attention is All you Need

Rules:
- If you are not confident the paper exists or you cannot uniquely identify it, output exactly: REFUSE
- Do NOT invent fields. If a field is unknown, omit it.
- Prefer including: title, author, year, and either doi or url.
- Output ONLY the BibTeX entry (no explanation, no markdown).

Ground-truth BibTeX

```
@inproceedings{vaswani2017attention,
  title = {Attention is All you Need},
  author = {Vaswani, Ashish and Shazeer, Noam and Parmar, Niki and
    Uszkoreit, Jakob and Jones, Llion and Gomez, Aidan N. and
    Kaiser, Lukasz and Polosukhin, I.},
  booktitle = {Advances in Neural Information Processing Systems},
  year = {2017},
  url = {https://www.semanticscholar.org/paper/
    204e3073870fae3d05bcbc2f6a8e263d9b72e776},
  note = {Publication date: 2017-06-12}
}
```

Model Response

```
@article{vaswani2017attention,
  title={Attention Is All You Need},
  author={Vaswani, Ashish and et al.},
  journal={Advances in Neural Information Processing Systems},
  year={2017},
  doi={10.48550/arXiv.1706.03762}
}
```

Scores (selected)

BibTeX parseable: `true`
Refusal detected: `false`
Title: unigram F1 = 1.00, bigram F1 = 1.00
Venue: F1 = 1.00
Year: exact match = 1
Authors: precision = 1.00, recall = 0.125, F1 = 0.222
DOI exact match (conditional): `null` (no gold DOI provided)

Interpretation. This is a favorable title-only example: the paper is prominent, predates the checkpoint by several years, and is uniquely identified by the exact title. At the structural level, the response satisfies the one-entry contract and is therefore parseable. The predicted entry type is `@article` and the venue is placed in `journal`, whereas the canonical record uses `@inproceedings` and `booktitle`. Those choices do not prevent parsing, and the normalized venue text still matches exactly. Capitalization differences in the supplied title also disappear under normalization, yielding unigram and bigram title $F_1 = 1$, while the year receives exact match $= 1$. These successes show correct preservation of visible and familiar fields, but they do not establish recovery of the complete record.

The author field exposes the substantive failure. The response retains only Ashish Vaswani and replaces the remaining seven authors with "et al." Although this abbreviation is conventional in prose, it is not a complete BibTeX author list and cannot support full attribution or reliable author indexing. Under the reported one-to-one author matching, the single recovered author is correct, so precision is 1; recall is only $1/8 = 0.125$, producing author $F_1 = 0.222$. Precision alone would therefore make the prediction appear perfect despite the loss of 87.5% of the authors, motivating the separate reporting of precision, recall, and F_1.

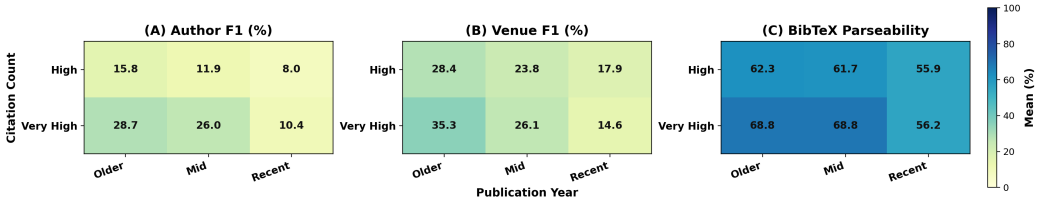

Figure 3: Citation-benchmark performance by paper characteristics. Rows denote citation-count buckets and columns denote publication-year buckets. Cell values are unweighted means across the checkpoint subsets eligible for each bucket; higher is better.

Identifier handling introduces a different uncertainty. The gold record contains a stable URL but no valid DOI, while the model supplies a DOI-shaped identifier associated with an arXiv record and omits the gold URL. DOI exact match is consequently `null`, rather than zero or one, because no gold DOI is available for deterministic comparison; HalluciBench neither rewards nor labels the supplied identifier as an exact-DOI error for this sample. Taken together, the response would look successful under parseability, title, venue, and year alone, yet it remains unsuitable as an authoritative citation because authorship is severely incomplete and the identifier cannot be verified under the available gold metadata. The example demonstrates why format validity, field-level fidelity, coverage, and missing-gold policies must be reported together.

A.7 Supplementary Results

A.7.1 Performance by Paper Characteristics

Figure 3 is descriptive rather than causal. Older and very-high-citation papers tend to have higher author fidelity, while recent buckets tend to have lower author/venue F_1 and parseability. One plausible explanation is exposure: older, prominent papers have accumulated more webpages, bibliographies, discussions, and repeated canonical citations from which a model could encounter their metadata. Their records may also be more standardized after years of use. HalluciBench cannot verify training inclusion or exposure frequency, however, and prominence may correlate with other properties such as author-list length, venue, topic, or the number of preprint and proceedings versions. The per-model two-year rule creates an additional compositional effect: later publication buckets are evaluated only by later-released eligible checkpoints, so model generation and paper age are entangled, and the recent buckets contain fewer eligible papers. Consequently, the figure cannot show that age or citation count causes better recall. The supported conclusion is narrower but still important: even within a prominence-biased pool deliberately designed to favor recall, metadata reconstruction remains far from saturated across paper groups.

A.8 Task Specific Results

Tables 4–8 and Tables 9–13 provide a more granular view of how model behavior changes across input conditions. In particular, they make it possible to separate improvements that come from copying metadata supplied by the task variant from improvements that require actual bibliographic recall. They also clarify which fields remain difficult even when partial metadata is provided, and which task variants mainly affect willingness to answer rather than citation fidelity itself.

A.9 Field-level Results

Tables 19 and 20 reveal several qualitative patterns. Author and venue errors have different shapes: for many models, author precision exceeds author recall, which is more consistent with incomplete author recovery than with systematic over-generation, whereas venue recall exceeds venue precision throughout the table, suggesting that venue errors more often involve partial or non-canonical matches than complete failures. The tables also show substantial per-example heterogeneity. Standard deviations are large for both author and venue metrics, indicating that average scores hide sharp variation across papers. Finally, models exhibit field-specialized profiles. `llama3.1:70b`

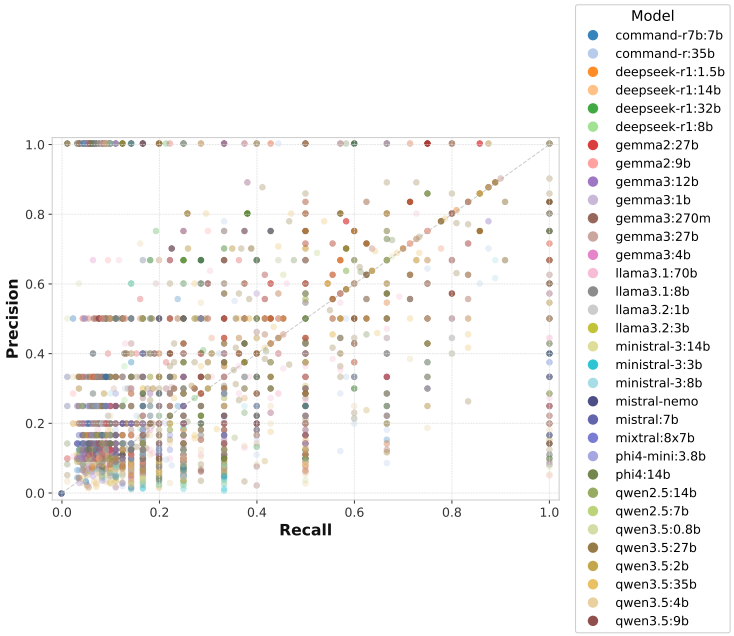

Figure 4: Per-sample author precision versus recall on the citation benchmark. Each point is one evaluated example; color denotes the generating model checkpoint. Points with failed inference are excluded.

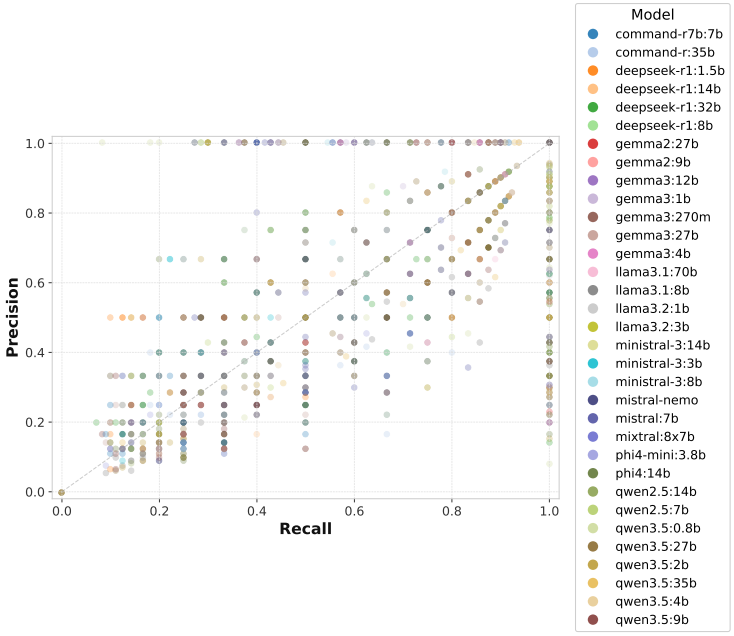

Figure 5: Per-sample title precision versus recall using unigram token overlap between predicted and gold titles. Each point is one evaluated example; color denotes the generating model checkpoint. Points with failed inference are excluded.

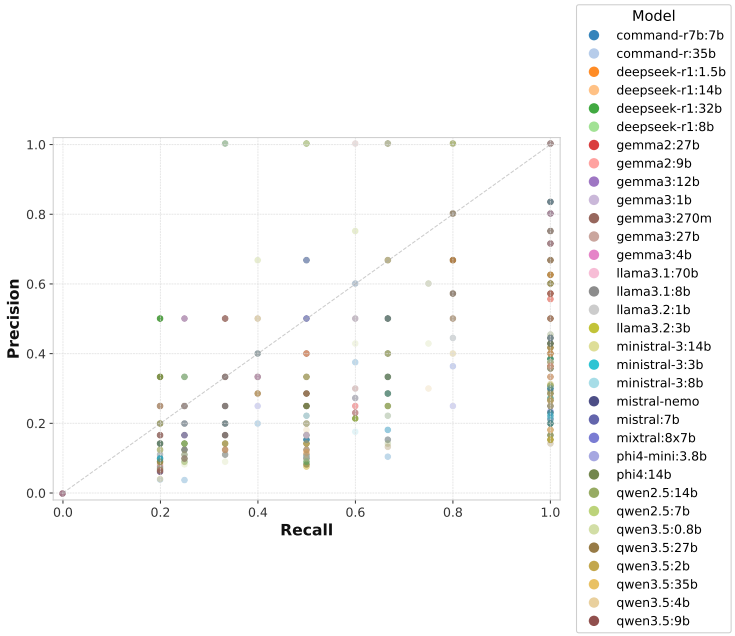

Figure 6: Per-sample venue precision versus recall on the citation benchmark. Each point is one evaluated example; color denotes the generating model checkpoint. Points with failed inference are excluded.

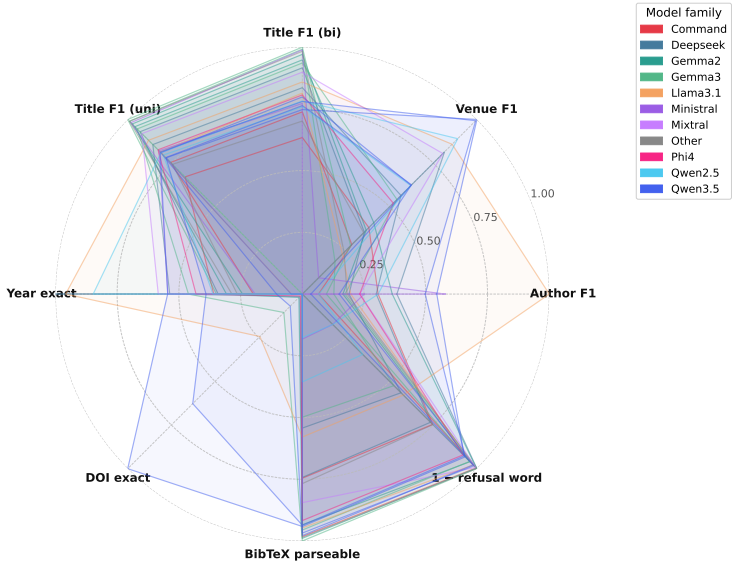

Figure 7: Radar comparison of citation-quality signals across models evaluated on HalluciBench. Each spoke is min–max normalized across the plotted models so that, for that metric, the worst observed value maps to 0 and the best to 1 (independent normalization per axis). Axes show mean token-level F1 for authors, venue, and title (bigram and unigram), plus rates for year match, DOI match (when evaluable), BibTeX parseability, and one minus the refusal-word indicator (higher is better). Only models with judge annotations in the released datapoints are included. Lines are colored by model family.

Table 4: Citation benchmark overall results for task variant: Title only. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
command-r:35b	88.6	11.3	11.3	96.5	10.1	24.5	49.1	0.0	1.3
command-r7b:7b	98.3	0.2	0.2	96.0	1.2	5.7	29.0	0.0	0.0
deepseek-r1:32b	75.8	25.1	25.1	97.7	9.3	13.9	45.6	0.0	0.9
deepseek-r1:14b	54.3	45.4	45.5	96.7	3.9	7.9	34.0	0.0	0.1
deepseek-r1:8b	0.8	0.1	0.1	0.8	0.1	0.4	0.6	0.0	0.1
gemma2:27b	99.2	0.3	0.3	98.9	12.3	2.0	39.8	0.2	1.4
gemma2:9b	80.8	12.4	12.4	91.7	3.5	3.0	28.7	0.0	0.1
gemma3:27b	**100.0**	**0.0**	**0.0**	99.5	6.9	0.8	40.5	1.3	0.1
gemma3:12b	99.3	0.1	0.1	99.3	2.0	0.9	32.2	0.0	0.2
gemma3:4b	85.6	10.9	10.9	95.9	0.3	0.5	14.1	0.0	0.0
llama3.1:70b	44.4	55.2	55.2	98.7	**50.5**	23.6	69.4	1.8	**3.1**
llama3.1:8b	95.9	0.7	0.7	95.9	4.9	1.9	29.8	0.0	0.3
ministral-3:14b	2.8	97.2	97.2	**100.0**	28.4	13.8	53.6	0.0	0.2
ministral-3:8b	0.6	99.4	99.4	**100.0**	16.7	0.0	**100.0**	0.0	0.0
mistral:7b	99.2	0.1	0.1	98.1	1.2	0.6	30.5	0.0	0.1
mistral-nemo	98.0	**0.0**	**0.0**	97.8	4.0	2.3	34.9	0.0	0.0
mixtral:8x7b	87.7	0.1	0.1	86.9	8.9	28.3	43.7	0.0	0.5
phi4:14b	80.3	19.3	19.3	98.2	8.3	16.7	38.8	0.0	0.5
qwen2.5:14b	94.8	4.8	4.8	98.3	9.4	3.1	46.8	0.0	0.6
qwen2.5:7b	20.0	79.8	79.8	96.1	11.3	25.1	52.0	0.0	0.1
qwen3.5:35b	94.4	5.7	5.7	98.0	17.5	40.5	47.5	**6.0**	1.9
qwen3.5:27b	96.1	4.0	4.0	97.0	19.6	**43.8**	42.4	2.6	2.6
qwen3.5:9b	97.6	2.3	2.3	97.9	4.6	23.4	28.8	0.7	0.2
qwen3.5:4b	97.0	0.7	0.8	95.6	1.6	17.5	23.6	0.0	0.2

Table 5: Citation benchmark overall results for task variant: Title only + DOI/URL requirement. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
command-r:35b	88.2	11.4	12.4	94.6	9.8	30.7	50.1	0.0	0.9
command-r7b:7b	99.0	0.1	0.6	95.7	1.3	16.0	30.4	0.0	0.3
deepseek-r1:32b	68.8	30.7	31.2	96.0	11.2	26.0	45.8	0.3	1.0
deepseek-r1:14b	46.9	52.0	52.1	95.5	3.6	12.5	36.9	0.0	0.8
deepseek-r1:8b	0.9	0.2	**0.4**	0.7	0.2	0.4	0.6	0.0	0.0
gemma2:27b	98.2	1.4	2.0	98.5	11.8	23.9	38.9	0.0	1.9
gemma2:9b	92.6	4.2	4.6	95.9	3.1	20.1	30.8	0.0	0.3
gemma3:27b	99.0	1.0	1.6	**99.5**	6.4	38.8	40.0	0.0	0.5
gemma3:12b	79.2	20.4	20.8	99.2	2.6	10.8	35.8	0.0	0.1
gemma3:4b	0.0	100.0	100.0	–	–	–	–	–	0.6
llama3.1:70b	62.8	37.0	37.9	98.0	**45.6**	44.1	**66.9**	1.6	**4.3**
llama3.1:8b	89.8	8.4	9.4	96.2	6.2	13.9	26.1	0.0	0.0
ministral-3:14b	2.5	97.5	97.5	99.4	35.1	14.6	64.0	0.0	1.0
ministral-3:8b	44.0	57.4	57.5	95.8	5.8	4.0	52.3	0.0	0.5
mistral:7b	95.4	1.7	2.5	94.0	1.2	15.7	30.7	0.0	0.0
mistral-nemo	**99.5**	**0.0**	0.5	99.3	3.6	11.7	31.8	0.0	0.0
mixtral:8x7b	86.3	1.8	2.8	86.8	9.3	34.6	43.8	0.0	0.6
phi4:14b	81.7	17.2	18.3	96.2	8.3	20.8	37.7	0.2	0.6
qwen2.5:14b	20.7	79.1	79.2	97.5	17.6	36.3	59.7	0.0	0.9
qwen2.5:7b	6.6	93.4	93.5	95.9	15.3	44.4	62.1	0.0	0.8
qwen3.5:35b	88.8	12.4	13.0	97.5	18.4	47.4	47.5	**6.5**	1.5
qwen3.5:27b	97.7	2.8	4.0	96.5	18.7	**47.7**	42.6	3.9	3.2
qwen3.5:9b	95.7	4.3	4.8	97.6	4.7	26.3	27.9	0.0	0.3
qwen3.5:4b	93.3	4.5	5.3	95.6	1.7	24.6	21.7	0.0	0.3

Table 6: Citation benchmark overall results for task variant: Title + author hint. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
command-r:35b	99.3	0.4	0.7	97.8	32.5	27.0	59.4	0.0	3.2
command-r7b:7b	97.2	1.5	1.9	97.7	26.0	15.6	45.7	0.0	0.8
deepseek-r1:32b	83.8	17.4	17.6	98.2	35.6	37.9	57.1	0.0	3.8
deepseek-r1:14b	62.4	37.0	36.9	96.5	29.9	29.1	45.1	0.0	1.8
deepseek-r1:8b	0.6	**0.0**	**0.0**	0.6	0.3	0.3	0.4	0.0	0.0
gemma2:27b	98.5	0.2	0.4	98.5	39.1	24.4	40.7	0.0	6.2
gemma2:9b	99.5	**0.0**	0.1	99.3	29.5	21.0	38.5	0.0	1.7
gemma3:27b	**99.9**	**0.0**	0.3	99.5	28.8	35.3	56.4	0.0	3.7
gemma3:12b	99.8	0.1	0.3	99.7	25.9	12.1	47.9	0.0	0.8
gemma3:4b	53.7	45.3	45.5	97.8	26.7	3.9	25.8	0.0	0.3
llama3.1:70b	52.4	47.6	47.8	99.3	56.6	40.3	**68.9**	0.0	7.1
llama3.1:8b	97.8	0.1	0.3	96.6	30.1	8.5	30.2	0.0	1.9
ministral-3:14b	0.0	100.0	100.0	–	–	–	–	–	0.0
ministral-3:8b	23.7	75.3	75.3	94.0	34.6	8.3	64.8	0.0	0.0
mistral:7b	23.2	88.4	88.5	90.3	27.4	8.6	51.5	0.0	0.1
mistral-nemo	97.2	**0.0**	0.1	97.0	30.8	11.1	44.2	0.0	0.2
mixtral:8x7b	93.2	0.2	0.8	92.4	26.7	39.1	59.2	0.0	2.7
phi4:14b	99.2	0.2	0.3	98.6	31.5	22.9	46.7	0.0	2.3
qwen2.5:14b	0.6	99.4	99.4	**100.0**	**66.7**	33.3	66.7	0.0	0.1
qwen2.5:7b	24.0	75.9	76.0	98.5	36.4	43.9	65.6	0.0	1.4
qwen3.5:35b	98.5	1.2	1.2	99.1	44.7	54.3	64.6	**0.4**	**12.8**
qwen3.5:27b	98.2	1.7	1.8	99.3	47.5	**60.7**	63.4	0.2	11.3
qwen3.5:9b	99.1	0.6	0.7	98.5	29.2	29.9	48.0	0.0	2.9
qwen3.5:4b	97.6	1.3	1.9	96.4	22.2	30.9	36.0	0.0	1.6

Table 7: Citation benchmark overall results for task variant: Title + venue + year. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
command-r:35b	13.4	86.5	86.5	92.6	16.6	90.1	94.8	0.0	0.6
command-r7b:7b	96.6	0.8	0.8	93.8	1.1	77.4	93.2	0.0	0.7
deepseek-r1:32b	68.8	32.1	32.1	98.1	9.8	87.6	98.5	0.0	3.4
deepseek-r1:14b	48.9	42.8	42.8	84.1	2.9	72.0	84.2	0.0	0.9
deepseek-r1:8b	0.6	0.6	0.6	0.5	0.1	0.5	0.6	0.0	0.0
gemma2:27b	99.0	0.2	0.5	98.7	13.0	92.5	97.5	0.0	5.7
gemma2:9b	95.2	3.3	3.4	98.0	3.4	89.2	97.8	0.0	1.1
gemma3:27b	**100.0**	0.4	0.5	98.9	7.1	92.6	98.8	0.2	1.8
gemma3:12b	99.7	**0.0**	**0.1**	99.6	2.5	**94.5**	99.3	0.0	0.4
gemma3:4b	50.6	40.6	40.6	85.0	0.1	80.4	84.7	0.0	1.2
llama3.1:70b	76.0	24.2	24.2	95.6	41.7	87.5	91.3	0.3	**18.7**
llama3.1:8b	90.4	2.3	2.3	87.4	5.1	83.2	84.6	0.0	1.5
ministral-3:14b	0.2	99.8	99.8	**100.0**	**57.7**	83.3	**100.0**	0.0	0.2
ministral-3:8b	10.9	91.5	91.5	97.3	10.4	65.7	94.5	0.0	0.4
mistral:7b	92.5	6.0	6.1	95.5	1.5	88.3	96.8	0.0	0.1
mistral-nemo	96.9	1.2	1.4	97.7	3.5	88.6	97.7	0.0	0.4
mixtral:8x7b	68.7	2.2	2.7	69.5	6.9	64.1	68.2	0.0	2.9
phi4:14b	97.2	3.2	3.3	95.2	6.3	86.1	96.8	0.0	2.8
qwen2.5:14b	63.8	40.6	40.8	94.6	10.5	91.9	89.7	0.0	2.9
qwen2.5:7b	27.3	72.6	72.6	97.6	8.4	93.9	97.8	0.0	1.4
qwen3.5:35b	89.3	13.7	14.0	97.3	19.6	87.3	91.2	**0.7**	6.8
qwen3.5:27b	77.3	26.4	26.7	97.4	20.8	91.2	86.8	0.2	5.7
qwen3.5:9b	96.9	3.5	4.0	96.0	5.2	88.0	87.2	0.0	1.1
qwen3.5:4b	96.9	2.1	2.5	96.4	1.5	89.8	95.3	0.0	0.2

Table 8: Citation benchmark overall results for task variant: Title + first author + year. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
command-r:35b	83.5	16.4	16.6	99.5	32.0	24.1	96.3	0.0	0.9
command-r7b:7b	98.8	0.4	0.5	98.9	25.1	7.7	98.2	0.0	0.2
deepseek-r1:32b	74.2	26.8	26.9	98.1	34.1	20.8	98.0	0.0	2.6
deepseek-r1:14b	60.0	38.3	38.4	95.3	31.3	8.2	95.2	0.0	0.5
deepseek-r1:8b	0.9	**0.0**	**0.0**	0.9	0.3	0.2	0.9	0.0	0.0
gemma2:27b	98.6	0.3	0.3	98.8	40.3	6.9	95.6	0.2	5.2
gemma2:9b	98.3	0.7	0.7	99.0	32.3	3.1	99.0	0.0	1.5
gemma3:27b	**100.0**	**0.0**	0.1	99.8	28.3	12.0	**100.0**	0.5	1.3
gemma3:12b	98.9	0.9	1.1	99.6	26.9	0.8	99.0	0.0	0.1
gemma3:4b	60.8	38.8	38.9	99.2	31.7	2.4	99.0	0.0	0.1
llama3.1:70b	54.8	45.1	45.1	99.3	**57.7**	29.1	96.5	1.1	6.7
llama3.1:8b	98.7	**0.0**	0.1	97.2	32.4	0.2	96.8	0.0	1.2
ministral-3:14b	0.0	100.0	100.0	–	–	–	–	–	0.0
ministral-3:8b	15.0	85.1	85.1	98.7	35.0	6.2	98.0	0.0	0.0
mistral:7b	73.5	29.4	29.4	97.8	29.0	6.5	97.9	0.0	0.3
mistral-nemo	99.8	**0.0**	0.1	99.8	36.2	0.1	99.8	0.0	0.4
mixtral:8x7b	86.1	0.1	0.2	86.0	24.7	34.7	84.5	0.0	1.8
phi4:14b	99.5	**0.0**	0.2	99.4	31.5	22.1	99.5	0.0	1.4
qwen2.5:14b	0.1	99.9	99.9	**100.0**	55.6	**66.7**	**100.0**	–	0.0
qwen2.5:7b	13.2	86.8	86.8	99.6	37.0	13.6	**100.0**	0.0	0.3
qwen3.5:35b	98.7	1.3	1.4	99.2	42.5	48.2	95.0	**5.3**	7.3
qwen3.5:27b	97.9	2.1	2.3	99.5	46.1	41.4	95.1	4.4	**8.5**
qwen3.5:9b	98.7	0.9	1.1	99.4	28.2	28.7	94.5	0.6	1.5
qwen3.5:4b	98.3	0.5	1.2	96.4	20.7	30.3	96.9	0.0	0.9

Table 9: Citation benchmark author and venue metrics by task variant. Task variant: Title only. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
command-r:35b	10.1 (22.1)	10.2 (22.9)	10.6 (22.9)	24.5 (36.9)	22.4 (34.3)	28.2 (41.9)
command-r7b:7b	1.2 (7.1)	1.2 (7.3)	1.3 (7.6)	5.7 (20.5)	5.2 (19.2)	6.7 (23.5)
deepseek-r1:32b	9.3 (22.3)	9.8 (23.6)	9.5 (22.7)	13.9 (29.6)	12.1 (26.3)	17.2 (36.0)
deepseek-r1:14b	3.9 (12.6)	4.6 (15.1)	3.9 (12.9)	7.9 (22.8)	6.9 (20.1)	9.9 (27.9)
deepseek-r1:8b	0.1 (2.3)	0.1 (2.4)	0.1 (2.2)	0.4 (5.7)	0.4 (5.3)	0.4 (6.5)
gemma2:27b	12.3 (26.1)	12.8 (26.9)	12.2 (26.2)	2.0 (13.1)	1.9 (12.4)	2.1 (14.2)
gemma2:9b	3.5 (13.2)	3.7 (13.8)	3.6 (13.8)	3.0 (14.4)	2.5 (12.3)	4.0 (18.6)
gemma3:27b	6.9 (18.1)	7.2 (18.8)	7.1 (18.8)	0.8 (7.8)	0.7 (7.0)	0.9 (8.8)
gemma3:12b	2.0 (9.5)	2.0 (9.9)	2.1 (10.2)	0.9 (8.9)	0.9 (8.7)	1.0 (9.5)
gemma3:4b	0.3 (3.8)	0.4 (4.5)	0.2 (3.8)	0.5 (5.8)	0.5 (5.7)	0.6 (6.4)
llama3.1:70b	**50.5 (34.4)**	**50.4 (34.7)**	**51.3 (34.3)**	23.6 (38.3)	21.5 (35.6)	26.9 (43.0)
llama3.1:8b	4.9 (15.4)	5.2 (16.4)	5.0 (15.8)	1.9 (11.9)	1.7 (10.7)	2.3 (14.1)
ministral-3:14b	28.4 (30.1)	29.4 (30.4)	28.3 (30.6)	13.8 (30.8)	12.2 (27.8)	16.1 (35.5)
ministral-3:8b	16.7 (25.8)	15.5 (24.1)	18.3 (28.6)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
mistral:7b	1.2 (7.3)	1.2 (7.8)	1.3 (8.0)	0.6 (6.8)	0.5 (6.2)	0.8 (8.2)
mistral-nemo	4.0 (14.9)	4.9 (17.2)	3.8 (14.5)	2.3 (11.6)	1.9 (10.3)	2.8 (14.2)
mixtral:8x7b	8.9 (20.0)	9.0 (20.2)	9.5 (21.3)	28.3 (38.0)	25.6 (35.2)	33.1 (43.7)
phi4:14b	8.3 (21.2)	8.9 (22.7)	8.4 (21.6)	16.7 (31.1)	14.6 (27.9)	20.4 (37.5)
qwen2.5:14b	9.4 (22.0)	9.6 (22.3)	9.8 (22.8)	3.1 (14.6)	2.7 (13.1)	4.1 (18.4)
qwen2.5:7b	11.3 (22.8)	12.2 (24.5)	11.2 (22.6)	25.1 (39.0)	23.2 (36.8)	28.4 (43.6)
qwen3.5:35b	17.5 (29.6)	17.8 (30.0)	17.8 (29.9)	40.5 (42.2)	36.4 (39.4)	47.4 (48.1)
qwen3.5:27b	19.6 (30.9)	20.1 (31.4)	19.9 (31.3)	**43.8 (41.4)**	**39.2 (38.6)**	**51.7 (47.8)**
qwen3.5:9b	4.6 (15.3)	4.8 (16.3)	5.0 (16.4)	23.4 (36.6)	21.3 (34.0)	27.0 (41.6)
qwen3.5:4b	1.6 (8.3)	1.6 (8.4)	1.8 (9.4)	17.5 (32.3)	15.5 (29.5)	21.2 (38.3)

Table 10: Citation benchmark author and venue metrics by task variant. Task variant: Title only + DOI/URL requirement. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
command-r:35b	9.8 (21.3)	9.8 (21.7)	10.4 (22.1)	30.7 (39.3)	28.1 (36.9)	35.3 (44.5)
command-r7b:7b	1.3 (7.5)	1.3 (7.9)	1.4 (8.6)	16.0 (31.0)	14.5 (29.0)	19.0 (36.0)
deepseek-r1:32b	11.2 (24.9)	11.7 (25.5)	11.2 (25.0)	26.0 (35.6)	22.1 (31.4)	33.0 (44.1)
deepseek-r1:14b	3.6 (12.2)	4.0 (13.1)	3.8 (13.2)	12.5 (28.0)	10.7 (24.6)	15.7 (34.7)
deepseek-r1:8b	0.2 (2.4)	0.2 (2.3)	0.2 (2.8)	0.4 (6.4)	0.4 (6.0)	0.5 (7.1)
gemma2:27b	11.8 (25.8)	12.4 (26.8)	11.7 (25.7)	23.9 (38.5)	22.0 (36.3)	26.8 (42.7)
gemma2:9b	3.1 (12.5)	3.3 (12.9)	3.3 (13.1)	20.1 (31.5)	16.6 (26.8)	26.6 (41.0)
gemma3:27b	6.4 (17.1)	6.9 (18.1)	6.5 (17.4)	38.8 (42.1)	35.6 (39.8)	44.0 (46.8)
gemma3:12b	2.6 (11.2)	2.8 (12.2)	2.6 (11.4)	10.8 (27.5)	10.1 (26.3)	12.0 (29.6)
gemma3:4b	–	–	–	–	–	–
llama3.1:70b	**45.6 (34.2)**	**45.7 (34.6)**	**46.6 (34.5)**	44.1 (43.5)	40.4 (41.0)	50.1 (48.5)
llama3.1:8b	6.2 (18.4)	6.4 (18.9)	6.2 (18.6)	13.9 (28.7)	12.3 (26.0)	17.0 (34.3)
ministral-3:14b	35.1 (40.9)	35.3 (40.8)	35.0 (41.0)	14.6 (32.5)	13.1 (29.8)	17.0 (37.3)
ministral-3:8b	5.8 (16.2)	5.9 (16.3)	6.1 (17.0)	4.0 (14.0)	3.4 (12.6)	5.5 (19.2)
mistral:7b	1.2 (7.1)	1.3 (7.4)	1.3 (7.6)	15.7 (31.2)	13.9 (28.5)	18.9 (36.6)
mistral-nemo	3.6 (13.7)	4.6 (16.7)	3.3 (13.0)	11.7 (27.3)	10.3 (24.7)	14.1 (32.6)
mixtral:8x7b	9.3 (20.0)	9.3 (20.0)	10.0 (21.5)	34.6 (39.5)	31.3 (36.6)	40.4 (45.4)
phi4:14b	8.3 (21.1)	8.6 (21.9)	8.5 (21.6)	20.8 (32.9)	18.1 (29.2)	25.8 (40.4)
qwen2.5:14b	17.6 (26.0)	18.6 (27.7)	17.7 (26.1)	36.3 (40.3)	32.1 (37.3)	44.2 (47.4)
qwen2.5:7b	15.3 (27.5)	16.5 (29.5)	15.3 (27.8)	44.4 (44.1)	42.2 (42.9)	48.2 (47.2)
qwen3.5:35b	18.4 (29.9)	18.5 (30.0)	18.8 (30.4)	47.4 (40.8)	41.9 (38.0)	**57.4 (47.4)**
qwen3.5:27b	18.7 (30.3)	19.4 (31.1)	18.7 (30.6)	**47.7 (41.8)**	**43.0 (39.6)**	55.7 (47.4)
qwen3.5:9b	4.7 (14.6)	5.0 (15.9)	4.9 (15.4)	26.3 (37.3)	23.6 (34.4)	31.2 (43.4)
qwen3.5:4b	1.7 (8.7)	1.7 (8.6)	1.8 (9.4)	24.6 (36.2)	22.0 (33.3)	29.4 (42.3)

Table 11: Citation benchmark author and venue metrics by task variant. Task variant: Title + author hint. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
command-r:35b	32.5 (25.8)	32.5 (26.5)	35.2 (27.5)	27.0 (38.4)	24.9 (36.2)	30.8 (43.2)
command-r7b:7b	26.0 (23.7)	26.9 (24.7)	27.6 (24.6)	15.6 (31.8)	14.2 (29.5)	18.1 (36.3)
deepseek-r1:32b	35.6 (26.8)	45.0 (33.4)	34.9 (28.3)	37.9 (37.9)	32.5 (33.9)	47.7 (46.2)
deepseek-r1:14b	29.9 (25.6)	43.0 (36.8)	27.9 (26.3)	29.1 (37.0)	25.0 (32.8)	36.4 (45.2)
deepseek-r1:8b	0.3 (4.3)	0.2 (4.3)	0.3 (4.4)	0.3 (4.8)	0.2 (4.2)	0.4 (5.7)
gemma2:27b	39.1 (29.1)	42.9 (31.0)	38.6 (29.4)	24.4 (39.7)	22.7 (37.5)	27.2 (43.6)
gemma2:9b	29.5 (24.9)	34.3 (29.7)	29.7 (25.5)	21.0 (32.6)	17.7 (28.5)	27.2 (41.2)
gemma3:27b	28.8 (24.3)	29.9 (24.7)	29.5 (25.0)	35.3 (42.1)	32.7 (40.0)	39.5 (46.4)
gemma3:12b	25.9 (23.4)	26.5 (24.5)	27.7 (23.9)	12.1 (29.1)	11.4 (28.0)	13.2 (31.1)
gemma3:4b	26.7 (19.7)	43.6 (36.2)	24.2 (19.8)	3.9 (15.7)	3.6 (14.8)	4.6 (17.9)
llama3.1:70b	56.6 (30.7)	56.5 (31.3)	58.3 (30.6)	40.3 (42.8)	36.7 (40.2)	46.1 (48.0)
llama3.1:8b	30.1 (25.9)	35.6 (30.5)	29.7 (26.5)	8.5 (24.0)	7.6 (22.2)	9.9 (27.2)
ministral-3:14b	–	–	–	–	–	–
ministral-3:8b	34.6 (30.6)	37.5 (33.1)	36.4 (31.8)	8.3 (20.8)	7.6 (19.2)	9.5 (24.3)
mistral:7b	27.4 (26.6)	32.6 (33.0)	28.2 (28.2)	8.6 (22.8)	7.4 (20.0)	10.8 (27.6)
mistral-nemo	30.8 (24.9)	43.5 (33.8)	27.9 (25.4)	11.1 (26.2)	9.8 (23.8)	13.3 (31.0)
mixtral:8x7b	26.7 (25.1)	27.3 (25.5)	28.1 (26.1)	39.1 (40.1)	35.7 (37.6)	45.1 (45.5)
phi4:14b	31.5 (24.9)	39.5 (32.1)	31.2 (26.1)	22.9 (34.3)	19.9 (30.6)	28.1 (41.7)
qwen2.5:14b	**66.7 (51.6)**	**66.7 (51.6)**	**66.7 (51.6)**	33.3 (51.6)	33.3 (51.6)	33.3 (51.6)
qwen2.5:7b	36.4 (28.8)	37.8 (30.1)	36.6 (28.6)	43.9 (43.7)	41.2 (42.3)	49.1 (47.3)
qwen3.5:35b	44.7 (29.9)	45.1 (30.3)	45.6 (29.8)	54.3 (42.2)	49.6 (40.3)	62.3 (46.7)
qwen3.5:27b	47.5 (29.8)	48.5 (30.2)	48.3 (29.9)	**60.7 (39.9)**	**55.5 (38.5)**	**69.2 (43.9)**
qwen3.5:9b	29.2 (25.4)	30.3 (27.2)	31.4 (26.4)	29.9 (39.3)	27.1 (36.7)	34.9 (44.8)
qwen3.5:4b	22.2 (22.2)	22.2 (23.8)	25.8 (24.8)	30.9 (38.6)	28.1 (36.2)	36.0 (44.1)

Table 12: Citation benchmark author and venue metrics by task variant. Task variant: Title + venue + year. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
command-r:35b	16.6 (27.4)	17.6 (28.9)	16.8 (27.7)	90.1 (23.1)	87.2 (24.7)	94.4 (22.6)
command-r7b:7b	1.1 (8.0)	1.1 (7.9)	1.3 (8.7)	77.4 (36.2)	74.7 (37.1)	82.7 (36.1)
deepseek-r1:32b	9.8 (22.7)	10.2 (23.4)	9.9 (23.3)	87.6 (15.6)	80.8 (20.4)	98.7 (10.9)
deepseek-r1:14b	2.9 (12.2)	3.3 (13.5)	2.9 (12.0)	72.0 (39.5)	68.2 (39.0)	77.7 (41.5)
deepseek-r1:8b	0.1 (1.1)	0.1 (1.6)	0.1 (1.0)	0.5 (6.8)	0.5 (6.4)	0.6 (7.8)
gemma2:27b	13.0 (26.3)	13.5 (27.0)	12.9 (26.3)	92.5 (15.1)	88.6 (18.9)	98.6 (11.9)
gemma2:9b	3.4 (13.4)	3.6 (13.9)	3.6 (14.1)	89.2 (19.1)	84.1 (21.8)	96.6 (17.8)
gemma3:27b	7.1 (18.8)	7.5 (19.3)	7.3 (19.6)	92.6 (13.6)	88.4 (17.8)	98.9 (10.4)
gemma3:12b	2.5 (10.3)	2.6 (10.8)	2.7 (11.0)	**94.5 (14.6)**	**92.1 (17.4)**	98.1 (13.0)
gemma3:4b	0.1 (1.6)	0.1 (1.4)	0.1 (2.1)	80.4 (39.6)	80.3 (39.6)	80.5 (39.7)
llama3.1:70b	41.7 (35.2)	41.9 (35.4)	42.5 (35.5)	87.5 (20.6)	82.0 (23.5)	96.2 (18.5)
llama3.1:8b	5.1 (16.6)	5.5 (17.8)	5.1 (16.9)	83.2 (34.8)	81.9 (35.2)	85.5 (34.9)
ministral-3:14b	**57.7 (59.8)**	**55.6 (62.8)**	**62.5 (53.0)**	83.3 (23.6)	75.0 (35.4)	**100.0 (0.0)**
ministral-3:8b	10.4 (23.0)	10.0 (22.6)	11.1 (24.0)	65.7 (36.7)	58.8 (35.8)	78.2 (40.7)
mistral:7b	1.5 (7.2)	1.6 (7.3)	1.7 (8.2)	88.3 (24.4)	84.7 (25.9)	93.7 (24.1)
mistral-nemo	3.5 (12.8)	4.6 (16.0)	3.2 (11.9)	88.6 (22.8)	84.9 (25.2)	94.3 (21.8)
mixtral:8x7b	6.9 (17.3)	6.9 (17.4)	7.6 (19.1)	64.1 (43.1)	60.4 (41.9)	69.7 (45.9)
phi4:14b	6.3 (18.4)	6.5 (18.8)	6.5 (19.1)	86.1 (17.5)	78.9 (21.5)	97.9 (14.2)
qwen2.5:14b	10.5 (22.5)	10.8 (23.1)	10.8 (23.3)	91.9 (19.7)	89.0 (22.4)	97.1 (16.7)
qwen2.5:7b	8.4 (20.4)	8.8 (21.2)	8.5 (20.3)	93.9 (14.9)	91.0 (17.8)	98.6 (12.0)
qwen3.5:35b	19.6 (31.1)	19.8 (31.3)	20.0 (31.5)	87.3 (18.5)	81.2 (22.8)	98.2 (13.0)
qwen3.5:27b	20.8 (31.2)	21.2 (31.6)	21.0 (31.6)	91.2 (14.1)	86.2 (19.1)	99.0 (9.3)
qwen3.5:9b	5.2 (16.2)	5.3 (16.9)	5.6 (17.2)	88.0 (22.7)	83.8 (24.9)	94.4 (21.9)
qwen3.5:4b	1.5 (7.5)	1.5 (7.8)	1.8 (8.7)	89.8 (21.8)	86.3 (24.5)	95.4 (20.2)

Table 13: Citation benchmark author and venue metrics by task variant. Task variant: Title + first author + year. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
command-r:35b	32.0 (21.8)	37.6 (28.9)	33.2 (23.1)	24.1 (36.7)	22.1 (34.3)	27.4 (41.5)
command-r7b:7b	25.1 (15.7)	36.9 (32.0)	24.7 (17.3)	7.7 (23.4)	6.9 (21.6)	9.0 (27.0)
deepseek-r1:32b	34.1 (23.3)	46.3 (33.9)	33.0 (24.5)	20.8 (33.8)	17.9 (29.9)	26.0 (41.4)
deepseek-r1:14b	31.3 (19.1)	60.5 (38.6)	25.6 (18.8)	8.2 (23.8)	7.1 (21.0)	10.1 (28.8)
deepseek-r1:8b	0.3 (4.1)	0.4 (5.3)	0.3 (4.2)	0.2 (3.8)	0.1 (3.5)	0.2 (4.5)
gemma2:27b	40.3 (26.5)	58.4 (35.3)	37.1 (27.9)	6.9 (23.9)	6.5 (22.6)	7.6 (26.1)
gemma2:9b	32.3 (18.2)	62.7 (37.9)	27.2 (18.9)	3.1 (15.4)	2.7 (13.9)	3.7 (18.3)
gemma3:27b	28.3 (19.8)	32.4 (23.9)	28.5 (21.4)	12.0 (29.2)	11.0 (27.1)	13.6 (32.7)
gemma3:12b	26.9 (16.3)	44.3 (36.1)	24.5 (16.2)	0.8 (7.8)	0.7 (7.0)	1.0 (8.8)
gemma3:4b	31.7 (16.1)	77.5 (36.2)	22.8 (14.2)	2.4 (13.4)	2.1 (11.8)	2.9 (15.8)
llama3.1:70b	**57.7 (29.9)**	57.5 (30.6)	59.6 (29.8)	29.1 (40.9)	26.5 (37.9)	33.2 (46.0)
llama3.1:8b	32.4 (20.1)	55.6 (37.5)	28.4 (20.3)	0.2 (2.7)	0.2 (2.5)	0.2 (3.2)
ministral-3:14b	–	–	–	–	–	–
ministral-3:8b	35.0 (24.2)	49.9 (35.6)	34.1 (27.1)	6.2 (18.2)	5.6 (16.5)	7.4 (21.8)
mistral:7b	29.0 (17.3)	55.7 (39.4)	24.1 (16.7)	6.5 (21.5)	5.8 (19.5)	7.7 (25.0)
mistral-nemo	36.2 (15.8)	**94.4 (20.8)**	23.8 (13.6)	0.1 (2.7)	0.1 (2.4)	0.1 (3.2)
mixtral:8x7b	24.7 (21.6)	31.4 (30.4)	25.0 (23.1)	34.7 (39.8)	31.5 (37.1)	40.1 (45.3)
phi4:14b	31.5 (20.9)	46.8 (34.4)	30.1 (23.1)	22.1 (34.4)	19.3 (30.8)	26.9 (41.3)
qwen2.5:14b	55.6 (n/a)	50.0 (n/a)	**62.5 (n/a)**	**66.7 (n/a)**	**57.1 (n/a)**	**80.0 (n/a)**
qwen2.5:7b	37.0 (27.4)	40.6 (31.1)	37.2 (27.5)	13.6 (32.2)	12.7 (30.7)	15.3 (35.6)
qwen3.5:35b	42.5 (28.9)	44.5 (30.3)	43.2 (29.3)	48.2 (42.5)	43.5 (40.1)	56.5 (48.0)
qwen3.5:27b	46.1 (28.4)	47.3 (28.9)	46.6 (29.0)	41.4 (42.8)	37.6 (40.1)	47.7 (48.4)
qwen3.5:9b	28.2 (21.4)	31.1 (26.4)	30.1 (22.5)	28.7 (39.5)	26.3 (37.1)	32.7 (44.3)
qwen3.5:4b	20.7 (16.3)	22.4 (22.1)	24.1 (19.1)	30.3 (38.5)	27.2 (35.6)	35.7 (44.7)

Table 14: Citation benchmark title metrics by task variant. Task variant: Title only. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
command-r:35b	96.5 (16.0)	96.5 (16.1)	96.7 (15.7)	95.7 (18.7)	95.7 (18.7)	95.9 (18.6)
command-r7b:7b	96.0 (17.9)	96.1 (17.9)	96.1 (17.7)	95.3 (20.2)	95.5 (20.1)	95.3 (20.3)
deepseek-r1:32b	97.7 (14.0)	97.8 (13.9)	97.8 (13.9)	97.4 (15.4)	97.5 (15.2)	97.5 (15.2)
deepseek-r1:14b	96.7 (17.1)	96.7 (17.2)	96.8 (16.8)	96.5 (17.6)	96.5 (17.7)	96.6 (17.4)
deepseek-r1:8b	0.8 (9.0)	0.8 (9.0)	0.8 (9.0)	0.8 (9.0)	0.8 (9.0)	0.8 (9.0)
gemma2:27b	98.9 (9.3)	98.9 (9.4)	98.9 (9.2)	98.6 (10.9)	98.7 (10.9)	98.6 (10.9)
gemma2:9b	91.7 (27.5)	91.7 (27.5)	91.7 (27.5)	91.5 (27.7)	91.5 (27.7)	91.5 (27.7)
gemma3:27b	99.5 (6.0)	99.5 (6.3)	99.6 (5.8)	99.3 (8.0)	99.2 (8.1)	99.3 (7.6)
gemma3:12b	99.3 (8.4)	99.3 (8.4)	99.3 (8.4)	99.2 (9.0)	99.2 (9.0)	99.2 (9.0)
gemma3:4b	95.9 (19.5)	96.0 (19.5)	95.9 (19.6)	95.9 (19.6)	96.0 (19.6)	95.9 (19.6)
llama3.1:70b	98.7 (10.9)	98.7 (10.9)	98.7 (11.0)	98.5 (11.7)	98.5 (11.8)	98.5 (11.7)
llama3.1:8b	95.9 (19.4)	95.9 (19.4)	96.0 (19.3)	95.7 (19.8)	95.7 (19.9)	95.8 (19.8)
ministral-3:14b	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**
ministral-3:8b	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**
mistral:7b	98.1 (12.1)	98.1 (12.3)	98.2 (11.7)	97.5 (14.3)	97.7 (14.2)	97.6 (14.2)
mistral-nemo	97.8 (14.3)	97.8 (14.2)	97.7 (14.3)	97.5 (15.0)	97.6 (14.9)	97.5 (15.0)
mixtral:8x7b	86.9 (33.3)	86.9 (33.3)	86.9 (33.2)	86.4 (33.9)	86.5 (33.8)	86.4 (33.9)
phi4:14b	98.2 (12.4)	98.1 (12.8)	98.4 (11.6)	97.8 (14.2)	97.7 (14.5)	98.0 (13.5)
qwen2.5:14b	98.3 (11.8)	98.2 (12.2)	98.4 (11.4)	98.0 (13.3)	97.9 (13.5)	98.0 (13.0)
qwen2.5:7b	96.1 (17.1)	96.7 (16.1)	96.0 (17.7)	95.4 (19.4)	95.8 (18.6)	95.4 (19.6)
qwen3.5:35b	98.0 (12.0)	98.0 (11.9)	98.0 (12.2)	97.4 (14.4)	97.4 (14.3)	97.4 (14.4)
qwen3.5:27b	97.0 (14.8)	96.9 (15.1)	97.2 (14.5)	96.2 (17.9)	96.2 (18.0)	96.4 (17.5)
qwen3.5:9b	97.9 (12.6)	98.0 (12.4)	98.0 (12.4)	97.3 (15.3)	97.4 (15.4)	97.4 (15.2)
qwen3.5:4b	95.6 (19.8)	95.7 (19.8)	95.7 (19.8)	95.3 (20.7)	95.3 (20.7)	95.4 (20.6)

Table 15: Citation benchmark title metrics by task variant. Task variant: Title only + DOI/URL requirement. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
command-r:35b	94.6 (19.2)	94.5 (19.4)	95.0 (18.9)	93.4 (22.8)	93.2 (22.9)	93.7 (22.5)
command-r7b:7b	95.7 (19.1)	95.7 (18.9)	95.7 (19.0)	95.0 (20.8)	95.1 (20.6)	95.0 (20.8)
deepseek-r1:32b	96.0 (18.3)	96.0 (18.3)	96.2 (18.2)	95.5 (20.1)	95.4 (20.2)	95.7 (19.8)
deepseek-r1:14b	95.5 (20.5)	95.5 (20.5)	95.5 (20.5)	95.4 (20.8)	95.5 (20.7)	95.4 (20.8)
deepseek-r1:8b	0.7 (8.4)	0.7 (8.4)	0.7 (8.4)	0.7 (8.4)	0.7 (8.4)	0.7 (8.4)
gemma2:27b	98.5 (11.1)	98.4 (11.3)	98.6 (10.7)	98.0 (13.5)	98.0 (13.6)	98.1 (13.3)
gemma2:9b	95.9 (19.6)	96.0 (19.5)	95.9 (19.6)	95.7 (20.0)	95.8 (20.0)	95.7 (20.0)
gemma3:27b	**99.5 (6.1)**	**99.4 (6.4)**	99.5 (5.5)	99.2 (8.3)	**99.2 (8.4)**	99.2 (8.0)
gemma3:12b	99.2 (8.4)	99.2 (8.5)	99.3 (8.1)	99.1 (9.4)	99.1 (9.4)	99.1 (9.4)
gemma3:4b	–	–	–	–	–	–
llama3.1:70b	98.0 (13.1)	98.0 (12.8)	98.0 (13.2)	97.7 (14.4)	97.7 (14.3)	97.6 (14.4)
llama3.1:8b	96.2 (18.5)	96.1 (18.6)	96.3 (18.3)	95.9 (19.0)	95.8 (19.2)	96.1 (18.8)
ministral-3:14b	99.4 (2.9)	99.0 (5.0)	**100.0 (0.0)**	**99.2 (4.0)**	99.0 (5.0)	**99.4 (2.9)**
ministral-3:8b	95.8 (18.4)	95.9 (18.6)	96.0 (18.2)	94.9 (20.7)	95.0 (20.7)	95.2 (20.3)
mistral:7b	94.0 (21.8)	94.0 (21.9)	94.2 (21.6)	93.0 (24.2)	93.0 (24.3)	93.2 (23.9)
mistral-nemo	99.3 (7.6)	99.3 (7.7)	99.3 (7.5)	99.1 (8.7)	99.1 (8.8)	99.1 (8.7)
mixtral:8x7b	86.8 (33.4)	86.8 (33.4)	86.9 (33.3)	86.3 (34.1)	86.4 (34.1)	86.4 (34.1)
phi4:14b	96.2 (18.0)	96.2 (18.1)	96.3 (17.7)	95.7 (19.6)	95.7 (19.6)	95.9 (19.4)
qwen2.5:14b	97.5 (14.4)	97.4 (14.3)	97.6 (14.4)	96.9 (15.9)	96.7 (16.1)	97.0 (15.7)
qwen2.5:7b	95.9 (16.1)	96.6 (15.4)	95.5 (16.9)	94.8 (19.3)	95.6 (18.6)	94.5 (20.1)
qwen3.5:35b	97.5 (12.0)	97.4 (12.4)	97.7 (11.6)	96.5 (15.7)	96.4 (16.0)	96.7 (15.4)
qwen3.5:27b	96.5 (15.1)	96.5 (15.3)	96.7 (14.9)	95.4 (18.6)	95.5 (18.6)	95.6 (18.4)
qwen3.5:9b	97.6 (13.7)	97.6 (14.0)	97.8 (13.3)	97.3 (15.2)	97.3 (15.4)	97.5 (14.8)
qwen3.5:4b	95.6 (19.4)	95.6 (19.4)	95.7 (19.3)	95.0 (20.8)	95.0 (20.8)	95.1 (20.7)

Table 16: Citation benchmark title metrics by task variant. Task variant: Title + author hint. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
command-r:35b	97.8 (12.0)	97.7 (12.2)	98.0 (11.6)	97.0 (15.2)	96.9 (15.4)	97.2 (14.9)
command-r7b:7b	97.7 (13.6)	97.6 (13.8)	97.8 (13.4)	97.1 (15.4)	97.1 (15.6)	97.2 (15.1)
deepseek-r1:32b	98.2 (10.6)	98.2 (10.8)	98.4 (10.2)	97.5 (13.3)	97.5 (13.5)	97.7 (12.9)
deepseek-r1:14b	96.5 (16.1)	96.3 (16.6)	96.9 (15.5)	95.5 (18.4)	95.4 (18.7)	95.8 (18.0)
deepseek-r1:8b	0.6 (7.8)	0.6 (7.8)	0.6 (7.8)	0.6 (7.8)	0.6 (7.8)	0.6 (7.8)
gemma2:27b	98.5 (11.9)	98.5 (12.0)	98.6 (11.8)	98.3 (12.5)	98.3 (12.6)	98.4 (12.3)
gemma2:9b	99.3 (7.6)	99.2 (8.0)	99.4 (7.1)	99.1 (8.7)	99.0 (9.1)	99.2 (8.0)
gemma3:27b	99.5 (5.4)	99.5 (5.6)	99.5 (5.2)	99.2 (7.7)	99.2 (7.9)	99.3 (7.2)
gemma3:12b	99.7 (4.6)	99.7 (4.8)	99.8 (4.5)	99.6 (5.8)	99.6 (6.0)	99.7 (5.6)
gemma3:4b	97.8 (13.9)	97.8 (14.1)	98.0 (13.5)	97.7 (14.4)	97.7 (14.5)	97.8 (14.2)
llama3.1:70b	99.3 (5.7)	99.3 (6.0)	99.5 (5.2)	98.9 (8.5)	98.9 (8.7)	99.0 (8.0)
llama3.1:8b	96.6 (17.3)	96.5 (17.6)	96.8 (17.0)	96.3 (18.1)	96.2 (18.3)	96.5 (17.6)
ministral-3:14b	–	–	–	–	–	–
ministral-3:8b	94.0 (21.9)	94.3 (21.7)	94.0 (22.1)	93.1 (23.8)	93.2 (23.5)	93.0 (23.9)
mistral:7b	90.3 (24.1)	90.7 (24.1)	90.5 (24.1)	87.5 (28.9)	87.9 (28.7)	87.8 (28.9)
mistral-nemo	97.0 (16.8)	97.0 (16.8)	97.0 (16.7)	96.8 (17.2)	96.8 (17.2)	96.9 (17.2)
mixtral:8x7b	92.4 (25.7)	92.2 (26.0)	92.9 (25.4)	92.0 (26.4)	91.8 (26.8)	92.6 (26.0)
phi4:14b	98.6 (10.3)	98.5 (10.8)	98.8 (9.6)	98.2 (12.3)	98.1 (12.7)	98.5 (11.6)
qwen2.5:14b	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**
qwen2.5:7b	98.5 (10.1)	98.4 (10.5)	98.6 (9.7)	97.9 (13.2)	97.9 (13.3)	97.9 (13.1)
qwen3.5:35b	99.1 (7.6)	99.1 (7.7)	99.2 (7.6)	98.9 (9.2)	98.8 (9.3)	98.9 (9.0)
qwen3.5:27b	99.3 (5.6)	99.3 (6.2)	99.5 (5.0)	99.0 (7.5)	99.0 (8.0)	99.2 (6.8)
qwen3.5:9b	98.5 (10.0)	98.5 (10.2)	98.6 (9.7)	97.9 (12.9)	98.0 (12.8)	98.0 (12.7)
qwen3.5:4b	96.4 (17.0)	96.3 (17.3)	96.7 (16.7)	95.9 (18.4)	95.7 (18.7)	96.2 (18.1)

Table 17: Citation benchmark title metrics by task variant. Task variant: Title + venue + year. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
command-r:35b	92.6 (22.9)	92.8 (22.8)	92.7 (22.8)	90.7 (28.0)	90.8 (27.8)	90.6 (28.2)
command-r7b:7b	93.8 (22.8)	93.8 (22.8)	93.9 (22.7)	93.1 (24.6)	93.1 (24.6)	93.2 (24.5)
deepseek-r1:32b	98.1 (12.1)	98.0 (12.6)	98.4 (11.6)	97.6 (14.2)	97.5 (14.6)	97.9 (13.5)
deepseek-r1:14b	84.1 (36.5)	84.1 (36.5)	84.0 (36.5)	84.0 (36.6)	84.0 (36.6)	84.0 (36.6)
deepseek-r1:8b	0.5 (7.0)	0.5 (6.9)	0.5 (7.0)	0.5 (6.9)	0.5 (6.8)	0.5 (7.0)
gemma2:27b	98.7 (10.1)	98.8 (10.0)	98.7 (10.2)	98.4 (11.9)	98.4 (11.8)	98.4 (11.8)
gemma2:9b	98.0 (13.6)	98.0 (13.5)	98.0 (13.6)	97.8 (14.2)	97.8 (14.1)	97.8 (14.3)
gemma3:27b	98.9 (8.2)	98.9 (8.4)	99.0 (7.8)	98.4 (11.2)	98.4 (11.4)	98.5 (10.8)
gemma3:12b	99.6 (5.6)	99.6 (5.6)	99.7 (5.6)	99.5 (6.5)	99.5 (6.6)	99.5 (6.5)
gemma3:4b	85.0 (35.5)	85.1 (35.6)	85.0 (35.5)	84.8 (35.7)	84.9 (35.7)	84.8 (35.7)
llama3.1:70b	95.6 (19.1)	95.6 (19.0)	95.7 (19.1)	94.8 (21.3)	94.8 (21.2)	94.8 (21.3)
llama3.1:8b	87.4 (32.1)	87.2 (32.3)	87.6 (31.9)	86.7 (33.0)	86.6 (33.2)	87.0 (32.9)
ministral-3:14b	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**	**100.0 (0.0)**
ministral-3:8b	97.3 (14.5)	97.5 (14.1)	97.1 (14.9)	96.8 (15.8)	96.9 (15.2)	96.6 (16.2)
mistral:7b	95.5 (18.6)	95.6 (18.8)	95.8 (18.2)	94.8 (20.7)	94.9 (20.8)	95.1 (20.2)
mistral-nemo	97.7 (14.3)	97.8 (14.3)	97.8 (14.2)	97.5 (15.1)	97.5 (15.1)	97.5 (15.0)
mixtral:8x7b	69.5 (45.9)	69.5 (45.9)	69.6 (45.9)	69.3 (46.0)	69.3 (46.0)	69.3 (46.0)
phi4:14b	95.2 (18.8)	95.0 (19.2)	95.8 (18.2)	94.1 (22.1)	93.9 (22.4)	94.7 (21.5)
qwen2.5:14b	94.6 (20.1)	94.6 (20.2)	94.8 (19.9)	93.3 (23.7)	93.4 (23.6)	93.4 (23.6)
qwen2.5:7b	97.6 (13.5)	97.5 (13.4)	97.8 (13.6)	97.0 (15.1)	96.9 (15.2)	97.3 (14.9)
qwen3.5:35b	97.3 (13.2)	97.3 (13.3)	97.5 (12.8)	96.2 (17.1)	96.2 (17.1)	96.4 (16.9)
qwen3.5:27b	97.4 (12.3)	97.4 (12.6)	97.6 (12.0)	96.4 (16.4)	96.3 (16.6)	96.6 (16.1)
qwen3.5:9b	96.0 (18.3)	95.9 (18.5)	96.2 (17.9)	95.4 (19.9)	95.3 (20.2)	95.6 (19.6)
qwen3.5:4b	96.4 (17.2)	96.5 (17.3)	96.6 (17.0)	95.9 (18.6)	95.9 (18.6)	96.0 (18.4)

Table 18: Citation benchmark title metrics by task variant. Task variant: Title + first author + year. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
command-r:35b	99.5 (5.8)	99.5 (5.8)	99.5 (5.7)	99.2 (7.9)	99.2 (8.1)	99.3 (7.7)
command-r7b:7b	98.9 (10.1)	99.0 (10.0)	98.8 (10.3)	98.7 (10.7)	98.9 (10.5)	98.7 (10.8)
deepseek-r1:32b	98.1 (12.7)	98.1 (12.7)	98.2 (12.8)	97.9 (13.7)	97.8 (13.6)	97.9 (13.7)
deepseek-r1:14b	95.3 (20.8)	95.2 (20.9)	95.4 (20.7)	95.2 (21.0)	95.1 (21.1)	95.3 (20.8)
deepseek-r1:8b	0.9 (9.2)	0.9 (9.2)	0.9 (9.2)	0.8 (8.9)	0.8 (8.9)	0.8 (8.9)
gemma2:27b	98.8 (10.9)	98.8 (10.9)	98.8 (10.9)	98.6 (11.5)	98.6 (11.5)	98.6 (11.5)
gemma2:9b	99.0 (10.0)	98.9 (10.0)	99.0 (10.0)	98.8 (10.6)	98.8 (10.7)	98.8 (10.6)
gemma3:27b	99.8 (2.1)	99.8 (2.8)	99.9 (1.0)	99.7 (4.4)	99.6 (4.8)	99.8 (3.5)
gemma3:12b	99.6 (5.7)	99.6 (5.7)	99.6 (5.8)	99.5 (6.8)	99.5 (6.7)	99.5 (6.8)
gemma3:4b	99.2 (8.3)	99.2 (8.5)	99.3 (8.2)	99.0 (9.5)	99.0 (9.7)	99.1 (9.2)
llama3.1:70b	99.3 (7.8)	99.2 (8.0)	99.3 (7.7)	98.9 (9.4)	98.9 (9.5)	98.9 (9.3)
llama3.1:8b	97.2 (15.2)	97.2 (15.3)	97.4 (15.1)	97.0 (16.1)	97.0 (16.2)	97.1 (15.9)
ministral-3:14b	–	–	–	–	–	–
ministral-3:8b	98.7 (10.3)	98.8 (9.8)	98.7 (10.6)	98.5 (11.8)	98.5 (11.8)	98.5 (11.8)
mistral:7b	97.8 (13.9)	98.0 (13.8)	97.8 (14.1)	97.5 (15.0)	97.7 (14.7)	97.4 (15.2)
mistral-nemo	99.8 (4.5)	99.8 (4.5)	99.8 (4.5)	99.6 (5.8)	99.6 (5.8)	99.6 (5.7)
mixtral:8x7b	86.0 (34.6)	86.0 (34.6)	86.0 (34.6)	85.8 (34.7)	85.8 (34.7)	85.8 (34.7)
phi4:14b	99.4 (7.3)	99.4 (7.4)	99.5 (7.1)	99.3 (8.1)	99.2 (8.2)	99.3 (7.8)
qwen2.5:14b	**100.0 (n/a)**	**100.0 (n/a)**	**100.0 (n/a)**	**100.0 (n/a)**	**100.0 (n/a)**	**100.0 (n/a)**
qwen2.5:7b	99.6 (4.3)	99.5 (5.8)	**100.0 (0.0)**	99.5 (5.8)	99.4 (7.0)	**100.0 (0.0)**
qwen3.5:35b	99.2 (6.2)	99.1 (6.7)	99.4 (5.8)	98.8 (8.2)	98.8 (8.6)	99.0 (7.5)
qwen3.5:27b	99.5 (4.7)	99.4 (5.3)	99.7 (4.1)	99.2 (6.8)	99.1 (7.3)	99.4 (6.1)
qwen3.5:9b	99.4 (7.0)	99.4 (7.1)	99.4 (6.9)	99.2 (8.3)	99.1 (8.3)	99.2 (8.2)
qwen3.5:4b	96.4 (17.4)	96.3 (17.6)	96.6 (17.2)	95.8 (18.9)	95.7 (19.1)	96.0 (18.6)

is the strongest author model by a large margin, whereas `qwen3.5:27b` and `qwen3.5:35b` are strongest on venue fidelity. `ministral-3:14b` is notably uneven, combining near-perfect title preservation and strong author recovery with weak venue fidelity.

A.10 Small Language Model Results

Tables 21–39 report results for the small language models not included in the main paper.

A.10.1 Overall behavior in the low-parameter regime

The small models are heterogeneous. Some are refusal-heavy, such as `gemma3:270m`, `llama3.2:1b`, and `phi4-mini:3.8b`, while others, such as `qwen3.5:0.8b` and `qwen3.5:2b`, answer aggressively and produce parseable BibTeX for most real-paper prompts but remain weak on recalled citation fields. `ministral-3:3b` is the strongest small model on recalled fields, but it still combines only moderate parseability with a high synthetic-refusal rate. Overall, the small models exhibit the same broad failure modes as the main-paper models, but at substantially weaker levels of citation fidelity.

A.10.2 Field-level fidelity at small scale

Tables 23, 24, and 21 show that small-model field fidelity is uneven across fields. The clearest pattern is the separation between title preservation and recalled metadata: some small models retain strong title fidelity while remaining weak on author, venue, year, and DOI. This indicates that small-scale improvements first appear in copying the visible structure of a citation rather than in recovering its bibliographic content.

The tables also show that low-parameter degradation is not uniform across models. Some small models preserve title much better than the remaining fields, whereas others are weak across all fields. Within this group, the strongest models improve recalled-field performance, but the gap

Table 19: Citation benchmark author and venue metrics. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
command-r:35b	21.1 (25.6)	22.4 (28.1)	22.4 (27.0)	28.9 (39.3)	26.7 (37.1)	32.8 (44.0)
command-r7b:7b	10.9 (18.4)	13.5 (24.4)	11.2 (19.2)	24.4 (39.6)	23.1 (38.2)	27.1 (42.8)
deepseek-r1:32b	20.6 (27.1)	25.4 (33.4)	20.3 (27.6)	36.5 (40.8)	32.4 (37.6)	43.8 (47.5)
deepseek-r1:14b	15.3 (22.1)	24.8 (36.4)	13.7 (21.3)	26.3 (39.5)	23.8 (36.9)	30.4 (44.6)
deepseek-r1:8b	0.2 (3.1)	0.2 (3.5)	0.2 (3.2)	0.4 (5.6)	0.3 (5.2)	0.4 (6.4)
gemma2:27b	23.3 (30.0)	28.0 (35.2)	22.5 (29.9)	29.9 (43.2)	28.3 (41.5)	32.5 (46.3)
gemma2:9b	14.8 (22.0)	22.2 (34.1)	13.9 (21.7)	27.6 (40.0)	25.0 (37.4)	32.0 (45.2)
gemma3:27b	15.6 (22.5)	16.8 (24.2)	15.8 (23.3)	35.9 (44.0)	33.7 (42.2)	39.4 (47.5)
gemma3:12b	12.4 (19.3)	16.2 (27.6)	12.3 (19.5)	24.4 (41.2)	23.6 (40.3)	25.7 (42.8)
gemma3:4b	13.0 (18.9)	27.1 (40.4)	10.4 (16.3)	19.6 (39.1)	19.5 (39.0)	19.9 (39.5)
llama3.1:70b	**49.6 (33.7)**	**49.5 (34.1)**	**50.8 (33.9)**	49.0 (44.4)	45.3 (42.3)	54.9 (48.7)
llama3.1:8b	15.9 (23.4)	21.9 (33.0)	15.0 (23.1)	21.7 (39.2)	20.8 (38.4)	23.1 (41.1)
ministral-3:14b	32.5 (36.0)	33.0 (36.1)	32.6 (36.2)	16.7 (33.5)	14.9 (30.7)	19.6 (38.7)
ministral-3:8b	18.3 (26.6)	21.3 (31.3)	18.9 (27.8)	12.3 (27.7)	11.1 (25.4)	15.0 (33.1)
mistral:7b	8.2 (16.7)	13.6 (29.6)	7.4 (16.0)	27.2 (41.7)	25.6 (40.1)	29.7 (44.8)
mistral-nemo	15.7 (22.4)	30.5 (41.6)	12.4 (19.8)	22.6 (38.9)	21.3 (37.4)	24.8 (42.1)
mixtral:8x7b	15.4 (22.7)	16.8 (25.4)	16.1 (24.0)	40.1 (42.0)	36.8 (39.6)	45.6 (46.9)
phi4:14b	17.9 (24.5)	23.1 (32.2)	17.6 (25.2)	34.7 (40.7)	31.0 (37.6)	40.9 (47.0)
qwen2.5:14b	11.0 (23.2)	11.3 (23.8)	11.3 (23.8)	38.8 (46.1)	37.0 (45.0)	42.0 (48.8)
qwen2.5:7b	21.0 (28.1)	22.3 (29.9)	21.1 (28.0)	50.5 (45.9)	48.2 (44.9)	54.5 (48.5)
qwen3.5:35b	29.1 (32.3)	29.7 (33.0)	29.6 (32.7)	55.2 (41.8)	50.2 (40.1)	**63.9 (46.5)**
qwen3.5:27b	31.0 (33.0)	31.8 (33.6)	31.4 (33.4)	**55.5 (42.0)**	**50.9 (40.4)**	63.2 (46.4)
qwen3.5:9b	14.5 (22.4)	15.4 (24.7)	15.5 (23.7)	39.2 (43.2)	36.4 (41.3)	44.0 (47.5)
qwen3.5:4b	9.6 (17.0)	9.9 (18.9)	11.1 (19.4)	38.7 (42.9)	35.9 (41.1)	43.6 (47.2)

Table 20: Citation benchmark title metrics. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
command-r:35b	96.9 (14.7)	96.9 (14.8)	97.1 (14.4)	96.1 (17.8)	96.0 (17.8)	96.3 (17.5)
command-r7b:7b	96.4 (17.3)	96.5 (17.3)	96.5 (17.3)	95.9 (19.0)	95.9 (19.0)	95.9 (19.0)
deepseek-r1:32b	97.7 (13.7)	97.6 (13.8)	97.8 (13.5)	97.2 (15.4)	97.2 (15.5)	97.4 (15.1)
deepseek-r1:14b	93.6 (23.8)	93.6 (23.9)	93.7 (23.7)	93.3 (24.3)	93.3 (24.4)	93.4 (24.2)
deepseek-r1:8b	0.7 (8.3)	0.7 (8.3)	0.7 (8.3)	0.7 (8.2)	0.7 (8.2)	0.7 (8.2)
gemma2:27b	98.7 (10.7)	98.7 (10.8)	98.7 (10.6)	98.4 (12.1)	98.4 (12.1)	98.4 (12.0)
gemma2:9b	96.9 (17.1)	96.9 (17.1)	96.9 (17.1)	96.7 (17.5)	96.7 (17.5)	96.7 (17.5)
gemma3:27b	99.4 (5.9)	99.4 (6.2)	99.5 (5.5)	99.1 (8.2)	99.1 (8.4)	99.2 (7.8)
gemma3:12b	99.5 (6.6)	99.5 (6.7)	99.5 (6.6)	99.4 (7.5)	99.4 (7.6)	99.4 (7.5)
gemma3:4b	94.6 (22.3)	94.7 (22.3)	94.7 (22.2)	94.5 (22.5)	94.6 (22.6)	94.5 (22.5)
llama3.1:70b	97.9 (13.1)	98.0 (13.0)	98.0 (13.1)	97.5 (14.8)	97.5 (14.8)	97.5 (14.7)
llama3.1:8b	94.6 (21.7)	94.5 (21.9)	94.8 (21.5)	94.3 (22.4)	94.2 (22.6)	94.5 (22.2)
ministral-3:14b	**99.7 (1.9)**	**99.6 (3.4)**	**100.0 (0.0)**	**99.6 (2.7)**	**99.6 (3.4)**	**99.7 (1.9)**
ministral-3:8b	96.0 (18.0)	96.1 (17.9)	96.0 (18.1)	95.2 (19.9)	95.3 (19.8)	95.3 (19.8)
mistral:7b	95.9 (17.8)	96.0 (17.9)	96.0 (17.6)	95.1 (20.1)	95.2 (20.1)	95.2 (19.9)
mistral-nemo	98.3 (12.4)	98.3 (12.4)	98.3 (12.4)	98.1 (13.1)	98.1 (13.1)	98.1 (13.1)
mixtral:8x7b	84.4 (36.0)	84.3 (36.0)	84.5 (35.9)	84.0 (36.3)	84.0 (36.4)	84.2 (36.3)
phi4:14b	97.5 (14.0)	97.5 (14.3)	97.8 (13.6)	97.0 (16.1)	97.0 (16.3)	97.3 (15.6)
qwen2.5:14b	96.9 (15.6)	96.8 (15.8)	97.0 (15.4)	96.2 (18.1)	96.2 (18.1)	96.3 (17.9)
qwen2.5:7b	97.7 (13.0)	97.8 (12.7)	97.8 (13.1)	97.1 (15.2)	97.2 (15.0)	97.2 (15.1)
qwen3.5:35b	98.3 (10.5)	98.2 (10.7)	98.4 (10.3)	97.6 (13.3)	97.6 (13.4)	97.7 (13.1)
qwen3.5:27b	98.0 (11.5)	97.9 (11.7)	98.2 (11.1)	97.3 (14.4)	97.3 (14.5)	97.5 (14.0)
qwen3.5:9b	97.9 (12.9)	97.9 (13.0)	98.0 (12.6)	97.4 (14.8)	97.4 (14.9)	97.5 (14.6)
qwen3.5:4b	96.1 (18.2)	96.0 (18.3)	96.3 (18.1)	95.6 (19.5)	95.5 (19.6)	95.8 (19.3)

Table 21: Citation-benchmark results for small models without LLM-judge coverage. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑
deepseek-r1:1.5b	16.4	9.0	16.9	3.0	3.3	11.8	0.0
gemma3:1b	33.0	22.8	42.5	5.8	9.1	22.9	0.0
gemma3:270m	1.3	76.3	5.5	0.4	0.4	1.5	0.0
llama3.2:3b	28.1	7.4	29.3	2.9	6.7	13.8	0.0
llama3.2:1b	10.2	58.6	21.1	4.4	1.8	16.1	0.0
ministral-3:3b	63.2	41.7	**97.7**	**15.2**	**37.8**	**64.2**	0.0
phi4-mini:3.8b Microsoft et al. (2025)	15.3	79.6	69.8	4.2	4.5	23.3	0.0
qwen3.5:2b	**95.5**	**0.2**	94.4	10.7	29.0	45.9	0.0
qwen3.5:0.8b	82.4	0.8	81.1	9.2	17.3	37.6	0.0

Table 22: Refusal-benchmark results for the small models without LLM-judge coverage. All metrics are reported on a 0–100 scale.

Model	Synthetic Refusal ↑	Parseable Synthetic Refusal ↓
deepseek-r1:1.5b	9.5	0.6
gemma3:1b	24.4	0.2
gemma3:270m	76.8	**0.1**
llama3.2:3b	20.1	0.1
llama3.2:1b	64.7	1.4
ministral-3:3b	72.7	1.6
phi4-mini:3.8b	**85.2**	1.1
qwen3.5:2b	0.2	0.2
qwen3.5:0.8b	0.9	0.6

between title fidelity and metadata recovery remains large. The small-model results therefore point to a structured low-parameter regime, in which citation behavior can look superficially well formed while still failing on the fields that require actual bibliographic recall.

Table 23: Citation benchmark title metrics for small models. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
deepseek-r1:1.5b	16.9 (37.3)	17.0 (37.4)	16.9 (37.4)	16.8 (37.2)	16.8 (37.3)	16.8 (37.2)
gemma3:1b	42.5 (49.3)	42.6 (49.4)	42.4 (49.3)	42.4 (49.3)	42.5 (49.4)	42.3 (49.2)
gemma3:270m	5.5 (22.8)	5.5 (22.8)	5.5 (22.8)	5.4 (22.6)	5.4 (22.6)	5.4 (22.6)
llama3.2:1b	21.1 (40.7)	21.1 (40.7)	21.1 (40.7)	21.0 (40.6)	21.0 (40.6)	21.0 (40.6)
llama3.2:3b	29.3 (45.4)	29.3 (45.5)	29.3 (45.5)	29.2 (45.4)	29.2 (45.4)	29.2 (45.4)
ministral-3:3b	**97.7 (14.1)**	**97.7 (14.1)**	**97.7 (14.1)**	**97.3 (15.6)**	**97.4 (15.6)**	**97.4 (15.5)**
phi4-mini:3.8b	69.8 (45.1)	70.0 (45.3)	70.0 (45.2)	69.1 (45.4)	69.3 (45.5)	69.3 (45.5)
qwen3.5:0.8b	81.1 (38.3)	81.0 (38.4)	81.4 (38.4)	80.5 (38.6)	80.4 (38.7)	81.0 (38.6)
qwen3.5:2b	94.4 (22.2)	94.4 (22.3)	94.6 (22.0)	94.0 (23.1)	94.0 (23.2)	94.2 (22.9)

A.10.3 Results By Task

Table 24: Citation benchmark author and venue metrics for small models. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
deepseek-r1:1.5b	3.0 (10.9)	6.4 (23.1)	2.5 (9.8)	3.3 (16.8)	3.1 (16.3)	3.5 (17.7)
gemma3:1b	5.8 (16.0)	13.8 (34.2)	4.0 (12.5)	9.1 (25.9)	8.9 (25.8)	9.6 (26.8)
gemma3:270m	0.4 (3.5)	1.1 (10.4)	0.2 (2.2)	0.4 (4.9)	0.4 (4.5)	0.5 (6.2)
llama3.2:1b	4.4 (13.9)	10.7 (30.7)	3.1 (10.7)	1.8 (13.0)	1.8 (13.0)	1.8 (13.2)
llama3.2:3b	2.9 (11.2)	5.4 (20.6)	2.4 (10.0)	6.7 (23.9)	6.4 (23.3)	7.2 (25.2)
ministral-3:3b	**15.2 (21.9)**	**29.3 (41.2)**	**12.8 (20.1)**	**37.8 (44.5)**	**35.8 (43.3)**	**41.1 (47.4)**
phi4-mini:3.8b	4.2 (11.8)	9.9 (28.5)	3.1 (9.1)	4.5 (17.3)	4.1 (16.2)	5.4 (20.4)
qwen3.5:0.8b	9.2 (18.3)	18.9 (36.9)	7.6 (15.9)	17.3 (34.3)	16.7 (33.7)	18.7 (36.4)
qwen3.5:2b	10.7 (19.1)	21.2 (37.8)	8.8 (16.8)	29.0 (42.0)	28.0 (41.2)	31.1 (44.3)

Table 25: Citation benchmark overall results for small models, task variant: Title only. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
deepseek-r1:1.5b	16.8	9.7	–	15.9	0.0	0.8	3.3	**0.0**	–
gemma3:1b	45.8	1.8	–	46.0	0.0	4.5	3.9	**0.0**	–
gemma3:270m	0.5	96.7	–	15.2	0.0	3.0	0.0	**0.0**	–
llama3.2:1b	0.4	99.6	–	57.1	0.0	0.0	14.3	**0.0**	–
llama3.2:3b	28.8	3.5	–	29.0	0.2	0.0	6.8	**0.0**	–
ministral-3:3b	21.7	78.1	–	**96.6**	**2.5**	11.6	**36.2**	**0.0**	–
phi4-mini:3.8b	59.6	30.2	–	78.5	0.1	3.1	11.5	**0.0**	–
qwen3.5:0.8b	80.6	0.7	–	79.3	0.2	5.6	7.9	**0.0**	–
qwen3.5:2b	**95.6**	**0.0**	–	94.9	0.3	**14.8**	12.5	**0.0**	–

Table 26: Citation benchmark overall results for small models, task variant: Title only + DOI/URL requirement. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
deepseek-r1:1.5b	2.9	11.7	–	2.6	0.0	0.1	0.9	**0.0**	–
gemma3:1b	5.4	86.8	–	40.9	0.1	3.2	2.3	**0.0**	–
gemma3:270m	4.4	9.8	–	4.9	0.0	0.1	0.3	**0.0**	–
llama3.2:1b	11.0	69.9	–	35.7	0.0	0.4	4.2	**0.0**	–
llama3.2:3b	42.8	17.3	–	51.1	0.3	7.9	11.4	**0.0**	–
ministral-3:3b	51.8	71.0	–	**97.9**	**0.9**	**24.4**	**26.7**	**0.0**	–
phi4-mini:3.8b	0.5	97.9	–	22.8	0.0	2.9	9.5	**0.0**	–
qwen3.5:0.8b	84.4	1.0	–	83.2	0.1	11.0	8.5	**0.0**	–
qwen3.5:2b	**94.2**	**0.4**	–	93.2	0.2	16.8	12.4	**0.0**	–

Table 27: Citation benchmark overall results for small models, task variant: Title + author hint. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
deepseek-r1:1.5b	13.9	7.6	–	14.5	3.7	1.2	3.0	**0.0**	–
gemma3:1b	32.8	17.2	–	39.5	3.6	4.8	3.6	**0.0**	–
gemma3:270m	0.1	99.5	–	20.0	0.0	0.0	0.0	**0.0**	–
llama3.2:1b	3.5	26.3	–	4.7	1.4	0.1	0.4	**0.0**	–
llama3.2:3b	19.1	6.9	–	20.4	6.5	1.2	3.4	**0.0**	–
ministral-3:3b	95.3	4.7	–	**98.4**	**31.6**	9.6	**39.9**	**0.0**	–
phi4-mini:3.8b	0.1	93.9	–	1.6	0.5	0.0	0.0	**0.0**	–
qwen3.5:0.8b	82.2	1.4	–	80.7	19.0	9.6	7.8	**0.0**	–
qwen3.5:2b	**96.7**	**0.4**	–	95.3	22.1	**20.3**	14.7	**0.0**	–

Table 28: Citation benchmark overall results for small models, task variant: Title + venue + year. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
deepseek-r1:1.5b	14.3	12.4	–	15.1	0.0	14.5	15.1	**0.0**	–
gemma3:1b	24.0	6.9	–	25.5	0.0	22.8	25.7	**0.0**	–
gemma3:270m	0.2	88.9	–	1.8	0.0	0.9	1.8	**0.0**	–
llama3.2:1b	13.6	95.5	–	86.5	0.0	24.2	77.3	**0.0**	–
llama3.2:3b	27.7	5.3	–	26.0	0.1	24.3	23.9	**0.0**	–
ministral-3:3b	**95.2**	6.6	–	**96.1**	**0.9**	**91.1**	**95.9**	**0.0**	–
phi4-mini:3.8b	1.6	95.2	–	30.9	0.0	28.9	30.6	**0.0**	–
qwen3.5:0.8b	82.0	1.0	–	80.5	0.1	58.1	80.5	**0.0**	–
qwen3.5:2b	94.8	**0.0**	–	93.1	0.1	87.5	93.6	**0.0**	–

Table 29: Citation benchmark overall results for small models, task variant: Title + first author + year. All metrics are reported on a 0–100 scale.

Model	BibTeX Parseability ↑	Strict Refusal ↓	Semantic Refusal ↓	Title F1 ↑	Author F1 ↑	Venue F1 ↑	Year ↑	DOI ↑	Judge ↑
deepseek-r1:1.5b	34.2	3.6	–	35.1	10.6	0.2	35.2	**0.0**	–
gemma3:1b	57.3	1.1	–	57.8	19.6	5.2	57.8	**0.0**	–
gemma3:270m	1.3	85.8	–	9.2	3.0	1.1	9.2	**0.0**	–
llama3.2:1b	22.7	1.5	–	18.3	8.8	0.0	22.1	**0.0**	–
llama3.2:3b	22.2	4.1	–	22.6	7.3	0.0	22.9	**0.0**	–
ministral-3:3b	52.0	47.9	–	**99.6**	**31.8**	**14.9**	**99.8**	**0.0**	–
phi4-mini:3.8b	14.7	81.0	–	73.0	22.3	5.4	75.0	**0.0**	–
qwen3.5:0.8b	82.7	0.3	–	81.7	25.6	2.8	81.3	**0.0**	–
qwen3.5:2b	**96.0**	**0.1**	–	95.7	30.5	5.5	95.4	**0.0**	–

Table 30: Citation benchmark author and venue metrics for small models by task variant. Task variant: Title only. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
deepseek-r1:1.5b	0.0 (0.2)	0.0 (0.1)	0.0 (0.5)	0.8 (6.0)	0.7 (5.3)	1.0 (7.3)
gemma3:1b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	4.5 (15.1)	4.3 (15.2)	5.1 (16.7)
gemma3:270m	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	3.0 (12.8)	2.3 (9.6)	4.5 (19.2)
llama3.2:1b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
llama3.2:3b	0.2 (3.0)	0.3 (4.4)	0.2 (2.7)	0.0 (0.5)	0.0 (0.3)	0.1 (1.6)
ministral-3:3b	**2.5 (10.2)**	**2.8 (11.5)**	**2.4 (9.6)**	11.6 (24.7)	9.9 (21.7)	14.5 (31.1)
phi4-mini:3.8b	0.1 (1.6)	0.1 (2.2)	0.1 (1.5)	3.1 (13.2)	2.8 (12.3)	3.8 (16.5)
qwen3.5:0.8b	0.2 (1.7)	0.2 (1.9)	0.2 (2.1)	5.6 (17.9)	5.4 (17.6)	6.5 (20.8)
qwen3.5:2b	0.3 (4.2)	0.4 (4.2)	0.4 (4.4)	**14.8 (30.5)**	**13.8 (29.2)**	**17.0 (34.5)**

Table 31: Citation benchmark author and venue metrics for small models by task variant. Task variant: Title only + DOI/URL requirement. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
deepseek-r1:1.5b	0.0 (0.6)	0.0 (0.5)	0.0 (0.8)	0.1 (2.1)	0.1 (2.0)	0.1 (2.4)
gemma3:1b	0.1 (1.2)	0.1 (1.1)	0.1 (1.2)	3.2 (13.5)	2.9 (12.8)	3.8 (15.0)
gemma3:270m	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.1 (2.6)	0.1 (2.3)	0.2 (3.5)
llama3.2:1b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.4 (3.8)	0.4 (4.2)	0.4 (3.7)
llama3.2:3b	0.3 (3.4)	0.3 (3.5)	0.3 (3.5)	7.9 (23.2)	7.0 (21.1)	9.4 (27.1)
ministral-3:3b	**0.9 (6.2)**	**1.1 (7.3)**	**0.9 (5.9)**	**24.4 (34.0)**	**20.8 (29.9)**	**31.0 (42.6)**
phi4-mini:3.8b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	2.9 (13.4)	2.1 (9.7)	4.8 (21.8)
qwen3.5:0.8b	0.1 (0.9)	0.1 (2.2)	0.1 (1.0)	11.0 (23.5)	10.0 (22.1)	13.4 (28.3)
qwen3.5:2b	0.2 (2.4)	0.2 (2.3)	0.3 (2.9)	16.8 (31.9)	15.5 (30.1)	19.2 (36.3)

Table 32: Citation benchmark author and venue metrics for small models by task variant. Task variant: Title + author hint. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
deepseek-r1:1.5b	3.7 (14.2)	5.2 (19.6)	3.6 (14.2)	1.2 (7.7)	1.1 (6.9)	1.6 (9.9)
gemma3:1b	3.6 (17.6)	5.0 (21.6)	3.4 (17.4)	4.8 (16.2)	4.5 (15.6)	5.5 (18.4)
gemma3:270m	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
llama3.2:1b	1.4 (10.3)	1.7 (11.5)	1.4 (10.2)	0.1 (3.3)	0.1 (2.9)	0.1 (3.7)
llama3.2:3b	6.5 (17.7)	9.7 (25.6)	5.8 (17.0)	1.2 (9.2)	1.1 (8.3)	1.4 (10.8)
ministral-3:3b	**31.6 (23.8)**	**60.1 (41.1)**	**27.3 (24.0)**	9.6 (24.2)	8.5 (22.5)	11.4 (28.0)
phi4-mini:3.8b	0.5 (3.7)	1.6 (12.8)	0.3 (2.1)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
qwen3.5:0.8b	19.0 (26.0)	30.1 (40.5)	18.3 (25.2)	9.6 (23.0)	8.9 (21.9)	11.3 (27.0)
qwen3.5:2b	22.1 (24.5)	30.8 (35.6)	21.7 (24.8)	**20.3 (33.9)**	**18.8 (32.3)**	**23.2 (38.4)**

Table 33: Citation benchmark author and venue metrics for small models by task variant. Task variant: Title + venue + year. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
deepseek-r1:1.5b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	14.5 (34.6)	14.1 (33.8)	15.1 (35.8)
gemma3:1b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	22.8 (42.0)	22.8 (42.0)	22.8 (42.0)
gemma3:270m	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.9 (9.4)	0.9 (9.4)	0.9 (9.4)
llama3.2:1b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	24.2 (42.6)	24.0 (42.4)	24.5 (43.0)
llama3.2:3b	0.1 (1.4)	0.1 (1.7)	0.1 (1.4)	24.3 (42.3)	23.9 (41.9)	25.1 (43.4)
ministral-3:3b	**0.9 (6.0)**	**1.0 (6.7)**	**0.9 (5.8)**	**91.1 (23.4)**	**88.8 (24.5)**	**94.3 (23.1)**
phi4-mini:3.8b	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	28.9 (44.2)	27.7 (42.9)	30.6 (46.6)
qwen3.5:0.8b	0.1 (1.6)	0.1 (1.5)	0.2 (2.1)	58.1 (47.3)	57.2 (47.2)	59.8 (48.3)
qwen3.5:2b	0.1 (1.8)	0.2 (2.5)	0.1 (1.6)	87.5 (30.5)	86.4 (31.1)	89.4 (30.3)

Table 34: Citation benchmark author and venue metrics for small models by task variant. Task variant: Title + first author + year. Values are shown as mean (std) on a 0–100 scale.

Model	Author F1 ↑	Author Precision ↑	Author Recall ↑	Venue F1 ↑	Venue Precision ↑	Venue Recall ↑
deepseek-r1:1.5b	10.6 (17.0)	25.1 (41.0)	8.4 (14.4)	0.2 (2.6)	0.1 (1.9)	0.3 (4.5)
gemma3:1b	19.6 (21.9)	49.7 (49.3)	12.9 (15.8)	5.2 (16.3)	4.9 (15.7)	6.0 (18.5)
gemma3:270m	3.0 (9.9)	9.2 (29.0)	1.8 (6.2)	1.1 (7.2)	1.0 (6.5)	1.3 (9.4)
llama3.2:1b	8.8 (17.8)	22.6 (41.8)	5.8 (12.7)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
llama3.2:3b	7.3 (15.2)	16.0 (34.4)	5.7 (12.2)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
ministral-3:3b	**31.8 (18.4)**	65.3 (38.5)	**25.2 (17.4)**	**14.9 (29.9)**	**13.2 (27.4)**	**18.2 (35.7)**
phi4-mini:3.8b	22.3 (18.4)	52.9 (45.7)	16.7 (14.6)	5.4 (17.5)	4.6 (15.4)	6.7 (21.8)
qwen3.5:0.8b	25.6 (19.1)	61.3 (45.2)	18.5 (14.5)	2.8 (13.2)	2.6 (12.6)	3.3 (15.6)
qwen3.5:2b	30.5 (18.6)	**74.2 (40.3)**	21.3 (14.7)	5.5 (19.8)	5.1 (18.7)	6.3 (22.1)

Table 35: Citation benchmark title metrics for small models by task variant. Task variant: Title only. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
deepseek-r1:1.5b	15.9 (36.2)	16.1 (36.5)	15.8 (36.3)	15.6 (36.1)	15.7 (36.3)	15.5 (36.1)
gemma3:1b	46.0 (49.6)	46.3 (49.8)	45.9 (49.5)	45.9 (49.5)	46.1 (49.8)	45.7 (49.4)
gemma3:270m	15.2 (36.4)	15.2 (36.4)	15.2 (36.4)	15.2 (36.4)	15.2 (36.4)	15.2 (36.4)
llama3.2:1b	57.1 (53.4)	57.1 (53.4)	57.1 (53.4)	57.1 (53.4)	57.1 (53.4)	57.1 (53.4)
llama3.2:3b	29.0 (45.3)	29.0 (45.3)	29.1 (45.4)	28.8 (45.3)	28.8 (45.3)	28.8 (45.3)
ministral-3:3b	**96.6 (17.8)**	**96.7 (17.7)**	**96.6 (17.8)**	**96.4 (18.6)**	**96.4 (18.6)**	**96.4 (18.6)**
phi4-mini:3.8b	78.5 (40.2)	78.7 (40.3)	78.6 (40.2)	77.6 (40.7)	77.8 (40.8)	77.6 (40.8)
qwen3.5:0.8b	79.3 (39.7)	79.2 (39.7)	79.6 (39.8)	78.9 (39.9)	78.7 (39.9)	79.2 (40.0)
qwen3.5:2b	94.9 (21.4)	94.8 (21.5)	95.1 (21.2)	94.4 (22.3)	94.4 (22.4)	94.7 (21.9)

Table 36: Citation benchmark title metrics for small models by task variant. Task variant: Title only + DOI/URL requirement. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
deepseek-r1:1.5b	2.6 (15.8)	2.7 (16.0)	2.6 (15.7)	2.6 (15.7)	2.6 (15.9)	2.5 (15.6)
gemma3:1b	40.9 (49.4)	40.9 (49.4)	40.9 (49.4)	40.9 (49.4)	40.9 (49.4)	40.9 (49.4)
gemma3:270m	4.9 (21.6)	4.9 (21.6)	4.9 (21.6)	4.9 (21.6)	4.9 (21.6)	4.9 (21.6)
llama3.2:1b	35.7 (47.8)	35.7 (47.8)	35.7 (47.9)	35.6 (47.8)	35.6 (47.8)	35.6 (47.8)
llama3.2:3b	51.1 (50.0)	51.1 (50.0)	51.1 (50.0)	50.9 (50.0)	51.0 (50.0)	50.9 (50.0)
ministral-3:3b	**97.9 (14.0)**	**98.0 (13.8)**	**97.9 (14.1)**	**97.8 (14.4)**	**97.9 (14.4)**	**97.8 (14.4)**
phi4-mini:3.8b	22.8 (41.9)	23.8 (43.6)	22.1 (41.1)	22.6 (41.8)	23.8 (43.6)	21.9 (40.9)
qwen3.5:0.8b	83.2 (36.8)	83.2 (36.9)	83.5 (36.8)	82.8 (37.2)	82.8 (37.2)	83.2 (37.1)
qwen3.5:2b	93.2 (24.5)	93.2 (24.6)	93.4 (24.4)	92.8 (25.3)	92.8 (25.4)	93.0 (25.1)

Table 37: Citation benchmark title metrics for small models by task variant. Task variant: Title + author hint. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
deepseek-r1:1.5b	14.5 (35.0)	14.4 (35.0)	14.5 (35.1)	14.4 (35.0)	14.4 (35.0)	14.5 (35.1)
gemma3:1b	39.5 (48.9)	39.6 (48.9)	39.5 (48.8)	39.4 (48.8)	39.4 (48.9)	39.3 (48.8)
gemma3:270m	20.0 (44.7)	20.0 (44.7)	20.0 (44.7)	20.0 (44.7)	20.0 (44.7)	20.0 (44.7)
llama3.2:1b	4.7 (20.9)	4.7 (21.0)	4.7 (20.9)	4.6 (20.9)	4.7 (21.0)	4.7 (20.9)
llama3.2:3b	20.4 (40.2)	20.5 (40.4)	20.4 (40.2)	20.3 (40.2)	20.5 (40.4)	20.3 (40.2)
ministral-3:3b	**98.4 (11.2)**	**98.4 (11.1)**	**98.4 (11.1)**	**97.9 (13.1)**	**98.0 (13.2)**	**98.0 (12.9)**
phi4-mini:3.8b	1.6 (12.8)	1.6 (12.8)	1.6 (12.8)	1.6 (12.8)	1.6 (12.8)	1.6 (12.8)
qwen3.5:0.8b	80.7 (38.4)	80.6 (38.5)	81.2 (38.5)	80.0 (38.8)	79.8 (38.9)	80.6 (38.8)
qwen3.5:2b	95.3 (19.9)	95.2 (20.2)	95.5 (19.6)	94.6 (21.7)	94.5 (21.9)	94.8 (21.4)

Table 38: Citation benchmark title metrics for small models by task variant. Task variant: Title + venue + year. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
deepseek-r1:1.5b	15.1 (35.6)	15.2 (35.7)	15.0 (35.5)	14.8 (35.4)	14.8 (35.4)	14.7 (35.3)
gemma3:1b	25.5 (43.5)	25.6 (43.6)	25.5 (43.4)	25.5 (43.5)	25.6 (43.6)	25.4 (43.4)
gemma3:270m	1.8 (13.3)	1.8 (13.3)	1.8 (13.3)	1.8 (13.3)	1.8 (13.3)	1.8 (13.3)
llama3.2:1b	86.5 (33.7)	86.4 (33.7)	86.7 (33.7)	86.4 (33.7)	86.3 (33.8)	86.7 (33.7)
llama3.2:3b	26.0 (43.7)	25.9 (43.7)	26.1 (43.8)	25.9 (43.7)	25.8 (43.7)	26.0 (43.8)
ministral-3:3b	**96.1 (18.2)**	**96.1 (18.2)**	**96.1 (18.2)**	**95.5 (20.0)**	**95.6 (20.0)**	**95.6 (19.9)**
phi4-mini:3.8b	30.9 (46.2)	31.1 (46.4)	30.8 (46.1)	30.7 (46.3)	30.9 (46.4)	30.6 (46.1)
qwen3.5:0.8b	80.5 (38.9)	80.5 (38.9)	80.7 (39.0)	79.9 (39.2)	79.9 (39.2)	80.1 (39.2)
qwen3.5:2b	93.1 (24.5)	93.1 (24.6)	93.2 (24.3)	92.8 (25.3)	92.7 (25.4)	92.9 (25.2)

Table 39: Citation benchmark title metrics for small models by task variant. Task variant: Title + first author + year. Values are shown as mean (std) on a 0–100 scale.

Model	Title Unigram F1 ↑	Title Unigram Precision ↑	Title Unigram Recall ↑	Title Bigram F1 ↑	Title Bigram Precision ↑	Title Bigram Recall ↑
deepseek-r1:1.5b	35.1 (47.7)	35.1 (47.7)	35.1 (47.7)	35.0 (47.6)	35.0 (47.6)	35.0 (47.6)
gemma3:1b	57.8 (49.4)	57.8 (49.4)	57.7 (49.3)	57.7 (49.4)	57.7 (49.4)	57.6 (49.4)
gemma3:270m	9.2 (29.0)	9.2 (29.0)	9.2 (29.0)	8.5 (28.0)	8.5 (28.0)	8.5 (28.0)
llama3.2:1b	18.3 (38.6)	18.3 (38.7)	18.3 (38.6)	18.2 (38.5)	18.2 (38.6)	18.1 (38.5)
llama3.2:3b	22.6 (41.8)	22.7 (41.8)	22.6 (41.8)	22.5 (41.7)	22.6 (41.8)	22.5 (41.7)
ministral-3:3b	**99.6 (6.2)**	**99.6 (6.2)**	**99.6 (6.2)**	**99.4 (7.6)**	**99.4 (7.6)**	**99.4 (7.6)**
phi4-mini:3.8b	73.0 (43.5)	73.1 (43.8)	73.8 (43.6)	72.5 (43.8)	72.9 (44.0)	73.5 (43.8)
qwen3.5:0.8b	81.7 (37.6)	81.4 (37.7)	82.3 (37.7)	81.1 (38.0)	80.8 (38.1)	81.7 (38.0)
qwen3.5:2b	95.7 (19.9)	95.7 (20.0)	95.7 (19.9)	95.5 (20.3)	95.5 (20.3)	95.5 (20.3)