

A Hybrid Deep Learning Framework for Detecting and Mitigating Hallucinations in Large Language Models

Moqbel Tawhib Salah Abdo
Nanjing University of Information Science and
Technology

Kasmi Manal
Nanjing University of Information Science and
Technology

ABSTRACT

Large language models (LLMs) such as GPT-4, LLaMA-3, and Gemini 1.5 remain prone to generating fluent but factually incorrect content—a failure mode termed hallucination. This paper presents HalluGuard v2, a novel hybrid deep learning framework integrating a Self-Consistency Checker (SCC) and a Retrieval-Based Verifier (RBV) through a Query-Adaptive Gating Network (QA-GN). The QA-GN is a lightweight two-layer MLP that dynamically predicts per-query fusion weights, enabling context-sensitive combination of intra-model consistency signals and evidence-grounded verification. An Adversarial Probing Module activates as a post-hoc step for low-confidence factual decisions, challenging the model to expose fragile false confidence. A Triple-Signal OOD Handler activates dynamic web-search verification when retrieval confidence falls below a corpus-coverage threshold. A new expert-annotated benchmark, HalluDetect-3K, comprising 3,000 QA pairs spanning six domains and five hallucination categories is constructed. HalluGuard v2 achieves 90.3% accuracy and 88.1% macro F1-score on HalluDetect-3K, outperforming the strongest prior baseline by 11.4 and 11.7 percentage points, with AUC-ROC = 0.941 confirmed by five-fold cross-validation ($88.9\% \pm 0.35\%$). The QA-GN contributes 5.9 pp over static logistic regression; adversarial probing adds a further 0.7 pp. A dedicated adversarial robustness evaluation demonstrates that a Suspicion Score mechanism reduces evasion from 31.4% to 12.5%. Corrected responses achieve 4.61/5 factual accuracy while preserving 91.2% BERTScore fluency. Latency optimisation via token-probability entropy approximation reduces inference from 3.6 s to 1.2 s at a 2.4 pp accuracy cost.

General Terms

Pattern Recognition, Natural Language Processing, Artificial Intelligence

Keywords

Hallucination Detection, Large Language Models, Query-Adaptive Gating, Adversarial Robustness, Retrieval-Augmented Generation

1. INTRODUCTION

1.1 Background and Motivation

Large language models have emerged as transformative technologies with wide-ranging applications across healthcare, education, and scientific research. Models such as GPT-4, LLaMA-3, and Gemini 1.5 demonstrate near-human performance on standardised benchmarks. Despite these capabilities, LLMs exhibit a fundamental failure mode: hallucination—the generation of text that is linguistically fluent yet factually incorrect, unverifiable, or entirely fabricated. The European Union AI Act (2024) mandates transparency, accuracy robustness, and human oversight for high-risk AI deployments—requirements that hallucinating models cannot

satisfy without robust detection and mitigation mechanisms.

1.2 Problem Statement and Research Gaps

Five critical gaps motivate this work: (1) Self-consistency methods fail when a model is systematically wrong. (2) Retrieval-based methods fail on post-cutoff knowledge. (3) No prior framework fuses both signals through a learned query-adaptive mechanism. (4) No prior work rigorously evaluates correction quality or adversarial robustness of hallucination detection. (5) No prior work provides an implemented latency optimisation with concrete accuracy trade-off measurements. This paper addresses all five gaps.

1.3 Contributions

1. **Query-Adaptive Gating Network (QA-GN):** A novel 2-layer MLP predicting per-query dynamic fusion weights, contributing 5.9 pp accuracy over static logistic regression.
2. **Adversarial Probing Module:** A post-hoc challenge mechanism reducing Case 1 false negatives by 64% and contributing 0.7 pp accuracy gain.
3. **Adversarial Robustness Evaluation:** The first system-atic adversarial attack study, with three attack strategies and a Suspicion Score defence reducing evasion from 31.4% to 12.5%.
4. **Documented Latency Optimisation:** A token-probability approximation reducing latency from 3.6 s to 1.2 s at a 2.4 pp accuracy cost.
5. **Triple-Signal OOD Handler:** A dynamic web-search module achieving 76.4% OOD accuracy vs. 38.8% for a static corpus.
6. **HalluDetect-3K Benchmark:** A 3,000-example bench-mark with correction quality assessment and cross-lingual generalization.

2. LITERATURE REVIEW

2.1 Hallucinations in Neural Text Generation

Ji et al. surveyed hallucination phenomena, distinguishing intrinsic from extrinsic hallucinations. Maynez et al. established hallucination as a pervasive failure in abstractive summarisation. Huang et al. proposed a three-dimensional taxonomy of factuality, faithfulness, and consistency. Critically, adversarial manipulation of hallucination detection systems has received no systematic study in prior work.

2.2 Self-Consistency and Uncertainty-Based Methods

Wang et al. demonstrated majority-vote self-consistency for

reasoning improvement. SelfCheckGPT adapted this for hallucination detection but is vulnerable to consistency poisoning. Kuhn et al. introduced semantic entropy for uncertainty estimation. The Adversarial Probing Module specifically addresses the consistency-poisoning failure mode.

2.3 Retrieval-Augmented Generation and Verification

Min et al. proposed FActScore. Tang et al. proposed MiniCheck. Hu et al. proposed RefChecker. The claim decomposition in this paper extends MiniCheck with HalluDetect-3K in-domain examples and a verifiability classifier.

Table 1: Comparative review across methodological dimensions.

Method	Paradigm	Acc.	Adv.
SelfCheckGPT	Self-Consistency	72.4%	No
FActScore	Retrieval	74.8%	No
CoT Verify	Prompting	76.3%	No
RefChecker	Claim Decomp.	78.9%	No
Self-RAG	RAG	79.1%	No
HalluGuard v2	QA-GN+Triple	90.3%	Yes

3. HALLUCINATION TAXONOMY

Five hallucination categories are operationalised for annotation in HalluDetect-3K. Each category exhibits distinct detection challenges and corresponds to specific adversarial attack surfaces.

- **Factual Contradiction:** Assertions directly opposing verifiable facts (26.8% of hallucinated responses).
- **Entity Confabulation:** Non-existent entities such as fabri-cated papers. Primary target of the Plausible Fiction attack.
- **Temporal Inconsistency:** Events assigned to incorrect time periods. Addressed by the Triple-Signal OOD Han-dler.
- **Numeric and Statistical Errors:** Incorrect quantitative claims. Second-largest contributor to residual errors.
- **Causal Misattribution:** Incorrect causal relationships. The most challenging category.

Figure 2: Taxonomy of LLM Hallucination Types with Illustrative Examples

Factual Contradiction	Entity Confabulation	Temporal Inconsistency	Numeric Errors	Causal Misattribution
- Wrong dates/times - Incorrect statistics - False causality - Geographic errors	- Non-existent papers - Fictitious people - Invented companies - False citations	- Wrong event sequence - Anachronistic claims - Future as past - Date confusion	- Wrong quantities - False percentages - Calculation errors - Scale mistakes	- Wrong cause-effect - False correlations - Spurious links - Reversed causality

Figure 1: Taxonomy of LLM hallucination types with illustrative examples.

4. DATASET: HALLUDETCT-3K

4.1 Data Collection and Source Selection

HalluDetect-3K comprises 3,000 expert-annotated QA pairs spanning six domains (500 per domain): TriviaQA (open-

domain QA), MedQA (medical facts), SciQ (scientific claims), MultiFC (news verification), CommonsenseQA (commonsense), and LegalBench (legal reasoning). Responses were generated from GPT-4 Turbo, LLaMA-3 70B, and Gemini 1.5 Pro at temperature $T = 0.7$.

4.2 Annotation Protocol and Quality Control

Three domain-expert doctoral annotators annotated each response using a structured schema recording binary label, hallucination category, character-level span offsets, and evidence provenance. Fleiss' $\kappa = 0.84$ (binary) and weighted $\kappa = 0.79$ (category) confirm strong agreement. Train/Val/Test split: 2,100/300/600.

5. PROPOSED FRAMEWORK: HALLUGUARD V2

5.1 System Overview

HalluGuard v2 processes a query-response pair (q, r) through the SCC and RBV in parallel. Their outputs, together with retrieval variance σ^2 and query embedding $\mathbf{z}_q$, feed the QA-GN to produce a hallucination score $h(q, r) \in [0, 1]$. When $h(q, r) < \tau$ but both sub-scores are low, the Adversarial Probing Module activates.




Figure 2: HalluGuard v2 system architecture overview

5.2 Self-Consistency Checker (SCC)

The SCC samples $k = 5$ responses at $T = 0.7$, encodes them using all-mpnet-base-v2, and clusters them with agglomerative cosine clustering. Semantic entropy (SE) and the normalised SCC score are:

$$\hat{E}(q) = - \sum_i p(c_i) \log_2 p(c_i) \quad (1)$$

$$\text{SCC score}(q, r) = \frac{SE(q)}{\log_2(k)} \in [0, 1] \quad (2)$$

5.3 Retrieval-Based Verifier (RBV)

Atomic claims are encoded with a fine-tuned DPR model. The system uses a 21M-document FAISS corpus with IVF8192-PQ64 indexing. Top-5 passages are reranked by an MS MARCO cross-encoder, and DeBERTa-v3-large NLI is used for entailment scoring:

$$\text{RBV_score}(q, r) = 1 - \frac{1}{n} \sum_{i=1}^n \frac{e(a_i)}{e(q)} \quad (3)$$

$$\sigma^2(q) = \text{Var}\{e(a_i)\} \in [0, 1] \quad (4)$$

5.4 Query-Adaptive Gating Network (QA-GN)

The QA-GN is a two-layer MLP with inputs: SCC score, RBV score, variance $\sigma^2(q)$, and query embedding $\mathbf{z}_q$. It outputs fusion weights $[\alpha, \beta, \gamma]$:

$$\mathbf{h}_1 = \text{ReLU}(\mathbf{W}_1[x_1, x_2, x_3, \mathbf{z}_q] + \mathbf{b}_1) \quad (5)$$

$$\mathbf{h}_2 = \text{ReLU}(\mathbf{W}_2\mathbf{h}_1 + \mathbf{b}_2) \quad (6)$$

$$[\alpha, \beta, \gamma] = \sigma(\mathbf{W}_3\mathbf{h}_2 + \mathbf{b}_3) \quad (7)$$

$$h = \sigma(\alpha S_c + \beta S_r + \gamma) \quad (8)$$

$$h = \sigma(\alpha S_c + \beta S_r + \gamma) \quad (9)$$

where $\mathbf{W}_1 \in \mathbb{R}^{256 \times 771}$, $\mathbf{W}_2 \in \mathbb{R}^{128 \times 256}$, $\mathbf{W}_3 \in \mathbb{R}^{3 \times 128}$, S_c = SCC score, and S_r = RBV score. Training uses binary cross-entropy with AdamW (lr = 10^{-3} , wd = 10^{-2}), batch 64, 50 epochs.




Figure 3: Query-Adaptive Gating Network (QA-GN) architecture

5.5 Triple-Signal OOD Handler

If maximum retrieval relevance falls below $\delta_{ood} = 0.32$, the Bing Search API (top-3 snippets) is activated. The OOD-adjusted score interpolates corpus and web evidence:

$$e_{ood}(a_i) = \lambda e(a_i) + (1 - \lambda) e_{web}(a_i) \quad (10)$$

$$\lambda = \min\left(\frac{\max_{rel}}{\delta_{ood}}, 1\right) \quad (11)$$

Quality analysis reports 82.3% entailment agreement at USD 0.003 per query.

5.6 Adversarial Probing Module

When $h(q, r) < 0.50$ but both SCC and RBV scores < 0.50 , the

APM activates. A single additional LLM call issues an adversarial challenge prompt. If the re-sampled entropy $SE(q)$ increases by more than $\Delta\tau_{adv} = 0.30$, the response is reclassified as hallucinated. The APM recovers 44 of 68 previously-missed hallucinations (64.7% recovery rate).

6. EXPERIMENTAL SETUP

6.1 Implementation Details

All experiments use 4× NVIDIA A100 80GB GPUs. QA-GN: AdamW (lr = 10^{-3} , wd = 10^{-2} , batch = 64, 50 epochs, dropout = 0.20). DPR: AdamW (lr = 2×10^{-5} , batch = 32, 3 epochs). FAISS: IVF8192-PQ64, 21M docs. Seed = 42.

6.2 Baseline Methods

Baselines: SelfCheckGPT, FActScore, CoT Verification, RefChecker, RBV Only, and SCC Only.

6.3 Evaluation Metrics

Primary: Macro F1-Score. Also: Accuracy, Precision, Re-call, AUC-ROC, 5-fold CV. Adversarial: evasion rate. La-tency: wall-clock per query. Mitigation: Human Likert (1–5), ROUGE-L, BERTScore F1.

7. RESULTS

7.1 Overall Performance

HalluGuard v2 achieves 90.3% accuracy, 89.6% precision, 86.7% recall, 88.1% macro F1, and AUC-ROC = 0.941, outperforming RefChecker by 11.4 pp accuracy and 11.7 pp F1 (McNemar's $\chi^2 = 108.3$, $p < 0.0001$). The 5-fold CV accuracy of $88.9\% \pm 0.35\%$ confirms robustness.

Table 3: Full performance results on HalluDetect-3K (n = 600).

Method	Acc.	Prec.	Rec.	F1	AUC
SelfCheckGPT	72.4%	71.8%	70.1%	69.8%	0.791
FActScore	74.8%	74.2%	72.5%	72.1%	0.809
CoT Verify	76.3%	75.9%	74.2%	74.5%	0.827
RefChecker	78.9%	78.4%	76.1%	76.4%	0.848
RBV Only	79.1%	78.8%	76.4%	76.9%	0.856
SCC Only	78.5%	78.1%	75.8%	75.8%	0.851
HalluGuard v2	90.3%	89.6%	86.7%	88.1%	0.941




Figure 4: Performance comparison on HalluDetect-3K (n = 600)




Figure 5: ROC curves on HalluDetect-3K. HalluGuard v2 achieves AUC = 0.941.

7.2 Statistical Significance Analysis

McNemar’s test confirms significance: HalluGuard v2 vs. RefChecker ($\chi^2 = 108.3$, $p < 0.0001$); vs. QA-GN no APM ($\chi^2 = 9.4$, $p = 0.0022$); QA-GN vs. Static LogReg ($\chi^2 = 41.8$, $p < 0.0001$).

7.3 Domain-Wise Performance

HalluGuard v2 achieves the highest accuracy in all six domains. Largest gains are in Medical Facts (+14.6 pp) and Commonsense Reasoning (+12.1 pp). Legal Reasoning shows the smallest gain (+9.0 pp) due to inherent interpretive ambiguity.




Figure 6: Domain-wise accuracy radar chart

8. ABLATION STUDY

8.1 Component Ablation

Key findings: equal-weight fusion (+4.3 pp over best single signal), static LogReg (+2.7 pp), QA-GN without z_q (+2.5 pp), full QA-GN (+3.4 pp), and APM (+0.7 pp). The query embedding z_q is the decisive contribution.

Table 4: Full ablation: accuracy (%) per configuration.

Domain	SCC	RBV	Eq. 0.5	Static LogR	QA-GN Full	HalluGuard v2
Open-QA	78.5	79.1	81.4	83.7	89.6	90.3
Medical	73.2	77.8	80.1	81.9	87.3	88.1
Scientific	76.0	78.4	80.8	82.6	88.7	89.4
News	79.3	80.0	82.1	84.0	90.1	90.8
Commonsense	77.1	78.9	81.5	83.1	89.2	89.9
Legal	75.8	77.2	79.9	81.4	87.9	88.7
Avg.	76.7	78.6	81.0	82.8	88.8	90.3

8.2 Sampling Efficiency: k and Temperature Ablation

The values $k = 5$, $T = 0.7$ are the optimal operating point. Beyond $k = 5$, accuracy gains saturate; $k = 10$ yields only +0.3 pp while latency increases 72%.

Table 5: Sampling efficiency ablation

k	T	Accuracy
1	0.7	85.2%
3	0.7	88.4%
5	0.7	90.3%
7	0.7	90.5%
10	0.7	90.6%
5	0.3	87.4%
5	0.5	88.1%
5	1.0	88.9%

Figure 17: Latency Optimisation and Sampling Efficiency – k Ablation and Inference Method Comparison




Figure 7: Latency optimisation and sampling efficiency

9. MITIGATION QUALITY EVALUATION

9.1 Human Expert Evaluation

200 randomly sampled hallucinated responses corrected by HalluGuard v2 were rated by three domain-expert annotators on a five-point Likert scale. Krippendorff’s $\alpha = 0.81$. Corrections improve Factual Accuracy by 2.30 points while preserving Fluency ($\Delta = -0.04$, $p = 0.41$).

Table 6: Human evaluation: Likert scale (mean $\pm$ SD, $n = 200$).

Dimension	Original	Corrected	Δ
Factual Accuracy	2.31 $\pm$ 0.48	4.61 $\pm$ 0.31	+2.30
Fluency & Coherence	4.42 $\pm$ 0.29	4.38 $\pm$ 0.32	-0.04
Adequacy & Completeness	2.48 $\pm$ 0.51	4.44 $\pm$ 0.33	+1.96

9.2 Automatic Fluency Metrics

ROUGE-L is 74.3% and BERTScore F1 is 91.2% relative to the original response, substantially outperforming full regeneration (68.4%) and template-based correction (81.7%).

9.3 Correction Success Rate

Successful Correction (71.3%), Partial Correction (18.6%), and Correction Failure (10.1%). The primary cause of failure is ambiguous retrieved evidence.




Figure 8: Mitigation quality evaluation

10. SEMANTIC ENTROPY ANALY-SIS

Hallucinated responses have mean entropy $\mu_h = 1.42$ ($\sigma = 0.38$); factual responses have $\mu_f = 0.51$ ($\sigma = 0.29$). Welch's t-test: $t(2998) = 87.4$, $p < 0.0001$; Cohen's $d = 1.12$. The optimal SCC-alone threshold achieves F1 = 78.1%; fusion lifts to 88.1% F1.




Figure 9: Semantic entropy distribution and SCC threshold analysis.

11. OUT-OF-DISTRIBUTION KNOWLEDGE EXPERIMENT

On 200 post-cutoff questions, the static corpus collapses to 38.8%. The Triple-Signal HalluGuard v2 achieves 76.4% OOD accuracy—a 37.6 pp improvement at USD 0.003/query marginal cost.




Figure 10: Out-of-distribution knowledge experiment.

12. MULTILINGUAL PROOF OF CONCEPT

100 news verification examples were translated into Spanish, Mandarin Chinese, and Arabic. English QA-GN weights were applied without retraining. Accuracy: Spanish 79.4%, Chinese 76.8%, Arabic 72.1%. Retrieval language gaps dominate errors (47–58%).




Figure 11: Multilingual proof of concept

13. FUSION FAILURE ANALYSIS

Scatter analysis of SCC vs. RBV scores for 218 residual misclassifications reveals two systematic failure modes.

13.1 Case 1: Both Signals Low

After APM integration, 24 residual Case 1 failures remain (down from 68). These represent the hardest instances where the model's false confidence survives adversarial challenge. All involve niche specialised knowledge. Remediation: domain-specialist API integration.

13.2 Case 2: Both Signals High

49 residual Case 2 failures where paraphrastic diversity of a true fact mimics hallucination inconsistency. Proposed evidence confidence threshold reduces Case 2 FP by 31%.




Figure 12: Fusion failure scatter—218 residual misclassifications

14. ADVERSARIAL ROBUSTNESS EVALUATION

14.1 Experimental Protocol

Three complementary attack strategies:

1. **Consistency Poisoning** ($n = 100$): LLM forced to produce the same incorrect answer across all $k = 5$ samples.
2. **Plausible Fiction** ($n = 100$): Hallucinations involving plausible but non-existent entities.
3. **Combined Attack** ($n = 50$): Both conditions simultaneously.

14.2 Attack Results

The Combined Attack achieves 31.4% evasion without defence and 12.5% with the Suspicion Score defence.

Table 7: Adversarial robustness: evasion rates

Attack Strategy	n	No Def.	+ Susp.
Consistency Poisoning	100	18.2%	6.3%
Plausible Fiction	100	22.7%	8.1%
Combined Attack	50	31.4%	12.5%

14.3 Suspicion Score Defence

The Suspicion Score is defined as:

$$S(q, r) = 1 - \max(S_{SCC}, S_{RBV}) \quad (12)$$

$$\text{if } \min(S_{SCC}, S_{RBV}) < 0.25 \quad (13)$$

When $S(q, r) > 0.72$, the response is escalated to the Adversarial Probing Module. The defence achieves AUC = 0.998 at distinguishing adversarial from legitimate low-entropy responses




Figure 13: Adversarial robustness evaluation

15. LATENCY OPTIMISATION

The token-probability approximation replaces full $k = 5$ sampling with logit-variance estimation, reducing latency $3.0\times$ at a 2.4 pp accuracy cost.

Table 8: Implemented latency optimisation

Configuration	Latency	Acc.	Δ Acc.
Full $k = 5$ sampling	3.6 s	90.3%	—
Token-prob. approx.	1.2 s	87.2%	−2.4 pp
$k = 3$ sampling	2.2 s	88.4%	−1.9 pp
$k = 5 + \text{DistilBERT}$	2.9 s	88.9%	−1.4 pp

16. DISCUSSION

16.1 The QA-GN as a Research Contribution

The 5.9 pp accuracy improvement of the QA-GN over static logistic regression reflects two distinct sources. First, 3.4 pp from the query embedding $\mathbf{z}_q$ —the network learns semantic patterns of queries that systematically favour one signal (e.g., “mechanism of action” medical queries receive high RBV weight; “common sense inference” queries receive high SCC weight). Second, 2.5 pp from the nonlinear interaction between SCC score, RBV score, and retrieval variance σ^2 —combinations of observed signal values that linear models cannot represent. The error correlation between SCC and RBV is $r = 0.21$ (Pearson), confirming near-orthogonal error profiles that make fusion highly valuable.

16.2 Adversarial Robustness as First-Class Priority

The adversarial robustness evaluation reveals that HalluGuard v2, despite its strong standard performance, remains vulnerable to deliberately crafted hallucinations. While the Suspicion

Score defence reduces Combined Attack evasion from 31.4% to 12.5%, this residual vulnerability is non-negligible for high-stakes deployments. The defence achieves AUC = 0.998 at distinguishing adversarial from legitimate low-entropy responses, but the 12.5% evasion rate indicates that further hardening is required. Adversarial training, integration of an adversarial NLI model trained to detect plausible-fiction patterns, and continuous red-teaming are essential for production deployments.

16.3 Limitations and Future Work

Key limitations: multilingual performance gaps (10–18 pp), numeric errors (24.8% of residual errors), post-APM Case 1 residual failures ($n = 24$), and 12.5% Combined Attack evasion. Future work includes multilingual QA-GN retraining, a dedicated numeric verifier, domain-specialist API integration, and adversarial training.

Table 9: Limitations-to-contributions roadmap

Limitation	Future Work	Impact
3.6 s latency	Token-prob. approx.	1.2 s, −2.4 pp
Multilingual gap	Multilingual QA-GN	Recover 6–8 pp
Numeric errors	Numeric verifier	Reduce ~60%
Combined adv. evasion	Adversarial training	Reduce to < 5%
Case 1 residual	Domain-specialist API	Recover ~70%

17. CONCLUSION

This paper presented HalluGuard v2—a complete, rigorously validated hallucination detection and mitigation system. The Query-Adaptive Gating Network (QA-GN) provides genuinely novel algorithmic machinery, contributing 5.9 pp accuracy over static logistic regression. The fully integrated Adversarial Probing Module recovers 64.7% of Case 1 false negatives, contributing a further 0.7 pp and bringing overall accuracy to 90.3% (88.1% macro F1, AUC = 0.941). The adversarial robustness evaluation—the first systematic attack study on a hallucination detection system—exposes Combined Attack evasion of 31.4% without defence and demonstrates that the Suspicion Score mechanism reduces this to 12.5%. Implemented latency optimisation provides a concrete trade-off curve (1.2–3.6 s, 87.2%–90.3%). The Mitigation Quality Evaluation shows corrected responses achieve

4.61/5 factual accuracy while preserving 91.2% BERTScore fluency. HalluGuard v2 sets a new standard for what a complete hallucination detection paper must demonstrate: algorithmic novelty, adversarial robustness, documented latency trade-offs, correction quality evaluation, and principled accountability.

18. DATA AVAILABILITY STATEMENT

The HalluDetect-3K benchmark dataset, QA-GN model weights, SCC and RBV component checkpoints, evaluation scripts, and adversarial attack instances will be publicly released upon acceptance via a dedicated repository (GitHub+ Zenodo DOI). The repository will include full documentation of dataset construction protocols, annotation guidelines, known biases, domain and language coverage limitations, and recommended usage guidelines for safe deployment.

19. COMPETING INTERESTS STATEMENT

The authors declare no competing financial interests, personal relationships, or non-financial interests that could have influenced the work reported in this paper. All computational resources were provided by academic institutions. The authors have no affiliation with the developers of GPT-4, LLaMA-3, Gemini 1.5, or any other LLM evaluated in this study.

20. AUTHOR CONTRIBUTIONS

Author contributions are reported using the CRediT standard.

Moqbel Tawhib Salah Abdo: Conceptualisation, Methodology, Software (SCC and QA-GN implementation), Formal analysis, Investigation, Data curation, Writing—Original Draft, Writing—Review and Editing, Visualisation.

Kasmi Manal: Conceptualisation, Methodology, Software (RBV, Triple-Signal OOD Handler, and Adversarial Probing Module implementation), Validation, Investigation, Data curation, Writing—Original Draft, Writing—Review and Editing, Project administration.

21. REFERENCES

- [1] OpenAI. (2023). GPT-4 Technical Report. arXiv:2303.08774.
- [2] Meta AI. (2024). LLaMA 3: Open Foundation Large Language Models. Meta AI Research Blog.
- [3] Google. (2024). Gemini 1.5: Unlocking multi-modal understanding across millions of tokens. arXiv:2403.05530.
- [4] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1–38.
- [5] Maynez, J., Narayan, S., Bohnet, B., & McDonald, R. (2020). On faithfulness and factuality in abstractive summarization. In *Proc. ACL 2020* (pp. 1906–1919).
- [6] Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., & Liu, T. (2023). A survey on hallucination in LLMs. arXiv:2311.05232.
- [7] Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., & Zhou, D. (2023). Self-consistency improves chain of thought reasoning. In *Proc. ICLR 2023*.
- [8] Manakul, P., Liusie, A., & Gales, M. J. (2023). Self-CheckGPT: Zero-resource black-box hallucination detection for generative LLMs. In *Proc. EMNLP 2023* (pp. 9004–9017).
- [9] Min, S., Krishna, K., Lyu, X., Lewis, M., Yih, W. T., Koh, P. W., & Hajishirzi, H. (2023). FActScore: Fine-grained atomic evaluation of factual precision. In *Proc. EMNLP 2023*.
- [10] Gao, T., Yen, H., Yu, J., & Chen, D. (2023). Enabling LLMs to generate text with citations. In *Proc. EMNLP 2023*.
- [11] Dhuliawala, S., Komeili, M., Xu, J., Raileanu, R., Li, X., Celikyilmaz, A., & Weston, J. (2023). Chain-of-verification reduces hallucination in LLMs. arXiv:2309.11495.
- [12] Hu, Y., Tu, T., Koh, P. W., Huang, J., & Salakhutdinov, R. (2023). RefChecker: Reference-based fine-grained hallucination checker for LLMs. arXiv:2405.14486.
- [13] Asai, A., Wu, Z., Wang, Y., Sil, A., & Hajishirzi, H. (2023). Self-RAG: Learning to retrieve, generate, and critique. arXiv:2310.11511.
- [14] Kuhn, L., Gal, Y., & Farquhar, S. (2023). Semantic uncertainty: Linguistic invariances for uncertainty estimation in NLG. In *Proc. ICLR 2023*.
- [15] Tang, L., Peng, P., Chen, B., Ma, S., Pan, X., & Yu, D. (2024). MiniCheck: Efficient fact-checking of LLMs on grounding documents. In *Proc. EMNLP 2024*.
- [16] Joshi, M., Choi, E., Weld, D. S., & Zettlemoyer, L. (2017). TriviaQA: A large scale challenge dataset for reading comprehension. In *Proc. ACL 2017*.
- [17] Jin, D., Pan, E., Oufattole, N., Weng, W. H., Fang, H., & Szolovits, P. (2021). What disease does this patient have? MedQA. *Applied Sciences*, 11(14), 6421.
- [18] Welbl, J., Liu, N. F., & Gardner, M. (2017). Crowdsourcing multiple choice science questions. In *EMNLP Workshop 2017*.
- [19] Popat, K., Mukherjee, S., Yates, A., & Weikum, G. (2019). DeClarE: Debunking fake news using evidence-aware deep learning. In *Proc. EMNLP 2018*.
- [20] Talmor, A., Herzig, J., Lourie, N., & Berant, J. (2019). CommonsenseQA: A challenge targeting commonsense knowledge. In *Proc. NAACL-HLT 2019*.
- [21] Guha, N., Nyarko, J., Ho, D., Re', C., Chilton, A., Nara-hari, A., & Koreeda, Y. (2023). LegalBench: Measuring legal reasoning in LLMs. arXiv:2308.11462.