International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 12-No.1 (2026) pp. 770-777
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

Closed-Loop Hallucination Mitigation in Generative Language Systems Through Adaptive Retrieval, Multi-Source Verification, and Judge-Guided Feedback

Ajay Athitya Ramanathan*

FourthSquare LLC, USA
* **Corresponding Author Email:** ajayathityaramanathan@gmail.com - **ORCID:** 0000-0002-1887-6650

Article Info:

DOI: 10.22399/ijcesen.4907
Received : 29 November 2025
Revised : 25 January 2026
Accepted : 02 February 2026

Keywords

Hallucination Mitigation,
Retrieval-Augmented Generation,
Claim Verification,
Adaptive Systems,
Generative Language Models

Abstract:

Generative language technologies have experienced remarkable transformation over recent years, evolving from laboratory prototypes into production-ready infrastructure serving enterprise analytics, strategic decision frameworks, and automated information services. Technical maturity notwithstanding, these platforms persistently generate linguistically sophisticated text lacking verifiable factual foundations—termed hallucinations within technical discourse. This phenomenon introduces considerable operational hazards across sensitive application contexts, particularly analytics transformation projects, automated documentation workflows, article compilation activities, and strategic advisory operations, wherein factually incorrect outputs may propagate through organizational systems without triggering detection mechanisms. Contemporary control strategies predominantly implement post-generation validation procedures or utilize static document retrieval architectures, addressing observable manifestations while fundamental causative factors remain unresolved. The architectural methodology introduced here reconceptualizes factual precision as an actively maintained operational characteristic rather than a discrete validation checkpoint executed after content generation. Synthesizing adaptive document retrieval procedures, multifaceted assertion assessment protocols, inter-source concordance analysis, and evaluation-guided regeneration mechanisms, this architectural framework facilitates continuous operational self-correction. The article transforms epistemic uncertainty from concealed system deficiency into an explicit, communicable system attribute, thereby enabling informed user interpretation. This design philosophy establishes that effective hallucination control necessitates integrated architectural planning rather than supplementary filtering layers, enabling trustworthy deployment across enterprise and research operational contexts.

1. Contextual Background and Technical Problem Definition

1.1 Factual Accuracy Challenges in Operational Deployments

Modern generative language platforms have achieved production-grade maturity, becoming integral infrastructure components throughout enterprise analytics ecosystems, strategic planning frameworks, and automated knowledge distribution systems [1]. These technologies demonstrate substantial text generation capabilities, producing human-equivalent linguistic output across varied organizational contexts and application scenarios. Despite these advances, a persistent technical limitation threatens deployment viability within mission-critical environments: the generation of linguistically fluent assertions lacking verifiable evidentiary grounding, designated as hallucinations throughout scholarly literature. This behavioral pattern materializes as confident declarative statements absent substantiation through retrieved documentation, training corpus references, or established factual knowledge bases. Within operational domains including analytics modernization initiatives, automated report generation, research synthesis pipelines, and policy development support frameworks, such fabricated content constitutes substantial operational liability, as erroneous assertions may cascade through organizational decision architectures and influence consequent actions without activating detection protocols.

1.2 Inadequacies of Prevailing Control Methodologies

Established methodologies for hallucination management predominantly operate as supplementary validation layers rather than foundational system components [2]. Conventional approaches depend upon probability-derived scoring mechanisms that analyze token likelihood distributions or extract uncertainty signals directly from model outputs themselves. These signals provide analytically useful indicators, but they don't always correlate with factual accuracy, especially when models make wrong claims and have high confidence metrics. Alternative frameworks utilize secondary classification architectures to classify outputs within binary safe-unsafe dimensions; however, these taxonomic methods are inadequate for identifying specific problematic assertions or recommending targeted corrective measures. Static document retrieval architectures compound these limitations through lacking structured recovery mechanisms when initial retrieval operations yield inadequate, temporally obsolete, or contextually misaligned documentary evidence. The resulting operational behavior is brittle, meaning that unsupported claims either go unnoticed through downstream processes or cause a complete rejection of responses without any specific fixes.

1.3 Transitioning Toward Active Control Architectures

Recent technical developments in retrieval-enhanced generation methodologies have established that anchoring language model outputs to external documentary sources reduces fabrication frequency through constraining response generation to retrieved evidentiary foundations. Predominant implementations nevertheless maintain static operational characteristics, relying upon single-pass retrieval sequences and coarse-grained confidence estimation heuristics. Such architectures encounter substantial difficulties processing nuanced assertions, reconciling contradictory source materials, accommodating temporal information currency fluctuations, and accounting for heterogeneous source reliability characteristics. The architectural framework advanced throughout this work represents a fundamental methodological transition from passive anomaly detection toward active operational control through closed-loop hallucination mitigation infrastructure. This approach reconceptualizes factual accuracy as a continuously evaluated system invariant rather than a singular post-generation validation event. Through integrating adaptive document retrieval procedures, logical entailment verification protocols, source credibility assessment mechanisms, and evaluation-guided content regeneration processes, this system architecture enables real-time operational self-correction that substantially diminishes unsupported assertions while preserving linguistic fluency and functional utility.

2. Foundational Architecture: Multifaceted Assertion Assessment

2.1 Composite Scoring Across Independent Analytical Dimensions

The architectural framework introduces a weighted assertion assessment protocol evaluating individual generated claims along multiple independent analytical axes rather than relying upon singular correctness indicators [3]. The system synthesizes evidence streams encompassing logical entailment verification, probabilistic confidence quantification, source credibility metrics, and temporal information currency scoring. Each analytical dimension adds weighted numerical values to composite reliability measurements that are based on individual assertions instead of groups of responses. This granular assessment methodology enables explicit computational reasoning regarding assertions, maintaining robust evidentiary certainty for those possessing qualified confidence characteristics and those resisting verification entirely given available source materials. The multidimensional scoring infrastructure establishes structured analytical foundations for subsequent adaptive determinations regarding retrieval expansion operations, targeted content regeneration procedures, and transparency implementation mechanisms.

2.2 Logical Entailment Verification Through Inference Models

Logical entailment verification procedures establish whether generated assertions receive semantic support from retrieved documentary materials through natural language inference methodologies [4]. Rather than employing superficial keyword matching algorithms or surface-level textual similarity calculations, the system applies trained neural inference architectures to determine whether source documents semantically entail generated assertions through logical implication relationships. This analytical procedure identifies unsupported claims even when maintaining topical coherence or stylistic consistency with the surrounding textual context. The entailment verification operates at sentence-level analytical granularity, generating binary support classifications alongside continuous probabilistic confidence metrics. Assertions failing defined entailment threshold criteria receive flagging for targeted intervention through focused retrieval expansion and constrained regeneration procedures, enabling precise correction without eliminating previously validated content segments.

2.3 Source Credibility and Information Currency Assessment

Credibility assessment procedures for sources prioritize retrieved documents according to editorial standards, specialized domain authority, citation pattern analysis, and historical reliability performance indicators. The system maintains dynamically updated reputation profiles characterizing individual information sources, applying elevated weighting coefficients to evidence originating from high-authority sources within composite reliability calculations. This adaptive weighting mechanism reduces false negative classification instances wherein accurate assertions from authoritative sources might otherwise face questioning due to surface-level ambiguity or contextual complexity. Temporal information currency scoring implements mathematical decay functions applied to temporally distant sources within rapidly evolving knowledge domains, adjusting confidence assessments according to information recency characteristics. These mechanisms collectively generate granular reliability profiles, enabling sophisticated reasoning regarding evidentiary support strength, transcending simplistic binary true-false classifications toward graduated probabilistic confidence assessments reflecting actual epistemic uncertainty conditions present within available documentary evidence.

3. Dynamic Operational Mechanisms: Targeted Retrieval and Precision Regeneration

3.1 Focused Retrieval for Deficient Confidence Assertions

A fundamental architectural innovation involves adaptive retrieval strategies conducting targeted document search operations concentrated exclusively upon identified problematic assertions rather than re-executing comprehensive retrieval pipelines [5]. When verification systems find claims with lower entailment scores or lower consensus metrics, focused retrieval operations start using the specific flagged assertion as a more precise query context. This surgical intervention approach enables operationally efficient remediation through concentrating computational resources upon specific areas requiring evidentiary reinforcement. The targeted retrieval increases the number of document collections by adding sources that are directly related to the flagged claims. This gives regeneration systems better contextual grounding materials. This operational strategy yields substantial efficiency improvements relative to complete response regeneration approaches, reducing unnecessary content variation within previously validated segments while strengthening deficient areas through precision-targeted intervention.

3.2 Cyclical Generation with Maintained Operational State

The content generation workflow operates through iterative cycles with preserved state maintenance across sequential correction operations [6]. Initial draft production proceeds utilizing preliminary document collections retrieved through standard operational mechanisms. A computational judge component subsequently evaluates discrete assertions, identifying regions exhibiting diminished entailment characteristics or reduced consensus metrics requiring targeted intervention. For each flagged area, the system dynamically expands retrieval operations, looking for additional or more authoritative documentary sources that are specifically relevant to the flagged assertions. Regeneration procedures are constrained to unsupported sections exclusively, preserving previously validated content while reinforcing identified problematic areas. This cyclical methodology produces response outputs converging toward elevated factual reliability through controlled iterative refinement rather than wholesale rejection strategies. The stateful operational architecture maintains contextual awareness regarding which assertions have undergone successful validation, which requires additional evidentiary support, and which remains unverifiable given currently available source materials.

3.3 Achieving Stability Through Controlled Cyclical Refinement

The iterative correction workflow operates with defined termination criteria, balancing reliability enhancement objectives against computational efficiency constraints. Following each regeneration cycle, the system re-evaluates modified assertions employing identical multidimensional scoring frameworks utilized during initial assessment.

When calculated reliability scores exceed predefined threshold parameters, assertion acceptance occurs and processing advances. When scores persist below threshold values following maximum permitted iteration counts, the system either eliminates the problematic assertion, substitutes conservative hedge formulations, or appends explicit flags indicating unverifiable status. This controlled iterative approach prevents indefinite correction loops while ensuring accepted content satisfies established reliability standards. The operational methodology mirrors expert human reasoning patterns wherein assertions undergo progressive refinement through evidence accumulation and revision cycles until reaching acceptable confidence thresholds or receiving explicit acknowledgment as uncertain given available information.

4. Operational Transparency and Challenge-Based Validation

4.1 Inter-Source Agreement Analysis Mechanisms

Addressing contradictions and ambiguity present across retrieved source materials, the framework incorporates inter-source concordance scoring, evaluating assertions according to agreement patterns among independent documentary sources [7]. Instead of hiding disagreements or picking one authoritative version by giving it more weight, the system makes sure that identified contradictions are clearly shown as structured qualifiers in the responses it generates. When sources furnish conflicting factual information, the framework presents observed disagreements alongside descriptive metadata characterizing relevant source attributes, enabling informed user interpretation and judgment. This methodological approach transforms epistemic uncertainty from a concealed operational failure mode into transparent system attributes supporting calibrated user interpretation. The concordance scoring system uses source credibility metrics to give more weight to agreement patterns that come from high-authority sources than to those that come from lower-credibility sources.

4.2 Visual Confidence Representation

The framework generates visual confidence representations accompanying textual response outputs, indicating calculated reliability levels at sentence or paragraph analytical granularity. These visualization elements enable rapid user identification of high-confidence assertions versus regions requiring interpretive caution or supplementary independent verification. Color-coded visual indicators or numerical confidence scores are embedded within user interface layers, rendering epistemic uncertainty immediately visible without requiring users to parse underlying technical metadata structures. This transparency mechanism proves particularly valuable within analytics and research operational contexts wherein partial epistemic uncertainty remains operationally acceptable provided clear communication occurs. Through exposing confidence metadata as first-class interface elements, the system cultivates calibrated user trust patterns wherein users may rely substantially upon high-confidence assertions while approaching low-confidence regions with appropriate interpretive skepticism.

4.3 Evolving Source Credibility Models

The framework incorporates machine learning mechanisms constructing dynamically updated source credibility graphs across extended operational timeframes [8]. Documentary sources undergo continuous evaluation along multiple analytical dimensions, including specialized domain expertise, editorial standard adherence, citation density patterns, publication venue prestige indicators, and historical factual accuracy within previous verification operations. The system tracks the statistical frequency with which assertions derived from each source successfully pass entailment validation protocols and consensus verification procedures. Sources consistently furnishing well-supported factual information receive progressively elevated credibility scores, while sources frequently associated with unsupported or contradicted assertions receive progressively diminished scores. This adaptive weighting mechanism enables continuous system evolution of trust models as new information environments are encountered during operation, improving statistical discrimination between high-credibility and low-credibility sources without requiring manual curation of source hierarchy taxonomies.

4.4 Challenge-Based Validation and Dialectical Verification

To stress-test generated outputs, an adversarial computational judge component attempts to identify plausible counterclaims, boundary cases, or absent qualifying caveats using identical documentary source collections available to primary generative systems. The adversarial component conducts systematic searches for evidence potentially contradicting or qualifying generated assertions, even when those assertions successfully pass standard entailment validation protocols. Through this dialectical process, the system identifies successful challenges and appends explicit assumptions, operational limitations, or alternative interpretive perspectives to the generated responses. This challenge-based validation methodology guarantees that generated answers are not just factually correct in isolation, but also withstand reasonable critical examination. The operational approach aligns system behavior with scholarly argumentation standards, which say that claims must be able to stand up to critical intellectual challenge and clearly recognize valid counterarguments or relevant boundary conditions.

5. Operational Applications and Broader Organizational Implications

5.1 Enterprise Analytics and Strategic Intelligence Platforms

Closed-loop hallucination mitigation architectures carry direct operational implications for enterprise analytics platforms and strategic decision intelligence systems [9]. Contemporary organizations increasingly deploy generative language technologies for automated reporting workflows, strategic insight generation, and executive decision support functions, yet persistent concerns regarding factual accuracy significantly constrain adoption within mission-critical operational contexts. The architectural framework proposed here enables operationally safe deployment through embedding continuous self-correction capabilities and transparency mechanisms directly within generative workflow infrastructures. Rather than depending upon resource-intensive manual review processes or accepting operational risks associated with unsupported factual claims, organizations may implement systems continuously validating their own generated outputs through embedded verification cycles. This architectural capability supports scalable operational automation while preserving rigorous epistemic standards, enabling systems to prove simultaneously intelligent and operationally accountable. The framework's transparency mechanisms further support regulatory compliance requirements and organizational audit protocols by providing explicit documentation trails of calculated confidence levels and the underlying evidentiary support.

5.2 Research Compilation and Automated Knowledge Synthesis

Within research synthesis and automated literature review application contexts, the framework addresses critical operational requirements for accurate information aggregation and rigorous source attribution [10]. Research professionals require systems capable of synthesizing empirical findings across multiple documentary sources while maintaining scholarly standards for evidence citation and claim substantiation. The multidimensional verification approach ensures synthesized assertions accurately reflect underlying source materials, while concordance scoring mechanisms appropriately handle observed disagreements within existing literature. The challenge-based validation component assists in identifying specific areas requiring human expert review through systematically surfacing potential contradictions or boundary cases automated systems cannot confidently resolve independently. Through rendering epistemic uncertainty explicit and operationally actionable, the framework enables productive human-machine collaborative workflows wherein automated systems handle well-supported synthesis operations while escalating ambiguous or contested areas for expert human judgment.

5.3 Ethical Information Distribution

From broader societal perspectives, such architectural approaches promote responsible information distribution practices by reducing the propagation of unsupported factual claims and explicitly communicating epistemic uncertainty to end users. As generative systems increasingly mediate knowledge access through conversational search interfaces, automated assistance agents, and algorithmic content generation platforms, embedding robust hallucination mitigation safeguards becomes a fundamental prerequisite for ethical technological deployment. The framework transparency mechanisms assist users in developing appropriate calibration in their trust of system outputs, enabling understanding of when to rely upon generated content and when to seek supplementary independent verification. This architectural approach contrasts sharply with systems projecting uniform confidence characteristics across all generated outputs regardless of underlying evidentiary support quality, which may inadvertently foster either excessive uncritical trust or blanket skepticism. By making graduated confidence characteristics clear in user interfaces, the framework helps people make truly informed decisions while still keeping the system useful for its main purpose.

5.4 Architectural Principles for Trustworthy Artificial Intelligence

The architectural framework proposed here shows that effective hallucination mitigation is not just a filtering problem that can be solved with post-hoc validation procedures. It is a fundamental architectural challenge that needs to be integrated into the design of the whole system. The closed-loop operational approach treats factual correctness as a continuously maintained system invariant rather than an emergent property evaluated exclusively after content generation. This architectural design perspective opens pathways for broader trustworthy artificial intelligence development wherein reliability constraints embed fundamentally within system architectures rather than applying as supplementary external guardrails. The multidimensional verification protocols, adaptive retrieval mechanisms, and challenge-based validation components furnish reusable architectural patterns applicable beyond hallucination mitigation to other artificial intelligence safety challenges requiring robust evidence handling capabilities and sophisticated uncertainty quantification.

Table 1: *Comparison of Hallucination Control Approaches [1][2]*

Control Methodology	Operational Characteristics	Primary Limitations	Recovery Mechanism	Granularity Level
Post-Generation Confidence Scoring	Analyzes token probability distributions after generation	Inconsistent correlation with factual accuracy	None-accepts or rejects entire response	Response-level only
Binary Safety Classification	Secondary model categorizes outputs as safe/unsafe	Cannot identify specific problematic assertions	Wholesale rejection without targeted correction	Response-level only
Static Retrieval Pipelines	Single-pass document retrieval before generation	No recovery when initial retrieval insufficient	Requires complete regeneration cycle	Document-level only
Closed-Loop Adaptive Control	Continuous evaluation during generation process	Increased computational overhead	Surgical remediation of specific assertions	Claim-level granularity
Judge-Guided Regeneration	Real-time entailment checks at sentence boundaries	Requires trained inference models	Targeted retrieval and constrained regeneration	Sentence-level precision

Table 2*: Multi-Dimensional Claim Verification Framework Components [3] [4]*

Verification Dimension	Assessment Criteria	Output Metric	Operational Threshold	Intervention Trigger
Logical Entailment	Semantic support from source documents	Binary classification with confidence score	0.75 minimum confidence	Below-threshold initiates targeted retrieval
Probabilistic Confidence	Token likelihood distribution analysis	Continuous probability value	0.65 minimum reliability	Low confidence flags for regeneration
Source Credibility	Editorial standards and domain authority	Weighted reputation score	0.70 minimum trust level	Low credibility triggers additional verification
Temporal Currency	Information recency with decay functions	Time-adjusted relevance score	0.60 minimum currency	Outdated sources prompt refresh cycle
Cross-Source Consensus	Agreement patterns among independent sources	Concordance percentage	60% minimum agreement	Conflicts surface as explicit qualifiers

Table 3*: Iterative Correction Process Workflow [5][6]*

Process Stage	Operational Action	System State	Evaluation Criteria	Next Stage Determination
Initial Draft	Produce content	Draft created	Overall response	Proceed to assertion

Generation	using preliminary document set	with mixed confidence levels	coherence	evaluation
Judge-Based Assessment	Evaluate individual assertions for entailment	Flagged claims identified	Entailment score below 0.75 threshold	Low-confidence claims trigger retrieval
Targeted Retrieval Expansion	Search for claim-specific documentary evidence	Expanded document collection available	New sources retrieved successfully	Proceed to surgical regeneration
Surgical Regeneration	Regenerate only flagged assertion sections	Partial content updated	Reliability improvement measured	Re-evaluate modified assertions
Convergence Validation	Check if reliability thresholds met	Iteration count tracked	Scores exceed thresholds or max iterations	Accept, hedge, or flag as unverifiable
State Preservation	Maintain validation status of all assertions	Complete validation map maintained	All assertions categorized	Final response assembly

Table 4: *Dynamic Source Reputation Assessment Framework [7][8]*

Credibility Dimension	Evaluation Metrics	Scoring Range	Weight Factor	Impact on Trust Model
Domain Expertise	Specialized knowledge depth and topical focus	0.0 to 1.0	0.25	High expertise increases entailment acceptance
Editorial Standards	Peer review processes and fact-checking protocols	0.0 to 1.0	0.20	Rigorous standards reduce verification cycles
Citation Density	Reference frequency in scholarly literature	0.0 to 1.0	0.15	High citation elevates source priority
Publication Venue Prestige	Journal rankings and conference tier classifications	0.0 to 1.0	0.15	Prestigious venues receive higher initial trust
Historical Accuracy	Past performance in entailment validation	0.0 to 1.0	0.25	Consistent accuracy progressively elevates scores.
Temporal Consistency	Stability of information across time periods	0.0 to 1.0	Variable	Recent contradictions trigger trust decay

6. Conclusions

Hallucinations persist as among the most significant technical obstacles to trustworthy adoption of generative language systems within high-stakes operational applications. The architectural framework advanced throughout this work transforms hallucination handling from passive anomaly detection into active, continuously operating adaptive control. The architecture integrates weighted assertion verification across multiple independent analytical dimensions, enabling granular reliability assessment at the individual claim level rather than the aggregate response level. Adaptive retrieval mechanisms enable surgical remediation focused upon identified problematic assertions without discarding previously validated content segments. Inter-source concordance scoring and challenge-based validation ensure operational robustness under critical examination while maintaining transparency regarding epistemic uncertainty and underlying evidentiary support. Instead of hiding uncertainty by filtering or leaving it out, the framework makes it clear, usable, and easy to understand through visualization tools and structured qualifying statements. The operational article demonstrates that effective hallucination mitigation requires fundamental architectural integration rather than supplementary external filtering, treating factual correctness as a continuously evaluated system property. As intelligent automation technologies mature and generative systems assume increasingly critical operational roles within enterprise infrastructures, research institutions, and broader societal information ecosystems, such self-correcting architectural approaches will prove essential to aligning generative capabilities with scholarly rigor, enterprise operational reliability, and broader societal trust. Future research investigations should explore extending these architectural mechanisms to multimodal content generation, real-time streaming output scenarios, and collaborative multi-agent operational systems wherein hallucination mitigation must operate across distributed architectural components with varying information access characteristics and heterogeneous reliability requirements.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

[1] Esmail Gumaan, "Theoretical Foundations and Mitigation of Hallucination in Large Language Models," arXiv (cs.CL), 20 July 2025. Available: https://arxiv.org/html/2507.22915v1?utm_source=copilot.com

[2] Yanyi Liu, et al., "Reducing Hallucinations of Large Language Models via Hierarchical Semantic Piece," Complex & Intelligent Systems (Springer, Vol. 11, Article 231), 02 April 2025. Available: https://link.springer.com/article/10.1007/s40747-025-01833-9?utm_source=copilot.com

[3] Bangrui Xu, et al., "Complex Claim Verification via Human Fact-Checking Imitation with Large Language Models," IEEE Transactions on Knowledge and Data Engineering, 2025. Available: https://ieeexplore.ieee.org/document/10799276?utm_source=copilot.com

[4] Miloš Košprdić, et al., "Scientific Claim Verification with Fine-Tuned NLI Models," Proceedings of SCITEPRESS, 2024. Available: https://www.scitepress.org/Papers/2024/129000/129000.pdf?utm_source=copilot.com

[5] Richard Shan, "LearnRAG: Implementing Retrieval-Augmented Generation for Adaptive Knowledge Systems," 2025 International Conference on Artificial Intelligence in Information and Communication (ICAIIC), 19 March 2025. Available: https://ieeexplore.ieee.org/document/10920869?utm_source=copilot.com

[6] Pit Pichappan, "Adaptive and Enhanced Retrieval Augmented Generation (RAG) Systems: A Summarized Survey," Journal of Digital Information Management (JDIM, Vol. 23, Issue 3), 2025. Available: https://www.dline.info/fpaper/jdim/v23i3/jdimv23i3_3.pdf?utm_source=copilot.com

[7] Henri Sohier, et al., "The Engineering of AI Evaluation and Scoring: Overview and Practices," 2025 IEEE International Systems Conference (SysCon), 30 May 2025. Available: https://ieeexplore.ieee.org/abstract/document/11014820?utm_source=copilot.com

[8] Rommel Gutierrez, William Villegas-Ch, and Jaime Govea, "Adaptive Consensus Optimization in Blockchain Using Reinforcement Learning and Validation in Adversarial Environments," Frontiers in Artificial Intelligence, Vol. 8, 30 September 2025. Available: https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2025.1672273/full?utm_source=copilot.com

[9] Nikhil Vishnu Vadlamudi, et al., authored the paper titled "Leveraging Generative AI Trained on Enterprise Data for Business Value in Large Enterprises," which will be presented at the 2025 IEEE Technology and Engineering Management Society (TEMSCON LATAM) on 24 November 2025. Available: https://ieeexplore.ieee.org/document/11238798?utm_source=copilot.com

[10] Murali Thirumalaisamy, et al., "Leveraging Generative AI for Mainframe Application Modernization to Cloud-Based Architectures," IEEE Dataport, 2025. Available: https://ieee-dataport.org/documents/leveraging-generative-ai-mainframe-application-modernization-cloud-based-architectures?utm_source=copilot.com