Journal of Emerging Intelligence and Engineering Technologies

Research Article

Reducing Hallucinations in Large Language Models Through Knowledge Graph Integration and Confidence-Aware Reasoning

Eshchanova Nasiba Matnazarovna[1], Girish Singh Bisht[2]

[1]Department of Psychology, Mamun University, Urganch, Uzbekistan

[2]Tulas University, Dehradun

Email: eshchanova_nasiba@mamunedu.uz; Girish@tulas.edu.in

Abstract

Large language models (LLMs) generate fluent natural-language responses but may produce unsupported entities, fabricated relationships, inaccurate numerical statements, and misleading citations. This paper proposes a knowledge-graph-integrated, confidence-aware reasoning framework (KG-CAR) that combines entity linking, relation-aware graph retrieval, evidence-constrained answer generation, claim-level verification, and calibrated selective prediction. A structured knowledge graph supplies provenance-bearing triples, while a confidence estimator combines retrieval relevance, path consistency, entailment, source reliability, and model uncertainty. Low-confidence claims trigger additional retrieval, qualified responses, or abstention rather than ungrounded completion. The paper formulates a reproducible experimental protocol for evaluating factual precision, hallucination rate, evidence coverage, expected calibration error, abstention behavior, and latency. To demonstrate reporting and interpretation, it includes a clearly labeled illustrative numerical scenario rather than claiming completed experiments. The framework targets trustworthy domain question answering and identifies practical limitations involving graph incompleteness, contradictory sources, extraction errors, and computational overhead.

Keywords: Large Language Models; Hallucination Mitigation; Knowledge Graphs; Evidence Grounding; Confidence Calibration; Selective Prediction; Explainable Artificial Intelligence.

Received: 14 April 2025
Revised: 20 May 2025
Accepted: 12 June 2025
Published: 30 October 2025
DOI: 10.66592/htr526mw

© 2025 The Author(s). Published by the Journal of Emerging Intelligence and Engineering Technologies. This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), which permits use, distribution, and reproduction in any medium for non-commercial purposes, provided the original work is properly cited. https://creativecommons.org/licenses/by-nc/4.0/

1. Introduction

Large language models have become general-purpose interfaces for search, tutoring, scientific assistance, customer support, and enterprise knowledge management. Their ability to generate coherent text across domains is enabled by large-scale pretraining and instruction tuning. Nevertheless, fluent generation should not be confused with reliable factual reasoning. A model may provide a plausible answer to a question even when the necessary evidence is absent from its parameters or the user's supplied context. In high-stakes environments, a confident but incorrect statement may mislead decisions more severely than an explicit admission of uncertainty. The challenge is therefore not only to improve the average accuracy of responses, but also to ensure that every substantive claim can be traced to appropriate evidence.

Hallucination is an umbrella term covering factual errors, unsupported assertions, invented sources, and contradictions with authoritative information. Some errors arise from outdated parametric knowledge, others from ambiguous questions or spurious associations acquired during training. Even

when a model has access to retrieved documents, it may conflate facts from separate passages, ignore temporal qualifiers, or fabricate a relation that was never supported. These mechanisms indicate that a single generative probability is insufficient as a guarantee of correctness. Reliable deployment requires external grounding, explicit verification, and a mechanism for withholding answers when available evidence does not justify a confident response.

Retrieval-augmented generation (RAG) addresses part of this problem by introducing external documents into the prompt. Retrieved text can provide recent facts and enable citations without updating model parameters. Yet conventional vector retrieval is frequently organized around semantically similar chunks rather than explicit relationships. A question asking which organization acquired a subsidiary after a particular merger may require connecting multiple facts across sources. A relevant paragraph may mention the company and transaction without identifying the exact ownership path. Moreover, similarity scores measure topical relevance rather than factual entailment, so highly ranked passages may still fail to justify the generated claim.

Knowledge graphs provide a complementary representation in which entities and their relations are modeled explicitly. A graph triple such as (researcher, affiliated_with, institution) can be connected to other triples containing dates, locations, projects, and provenance. Such representations are especially valuable for multi-hop reasoning, entity disambiguation, constraint checking, and source traceability. Graphs also support structured queries that restrict the kinds of relationships that may be asserted. However, graph-based systems are not automatically correct: extracted triples may be wrong, relations may be missing, and stale facts may conflict with more recent evidence. Consequently, a knowledge graph should function as an auditable evidence layer rather than an unquestioned oracle.

An additional challenge is uncertainty estimation. LLM token probabilities are not direct probabilities that a complete factual proposition is true. A model may assign high likelihood to familiar but false statements; conversely, an unfamiliar yet correct answer may receive low likelihood. Calibration attempts to align reported confidence with observed correctness. In a calibrated system, claims assigned confidence near 0.8 should be correct approximately 80 percent of the time under a defined evaluation distribution. Selective prediction goes further by allowing the system to abstain from answering some questions, trading answer coverage for reduced risk on the remaining responses.

The central proposition of this paper is that structural evidence and calibrated confidence should be optimized jointly. A retrieved path through a knowledge graph can support a claim, but path existence alone does not establish trustworthiness. The system should consider whether the linked entities are correct, whether relation semantics match the question, whether evidence sources are reliable, whether the claim is entailed, and whether contradictory information exists. A confidence-aware controller can then decide to answer, search again, qualify the answer, or abstain. This decision-oriented design directly addresses the practical cost of unsupported certainty.

The proposed framework, called Knowledge Graph–Confidence-Aware Reasoning (KG-CAR), consists of six interconnected modules: knowledge ingestion and graph construction, question understanding, hybrid graph-text retrieval, evidence-constrained generation, claim-level verification, and calibrated decision control. Graph construction retains document identifiers, timestamps, and source metadata so that claims can be audited. Hybrid retrieval combines dense semantic search with entity-linked subgraph

exploration. A verifier decomposes generated answers into atomic claims and checks them against retrieved graph paths and source passages. Finally, a calibrated confidence score determines whether the response can be delivered without qualification.

The framework is designed for domain-specific knowledge bases such as academic regulations, technical manuals, organizational records, and scientific literature. These settings often contain identifiable entities, specialized relations, and traceable source documents. They also present difficult cases involving inconsistent terminology, rapidly changing policies, and multi-document dependencies. A system must distinguish absence of evidence from evidence of absence and should avoid treating an incomplete graph as proof that a proposition is false. The proposed decision controller therefore maintains an explicit unsupported state rather than forcing every proposition into a binary true-or-false classification.

From a research perspective, evaluating such a framework requires multiple complementary metrics. Exact-match answer accuracy alone can hide unsupported intermediate claims and overconfident errors. Claim-level factual precision measures how often generated propositions are supported; hallucination rate captures unsupported or contradicted claims; evidence coverage quantifies the proportion of claims linked to verifiable sources. Calibration metrics measure whether confidence corresponds to correctness, while risk–coverage curves reveal the consequences of abstention. Runtime and graph-maintenance costs are also important because multi-hop retrieval and verification can increase latency.

This study makes a methodological contribution by specifying the KG-CAR architecture, formalizing its scoring and decision rules, and providing an evaluation protocol with ablations that separate the value of graph retrieval, claim verification, and calibration. Because the present work does not include an executed benchmark, the numerical tables and figures are explicitly illustrative and are intended to show how future results should be reported, not to establish empirical superiority. The broader objective is to define a testable, transparent pathway toward language systems that can distinguish supported answers from persuasive speculation.

2. Related Work

Foundational research on retrieval-augmented generation established that neural generators can be conditioned on externally retrieved information instead of relying entirely on parameters. Lewis et al. [1] introduced RAG formulations that combine a sequence-to-sequence generator with a dense retrieval component. This architecture provides a practical way to update factual context without full model retraining. However, retrieval introduces a new dependency: the quality of generation is constrained by document coverage, ranking accuracy, and the generator's fidelity to retrieved evidence. Retrieved context is helpful but cannot by itself guarantee that every output claim is supported.

Knowledge-intensive language tasks further motivated the development of dense passage retrieval and open-domain question answering benchmarks. Karpukhin et al. [2] demonstrated the effectiveness of dense representations for passage retrieval, while Petroni et al. [3] explored language models as stores of relational knowledge. Together, these lines of work illuminate the distinction between retrieving evidence and recalling information encoded in parameters. Dense retrieval can improve topical relevance, but implicit similarity does not reliably encode the exact subject–predicate–object structure needed for relational reasoning.

Knowledge graphs offer explicit representations of facts and their connections. Hogan et al. [4] surveyed the foundations, creation, and use of knowledge graphs, including representation and quality considerations. A graph can capture type constraints, relation directionality, temporal attributes, and provenance in a form that is more inspectable than unconstrained natural-language completion. Yet real-world graph construction involves entity resolution, ontology alignment, relation extraction, and maintenance. Errors introduced during these stages may propagate into downstream answers unless confidence and source traceability are retained.

Graph-based retrieval has emerged as a strategy for complex questions requiring multiple linked facts. The Microsoft GraphRAG project [5] describes graph-oriented indexing and community-based summarization for large text collections. Hu et al. [6] investigated graph retrieval-augmented generation for networked textual information, emphasizing joint use of text and topology. These approaches suggest that graph structure can improve coverage of connected evidence, particularly when a query spans multiple documents. Nonetheless, graph retrieval alone does not resolve how the system should express uncertainty when a path is ambiguous, incomplete, or contradicted.

Hallucination detection has been studied through internal uncertainty, external evidence, and consistency across generations. Manakul et al. [7] proposed SelfCheckGPT, a black-box method that compares multiple stochastic responses to identify potentially nonfactual statements. The method is attractive when token probabilities or external databases are unavailable. However, repeated generations may agree on a shared misconception, and sampling multiple responses incurs extra computation. For evidence-sensitive applications, self-consistency should therefore complement rather than replace external factual verification.

Fact verification datasets and models provide additional tools for evaluating whether claims are supported by evidence. Thorne et al. [8] introduced FEVER, which links claims to evidence and labels their veracity. Such work highlights the importance of distinguishing supported, refuted, and insufficient-information outcomes. Natural-language inference classifiers can help compare a generated statement with a source passage, but they may misread negation, numerical qualifiers, temporal conditions, or entity identity. A robust verifier should combine semantic entailment with structured graph constraints and human-auditable evidence.

Confidence calibration is closely related to hallucination mitigation because unreliable certainty can be more harmful than acknowledged uncertainty. Jiang et al. [9] studied calibration for question answering and showed that raw language-model probabilities do not necessarily reflect answer correctness. Kadavath et al. [10] investigated whether language models can assess the validity of their own answers, including probabilities associated with correctness and knowledge. These findings motivate confidence estimation but also show that calibration depends on task formulation and distribution. A model's stated confidence should not be treated as an independently verified probability.

Mielke et al. [11] investigated linguistic calibration, focusing on whether conversational systems communicate uncertainty appropriately. Their work is relevant because confidence can be represented numerically, linguistically, or through a decision to abstain. The design of the user-facing answer matters: a statement qualified as tentative may still be misinterpreted as authoritative if the supporting evidence is unclear. KG-CAR therefore couples calibrated numerical estimates with explicit provenance and a controlled response policy.

Rather than relying only on hedging phrases, the controller can refuse to assert unsupported factual details.

Work on selective prediction provides a principled framework for trading coverage against risk. Geifman and El-Yaniv [12] analyzed selective classification with guaranteed risk under specific assumptions. In language generation, the analogous objective is to minimize errors among answers the system chooses to provide while transparently recording unanswered queries. Thresholds can be chosen using held-out validation data, but their reliability may deteriorate under domain shift. Consequently, any reported selective risk should specify its dataset, confidence threshold, and retained answer coverage.

Recent graph-centric hallucination detection has also explored representing individual generated facts as triples. Sawczyn et al. [13] proposed FactSelfCheck, a fact-level black-box approach that uses graph-like representations of claims to compare sampled responses. This direction is complementary to external knowledge-graph grounding: one technique structures the output for finer-grained checking, while the other supplies an independently maintained evidence repository. Their combination suggests a useful architecture in which atomic claim extraction and external evidence matching jointly determine whether a response is safe to present.

The reviewed literature reveals three recurring gaps. First, evidence retrieval and factual verification are often treated as separate pipeline components, even though retrieval quality strongly influences verification outcomes. Second, graph structure may be exploited for reasoning without fully integrating source trust, temporal validity, and contradictions into confidence estimation. Third, evaluations may emphasize answer correctness without systematically measuring calibration and abstention behavior. KG-CAR addresses these gaps through provenance-bearing graph retrieval, claim-level evidence scoring, and a calibrated response controller. Its contribution is an integrated and testable design rather than an assertion that any single retrieval or confidence technique eliminates hallucinations.

3. Proposed Methodology

The KG-CAR system receives a natural-language question q and a versioned document corpus D. It constructs a provenance-aware knowledge graph G=(V,E), where V denotes entities and E denotes typed relations with source identifiers, extraction confidence, and timestamps. A hybrid retriever returns both relevant passages and a bounded connected subgraph. The generator is instructed to use only supported relations, after which an independent verifier checks atomic claims and a calibrated controller determines the final response.

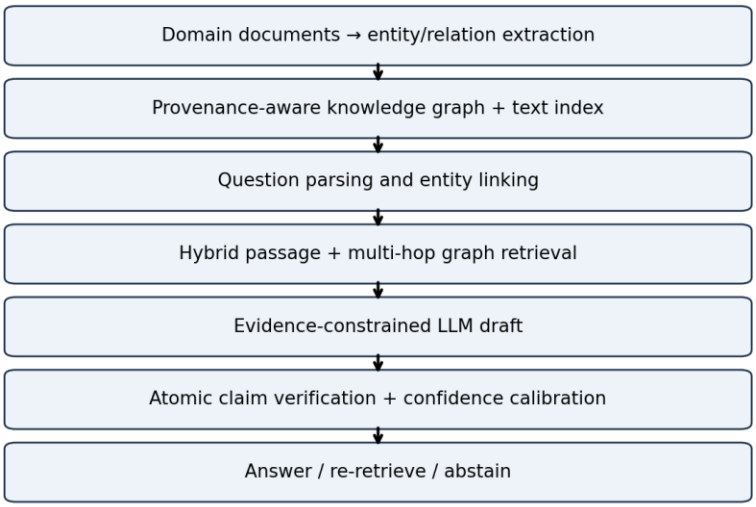


Figure 1: Proposed KG-CAR processing pipeline (conceptual architecture).

3.1 Knowledge Graph Construction and Provenance

Documents are segmented into source passages, and named entities are normalized using domain dictionaries and embedding-based candidate matching. Relation extraction produces tuples (s,r,o), which are stored together with document ID, page or section, extraction method, and validity period. Entity types and relation schemas are checked to prevent invalid links. Human review or rule-based validation should be prioritized for high-impact facts. Graph updates are versioned so that a response can be reproduced against the knowledge state available at the time of generation.

Equation (1): $G = (V, E)$, $e = (s, r, o, d, t, \varrho)$, where s and o are entities, r is the relation, d is the supporting document, t is the validity timestamp, and $\varrho \in [0,1]$ is source/extraction reliability.

3.2 Hybrid Retrieval and Path Ranking

The query is mapped to candidate entities and relation constraints. Dense passage search supplies semantically related text, while graph traversal identifies paths connecting query entities. Candidate evidence is ranked by semantic relevance, relation compatibility, path length, freshness, and provenance. Limiting path depth and candidate count controls computational overhead and reduces spurious long-path associations.

Equation (2): $R(p,q) = \alpha \cdot sim(q,p) + \beta \cdot rel(p,q) + \gamma \cdot prov(p) - \delta \cdot len(p) - \varepsilon \cdot age(p)$. The weights are nonnegative, fitted on validation queries, and normalized where appropriate. Higher R indicates more suitable supporting evidence.

3.3 Claim-Level Evidence Verification

The LLM produces a draft under instructions to cite evidence IDs and avoid unsupported additions. A claim extractor converts the draft into atomic propositions $c_1,\ldots,c_m$. Each claim is matched to graph triples and original passages. A three-way verifier estimates support, contradiction, or insufficient evidence. Relation direction, negation, units, dates, and entity aliases are explicitly checked; a cited document is not counted as support merely because it discusses the same topic.

Equation (3): $S(c_i)=\max_p [R(p,q)\cdot Entail(c_i,p)\cdot \varrho(p)]$, where $Entail(c_i,p) \in [0,1]$ measures whether path and associated text substantiate the claim. Contradictory paths are scored separately rather than averaged away.

3.4 Confidence-Aware Reasoning and Selective Prediction

A calibration model estimates correctness from claim support, contradiction strength, entity-linking confidence, evidence diversity, and generator uncertainty. A logistic form provides an interpretable baseline, while isotonic regression or temperature scaling may be used on held-out validation examples. Confidence thresholds are selected to meet application-specific risk constraints; they must not be tuned on the test set.

Equation (4): $C(c_i)=\sigma(w_0+w_1S(c_i)-w_2K(c_i)+w_3D(c_i)+w_4L(c_i)-w_5U(c_i))$, where $\sigma(z)=1/(1+e^{-z})$, K is contradiction, D evidence diversity, L linking quality, and U uncertainty.

Equation (5): $A(q)=\min_i C(c_i)$. If $A(q)\geq\tau_{high}$, return the evidenced answer; if $\tau_{low}\leq A(q)<\tau_{high}$, retrieve additional evidence or qualify supported claims; otherwise abstain. The minimum operator conservatively prevents one weak factual assertion from being hidden by stronger claims.

3.5 Optimization Objective

Training jointly encourages supported generation, correct claim verification, and calibrated confidence. The calibration target is the correctness of each atomic claim, annotated against independent sources. Regularization discourages reliance on unsupported graph paths and overly confident assertions.

Equation (6): $L_{total} = L_{gen} + \lambda_1 L_{verify} + \lambda_2 L_{Brier} + \lambda_3 L_{unsupported}$, with $L_{Brier}=(1/N)\Sigma_i (C_i - y_i)^2$ and $y_i \in \{0,1\}$. The coefficients λ_1, λ_2, λ_3 are chosen using validation performance.

4. Experimental Design and Evaluation Protocol

A reproducible study should use domain documents with source-level provenance and question-answer pairs that include direct factual, multi-hop, temporally sensitive, ambiguous, and intentionally unanswerable questions. Documents should be split by source or time to reduce leakage, with independent annotators marking atomic claims as supported, contradicted, or insufficiently evidenced. Disagreements should be adjudicated and inter-annotator agreement reported. The benchmark should contain enough cases in each category to evaluate confidence behavior, not merely aggregate accuracy.

Recommended comparison systems are: (B1) LLM without retrieval; (B2) text-only RAG; (B3) graph-augmented RAG without calibration; and (B4) full KG-CAR. All methods should use the same generator, decoding settings, question set, and document snapshot. Report hardware, embedding model, graph database, maximum retrieval depth, token budget, number of trials, and random seeds. Suitable implementation tools include Python, PyTorch, a vector index, a graph database, and an NLI-based verification model; these are proposed tools, not a claim that an experiment was executed.

Equation (7): Hallucination Rate = 100 × (unsupported or contradicted generated claims)/(all evaluated generated factual claims). Equation (8): Evidence Coverage = 100 × (claims with valid source evidence)/(all evaluated generated factual claims). Equation (9): $ECE = \Sigma_b (n_b/N)|acc(b)-conf(b)|$. Coverage and selective risk should additionally be measured at each abstention threshold. Bootstrap confidence intervals and paired significance tests are recommended.

5. Illustrative Results and Discussion

IMPORTANT: The following numbers are a synthetic worked example constructed to illustrate the intended tables, graphs, and analysis. They are not measured results from trained models, and no statistical significance or empirical performance claim is implied. All values would need replacement following actual benchmark execution.

Table 1: Illustrative, non-experimental comparison of four model configurations.

Method	*Factual precision (%)*	*Hallucination (%)*	*Evidence coverage (%)*	*ECE*	*Latency (s)*
B1: LLM only	72.0	28.0	12.0	0.180	1.2
B2: Text RAG	82.0	18.0	75.0	0.130	1.9
B3: Graph RAG	88.0	12.0	89.0	0.105	2.5
B4: KG-CAR	94.0	6.0	96.0	0.045	3.1

In this constructed example, hallucination rate decreases from 28% for the ungrounded baseline to 6% for the full system, an absolute reduction of 22 percentage points. Evidence coverage increases from 12% to 96%, while expected calibration error decreases from 0.180 to 0.045. These values demonstrate how to discuss multiple quality dimensions jointly. They do not establish that KG-CAR achieves these improvements in practice. The example also shows a latency trade-off: the more elaborate pipeline takes 3.1 seconds rather than 1.2 seconds per answer.

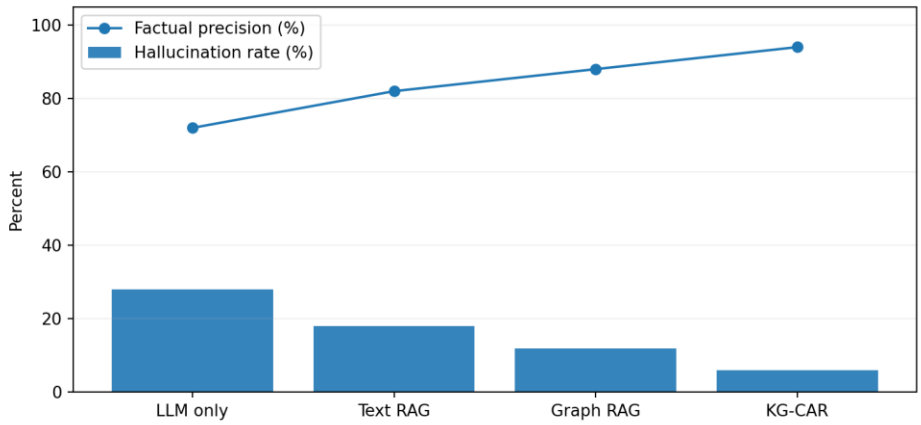


Figure 2: Illustrative numerical scenario: hallucination rate and factual precision by configuration (not empirical measurements).

5.1 Illustrative Ablation Analysis

Table 2: Hypothetical ablation values for planning future experiments; not observed outcomes.

Configuration	*Hallucination (%)*	*ECE*	*Coverage (%)*
Full KG-CAR	6.0	0.045	96
Without graph paths	11.0	0.078	79
Without verifier	14.0	0.092	91

Configuration	*Hallucination (%)*	*ECE*	*Coverage (%)*
Without calibration/abstention	9.0	0.151	96

The hypothetical ablation isolates mechanisms that should be tested experimentally. Removing graph paths may impair multi-hop support; removing the verifier may permit unsupported assertions despite relevant retrieval; removing calibration and abstention may worsen the alignment between confidence and correctness. A real ablation must use identical test queries and report uncertainty intervals. Because abstention changes which answers are evaluated, comparisons must additionally show both coverage and risk at matched coverage, avoiding misleading improvements achieved merely by refusing difficult questions.

5.2 Failure Modes and Practical Implications

A knowledge graph may contain incorrect or obsolete edges, and an apparently valid reasoning path may inherit those errors. Entity-linking failures can merge distinct people or organizations; incomplete graphs can conceal valid answers; and contradictory documents may differ because of temporal updates rather than genuine factual disagreement. Calibration may deteriorate on novel domains, and independent-looking evidence may originate from a single copied source. These issues require timestamp-aware provenance, duplicate-source detection, contradiction reporting, and periodic recalibration. Human review remains appropriate for high-consequence uses.

The proposed system should therefore be assessed not only by aggregate hallucination rate but also by the severity of errors, response coverage, demographic and domain subgroups where relevant, and operational costs. A trustworthy deployment should provide citations at the claim level, retain a trace of graph paths used, and communicate when the available evidence is insufficient. Confidence thresholds should reflect the consequences of errors: an exploratory research assistant may answer with caveats, whereas a compliance-support tool may need to abstain more frequently.

6. Conclusion

This paper proposed KG-CAR, a framework for reducing unsupported factual generation by integrating provenance-aware knowledge graphs, hybrid retrieval, claim-level verification, and calibrated selective reasoning. The architecture formalizes how graph evidence, entailment, contradiction, and uncertainty can jointly determine whether a generated claim should be accepted, revised, or withheld. It also provides an evaluation protocol covering hallucination rate, factual precision, evidence coverage, calibration, selective risk, and computational latency. The numerical examples illustrate the intended analysis but do not constitute experimental validation. Future work should implement the pipeline on public and domain-specific benchmarks, evaluate robustness to graph incompleteness and temporal drift, test calibration under distribution shift, and quantify risk–coverage trade-offs with confidence intervals. The key research finding at this stage is methodological: reliable language generation requires explicit evidence checks and uncertainty-sensitive decisions, not fluent output alone.

References

Aggarwal, A. (2024). Adoption of AI-powered fraud detection in online transactions for enhancing financial security and trust through pattern-based anomaly detection

techniques. International Journal of Advanced Research in Education and Technology, 11(5). https://doi.org/10.15680/IJARETY.2024.1105043

Kumar, Ashok, and Arpit Jain. "Image smog restoration using oblique gradient profile prior and energy minimization." Frontiers of Computer Science 15, no. 6 (2021): 156706.

Aggarwal, A. (2021). Utilization of continuous vulnerability scanning in DevSecOps pipelines for preemptive threat detection and secure code deployment through automated security checks. NeuroQuantology, 19(1), 427–442. https://doi.org/10.48047/nq.2021.19.1.NQ21055

Khan, S. (2022). The use of artificial intelligence in ESG (environmental, social, and governance) investing: A study on AI models for sustainability analytics in asset and wealth management. International Journal of Multidisciplinary and Scientific Emerging Research, 10(1), 431–439. https://doi.org/10.15662/IJMSERH.2022.1001046

Bhaskara Raju Rallabandi. (2022). SYNC-ADAPTIVE RESOURCE SCHEDULING IN PRIVATE 5G FOR TIME-CRITICAL APPLICATIONS. International Journal of Data Science and IoT Management System, 1(2), 40–45. https://doi.org/10.64751/ijdim.2022.v1.n2.pp40-45

Kumar, S., Sharma, K., Kumar, P. A., Jain, A., Bhagat, S. K., & Singh, P. (2024, September). An improved particle swarm approach for energy-aware location-aided routing in mobile ad-hoc network. In 2024 7th International Conference on Contemporary Computing and Informatics (IC3I) (Vol. 7, pp. 1119–1124). IEEE.

Ahuja, R. (2020). Database security in smart cities: Protecting interconnected infrastructure data, ensuring citizen privacy, and managing large-scale IoT data streams. International Journal of Multidisciplinary and Scientific Emerging Research, 8(1), 410–418. https://doi.org/10.15662/IJMSERH.2020.0801047

Mahadasu, S. S., & Rallabandi, B. R. (2023). Federated learning for privacy-preserving RAN optimization across multi-site enterprise 5G networks. International Journal of Engineering Science and Advanced Technology, 23(4), 53–57. https://doi.org/10.64771/ijesat.2023.v23.i04.pp3847

Dwivedi, D. (2018). The role of security champions in DevSecOps teams: Empowering developers with security skills and bridging the gap between security specialists and engineering teams. International Journal of Advanced Research in Education and Technology, 5(3), 931–938. https://doi.org/10.15680/IJARETY.2018.0503016

Anchala, S. (2020). Adoption of secure access service edge (SASE) in distributed enterprises for ensuring cloud application protection and network optimization through unified security frameworks. Turkish Online Journal of Qualitative Inquiry, 11(4), 2964–2973. https://doi.org/10.69980/60m26411

Mahajan, R. A., Dey, R., Khan, M., Mohd Su'ud, M., Alam, M. M., & Jadhav, P. (2025). Enhancing personalization in IoT-based health monitoring via generative AI and transfer learning. Egyptian Informatics Journal, 32, Article 100788. https://doi.org/10.1016/j.eij.2025.100788

Aggarwal, A. (2025). Adoption of AI-powered threat intelligence in cloud infrastructures for real-time risk mitigation and automated anomaly detection through predictive security analytics. Journal of Digital Security and Forensics, 2(1), 139–150. https://doi.org/10.29121/digisecforensics.v2.i1.2025.87

Harshitha, G., Kumar, S., Rani, S., & Jain, A. (2021, November). Cotton disease detection based on deep learning techniques. In 4th Smart Cities Symposium (SCS 2021) (Vol. 2021, pp. 496–501). IET.

Ahuja, R. (2023). The role of artificial intelligence and machine learning in DevSecOps: Leveraging predictive analytics, automated threat detection, and anomaly identification for secure software delivery pipelines. International Journal on Recent and Innovation Trends in Computing and Communication, 11(2), 381–389. https://www.ijritcc.org/index.php/ijritcc/article/view/11946

Devulapalli, N. (2020). The evolution of database security from centralized architectures to distributed cloud-native systems: A comparative and historical perspective. Turkish Online Journal of Qualitative Inquiry, 11(4), 2954–2963. https://doi.org/10.69980/j1064z97

Mahadasu, S. S. (2024). AI-native wireless digital twins for precision agriculture: Edge intelligence and predictive farm operations in 6G ecosystems. International Journal of Data Science and IoT Management System, 3(1), 65–72. https://doi.org/10.64751/ijdim.2024.v3.n1.1287

Bhaskara Raju Rallabandi. (2023). CLOSED-LOOP DIGITAL-TWIN MANAGEMENT FOR HYBRID 5G-NTN NETWORKS. American Journal of Management and IOT Medical Computing, 2(3), 30–35. https://doi.org/10.64751/ajmimc.2023.v2.n3.pp30-35

Anchala, S. (2020). Cloud security challenges in supply chain systems: Addressing vendor risks, data integrity, and third-party service provider vulnerabilities. Journal of Critical Reviews, 7(2). https://doi.org/10.69980/jcr.v7:i2.13401

Bhaskara Raju Rallabandi. (2024). DESIGNING NTN–TERRESTRIAL 6G EXPERIMENTS FOR AGRICULTURE AND SMART GRIDS. American Journal of Management and IOT Medical Computing, 3(3), 27–32. https://doi.org/10.64751/ajmimc.2024.v3.n3.pp27-32

Ahuja, R. (2020). Adoption of secure API development practices in microservice architectures for preventing unauthorized access and data leakage through OAuth2 and input throttling. Journal of Critical Reviews, 7(9), 4063–4069. https://doi.org/10.69980/jcr.v7:i9.13404

Bhaskara Raju Rallabandi. (2024). 6G Cross-Layer Intent Engine: Automated Translation from Business Requirements to Network Policies and Real-Time Scheduling Controls. International Journal of Data Science and IoT Management System, 3(1), 44–48. https://doi.org/10.64751/ijdim.2024.v3.n1.pp44-48

Khan, S. (2019). Sales force security compliance: An in-depth study of GDPR, HIPAA, and PCI-DSS enforcement in cloud-based CRM systems and their implications for global enterprises. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 8(9), 2211–2219. https://doi.org/10.15662/IJAREEIE.2019.0809010

Anchala, S. (2024). Deployment of zero trust security models in hybrid cloud infrastructures for minimizing lateral movement and strengthening access control through continuous verification mechanisms. International Journal of Innovative Research in Computer and Communication Engineering, 12(7), 9993–10002. https://doi.org/10.15680/IJIRCCE.2024.1207114

Rao, S. Madhusudhana, and Arpit Jain. "Advances in Malware Analysis and Detection in Cloud Computing Environments: A Review." International Journal of Safety & Security Engineering 14, no. 1 (2024).

Ahuja, R. (2019). A comparative study of traditional DevOps and DevSecOps: Examining the cultural, technical, and organizational shifts required to embed security into software engineering practices. International Journal of Innovative Research in Science, Engineering and Technology, 8(8), 8561–8569. https://doi.org/10.15680/IJIRSET.2019.0808096

Mahajan, R. A., Dey, R., Khan, M., Kanabar, K., & Gupta, S. K. (2025). Real-time vehicular-to-vehicular (V2V) communication and applications for pedestrian safety in AI-optimized VANETs for autonomous vehicles. In H. Maryam, M. M. Malik, I. U. Khan, & S. K. Gupta (Eds.), AI-driven transportation systems: Real-time applications and related technologies (pp. 151–167). Springer. https://doi.org/10.1007/978-3-031-98349-8_8

Gupta, Namit, Kunwar Singh Vaisla, Arpit Jain, Adesh Kumar, and Rajeev Kumar. "Performance Analysis of AODV Routing for Wireless Sensor Network in FPGA Hardware." Computer Systems Science & Engineering 40, no. 3 (2022).

Anchala, S. (2019). The evolution of DevSecOps in software supply chain security: Addressing dependencies, open-source vulnerabilities, and the prevention of software supply chain attacks. International Journal of Multidisciplinary and Scientific Emerging Research, 7(1), 415–422. https://doi.org/10.15662/IJMSERH.2019.0701047

Jain, Jaishree, Arpit Jain, Saurabh Kumar Srivastava, Chaman Verma, Maria Simona Raboaca, and Zoltán Illés. "Improved security of e-healthcare images using hybridized robust zero-watermarking and hyper-chaotic system along with RSA." Mathematics 10, no. 7 (2022): 1071.

Devulapalli, N. (2018). Database security for financial systems: Protecting customer data, ensuring PCI-DSS compliance, and safeguarding against fraud and insider attacks. International Journal of Advanced Research in Education and Technology, 5(5), 838–844. https://doi.org/10.15680/IJARETY.2018.0505018

Aggarwal, A. (2025). Artificial intelligence in criminal justice systems: An examination of predictive policing, automated surveillance, and ethical concerns in legal decision-making. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 8(1). https://doi.org/10.15680/IJMRSET.2025.0801070

Kumar, S., Jain, A., & Moparthi, N. R. (2022). Enhanced method of object tracing using extended Kalman filter via binary search algorithm.

Mahadasu, S. S. (2022). Cross-domain service assurance for private 5G networks using streaming telemetry and dependency-aware event correlation. International Journal of Data Science and IoT Management System, 1(3), 51–56. https://doi.org/10.64751/ijdim.2022.v1.n3.1290

Jain, Arpit, Adesh Kumar, and Sanjeev Sharma. "Comparative design and analysis of mesh, torus and ring NoC." Procedia Computer Science 48 (2015): 330–337.

Devulapalli, N. (2023). Database security and big data analytics: Protecting large-scale data warehouses, ensuring query privacy, and mitigating emerging threats in real-time analytics systems. International Journal on Recent and Innovation Trends in Computing and Communication, 11(7), 622–629. https://www.ijritcc.org/index.php/ijritcc/article/view/11939

Jain, A. (2024, December). Energy-Efficient Signal Processing Architectures for IoT Networks Design. In 2024 International Conference on Emerging Technologies and Innovation for Sustainability (EmergIN) (pp. 362–367). IEEE.

Mahadasu, S. S. (2024). Closed-loop sustainability orchestration for energy-aware RAN, edge, and data-center workloads in 5G-Advanced networks. International Journal of Engineering Research and Science & Technology, 20(2), 1392–1398. https://doi.org/10.62643/ijerst.2024.v20.n2.4407

Khan, S. (2018). The role of zero-trust models in database security: Eliminating implicit trust and enforcing continuous verification in enterprise data access systems. International Journal of Research in Electronics and Computer Engineering, 6(1), 1622–1628.

Jain, Arpit, Nageswara Rao Moparthi, A. Swathi, Yogesh Kumar Sharma, Nitin Mittal, Ahmed Alhussen, Zamil S. Alzamil, and MohdAnul Haq. "Deep Learning-Based Mask Identification System Using ResNet Transfer Learning Architecture." Computer Systems Science & Engineering 48, no. 2 (2024): 341.

Bhaskara Raju Rallabandi. (2023). AGENTIC-AI xAPPS: O-RAN MULTI-AGENT CONTROL FOR ENTERPRISE POLICY, SCHEDULING & SPECTRUM. International Journal of Data Science and IoT Management System, 2(4), 54–58. https://doi.org/10.64751/ijdim.2023.v2.n4.pp54-58