knowledge

MDPI

Review

Augmenting Large Language Models with External Data Sources: A Systematic Review of Methodologies, Performance Metrics, and Information Fidelity

Soham Mukherjee *, John Le * and Chau Nguyen

School of Computing and Information Technology, University of Wollongong, Wollongong, NSW 2522, Australia; chaun@uow.edu.au

* Correspondence: sm765@uowmail.edu.au (S.M.); johnle@uow.edu.au (J.L.)

Abstract

Large Language Models (LLMs) have emerged as transformative tools across various domains, exhibiting remarkable capabilities in natural language processing and generation. However, their reliance on static pre-training data limits their ability to access up-to-date and domain-specific information. The existing research often treats augmentation strategies in isolation, and limited efforts have been made to systematically compare them through the lens of information integrity. This review focuses specifically on Retrieval-Augmented Generation (RAG) and fine-tuning, identifying them as the two dominant paradigms for integrating external knowledge: RAG for retrieval-based context injection and fine-tuning for parametric knowledge adaptation. While existing surveys predominantly focus on performance metrics like accuracy or latency, this paper addresses the critical gap of data fidelity—the preservation of truthfulness, integrity, and fairness during augmentation. We systematically synthesize empirical findings from diverse methodologies to determine how each approach mitigates hallucinations and bias. By comparing the trade-offs between retrieval-based context injection and parametric knowledge adaptation, this survey provides unique value to readers by providing a structured taxonomy, a unified evaluation framework, and actionable insights to guide future research and practical deployment of robust, high-fidelity LLMs.

Keywords: large language models; Retrieval-Augmented Generation; fine-tuning; data fidelity; Information Retrieval; AI safety; factuality

Check for updates

Academic Editor: Sotiris Kotsiantis

Received: 2 April 2026
Revised: 14 June 2026
Accepted: 17 June 2026
Published: 25 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Large Language Models (LLMs), such as those detailed by Zhang et al. [1] and Sun et al. [2], signify a substantial leap in artificial intelligence. While these models demonstrate sophisticated capabilities in natural language processing (NLP), their reliance on static pre-training corpora introduces a significant "knowledge cut-off." This limitation renders LLMs incapable of processing real-time information, frequently leading to the generation of factually incorrect or outdated content—a phenomenon widely recognized as "hallucination" [3,4].

To address these limitations, two primary augmentation paradigms have emerged: Retrieval-Augmented Generation (RAG) and fine-tuning. RAG systems, as described by Xue et al. [3] and Chauhan et al. [5], mitigate the knowledge cut-off by retrieving relevant external document chunks to condition the model's generation. Conversely, fine-tuning

involves continued training on domain-specific datasets to adjust internal parameters, deeply embedding specific knowledge or stylistic nuances [6]. While advanced variations exist—such as Dynamic RAG [4] and Knowledge-Graph enhanced pre-training [1], current research largely evaluates these methods based on surface-level performance metrics rather than the integrity of information assimilation. Figure 1 provides a timeline illustrating the rapid evolution of these LLM augmentation methodologies.

A critical oversight in existing literature is the lack of comprehensive analysis through the lens of Data Fidelity. Current evaluations often reveal a misalignment between retrieval metrics and downstream performance [7], posing a risk of degrading public knowledge if AI outputs are uncritically accepted as a single source of truth [8].

In this review, we explicitly define Data Fidelity as a multidimensional construct that encompasses three interactive layers: the properties of the external data sources (accuracy and curation), the interaction between the retrieval mechanism and the model (contextual relevance and grounding), and the properties of the final model outputs (truthfulness and structural coherence). We establish that Data Fidelity is operationalized through specific quantitative metrics mapped to three core pillars:

1. Truthfulness and Factuality: The minimization of stochastic hallucinations and adherence to verifiable facts, operationalized via factual consistency scoring and hallucination rates.
2. Integrity and Curation: The resilience of the system against data poisoning [3] and context degradation, measured via retrieval accuracy (MRR) and grounding correctness (Exact Match).
3. Representational Fairness: The mitigation of algorithmic biases inherited from source data or amplified during augmentation [9], quantified through demographic parity and equalized odds indicators.

Prior reviews typically focus on isolated aspects of the ecosystem, such as datasets [8], general data-augmentation methods [10], RAG-only architectures [11], or broad augmented LLM frameworks without a specific focus on integrity. In contrast, this review brings distinct value to the reader by offering:

- A Novel Perspective: Unifying disparate augmentation methods (RAG, fine-tuning, Hybrid) under the singular, critical umbrella of Data Fidelity and safety.
- A Structured Taxonomy: Clarifying the relationships, evolution, and architectural nuances of modern approaches through conceptual diagrams and structured categorization.
- Actionable Insights: Providing a synthesized comparative framework to guide researchers and practitioners in selecting the optimal architecture based on their specific constraints (e.g., latency vs. provenance).

To this end, this survey addresses the following review questions, mapped to the corresponding sections of this paper:

- RQ-A: What are the main methodologies for integrating external data into LLMs (addressed in Section 4, providing a taxonomy of RAG, fine-tuning, and Hybrid approaches).
- RQ-B: How do RAG, fine-tuning, and Hybrid strategies compare in terms of fidelity and performance (addressed in Section 6, which synthesizes comparative performance metrics)?
- RQ-C: What open challenges and research gaps remain for developing trustworthy and high-fidelity augmented LLMs (addressed in Section 7, outlining future research categories)?



Figure 1. A timeline illustrating the rapid evolution of LLM augmentation methodologies, tracing the development from early parameter-heavy fine-tuning to modern Agentic RAG and dynamic Hybrid systems.

The remainder of this paper is structured to follow a clear logical progression from foundational theory to practical application: Section 2 establishes the Conceptual Framework motivating augmentation and defining fidelity threats. Section 3 outlines the systematic review process. Section 4 presents a Method Taxonomy, categorizing the core

architectures. Section 5 outlines the Evaluation Metrics used to measure success. Section 6 conducts a Comparative Analysis, synthesizing the findings to evaluate strengths and weaknesses. Finally, Section 7 outlines future directions, and Section 8 concludes.

2. Conceptual Framework and Background

The objective of this section is to establish the theoretical foundations of LLM augmentation and delineate the cross-cutting concerns that threaten data fidelity. Understanding these underlying vulnerabilities is crucial before evaluating the technical methods designed to address them. The three-pillar framework of Data Fidelity—comprising Truthfulness, Integrity, and Fairness—was derived through a thematic synthesis of these core vulnerabilities identified across the current literature landscape.

2.1. The Need for Augmentation

Standard LLMs are limited by their training data cut-off and the stochastic nature of token generation. Hallucination, which is the generation of plausible but incorrect information, remains a primary vulnerability [12]. Research indicates that an LLM's internal state may "know" when it is fabricating information, yet the token-by-token generation process can still result in a false output [12]. Augmentation strategies aim to ground these models in external reality. RAG achieves this via a non-parametric memory (external databases), while fine-tuning relies on parametric updates. However, both methods introduce unique vectors for information loss and fidelity degradation. A high-level overview of these external data augmentation strategies is presented in Figure 2.

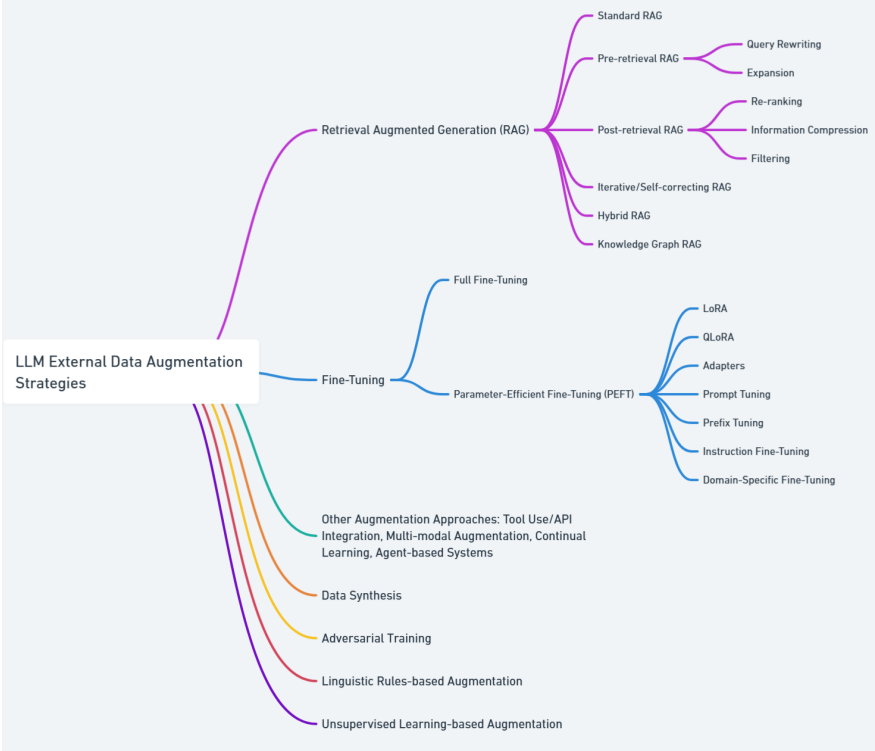


Figure 2. Overview of LLM external data augmentation strategies. The different colors delineate the primary taxonomic categories, distinguishing retrieval-based methods (purple) from parameter-based fine-tuning (blue) and other alternative approaches.

2.2. Formal Definition of Data Fidelity

The term data fidelity as employed in this review is not a novel atomic concept but a composite, system-level property defined over established constructs in the liter-

ature. This subsection provides a formal definition and characterizes the nature of its constituent pillars.

2.2.1. Disambiguating Data Fidelity from Related Constructs

Several closely related but narrower constructs appear throughout the augmented-LLM literature. To prevent conflation, we draw the following distinctions:

- Faithfulness (attribution): whether a generated claim is logically entailed by the source passage provided to the model [11,13]. It is a generation-level property.
- Factual Consistency: the degree to which generated statements align with verifiable real-world facts, regardless of whether those facts were supplied in the prompt [14].
- Grounding: the mechanism by which a model anchors its outputs to retrieved evidence, typically assessed through citation accuracy and source traceability [7].
- Trustworthiness: the broadest term, encompassing safety, privacy, fairness, robustness, and ethical alignment, as evaluated by frameworks such as TrustLLM [15].

Data fidelity is distinguished from these constructs in two respects. First, it is pipeline-spanning: rather than evaluating a single stage (retrieval or generation), it tracks information integrity across the full augmentation pipeline, from source curation through retrieval interaction to final output. Second, it is operationalized through three interdependent pillars (Truthfulness, Integrity, and Representational Fairness) that jointly capture how information is preserved, distorted, or lost during augmentation.

To explicitly clarify these boundaries, Table 1 maps related constructs to our framework, while Figure 3 provides a visual mapping that positions Data Fidelity relative to existing evaluation concepts.

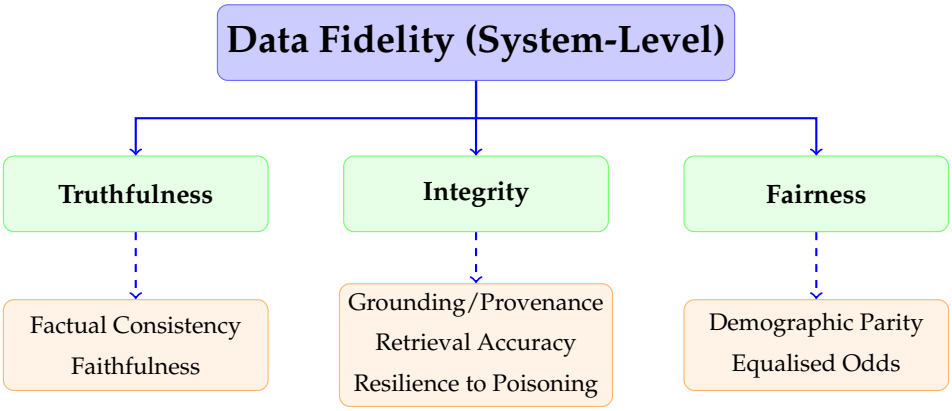


Figure 3. Visual mapping positioning Data Fidelity as the overarching system-level construct. It encompasses three core pillars (Truthfulness, Integrity, Fairness), which are structurally distinct from, but operationalized through, established generative and retrieval-based evaluation concepts.

Table 1. Mapping of related constructs to the Data Fidelity framework.

Construct	Scope	Stage Evaluated	Relation to Data Fidelity
Faithfulness	Output entails source	Generation	Subset of Truthfulness pillar
Factual Consistency	Output matches world facts	Generation	Core metric of Truthfulness pillar
Grounding	Output cites retrieved evidence	Generation—Retrieval	Core metric of Integrity pillar
Trustworthiness	Safety, fairness, robustness	System-wide	Partial overlap; extends beyond fidelity
Data Fidelity	Information preservation across augmentation	End-to-end pipeline	Unifying construct (this review)

2.2.2. Compositional Definition

Let an augmented LLM system be defined by a retrieval component $\mathcal{R}$, a generation component $\mathcal{G}$, and an external knowledge source $\mathcal{K}$. Because Data Fidelity is an emergent property of the entire pipeline, we formally define Data Fidelity $\mathcal{F}$ as a composite function of three pillar indices, where each pillar is inherently dependent on the interaction of the system's sub-components:

$$\mathcal{F} = f(T(\mathcal{R}, \mathcal{G}, \mathcal{K}), I(\mathcal{R}, \mathcal{G}, \mathcal{K}), \Phi(\mathcal{R}, \mathcal{G}, \mathcal{K})) \tag{1}$$

where:

- T (Truthfulness) captures the degree to which generated outputs are entailed by the retrieved context and consistent with verifiable world knowledge. It is defined as a function of faithfulness and factual consistency:

$$T = f_T(\text{faithfulness}(\mathcal{G}, \mathcal{R}(\mathcal{K})), \text{factual_consistency}(\mathcal{G})) \tag{2}$$

- I (Integrity) captures the preservation of information across the pipeline, including retrieval precision, correct citation, and resilience to data poisoning or context degradation:

$$I = f_I(\text{grounding}(\mathcal{G}, \mathcal{R}), \text{retrieval_accuracy}(\mathcal{R}, \mathcal{K}), \text{resilience}(\mathcal{K})) \tag{3}$$

- Φ (Representational Fairness) captures the mitigation of biases inherited from $\mathcal{K}$ or amplified during retrieval and generation:

$$\Phi = f_\Phi(\text{demographic_parity}(\mathcal{R}, \mathcal{G}), \text{equalised_odds}(\mathcal{R}, \mathcal{G})) \tag{4}$$

While specific mathematical aggregation functions (such as threshold-gated geometric means) and the empirical calibration of these parameters present a critical direction for future technical research, this compositional structure provides the theoretical scaffold required to systematically evaluate the current literature.

2.2.3. Novelty and Practical Advantages of the Framework

Our proposed Data Fidelity framework introduces a critical paradigm shift from isolated performance tracking to holistic pipeline evaluation. Its primary novelty lies in bridging the gap between generation-centric metrics (like factual consistency) and retrieval-centric metrics (like MRR) under a single, unified compositional umbrella. The practical advantages of this framework are threefold: (1) Comprehensive Vulnerability Detection: By evaluating the system end-to-end, it prevents high retrieval accuracy from masking critical failures in downstream generation or fairness. (2) Standardized Benchmarking: It provides developers with a clear, theoretically sound scaffold ($\mathcal{F}$) to objectively compare diverse architectures (RAG, fine-tuning, Hybrid) against safety thresholds. (3) Deployment Readiness: It equips enterprise practitioners with actionable criteria to guarantee that LLM deployments remain truthful, unbiased, and resilient to adversarial data poisoning in high-stakes environments.

2.3. Cross-Cutting Fidelity Concerns

Ensuring high fidelity requires addressing the intersection of security, privacy, and ethics.

2.3.1. Security: Vulnerabilities and Integrity

The operational security of LLMs directly impacts data integrity, with research identifying numerous vulnerabilities. Raj et al. [16] consider prompt injection to be the primary threat, where malicious entities bypass safety filters by manipulating inputs. Prompt injection attacks can be broadly categorized into direct injections (jailbreaking the system prompt directly via user input) and indirect injections (embedding malicious instructions within external documents). Indirect prompt injection is particularly threatening to RAG architectures: an adversary can poison the external retrieval corpus, effectively hijacking the generator's output and compromising the entire system's integrity via malicious retrieved contexts [17]. Furthermore, M'Lisha et al. [18] found that models frequently hallucinate details about zero-day vulnerabilities, highlighting the risk of relying on un-augmented models for security tasks.

2.3.2. Privacy and Confidentiality

Protecting Personally Identifiable Information (PII) is a fundamental challenge in deploying LLMs due to the vastness of datasets, which may contain such information inadvertently. Rathod et al. [19] highlight that LLMs can inadvertently leak training data verbatim. Beyond direct leakage, LLMs are vulnerable to inference attacks, which allow an attacker to infer sensitive attributes about individuals even if their data is not explicitly reproduced [19].

In RAG, privacy risks are exacerbated if the retriever accesses sensitive documents without proper access control, leading to "context leakage" where a user is presented with unauthorized data. Privacy-preserving techniques, such as Differential Privacy (DP) and Federated Learning (FL), are increasingly critical for high-fidelity fine-tuning [20]. FL allows for model training on decentralized user data without centralizing it, while DP adds statistical noise to limit the data shared during training.

2.3.3. Bias and Fairness

One of the most extensively studied issues in LLMs is their tendency to inherit and amplify societal biases from their training data [21]. This "representational fairness" is a core pillar of fidelity. Sharma et al. [22] found that conversational LLMs can create "generative echo chambers" by selectively reinforcing a user's existing beliefs and filtering out opposing views, which can increase opinion polarization over time.

In RAG, bias can stem from the retrieval algorithm itself ("ranking bias"), prioritizing popular but potentially biased sources over neutral ones ("popularity bias") or showing a preference for LLM-generated content ("source bias") [23]. Achieving true data fidelity requires mitigating these biases across both the model parameters and the retrieval corpus [24].

Summary of Conceptual Background: The fundamental limitation of static LLMs is their propensity for hallucination and knowledge decay. Augmenting these models is necessary, but introduces complex vectors for fidelity degradation. To maintain trust, augmentation architectures must be designed to actively resist data poisoning, preserve data privacy, and counteract inherent algorithmic bias.

3. Review Methodology

This study conducts a systematic literature review to map the landscape of LLM augmentation techniques. The review process follows the guidelines set out by PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) to ensure transparency and reproducibility. Rather than an inductive thematic analysis, this review adopts a deductive narrative synthesis approach. The structural taxonomy (RAG, fine-tuning,

Hybrid) and evaluation criteria (accuracy, robustness, latency, interpretability) were established *a priori* based on existing foundational frameworks [11,25] to systematically extract and answer the formulated research questions.

3.1. Search Strategy and Selection

We searched major academic databases (ACM Digital Library, IEEE Xplore, Springer Link, Science Direct) and preprint archives (arXiv). The Boolean search strings utilized across all repositories were ("Large Language Model" OR "LLM") AND ("Retrieval-Augmented Generation" OR "RAG" OR "Fine-tuning") AND ("Data Fidelity" OR "Hallucination" OR "Factuality" OR "Trustworthiness"). While the core of our comparative analysis focuses on the refined subset of articles, a broader corpus of literature was analyzed during the screening phase to establish historical context, evaluate edge-case vulnerabilities, and infer the broader trajectory of systemic augmentation trends [26–70].

3.2. Inclusion and Exclusion Criteria

Our initial search yielded 818 articles (Figure 4). These initial records were distributed across the databases as follows: IEEE Xplore (n = 210), arXiv (n = 208), ACM Digital Library (n = 145), Science Direct (n = 135), and Springer Link (n = 120). To ensure the relevance and quality of the reviewed literature, we applied the following specific inclusion and exclusion criteria:

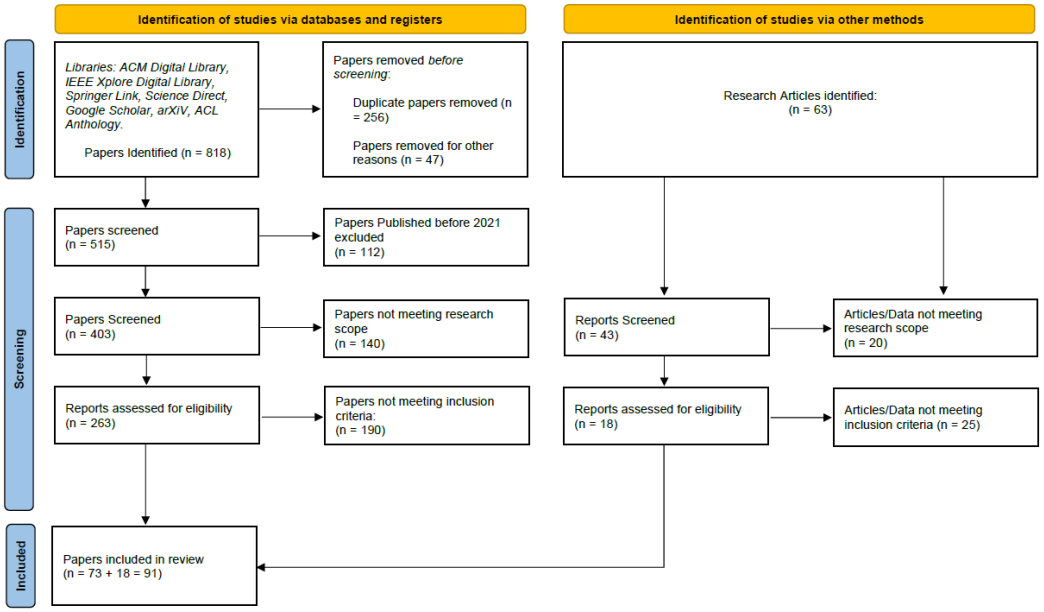


Figure 4. Paper selection and refinement process represented in a PRISMA Flow.

- **Inclusion Criteria:**
  - *Publication Date:* Articles published between 1 January 2021, and 31 December 2025, to capture the most recent advancements in LLM capabilities.
  - *Relevance:* Papers explicitly discussing "Retrieval-Augmented Generation," "Fine-Tuning," or "External Data Augmentation" in the context of Large Language Models.
  - *Focus on Fidelity:* Studies that include evaluations or discussions regarding data fidelity, operationally defined as containing empirical measurements of factual consistency, hallucination mitigation, or information loss.
  - *Source Type:* Peer-reviewed journal articles and conference proceedings (e.g., ACL, NeurIPS, ICLR) to ensure academic rigor.

- **Exclusion Criteria:**
  - *Language:* Non-English publications.
  - *Scope:* Papers focusing solely on pre-training architectures without addressing external data integration.
  - *Format:* Opinion pieces, editorials, and non-peer-reviewed preprints unless they represented heavily cited, foundational architectures (e.g., LongRAG). We explicitly acknowledge that reliance on non-peer-reviewed preprints constitutes a limitation, and findings drawn from them should be interpreted with appropriate caution regarding epistemic robustness.
  - *Redundancy:* Duplicate studies or earlier versions of papers that have been subsequently published in journals.

Screening was independently conducted by two reviewers, achieving a high inter-rater reliability (Cohen's $\kappa = 0.85$). Any disagreements were resolved through discussion and consensus. After filtering for duplicates and relevance, 36 articles were selected for in-depth analysis. These were categorized into RAG, fine-tuning, and Hybrid architectures to address the review questions (RQ-A, RQ-B, RQ-C).

4. Taxonomy of Augmentation Methodologies

The objective of this section is to categorize the primary methodologies for integrating external data into LLMs, providing a structured taxonomy (RQ-A). Rather than merely describing systems, we classify these methods by analyzing them against clear comparison criteria: accuracy (precision of facts), robustness (resistance to noise), latency (computational efficiency), and interpretability (provenance tracking).

A structural comparison between Retrieval-Augmented Generation (RAG) and Fine-Tuning reveals divergent capabilities regarding knowledge management, hallucination mitigation, and data privacy [10,11,25]. Fine-tuning deeply embeds domain-specific patterns into the parametric memory of the model, which is highly effective for adopting specific structural formats or stylistic nuances. However, it is generally less suited for maintaining dynamic knowledge, as updating facts requires resource-intensive retraining. Conversely, RAG directly queries external databases at inference time, enabling real-time knowledge updates without altering the underlying model weights. Crucially, RAG frameworks are generally less prone to hallucinations under typical retrieval settings because generated responses are grounded in explicitly retrieved evidence. While fine-tuning can reduce hallucinations on familiar domain data, models remain highly susceptible to fabricating information when presented with unfamiliar inputs. Furthermore, RAG provides granular access control for privacy and ethical compliance through database-level security, whereas fine-tuning risks embedding sensitive or copyrighted material permanently into the model parameters. Table 2 highlights these fundamental conceptual differences.

Table 2. Qualitative comparison of RAG and fine-tuning paradigms [10,11].

Feature	Retrieval-Augmented Generation (RAG)	Fine-Tuning
Knowledge Updates	Directly updates the retrieval knowledge base, ensuring information remains current without retraining. Highly suitable for dynamic data.	Requires continuous retraining to update knowledge. Not suitable for adding new knowledge or scenarios requiring rapid iteration.
Reducing Hallucinations	Generally less prone to hallucinations as each answer is grounded in verifiable retrieved evidence.	Reduces hallucinations on specific domain data but may still exhibit severe hallucinations when faced with unfamiliar input.

Table 2. *Cont.*

Feature	Retrieval-Augmented Generation (RAG)	Fine-Tuning
Ethical and Privacy	Privacy concerns arise primarily from storing and retrieving text from external databases, which can be secured via access controls.	Ethical and privacy concerns arise from permanently embedding sensitive content directly into the model's training data parameters.

We categorize methodologies based on their architectural approach to fidelity: Retrieval-Based (RAG), Parameter-Based (Fine-tuning), and Hybrid/Agentic approaches. Figure 5 illustrates the conceptual relationships and overlaps between these dominant paradigms.

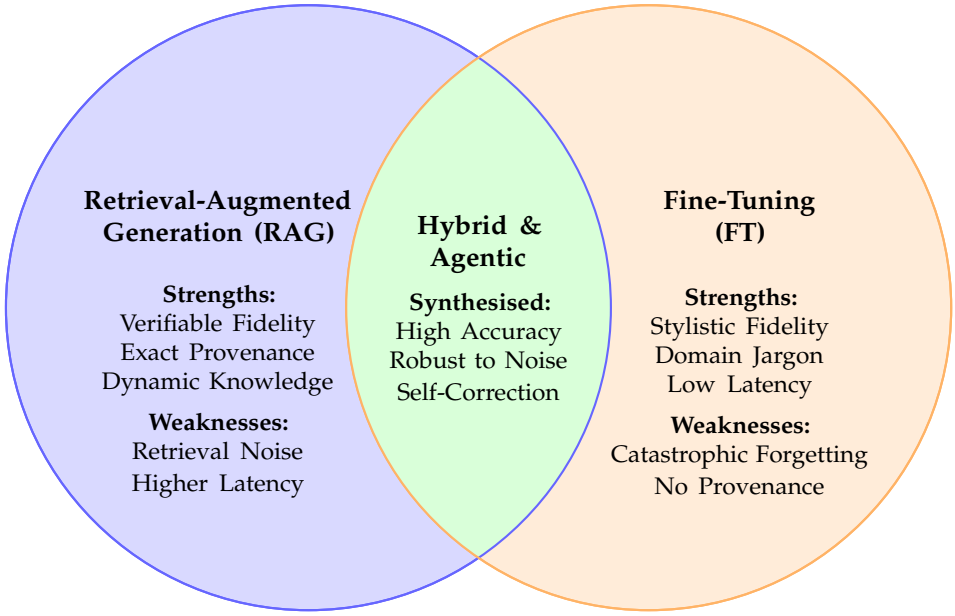


Figure 5. A conceptual Venn diagram demonstrating the overlapping relationships, shared benefits, and distinct capabilities of RAG, fine-tuning, and Hybrid augmentation approaches.

4.1. Retrieval-Augmented Generation (RAG)

RAG integrates external data by retrieving relevant document chunks to condition the LLM's generation, akin to an "open-book" exam [11]. This process allows the model to retrieve current factual data to substantiate its responses, thereby enhancing accuracy and interpretability.

4.1.1. Naive RAG

The "Naive" approach is the most basic implementation. The system performs a simple semantic search (typically cosine similarity) to retrieve top-k documents from a vector database and concatenates them with the user prompt [71]. While easy to implement and highly interpretable, it suffers from low contextual robustness. It can struggle with getting the retrieval step right, often pulling in text that seems related on the surface but is not actually helpful. Another issue is the "lost in the middle" phenomenon, where the LLM tends to ignore information buried in long contexts, leading to fidelity degradation [72]. However, practical implementations in enterprise environments demonstrate that combining RAG with automated extract-transform-load (ETL) processes is essential to maintain data currency [73].

4.1.2. Advanced RAG

To overcome the limitations of naive RAG, advanced techniques aim to improve the robustness and fidelity of the retrieved context:

- Query Transformation and Decomposition: Techniques like Query Decomposition break complex questions into sub-queries. The system finds documents for each sub-question and pieces the information together. HyDE (Hypothetical Document Embeddings) takes a different approach by asking the LLM to generate a theoretical answer to guide retrieval, ensuring the semantic search aligns with the answer space rather than just the question space [74].
- Re-Ranking: A second-stage process using cross-encoders to score documents based on true relevance. RankRAG [75] demonstrates that re-ranking significantly reduces noise. Unlike bi-encoders used in the initial search, cross-encoders process the query and document together, allowing for a deeper analysis of relevance. This directly improves the factual accuracy of the generation by filtering out irrelevant "distractors."
- Knowledge Fusion: Methods like RAG-Fusion [76] generate multiple query perspectives and fuse the results using Reciprocal Rank Fusion (RRF). This provides a richer context by combining diverse sets of retrieved documents. However, this negatively impacts the system's latency due to the multiple retrieval operations required.

4.2. Fine-Tuning Strategies

Fine-tuning modifies the model's internal weights to internalize knowledge or adapt behaviour. This involves continued training on domain-specific datasets, fundamentally altering the parametric memory of the model.

4.2.1. Knowledge Injection and Domain Adaptation

Approaches like Synthetic Knowledge Ingestion (SKI) [74] generate high-quality synthetic datasets to inject facts into the model. Strategies include fine-grained synthesis and interleaved generation to prepare data optimized for supervised fine-tuning. While highly effective for improving latency (as it removes the retrieval bottleneck), fine-tuning suffers from "catastrophic forgetting"—the degradation of previously learned knowledge—posing a severe risk to integrity fidelity and offering near-zero interpretability.

4.2.2. Parameter-Efficient Fine-Tuning (PEFT)

To address efficiency and forgetting issues, PEFT methods like LoRA (Low-Rank Adaptation) freeze pre-trained weights and inject trainable rank decomposition matrices [77]. This reduces the number of trainable parameters by up to 10,000× while matching full fine-tuning performance. RAFT (Reward Ranked Fine-tuning) [78] aligns the model to prefer factual responses by training on "winning" samples ranked by a reward model, effectively steering the model toward truthfulness fidelity without the complexity of full reinforcement learning pipelines. Table 3 outlines key literature, applications, and limitations associated with these fine-tuning and PEFT methodologies.

Table 3. Literature on fine-tuning and PEFT: Methodology and Applications.

Paper	Methodology	Application/Benchmark	Key Results	Limitations & Future Scope
LoRA [77]	Freezes pre-trained weights, injects trainable low-rank matrices.	GPT-3 175B on RoBERTa/DeBERTa tasks.	Reduces trainable parameters by 10,000×, matches FFT performance.	Rank r is a sensitive hyperparameter.
RAFT [78]	Rejection sampling FT via reward model ranking.	Generative alignment on LLaMA-7B.	More stable than PPO-based RLHF for alignment tasks.	Overhead of generating k samples during training.

Table 3. *Cont.*

Paper	Methodology	Application/Benchmark	Key Results	Limitations & Future Scope
LoFiT [79]	Fine-tunes sparse subset (3–10%) of attention heads.	TruthfulQA on Llama2-7B.	Improved TruthfulQA accuracy from 62.2% to 74.4%.	Locating optimal "locus" of heads is non-trivial.
Multi-Fidelity FT [80]	Sequential training: low-fidelity then high-fidelity data.	Noise robustness benchmarks.	Outperforms naive mixing, leverages noisy data effectively.	Risk of catastrophic forgetting of initial stage.

4.2.3. Continual Learning with New Data

Continual Learning (CL) addresses the need to update models with new external information without complete retraining. Ref. [81] demonstrates that "Replay to Remember" mechanisms, coupled with LoRA, can stabilize performance and facilitate partial recovery of domain-specific knowledge when exposed to streaming data. However, despite advancements, current CL techniques do not completely eradicate catastrophic forgetting, especially when new data streams are highly dissimilar to previously learned domains.

4.3. Hybrid, Agentic, and Structured Approaches

The field is moving towards systems that combine the strengths of both paradigms to optimize across all comparison criteria.

4.3.1. Knowledge Graph Integration

Graph-enhanced LLMs, such as GRAG [82], retrieve structured subgraphs rather than text chunks. By linearizing these graphs into text, they preserve relational information that flat text retrieval often loses, improving structural fidelity and robustness. ERNIE [1] and KnowBERT [83] integrate knowledge graph embeddings directly into the LLM architecture using techniques like TransE to fuse entity info.

4.3.2. Agentic Frameworks and Plugins

Agentic systems, such as AgentFusion [84], employ autonomous agents to handle retrieval, validation, and generation separately. Tool-using agents can query APIs or SQL databases dynamically. Frameworks like MOYA [85] employ multi-agent systems for CloudOps, where distinct agents handle security scanning and reporting. These systems offer high adaptability but introduce new security risks regarding autonomous execution and data access controls. The idea of an LLM acting as an intelligent agent is often achieved through "Reasoning and Acting" (ReAct) loops, where the model decides on its own when to use a function or call an API. Key contributions and limitations of recent literature concerning knowledge graphs, structured data, and agentic RAG frameworks are summarized in Table 4.

Table 4. Literature on Knowledge Graph, Structured Data, and Agentic RAG.

Paper	Methodology	Application/Benchmark	Key Results	Limitations & Future Scope
ERNIE [1]	Pre-training with KGs, fuses entity info via TransE.	Entity Typing, Relation Classification.	Significant improvements on knowledge-driven tasks.	Static knowledge injection requires costly retraining.
GRAG [82]	Retrieves k-hop ego-graphs, "soft pruning".	Graph QA (WebQSP).	Outperforms standard RAG by preserving topological structure.	Subgraph retrieval is computationally intensive.

Table 4. *Cont.*

Paper	Methodology	Application/Benchmark	Key Results	Limitations & Future Scope
BIORAG [6]	Domain-Specific RAG, MeSH hierarchy refinement.	Biomedical QA.	73-90% accuracy on medical datasets.	Highly domain-coupled maintenance.
AgentFusion [84]	Multi-agent collaboration for retrieval/validation.	Technical documentation.	Significantly reduces hallucinations in non-English contexts.	High complexity in agent orchestration.

Summary of Method Taxonomy: Retrieval-based methods (RAG) excel in grounding outputs and interpretability but are susceptible to retrieval noise and high latency. Parameter-based methods (Fine-tuning) deeply embed domain logic and offer excellent latency, but struggle with rapid knowledge decay, lack provenance, and suffer from catastrophic forgetting. Hybrid architectures aim to balance these trade-offs by using retrieval for factual grounding and fine-tuning for domain-specific reasoning and stylistic alignment.

5. Evaluation Metrics for Fidelity

The objective of this section is to outline how the theoretical concept of data fidelity—as defined in the Introduction—is practically and quantitatively measured across the distinct stages of the augmentation pipeline. Evaluating augmented LLMs requires assessing performance at three distinct levels: the retriever, the generator, and the end-to-end system.

5.1. Retrieval-Level Metrics

Fidelity begins with retrieval. Key metrics include Hit Rate (presence of ground truth in top-k) and Mean Reciprocal Rank (MRR). High retrieval accuracy is essential to prevent "hallucination due to omission" [86]. If the retriever fails to fetch the relevant documents, even the most advanced LLM will struggle to produce a correct response, thereby severing the foundation of data fidelity.

5.2. Generation-Level Metrics

Metrics here directly quantify the Truthfulness pillar of data fidelity by evaluating the output against the retrieved context. Factual Consistency measures whether the generated claim is supported by the source, often using automated natural language inference (NLI) models [87], serving as a direct proxy for the hallucination rate. Grounding Correctness evaluates how well the generated text cites the provided sources. Traditional NLP metrics like BLEU and ROUGE are often insufficient as they measure n-gram overlap but fail to capture semantic correctness or factual hallucinations.

5.3. End-to-End Evaluation

Downstream performance is measured via Exact Match (EM) and F1-Score. However, Salemi et al. [7] note a misalignment between retrieval scores and generation quality, suggesting that simply fetching the "right" documents is not enough. The information must also be presented in a way that is usable by the LLM, necessitating holistic human evaluation and bias/fairness indicators to capture nuances of tone, safety, and representational fairness that automated metrics miss.

5.4. Computational Efficiency of Evaluation

A critical challenge in evaluating the data fidelity of RAG frameworks is the immense computational overhead required for End-to-End (E2E) evaluation. Traditional E2E evaluation processes feed the entirety of the retrieved document list into the LLM simultaneously,

resulting in quadratic scaling of computational costs relative to input length. As detailed in Table 5, benchmarking standard RAG architectures using E2E methods on datasets like Natural Questions (NQ) and FEVER can consume between 46 GB and 75 GB of GPU memory and require over 3000 s of runtime. Advanced evaluation methodologies, such as eRAG [7], mitigate this bottleneck by assessing documents individually. This document-level approach not only achieves higher correlation with true downstream performance but also reduces GPU memory consumption by up to 50 times (requiring only 1.5 GB uniformly) and accelerates runtime by a factor of 1.2–3.2. Efficient evaluation metrics are therefore paramount for the rapid iteration and deployment of secure augmented language models.

Table 5. Computational Efficiency: End-to-End (E2E) vs. Document-Level (eRAG) Evaluation.

Dataset	Runtime E2E (s)	Runtime eRAG (s)	Memory E2E (GB)	Memory eRAG-Doc (GB)
Natural Qs (NQ)	918	351	75.0	1.5
TriviaQA	1819	686	46.2	1.5
HotpotQA	1844	712	52.4	1.5
FEVER	3395	1044	66.5	1.5
Wizard of Wiki	912	740	47.9	1.5

5.5. Evaluation Biases and Benchmark Limitations

It is crucial to acknowledge the inherent limitations and potential biases within current evaluation benchmarks. For example, annotation bias in QA datasets—where human labellers may favour specific phrasing or sentence structures—can misalign with genuine factual correctness. Furthermore, the construction of "retrieval-friendly" benchmarks can artificially inflate the perceived performance of RAG paradigms by failing to simulate the adversarial noise found in real-world data. Therefore, achieving robust data fidelity requires a holistic evaluation framework that integrates both automated NLP metrics and domain-specific human evaluation to counteract these benchmark-specific vulnerabilities.

Summary of Evaluation Metrics: Evaluating an augmented LLM requires a multi-tiered approach. Excellent retrieval metrics (MRR) do not guarantee high fidelity if generation metrics (Factual Consistency) reveal that the model ignores the retrieved context to hallucinate answers. Furthermore, shifting to document-level evaluation provides significant computational efficiency improvements over standard end-to-end approaches.

6. Comparative Analysis and Discussion

The objective of this section is to systematically synthesize the literature, moving beyond individual study summaries to derive actionable insights regarding the strengths, limitations, and optimal use cases for each augmentation paradigm. This directly addresses RQ-B.

Based on the aggregated performance metrics across various RAG and fine-tuning models, we analyze the trade-offs between these paradigms across our core comparison criteria: accuracy, robustness, latency, and interpretability.

6.1. Accuracy vs. Factual Fidelity

Accuracy stands as a fundamental metric for evaluating the overall correctness of the generated responses. As synthesized in Table 6, RAG systems achieve Verifiable Fidelity by anchoring generation to retrieved evidence, yielding high interpretability. For instance, TC-RAG achieves high accuracy (83.15% on MMCU-Medical) by leveraging specialized retrieval (Table 7). This high accuracy demonstrates that domain-specific indexing is a prerequisite for high information fidelity in specialized fields; general-purpose retrieval often fails to capture the necessary semantic nuance.

In contrast, fine-tuning achieves Stylistic Fidelity—internalizing domain jargon—but risks lower Factual Fidelity on novel facts due to hallucination on unseen data. While fine-tuned models like T5-Large can achieve high accuracy on static benchmarks (0.91 on BoolQ), fine-tuning on new facts can often increase the tendency to hallucinate on adjacent topics. This suggests a core trade-off: achieving high accuracy on a specific set of newly learned facts via fine-tuning may come at the cost of degrading the model's broader information robustness.

While fine-tuning is traditionally viewed as a mechanism for knowledge injection, empirical evidence from Gekhman et al. [88] utilizing the PaLM 2-S architecture suggests that acquiring novel factual knowledge via supervised fine-tuning actively degrades a model's pre-existing factual fidelity. When the PaLM 2-S model is fine-tuned on facts categorized as "Unknown" (information entirely absent from its pre-training), it exhibits severe overfitting. As demonstrated in Table 8, extending fine-tuning to convergence on unknown datasets ($D_{Unknown}$) results in catastrophic performance drops on the model's previously "HighlyKnown" facts, falling from 95.6% accuracy at early stopping to merely 55.8% at convergence. This indicates that attempting to force an LLM to internalize new facts linearly increases its tendency to hallucinate regarding its established parametric knowledge. Therefore, fine-tuning should primarily be utilized to expose and structure pre-existing knowledge rather than to inject entirely new factual databases. It is important to note, however, that these findings are drawn primarily from a single study utilizing the PaLM 2-S architecture; while indicative of the risks of parametric knowledge injection, further empirical validation across diverse model families (e.g., LLaMA, Mistral) is required before establishing this as a universal rule. Practically, these results caution enterprise developers against using fine-tuning as a primary method for injecting rapidly changing facts, as doing so may inadvertently corrupt the model's foundational knowledge base and increase hallucination rates on previously secure tasks.

Table 6. Comparative Synthesis of LLM Augmentation Paradigms based on Data Fidelity Criteria.

Paradigm	Accuracy & Factual Fidelity	Robustness (To Noise/Decay)	Latency & Cost	Interpretability & Provenance
Naive RAG	High for retrieved facts, prone to hallucination if retrieval fails.	Low, highly susceptible to irrelevant "distractor" documents.	Medium, depends on vector search speed.	High, directly traces output to source chunks.
Advanced RAG (e.g., RankRAG)	Very High, re-ranking filters noise before generation.	High, robust against poorly phrased queries via query expansion.	High, re-ranking and multi-queries add significant delay.	High, maintains clear provenance to filtered sources.
Fine-Tuning (PEFT)	Medium, good for stylistic alignment but struggles with novel facts.	Low to Medium, prone to catastrophic forgetting over time.	Low, no retrieval step required at inference.	Low, knowledge is parametrically embedded (black-box).
Hybrid/Agentic	Very High, combines parametric reasoning with external facts.	Very High, agents can verify and self-correct retrieved data.	Very High, multiple LLM calls and tool uses required.	Medium to High, execution traces provide some transparency.

Table 7. Methodological Evaluation of Specialized RAG Architectures.

Framework	Target Domain	Evaluation Dataset	Evaluation Metric	Reported Performance
TC-RAG [89]	Medical QA	MMCU-Medical	Exact Match (EM)	83.15%
R2AG [90]	Open-domain QA	NQ-10	Exact Match (EM)	69.30%
VUL-RAG [91]	Software Security	PairVul	Precision	61.00%

Table 8. Mean Impact of Fine-Tuning on Pre-existing Knowledge Fidelity (PaLM 2-S Architecture) [88].

Training Dataset Variant	HighlyKnown (Test)	MaybeKnown (Test)	WeaklyKnown (Test)	Unknown (Test)
Early Stopping (Optimal Generalization)				
Trained on $D_{HighlyKnown}$	98.7%	60.1%	9.0%	0.6%
Trained on $D_{Unknown}$	95.6%	52.9%	6.5%	0.6%
Convergence (Overfitting on New Knowledge)				
Trained on $D_{HighlyKnown}$	98.4%	58.8%	8.5%	0.7%
Trained on $D_{Unknown}$	55.8%	36.6%	12.2%	3.2%

In Table 9, the comparative results highlight a crucial operational insight for practitioners: neither RAG nor Fine-Tuning alone is strictly optimal for complex tasks. While fine-tuning frequently struggles to internalize novel facts, combining it with RAG in a Hybrid architecture reliably yields the highest overall accuracy. This indicates that enterprise deployments should leverage fine-tuning primarily for aligning domain logic and syntax, while relying strictly on RAG for verifiable factual grounding. Hybrid approaches often outperform individual strategies. On the Llama3 8B model, RAG alone achieves 0.813 accuracy, while fine-tuning alone achieves only 0.569. The Hybrid approach improves this to 0.833, highlighting that RAG is the primary driver of factual accuracy for novel data, whereas fine-tuning alone often fails to effectively internalize new facts.

Table 9. Trade-offs Between RAG, Fine-Tuning, and Hybrid Paradigms (POPQA Dataset Mean Accuracy).

Model Architecture	RAG Only (Ideal)	Fine-Tuning (PEFT) Only	Hybrid Approach
StableLM2 (1.6B)	0.761	0.217	0.821
Llama3 (8B)	0.813	0.569	0.833

6.2. Contextual Robustness and Recall

RAG offers superior recall for open-domain queries but faces challenges with fragmentation (context loss). Fine-tuning generally has lower recall for new facts but higher coherence. Table 10 illustrates that methods like RankRAG significantly boost recall by refining the retrieval stage, increasing robustness against noisy raw data.

Traditional RAG systems often require retrieving hundreds of short passages to achieve high recall (e.g., over 90%), which exacerbates context fragmentation. By contrast, LONGRAG demonstrates remarkable efficiency by processing Wikipedia into much larger 4K-token units, achieving a strong Answer Recall of 71.7% on NQ by retrieving just a single document (Recall@1). By retrieving fewer but longer, contextually complete documents, it effectively mitigates the "lost in the middle" problem while reducing the burden on the retriever. As evidenced in Table 10, transitioning from traditional short-passage retrieval to extended-context approaches dramatically improves retrieval completeness at low depths (Recall@1). For developers, this signifies that maximizing the context window of modern LLMs reduces the search burden on the retriever, thereby minimizing the risk of missing critical information ("hallucination by omission") during the generative phase. In contrast, fine-tuning typically struggles to inject new knowledge and consistently underperforms RAG in recall-heavy tasks, making it a risky strategy for applications in rapidly evolving domains.

Table 10. Retrieval Completeness (Recall@*k*).

Model/Method	Dataset	Metric	Score (%)
LONGRAG [72]	Natural Questions (NQ)	Answer Recall@1	71.7
LONGRAG [72]	HotpotQA	Answer Recall@2	72.5
RankRAG (8B) [75]	TriviaQA	Recall@5	93.2
RankRAG (8B) [75]	TriviaQA	Recall@10	95.4

6.3. Interpretability and Safety (Exact Match)

To systematically assess information fidelity, it is necessary to benchmark the baseline hallucination propensity of foundation models across diverse generative tasks. Table 11 presents hallucination evaluation results utilizing the TrustLLM framework [15] across multiple-choice (MC), open-ended question-answering (QA), knowledge-grounded dialogue (KGD), and text summarization (SUM) tasks. The data reveals that state-of-the-art models exhibit highly task-dependent fidelity. For instance, GPT-4 achieves superior performance in objective tasks, scoring 0.835 in MC and 0.760 in summarization, yet it struggles significantly with knowledge-grounded dialogue, dropping to an accuracy of 0.150. Conversely, models optimized via advanced alignment training, such as ChatGLM2, demonstrate robust performance in dialogue (0.500) and QA (0.600). The widespread failure of baseline LLMs to consistently surpass 0.600 accuracy across these generation tasks underscores the absolute necessity of external augmentation techniques like RAG to ensure high-fidelity outputs.The hallucination metrics presented in Table 11 reveals a stark vulnerability in unaugmented foundation models. Even state-of-the-art models exhibit significant performance drops in dynamic, knowledge-grounded dialogue and open QA, frequently falling well below acceptable safety thresholds. This empirical data practically reinforces the absolute necessity of external data augmentation to achieve the truthfulness required for production-level fidelity.

For precision-critical tasks, RAG frameworks demonstrate superior Interpretability, allowing for atomic verification of claims against source documents. Table 12 shows that advanced RAG methods like RankRAG achieve high Exact Match scores (up to 0.829 on TriviaQA), significantly outperforming standard baselines in precision-critical scenarios. By incorporating re-ranking mechanisms and semantic gap bridging (e.g., R2AG), these systems actively filter noise and enforce strict adherence to the retrieved context. This is highly advantageous for precision-critical domains, such as legal or medical QA, where tracing provenance and maintaining factual integrity are non-negotiable.

While RAG systems allow for verification, fine-tuning lacks this traceability. Once knowledge is encoded, the link to the source is severed. However, RAG introduces Input Fidelity Risks, where the system is susceptible to retrieving "poisoned" or irrelevant contexts that immediately degrade the output, a vulnerability less prevalent in rigorously curated fine-tuning datasets.

Table 11. Comprehensive Hallucination Evaluation Across Generation Tasks (Accuracy) [15].

Model Architecture	Multi-Choice (MC)	Open QA	Dialogue (KGD)	Summarisation
GPT-4	0.835	0.320	0.150	0.760
ChatGPT (GPT-3.5)	0.557	0.500	0.430	0.630
ChatGLM2	0.557	0.600	0.500	0.510
Llama2-70B	0.256	0.370	0.440	0.540
Mistral-7B	0.412	0.480	0.450	0.490
Vicuna-33B	0.412	0.410	0.420	0.450

Table 12. Performance on Precision-Critical QA.

Model/Method	Base LLM	Dataset	Metric	Score
R2AG [90]	LLaMA-2-7B	HotpotQA	Accuracy	0.667
R2AG (w/RAFT) [90]	LLaMA-2-7B	HotpotQA	Accuracy	0.735
RankRAG [75]	LLaMA3-8B	TriviaQA	Exact Match (EM)	0.829

6.4. Decision Framework: When to Use What?

Synthesizing the comparative performance across these criteria reveals distinct operational domains for practitioner constraints. Researchers should use the following decision framework to select the appropriate architecture based on their fidelity requirements and computational constraints:

- Use RAG when: Data changes frequently (News, Stock prices), provenance, interpretability, and citation are required, the knowledge base is vast (Web-scale). RAG is essential for dynamic, "knowledge-intensive" tasks where output must be strictly verified.
- Use Fine-Tuning when: The domain is static (Medical terminology, Legal syntax), low latency is critical (eliminating retrieval steps), the goal is stylistic adaptation or instruction following. Fine-tuning offers superior stylistic robustness and latency, but poor interpretability.
- Use Hybrid when: High precision is required on complex, domain-specific tasks that demand both deep structural understanding (FT) and up-to-date facts (RAG).

Summary of Comparative Analysis: The literature demonstrates that RAG is the primary driver of factual accuracy and interpretability for novel data, while fine-tuning alone fails to effectively internalize new facts without suffering from hallucination. Hybrid approaches reliably yield the highest performance by offsetting the limitations of each individual paradigm.

7. Future Research Directions

This section addresses RQ-C: What open challenges remain?

We identify critical gaps that serve as a roadmap for future research into high-fidelity augmentation.

7.1. Quantifying Data Fidelity

There is a lack of standardized metrics for "fidelity." Future research should develop unified frameworks that quantify the three pillars (truthfulness, integrity, fairness) simultaneously, rather than treating them as separate optimization targets. Metrics like FactScore and TruLens are a step in the right direction, but broader adoption is needed.

7.2. Continual Learning Without Forgetting

While PEFT mitigates catastrophic forgetting, it does not eliminate it. Developing robust Continual Learning (CL) algorithms that can integrate streaming data updates without degrading existing knowledge remains a "holy grail" for parametric augmentation. Future work should explore adaptive "switch" mechanisms that dynamically determine whether to retrieve external context or rely on parametric knowledge.

7.3. Privacy-Preserving Augmentation

As RAG systems are deployed in enterprise environments, "Private RAG" becomes critical. Future work must explore homomorphic encryption and secure enclaves to allow

retrieval of sensitive data without exposing the raw content to the LLM provider or the retrieval index. This is essential to ensure that augmenting models with sensitive external data does not lead to inadvertent data leakage.

7.4. Agentic Reliability and Safety

As systems move towards autonomous agents (Agentic RAG), ensuring the reliability of tool use is paramount. Research is needed into "guardrails" that prevent agents from executing harmful API calls or accessing restricted data during the augmentation loop.

7.5. Multimodal Augmentation Fidelity

Current augmentation evaluations predominantly focus on text-based enterprise QA. The expansion of RAG and fine-tuning paradigms into multimodal scenarios (e.g., vision-language models or complex tabular parsing) presents a nascent but crucial research frontier. Future studies must develop fidelity benchmarks specific to non-textual knowledge integration to ensure visual or structured data is not hallucinated during the generation process.

Furthermore, we explicitly note that several future directions identified in this section, particularly concerning agentic RAG and continual learning, are currently supported primarily by preprint literature. These areas require rigorous peer-reviewed validation before definitive conclusions can be drawn regarding their long-term fidelity.

8. Conclusions

This systematic review mapped the landscape of methodologies for augmenting large language models with external data, specifically contrasting Retrieval-Augmented Generation (RAG) and fine-tuning through the lens of data fidelity. Our analysis identified a distinct dichotomy in how these paradigms handle information integrity. RAG consistently excels in maintaining verifiable, source-grounded factual fidelity, making it indispensable for dynamic, knowledge-intensive tasks. Conversely, fine-tuning is highly effective for stylistic adaptation and deeply embedding static domain nuances, but inherently risks catastrophic forgetting and increased hallucinations when exposed to novel facts. A key finding is that optimal data fidelity is rarely achieved through a single method. Rather, state-of-the-art performance increasingly relies on hybrid architectures and agentic frameworks that synergize retrieval precision with parametric understanding.

To advance the development of trustworthy AI, future research must shift from isolated performance metrics to comprehensive evaluations of truthfulness, integrity, and representational fairness. We recommend prioritizing the standardization of data fidelity benchmarks and the development of robust, privacy-preserving augmentation techniques.

Author Contributions: Conceptualization, S.M.; data curation, S.M.; formal analysis, S.M., J.L. and C.N.; investigation, S.M., J.L. and C.N.; project administration, S.M., J.L. and C.N.; resources, S.M., J.L. and C.N.; supervision, J.L. and C.N.; validation, S.M., J.L. and C.N.; visualization, S.M., J.L. and C.N.; writing—original draft, S.M., J.L. and C.N.; writing—review and editing, S.M., J.L. and C.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
CL	Continual Learning
DP	Differential Privacy
EM	Exact Match
KG	Knowledge Graph
LLM	Large Language Model
LoRA	Low-Rank Adaptation
PEFT	Parameter-Efficient Fine-tuning
RAG	Retrieval-Augmented Generation
ReAct	Reasoning and Acting

References

1. Zhang, Z.; Han, X.; Liu, Z.; Jiang, X.; Sun, M.; Liu, Q. ERNIE: Enhanced Language Representation with Informative Entities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*; Korhonen, A., Traum, D., Màrquez, L., Eds.; Association for Computational Linguistics: Florence, Italy, 2019; pp. 1441–1451. [CrossRef]
2. Sun, Y.; Wang, S.; Feng, S.; Ding, S.; Pang, C.; Shang, J.; Liu, J.; Chen, X.; Zhao, Y.; Lu, Y.; et al. ERNIE 3.0: Large-scale Knowledge Enhanced Pre-training for Language Understanding and Generation. *arXiv* **2021**, arXiv:2107.02137.
3. Xue, J.; Zheng, M.; Hu, Y.; Liu, F.; Chen, X.; Lou, Q. BadRAG: Identifying Vulnerabilities in Retrieval Augmented Generation of Large Language Models. *arXiv* **2024**, arXiv:2406.00083. [CrossRef]
4. Weihang, S.; Tang, Y.; Ai, Q.; Wu, Z.; Liu, Y. DRAGIN: Dynamic Retrieval Augmented Generation based on the Real-time Information Needs of Large Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; Association for Computational Linguistics: Bangkok, Thailand, 2024; pp. 12991–13013. [CrossRef]
5. Chauhan, P.; Sahani, R.K.; Datta, S.; Qadir, A.; Raj, M.; Ali, M.M. Evaluating Top-k RAG-based approach for Game Review Generation. In Proceedings of the 2024 IEEE International Conference on Computing, Power and Communication Technologies (IC2PCT), Greater Noida, India, 9–10 February 2024; Volume 5, pp. 258–263. [CrossRef]
6. Wang, C.; Long, Q.; Xiao, M.; Cai, X.; Wu, C.; Meng, Z.; Wang, X.; Zhou, Y. BioRAG: A RAG-LLM Framework for Biological Question Reasoning. *arXiv* **2024**, arXiv:2408.01107. [CrossRef]
7. Salemi, A.; Zamani, H. Evaluating Retrieval Quality in Retrieval-Augmented Generation. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, Washington, DC, USA, 14–18 July 2024; SIGIR '24, pp. 2395–2400. [CrossRef]
8. Spennemann, D.H.R. Generative Artificial Intelligence and the Future of Public Knowledge. *Knowledge* **2025**, *5*, 20. [CrossRef]
9. Koo, M. ChatGPT Research: A Bibliometric Analysis Based on the Web of Science from 2023 to June 2024. *Knowledge* **2025**, *5*, 4. [CrossRef]
10. Ovadia, O.; Brief, M.; Mishaeli, M.; Elisha, O. Fine-Tuning or Retrieval? Comparing Knowledge Injection in LLMs. In Proceedings of the Conference on Empirical Methods in Natural Language Processing, Singapore, 6–10 December 2023.
11. Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; Guo, Q.; Wang, M.; et al. Retrieval-Augmented Generation for Large Language Models: A Survey. *arXiv* **2023**, arXiv:2312.10997. [CrossRef]
12. Azaria, A.; Mitchell, T. The Internal State of an LLM Knows When It's Lying. *arXiv* **2023**, arXiv:2304.13734.
13. Madabushi, H.T. FS-RAG: A Frame Semantics Based Approach for Improved Factual Accuracy in Large Language Models. *arXiv* **2024**, arXiv:2406.16167.
14. Wei, J.; Yang, C.; Song, X.; Lu, Y.; Hu, N.Z.; Huang, J.; Tran, D.; Peng, D.; Liu, R.; Huang, D.; et al. Long-form factuality in large language models. In Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024.
15. Huang, Y.; Sun, L.; Wang, H.; Wu, S.; Zhang, Q.; Li, Y.; Gao, C.; Huang, Y.; Lyu, W.; Zhang, Y.; et al. Position: TRUSTLLM: Trustworthiness in large language models. In Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria, 21–27 July 2024; ICML'24.
16. Raj, G.; Hamzah; Raj, N.; Ranjan, N. Hacking LLMs: A Technical Analysis of Security Vulnerabilities and Defense Mechanisms. In Proceedings of the 2025 2nd International Conference on Computational Intelligence, Communication Technology and Networking (CICTN), Ghaziabad, India, 6–7 February 2025; pp. 555–560. [CrossRef]

17. Gummadi, V.; Udayaraju, P.; Sarabu, V.R.; Ravulu, C.; Seelam, D.R.; Venkataramana, S. Enhancing Communication and Data Transmission Security in RAG Using Large Language Models. In Proceedings of the 2024 4th International Conference on Sustainable Expert Systems (ICSES), Pokhara, Nepal, 15–17 October 2024; pp. 612–617. [CrossRef]
18. Lisha, M.; Agarwal, V.; Kamthania, S.; Vutkur, P.; Madhusoodhana Chari, S. Benchmarking LLM for Zero-day Vulnerabilities. In Proceedings of the 2024 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT), Bangalore, India, 12–14 July 2024; pp. 1–6. [CrossRef]
19. Rathod, V.; Nabavirazavi, S.; Zad, S.; Iyengar, S.S. Privacy and Security Challenges in Large Language Models. In Proceedings of the 2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC), Las Vegas, NV, USA, 6–8 January 2025; pp. 00746–00752. [CrossRef]
20. Pan, Q.; Wu, J. Selective Privacy-Preserving Federated Learning for Large Language Model Fine-Tuning. In Proceedings of the 2025 International Wireless Communications and Mobile Computing (IWCMC), Ayia Napa, Cyprus, 16–20 June 2025; pp. 1626–1631. [CrossRef]
21. Gallegos, I.O.; Rossi, R.A.; Barrow, J.; Tanjim, M.M.; Kim, S.; Dernoncourt, F.; Yu, T.; Zhang, R.; Ahmed, N.K. Bias and Fairness in Large Language Models: A Survey. *Comput. Linguist.* **2024**, *50*, 1097–1179. [CrossRef]
22. Sharma, N.; Liao, Q.V.; Xiao, Z. Generative Echo Chamber? Effect of LLM-Powered Search Systems on Diverse Information Seeking. In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, 11–16 May 2024; CHI '24. [CrossRef]
23. Sakurai, T.; Shiramatsu, S.; Kinoshita, R. LLM-based Agent for Recommending Information Related to Web Discussions at Appropriate Timing. In Proceedings of the 2024 IEEE International Conference on Agents (ICA), Kyoto, Japan, 25–27 September 2024; pp. 120–123. [CrossRef]
24. Zhang, W.; Liu, H.; Dong, Z.; Du, Y.; Zhu, C.; Song, Y.; Zhu, H.; Wu, Z. Bridging the Information Gap Between Domain-Specific Model and General LLM for Personalized Recommendation. In *Proceedings of the Web and Big Data*; Zhang, W., Tung, A., Zheng, Z., Yang, Z., Wang, X., Guo, H., Eds.; Springer: Singapore, 2024; pp. 280–294.
25. Mialon, G.; Dessì, R.; Lomeli, M.; Nalmpantis, C.; Pasunuru, R.; Raileanu, R.; Rozière, B.; Schick, T.; Dwivedi-Yu, J.; Celikyilmaz, A.; et al. Augmented Language Models: A Survey. *arXiv* **2023**, arXiv:2302.07842.
26. Hagen, M.; Völske, M.; Göring, S.; Stein, B. Axiomatic Result Re-Ranking. In Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, Indianapolis, IN, USA, 24–28 October 2016; CIKM '16, pp. 721–730. [CrossRef]
27. Dai, S.; Xu, C.; Xu, S.; Pang, L.; Dong, Z.; Xu, J. Bias and Unfairness in Information Retrieval Systems: New Challenges in the LLM Era. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Barcelona, Spain, 25–29 August 2024; KDD '24, pp. 6437–6447. [CrossRef]
28. Razuvayevskaya, O.; Wu, B.; Leite, J.A.; Heppell, F.; Srba, I.; Scarton, C.; Bontcheva, K.; Song, X. Comparison between parameter-efficient techniques and full fine-tuning: A case study on multilingual news article classification. *PLoS ONE* **2024**, *19*, e0301738. [CrossRef] [PubMed]
29. Huang, C.W.; Chen, Y.N. FactAlign: Long-form Factuality Alignment of Large Language Models. *arXiv* **2024**, arXiv:2410.01691.
30. Wang, Y.; Wang, M.; Manzoor, M.A.; Liu, F.; Georgiev, G.N.; Das, R.J.; Nakov, P. Factuality of Large Language Models: A Survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*; Al-Onaizan, Y., Bansal, M., Chen, Y.N., Eds.; Association for Computational Linguistics: Miami, FL, USA, 2024; pp. 19519–19529. [CrossRef]
31. He, X.; Tian, Y.; Sun, Y.; Chawla, N.V.; Laurent, T.; LeCun, Y.; Bresson, X.; Hooi, B. G-retriever: Retrieval-augmented generation for textual graph understanding and question answering. In Proceedings of the 38th International Conference on Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; NIPS '24.
32. Han, H.; Wang, Y.; Shomer, H.; Guo, K.; Ding, J.; Lei, Y.; Halappanavar, M.; Rossi, R.A.; Mukherjee, S.; Tang, X.; et al. Retrieval-Augmented Generation with Graphs (GraphRAG). *arXiv* **2025**, arXiv:2501.00309.
33. Gao, L.; Ma, X.; Lin, J.; Callan, J. Precise Zero-Shot Dense Retrieval without Relevance Labels. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; Rogers, A., Boyd-Graber, J., Okazaki, N., Eds.; Association for Computational Linguistics: Toronto, ON, Canada, 2023; pp. 1762–1777. [CrossRef]
34. Tian, S.; Luo, Y.; Xu, T.; Yuan, C.; Jiang, H.; Wei, C.; Wang, X. KG-Adapter: Enabling Knowledge Graph Integration in Large Language Models through Parameter-Efficient Fine-Tuning. In *Proceedings of the Findings of the Association for Computational Linguistics: ACL 2024*; Ku, L.W., Martins, A., Srikumar, V., Eds.; Association for Computational Linguistics: Bangkok, Thailand, 2024; pp. 3813–3828. [CrossRef]
35. Raza, S.; Raval, A.; Chatrath, V. MBIAS: Mitigating Bias in Large Language Models While Retaining Context. In *Proceedings of the 14th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*; De Clercq, O., Barriere, V., Barnes, J., Klinger, R., Sedoc, J., Tafreshi, S., Eds.; Association for Computational Linguistics: Bangkok, Thailand, 2024; pp. 97–111. [CrossRef]

36. Wang, Y.; Shi, X.; Zhao, X. MLLM4Rec: Multimodal information enhancing LLM for sequential recommendation. *J. Intell. Inf. Syst.* **2025**, *63*, 745–761. [CrossRef]
37. Lee, C.; Roy, R.; Xu, M.; Raiman, J.; Shoeybi, M.; Catanzaro, B.; Ping, W. NV-Embed: Improved Techniques for Training LLMs as Generalist Embedding Models. *arXiv* **2025**, arXiv:2405.17428.
38. Bin Islam, S.; Rahman, M.; Hossain, K.; Hoque, E.; Joty, S.; Parvez, M.R. Open-RAG: Enhanced Retrieval Augmented Reasoning with Open-Source Large Language Models. In *Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2024*; Association for Computational Linguistics: Miami, FL, USA, 2024; pp. 14231–14244. [CrossRef]
39. Jiang, W.; Zhang, S.; Han, B.; Wang, J.; Wang, B.; Kraska, T. PipeRAG: Fast Retrieval-Augmented Generation via Adaptive Pipeline Parallelism. In Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1, Toronto, ON, Canada, 10–14 August 2025; KDD '25, pp. 589–600. [CrossRef]
40. Wang, L.; Chen, S.; Jiang, L.; Pan, S.; Cai, R.; Yang, S.; Yang, F. Parameter-efficient fine-tuning in large language models: A survey of methodologies. *Artif. Intell. Rev.* **2025**, *58*, 227. [CrossRef]
41. Seo, S.; Noh, S.; Lee, J.; Lim, S.; Lee, W.H.; Kang, H. REVECA: Adaptive planning and trajectory-based validation in cooperative language agents using information relevance and relative proximity. *Proc. AAAI Conf. Artif. Intell.* **2025**, *39*, 23295–23303. [CrossRef]
42. Chan, C.M.; Xu, C.; Yuan, R.; Luo, H.; Xue, W.; Guo, Y.; Fu, J. RQ-RAG: Learning to Refine Queries for Retrieval Augmented Generation. In Proceedings of the First Conference on Language Modeling, Philadelphia, PA, USA, 7–9 October 2024.
43. Wei Jie, Y.; Ferdinan, T.; Kazienko, P.; Satapathy, R.; Cambria, E. Self-training Large Language Models through Knowledge Detection. In *Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2024*; Al-Onaizan, Y., Bansal, M., Chen, Y.N., Eds.; Association for Computational Linguistics: Miami, FL, USA, 2024; pp. 15033–15045. [CrossRef]
44. Cosme, D.; Galvao, A.; e Abreu, F.B. A Systematic Literature Review on LLM-Based Information Retrieval: The Issue of Contents Classification. In Proceedings of the International Conference on Knowledge Discovery and Information Retrieval, Porto, Portugal, 17–19 November 2024.
45. Choi, Y.; Asif, M.A.; Han, Z.; Willes, J.; Krishnan, R.G. Teaching LLMs How to Learn with Contextual Fine-Tuning. *arXiv* **2025**, arXiv:2503.09032.
46. Wu, Z.; Hao, Y.; Mou, L. ULPT: Prompt Tuning with Ultra-Low-Dimensional Optimization. *arXiv* **2025**, arXiv:2502.04501.
47. Ning, L.; Liu, L.; Wu, J.; Wu, N.; Berlowitz, D.; Prakash, S.; Green, B.; O'Banion, S.; Xie, J. User-LLM: Efficient LLM Contextualization with User Embeddings. In Proceedings of the Companion Proceedings of the ACM on Web Conference 2025, Sydney, Australia, 28 April–2 May 2025; WWW '25, pp. 1219–1223. [CrossRef]
48. Zhang, Y.; Wu, Y.; Hua, W.; Lu, X.; Hu, X. Understanding Dynamic Diffusion Process of LLM-based Agents under Information Asymmetry. *arXiv* **2025**, arXiv:2502.13160.
49. Li, Y.; Tan, Z.; Xiao, W. LLM for Uniform Information Extraction Using Multi-task Learning Optimization. In *Asia-Pacific Web (APWeb) and Web-Age Information Management (WAIM) Joint International Conference on Web and Big Data*; Zhang, W., Tung, A., Zheng, Z., Yang, Z., Wang, X., Guo, H., Eds.; Springer: Singapore, 2025; pp. 17–29.
50. Cambon, A.; Hecht, B.; Edelman, B.; Ngwe, D.; Jaffe, S.; Heger, A.; Vorvoreanu, M.; Peng, S.; Hofman, J.; Farach, A.; et al. Early LLM-based Tools for Enterprise Information Workers Likely Provide Meaningful Boosts to Productivity. Technical Report MSR-TR-2023-43, Microsoft. 2023. Available online: https://researchdmr.com/files/AI-and-Productivity-Report-First-Edition2023.pdf (accessed on 1 April 2026).
51. Purohit, S.; Chin, G.; Mackey, P.S.; Cottam, J.A. GraphAide: Advanced Graph-Assisted Query and Reasoning System. In Proceedings of the 2024 IEEE International Conference on Big Data (BigData), Washington, DC, USA, 15–18 December 2024; pp. 3485–3493. [CrossRef]
52. Patel, V.; Tejani, P.; Parekh, J.; Huang, K.; Tan, X. Developing A Chatbot: A Hybrid Approach Using Deep Learning and RAG. In Proceedings of the 2024 IEEE/WIC International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT), Bangkok, Thailand, 10–13 December 2024; pp. 273–280. [CrossRef]
53. Renney, H.; Nethercott, M.; Williams, O.; Evetts, J.; Lang, J. Reimagining the Data Landscape: A Multi-Agent Paradigm for Data Interfacing. In Proceedings of the 2025 8th International Conference on Data Science and Machine Learning Applications (CDMA), Riyadh, Saudi Arabia, 11–13 February 2025; pp. 114–119. [CrossRef]
54. Xu, J.; Wang, J.; Leung, J.; Gu, J. GRASP: Municipal Budget AI Chatbots for Enhancing Civic Engagement. In Proceedings of the 2024 IEEE International Conference on Big Data (BigData), Washington, DC, USA, 15–18 December 2024; pp. 7438–7442. [CrossRef]
55. Xu, X.; Zhang, D.; Liu, Q.; Lu, Q.; Zhu, L. Agentic RAG with Human-in-the-Retrieval. In Proceedings of the 2025 IEEE 22nd International Conference on Software Architecture Companion (ICSA-C), Odense, Denmark, 31 March–4 April 2025; pp. 498–502. [CrossRef]

56. Honnalli, R.; Farooq, J. LLM-Powered Agentic AI Approach to Securing EV Charging Systems Against Cyber Threats. In Proceedings of the 2025 IEEE 26th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), Costa Navarino, Greece, 9–12 June 2025; pp. 266–274. [CrossRef]
57. Afsharmazayejani, R.; Shahmiri, M.M.; Link, P.; Pearce, H.; Tan, B. Toward Hardware Security Benchmarking of LLMs. In Proceedings of the 2024 IEEE LLM Aided Design Workshop (LAD), San Jose, CA, USA, 23–24 May 2024; pp. 1–7. [CrossRef]
58. Fasha, M.; Rub, F.A.; Matar, N.; Sowan, B.; Al Khaldy, M.; Barham, H. Mitigating the OWASP Top 10 For Large Language Models Applications using Intelligent Agents. In Proceedings of the 2024 2nd International Conference on Cyber Resilience (ICCR), Dubai, UAE, 26–28 February 2024; pp. 1–9. [CrossRef]
59. Allan, K.; Azcona, J.; Sripada, S.; Leontidis, G.; Sutherland, C.A.M.; Phillips, L.H.; Martin, D. Stereotypical bias amplification and reversal in an experimental model of human interaction with generative artificial intelligence. *R. Soc. Open Sci.* **2025**, *12*, 241472. [CrossRef] [PubMed]
60. Mondal, D.; Lipizzi, C. Mitigating Large Language Model Bias: Automated Dataset Augmentation and Prejudice Quantification. *Computers* **2024**, *13*, 141. [CrossRef]
61. Su, C.; Wen, J.; Kang, J.; Wang, Y.; Su, Y.; Pan, H.; Zhong, Z.; Shamim Hossain, M. Hybrid RAG-Empowered Multimodal LLM for Secure Data Management in Internet of Medical Things: A Diffusion-Based Contract Approach. *IEEE Internet Things J.* **2025**, *12*, 13428–13440. [CrossRef]
62. Lei, Y.; Ding, L.; Cao, Y.; Zan, C.; Yates, A.; Tao, D. Unsupervised Dense Retrieval with Relevance-Aware Contrastive Pre-Training. In *Proceedings of the Findings of the Association for Computational Linguistics: ACL 2023*; Rogers, A., Boyd-Graber, J., Okazaki, N., Eds.; Association for Computational Linguistics: Toronto, ON, Canada, 2023; pp. 10932–10940. [CrossRef]
63. Shi, Y.; Zi, X.; Shi, Z.; Zhang, H.; Wu, Q.; Xu, M. Enhancing Retrieval and Managing Retrieval: A Four-Module Synergy for Improved Quality and Efficiency in RAG Systems. *arXiv* **2024**, arXiv:2407.10670.
64. Yuan, X.J.; Guo, Q.; Dadson, Y.A.; Goodarzi, M.; Jung, J.; Dong, Y.; Albert, N.; Bennett Gayle, D.; Sharma, P.; Ogunbayo, O.T.; et al. A Review of Ethical Challenges in AI for Emergency Management. *Knowledge* **2025**, *5*, 21. [CrossRef]
65. Tsoulos, I.G.; Charilogis, V. Gen2Gen: Efficiently Training Artificial Neural Networks Using a Series of Genetic Algorithms. *Knowledge* **2025**, *5*, 17. [CrossRef]
66. Soudani, H.; Kanoulas, E.; Hasibi, F. Fine Tuning vs. Retrieval Augmented Generation for Less Popular Knowledge. In Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, Tokyo, Japan, 19–21 November 2024; SIGIR-AP 2024, pp. 12–22. [CrossRef]
67. Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.N.; Truitt, S.; Larson, J. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. *arXiv* **2024**, arXiv:2404.16130.
68. Wang, S.; Xiang, Y. Research on Data Augmentation Techniques for Text Classification Based on Antonym Replacement and Random Swapping. In Proceedings of the International Conference on Modeling, Natural Language Processing and Machine Learning, New York, NY, USA, 17–19 May 2024; CMNM '24, pp. 103–108. [CrossRef]
69. Liu, Y.; Cao, J.; Liu, C.; Ding, K.; Jin, L. Datasets for Large Language Models: A Comprehensive Survey. *arXiv* **2024**, arXiv:2402.18041.
70. Nia, N.; Amiri, A.; Luo, Y.; Kline, A. Ethical perspectives on deployment of large language model agents in biomedicine: A survey. *AI Ethics* **2025**, *6*, 32. [CrossRef]
71. Jin, B.; Yoon, J.; Han, J.; Arik, S.O. Long-Context LLMs Meet RAG: Overcoming Challenges for Long Inputs in RAG. *arXiv* **2024**, arXiv:2410.05983.
72. Jiang, Z.; Ma, X.; Chen, W. LongRAG: Enhancing Retrieval-Augmented Generation with Long-context LLMs. *arXiv* **2024**, arXiv:2406.15319.
73. Rivera, J.; Zapata, S.; Pizarro, R.; Keith, B. Enhancing Chatbot Performance in a SaaS Platform Through Retrieval-Augmented Generation and Prompt Engineering: A Case Study in Behavioral Safety Analysis. *Knowledge* **2025**, *5*, 25. [CrossRef]
74. Zhang, J.; Cui, W.; Huang, Y.; Das, K.; Kumar, S. Synthetic Knowledge Ingestion: Towards Knowledge Refinement and Injection for Enhancing Large Language Models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*; Al-Onaizan, Y., Bansal, M., Chen, Y.N., Eds.; Association for Computational Linguistics: Miami, FL, USA, 12–16 November 2024; pp. 21456–21473. [CrossRef]
75. Yu, Y.; Ping, W.; Liu, Z.; Wang, B.; You, J.; Zhang, C.; Shoeybi, M.; Catanzaro, B. RankRAG: Unifying context ranking with retrieval-augmented generation in LLMs. In Proceedings of the 38th International Conference on Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; NIPS '24.
76. Rackauckas, Z. Rag-Fusion: A New Take on Retrieval Augmented Generation. *Int. J. Nat. Lang. Comput.* **2024**, *13*, 37–47. [CrossRef]
77. Hu, E.J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W. LoRA: Low-Rank Adaptation of Large Language Models. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.

78. Dong, H.; Xiong, W.; Goyal, D.; Pan, R.; Diao, S.; Zhang, J.; Shum, K.; Zhang, T. RAFT: Reward rAnked FineTuning for Generative Foundation Model Alignment. *arXiv* **2023**, arXiv:2304.06767.
79. Yin, F.; Ye, X.; Durrett, G. LOFIT: Localized fine-tuning on LLM representations. In Proceedings of the 38th International Conference on Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; NIPS '24.
80. Tonolini, F.; Massiah, J.; Aletras, N.; Kazai, G. Multi-Fidelity Fine-Tuning of Pre-Trained Language Models. 2024. Available online: https://openreview.net/forum?id=KGhkIySg3h (accessed on 1 April 2026).
81. Pillai, S. Replay to Remember: Retaining Domain Knowledge in Streaming Language Models. *arXiv* **2025**, arXiv:2504.17780.
82. Hu, Y.; Lei, Z.; Zhang, Z.; Pan, B.; Ling, C.; Zhao, L. GRAG: Graph Retrieval-Augmented Generation. In *Proceedings of the Findings of the Association for Computational Linguistics: NAACL 2025*; Chiruzzo, L., Ritter, A., Wang, L., Eds.; Association for Computational Linguistics: Albuquerque, NM, USA, 2025; pp. 4145–4157. [CrossRef]
83. Peters, M.E.; Neumann, M.; Logan, R.; Schwartz, R.; Joshi, V.; Singh, S.; Smith, N.A. Knowledge Enhanced Contextual Word Representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*; Inui, K., Jiang, J., Ng, V., Wan, X., Eds.; Association for Computational Linguistics: Hong Kong, China, 2019; pp. 43–54. [CrossRef]
84. Saeid, Y.; Kopinski, T. AgentFusion: A Multi-Agent Approach to Accurate Text Generation. In Proceedings of the 2024 International Conference on Electrical and Computer Engineering Researches (ICECER), Gaborone, Botswana, 4–6 December 2024; pp. 1–8. [CrossRef]
85. Parthasarathy, K.; Vaidhyanathan, K.; Dhar, R.; Krishnamachari, V.; Kakran, A.; Akshathala, S.; Arun, S.; Karan, A.; Muhammed, B.; Dubey, S.; et al. Engineering LLM Powered Multi-Agent Framework for Autonomous CloudOps. In Proceedings of the 2025 IEEE/ACM 4th International Conference on AI Engineering—Software Engineering for AI (CAIN), Ottawa, ON, Canada, 27–29 April 2025; pp. 201–211. [CrossRef]
86. Afzal, A.; Vladika, J.; Fazlija, G.; Staradubets, A.; Matthes, F. Towards Optimizing a Retrieval Augmented Generation using Large Language Model on Academic Data. In Proceedings of the 2024 8th International Conference on Natural Language Processing and Information Retrieval, Bangkok, Thailand, 18–20 December 2024; NLPIR '24, pp. 250–257. [CrossRef]
87. Roychowdhury, S.; Krema, M.; Mahammad, A.; Moore, B.; Mukherjee, A.; Prakashchandra, P. ERATTA: Extreme RAG for enterprise-Table To Answers with Large Language Models. In Proceedings of the 2024 IEEE International Conference on Big Data (BigData), Washington, DC, USA, 15–18 December 2024; pp. 4605–4610. [CrossRef]
88. Gekhman, Z.; Yona, G.; Aharoni, R.; Eyal, M.; Feder, A.; Reichart, R.; Herzig, J. Does Fine-Tuning LLMs on New Knowledge Encourage Hallucinations? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*; Al-Onaizan, Y., Bansal, M., Chen, Y.N., Eds.; Association for Computational Linguistics: Miami, FL, USA, 2024; pp. 7765–7784. [CrossRef]
89. Jiang, X.; Fang, Y.; Qiu, R.; Zhang, H.; Xu, Y.; Chen, H.; Zhang, W.; Zhang, R.; Fang, Y.; Ma, X.; et al. TC–RAG: Turing–Complete RAG's Case study on Medical LLM Systems. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; Che, W., Nabende, J., Shutova, E., Pilehvar, M.T., Eds.; Association for Computational Linguistics: Vienna, Austria, 2025; pp. 11400–11426. [CrossRef]
90. Ye, F.; Li, S.; Zhang, Y.; Chen, L. R2AG: Incorporating Retrieval Information into Retrieval Augmented Generation. *arXiv* **2024**, arXiv:2406.13249.
91. Du, X.; Zheng, G.; Wang, K.; Feng, J.; Deng, W.; Liu, M.; Chen, B.; Peng, X.; Ma, T.; Lou, Y. Vul-RAG: Enhancing LLM-based Vulnerability Detection via Knowledge-level RAG. *arXiv* **2024**, arXiv:2406.11147. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.