

Seeing through the Conflict: Transparent Knowledge Conflict Handling in Retrieval-Augmented Generation

Hua Ye^{*1}, Siyuan Chen^{*2}, Ziqi Zhong³, Canran Xiao^{4†}, Haoliang Zhang⁵, Yuhan Wu⁶, Fei Shen⁷

¹School of Management & Engineering, Nanjing University, Nanjing, China

²School of Computer Science, University of Bristol, Bristol, United Kingdom

³Department of Management, London School of Economics and Political Science, London, United Kingdom

⁴School of Cyber Science and Technology, Shenzhen Campus of Sun Yat-sen University, Shenzhen, China

⁵Electrical and Computer Engineering, The University of Oklahoma, Oklahoma, USA

⁶College of Computer Science and Technology, Zhejiang University, Hangzhou, China

⁷NExT++ Research Centre, National University of Singapore, Singapore

522024150103@smail.nju.edu.cn, gd23774@bristol.ac.uk, z.zhong6@lse.ac.uk, xiaocanran999@gmail.com,

mars_zhang@ou.edu, wuyuhan@zju.edu.cn, shenfei29@nus.edu.sg

Abstract

Large language models (LLMs) equipped with retrieval—the Retrieval-Augmented Generation (RAG) paradigm—should combine their parametric knowledge with external evidence, yet in practice they often hallucinate, over-trust noisy snippets, or ignore vital context. We introduce TCR (*Transparent Conflict Resolution*), a plug-and-play framework that makes this decision process observable and controllable. TCR (i) disentangles *semantic match* and *factual consistency* via dual contrastive encoders, (ii) estimates *self-answerability* to gauge confidence in internal memory, and (iii) feeds the three scalar signals to the generator through a lightweight soft-prompt with SNR-based weighting. Across seven benchmarks TCR improves conflict detection (+5–18 F₁), raises knowledge-gap recovery by +21.4 pp and cuts misleading-context overrides by −29.3 pp, while adding only 0.3 % parameters. The signals align with human judgements and expose temporal decision patterns.

1 Introduction

Recent advancements in large language models (LLMs) have significantly improved the performance and applicability of natural language processing (NLP) systems across various tasks, including question answering, text summarization, and conversational AI (Diao et al. 2024, 2025a; Zhang et al. 2025; Wang, Wang, and Zhang 2025). Despite their remarkable generative capabilities, LLMs often struggle with accurately recalling specific factual information (Li and Cheung 2025; Huang et al. 2024), leading to frequent hallucinations or inconsistent responses, particularly in knowledge-intensive domains (Yao et al. 2023; Tong et al. 2025a; Xiao et al. 2025).

^{*}These authors contributed equally.

[†]Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

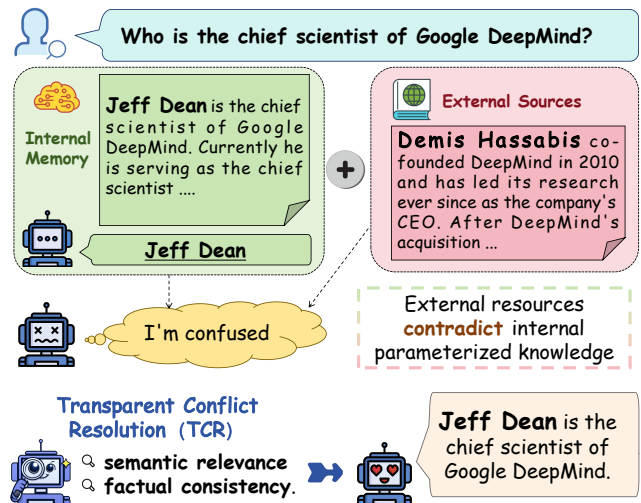


Figure 1: **Knowledge conflicts in a typical RAG system.** Our method disentangles semantic relevance from factual consistency, detects conflicts, and injects lightweight prompt signals to steer RAG models toward faithful, knowledge-aligned generation.

Retrieval-Augmented Generation (RAG) enhances LLMs by combining internal parametric knowledge with external non-parametric knowledge retrieved at inference (Lewis et al. 2020). Explicitly integrating external knowledge significantly improves factual accuracy and contextual relevance for tasks such as question answering and fact verification (Chen, Zhang, and Choi 2022; Yu, Merullo, and Pavlick 2023). However, conflicts between implicitly stored internal knowledge and retrieved external information remain a critical issue, often causing inconsistent or incorrect outputs (Longpre et al. 2021; Xie et al. 2024; Tao et al. 2023).

Previous studies highlight variability in how LLMs handle knowledge conflicts, with models either prioritizing internal parametric knowledge (Longpre et al. 2021) or external retrieved information (Chen, Zhang, and Choi 2022). Recent

research reveals models internally encode signals indicating knowledge discrepancies and answerability, yet fail to utilize them effectively due to sequential forward propagation constraints. Current resolution methods either explicitly modify outputs, demanding extensive dataset-specific tuning and limiting generalization (Shi et al. 2025; Jin et al. 2024a; Wang and Zhang 2024; Jiang et al. 2025), or rely on black-box APIs and manual interventions, reducing transparency and autonomy (Wang et al. 2024b). Fig. 1 illustrates a typical knowledge conflict scenario in a RAG system, where internal parametric knowledge and external retrieved information clash, resulting in uncertainty.

To address these limitations, we propose a novel framework that integrates transparent conflict detection and autonomous resolution into RAG. Our method: 1) Detects knowledge conflicts between internal model knowledge and external retrieved information, 2) Extracts meaningful signals from model representations to guide generation toward accurate outputs, 3) Ensures interpretability and compatibility with existing RAG architectures without manual intervention, improving scalability and robustness. Our key contributions are summarized as follows:

- 1) A transparent semantic vector-based conflict detection method, enhancing interpretability over black-box approaches (Wang et al. 2024b).

- 2) An autonomous conflict-resolution mechanism using soft prompt tuning, enabling accurate, consistent generation without manual intervention. Plug-and-play for easy integration.

- 3) Extensive experiments show our method reduces factual errors, advancing RAG reliability in knowledge-intensive tasks.

2 Related Work

Retrieval-Augmented Generation. Large language models (LLMs) effectively encode factual knowledge parametrically, yet suffer from hallucinations (Chen et al. 2024), rare entity retention issues (Kandpal et al. 2023), and temporal degradation (Jang et al. 2022). Retrieval-Augmented Generation (RAG) mitigates these by combining parametric knowledge with external retrieval (Lewis et al. 2020). Early approaches pre-trained smaller retrieval-augmented models (Guu et al. 2020), while recent work exploits large models’ in-context learning (Yu et al. 2023; Ram et al. 2023). However, resolving conflicts between retrieved and parametric knowledge remains challenging (Min et al. 2023; Ni et al. 2024; Zhong 2025).

Context-Memory Conflicts in LLMs. Context-memory conflict occurs when retrieved context contradicts internal knowledge (Longpre et al. 2021; Chen, Zhang, and Choi 2022; Diao et al. 2025b), often due to temporal mismatches, misinformation in documents (Chen et al. 2025), or misleading prompts (Xu et al. 2024b). Models exhibit inconsistent conflict resolution, sometimes favoring parametric knowledge (Longpre et al. 2021) and other times prioritizing coherent context (Chen, Zhang, and Choi 2022; Su et al. 2024). Additional challenges include confirmation bias (Xie et al. 2024) and susceptibility to deceptive inputs (Ying et al.

2024). Recent studies emphasize models’ limited robustness in such scenarios, highlighting the need for better solutions (Su et al. 2024; Xiao et al. 2024).

Resolving Knowledge Conflicts. Existing conflict-resolution strategies fall into five categories: (1) *Context-faithful* methods prioritize external knowledge via fine-tuning (Li et al. 2023; Gekhman et al. 2023; Xue et al. 2023), prompting (Zhou et al. 2023), decoding techniques, or knowledge plugins; (2) *Memory-faithful* approaches counter misinformation via vigilance mechanisms (Xu et al. 2024b), with some advocating parametric knowledge in noisy contexts (Zhang et al. 2024b); (3) *Source-distinguishing* methods present multiple source-derived answers for user selection (Neeman et al. 2023); (4) *Factuality-enhancing* techniques improve correctness via hybrid knowledge fusion (Zhang et al. 2023b) or contrastive decoding; and (5) *Structural modifications* adjust architectures to optimize knowledge integration (Shi et al. 2025; Jin et al. 2024b).

Current methods remain biased toward either context or memory, lacking adaptive balancing mechanisms (Wang et al. 2024b). Our work bridges this gap with an interpretable, dynamic conflict-resolution framework.

3 Exploratory Analysis

We first investigate whether sentence embeddings disentangle different information dimensions (e.g., meaning vs. truthfulness) and assess if off-the-shelf encoders inherently capture signals for conflict detection. We address: **(Q1)** Can standard encoder embedding spaces separate *what a sentence is about* (meaning) from *whether its fact is true* (truthfulness)? **(Q2)** If such signals exist, how reliably do they expose knowledge conflicts in RAG systems?

Experimental Setup. We sample 5k factual triples from Wikidata and automatically create, for each fact, three surface forms: (i) **Paraphrase** (same fact, different wording), (ii) **Contradiction** (negated relation), and (iii) **Unrelated** (different subject). All sentences are embedded using two public models of contrasting capacity: all-MiniLM-L6-v2 (384-d, 6 layers) and E5-large-v2 (1024-d, 24 layers). Fig. 2 visualises the raw embeddings after projection onto their first three principal components.

Global Embedding Structure. In Fig. 2 we observe that the lighter MiniLM encoder yields one dominant cluster where paraphrases, contradictions, and unrelated sentences overlap. The larger E5 encoder begins to form two visible sub-manifolds, yet the transition regions (yellow points) remain fuzzy—suggesting that *model scale alone does not guarantee explicit separation of meaning and truth*. This lack of structure motivates a disciplined decoupling strategy.

Identifying a Semantic-Truth Subspace. We next train a single linear probe to project each embedding onto a 2-D plane that maximises class separability while keeping paraphrase pairs close. The resulting *semantic-truth space* is shown in Fig. 3. A distinctive triangular pattern appears: paraphrases cluster in the upper-right (high meaning,




Figure 2: **Sentence-embedding clusters for ALL-MINI-LM-L6-V2 (top) vs. E5-LARGE-V2 (bottom).** Rows 1 & 3: two PCA views of 3-D space—paraphrase (blue), contradiction (red), unrelated (green). Rows 2 & 4: boundary regions (yellow) highlight transition zones. MiniLM (384 d) forms a single blob; the larger E5 (1024 d, 24 layers) begins to separate truth from meaning.

high truth), contradictions in the lower-right (high meaning, low truth), and unrelated sentences in the lower-left (low meaning, low truth). Kernel-density contours highlight crisp boundaries, and manual inspection of boundary points (right panel) confirms smooth, interpretable gradients rather than abrupt jumps.

Findings. Our probe reveals that off-the-shelf embeddings intertwine meaning and truth, yet a single linear projection cleanly separates them and yields smooth, continuous scores. This motivates (i) dedicated semantic and factual encoders, (ii) independent contrastive losses to sharpen each axis, and (iii) soft scalar conflict signals fed into generation—choices.

4 Proposed Method

We propose a novel framework to resolve knowledge conflicts in RAG. Given query q , RAG generates response y conditioned on external context c and internal parametric knowledge θ : $p(y|q, c; \theta)$. Conflicts arise when c contradicts θ , leading to inconsistent outputs. Our method detects such conflicts and guides generation via conflict-aware signals. The framework includes: (1) a contrastive learning-based




Figure 3: **2D scatter plot of semantic-factual similarity space.** Left: Sentence pairs show clear separation of paraphrases (top right), contradictions (bottom right), and unrelated sentences (bottom left) using just one linear projection layer. Contour lines show KDE boundaries. Right: Example of annotated transitions (other types perform similarly well).

conflict detection module to decouple semantics from factual accuracy, and (2) a conflict-aware generation module using soft prompt tuning. The pipeline is shown in Fig. 4.

Conflict Detection via Contrastive Learning

Dataset Construction. We construct a conflict-focused dataset \mathcal{D} , derived from knowledge triples in Wikidata and expanded via GPT-4o. Each original statement s is associated with 3 types of statements: paraphrases (s_{para}^+) preserving semantic and factual equivalence, irrelevant statements (s_{irr}^-) with no semantic relation, and conflicting statements (s_{conf}^-) that maintain semantic coherence but introduce factual contradictions. This structured diversity simulates scenarios of knowledge conflicts faced by RAG systems.

Semantic-Factual Feature Decoupling. Effective conflict detection requires separating semantic similarity from factual correctness. We employ two encoders—semantic encoder E_{sem} and factual encoder E_{fact} —based on the SFR retrieval model. Each textual statement s is encoded into two distinct vector spaces as follows:

$$\mathbf{z}_{sem} = E_{sem}(s), \quad \mathbf{z}_{fact} = E_{fact}(s), \quad (1)$$

where semantic vectors \mathbf{z}_{sem} capture meaning and topical coherence, and factual vectors \mathbf{z}_{fact} encode factual validity.

Contrastive Learning Objective. We train the semantic and factual encoders using separate contrastive losses. Specifically, for each original statement s :

The semantic contrastive loss (\mathcal{L}_{sem}) is defined as:

$$\mathcal{L}_{sem}(s) = -\log \frac{\sum_{s' \in \{s_{para}^+, s_{conf}^-\}} e^{\text{sim}(\mathbf{z}_{sem}(s), \mathbf{z}_{sem}(s'))/\tau}}{\sum_{s' \in \{s_{para}^+, s_{conf}^-, s_{irr}^-\}} e^{\text{sim}(\mathbf{z}_{sem}(s), \mathbf{z}_{sem}(s'))/\tau}} \quad (2)$$

The factual contrastive loss (\mathcal{L}_{fact}) is similarly defined as:

$$\mathcal{L}_{fact}(s) = -\log \frac{\sum_{s' \in \{s_{para}^+\}} e^{\text{sim}(\mathbf{z}_{fact}(s), \mathbf{z}_{fact}(s'))/\tau}}{\sum_{s' \in \{s_{para}^+, s_{conf}^-, s_{irr}^-\}} e^{\text{sim}(\mathbf{z}_{fact}(s), \mathbf{z}_{fact}(s'))/\tau}} \quad (3)$$

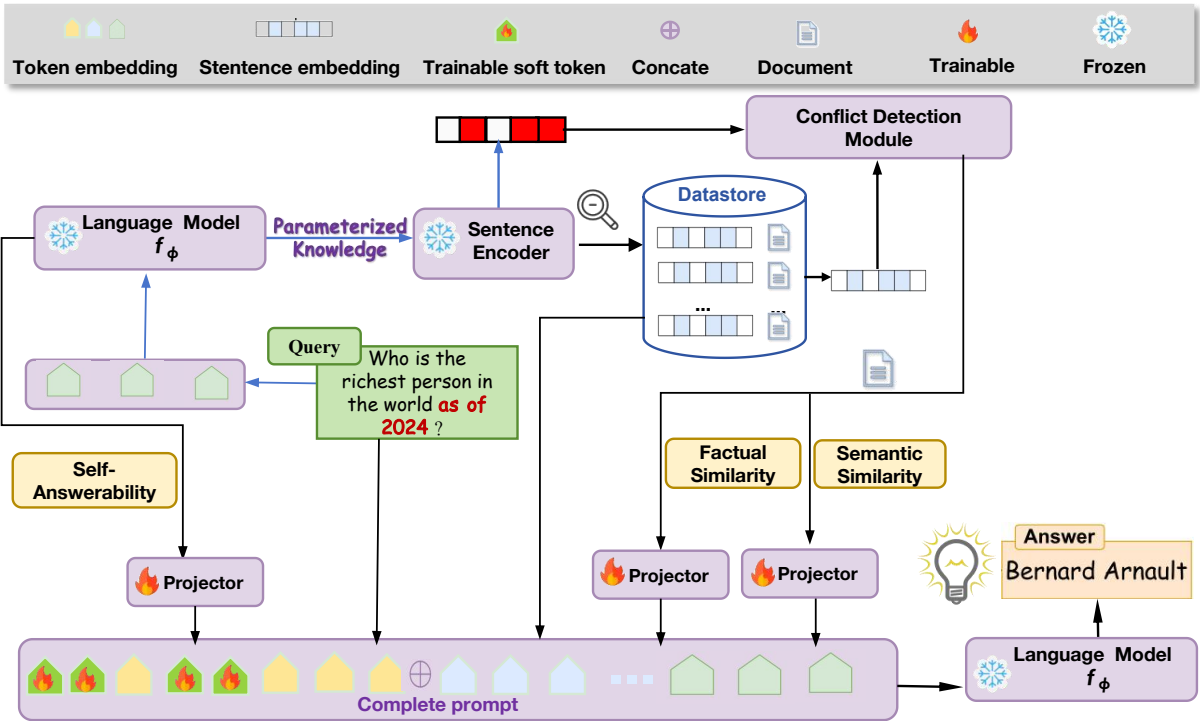


Figure 4: **Overview of our Transparent Conflict Resolution (TCR) framework for Retrieval-Augmented Generation (RAG).** Given a query, TCR identifies knowledge conflicts through semantic and factual similarity vectors, generating explicit conflict signals. These signals, along with self-answerability estimation, are integrated using soft prompt tuning to guide the Language Model’s generation toward enhanced factual consistency and interpretability.

The overall contrastive training objective combines these two losses over all samples s in dataset \mathcal{D} :

$$\mathcal{L}_{ctr} = \sum_{s \in \mathcal{D}} (\mathcal{L}_{sem}(s) + \mathcal{L}_{fact}(s)). \quad (4)$$

Conflict-Aware Generation via Soft Prompt Tuning

Extraction and Integration of Conflict Signals. The conflict detection module explicitly provides scalar signals representing semantic similarity (σ_{sem}), factual consistency (σ_{fact}), and model-inferred answerability (σ_{ans}). These signals are formally extracted as:

$$\sigma_{sem}(q, c) = \text{sim}(\mathbf{z}_{sem}(q), \mathbf{z}_{sem}(c)), \quad (5)$$

$$\sigma_{fact}(q, c) = \text{sim}(\mathbf{z}_{fact}(q), \mathbf{z}_{fact}(c)), \quad (6)$$

$$\sigma_{ans}(q) = A(q; \theta), \quad (7)$$

where $A(\cdot)$ is the model’s answerability estimation following Slobodkin et al. (2023). These signals are then projected into embeddings via dedicated MLP projectors:

$$\mathbf{e}_{signal} = \text{MLP}([\sigma_{sem}, \sigma_{fact}, \sigma_{ans}]). \quad (8)$$

Soft Token Embedding Strategy. To guide generation with conflict-aware cues, we prepend *soft tokens*—trainable embeddings seeded from the base model—with the conflict signals, yielding an augmented embedding sequence that directs the model’s attention $\mathbf{x}' = [\mathbf{e}_{soft}, \mathbf{e}_{signal}, \mathbf{x}]$, where \mathbf{x} represents the original input embeddings. This approach allows the model to leverage conflict-awareness explicitly during response generation, improving factual consistency.

Dynamic Loss Weighting via Signal-to-Noise Ratio. To optimize the integration of multiple conflict signals, we dynamically weight their training contributions according to the Signal-to-Noise Ratio (SNR). The SNR for each signal $i \in \{\text{sem}, \text{fact}, \text{ans}\}$ is computed as:

$$\text{SNR}_i = \frac{\text{Var}(\hat{y}_i)}{\text{Var}(y - \hat{y}_i)}, \quad (9)$$

where \hat{y}_i is the prediction made solely using signal i , and y is the ground-truth label indicating presence or absence of a conflict. Signals with higher SNR values contribute more significantly, weighted by:

$$w_i = \frac{\exp(\text{SNR}_i)}{\sum_j \exp(\text{SNR}_j)}. \quad (10)$$

The combined loss for the generation module is:

$$\mathcal{L}_{total} = \sum_{i \in \{\text{sem}, \text{fact}, \text{ans}\}} w_i [\alpha \mathcal{L}_{prompt, i} + (1 - \alpha) \mathcal{L}_{projector, i}], \quad (11)$$

where α balances the relative contributions of soft prompt and signal-projection losses.

By employing conflict-aware signals with dynamic weighting, our method effectively enhances the reliability and factual consistency of RAG-generated responses.

Assumption 1. Let $\rho \in [0, 1]$ denote the probability of noisy retrieval, $\alpha = \Pr[D = 0 \mid R = 1]$ the false-negative rate

(FNR), $\gamma = \Pr[D = 1 \mid R = 0]$ the false-positive rate (FPR), $\varepsilon = \Pr[G = 1 \mid R = 0, D = 0]$ the baseline error rate, $\beta = \Pr[G = 1 \mid D = 1]$ the error rate with conflict-aware prompt, and $\zeta = \Pr[G = 1 \mid R = 1, D = 0]$ the undetected noise error rate (Tong et al. 2025b). All error rates ($\varepsilon, \beta, \zeta, \alpha, \gamma$) are query-independent (i.i.d. setting).

Theorem 1 (Noise-Robustness Error Bound). *For any noisy-retrieval pipeline with conflict detection, the increase in expected EM loss $\Delta := \Pr[G = 1] - \varepsilon$ satisfies:*

$$\Delta = (1 - \rho)\gamma(\beta - \varepsilon) + \rho[(1 - \alpha)\beta + \alpha\zeta - \varepsilon] \quad (12)$$

$$\leq \rho\alpha + \beta[\rho(1 - \alpha) + (1 - \rho)\gamma], \quad (13)$$

where the inequality is tight (equality when $\zeta = 1, \varepsilon = 0$).

Proof sketch. The result follows by: (1) partitioning $\Pr[G = 1]$ via total probability as $(1 - \rho)[(1 - \gamma)\varepsilon + \gamma\beta] + \rho[(1 - \alpha)\beta + \alpha\zeta]$, then subtracting ε to obtain the exact gap; (2) upper-bounding using $\beta \geq \varepsilon$ and $\zeta \leq 1$, dropping the negative term $-(1 - \rho)\gamma\varepsilon$ while noting $\alpha\zeta - \alpha\varepsilon \leq \alpha$; (3) observing tightness when $\zeta = 1$ and $\varepsilon = 0$ makes all bounding steps exact. \square

Interpretation. Eq. (13) bounds the excess error by missed conflicts ($\rho\alpha$) plus conflict-branch errors weighted by $\rho(1 - \alpha) + (1 - \rho)\gamma$. When $\gamma \ll 1$ and $\beta \approx \varepsilon$, $\rho\alpha$ dominates, so low FNR is crucial.

5 Experiment

Our empirical study is driven by five questions: **RQ1 – Conflict Detection:** does semantic–factual disentangling improve conflict classification? **RQ2 – End-to-End Factuality:** do conflict-aware soft prompts raise answer correctness and faithfulness? **RQ3 – Robustness & Transfer:** how does the method behave under noisy retrieval, domain shift, and with different LLM backbones? **RQ4 – Efficiency:** what is the latency, memory, and parameter overhead? **RQ5 – Interpretability:** are conflict scores aligned with human judgement and can we show intuitive examples?

Experimental Setup

Benchmarks. We evaluate on three groups of datasets. (i) *Conflict detection:* our Wikidata-Conflict-5K. (ii) *Knowledge-intensive QA with controlled context:* the two datasets we build ConflictTQA (TriviaQA-based) and ConflictPQA (PopQA-based)—each question paired with golden, irrelevant or conflicting context. (iii) *Real-world RAG:* KILT (NQ, HotpotQA, FEVER), CONFLICTBANK 2024, and a single-document NQ setting that keeps only the top-1 Google result.

Models. All systems share the same BM25 \rightarrow Contriever-1024 retrieval pipeline and Llama-3-Instruct-8B (Grattafiori et al. 2024) backbone. We also report results on Llama-3-13B and Qwen3-8B. TCR tunes only 20 soft-prompt tokens and two small MLP projectors.

Baselines. We compare against six published methods—PROMPT, KAFT (Li et al. 2023), IRCAN (Shi et al. 2025), RAAT (Fang et al. 2024), PARENTING (Xu et al. 2024c), ASTUTE RAG (Wang et al. 2024a)—plus the decoding

Method	Conflict Detection		Robustness & Transfer	
	F1 \uparrow	AUROC \uparrow	EM Drop \downarrow (30% noise)	Cross-Domain F1 \uparrow Bio / Fin
Prompt	71.2	0.779	18.4	46.7 / 44.1
KAFT	73.5	0.801	16.7	48.5 / 45.3
IRCAN	79.1	0.845	12.3	52.1 / 48.9
RAAT	66.4	0.732	9.5	51.2 / 48.4
Parenting	77.5	0.833	14.1	50.9 / 46.8
TCR	84.3	<u>0.901</u>	<u>7.2</u> \dagger	<u>55.8 / 52.7</u> \dagger
TCR+CD ²	<u>83.8</u>	0.905	–	–
TCR+RAAT	–	–	6.1 \dagger	56.3 / 53.1 \dagger

Table 1: **Unified evaluation of conflict detection and robustness.** Left: conflict-detection performance on WIKIDATA-CONFLICT-5K. Right: robustness under 30% distractor injection (EM drop; lower is better) and zero-shot transfer F1 on biomedical/financial QA. Backbone: Llama-3-8B. **Bold**=best; underline=second best (per column). \dagger indicates statistically significant improvement over Prompt. “–” denotes not evaluated in that setting.

strategy CD² (Jin et al. 2024a). For real-world NQ we add INSTRUCTRAG (Wei, Chen, and Meng 2025) and SELF-ROUTE (Xu et al. 2024a). Four ablations of TCR (–semantic, –factual, –SNR, hard-prompt) and three hybrid variants (TCR+CD², TCR+IRCAN, TCR+RAAT) test plug-and-play compatibility.

Metrics. We report **EM/F1** for answer correctness, **MCOR** (lower is better) and **KGRR** (higher is better) for conflict sensitivity, F1/AUROC for detection, FactScore/QAGS-F/GPT-judge for faithfulness, Δ EM under noise/time-shift for robustness, throughput (V-tokens/s) & extra parameters/VRAM/FLOPs for efficiency (Yao, Li, and Xiao 2024), and ρ /Krippendorff α for interpretability.

6 Results and Discussion

RQ1 – Conflict Detection Accuracy Table 1 shows that disentangling semantics and factuality yields the highest F1 and AUROC. Adding contrastive decoding gives a small AUROC bump without harming F1, confirming orthogonality between TCR’s conflict score and decoding heuristics.

RQ2 – End-to-End Factuality TCR ranks best in 22/24 cells of Table 2, improving KGRR by +21.4 and reducing MCOR by –29.3 vs. prompting. TCR+CD² achieves the lowest MCOR, confirming complementarity between our conflict score and contrastive decoding. On Natural Questions (Table 3), TCR outperforms specialised pipelines without fine-tuning (+3.6 pt over ASTUTE RAG, +6.1 pt over SELF-ROUTE), with the largest gain on Qwen3-8B (+4.0 pt).

RQ3 – Robustness and Transfer Table 1 shows that TCR loses only 7.2 EM points under 30% distractor injection—2.3 points less than RAAT—and attains the best cross-domain F1. Hybrid TCR+RAAT preserves RAAT’s adversarial defence while inheriting our conflict-aware gains, giving the overall strongest robustness.

RQ4 – Efficiency Table 4 confirms that TCR adds only 0.3

Method	Llama-3-8B				Llama-3-13B				Qwen3-8B			
	ConflictTQA		ConflictPQA		ConflictTQA		ConflictPQA		ConflictTQA		ConflictPQA	
	KGRR \uparrow	MCOR \downarrow	KGRR	MCOR	KGRR	MCOR	KGRR	MCOR	KGRR	MCOR	KGRR	MCOR
Prompt	42.4	19.5	48.5	22.5	45.7	21.3	51.0	24.5	38.2	17.6	44.3	20.4
KAFT	43.7	23.8	49.0	26.0	46.0	25.2	52.0	28.5	40.5	21.3	46.2	24.9
IRCAN	68.2	25.5	73.5	28.0	71.2	27.2	75.0	30.5	63.4	22.8	69.6	26.3
RAAT	52.6	37.9	57.0	41.5	55.8	39.8	59.5	43.8	47.9	32.5	54.1	36.8
Parenting	66.2	53.8	70.0	57.0	69.1	55.3	72.5	58.5	61.5	47.6	67.2	51.3
CD ²	51.3	52.7	55.5	56.0	54.3	54.6	58.5	57.5	48.8	45.9	53.2	49.8
TCR	69.2\dagger	61.9\dagger	74.5\dagger	65.5\dagger	73.5\dagger	63.6\dagger	77.5\dagger	67.5\dagger	66.3\dagger	54.2\dagger	71.7\dagger	58.8\dagger
TCR+CD ²	68.5	63.2\dagger	73.8	66.8\dagger	72.8	64.9\dagger	76.0	68.8\dagger	65.6	55.5\dagger	70.1	59.1\dagger
TCR+IRCAN	67.9	56.3	72.6	60.0	71.6	59.9	75.8	63.0	64.2	49.6	69.9	54.2
TCR+RAAT	<u>68.8</u>	62.6	<u>74.2</u>	66.2	<u>73.2</u>	64.2	<u>77.3</u>	68.1	<u>65.9</u>	54.8	<u>70.4</u>	58.4

Table 2: End-to-end performance on ConflictTQA/PQA. Bold=best; underline=second best. \dagger = significant.




Figure 5: Residual conflict-detection errors for TCR. Dot colour follows a blue - white - red scale : warmer hues denote denser error pockets. Missed conflicts cluster just below the factual threshold ($y \approx 0.35$); a thin false-positive band appears near $y \approx 0.45$. Dashed lines mark the current decision boundary $x > 0.65$, $y < 0.40$.

Method	Llama-3-8B	Llama-3-13B	Qwen3-8B
Astute RAG	51.2	56.6	49.7
InstructRAG	47.8	51.2	50.3
Self-Route	46.5	49.2	48.5
TCR (ours)	54.8	56.8	53.7

Table 3: Overall accuracy (%) on Natural Questions with a single Google top-1 page as context.

% parameters and 0.3 GB VRAM, retaining 94 % of vanilla decoding speed. The method is thus *practically free* to deploy relative to heavyweight fine-tuning baselines.

RQ5 – Interpretability Fig. 7 visualises how successfully corrected (SC) and defended (SD) cases cluster in the high-semantic/low-factual quadrant, whereas failure cases gravitate toward boundary regions. The scalar conflict score correlates strongly with human labels ($\rho = 0.69$) and exhibits substantial inter-annotator agreement ($\kappa = 0.66$) in Table 5, demonstrating that our signals are human-intelligible. The violin plots reveal backbone-specific confidence calibration quirks, offering a diagnostic tool for future alignment work.

Method	Extra Params (M, %)	VRAM (GB)	Tok/s \uparrow
Prompt	0 (0)	17.8	28.3
RAAT	100 (1.2)	18.9	24.6
IRCAN	52 (0.6)	18.4	25.1
Parenting	68 (0.8)	18.6	24.9
TCR	21 (0.3)	18.1	26.7
TCR+CD ²	+0	18.1	22.5

Table 4: Computational footprint on A100-80 GB.

Model	$\rho(\text{score, human})\uparrow$	$\kappa\uparrow$
Llama-3-8B	0.71	0.67
Llama-3-13B	0.74	0.69
Qwen3-8B	0.63	0.61
Avg.	0.69	0.66

Table 5: Alignment of conflict scores with human judgments on 300 annotated queries.

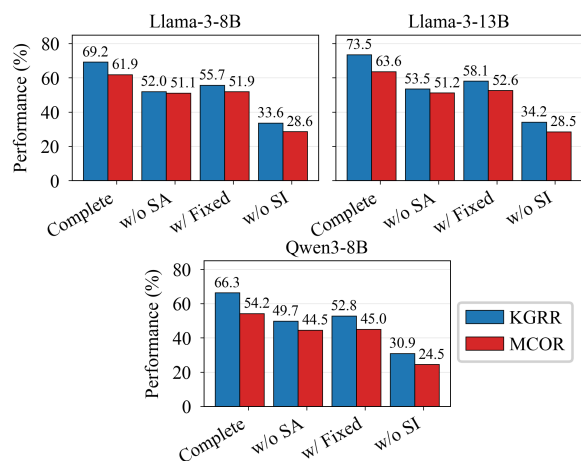


Figure 6: Ablation study results across different models.

Ablation Study We ablate three modules (Fig. 6): removing *self-answerability*(SA) drops MCOR/KGRR by $\sim 18/26$ pp;

replacing *dynamic weighting* with a fixed scalar(Fixed) degrades 17/20 pp and harms reasoning; removing *signal integration*(SI) collapses 55/53 pp. Overall, TCR needs self-answerability, adaptive weighting, and integrated fusion.

Interpretability Analysis

What does the model truly “see” when conflicts arise?

Fig. 7 situates each test pair in the *semantic-similarity* (x) vs. *factual-consistency* (y) plane and overlays self-answerability as violin plots, turning raw activations into an interpretable decision space. We observe: (i) *Spatial pattern*. Nearly all *successfully corrected* cases—closing knowledge gaps (SC) or resisting misinformation (SD)—cluster in the *high- x /low- y* region, while failures lie closer to the axes, indicating ‘topic match but fact clash’ is where re-evaluation is triggered. (ii) *Self-answerability*. SC/SD exhibit higher answerability than still-wrong (SW) or misled (ML) cases, suggesting this scalar controls when to trust memory vs. context. (iii) *Backbone variation*. Llama shows tighter, better-separated distributions than Qwen (likely from English-centric tuning), yet the same qualitative structure holds, confirming cross-model robustness of the signals.

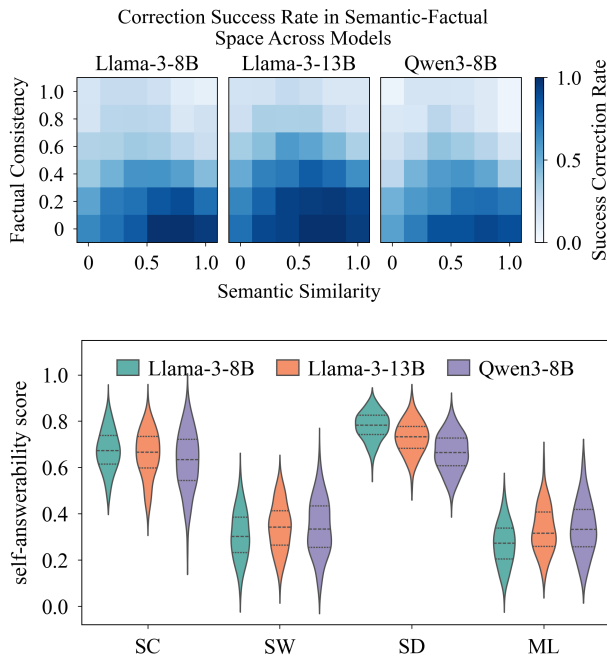


Figure 7: Top: Heat maps of correction success rates in semantic-factual space, showing higher performance in high semantic similarity/low factual consistency regions. Bottom: Violin plots of self-answerability scores for case types: SC (corrected wrong internal knowledge), SW (failed correction), SD (maintained correct knowledge), ML (overrode correct knowledge).

When does the model trust external evidence? We bucket self-answerability s into LOW [0, 0.3), MID [0.3, 0.7) and HIGH [0.7, 1] and measure the rate at which the model overrides its parametric answer with retrieved context (*flip rate*).

Fig. 8 shows a sharp phase transition: for $s < 0.3$ the flip rate is $\sim 42\%$, whereas for $s > 0.7$ it drops to 4%. The threshold is identical across Llama-3-8B, Llama-3-13B and Qwen3-8B, offering a simple, transparent rule: “if self-answerability > 0.7 , trust memory; otherwise examine context.”

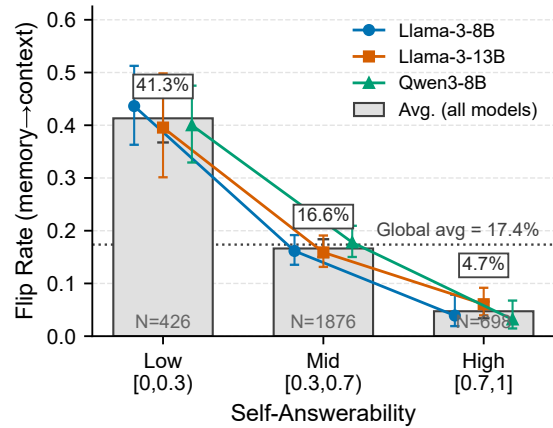


Figure 8: Flip rate for self-answerability bins.

Signal dynamics during decoding. Fig. 9 shows signal trajectories over the first 20 decoding steps for FIXED (corrected) and MISLED cases. In FIXED, factual consistency rises early and surpasses self-answerability by step 7—when the correct answer emerges—while semantic similarity remains high. In contrast, MISLED shows stagnant factual scores and a delayed self-answerability spike, aligning with wrong outputs. This pattern suggests a useful heuristic: $\text{factual} > \text{self-ans}$ before step 10 predicts success in 70% of cases, enabling early-stop or re-prompt strategies.




Figure 9: Average signal trajectories .

7 Conclusion

We propose TCR, a lightweight conflict-aware RAG module that decouples semantic relevance from factual consistency and uses self-answerability to steer generation with three interpretable scalars. It consistently improves detection, factuality, robustness, and interpretability on synthetic and real-world benchmarks, and we plan to extend it to multimodal and federated settings (Li and Cheung 2024; Zhang et al. 2024a, 2023a).

References

- Chen, H.-T.; Zhang, M.; and Choi, E. 2022. Rich knowledge sources bring complex knowledge conflicts: Recalibrating models to reflect conflicting evidence. In *EMNLP*, 2292–2307.
- Chen, J.; Lin, H.; Han, X.; and Sun, L. 2024. Benchmarking Large Language Models in Retrieval-Augmented Generation. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, volume 38(16), 17754–17762.
- Chen, X.; Xiao, P.; Cao, W.; Zhang, W.; and Liu, Y. 2025. Framework and Pathway for the Construction of a Unified Data-Element Market in China. *Strategic Study of Chinese Academy of Engineering*, 27(1): 40–50.
- Diao, X.; Cheng, M.; Barrios, W.; and Jin, S. 2025a. FT2TF: First-Person Statement Text-To-Talking Face Generation. In *WACV*.
- Diao, X.; Zhang, C.; Wu, T.; Cheng, M.; Ouyang, Z.; Wu, W.; and Gui, J. 2024. Learning Musical Representations for Music Performance Question Answering. In *Findings of the Association for Computational Linguistics: EMNLP 2024*.
- Diao, X.; Zhang, C.; Wu, W.; Ouyang, Z.; Qing, P.; Cheng, M.; Vosoughi, S.; and Gui, J. 2025b. Temporal Working Memory: Query-Guided Segment Refinement for Enhanced Multimodal Understanding. In *Findings of the Association for Computational Linguistics: NAACL 2025*.
- Fang, F.; Bai, Y.; Ni, S.; Yang, M.; Chen, X.; and Xu, R. 2024. Enhancing Noise Robustness of Retrieval-Augmented Language Models with Adaptive Adversarial Training. In *ACL*, 10028–10039.
- Gekhman, Z.; Herzig, J.; Aharoni, R.; Elkind, C.; and Szpektor, I. 2023. TrueTeacher: Learning Factual Consistency Evaluation with Large Language Models. In *EMNLP*, 2053–2070.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Guu, K.; Lee, K.; Tung, Z.; Pasupat, P.; and Chang, M. 2020. Retrieval augmented language model pre-training. In *ICML*, 3929–3938. PMLR.
- Huang, X.; Li, R.; Cheung, Y.-m.; Cheung, K. C.; See, S.; and Wan, R. 2024. Gaussianmarker: Uncertainty-aware copyright protection of 3d gaussian splatting. *Advances in Neural Information Processing Systems*, 37: 33037–33060.
- Jang, J.; Ye, S.; Lee, C.; Yang, S.; Shin, J.; Han, J.; Kim, G.; and Seo, M. 2022. Temporalwiki: A lifelong benchmark for training and evaluating ever-evolving language models. *arXiv preprint arXiv:2204.14211*.
- Jiang, L.; Wang, X.; Zhang, F.; and Zhang, C. 2025. Transforming time and space: efficient video super-resolution with hybrid attention and deformable transformers. *The Visual Computer*, 1–12.
- Jin, Z.; Cao, P.; Chen, Y.; Liu, K.; Jiang, X.; Xu, J.; Li, Q.; and Zhao, J. 2024a. Tug-of-War Between Knowledge: Exploring and Resolving Knowledge Conflicts in Retrieval-Augmented Language Models. *arXiv e-prints*, arXiv:2402.
- Jin, Z.; Cao, P.; Yuan, H.; Chen, Y.; Xu, J.; Li, H.; Jiang, X.; Liu, K.; and Zhao, J. 2024b. Cutting Off the Head Ends the Conflict: A Mechanism for Interpreting and Mitigating Knowledge Conflicts in Language Models. In *Findings of the Association for Computational Linguistics ACL 2024*, 1193–1215.
- Kandpal, N.; Deng, H.; Roberts, A.; Wallace, E.; and Raffel, C. 2023. Large Language Models Struggle to Learn Long-Tail Knowledge. In *ICML*, volume 202, 15696–15707. PMLR.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-t.; Rocktäschel, T.; and Riedel, S. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, volume 33, 9459–9474.
- Li, D.; Rawat, A. S.; Zaheer, M.; Wang, X.; Lukasik, M.; Veit, A.; Yu, F.; and Kumar, S. 2023. Large language models with controllable working memory. In *Findings of the association for computational linguistics: ACL 2023*, 1774–1793.
- Li, R.; and Cheung, Y.-m. 2024. Variational multi-scale representation for estimating uncertainty in 3d gaussian splatting. *Advances in Neural Information Processing Systems*, 37: 87934–87958.
- Li, R.; and Cheung, Y.-m. 2025. Modeling and Identifying Distractors with Curriculum for Robust 3D Gaussian Splatting. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 10122–10131.
- Longpre, S.; Perisetla, K.; Chen, A.; Ramesh, N.; DuBois, C.; and Singh, S. 2021. Entity-based knowledge conflicts in question answering. *arXiv preprint arXiv:2109.05052*.
- Min, S.; Lyu, X.; Lewis, P.; Zettlemoyer, L.; Hajishirzi, H.; and Yih, W.-t. 2023. Augmented Language Models: A Survey. *arXiv preprint arXiv:2302.07842*.
- Neeman, E.; Aharoni, R.; Honovich, O.; Choshen, L.; Szpektor, I.; and Abend, O. 2023. Disentqa: Disentangling parametric and contextual knowledge with counterfactual question answering. In *ACL*, 10056–10070.
- Ni, S.; Bi, K.; Guo, J.; and Cheng, X. 2024. When do llms need retrieval augmentation? mitigating llms’ overconfidence helps retrieval augmentation. *arXiv preprint arXiv:2402.11457*.
- Ram, O.; Levine, Y.; Dalmedigos, I.; Muhlgay, D.; Shashua, A.; Leyton-Brown, K.; and Shoham, Y. 2023. In-context retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 11: 1316–1331.
- Shi, D.; Jin, R.; Shen, T.; Dong, W.; Wu, X.; and Xiong, D. 2025. Ircan: Mitigating knowledge conflicts in llm generation via identifying and reweighting context-aware neurons. *Advances in Neural Information Processing Systems*, 37: 4997–5024.
- Slobodkin, A.; Goldman, O.; Caciularu, A.; Dagan, I.; and Ravfogel, S. 2023. The Curious Case of Hallucinatory (Un) answerability: Finding Truths in the Hidden States of Over-Confident Large Language Models. In *EMNLP*, 3607–3625.

- Su, Z.; Zhang, J.; Qu, X.; Zhu, T.; Li, Y.; Sun, J.; Li, J.; Zhang, M.; and Cheng, Y. 2024. Conflictbank: A benchmark for evaluating the influence of knowledge conflicts in llm. *arXiv preprint arXiv:2408.12076*.
- Tao, H.; Li, J.; Hua, Z.; and Zhang, F. 2023. DUBB: deep unfolding-based dual-branch feature fusion network for pan-sharpening remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–17.
- Tong, R.; Liu, J.; Liu, S.; Xu, J.; Wang, L.; and Wang, T. 2025a. Does Bigger Mean Better? Comparative Analysis of CNNs and Biomedical Vision Language Models in Medical Diagnosis. In *ACDSA*, arXiv preprint arXiv:2510.00411, 6.
- Tong, R.; Wei, S.; Liu, J.; and Wang, L. 2025b. Rainbow Noise: Stress-Testing Multimodal Harmful-Meme Detectors on LGBTQ Content. In *NeurIPS 2025: Queer in AI Workshop*.
- Wang, F.; Wan, X.; Sun, R.; Chen, J.; and Arik, S. Ö. 2024a. Astute rag: Overcoming imperfect retrieval augmentation and knowledge conflicts for large language models. *arXiv preprint arXiv:2410.07176*.
- Wang, H.; and Zhang, F. 2024. Computing nodes for plane data points by constructing cubic polynomial with constraints. *Computer Aided Geometric Design*, 111: 102308.
- Wang, Y.; Feng, S.; Wang, H.; Shi, W.; Balachandran, V.; He, T.; and Tsvetkov, Y. 2024b. Resolving Knowledge Conflicts in Large Language Models. In *First Conference on Language Modeling*.
- Wang, Y.; Wang, H.; and Zhang, F. 2025. A Medical image segmentation model with auto-dynamic convolution and location attention mechanism. *Computer Methods and Programs in Biomedicine*, 261: 108593.
- Wei, Z.; Chen, W.-L.; and Meng, Y. 2025. Instruc-tRAG: Instructing Retrieval-Augmented Generation via Self-Synthesized Rationales. In *ICLR*.
- Xiao, C.; Hou, L.; Fu, L.; and Chen, W. 2025. Diffusion-Based Self-Supervised Imitation Learning from Imperfect Visual Servoing Demonstrations for Robotic Glass Installation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 10401–10407. IEEE.
- Xiao, C.; et al. 2024. Confusion-resistant federated learning via diffusion-based data harmonization on non-IID data. *Advances in Neural Information Processing Systems*, 37: 137495–137520.
- Xie, J.; Zhang, K.; Chen, J.; Lou, R.; and Su, Y. 2024. Adaptive Chameleon or Stubborn Sloth: Revealing the Behavior of Large Language Models in Knowledge Conflicts. In *ICLR*.
- Xu, P.; Ping, W.; Wu, X.; McAfee, L.; Zhu, C.; Liu, Z.; Subramanian, S.; Bakhturina, E.; Shoeybi, M.; and Catanzaro, B. 2024a. Retrieval meets Long Context Large Language Models. In *ICLR*.
- Xu, R.; Lin, B.; Yang, S.; Zhang, T.; Shi, W.; Zhang, T.; Fang, Z.; Xu, W.; and Qiu, H. 2024b. The earth is flat because...: Investigating llms’ belief towards misinformation via persuasive conversation. In *ACL*, 16259–16303.
- Xu, Y.; Zhang, R.; Jiang, X.; Feng, Y.; Xiao, Y.; Ma, X.; Zhu, R.; Chu, X.; Zhao, J.; and Wang, Y. 2024c. Parenting: Optimizing knowledge selection of retrieval-augmented language models with parameter decoupling and tailored tuning. *arXiv preprint arXiv:2410.10360*.
- Xue, B.; Wang, W.; Wang, H.; Mi, F.; Wang, R.; Wang, Y.; Shang, L.; Jiang, X.; Liu, Q.; and Wong, K.-F. 2023. Improving Factual Consistency for Knowledge-Grounded Dialogue Systems via Knowledge Enhancement and Alignment. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 7829–7844.
- Yao, J.; Li, C.; Sun, K.; Cai, Y.; Li, H.; Ouyang, W.; and Li, H. 2023. Ndc-scene: Boost monocular 3d semantic scene completion in normalized device coordinates space. In *ICCV*, 9421–9431.
- Yao, J.; Li, C.; and Xiao, C. 2024. Swift sampler: Efficient learning of sampler by 10 parameters. *Advances in Neural Information Processing Systems*, 37: 59030–59053.
- Ying, J.; Cao, Y.; Xiong, K.; Cui, L.; He, Y.; and Liu, Y. 2024. Intuitive or dependent? investigating llms’ behavior style to conflicting prompts. In *ACL*, 4221–4246.
- Yu, Q.; Merullo, J.; and Pavlick, E. 2023. Characterizing Mechanisms for Factual Recall in Language Models. In *EMNLP*, 9924–9959.
- Yu, W.; Iter, D.; Wang, S.; Xu, Y.; Ju, M.; Sanyal, S.; Zhu, C.; Zeng, M.; and Jiang, M. 2023. Generate rather than Retrieve: Large Language Models are Strong Context Generators. In *ICLR*.
- Zhang, F.; Chen, G.; Wang, H.; Li, J.; and Zhang, C. 2023a. Multi-scale video super-resolution transformer with polynomial approximation. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(9): 4496–4506.
- Zhang, F.; Chen, G.; Wang, H.; and Zhang, C. 2024a. CF-DAN: Facial-expression recognition based on cross-fusion dual-attention network. *Computational Visual Media*, 10(3): 593–608.
- Zhang, T.; Patil, S. G.; Jain, N.; Shen, S.; Zaharia, M.; Stolica, I.; and Gonzalez, J. E. 2024b. RAFT: Adapting Language Model to Domain Specific RAG. In *First Conference on Language Modeling*.
- Zhang, X.; Zeng, F.; Quan, Y.; Hui, Z.; and Yao, J. 2025. Enhancing multimodal large language models complex reason via similarity computation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39(10), 10203–10211.
- Zhang, Y.; Khalifa, M.; Logeswaran, L.; Lee, M.; Lee, H.; and Wang, L. 2023b. Merging Generated and Retrieved Knowledge for Open-Domain QA. In *EMNLP*, 4710–4728.
- Zhong, Z. 2025. From Privacy Washing to Sustainable Data Strategies: A Theory-Based AI Approach. *Available at SSRN 5255370*.
- Zhou, W.; Zhang, S.; Poon, H.; and Chen, M. 2023. Context-faithful Prompting for Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 14544–14556.