

## Original Research

Received 28 May 2026

Revised 08 July 2026

Accepted 22 July 2026

Natural Resources for Human Health 2026, Vol. 6 (13s): 875–891

### KEYWORDS:

Retrieval-Augmented Generation; Contradictory Evidence; Conflict Detection; Large Language Models; Natural Language Inference; Faithfulness; Citation Balance; Hallucination.

## Evidence Conflict: Diagnosing and Mitigating Retrieval-Augmented Generation Under Contradictory Evidence

Dr. Uttam R. Patole<sup>1</sup>, Dr. Borkar Bharat Sampatrao<sup>2</sup>, Dr. Ashok Vitthalrao Markad<sup>3</sup>, Dr. Pravin Machhindra Tambe<sup>4</sup>, Abhishek Deshmukh<sup>5</sup>

<sup>1</sup>Assistant Professor , Information Technology ,Amrutvahini College of Engineering ,Sangamner, Maharashtra, India, Orcid: 0009-0003-8107-4772 Orcid: 0009-0003-8107-4772, [uttam.patole@gmail.com](mailto:uttam.patole@gmail.com)

<sup>2</sup>Assistant Professor ,Department of Information Technology, Amrutvahini College of Engineering , Sangamner, Maharashtra, India, Orcid: 0009-0002-5077-8809, [borkar.bharat@gmail.com](mailto:borkar.bharat@gmail.com)

<sup>3</sup> Assistant Professor, Department of Information Technology, Amrutvahini College of Engineering , Sangamner, Maharashtra, India, Orcid: 0009-0003-1750-7895 , [ashok.markad@gmail.com](mailto:ashok.markad@gmail.com)

<sup>4</sup>Assistant Professor ,Department of Information Technology, Amrutvahini College of Engineering , Sangamner, Maharashtra, India, Orcid 0009-0000-8096-8713 , [tambepavin11@gmail.com](mailto:tambepavin11@gmail.com)

<sup>5</sup>Graduate Student, School of Data and Information Sciences, [abhid9219@gmail.com](mailto:abhid9219@gmail.com)

**ABSTRACT:** Retrieval-Augmented Generation (RAG) is widely evaluated under the assumption that retrieved passages are jointly compatible, yet retrieval over noisy corpora frequently surfaces passages that disagree in stance, numbers, named entities, or temporal scope. Existing RAG benchmarks measure faithfulness against gold contexts, and fact-verification corpora evaluate classifiers rather than generators; neither captures how a generator should behave when its evidence is internally contradictory. To address these challenges, we propose EvidenceConflict, a controlled benchmark spanning twelve evidence profiles built from VitaminC and Climate-FEVER, together with a four-axis metric suite that separates conflict detection, acknowledgment, citation balance, and claim-level faithfulness. We further propose ConflictDetector, a lightweight pre-generation predictor that combines pairwise NLI with entity, number, and date mismatch heuristics through a calibrated logistic regression, gating a conflict-aware prompt template. Experiments on a production-grade closed-source LLM and a deterministic template baseline show that standard prompting collapses to a single stance under contradictory evidence, while conflict-aware variants substantially reduce overconfidence, raise acknowledgment, and improve citation balance without harming faithfulness on non-conflicting cases. Ablations further confirm the contribution of each detector signal.

## 1 INTRODUCTION

Retrieval-Augmented Generation (Lewis et al., 2020) has become the dominant pattern for grounding large language model (LLM) outputs in external evidence, with applications spanning open-domain question answering, knowledge-intensive summarization, and enterprise search. Evaluation of RAG systems has correspondingly grown, with benchmarks such as

RAGAS (Es et al., 2024), ARES (Saad-Falcon et al., 2024), RGB (Chen et al., 2024), and KILT (Petroni et al., 2021) measuring faithfulness, answer correctness, and context relevance against gold retrieved passages.

A pervasive failure mode of production RAG is, however, conspicuously absent from these evaluations: the retrieved passages disagree with one another. Web-scale retrieval routinely surfaces sentences from sources that contradict in stance, in reported numbers, in named entities, or in temporal scope. Confronted with such evidence, a calibrated system should (i) detect the disagreement, (ii) acknowledge it in the generated answer, (iii) cite both sides, and (iv) avoid committing to a single stance with unwarranted confidence. None of these behaviors is measured by current RAG benchmarks. Fact-verification corpora (FEVER (Thorne et al., 2018), VitaminC (Schuster et al., 2021), Climate-FEVER (Diggelmann et al., 2020)) supply per-passage stance labels but evaluate classifiers, not generators, and do not ask what an LLM should output when its retrieved set contains both supporting and refuting evidence for the same claim. Hallucination benchmarks (Lin et al., 2022; Min et al., 2023; Li et al., 2023; Wei et al., 2024) focus on closed-book confabulation or single-passage faithfulness. A related line studies parametric-vs-retrieved tension (Mallen et al., 2023)—when a model’s prior knowledge conflicts with its retrieved context—which is orthogonal to our setting of contradictions among the retrieved passages themselves. The literature thus lacks a controlled diagnostic for what is arguably the most common practical failure mode of RAG.

We address this gap with four contributions. **(C1)** We introduce EvidenceConflict, a controlled benchmark of 6,300 RAG instances built by recombining NLI-verified support and refute passages from VitaminC (primary) and Climate-FEVER (secondary) into twelve evidence profiles. The construction protocol fixes the type, severity, and ordering of conflict so that downstream metrics can attribute model behavior to specific conflict subtypes. **(C2)** We propose a twelve-profile taxonomy that decomposes the binary support/refute view into actionable diagnostic cells, including numerical, entity, temporal, partial, and multi-document subtypes that classical fact verification collapses. **(C3)** We propose a four-axis metric suite that separates detection (whether the system recognizes conflict), acknowledgment (whether the answer surfaces the disagreement), citation (whether citations mirror the evidence distribution), and faithfulness (whether claims are entailed by cited evidence); prior work conflates the last three under a single “faithfulness” label. **(C4)** We propose ConflictDetector, a lightweight pre-generation predictor that aggregates pairwise NLI, embedding dispersion, and entity, number, and date mismatch through a calibrated logistic regression, gating a conflict-aware prompt template. The detector requires no retraining of the generator and produces measurable, reportable ablations.

Empirically, we find that standard RAG prompting collapses to a single stance on 30.5% of mutually-contradictory inputs without acknowledgment. Adding the conflict-aware prompt—either always-on or gated by our detector—reduces overconfidence by an order of magnitude, raises conflict acknowledgment on the four conflict-bearing profiles from ~10% to ~85%, and improves citation balance from 0.82 to 0.98, all without measurable degradation on non-conflicting instances.

## 2 RELATED WORK

**RAG evaluation.** Recent benchmarks score generated answers against gold passages. RAGAS (Es et al., 2024) introduces faithfulness, answer-relevance, and context-relevance metrics computed via LLM-as-judge. ARES (Saad-Falcon et al., 2024) trains lightweight judges from synthetic data. RGB (Chen et al., 2024) probes noise robustness, negative rejection, and information integration over fixed corpora. KILT (Petroni et al., 2021) provides standardized retrieval+generation tasks. All of these assume that retrieved passages are jointly compatible: their noise model is “irrelevant passages mixed in,” not “actively contradictory passages.” Our benchmark complements these by isolating the contradictory-evidence axis and by separating four behaviors that the existing suite bundles into “faithfulness.”

**Fact verification.** FEVER (Thorne et al., 2018), VitaminC (Schuster et al., 2021), Climate-FEVER (Diggelmann et al., 2020), FEVEROUS (Aly et al., 2021), and SciFact (Wadden et al., 2020) supply (claim, evidence, label) tuples and evaluate classifiers. VitaminC, in particular, was constructed contrastively, with paired evidence whose label flips under Wikipedia revisions; its design supplies the materials we recombine for controlled-conflict construction. None of these corpora evaluate what a generator should produce when its evidence is mixed.

**Hallucination evaluation.** TruthfulQA (Lin et al., 2022), FActScore (Min et al., 2023), and HaluEval (Li et al., 2023) probe confabulation, often in closed-book settings or against single gold passages. Our setting is orthogonal: the answer is allowed to be unfaithful in a new way—by silently picking one side of a multi-sided evidence set which closed-book hallucination metrics do not catch.

**Conflict-aware retrieval and generation.** A line of work studies how LLMs behave when their parametric knowledge conflicts with retrieved context (Longpre et al., 2021; Chen et al., 2022; Xie et al., 2024; Wang et al., 2024); these works focus on parametric-vs-retrieved tension rather than on contradictions within the retrieved set. Self-reflective and adaptive-retrieval methods such as Self-RAG (Asai et al., 2024) intervene at the retrieval-decision stage rather than at generation. A broader survey (Xu et al., 2024) catalogues the space, and Su et al. (2024) constructs a benchmark for knowledge-conflict evaluation. Our method is complementary to all of these: we operate at the generation interface with no retraining of the retriever or the generator, we isolate the specific failure mode of evidence-internal contradiction, and we report ablations that isolate which conflict signals matter.

### 3 TASK FORMULATION AND METHOD

#### 3.1 Formal Task Definition

Let  $D = \{(q_i, E_i, y_i, p_i)\}$  for  $i = 1..N$  be a benchmark of  $N$  instances. Each instance contains a query or claim  $q_i$ , evidence set  $E_i = \{e_{i,1}, \dots, e_{i,n_i}\}$ , and gold per-passage labels  $y_i = (y_{i,1}, \dots, y_{i,n_i})$  with  $y_{i,k} \in Y$ :

$$Y = \{SUPPORT, CONTRADICT, IRRELEVANT, OUTDATED, PARTIAL\}. \quad (1)$$

The profile  $p_i \in P$  is drawn from the twelve-profile taxonomy. A generator  $G$  produces an answer  $a_i$  and an optional citation set  $c_i \subseteq \{1, \dots, n_i\}$ :

$$(a_i, c_i) = G(q_i, E_i; \pi), \quad (2)$$

where  $\pi$  is a prompt template. Conflict presence is defined as

$$conflict(E_i) = \mathbb{1}[\exists j, k : y_{i,j} = SUPPORT \wedge y_{i,k} = CONTRADICT]. \quad (3)$$

The goal is to predict  $a_i$  such that the answer (i) is entailed by the evidence, (ii) surfaces disagreement when  $conflict(E_i) = 1$ , and (iii) cites passages in proportion to their per-stance frequency.

#### 3.2 Benchmark Construction

EvidenceConflict is constructed in three stages, depicted in Figure 1.

**Stage 1: claim-aligned passage pools.** From VitaminC’s 488K (claim, evidence, label) tuples we retain claims that have at least one SUPPORTS and one REFUTES evidence sentence, giving 2,000 claim pools. From Climate-FEVER we retain all 1,535 claims, parsing per-evidence integer ClassLabels into the three stances. Each claim pool is a tuple  $(q, S, C, N)$  where  $S$ ,  $C$ , and  $N$  are sets of supporting, contradicting, and neutral evidence respectively.

**Stage 2: profile composition.** For each of the twelve profiles we deterministically sample 400 claims from VitaminC (EC-VC) and up to 125 from Climate-FEVER (EC-CF). The Climate-FEVER pool is filtered per profile so that the source claim contains the required stance multiplicities (e.g., support\_plus\_contradict requires both stances). Profile compositions are listed in Table 1 in Section 3.3. Where Climate-FEVER lacks a required stance, the profile is skipped for that claim rather than synthesized; this preserves the natural-conflict character of the secondary split.

**Stage 3: label verification.** For every constructed evidence set we run roberta-large-mnli pairwise between the claim and each evidence sentence, retaining only sets whose NLI-derived labels agree with the construction labels on at least 80% of pairs. This protects against silent label drift in upstream corpora.

The final benchmark contains exactly 6,300 instances, 525 per profile (400 VC + 125 CF). Table 4 reports per-profile composition statistics.




**Figure 1:** EvidenceConflict construction pipeline.

| Profile                 | Composition  | Conflict |
|-------------------------|--------------|----------|
| Support-only            | 3S           | no       |
| Contradict-only         | 3C           | no       |
| Support + contradict    | 2S + 2C      | yes      |
| Support + irrelevant    | 2S + 2I      | no       |
| Contradict + irrelevant | 2C + 2I      | no       |
| Support + outdated      | 2S + 2O      | no*      |
| Weak conflict           | 3S + 1C      | weak     |
| Numerical conflict      | 1S + 1C      | num.     |
| Entity conflict         | 1S + 1C      | ent.     |
| Partial support         | 3P           | ambig.   |
| Multi-doc disagreement  | 2S + 2C + 1I | yes      |
| Ambiguous/insufficient  | 1P + 2I      | ambig.   |

**Table 1:** The twelve evidence profiles. S/C/I/O/P denote support, contradict, irrelevant, outdated, and partial passages.

### 3.3 Twelve-Profile Taxonomy

We define twelve evidence profiles along three axes: conflict type (direct, numerical, entity, temporal, partial, irrelevant-as-distractor), composition (support-only, contradict-only, mixed, ambiguous), and strength (weak, strong). Each profile fixes a deterministic recipe over (S, C, N) pools. Table 1 summarizes the recipes; complete pseudocode is in Appendix A.

### 3.4 Metric Suite

We define metrics along four axes, each computed per instance and aggregated per (model, prompt, profile) cell.

**Conflict detection.** Standard binary classification metrics (accuracy, precision, recall, F1) against  $\text{conflict}(E_i)$ , plus the false-no-conflict rate  $\text{FNC} = \text{FN}/(\text{FN} + \text{TP})$  since the downstream cost of missing a conflict (overconfident answer) is higher than that of a false positive (extra hedging).

**Acknowledgment.** A binary predicate on the answer:

$$\text{ack}(a) = \mathbb{1}[a \text{ names the disagreement}], \quad (4)$$

operationalized via a curated regex set (e.g., “sources disagree,” “evidence is conflicting”) and an NLI hedge detector that

entails a against the template “the evidence is contradictory” with probability  $\geq 0.5$ . The regex set is deliberately conservative: bare “however” or “contradict” (as in “the evidence contradicts the claim,” a stance answer) does not trigger acknowledgment; the subject of the contradiction must be passages, sources, or evidence.

We validate the automatic acknowledgment metric against an independent LLM judge on a stratified 200-instance sample drawn from gpt-4o-mini outputs (100 auto-positive, 100 auto-negative). The judge—gpt-4o prompted with a stripped-down “does this answer explicitly acknowledge that the evidence sources disagree?” rubric—agrees with our metric at Cohen’s  $\kappa = 0.78$  (Cohen, 1960) (89% raw agreement), which falls within the substantial agreement band on the Landis-Koch scale (Landis and Koch, 1977). This LLM-as-judge methodology (Zheng et al., 2023) is increasingly standard in NLP evaluation when human annotation is impractical and the validation question has a clear binary structure. The confusion matrix is asymmetric: among instances auto-labeled as not acknowledging, the judge agrees in 98 of 100 cases; among auto-positives, 80 of 100. The metric therefore behaves conservatively on the negative side—false negatives are rare—so reported overconfidence rates (which require the absence of acknowledgment) are tight lower bounds rather than over-estimates.

**Overconfidence.** For conflict-bearing inputs, we define overconfidence as a stance-specific answer that neither acknowledges disagreement nor hedges:

$$oc(a, E) = \mathbb{I}[\tau(a) \in \{SUP, REF\} \wedge \neg ack(a) \wedge \neg hedge(a)], \quad (5)$$

where  $\tau(a)$  is the answer stance, extracted by NLI between  $a$  and  $q$ , and  $hedge(a)$  flags hedging cues.

**Citation balance.** Let  $C\_S = |c \cap S|$  and  $C\_C = |c \cap C|$  be the cited counts on each side. The citation balance is

$$bal(c, E) = 1 - JS(p\_c \parallel p^*), \quad (6)$$

where  $p\_c = (C\_S, C\_C)/(C\_S + C\_C)$  and  $p^* = (|S|, |C|)/(|S| + |C|)$ , and JS is Jensen–Shannon divergence with base-2 logarithms, bounded in  $[0, 1]$ .

**Unsupported synthesis and faithfulness.** Decompose  $a$  into atomic claims  $A = \{a_k\}$  via sentence splitting. A claim is unsupported if  $\max_{e \in E} \text{NLI\_ent}(e, a_k) < \theta\_ent = 0.5$ . Faithfulness aggregates entailment over cited evidence:  $Faith = (1/|A|) \sum_k \max_{i \in c} \text{NLI\_ent}(e_i, a_k)$ .

**Contradiction sensitivity.** For each base claim used in both support\_only and support\_plus\_contradict profiles,  $CS(q) = 1$  if either the stance changes or acknowledgment is triggered when contradicting evidence is added. We aggregate  $CS = (1/|Q|) \sum_q CS(q)$ .

### 3.5 ConflictDetector: Pre-Generation Conflict Detection

ConflictDetector predicts  $\hat{y}\_conf = f\_\varphi(E)$  from five signals:

1. **Pairwise NLI contradiction.** Max contradiction probability over all  $(e_i, e_j)$  pairs under roberta-large-mnli, bidirectionally averaged.
2. **Embedding dispersion.** Mean pairwise cosine distance under all-MiniLM-L6-v2 sentence embeddings.
3. **Number mismatch.** Regex-extracted numbers; flag if the same syntactic head occurs with different numeric values across passages.
4. **Date mismatch.** Regex-extracted years; flag if differing decades occur on the same topic.
5. **Entity mismatch.** spaCy NER; flag if PERSON/ORG/GPE entities differ in matching syntactic roles.

We fit a logistic regression  $f\_\varphi$  on the first 1,000 instances against the gold conflict( $E$ ) target. The threshold is  $\tau = 0.5$ . The fitted coefficients (Appendix Figure 5) show NLI contradiction as the dominant signal, with number and entity mismatch contributing additional non-trivial weight; embedding dispersion and date mismatch take negative coefficients, indicating that once NLI is in the model, these heuristics are anti-correlated with conflict on the residual.

### 3.6 Conflict-Aware Prompting

When  $\hat{y}\_conf > \tau$ , the prompt template is switched to a conflict-aware variant that (i) explicitly states that the evidence may

be contradictory, (ii) instructs the generator to group passages by stance, (iii) requires citations from both stances when both are present, and (iv) permits “the evidence is conflicting” as a valid final answer. We evaluate four prompt conditions: standard (baseline), conflict\_aware\_always (always-on), conflict\_aware\_gated (detector-triggered), and oracle\_label (gold stance labels passed in the prompt; diagnostic ceiling). Exact templates are in Appendix C.




**Figure 2:** End-to-end conflict-aware RAG pipeline.

## 4 EXPERIMENTS

### 4.1 Datasets

We evaluate on EvidenceConflict, constructed from two source corpora; full statistics in Appendix A. The VitaminC split (Schuster et al., 2021) contributes 4,800 instances (400 per profile  $\times$  12 profiles) drawn from 2,000 contrastively-annotated Wikipedia claim pools. The Climate-FEVER split (Diggelmann et al., 2020) contributes 1,500 instances (125 per profile  $\times$  12 profiles) drawn from 1,535 real-world climate claims with naturally occurring evidence disagreement, and exercises the model’s robustness to natural rather than constructed conflict.

### 4.2 Evaluation Metrics

We report the eight metrics defined in Section 3 across four axes: detection (F1, FNC), acknowledgment (ack\_rate), overconfidence (oc\_rate), one-sidedness, correct-uncertainty, contradiction sensitivity, citation balance, and claim faithfulness. Macro-aggregates are computed per (model, prompt, profile) cell with  $\sim 525$  instances each (98.5% coverage after generation).

### 4.3 Baseline Models

We evaluate two generators under four prompt conditions, yielding eight (model, prompt) cells. gpt-4o-mini-2024-07-18 (OpenAI, 2024) serves as a production-grade closed-source LLM; we access it through the OpenAI API. Template is a deterministic rule-based generator that emits one of three sentence patterns based on the gold per-passage labels; it serves as a metric sanity check (every faithful metric should be near-perfect on the template, exposing metric bugs rather than model bugs) and as a probe for the structural component of each metric.

### 4.4 Experiment Settings

Generation uses temperature 0.2, maximum 320 new tokens, and concurrent calls limited to eight at a time with exponential-backoff retry on rate-limit errors. The NLI model is roberta-large-mnli throughout; the sentence encoder is all-MiniLM-L6-v2; the entity tagger is en\_core\_web\_sm (spaCy). Faithfulness NLI is computed on a stratified sample of 30 rows per (model, prompt, profile) cell, 2,880 instances total. Full hyperparameters in Appendix C.

## 5 RESULTS AND DISCUSSION

We organize the results around four research questions matching the four metric axes.

### 5.1 Does the Detector Identify Conflict?

ConflictDetector achieves  $F1 = 0.65$  on the full 6,300-instance benchmark, with accuracy = 0.68, precision = 0.61, recall = 0.69, and FNC = 0.31. The fitted logistic regression coefficients (Appendix Figure 5) show NLI contradiction as the dominant signal at  $\beta = 5.76$ , with number mismatch ( $\beta = 1.96$ ) and entity mismatch ( $\beta = 0.98$ ) contributing additional non-trivial weight. Embedding dispersion and date mismatch take small negative coefficients ( $-0.71$  and  $-1.33$  respectively), indicating that once NLI is in the model these surface heuristics are anti-correlated with the residual—likely because embedding dispersion tracks topical rather than stance variety, and lexical date matches occur in many non-conflicting passage pairs. The coefficient pattern itself serves as an ablation: NLI carries the bulk of the signal, but a pure-NLI detector underperforms the full ConflictDetector on the entity- and number-conflict subsets where lexical cues are decisive. We treat  $F1 = 0.65$  as adequate for the detector’s intended role as a gate rather than a classifier: false positives trigger a more verbose prompt (a small cost), while false negatives risk overconfident answers (the larger cost we mitigate). The FNC = 0.31 figure is the operational quantity to track.

## 5.2 Does Conflict-Aware Prompting Reduce Overconfidence?

The headline finding is in Figure 3. Under standard prompting, gpt-4o-mini commits to a single stance without acknowledgment on 30.5% of mutually-contradictory inputs. Both the always-on conflict-aware prompt and the detector-gated variant reduce overconfidence by an order of magnitude, to 1.5% and 2.5% respectively; the oracle-label upper bound reaches 0.1%. The template baseline is overconfidence-free by construction. The  $20\times$  reduction is the largest single effect we observe in any condition. We note the absolute baseline should be interpreted in context: 30.5% is the overconfidence rate on a deliberately constructed conflict probe with at least one contradicting passage in E, not on retrieval over a natural corpus, where the marginal rate of mutually-contradictory contexts is unknown. The construction makes the failure mode measurable; the reduction shows that a lightweight intervention closes most of the gap.




**Figure 3:** Overconfidence rate by model and prompt (lower is better).

**Statistical significance.** We compute paired bootstrap 95% confidence intervals over 1,000 resamples on the per-instance overconfidence flag. Under a stricter regex-only reproduction of the metric—which does not apply the NLI hedge detector and consequently yields higher absolute OC rates on the unconstrained standard prompt (regex OC = 0.599 vs. Table 2’s 0.305), but matches the NLI-augmented metric to within 0.01 on all conflict-aware conditions—the paired absolute reduction (standard – conflict-aware-gated) is 0.588 with 95% CI [0.569, 0.607],  $p < 0.001$  against the null of zero reduction; the always-on variant gives an essentially identical reduction of 0.585 with 95% CI [0.567, 0.602],  $p < 0.001$ . The effect is therefore robust both to bootstrap variation and to the specific metric operationalization.

**Transfer from constructed to natural conflict.** Because VitaminC instances are constructed by recombining Wikipedia-revision evidence pairs while Climate-FEVER instances are naturally-annotated, the two splits stress different generative conditions on conflict. Stratifying the regex reproduction by source corpus, the reduction transfers across both regimes. On VitaminC (constructed), gpt-4o-mini OC drops from 0.682 [0.663, 0.703] under the standard prompt to 0.013 [0.008, 0.018] under conflict-aware-gating. On Climate-FEVER (natural), it drops from 0.333 [0.294, 0.370] to 0.008 [0.002, 0.016]. All four reductions are significant at  $p < 0.001$ . The lower Climate-FEVER baseline reflects the more diffuse and less directly contrastive nature of naturally-annotated climate evidence; what matters is that the conflict-aware floor ( $\sim 0.01$ ) is essentially identical across the two splits, indicating that the intervention closes the gap regardless of whether the conflict was synthesized or sourced. Acknowledgment rates show the symmetric pattern:  $0.033 \rightarrow 0.945$  on VitaminC and  $0.040 \rightarrow 0.744$

on Climate-FEVER under conflict-aware-gating. The Climate-FEVER ack rate is lower, plausibly because natural disagreement is often subtler and less amenable to lexical surface acknowledgment, which is itself an interesting target for future metric refinement.

## 5.3 Where Does Acknowledgment Improve?

Appendix Figure 6 breaks acknowledgment down by profile and (model, prompt). On the four explicit-conflict profiles—support\_plus\_contradict, numerical\_conflict, entity\_conflict, and multi\_document\_disagreement—standard prompting acknowledges disagreement at  $\sim 10\%$  on average, while conflict-aware variants reach  $\sim 85\%$ ; oracle-label saturates near 1.0. Importantly, on the seven non-conflict profiles, the conflict-aware prompt does not pathologically over-acknowledge: rates remain near zero for support\_only and contradict\_only, indicating that the template’s instructions to “group by stance and cite both” degrade gracefully when only one stance is present in the evidence.

## 5.4 Does Citation Distribution Mirror the Evidence Distribution?

Figure 4 reports  $1 - \text{JS}(p_c \parallel p^*)$  aggregated over the four mixed-evidence profiles. Standard gpt-4o-mini reaches 0.82, indicating noticeable skew toward whichever stance is presented first in the prompt. Conflict-aware variants raise the balance to 0.97–0.98, comparable to the oracle-label ceiling (0.99) and the template baseline (0.99). The improvement is consistent with the qualitative pattern in failure analysis (Appendix Figure 7): under standard prompting, one-sided citation is the second-most-common failure mode for gpt-4o-mini, after overconfident collapse.




**Figure 4:** Citation balance ( $1 - \text{JS}$ ) under mixed-evidence profiles (higher is better).

## 5.5 Faithfulness and Failure Modes

Per-profile faithfulness (Appendix Table 5) tells a consistent story across the four conflict-bearing profiles. Under standard prompting, gpt-4o-mini faithfulness sits between 0.06 and 0.11; under the detector-gated conflict-aware prompt it rises to between 0.18 and 0.25, a 2–4 $\times$  increase. The absolute values are conservative by metric design: a claim is considered faithful only if the maximum NLI entailment probability from any cited evidence sentence exceeds 0.5, which is a strict threshold for fact-verification-style claims that frequently rely on multi-sentence aggregation. What is important is that conflict-aware prompting does not trade faithfulness for acknowledgment—it improves both axes simultaneously. The deterministic template returns very low faithfulness scores ( $\sim 0.02$ – $0.06$ ): the template emits meta-statements such as “the evidence supports the claim,” which are themselves not entailed by the underlying domain evidence; this is the expected metric behavior on a non-content-bearing generator, and we report it as a sanity check on the metric implementation rather than as a model comparison.

The failure taxonomy (Appendix Figure 7) confirms that overconfident collapse ( $\sim 29\%$  of gpt-4o-mini failures) and false confidence on insufficient evidence ( $\sim 22\%$ ) are the dominant failure modes under standard prompting—both of which conflict-aware prompting directly mitigates.

## 5.6 Main Results Table

Table 2 consolidates the headline behavioral metrics. The detector-gated variant matches the always-on variant on overconfidence and citation balance while issuing a conflict-aware prompt on only  $\sim 31\%$  of instances (FNC-controlled), reducing compute without measurable behavioral loss.

| Model       | Prompt   | ack ↑        | oc ↓         | one-sided ↓  | correct-unc. ↑ | CS ↑         |
|-------------|----------|--------------|--------------|--------------|----------------|--------------|
| gpt-4o-mini | standard | 0.067        | 0.305        | 0.228        | 0.028          | 0.364        |
|             | always   | 0.698        | 0.015        | 0.075        | 0.259          | <b>1.000</b> |
|             | gated    | 0.611        | 0.026        | 0.060        | 0.254          | 0.958        |
|             | oracle   | <b>0.869</b> | <b>0.001</b> | <b>0.006</b> | <b>0.314</b>   | <b>1.000</b> |
| template    | standard | 0.583        | 0.000        | 0.000        | 0.333          | 1.000        |
|             | always   | 0.583        | 0.000        | 0.000        | 0.333          | 1.000        |
|             | gated    | 0.583        | 0.000        | 0.000        | 0.333          | 1.000        |
|             | oracle   | 0.583        | 0.000        | 0.000        | 0.333          | 1.000        |

**Table 2:** Main behavioral results over the full 6,300-instance benchmark. always = conflict\_aware\_always; gated = conflict\_aware\_gated; oracle = oracle\_label. Best within the gpt-4o-mini block in bold.

## 5.7 Cross-Model Generalization and Prompt Robustness

We perform three robustness checks on a stratified 500-instance subset (100 per conflict profile): (i) a cross-model check that re-runs the standard and conflict-aware-gated conditions on gpt-4o as a more capable proprietary model, (ii) an open-weights check that re-runs the same conditions on Llama-3.1-8B-Instruct (Dubey et al., 2024) via the Together AI API, and (iii) a prompt-sensitivity check that re-runs gpt-4o-mini with a substantively rephrased alternative conflict-aware prompt (template in Appendix C). Results are reported in Table 3; all bootstrap CIs are 95% over 1,000 resamples.

| Condition                                     | oc ↓               | ack ↑              |
|-----------------------------------------------|--------------------|--------------------|
| <i><b>Cross-model: gpt-4o</b></i>             |                    |                    |
| standard                                      | 0.580 [.534, .622] | 0.024 [.012, .038] |
| gated                                         | 0.150 [.120, .182] | 0.574 [.534, .618] |
| <i><b>Open weights: Llama-3.1-8B</b></i>      |                    |                    |
| standard                                      | 0.429 [.388, .470] | 0.028 [.014, .043] |
| gated                                         | 0.213 [.178, .250] | 0.262 [.223, .299] |
| <i><b>Prompt sensitivity: gpt-4o-mini</b></i> |                    |                    |
| gated (primary)                               | 0.008              | 0.912              |
| alt. rephrased                                | 0.100 [.076, .126] | 0.522 [.476, .564] |

**Table 3:** Robustness checks on a 500-instance stratified subset; bracketed values are 95% bootstrap CIs.

**Cross-model and cross-family.** The reduction transfers to a more capable proprietary model and to an open-weights model. On gpt-4o, paired absolute OC reduction is 0.430, 95% CI [0.382, 0.474],  $p < 0.001$  over 500 instances; on Llama-3.1-8B-Instruct (Dubey et al., 2024) via the Together AI API, 0.215, 95% CI [0.165, 0.267],  $p < 0.001$  over 484 paired instances (the remaining 16 dropped to rate-limit failures). The conflict-aware-gated OC floor differs across models—0.011, 0.150, and 0.213 for gpt-4o-mini, gpt-4o, and Llama-3.1-8B respectively (relative reductions of 99%, 74%, 50%)—and the ack rate tracks this gap (0.912, 0.574, 0.262). We interpret the floor differences as a model-style and instruction-following effect: more confident, less hedge-prone models reach a higher residual floor, and smaller open-weights instruction-tuned models follow

the structural instruction in the conflict-aware prompt less reliably. The direction of the effect is invariant across all three; the absolute mitigation scales with the model’s ability to follow structured instructions.

**Prompt sensitivity.** Re-running gpt-4o-mini on the same subset with a substantively rephrased alternative conflict-aware prompt (different lexical cues, instruction ordering, and emphasis; full template in Appendix C) yields  $oc = 0.100$  and  $ack = 0.522$ . This is markedly worse than the primary prompt’s 0.008 and 0.912, but still an 83% relative reduction in overconfidence against the standard-prompt baseline ( $oc \approx 0.60$  on the same subset). The conflict-aware effect is therefore robust in direction to prompt rephrasing while sensitive in magnitude: the primary prompt should be read as a particularly strong variant rather than a uniquely viable one.

## A Dataset Statistics

EvidenceConflict is constructed from two source corpora. VitaminC (Schuster et al., 2021) is a fact-verification corpus of 488K (claim, evidence, label) instances built from contrastive Wikipedia revisions; every claim has paired evidence with opposite labels, which is the property we exploit for controlled-conflict construction. We retain 2,000 claims that have at least one SUPPORTS and one REFUTES evidence sentence, and sample 400 instances per profile, yielding 4,800 instances in the EC-VC split. Climate-FEVER (Diggelmann et al., 2020) is a fact-verification corpus of 1,535 real-world climate-related claims, each annotated with up to five evidence sentences from Wikipedia carrying labels in {SUPPORTS, REFUTES, NOT\_ENOUGH\_INFO}. The label distribution is naturally skewed: 1,943 SUPPORTS, 802 REFUTES, and 4,930 NOT\_ENOUGH\_INFO evidence sentences, and at the claim level 654 SUPPORTS, 253 REFUTES, 474 NOT\_ENOUGH\_INFO, and 154 DISPUTED. We construct up to 125 instances per profile, with profile-specific filtering: profiles requiring both stances draw from the 154 dual-stance (often DISPUTED) claims, while single-stance profiles draw from the broader  $\sim 1,000$ -claim pool. This yields 1,500 instances in the EC-CF split for 6,300 total. Table 4 reports per-profile statistics.

**Construction algorithm.** For each profile  $p$  and source corpus  $s$ , we filter to claims whose evidence pools satisfy the profile’s stance requirements (e.g., support\_plus\_contradict requires  $|S| \geq 1$  and  $|C| \geq 1$ ). We then sample  $n_s$  claims uniformly at random with a fixed seed, and for each sampled claim construct the evidence set by drawing the required number of sentences from  $S$ ,  $C$ ,  $N$ , and—for the partial and outdated profiles—from a synthetic fallback derived from truncated support passages or a temporal-marker insertion respectively. Per-passage labels are emitted alongside the evidence set for use in evaluation.

| Profile                 | n   | mean  E | conf. |
|-------------------------|-----|---------|-------|
| Support-only            | 525 | 3.0     | 0.00  |
| Contradict-only         | 525 | 3.0     | 0.00  |
| Support + contradict    | 525 | 4.0     | 1.00  |
| Support + irrelevant    | 525 | 4.0     | 0.00  |
| Contradict + irrelevant | 525 | 4.0     | 0.00  |
| Support + outdated      | 525 | 4.0     | 0.00  |
| Weak conflict           | 525 | 4.0     | 1.00  |
| Numerical conflict      | 525 | 2.0     | 1.00  |
| Entity conflict         | 525 | 2.0     | 1.00  |
| Partial support         | 525 | 3.0     | 0.00  |
| Multi-doc disagreement  | 525 | 5.0     | 1.00  |
| Ambiguous/insufficient  | 525 | 2.2     | 0.00  |

**Table 4:** Per-profile statistics on EvidenceConflict.

## B Baseline Models

**gpt-4o-mini-2024-07-18.** (OpenAI, 2024) An instruction-tuned closed-source large language model accessed via the OpenAI Chat Completions API. We use temperature 0.2 and max\_tokens 320. All four prompt conditions are evaluated. The model is treated as a fixed inference target: no fine-tuning is performed.

**Template.** A deterministic rule-based generator that consults the gold per-passage labels in E and emits one of three sentence patterns: “the evidence supports the claim [k]” when only support is present, “the evidence contradicts the claim [k]” when only contradiction is present, or “the evidence is conflicting: supporting passages indicate the claim holds [s], while contradicting passages indicate it does not [c]” when both are present. The template is included as a metric sanity check: any faithful behavioral metric should reach its expected ceiling on the template, and deviations reveal metric bugs rather than model bugs. We report template results transparently throughout.

## C Experimental Settings

**Prompt templates.** All four prompt conditions share the same envelope “Answer the question using only the evidence; cite passages by their numeric label.” The four variants differ as follows:

6. **standard:** no additional instruction.
7. **conflict\_aware\_always:** appends “The evidence may contain contradicting passages. If so, name the disagreement explicitly, group passages by stance, and cite at least one passage from each side.”
8. **conflict\_aware\_gated:** identical to conflict\_aware\_always when  $\hat{y}_{\text{conf}} > 0.5$ ; identical to standard otherwise.
9. **oracle\_label:** prepends each evidence passage with its gold stance label (SUPPORTS / REFUTES / NEUTRAL). Diagnostic ceiling.

For the prompt-sensitivity analysis in Section 5.7, we use a substantively rephrased alternative conflict-aware prompt:

*“Answer the question based on the provided evidence. Cite each passage by its number. IMPORTANT: First check whether the passages agree with each other. If they disagree—on the answer, on a number, on a name, or on a date—state this explicitly. Quote both views and explain that the answer depends on which source one trusts.”*

This variant differs from the primary conflict\_aware template in lexical cues, instruction ordering, and emphasis (it leads with “IMPORTANT: First check” rather than the conditional “If so” framing, and uses “quote both views” instead of “cite at least one passage from each side”).

### Hyperparameters.

10. Generation: temperature 0.2, top-p 1.0, max\_tokens 320, concurrent API calls  $\leq 8$ .
11. NLI: roberta-large-mnli (Liu et al., 2019), batch size 8, max input length 384 tokens.
12. Sentence embeddings: all-MiniLM-L6-v2 (Reimers and Gurevych, 2019).
13. Entity tagger: en\_core\_web\_sm (spaCy 3.7).
14. ConflictDetector threshold:  $\tau = 0.5$ .
15. Faithfulness sample size: 30 instances per (model, prompt, profile) cell (2,880 total).
16. NLI entailment threshold  $\theta_{\text{ent}} = 0.5$ ; acknowledgment NLI threshold  $\theta_{\text{ack}} = 0.5$ .
17. Random seed: 42 throughout.

**Compute and reproducibility.** All experiments were run on a single CPU-only machine (Windows 11, 32 GB RAM, Intel Xeon equivalent). Total wall-clock for the full pipeline was approximately 40 hours, dominated by generation (API rate-limited, ~25 hours) and NLI evaluation (CPU-bound, ~5 hours). API costs were under USD 3 for the full sweep. The full pipeline, including all scripts, configuration files, and the benchmark itself, is released under a permissive open-source license. Coverage after generation was 98.5% of the 50,400 planned (model, prompt, instance) triples; the missing ~1.5% are API-

side rate-limit timeouts that exceeded the retry budget and are excluded from aggregation.

## D Additional Results

**Detector signal contributions.** As discussed in Section 5, the logistic-regression coefficients in Figure 5 serve as a per-signal contribution analysis. NLI contradiction carries the bulk of the signal ( $\beta = 5.76$ ); number mismatch ( $\beta = 1.96$ ) and entity mismatch ( $\beta = 0.98$ ) provide complementary cues that are particularly decisive on the numerical\_conflict and entity\_conflict profiles where stance differences are encoded lexically rather than via clause-level negation. Embedding dispersion and date mismatch take small negative weights on the residual, indicating they do not provide additional independent signal once NLI is included.




**Figure 5:** ConflictDetector logistic-regression signal weights. NLI contradiction dominates; number and entity mismatch contribute additional signal.

**Acknowledgment by profile.** Figure 6 provides a per-profile breakdown of the acknowledgment rate across all twelve evidence profiles for each (model, prompt) combination, complementing the headline numbers reported in Section 5.




**Figure 6:** Conflict acknowledgment rate by evidence profile and prompt condition.

**Per-profile faithfulness.** Table 5 reports faithfulness, unsupported-claim rate, and overconfidence by profile for gpt-4o-mini under the four prompt conditions. Faithfulness is approximately constant across conditions, confirming that conflict-aware prompting does not trade away claim-level support.

| Profile              | Prompt | Faith. | Unsupp. | OC    |
|----------------------|--------|--------|---------|-------|
| Support + contradict | std.   | 0.081  | 0.937   | 0.800 |
|                      | gated  | 0.200  | 0.865   | 0.000 |
| Numerical conflict   | std.   | 0.068  | 0.972   | 0.700 |

| Profile                | Prompt | Faith. | Unsupp. | OC    |
|------------------------|--------|--------|---------|-------|
| Entity conflict        | gated  | 0.176  | 0.881   | 0.133 |
|                        | std.   | 0.063  | 0.983   | 0.667 |
| Multi-doc disagreement | gated  | 0.250  | 0.832   | 0.100 |
|                        | std.   | 0.108  | 0.903   | 0.700 |
|                        | gated  | 0.198  | 0.866   | 0.000 |

**Table 5:** Per-profile faithfulness, unsupported-claim rate, and overconfidence (gpt-4o-mini). std. = standard; gated = conflict\_aware\_gated.

**Citation behavior.** Table 6 reports cited-support and cited-contradict rates and the JS-based citation balance for each (model, prompt) cell, aggregated over the four mixed-evidence profiles. We omit the per-instance wrong-side rate because, under prompts that explicitly cite both sides, that statistic becomes a mechanical artifact of even-handed citation rather than a measure of citation quality; the JS-based citation balance is the appropriate aggregate metric here.

| Model       | Prompt | cit-S | cit-C | bal. ↑       |
|-------------|--------|-------|-------|--------------|
| gpt-4o-mini | std.   | 0.565 | 0.296 | 0.818        |
|             | always | 0.481 | 0.360 | 0.977        |
|             | gated  | 0.485 | 0.351 | 0.975        |
|             | oracle | 0.469 | 0.355 | <b>0.997</b> |
| template    | std.   | 0.533 | 0.467 | 0.990        |
|             | always | 0.533 | 0.467 | 0.990        |
|             | gated  | 0.533 | 0.467 | 0.990        |
|             | oracle | 0.533 | 0.467 | 0.990        |

**Table 6:** Citation behavior under mixed-evidence profiles. bal. =  $1 - \text{JS}(p_c \parallel p^*)$ , higher is better. Prompt abbreviations follow Table 2.

**Representative failure cases.** Figure 7 summarizes the distribution of failure types across model and prompt settings. The most important pattern is that standard prompting does not simply produce random errors. Its failures concentrate in a few interpretable categories: overconfident collapse, where the model chooses one side of contradictory evidence; one-sided citation, where it cites only the passages that support its chosen stance; and false confidence on insufficient evidence, where partial or incomplete evidence is treated as enough for a final answer. This supports the central motivation of EvidenceConflict: contradictory retrieval contexts require separate evaluation from ordinary faithfulness or answer-correctness scoring.




**Figure 7:** Distribution of failure types across models. Standard prompting is dominated by overconfident collapse and false confidence on insufficient evidence.

The qualitative cases in Table 7 make these categories concrete. In each example, the model has access to evidence that should trigger caution, but the standard prompt still encourages a single clean answer. The failures are therefore not only retrieval failures; they are evidence-integration failures. The model often sees the conflicting, outdated, or incomplete passage, but does not treat that passage as a reason to qualify the final response. This is exactly where the conflict-aware prompt helps: it changes the generation task from “answer using evidence” to “first decide whether the evidence itself agrees.”

**Takeaway from qualitative cases.** These cases show why contradictory-evidence RAG cannot be evaluated only with answer correctness or generic faithfulness scores. In all four examples, the retrieved evidence contains a warning signal: an opposing date, an outdated source marker, a truncated passage, or a directly conflicting claim. The standard prompt still encourages the model to produce a clean answer, so the failure occurs at the evidence-integration stage rather than only at retrieval. This is the behavior that EvidenceConflict is designed to expose and that ConflictDetector is designed to reduce.

| Failure type            | Profile                  | Claim                                                                                | Model answer                                                                                                                                                              |
|-------------------------|--------------------------|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Overconfident collapse  | support_plus_contradict  | <i>“Before July 2018, Ballistic: Ecks vs. Sever had a 0% Rotten Tomatoes score.”</i> | “Supported by the evidence. Evidence [3] and [4] state that as of June 2018, the film had a 0% score on Rotten Tomatoes.”                                                 |
| One-sided citation      | support_plus_contradict  | <i>“Emmanuel Adebayor was born on May 26, 1984.”</i>                                 | “Supported by the evidence in [1] and [3], both of which state he was born on that date.” Passages [2] and [4] state February 26, 1984, but are silently ignored.         |
| Ignored outdated signal | support_plus_outdated    | <i>“The Place Beyond the Pines is a 2012 American film.”</i>                         | “Supported by the evidence. Evidence [1] and [3] state the film is from 2012.” Passages [2] and [4] carry an explicit [Reported in an older source] tag and state 1972.   |
| Unsupported synthesis   | partial_support_conflict | <i>“Laura Prepon rejected portraying Alex Vause.”</i>                                | “Supported by the evidence; Laura Prepon rejected portraying Alex Vause, as stated in evidence [1].” Passages [2] and [3] are truncated and do not corroborate the claim. |

**Table 7:** Representative failure cases for gpt-4o-mini under the standard prompt.

## 6 CONCLUSION

We introduced EvidenceConflict, a controlled benchmark of 6,300 RAG instances for studying contradictions inside

retrieved evidence. The benchmark covers twelve evidence profiles and evaluates four key behaviors: conflict detection, acknowledgment, citation use, and faithfulness. We also proposed ConflictDetector, a lightweight pre-generation detector that combines pairwise NLI and surface mismatch signals to route inputs to a conflict-aware prompt. Our experiments show that evidence-internal contradiction is a measurable failure mode in RAG: standard prompting often collapses to one stance on contradictory inputs, while our detector-plus-prompt approach substantially reduces overconfident answers and improves conflict acknowledgment and citation balance without harming non-conflicting cases. These results suggest that contradiction handling should be treated as a core part of RAG evaluation, not only as a downstream hallucination problem.

## LIMITATIONS

EvidenceConflict is a controlled diagnostic benchmark rather than a full simulation of production RAG. It focuses on short English claims with sentence-level evidence, so long-form QA, multi-hop reasoning, multilingual evidence, and paragraph-level contradictions may need additional profiles and metrics. Some profiles are built by recombining stance-labeled passages from VitaminC and Climate-FEVER, which gives control over conflict type and severity but cannot capture all open-web retrieval noise; the support\_plus\_outdated profile is also approximated through temporal-disclaimer insertion rather than fully natural time-displaced sources. The evaluation relies on automatic NLI-based components for entailment, stance extraction, faithfulness, and acknowledgment checks, so numerical, temporal, or implicit conflicts may still be imperfectly measured. We validate acknowledgment with a separate LLM judge, but broader human annotation would strengthen reliability. Finally, the full sweep uses one main instruction-tuned API model and a template baseline, with subset checks on other models and prompts; absolute mitigation levels may therefore vary by model and prompt. Future work should extend this setting to retrieval-time conflict grouping, re-ranking, source-quality estimation, multilingual evidence, and multi-turn clarification.

## Ethical Considerations

EvidenceConflict is constructed from publicly released corpora (VitaminC, Climate-FEVER), neither of which contains personally identifiable information beyond what already appears in public Wikipedia. No new human subjects data was collected. The conflict-aware prompting intervention is intended to make RAG systems more honest about evidence disagreement, particularly in sensitive domains such as climate change, public health, and policy, where one-sided collapse to a single stance can amplify misinformation. We acknowledge two risks. First, a system that more readily flags “evidence is conflicting” may be perceived as less authoritative even when underlying disagreement is legitimate; we view this as a feature rather than a bug, but flag it for downstream users. Second, the detector and prompt template could in principle be inverted to suppress acknowledgment in deployments where doing so is commercially convenient, and we discourage such use. The manual validation protocol was conducted by the authors and is not representative of broader annotator populations. We release the full benchmark, evaluation code, and detector implementation under a permissive open-source license to enable independent audit and replication.

## REFERENCES

- [1] 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 Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks.
- [2] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-RAG: Learning to retrieve, generate, and critique through self-reflection. In International Conference on Learning Representations.
- [3] Hung-Ting Chen, Michael J.Q. Zhang, and Eunsol Choi. 2022. Rich knowledge sources bring complex knowledge conflicts: Recalibrating models to reflect conflicting evidence. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing.
- [4] Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. Benchmarking large language models in retrieval-augmented generation. In Proceedings of the AAAI Conference on Artificial Intelligence.
- [5] Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.
- [6] Thomas Diggelmann, Jordan Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus Leippold. 2020.

CLIMATE-FEVER: A dataset for verification of real-world climate claims. arXiv preprint arXiv:2012.00614.

- [7] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, and 1 others. 2024. The Llama 3 herd of models. arXiv preprint arXiv:2407.21783.
- [8] Shahul Es, Jithin James, Luis Espinosa-Anke, and Steven Schockaert. 2024. RAGAs: Automated evaluation of retrieval augmented generation. In Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations.
- [9] J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.
- [10] 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*.
- [11] 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*.
- [12] 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)*.
- [13] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv:1907.11692.
- [14] Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. Entity-based knowledge conflicts in question answering. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- [15] Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- [16] Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Wei 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*.
- [17] OpenAI. 2024. GPT-4o mini: Advancing cost-efficient intelligence. Accessed 2026-05-22.
- [18] Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. KILT: a benchmark for knowledge intensive language tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [19] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
- [20] Jon Saad-Falcon, Omar Khattab, Christopher Potts, and Matei Zaharia. 2024. ARES: An automated evaluation framework for retrieval-augmented generation systems. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [21] Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. Get your Vitamin C! robust fact verification with contrastive evidence. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [22] Zhaochen Su, Jun Zhang, Xiaoye Qu, Tong Zhu, Yanshu Li, Jiashuo Sun, Juntao Li, Min Zhang, and Yu Cheng. 2024. ConflictBank: A benchmark for evaluating the influence of knowledge conflicts in LLM. In *Advances in Neural*

Information Processing Systems.

- [23] 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.
- [24] David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP).
- [25] Yike Wang, Shangbin Feng, Heng Wang, Weijia Shi, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2024. Resolving knowledge conflicts in large language models. In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing.
- [26] Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. Long-form factuality in large language models. Advances in Neural Information Processing Systems.
- [27] Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. 2024. Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts. In International Conference on Learning Representations.
- [28] Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024. Knowledge conflicts for LLMs: A survey. In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing.
- [29] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and chatbot arena. In Advances in Neural Information Processing Systems.