

Relevant Is Not Warranted: Evidence-Force Calibration for Cited RAG

Pin Qian^{1,*} Su Wang^{1,*} Yihang Chen^{2,*} Wenxuan Xu³ Xiaoyuan Wang^{1,*}
Qiaolin Yu⁴ Shuhuai Lin¹ Sipeng Zhang⁵ Junxian You⁶ Xinpeng Wei^{2,†}

¹Carnegie Mellon University, Pittsburgh, PA, USA

²Georgia Institute of Technology, Atlanta, GA, USA

³Dartmouth College, Hanover, NH, USA ⁴Cornell University, Ithaca, NY, USA

⁵University of California San Diego, San Diego, CA, USA

⁶University of Glasgow, Glasgow, UK

pqian@alumni.cmu.edu william.xp.wei@gatech.edu

*Equal contribution. [†]Corresponding authors.

Abstract

Cited RAG evaluation often treats visible sources as a grounding signal, but a real, topically relevant citation can still under-warrant the attached wording. We study this diagnostic failure as *citation laundering*: a related source is presented as warrant for an over-strong claim. We introduce ForceBench, a contrastive stress test for evidence-force calibration. Each item holds a cited passage fixed and pairs an evidence-calibrated claim with a localized force-raised variant across five operational axes: relation, modality, scope, temporal validity, and numeric specificity. A calibrated evaluator should score the evidence-calibrated claim higher. Headline experiments use a fixed, locality-filtered 198-pair evaluation set. A citation-presence sanity check is uninformative by design; token and entity overlap still violate monotonicity on 32.8–36.4% of pairs. Across four hosted model judges, standard generic support prompting is insufficient for this force-calibration stress test (aggregate monotonicity violation rate (MVR) 47.2%), while explicit warrant-strength prompting lowers MVR to 24.5% but remains imperfect. Three smaller open-weight judges show the same direction, with aggregate MVR falling from 56.6% to 37.4%. We release the benchmark, prompts, outputs, and plug-in pipeline so citation evaluators can report MVR and force sensitivity alongside conventional support metrics.

1 Introduction

Retrieval-augmented generation connects language generation to external evidence and provenance (Lewis et al., 2020). In cited RAG, citations have become a visible grounding interface: browser-assisted QA, quote-supported QA, and citation-augmented generation systems attach references or supporting evidence to generated answers, inviting users to treat cited claims as grounded (Nakano

et al., 2021; Menick et al., 2022; Gao et al., 2023; Rashkin et al., 2023). This presentation can compress a hard verification problem into a visible cue: this sentence has a source. The cue is useful, but incomplete. A citation can be present, relevant, and still too weak, narrow, stale, or imprecise for the sentence it is asked to support.

Consider an answer that says a drug exposure *causes* an adverse event while the cited source reports only an observational association and explicitly notes that causality could not be established. Association-to-causation shifts are a known form of scientific overstatement (Sumner et al., 2014); in high-stakes health settings, the distinction licenses different downstream inferences. The citation is real. The entities match. A citation-presence metric passes it, and many retrieval or overlap checks see the right topic. The evidence licenses a weaker claim: association, not causation. The generated answer has increased the *force* of the cited evidence.

The same pattern appears outside medical causality. In our adjudicated examples, one cited passage says that men *may* have been more susceptible to a historical plague; the answer cites it to state that men *were* more susceptible. A public-health passage says that Ebola’s R_0 typically sits around 1.5 to 2.0; the answer cites it for the point claim that Ebola’s R_0 is 2.0. A clinical passage reports a trial in children with convulsive status epilepticus; the answer cites it for a claim about status epilepticus patients generally. In each case, the citation is related enough to appear plausible, but the claim asks it to carry too much evidential force.

We call this failure **citation laundering**, using the term narrowly for cases where a related citation makes an over-strong claim appear warranted. The citation exists, is recoverable, and is sufficiently related to pass inspection, yet it is too weak for the claim’s relation, modality, scope, temporal status,




Figure 1: **ForceBench evaluates citation laundering through contrastive monotonicity.** The cited evidence is held fixed while the benchmark contrasts a strongest warranted claim with a minimally stronger overclaim. The same local test is instantiated across five force axes: relation, modality, scope, temporal validity, and numeric specificity. A calibrated evaluator should assign higher support to the warranted claim than to the overclaim.

or numeric precision. The failure mode is that surface relevance can mask insufficient warrant.

Relation to hallucination. Citation laundering is a factuality failure in the claim-evidence relation. It differs from factuality and hallucination benchmarks that evaluate false, unsupported, or unverifiable model outputs more broadly (Lin et al., 2022; Li et al., 2023): here the source exists, is recoverable, and is topically related, while the displayed citation under-warrants the force of the claim. A detector that only checks whether each claim has some cited source will pass every such claim by construction. We therefore frame citation laundering at the claim-citation pair level.

This paper introduces **evidence-force calibration** for cited RAG. Given cited evidence e and claim c , we ask whether the force of c is no greater than the force licensed by e . We operationalize this with a contrastive monotonicity test (Figure 1): if the same evidence supports an evidence-calibrated claim c_{safe} , then a localized force-raised variant c_{raised} should receive lower support:

$$score(e, c_{\text{safe}}) > score(e, c_{\text{raised}}). \quad (1)$$

If a judge scores “may help” and “will help” equally against evidence that only says “may,” it loses the evidence-force ordering even if it recognizes the topic.

This makes the evaluation target operational. A useful cited-generation system should preserve citation-warranted wording such as “may help”

while rejecting stronger wording such as “will ensure,” because ranking, filtering, and repair need calibrated claims alongside unsupported-claim warnings.

ForceBench is deliberately a compact diagnostic stress test, not a survey of naturally occurring RAG outputs. Its candidate selection targets visible warrant boundaries, so the reported error rates characterize evaluator behavior on these controlled contrasts; they do not estimate how prevalent citation laundering is in deployed systems. Likewise, per-axis results are descriptive diagnostics for this dataset rather than population-level rankings of which force gaps are generally hardest.

Our contributions are:

1. We formalize citation laundering as a mismatch between claim force and evidence force, distinct from missing citations, irrelevant retrieval, and ordinary unsupported generation.
2. We introduce ForceBench, a diagnostic contrastive benchmark of local force contrasts across relation, modality, scope, temporal, and numeric axes.
3. We define monotonicity violation rate and force sensitivity for evaluating whether citation evaluators respect local increases in claim force.
4. We construct an adjudicated benchmark and report headline results on a fixed locality-filtered evaluation set.
5. We evaluate deterministic baselines, four hosted

judges, and three smaller open-weight judges, finding that generic citation-support prompting is insufficient for this stress test; explicit warrant-strength prompting helps but remains imperfect.

2 Evidence-Force Calibration

Attribution and citation evaluation typically ask whether cited evidence supports a claim or answer (Rashkin et al., 2023; Gao et al., 2023; Yue et al., 2023; Xu et al., 2025). ForceBench asks a stricter question: does the evidence warrant the claim at the same level of relation, modality, scope, temporal validity, and numeric specificity?

We model claim force as a structured categorical tuple:

$$\text{claim_force}(c) = \{\text{relation, modality, scope, temporal, numeric}\}. \quad (2)$$

Evidence force is the warrant licensed by the cited passage, including stated limitations and context. A claim is force-calibrated if:

$$\text{claim_force}(c) \leq \text{evidence_force}(e), \quad (3)$$

where $\leq$ is an evidence-relative partial order: the claim does not exceed the cited passage on any included axis. Many claims are incomparable. ForceBench only constructs pairs that preserve the same core proposition and differ by a localized force increase on one primary axis.

Axes. The five axes are operational and intentionally non-exhaustive. We select dimensions that satisfy a practical test: holding the cited evidence fixed, a claim can remain topically related while exceeding what the citation warrants through a local increase in relation strength, certainty, scope, temporal status, or numeric specificity. The axes draw on adjacent evaluation lines without treating any one taxonomy as exhaustive. Attribution and citation evaluation motivate the local support question (Rashkin et al., 2023; Xu et al., 2025). Atomic and long-form factuality motivate claim-level granularity (Min et al., 2023; Wei et al., 2024). Hedge and uncertainty work informs the modality axis (Farkas et al., 2010; Yona et al., 2024). Temporal annotation and event factuality motivate temporal contrasts (Saurí and Pustejovsky, 2009; Pustejovsky et al., 2003); fact-verification benchmarks motivate evidential-status contrasts (Thorne et al., 2018; Aly et al., 2021). Scientific overstatement

work motivates relation-strength contrasts such as association-to-causation (Sumner et al., 2014). We exclude dimensions such as source authority, source freshness as a property of the citation itself, citation placement, and aggregation across multiple sources because they require answer-level or source-set judgments beyond a local single-citation force test. ForceBench makes the included axes citation-specific by asking whether the displayed source licenses the stronger wording.

ForceBench focuses on five axes:

- **Relation:** association, mention, or risk becomes causation, prevention, proof, or obligation.
- **Modality:** possible, conditional, preliminary, or suggestive evidence becomes definite, necessary, or guaranteed.
- **Scope:** a claim licensed for a subgroup, jurisdiction, study population, version, or institution becomes general.
- **Temporal:** dated, predicted, or as-of evidence becomes a current or timeless claim.
- **Numeric:** approximate, ranged, or bounded evidence becomes an exact endpoint or point estimate.

Boundary. Citation laundering covers the subset of citation errors where a related source underwarrants a localized force increase. Fabricated quotations, wrong-entity claims, missing links, and broad unsupported additions are important failures with different diagnostic structure. ForceBench isolates the local property that support should decrease when claim force increases while the cited evidence remains fixed.

3 ForceBench

ForceBench uses one cited passage and one atomic claim as its unit. Each item contains the cited passage, a decisive evidence span, an evidence-calibrated claim, a localized force-raised claim, a primary force axis, severity, and a repair. The calibrated claim is wording that annotators judge to be licensed by the cited passage. The force-raised claim preserves the same local proposition but exceeds the evidence on one primary axis.

3.1 Construction

ForceBench is built through a single local-warrant annotation protocol. We start from cited claim-evidence contexts sampled from existing cited

QA and attribution resources. One source is AttributionBench-derived rows (Li et al., 2024), spanning ExpertQA (Malaviya et al., 2024), LFQA-style long-form QA (Chen et al., 2024), and Stanford-GenSearch-style verifiability settings (Liu et al., 2023). We also include contexts from AttributedQA (Bohnet et al., 2023), GaRAGe (Sorodoc et al., 2025), and additional ExpertQA examples. The benchmark is a stress test, with sampling targeted at contexts whose cited passages expose local boundaries in relation, modality, scope, temporal status, or numeric specificity.

For each selected context, we draft a candidate row containing the cited passage, a decisive evidence span, an evidence-calibrated claim, a localized force-raised claim, a repair, a rationale, and proposed axis and severity metadata. These proposed fields are pre-annotation suggestions, not benchmark labels. Every candidate row is independently annotated by two research assistants under the same local-warrant guideline. Their reviews are reconciled through adjudication, including disagreements over row acceptance, evidence spans, claim wording, repair, axis, or severity. Rows that pass adjudication then undergo quality control of those fields; only rows passing both adjudication and quality control are eligible for inclusion.

The protocol asks annotators to answer one local question: what wording does this displayed citation license? Annotators identify the decisive evidence span, verify an evidence-calibrated claim, test a stronger variant against the same passage, assign one primary force axis when the contrast is local, and write a repair that restores the citation-warranted wording. The central guideline is to ignore external truth: a claim may be true in the world and still be a ForceBench force gap if the displayed citation does not warrant it.

We use *localized* as an operational criterion rather than a formal edit-distance guarantee. A pair is admissible only if the two claims preserve the same cited passage, answer context, main entity, event, and topical relation, while differing primarily in one force dimension. Candidates are rejected if the calibrated claim is not supported, the stronger claim is already supported, the contrast changes entities or events, combines multiple force shifts, introduces an unrelated unsupported fact, requires outside knowledge, or lacks a clear evidence span.

Axis	Retained	Eval	S1	S2	S3
Relation	48	40	0	38	2
Modality	48	39	5	33	1
Scope	51	40	1	34	5
Temporal	39	39	10	27	2
Numeric	43	40	16	23	1
Total	229	198	32	155	11

Table 1: **ForceBench data summary.** Retained counts are adjudicated pairs kept for the benchmark pool after capping 52 surplus relation candidates and removing two nonlocal rows; Eval is the fixed locality-filtered evaluation set. S1–S3 are severity counts in Eval.

3.2 Quality Control

Force gaps are close to their positive counterparts, so we use two independent annotations for every candidate row, followed by adjudication, evidence-span checks, locality checks, and provenance/duplicate checks. On pre-adjudication row-level decisions, the two research assistants achieve Cohen’s $\kappa = 0.78$. The audit asks whether the cited evidence warrants the calibrated formulation while under-warranting the stronger formulation.

Final promotion is deliberately conservative. Of 433 double-annotated candidates, 283 remain after adjudication and quality control. Capping 52 surplus relation-axis rows and removing two nonlocal rows in a final audit yield a 229-pair retained pool; a stratified locality-filtered subset of 198 pairs is the fixed headline evaluation set (Table 1). A row is retained only if the evidence span is decisive, the calibrated claim is supported, the stronger claim is under-warranted, and the contrast is a local force increase with no ordinary unsupported addition. *Locality* is an operational, adjudicated semantic judgment rather than a formal edit-distance threshold. When forming retained and evaluation sets, we stratify by force axis and source provenance and use a stable hash for tie-breaking, separating candidate collection from the metric denominator. Appendix B gives the detailed accounting, split/filter rule, and annotation statistics; Appendix C gives the human annotation guidelines.

4 Metrics

For an evaluator that assigns support scores, a monotonicity violation occurs when the force-raised claim receives support greater than or equal to the evidence-calibrated claim. We define mono-

tonicity violation rate:

$$\text{MVR} = \frac{1}{N} \sum_i \mathbf{1}[s_i^{\text{raised}} \geq s_i^{\text{safe}}]. \quad (4)$$

Here $s_i^{\text{safe}} = \text{score}(e_i, c_{\text{safe},i})$ and $s_i^{\text{raised}} = \text{score}(e_i, c_{\text{raised},i})$. Lower is better. We also report force sensitivity:

$$\text{FS} = \frac{1}{N} \sum_i (s_i^{\text{safe}} - s_i^{\text{raised}}). \quad (5)$$

Higher is better. MVR captures whether the evaluator preserves strict ordering; FS captures the average margin. The strict inequality is deliberate: a tie between the evidence-calibrated claim and the force-raised variant gives no preference to the warranted wording for ranking, filtering, or repair. Both metrics are therefore necessary: a model may increase margins while still assigning both claims the same discrete label.

Categorical judges are mapped to scores. For generic support prompts, `fully_supported` receives 1, `partially_supported` receives 0.5, and `unsupported`, `contradicted`, or `irrelevant` labels receive 0. For force-aware prompts, `force_calibrated` receives 1, while `force-gap`, `contradicted`, or `irrelevant` labels receive 0. If a model returns an unparseable response or refuses after retry for either member of a pair, we count that pair as a monotonicity violation with margin 0. This conservative policy keeps denominators fixed and treats evaluator non-compliance as a failure to provide an actionable preference. Because the generic and force-aware rubrics use different categorical score mappings, their FS values are prompt-conditioned margins and should not be compared as if they were on one absolute scale. MVR directly tests the within-rubric pair ordering; changes in FS provide supporting evidence about separation under each rubric.

5 Experiments

We evaluate whether common citation signals and strong LLM judges preserve force monotonicity.

5.1 Baselines

We report citation presence, token overlap, topical entity overlap, a rule-based ForceCheck, an axis-aware lexical ForceCheck prototype, and four hosted model judges. The axis-aware prototype starts from token overlap and applies penalties




Figure 2: **Main MVR-FS tradeoff on the 198-pair evaluation set.** Lower MVR and higher FS are better. Circles denote deterministic baselines, squares denote the four-model aggregate, and the arrow connects the generic and force-aware prompt conditions. The dashed line marks the 0.5 random-order MVR reference.

when the claim removes limiters visible in the evidence, such as changing “may” to “will,” a range to a point estimate, or a dated statement to a current claim. It is an interpretable diagnostic baseline for testing whether visible force markers matter. Its limitations include paraphrase, implicit scope, and context-dependent evidence boundaries.

Model judge selection. We evaluate four high-capability LLM judges: Claude Sonnet 4.6, GPT-5.5, GLM 5.1, and Qwen 3.6 Flash. The panel is intentionally limited. It targets capable deployed judges and asks whether they preserve calibrated-over-force-raised ordering under fixed evidence when prompted with a standard support question. Using multiple model families gives a limited cross-interface check of the observed pattern.

These four models were evaluated through API calls recorded in the run manifests. Each JSON-only prompt is applied to both the evidence-calibrated claim and the force-raised claim for every pair on the same 198-pair locality-filtered evaluation set. Each item is submitted as a fresh deterministic, stateless request with tool access disabled and a 1024-token output cap. Full model identifiers, API settings, run dates, parser settings, and retry policy are provided in Appendix G.

This evaluator suite targets the paper’s narrow local-ordering claim under fixed evidence. MVR and FS are plug-in stress-test metrics: any attribution or citation evaluator can be run on the same contrastive pairs. The reported suite spans the main

Evaluator	Prompt	MVR	FS
<i>Deterministic baselines</i>			
Citation presence	–	1.000	0.000
Token overlap	–	0.328	0.066
Entity overlap	–	0.364	0.066
Rule ForceCheck	–	0.298	0.156
Axis ForceCheck	–	0.232	0.248
<i>Reported model judges</i>			
Model aggregate	Generic support	0.472	0.333
Model aggregate	Force-aware	0.245	0.754

Table 2: **Main MVR and FS on the 198-pair evaluation set.** Lower MVR and higher FS are better. Model rows report aggregate results across the four hosted judges. Parse/refusal pairs are counted as monotonicity violations with zero margin.

signals needed for this diagnostic: citation appearance, topical overlap, visible force-shift heuristics, generic support judging, and explicit force-aware judging. Broader learned or task-specific citation evaluators remain important future plug-ins.

5.2 Main Results

Figure 2 plots monotonicity violation against force sensitivity. Citation presence is a sanity check: it never distinguishes the paired claims and therefore has 100% MVR by construction. Token and entity overlap test whether topical similarity is enough to preserve monotonicity; both remain low-sensitivity despite using more information than citation presence. The axis-aware lexical prototype is the strongest deterministic baseline (MVR 0.232), but it is hand-aligned with visible force markers and still leaves substantial failures, especially when force is implicit.

Table 2 summarizes deterministic baselines and aggregate model-judge results for the two reported model prompts. The pooled model-judge result is the central stress test. With a generic support prompt, the four-judge panel has MVR 0.472 (374/792), and force sensitivity is 0.333. Under an explicitly force-aware prompt, MVR drops to 0.245 (194/792), and force sensitivity increases to 0.754. This paired comparison should be read as a prompt-conditioned diagnostic, because the generic and force-aware prompts ask different questions. It shows that ordinary support wording is insufficient for this force-calibration stress test, and that making warrant strength explicit improves ordering on the same model-pair units.

Model	Generic		Force-aware	
	MVR (%)	FS	MVR (%)	FS
Qwen3.5-4B	53.0	0.26	34.8	0.52
Llama-3.2-3B-Instruct	60.6	0.22	41.9	0.48
Gemma-3-4B-IT	56.1	0.24	35.4	0.55
Aggregate	56.6	0.240	37.4	0.517

Table 3: **Open-weight judge results.** Lower MVR and higher FS are better. All values use the same 198-pair evaluation set; the aggregate pools the three judges. FS is prompt-conditioned because the two rubrics use different label mappings.




Figure 3: **Per-model prompt-rubric MVR.** Each model has four bars: generic support, axis-free force, axis-list force-aware, and dummy-axis control. From left to right, models are Claude Sonnet 4.6, GPT-5.5, GLM 5.1, and Qwen 3.6 Flash. Lower MVR is better.

5.3 Open-Weight Judges

To broaden coverage, Table 3 reports three smaller open-weight judges evaluated on the same 198-pair set under the generic and force-aware prompt conditions. The force-aware condition lowers MVR and raises FS for every model, but residual MVR remains 34.8–41.9%. We report these models separately from the four-model headline aggregate so the original comparison remains auditable.

5.4 Prompt-Rubric Ablation

The reported force-aware prompt names the same five force axes used to construct ForceBench, so we test whether axis names drive the result. Table 4 compares two controls: a minimal axis-free force prompt that keeps the warrant-strength instruction while omitting the benchmark taxonomy, and a dummy-axis control that uses irrelevant rubric dimensions. The axis-free prompt reduces aggregate MVR from 47.2% to 24.5%, while the axis-list prompt also yields 24.5% MVR and raises FS to 0.754. The dummy-axis control rises to

Prompt	MVR (%)	FS
Generic support	47.2	0.333
Minimal axis-free force	24.5	0.753
Axis-list force-aware	24.5	0.754
Dummy-axis control	50.9	0.480

Table 4: **Aggregate prompt-rubric ablation.** Lower MVR and higher FS are better. The axis-free prompt omits the benchmark axis taxonomy; the dummy-axis control uses irrelevant rubric dimensions.




Figure 4: **Descriptive axis-level MVR on the 198-pair evaluation set.** Darker and warmer cells indicate higher violation rates. Deterministic rows use 39–40 pairs per axis; the aggregate generic-support model row uses 156–160 model-pair units per axis. These selected stress-test slices are diagnostic, not population estimates.

50.9%. The prompt gain therefore tracks explicit warrant-strength framing. We interpret it as prompt-conditioned task framing for this stress test.

Figure 3 shows the per-model prompt-rubric pattern. Both force prompts improve over generic support prompting for all four judges, while the dummy-axis control often worsens MVR. Table 4 reports the corresponding aggregate MVR and FS values. We interpret this as a prompt-conditioned result for this four-judge panel.

5.5 Axis Behavior

Axis behavior is reported descriptively in Figure 4. On this selected set, relation shifts are comparatively visible because an association-to-causation change is often lexical; the aggregate generic-support row has its lowest MVR on relation. Numeric, temporal, and modality shifts often require attending to ranges, dates, hedges, or stated uncertainty rather than matching entities. Because each axis contains only 39–40 pairs and sampling targets visible warrant boundaries, we do not infer a general difficulty ranking from these differences.

6 Discussion

Warrant calibration. The central lesson concerns warrant calibration. A citation evaluator that rejects both variants may be safer than one that accepts both, yet it still cannot recover citation-warranted wording. The force-aware prompt improves ordering overall while still leaving residual violations. Evidence-force calibration asks for a precise capability: keep the warranted claim and reject only the extra force.

Residual prompt-conditioned errors. In this fixed evaluation set, with these prompts and single-run model calls, the aggregate force-aware MVR remains 0.245; the open-weight judges likewise retain aggregate MVR 0.374. Prompting therefore exposes and partially improves the target behavior but does not solve it. A next step is evaluator-side supervision or fine-tuning on a training set separated from ForceBench evaluation pairs, together with training data that directly contrasts calibrated and force-raised wording; analogous fine-grained preference optimization separates grounding and reasoning signals in vision-language training (Shen et al., 2025). Architectures that represent evidence limitations and compare claim force to those limitations may offer a stronger inductive bias than a single absolute support label. Broader post-training work studies entropy interventions for reinforcement learning and experience-driven agent adaptation (Hao et al., 2026a,b). ForceBench offers a held-out diagnostic for testing whether such analogous mechanisms improve warrant-sensitive evaluation without training on its evaluation pairs.

Limits of the diagnostic heuristic. The axis-aware ForceCheck result tests whether visible limiters are informative. The heuristic is hand-aligned with the benchmark axes and depends on explicit lexical cues; it performs well on relation and modality cases with overt markers, and still violates monotonicity on 35.0% of numeric and 43.6% of temporal examples. This residual pattern is important: many force gaps involve a relationship between the wording of the claim and the limitations of the cited passage.

Why contrastive pairs matter. Absolute support labels are noisy because annotators and models differ in how much context they require. The contrastive design holds evidence fixed and asks only for an ordering. Observed cited-output examples motivate the failure pattern, while source-

grounded local contrasts remove retrieval quality, answer style, and external truth as explanations. This makes the claim crisp: a citation evaluator should give lower support to a stronger claim when the extra force exceeds what the source licenses.

A layered citation audit. Citation laundering complements, rather than replaces, fabricated-citation detection. A practical pipeline should first verify that a citation exists and resolves to the intended source, then check relevance and provenance, and finally test whether that source warrants the attached claim as written. Conflict-aware RAG can separate semantic relevance from factual consistency (Ye et al., 2026). Complementary agent-memory work studies cascade repair after faulty state is diagnosed (Zhao et al., 2026). ForceBench isolates the narrower claim-level warrant check for a displayed source. The last stage is necessary precisely because a real, relevant citation can still be asked to support excessive force.

Beyond single-source detection. Future work should extend the strongest-warrant question to multi-source and multi-hop answers, including sets of passages that are individually partial, jointly sufficient, or mutually conflicting. Such evaluations should test whether provenance and warrant are preserved when several pieces of evidence must be composed. Retrieval augmentation also extends to streaming cross-modal generation, as in simultaneous speech translation where partial speech retrieves terminology hints; extending warrant evaluation to this setting would require aligning retrieved-evidence use with the incremental outputs it informs (Luo et al., 2026). Coverage should also expand across domains and languages, where evidential conventions and force markers differ. Finally, the human-written repair field permits a generation task: given evidence and a force-raised claim, produce calibrated wording and evaluate whether the repair preserves the useful proposition while removing only unsupported force.

7 Related Work

ForceBench is complementary to attribution and citation evaluation. AIS defines attribution as verification against identified sources (Rashkin et al., 2023); ALCE evaluates citation-augmented long-form QA along correctness and citation quality dimensions (Gao et al., 2023); and automatic attribution evaluators such as AttrScore investigate

LLM and finetuned support judgments for cited claims (Yue et al., 2023). AttributionBench-style data (Li et al., 2024), broad RAG diagnostics such as RAGChecker (Ru et al., 2024), and citation-specific evaluators such as CiteEval (Xu et al., 2025) provide broader citation-quality testbeds. CiteEval is especially close in motivation because it makes citation assessment more fine-grained; ForceBench fixes the cited passage and measures a local monotonicity property under a controlled force increase. Citation-granularity studies also show that attribution behavior and system comparison depend on how evidence is segmented (Wang et al., 2026). Those studies vary the citation unit; ForceBench instead fixes one passage and tests whether it warrants a controlled claim-force change. Citation-context and citation-intent work studies how scientific papers use prior work (Cohan et al., 2019); ForceBench asks whether a displayed citation warrants the force of a generated claim. Claim-level factuality and hallucination benchmarks ask whether generated content is truthful, hallucinated, or factually precise, including TruthfulQA, HaluEval, FActScore, and SAFE/Long-form factuality (Lin et al., 2022; Li et al., 2023; Min et al., 2023; Wei et al., 2024). NLI and fact-verification datasets formalize support as entailment, contradiction, or neutrality (Bowman et al., 2015; Thorne et al., 2018; Aly et al., 2021), and hedge/uncertainty work motivates force distinctions (Farkas et al., 2010; Yona et al., 2024). These lines target absolute truth, support classification, or uncertainty expression. ForceBench conditions on a real, recoverable, topically related citation and asks whether support is monotonic when only claim force increases. Appendix A summarizes this boundary with adjacent evaluation lines.

The closest conceptual line is scientific overstatement, especially health-science exaggeration analysis (Sumner et al., 2014). ForceBench transfers that concern to generated cited answers, where an inline citation may make a claim appear warranted. The unit is the claim-citation pair in RAG answers. The label target is whether the displayed citation licenses the stronger wording. The metric is therefore contrastive monotonicity under fixed cited evidence.

8 Conclusion

ForceBench targets a missing layer in cited RAG evaluation: evidence warrant. A citation can be real

and relevant while still failing to license the force of the attached claim. By holding evidence fixed and applying a localized force increase, ForceBench turns this failure into a direct monotonicity test. On the fixed locality-filtered 198-pair evaluation set, overlap baselines violate monotonicity on 32.8–36.4% of pairs. Explicit warrant-strength prompting reduces MVR from 47.2% to 24.5% for the four hosted judges and from 56.6% to 37.4% for the three open-weight judges, but residual failures remain. Practitioners should treat citation presence as a minimal sanity check and supplement conventional support metrics with MVR and FS on warrant-sensitive stress tests. The remaining errors motivate evaluator-side training and supervision.

Limitations

ForceBench is a compact diagnostic benchmark. Its MVR numbers are stress-test error rates under fixed evidence, with prevalence estimation outside scope. The examples are selected for visible local warrant boundaries, so prompt gains reflect this targeted evaluation set. The force-aware and axis-free prompts explicitly define the warrant-strength target; their gains show the effect of prompt-conditioned task framing in this stress test. Larger releases should broaden domain, source, and language coverage. Future work should also extend evidence-force calibration beyond text-only claim-passage pairs to multimodal cited outputs, where citation warrant may depend jointly on visual evidence, textual claims, and domain-specific interpretation. Open-vocabulary segmentation provides a localization setting in which textual concepts must align with foreground and background visual evidence (Yao et al., 2026). Representative future settings include image–text capacity imbalance and localized video-artifact judgments (Xu et al., 2026; Fu et al., 2025), as well as multimodal physical-dynamics modeling and vision-foundation-model adaptation (Li et al., 2026a,b); together, they illustrate support regimes spanning textual constraints and visual evidence. Locality is an operational annotation judgment rather than a formal edit-distance guarantee; a final audit removed two nonlocal rows, and the release reports those IDs and reasons, but borderline cases may remain. The evaluator suite covers citation presence, lexical/entity overlap, rule/axis ForceCheck prototypes, four hosted judges, and three open-weight judges; CiteEval-style, AutoAIS-style, and learned

attribution baselines remain useful broader coverage (Xu et al., 2025; Gao et al., 2023; Yue et al., 2023). Model-judge results are single-run measurements under the prompts reported here; model or interface drift, decoding changes, and prompt-target alignment can change absolute and relative MVR values. Repair fields exist for accepted items, while repair success is outside the main evaluation. ForceBench isolates single English claim-citation pairs; coverage excludes multi-hop synthesis, aggregation across multiple sources, source authority, citation recency as a source property, citation placement, and answer-level usefulness.

Ethical Considerations

ForceBench is an evaluation resource for identifying cases where cited evidence under-warrants the force of a generated claim. Use of the benchmark should stay within citation-warrant evaluation, with population prevalence outside scope. Some examples involve medical, legal, or policy-sensitive content; labels concern citation warrant rather than real-world truth or advice. Released rows preserve source provenance and use citation snippets for research evaluation; users should respect the terms of the source collections when redistributing passages. The research assistants judged only local citation warrant and were instructed to avoid medical, legal, or policy advice. Data release should preserve annotation decisions, rejected-boundary labels, and the distinction between final adjudicated labels and diagnostic artifacts.

References

- Rami Aly, Zhijiang Guo, Michael Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. [FEVEROUS: Fact extraction and verification over unstructured and structured information](#). In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*.
- Bernd Bohnet, Vinh Q. Tran, Pat Verga, Roei Aharoni, Daniel Andor, Livio Baldini Soares, Massimiliano Ciaramita, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, Kai Hui, Tom Kwiatkowski, Ji Ma, Jianmo Ni, Lierni Sestorain Saralegui, Tal Schuster, William W. Cohen, Michael Collins, Dipanjan Das, and 3 others. 2023. [Attributed question answering: Evaluation and modeling for attributed large language models](#). *Preprint*, arXiv:2212.08037.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#).

- In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Hung-Ting Chen, Fangyuan Xu, Shane Arora, and Eunsol Choi. 2024. [Understanding retrieval augmentation for long-form question answering](#). In *Proceedings of the 1st Conference on Language Modeling*. COLM 2024.
- Arman Cohan, Waleed Ammar, Madeleine van Zuylen, and Field Cady. 2019. [Structural scaffolds for citation intent classification in scientific publications](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3586–3596, Minneapolis, Minnesota. Association for Computational Linguistics.
- Richárd Farkas, Veronika Vincze, György Móra, János Csirik, and György Szarvas. 2010. [The CoNLL-2010 shared task: Learning to detect hedges and their scope in natural language text](#). In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning – Shared Task*, pages 1–12, Uppsala, Sweden. Association for Computational Linguistics.
- Xingyu Fu, Siyi Liu, Yinuo Xu, Pan Lu, Guangqiuse Hu, Tianbo Yang, Taran Anantasagar, Christopher Shen, Yikai Mao, Yuanzhe Liu, Keyush Shah, Chung Un Lee, Yejin Choi, James Zou, Dan Roth, and Chris Callison-Burch. 2025. Learning human-perceived fakeness in AI-generated videos via multimodal LLMs. *arXiv preprint arXiv:2509.22646*.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488, Singapore. Association for Computational Linguistics.
- Zhezhen Hao, Hong Wang, Haoyang Liu, Jian Luo, Jiarui Yu, Hande Dong, Qiang Lin, Can Wang, and Jiawei Chen. 2026a. [Rethinking entropy interventions in RLVR: An entropy change perspective](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31105–31133, San Diego, California, United States. Association for Computational Linguistics.
- Zhezhen Hao, Hong Wang, Jian Luo, Jianqing Zhang, Yuyan Zhou, Qiang Lin, Can Wang, Hande Dong, and Jiawei Chen. 2026b. [ReCreate: Reasoning and creating domain agents driven by experience](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31018–31046, San Diego, California, United States. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33*, pages 9459–9474.
- Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. [HaluEval: A large-scale hallucination evaluation benchmark for large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6449–6464, Singapore. Association for Computational Linguistics.
- Yifei Li, Xiang Yue, Zeyi Liao, and Huan Sun. 2024. [AttributionBench: How hard is automatic attribution evaluation?](#) In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14919–14935, Bangkok, Thailand. Association for Computational Linguistics.
- Yuqi Li, Junhao Dong, Jiao Liu, Piotr Koniusz, Han-sheng Zeng, Chuanguang Yang, Junming Liu, Yingli Tian, Tingwen Huang, and Hao Wu. 2026a. [Evolving multimodal models for physical dynamics: A multi-objective neuroevolution approach](#). *IEEE Transactions on Evolutionary Computation*.
- Yuqi Li, Xi Xiao, Yunbei Zhang, Lin Zhao, Yu Li, Aiden Zhao, Tianyang Wang, Hao Xu, and Yingli Tian. 2026b. Rethinking layer-wise information allocation for vision foundation model adaptation. *arXiv preprint arXiv:2607.21973*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring how models mimic human falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Nelson Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7001–7025, Singapore. Association for Computational Linguistics.
- Jiaxuan Luo, Siqi Ouyang, Jiaxing Xu, and Lei Li. 2026. RASST: Retrieval-augmented simultaneous speech translation. *arXiv preprint arXiv:2601.22777*.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2024. [ExpertQA: Expert-curated questions and attributed answers](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3025–3045, Mexico City, Mexico. Association for Computational Linguistics.
- Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick,

- Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, and Nat McAleese. 2022. [Teaching language models to support answers with verified quotes](#). *arXiv preprint arXiv:2203.11147*.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FACTScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021. [WebGPT: Browser-assisted question-answering with human feedback](#). *arXiv preprint arXiv:2112.09332*.
- James Pustejovsky, José Castaño, Robert Ingria, Roser Saurí, Robert Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir Radev. 2003. [TimeML: Robust specification of event and temporal expressions in text](#). In *New Directions in Question Answering: Papers from the 2003 AAAI Spring Symposium*. AAAI Press.
- Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2023. [Measuring attribution in natural language generation models](#). *Computational Linguistics*, 49(4):777–840.
- Dongyu Ru, Lin Qiu, Xiangkun Hu, Tianhang Zhang, Peng Shi, Shuaichen Chang, Cheng Jiayang, Cunxiang Wang, Shichao Sun, Huanyu Li, Zizhao Zhang, Binjie Wang, Jiarong Jiang, Tong He, Zhiguo Wang, Pengfei Liu, Yue Zhang, and Zheng Zhang. 2024. [RAGChecker: A fine-grained framework for diagnosing retrieval-augmented generation](#). In *Advances in Neural Information Processing Systems 37*.
- Roser Saurí and James Pustejovsky. 2009. [FactBank: A corpus annotated with event factuality](#). *Language Resources and Evaluation*, 43(3):227–268.
- Yifan Shen, Yuanzhe Liu, Jingyuan Zhu, Xu Cao, Xiaofeng Zhang, Yixiao He, Wenming Ye, James Rehg, and Ismini Lourentzou. 2025. [Fine-grained preference optimization improves spatial reasoning in VLMs](#). In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pages 17929–17960. Curran Associates, Inc.
- Ionut Teodor Sorodoc, Leonardo F. R. Ribeiro, Rexhina Blloshmi, Christopher Davis, and Adrià de Gispert. 2025. [GaRAGE: A benchmark with grounding annotations for RAG evaluation](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 17030–17049, Vienna, Austria. Association for Computational Linguistics.
- Petroc Sumner, Solveiga Vivian-Griffiths, Jacky Boivin, Andy Williams, Christos A. Venetis, Aimée Davies, Jack Ogden, Leanne Whelan, Bethan Hughes, Bethan Dalton, Fred Boy, and Christopher D. Chambers. 2014. [The association between exaggeration in health related science news and academic press releases: Retrospective observational study](#). *BMJ*, 349:g7015.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: A large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Hexuan Wang, Jingyu Zhang, Benjamin Van Durme, and Daniel Khashabi. 2026. Are finer citations always better? rethinking granularity for attributed generation. *arXiv preprint arXiv:2604.01432*.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#). In *Advances in Neural Information Processing Systems 37*.
- Junhao Xu, Zheng Liu, Mengqi Zhang, Guangyuan Dong, and Zhumin Chen. 2026. [Lever can move the earth: Towards adaptive semantic capacity balance for image-text retrieval](#). *IEEE Transactions on Multimedia*.
- Yumo Xu, Peng Qi, Jifan Chen, Kunlun Liu, Rujun Han, Lan Liu, Bonan Min, Vittorio Castelli, Arshit Gupta, and Zhiguo Wang. 2025. [CiteEval: Principle-driven citation evaluation for source attribution](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 32759–32778, Vienna, Austria. Association for Computational Linguistics.
- Zhen Yao, Xin Li, Taotao Jing, Shuai Zhang, and Mooi Choo Chuah. 2026. Disa: Saliency-aware foreground-background disentangled framework for open-vocabulary semantic segmentation. *arXiv preprint arXiv:2601.20064*.
- Hua Ye, Siyuan Chen, Ziqi Zhong, Canran Xiao, Hao-liang Zhang, Yuhuan Wu, and Fei Shen. 2026. [Seeing through the conflict: Transparent knowledge conflict handling in retrieval-augmented generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(40):34423–34431.
- Gal Yona, Roei Aharoni, and Mor Geva. 2024. [Can large language models faithfully express their intrinsic uncertainty in words?](#) In *Proceedings of the 2024*

Conference on Empirical Methods in Natural Language Processing, pages 7752–7764, Miami, Florida, USA. Association for Computational Linguistics.

Xiang Yue, Boshi Wang, Ziru Chen, Kai Zhang, Yu Su, and Huan Sun. 2023. [Automatic evaluation of attribution by large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4615–4635, Singapore. Association for Computational Linguistics.

Yang Zhao, Chengxiao Dai, Mengying Kou, and Yue Xiu. 2026. MEMOREPAIR: Barrier-first cascade repair in agentic memory. *arXiv preprint arXiv:2605.07242*.

A Adjacent Evaluation Lines

Attribution. Attribution work asks whether generated claims are supported by identified sources (Rashkin et al., 2023; Yue et al., 2023); ForceBench holds the source fixed and tests whether support decreases when only claim force becomes stronger.

RAG diagnostics. RAG diagnostics evaluate retrieval quality, answer correctness, and citation quality across full outputs (Gao et al., 2023; Ru et al., 2024; Xu et al., 2025); ForceBench isolates a local failure that can survive relevant retrieval and plausible citation placement.

Source audits. Source-audit settings check whether a citation is accessible, topically relevant, or factually consistent with the cited content; ForceBench assumes the passage is available, then asks whether the attached claim is calibrated to that passage’s warrant.

Overstatement. Scientific-overstatement work studies exaggeration or evidential proportionality relative to source evidence (Sumner et al., 2014); ForceBench transfers that concern to generated cited answers and operationalizes it as local contrastive pairs under one fixed citation.

B Dataset and Annotation

Sampling and candidate drafting. Sampling is designed to build stress-test examples, with prevalence estimation outside scope. We sample cited claim-evidence contexts from existing cited QA and attribution resources and keep contexts only when the displayed passage contains a visible local boundary in relation, modality, scope, temporal status, or numeric specificity. Candidate drafting then normalizes each selected context into the same row format: question or answer context, cited passage, decisive evidence span, evidence-calibrated claim, force-raised claim, repair, rationale, and proposed axis/severity metadata. Pre-annotation drafts are excluded before they become candidate rows if the passage is unavailable, the proposed contrast requires outside knowledge, the claims change the entity or answer target, or the contrast is an ordinary unsupported addition rather than a local force increase. All candidate rows that enter the annotation pool are double-annotated.

Candidate accounting. The benchmark is the result of one candidate-drafting, double-annotation, and adjudication workflow under a single local-warrant protocol. The same annotation pass is used regardless of source or date: every candidate row is independently reviewed by two research assistants before adjudication. In total, 433 candidate rows enter the annotation pool and 283 are accepted after adjudication and quality control. We then cap surplus accepted rows from overrepresented axes before benchmark reporting, select

Counted set	Rows
Candidate rows entering double annotation	433
Accepted after adjudication and quality control	283
Headline evaluation set before final locality filter	200
Final headline evaluation set after locality filter	198

Table 5: **Construction flow.** All candidate rows use the same two-annotator workflow. The final row is the denominator for headline model and baseline results.

the headline evaluation set by stratifying over force axis and source provenance, and apply a final locality filter. Table 5 gives the compact accounting for the annotation pool and headline denominator.

Human review statistics. Candidate suggestions become benchmark labels only after human review. All 433 candidate rows that enter the annotation pool are independently annotated by two research assistants and then adjudicated before release; 283 are accepted and 150 are rejected. On pre-adjudication row-level decision labels, the annotators achieve Cohen’s $\kappa = 0.78$.

Every row in the final benchmark is therefore backed by two independent local-warrant judgments and final adjudication. Disagreements over accept/reject status, evidence-span length, safe-claim wording, overclaim wording, repair, primary axis, or severity are resolved before export. We treat axis and severity as adjudicated metadata: the axis is used for analysis and splitting, while severity is released for diagnosis; MVR and FS use only pair ordering.

C Annotation Guidelines

The annotation guideline given to the research assistants defines the task as local citation-warrant judgment. Annotators are told to use only the displayed cited passage and to ignore whether a claim might be true according to outside knowledge.

Annotation unit. Each row contains a question or answer context, a cited passage, a proposed evidence-calibrated claim, a proposed force-raised claim, and candidate metadata. Annotators may correct the evidence span, calibrated claim, force-raised claim, repair, primary axis, and severity in the `annotation_*` fields, or reject the row.

Annotation interface. Annotators see the question or answer context and the cited passage, which is the only evidence allowed. They also see pre-annotation suggestions for the evidence span, calibrated claim, force-raised claim, repair, rationale, axis, and severity. Annotators write the release-facing fields in `annotation_*`: row decision (`accept_force_gap`, `reject_not_force_gap`, or `needs_discussion`), corrected evidence span, calibrated claim, force-raised claim, repair, primary axis, severity, boundary type, and notes. Final labels come from these human corrections and adjudication.

Guideline: Decision Procedure

Annotators follow the same sequence for every candidate:

1. Read the cited passage and mark the shortest span that determines the support boundary.
2. Decide whether the calibrated claim is directly warranted by that span using only the displayed passage.
3. Decide whether the force-raised claim preserves the

same entity, event, source, and answer context.

4. Decide whether the force-raised claim exceeds the passage's warrant on one primary axis.
5. Assign relation, modality, scope, temporal, or numeric as the primary axis; mark severity; and write a repair that restores citation-warranted wording.

Guideline: Accept Criteria

A row is accepted only when the calibrated claim is supported, the force-raised claim is under-warranted by the cited passage, and the contrast is local: it keeps the same core proposition while increasing force along one dominant axis. The annotator must be able to explain the contrast in the form: the evidence licenses X, but the force-raised claim states stronger Y.

Guideline: Reject Criteria

Annotators reject rows when the calibrated claim is unsupported; the force-raised claim is actually supported; the pair changes the main entity, event, time, or answer target; the row combines multiple force shifts with no clear primary axis; the stronger claim adds an unrelated unsupported fact; the judgment depends on outside knowledge; or the cited passage lacks a decisive evidence span.

Decision labels. Annotators choose one row-level decision. `accept_force_gap` means the citation supports the calibrated claim, the stronger claim exceeds that same citation, and the contrast is local. `reject_not_force_gap` means the row is unsupported, already supported, nonlocal, wrong-entity/topic, or otherwise outside scope. `needs_discussion` is reserved for genuinely unclear boundaries and is resolved during adjudication before release.

Axis definitions. Relation errors turn association, mention, risk, or correlation into causation, proof, prevention, or obligation. Modality errors turn possibility, uncertainty, suggestion, or conditional evidence into certainty or necessity. Scope errors generalize from a bounded population, product, jurisdiction, study, source, or version to a broader class. Temporal errors turn dated, future, predicted, or as-of evidence into current or timeless claims. Numeric errors turn approximate, ranged, bounded, or comparative quantities into exact point claims.

Severity and repair. Severity 1 marks a subtle wording shift that can still mislead an evaluator; severity 2 marks a clear force mismatch; severity 3 marks a high-impact or strongly misleading mismatch. Repairs must preserve only what the cited passage warrants, usually by restoring a qualifier, scope restriction, time marker, approximate quantity, or weaker relation.

D Benchmark Examples

Each example below shows one retained benchmark row after the final locality filter. The cited evidence is fixed; the evidence-calibrated claim is licensed by that evidence, while the force-raised claim asks the citation to support stronger force than it states.

Relation. Evidence cue: “associated with up to a 40 per cent risk.” Calibrated claim: The use of sodium valproate during pregnancy is associated with up to a 40 percent risk

of neurodevelopmental disorders and a 10 percent risk of physical disabilities for an unborn child. Force-raised claim: The use of sodium valproate during pregnancy results in up to a 40 percent risk of neurodevelopmental disorders and a 10 percent risk of physical disabilities for an unborn child.

Modality. Evidence cue: “evidence that men may have been more susceptible.” Calibrated claim: Some evidence suggests men may have been more susceptible to the 1361–1362 plague. Force-raised claim: Men were more susceptible to the 1361–1362 plague.

Scope. Evidence cue: “convulsive status epilepticus in children.” Calibrated claim: A trial found levetiracetam was not superior to phenytoin as a second-line agent for convulsive status epilepticus in children. Force-raised claim: A trial found levetiracetam was not superior to phenytoin as a second-line agent for convulsive status epilepticus patients.

Temporal. Evidence cue: “As of 2019, New Jersey has nine casinos.” Calibrated claim: As of 2019, New Jersey had nine casinos, all in Atlantic City. Force-raised claim: There are currently nine casinos in Atlantic City, New Jersey.

Numeric. Evidence cue: “typically sits around 1.5 to 2.0.” Calibrated claim: Ebola's R0 typically sits around 1.5 to 2.0. Force-raised claim: Ebola's R0 is 2.0.

E Boundary Accounting

The main text treats rejected rows as boundary evidence that explains exclusions. These rows are important because they show that the guideline separates local force gaps from ordinary unsupported claims, supported stronger claims, and candidate-construction artifacts.

Candidate boundary checks. Candidate rows are screened for local admissibility before release. Rejection categories include unsupported calibrated claims, supported force-raised claims, reversed or nonlocal relation rewrites, duplicate rows, safe or underclaimed rows, ordinary unsupported additions, ambiguous force gaps, and evidence spans that are not decisive.

Independent annotation checks. All candidate rows that enter the annotation pool are independently double-annotated before adjudication; pre-adjudication row-level agreement is Cohen's $\kappa = 0.78$. Accept/reject disagreements are resolved before inclusion. Axis checks are computed only after both annotators accept a row, so they are treated as conditional consistency checks over accepted rows.

These checks document the boundary between local citation-force gaps and other citation or candidate-construction failures; prevalence estimation remains outside scope.

F Judge Prompts

The exact machine-readable prompt files are released with the benchmark. The templates below show the prompts used for the reported model-judge results.

Generic Support Judge

System. “You are evaluating whether a cited passage supports a claim. Use only the cited passage. Do not use

outside knowledge.”
User. question, cited passage, and claim, followed by: “Does the cited passage support the claim? Choose exactly one label: `fully_supported`, `partially_supported`, `unsupported`, `contradicted`, `irrelevant`. Return JSON only with fields `label`, `confidence`, and `rationale`.”

Force-Aware Judge
System. “You are evaluating evidence-force calibration. Use only the cited passage. Do not use outside knowledge. A claim is force-calibrated only if the cited evidence warrants the claim at the same or stronger level of relation, modality, scope, temporal validity, and numeric specificity.”
User. question, cited passage, and claim, followed by: “Identify the strongest claim warranted by the passage; decide whether the given claim exceeds the passage’s warranted force; if it exceeds the evidence, identify the primary force-gap axis. Choose exactly one label: `force_calibrated`, `force_gap`, `contradicted`, `irrelevant`. If `force_gap`, choose one primary axis: relation, modality, scope, temporal, numeric. Return JSON only with fields `label`, `primary_axis`, `severity`, `warranted_claim`, and `rationale`.”

Axis-Free Force Ablation
System. “You are evaluating whether a cited passage warrants a claim as written. Use only the cited passage. Do not use outside knowledge.”
User. question, cited passage, and claim, followed by: “Decide whether the cited passage licenses the claim exactly as worded. Mark `force_gap` if the claim is related to the passage but says more than the passage licenses. Mark `force_calibrated` only if the cited passage licenses the full strength of the claim as written. Do not use a predefined taxonomy. Return JSON only with fields `label`, `primary_axis`, `severity`, `warranted_claim`, and `rationale`.”

Dummy-Axis Rubric Control
System. “You are a citation judge. Use only the cited passage. Do not use outside knowledge.”
User. question, cited passage, and claim, followed by: “Decide whether the cited passage supports the claim. When checking support, focus on source authority, writing clarity, citation placement, topical completeness, and fluency. Choose exactly one label: `force_calibrated`, `force_gap`, `contradicted`, `irrelevant`. Return JSON only with fields `label`, `primary_axis`, `severity`, `warranted_claim`, and `rationale`.”

G Hosted Model and Decoding Configuration

All reported hosted-model API calls were executed on May 17, 2026 UTC. The four hosted judges were called through their API endpoints. In the same order as the main text and Figure 3, the exact model identifiers were `claude-sonnet-4-6`, `gpt-5.5`, `glm-5.1`, and `qwen3.6-flash`.
For reproducibility, every claim judgment was submitted as a fresh stateless API request with deterministic decoding, JSON-only output instructions, and a 1024-token output cap. GPT-5.5 used the Responses-style `max_output_tokens=1024` setting; the other API calls used `max_tokens=1024`. For chat-style APIs, deterministic decoding used `temperature=0.0`; for provider interfaces with a named deterministic mode, we used that mode. The same prompt templates, parser, scoring script, retry policy, and failure-handling rules were used for all models and prompts. Run manifests record the API target, exact model identifier, prompt file, request count, output cap, decoding settings, run timestamp, parser version, retry counts, and sanitized output paths.