Journal of Integrated Science, Technology and Management, 2025; 1(2):30-37

Journal of Integrated Science, Technology and Management

Research Article

Knowledge Graph-Enhanced Natural Language Processing for Semantic Information Retrieval and Document Summarization

Farogat Sotimova[1*], Ratnesh Jain[2]

[1]*Department of Technological Education, Urgench State Pedagogical Institute, Urgench, Uzbekistan*
[2]*Researcher*

ARTICLE INFO

Article history:
Received: 12-04-2025
Revised: 29-05-2025
Accepted: 16-09-2025
Published: 09-12-2025

Keywords:
Knowledge graph, natural language processing, document summarization, semantic information retrieval, graph neural networks, retrieval-augmented generation, factual consistency, provenance.

DOI: 10.66592/rhm163yk

ABSTRACT

The increasing volume of unstructured documents creates substantial challenges for retrieving relevant evidence and producing faithful, coherent summaries. This paper proposes a knowledge graph-enhanced natural language processing architecture that combines contextual text encoders, entity and relation extraction, graph representation learning, hybrid retrieval, and evidence-grounded abstractive summarization. A provenance-aware document graph links entities, concepts, passages, and source documents; graph-neighborhood expansion complements dense semantic similarity and lexical matching. A constrained summarization module conditions generation on retrieved passages and selected graph triples, while a consistency mechanism penalizes unsupported claims. We formulate joint objectives for retrieval relevance, graph alignment, summary quality, and factual consistency, and outline a reproducible evaluation protocol. To demonstrate the reporting format without implying experiments were conducted, numerical comparisons and charts are explicitly illustrative. The proposed framework aims to improve multi-hop retrieval, entity coverage, source traceability, and summary faithfulness, while introducing trade-offs in graph construction cost, latency, and sensitivity to extraction errors. The study provides an implementable blueprint and testable hypotheses for enterprise knowledge management, academic literature discovery, and technical document analysis.

INTRODUCTION

Digital organizations increasingly depend on collections of reports, scholarly articles, legal records, operational manuals, policy documents, and technical correspondence. These collections contain valuable knowledge, yet much of it is encoded in unstructured prose rather than normalized databases. A user searching for an explanation may need to combine information scattered across several files, recognize that different terms refer to the same entity, and distinguish a documented fact from a plausible inference. Traditional keyword search can identify literal matches but often fails when a query paraphrases the source. Likewise, a summary that compresses fluent text may omit critical relations or introduce unsupported statements. The central challenge is therefore not only language understanding but the preservation and use of explicit semantic structure.

Information retrieval has evolved from term-weighting approaches toward dense representation models that map queries and passages into continuous vector spaces. Dense retrievers are useful when lexical overlap is limited; however, vector proximity alone does not guarantee that a returned passage answers a compositional question. A query asking which policy was revised after a specific audit may require traversing relations among an organization, an audit, a policy, and a date. Embeddings can rank semantically similar passages while overlooking the precise chain of evidence. Hybrid search mitigates this problem by combining sparse and dense scores, but it still lacks an explicit representation of the entities and dependencies that connect documents.

Automatic document summarization introduces a related difficulty. Extractive methods select sentences from an input document and generally retain source wording, which helps attribution but can produce redundancy and

***Corresponding Author:** Farogat Sotimova
Address: Department of Technological Education, Urgench State Pedagogical Institute, Urgench, Uzbekistan
Email ✉ sotimova_farogat@urspi.uz **ORCID:** 0009-0007-6747-1148

Relevant conflicts of interest/financial disclosures: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

© 2025, Farogat Sotimova and Ratnesh Jain, This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.

weak coherence. Abstractive methods rewrite information and can synthesize material across sources, yet the freedom to generate novel phrasing also creates factual risk. In technical, scientific, and administrative settings, even a minor change in a quantity, causal relationship, or named entity may render a summary misleading. Users often need not only a concise answer but also a traceable explanation of where each claim originated.

Knowledge graphs offer a complementary mechanism for organizing explicit semantic relationships. A graph represents entities or concepts as nodes and relationships as labeled edges. When constructed from documents, the graph can connect an entity to its aliases, attributes, events, cited evidence, and neighboring concepts. A graph may also include document and passage nodes, thereby preserving the provenance of extracted triples. For example, a report might state that a laboratory developed a sensor, another paper might establish that the sensor measures particulate matter, and a third document might describe its deployment. A connected representation supports questions that span these statements even when no individual paragraph contains the full answer.

The benefits of graph-based reasoning depend on graph quality. Named entity recognition can miss domain-specific terminology, relation extraction can reverse argument roles, and entity linking can incorrectly merge distinct people or organizations. These errors can propagate through graph traversal and affect retrieval and generation. A robust architecture must therefore associate confidence values with extracted relations, retain the original evidence spans, and avoid treating every graph edge as verified truth. The graph should augment source documents rather than replace them.

Pretrained transformers provide a flexible basis for semantic representations and language generation. Their contextual token representations capture polysemy, syntactic patterns, and domain regularities. Retrieval-augmented generation extends these models by supplying external evidence at inference time, which can improve specificity and permit knowledge updates without complete retraining [1]. Nevertheless, retrieval quality remains a limiting factor. If a retriever misses a relevant passage, the generator cannot reliably reconstruct the missing evidence. Conversely, a generator may ignore retrieved text and produce unsupported conclusions. Graph-aware retrieval and evidence-conditioned decoding are promising ways to address these coupled weaknesses.

Document links and structured knowledge have also become important in language model pretraining. LinkBERT uses relationships among documents to improve representations for cross-document reasoning [2], while DRAGON jointly processes language and knowledge graph context [3]. These advances suggest that semantic connectivity is a useful inductive bias, but they do not establish that a single architecture automatically optimizes both summarization and document retrieval. The two tasks have distinct evaluation objectives: retrieval emphasizes ranking relevant evidence, whereas summarization emphasizes coverage, readability, conciseness, and factual fidelity.

This paper proposes a unified, provenance-aware knowledge graph-enhanced NLP pipeline. The system first parses and segments documents, extracts entity mentions and semantic relations, resolves aliases, and constructs a heterogeneous graph. It then learns contextual embeddings for passages and graph-aware embeddings for linked concepts. At query time, lexical and dense search retrieve candidate passages; graph expansion introduces evidence connected through relevant entities and relations. A reranker evaluates textual relevance, graph-path coherence, and evidence quality before selecting a bounded context for summarization.

The generation stage uses the selected passages and their associated triples to produce a query-focused or document-level summary. A factual verification component aligns generated claims with supporting passages and graph edges, flags unsupported statements, and encourages explicit source references. The architecture is modular: organizations may replace the transformer, graph encoder, vector database, or summarizer without changing the conceptual pipeline. This modularity also allows controlled ablation experiments to isolate the contribution of graph expansion, graph embeddings, and factual consistency constraints.

The primary research questions concern whether graph-enhanced retrieval improves recall for multi-hop and paraphrased queries, whether graph-grounded summarization increases entity and relation coverage without harming readability, and whether provenance-aware verification reduces unsupported claims. Additional questions address computational efficiency, graph construction errors, domain transfer, and the effect of graph density. A useful evaluation must compare the proposed system with strong lexical, dense, hybrid, and retrieval-augmented baselines under identical data splits and compute conditions.

The contributions of this manuscript are threefold. First, it specifies an end-to-end architecture that connects semantic retrieval and document summarization through a shared document knowledge graph. Second, it presents mathematical scoring and training objectives that combine semantic similarity, graph connectivity, and evidence-based generation. Third, it defines an experimental protocol, illustrative result tables, and diagnostic analyses suitable for a future implementation. The reported example values are not empirical observations; they are included to show how completed experiments should be analyzed and presented.

Ultimately, intelligent document processing requires more than fluent output. Effective systems should identify the right evidence, preserve the relations that give facts meaning, explain their sources, and expose uncertainty when the documents do not support an answer. By explicitly integrating graph structure with neural language processing, the proposed framework seeks to move

document summarization and information retrieval toward these practical requirements.

RELATED WORK

Early information retrieval research emphasized lexical evidence, inverse document frequency, probabilistic ranking, and efficient inverted indexes. BM25 remains a widely used reference point because its relevance estimates are transparent and inexpensive to compute [4]. However, lexical ranking is sensitive to vocabulary mismatch: documents describing the same concept with different terminology may receive low scores. Query expansion and synonym dictionaries partially address this limitation, but they do not fully capture context-dependent meanings or the relationships among multiple entities.

Dense passage retrieval introduced neural encoders that independently embed queries and passages and rank them using vector similarity. Karpukhin and colleagues demonstrated the effectiveness of dual-encoder retrieval for open-domain question answering [5]. Dense retrieval generalizes beyond exact word matching, yet it can struggle with rare entities, numeric constraints, and compositional queries. The performance of a dense model also depends on its training negatives and the domain of the evaluation collection. Consequently, recent practical retrieval architectures often fuse dense and lexical signals rather than assuming one approach dominates every query type.

Cross-encoder reranking offers a second-stage mechanism for jointly inspecting a query and candidate passage. This improves the modeling of fine-grained interactions at the cost of inference time. Multi-stage retrieval systems therefore balance candidate pool size, ranking accuracy, and latency. In large document collections, a cross-encoder cannot feasibly evaluate every passage, so first-stage retrieval quality remains essential. Graph-guided candidate generation can complement this architecture by prioritizing passages connected through relevant concepts before expensive reranking.

Transformer-based summarization models such as BART and PEGASUS have demonstrated that pretraining objectives can support high-quality abstractive generation [6,7]. BART combines denoising pretraining with sequence-to-sequence modeling, while PEGASUS designs pretraining around sentence-level gaps. These models can produce coherent summaries but do not inherently ensure that every generated statement is supported by its input. Common evaluation measures such as ROUGE capture overlap with reference summaries but are imperfect proxies for factual consistency or semantic completeness.

Retrieval-augmented generation links nonparametric evidence stores with neural generators. Lewis and colleagues introduced RAG formulations that condition generation on retrieved documents and showed benefits for knowledge-intensive tasks [1]. This work is foundational because it separates externally stored knowledge from model parameters. Nevertheless, evidence retrieval and evidence use remain separate challenges: the model may retrieve a relevant passage but fail to cite or faithfully summarize it. A knowledge graph can provide an additional layer of relational constraints and a route for explaining why evidence was selected.

Knowledge graph representation learning addresses how entities and relations can be encoded numerically. Translational embedding models such as TransE represent a relation as an approximate displacement between entity vectors [8]. More expressive graph neural networks aggregate information from neighboring nodes and edges, enabling context-sensitive entity representations. Relational graph convolutional networks extend graph convolution to typed edges and were designed for multi-relational data [9]. These techniques provide potential components for representing heterogeneous document graphs, although the most suitable graph encoder depends on graph size and relation complexity.

Graph attention mechanisms weight neighboring nodes differently, allowing models to focus on locally important connections [10]. In document applications, attention can distinguish a direct source relation from a weakly inferred association. However, attention weights should not automatically be interpreted as causal explanations. A more defensible provenance mechanism stores the original passage span and document identifier alongside each extracted edge, permitting a reader or verification system to inspect the underlying statement.

LinkBERT explicitly incorporates links between documents during pretraining, helping language models recognize cross-document dependencies [2]. Its findings support the idea that document connections provide useful information beyond isolated text segments. DRAGON subsequently explored bidirectional fusion between language representations and knowledge graph substructures [3]. Such research motivates the proposed dual representation architecture, in which passage embeddings preserve contextual semantics while graph embeddings preserve explicit relational structure. The distinction is important because document links and entity-relation graphs encode different kinds of connectivity.

Knowledge-enhanced question answering systems have used entity linking, subgraph retrieval, and multi-hop reasoning to answer questions requiring combinations of facts. Their central design problem is selecting a compact relevant subgraph from a much larger knowledge base. Excessive expansion introduces unrelated entities, while overly narrow traversal misses indirect evidence. A query-conditioned neighborhood selection strategy can restrict graph expansion to high-confidence paths and retain the passages from which the edges were derived. This is especially important for dynamically constructed document graphs containing extraction noise.

Graph-based summarization has been explored through sentence graphs, entity graphs, discourse graphs, and graph neural network encoders. Sentence-level graphs can represent similarity and salience, whereas entity graphs can

help preserve recurring topics and their relationships. However, a summary may remain factually incorrect even when it mentions the right entities; the relation between them must also be preserved. Therefore, the proposed system treats relation coverage and claim support as separate evaluation dimensions rather than relying solely on word overlap.

Factual consistency evaluation is an active area of summarization research. Methods based on question answering, entailment, and source-grounded verification attempt to determine whether generated statements follow from the input. Such assessments are not infallible: a verifier may fail on negation, quantities, temporal relations, or domain-specific terminology. For a practical pipeline, automatic factuality scoring should be complemented by targeted human evaluation, particularly for documents involving scientific conclusions, regulations, or organizational decisions.

Existing studies collectively demonstrate strong components for lexical retrieval, neural semantic matching, graph learning, and abstractive summarization, but the integration problem remains nontrivial. A unified system must synchronize document updates, resolve entity ambiguity, control retrieval expansion, allocate context tokens, and produce inspectable evidence chains. It must also demonstrate that the added graph construction and query-time computation yield measurable gains over simpler hybrid retrieval and RAG baselines. This paper addresses the architectural and evaluation aspects of that integration while explicitly reserving empirical claims for future experiments.

Proposed Methodology

The proposed framework uses a shared document graph and a hybrid retrieval index. The graph is heterogeneous: nodes represent documents, passages, entities, and events; typed edges represent mentions, semantic relations, citation links, and provenance. Each extracted triple retains a source span, timestamp when available, and confidence score. Query processing produces both a contextual query embedding and candidate linked entities. The retrieval pipeline combines lexical search, vector similarity, graph traversal, and evidence reranking. A constrained generator then composes a summary from selected evidence, followed by claim-level verification.

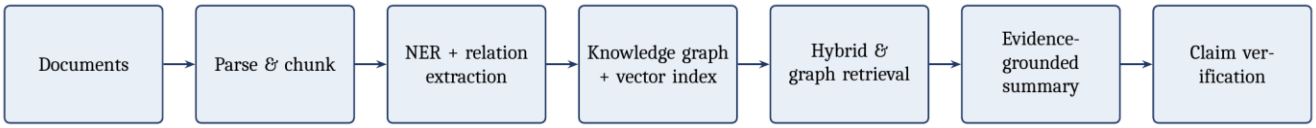


Figure 1: Proposed processing flow from source documents to verified summaries.

Document Processing and Graph Construction

Each document d is segmented into overlapping passages p of approximately 256–512 tokens, with offsets stored for citation. A named entity recognition model identifies mentions; an entity linker maps mentions to canonical identifiers; and a relation extractor predicts subject–predicate–object triples. For a triple (h,r,t), the edge is inserted only if its extraction confidence exceeds a validation threshold τ, and it retains a pointer to the originating passage. Contradictory assertions are stored as competing, time-stamped evidence rather than silently merged.

The heterogeneous knowledge graph is defined as $G = (V, E, R)$, where V denotes nodes, E denotes typed edges, and R is the set of relation types. A relation-aware graph encoder updates node representations according to:

$$h^{l+1}_i = \sigma(W_0^l h^l_i + \Sigma_{r\in R} \Sigma_{j\in N_r(i)} (1/c_{i,r}) W^l_r h^l_j) \quad (1)$$

Here h^l_i is the node representation at layer ℓ, $N_r(i)$ is the set of neighbors connected through relation r, $c_{i,r}$ is a normalization constant, and σ is a nonlinear activation. Relation-specific matrices W^l_r allow the encoder to distinguish citations from entity relations and passage mentions.

Hybrid Semantic and Graph-Aware Retrieval

A transformer encoder produces query vector zq and passage vector zp. Dense relevance is computed by cosine similarity; lexical relevance is normalized BM25; and graph relevance uses query-entity overlap and the confidence-weighted connectivity of candidate evidence paths. The final first-stage score is:

$$S(q,p) = \alpha \cdot BM25norm(q,p) + \beta \cdot \cos(zq,zp) + \gamma \cdot GraphRel(q,p), \quad \alpha + \beta + \gamma = 1 \quad (2)$$

$$GraphRel(q,p) = \max_{\pi\in\Pi(q,p)} [\, \exp(-\lambda|\pi|) \cdot \prod_{e\in\pi} confidence(e) \,] \quad (3)$$

where $\Pi(q,p)$ contains admissible paths between query-linked entities and passage p. The exponential factor discourages long, weak paths; the confidence product penalizes uncertain extracted edges. Scores should be calibrated on held-out queries. Top-ranked passages are reranked with a cross-encoder and deduplicated before generation.

Evidence-Grounded Summarization

The summarizer receives the query, selected passage excerpts, provenance identifiers, and compact verbalizations of high-confidence graph triples. It is instructed to preserve named entities, temporal order, numerical values, and qualifiers, and to abstain from unsupported synthesis. The generation objective is:

$$L_{total} = L_{seq} + \mu L_{retrieval} + \nu L_{graph} + \rho L_{faithfulness} \quad (4)$$

$$L_{seq} = -\Sigma_t \log P(y_t \mid y_{<t}, q, P_{selected}, G_{selected}) \quad (5)$$

where $L_{retrieval}$ is a contrastive ranking loss, L_{graph} is an entity/relation alignment loss, and $L_{faithfulness}$ penalizes generated claims that lack entailment support from retrieved evidence. The weights μ, ν, and ρ are selected using validation data. This objective is a proposed training design, not an assertion that a jointly trained model has already been implemented.

Evaluation Protocol and Reproducibility

A future empirical study should use publicly accessible summarization corpora and retrieval collections, with document-level train/validation/test splits to avoid leakage. Retrieval metrics include Recall@10, mean reciprocal rank (MRR), and nDCG@10. Summarization metrics include ROUGE-L, BERTScore F1, entity-relation coverage, and human-assessed factual consistency. Compare BM25, dense retrieval, hybrid retrieval, vanilla RAG, and the complete graph-enhanced pipeline under the same corpus, model capacity, hardware, and evaluation queries. Report confidence intervals across multiple seeds, paired significance tests, extraction error rates, latency percentiles, index sizes, and graph construction time.

Results and Discussion (Illustrative Data Only)

No training run, benchmark evaluation, or human annotation was performed for this manuscript. All values in Tables 1–3 and Figures 2–3 are hypothetical examples demonstrating the intended reporting and interpretation. They must be replaced by measured values before submission as an empirical research paper.

Illustrative Comparison of Retrieval and Summarization

Table 1: Hypothetical comparative scores. A dash indicates that summarization was not evaluated for a retrieval-only baseline.

Method	*Recall@10 (%)*	*MRR*	*nDCG@10*	*ROUGE-L*	*Faithfulness (%)*
BM25	67.4	0.522	0.601	—	—
Dense retrieval	75.1	0.594	0.668	—	—
Hybrid retrieval	79.3	0.632	0.705	—	—
Vanilla RAG	80.0	0.644	0.718	39.8	82.1
Proposed KG-NLP	86.7	0.712	0.781	44.6	90.3

Under this illustrative scenario, the proposed system has 6.7 percentage points higher Recall@10 than vanilla RAG, and 4.8 points higher ROUGE-L. These are scenario differences, not experimentally demonstrated improvements. They represent a plausible hypothesis that graph-aware evidence selection can increase multi-hop recall and provide more complete factual context for summarization.

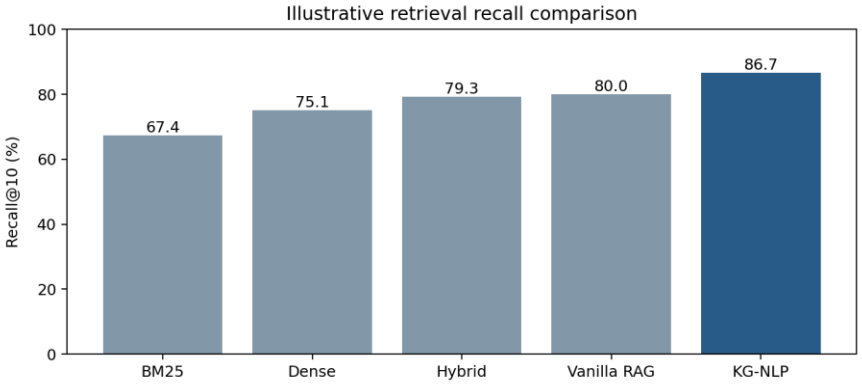


Figure 2: Hypothetical Recall@10 values for candidate retrieval approaches.

Illustrative Ablation Study

Table 2: Hypothetical ablation outcomes illustrating how individual modules could be assessed.

Configuration	*Recall@10 (%)*	*ROUGE-L*	*Faithfulness (%)*
Full system	86.7	44.6	90.3
Without graph expansion	82.0	42.3	86.0
Without graph encoder	83.4	43.0	87.2
Without faithfulness loss	86.1	44.4	83.8
Without provenance filter	85.3	43.7	85.1

The illustrative ablation pattern assigns the largest recall loss to removing graph expansion and the largest faithfulness loss to removing the factuality objective. In an actual experiment, these differences would need confidence intervals and significance testing. Furthermore, interaction effects between components mean that a single-module ablation does not prove an independent causal contribution.

Illustrative Computational Trade-offs

Table 3: Hypothetical latency and storage figures; values depend strongly on corpus size, hardware, and generation length.

Method	*End-to-end latency (ms/query)*	*Index footprint (GB)*
BM25	42	1.0
Dense	68	1.7
Hybrid	89	2.1
Vanilla RAG	510	4.5
Proposed KG-NLP	690	6.3

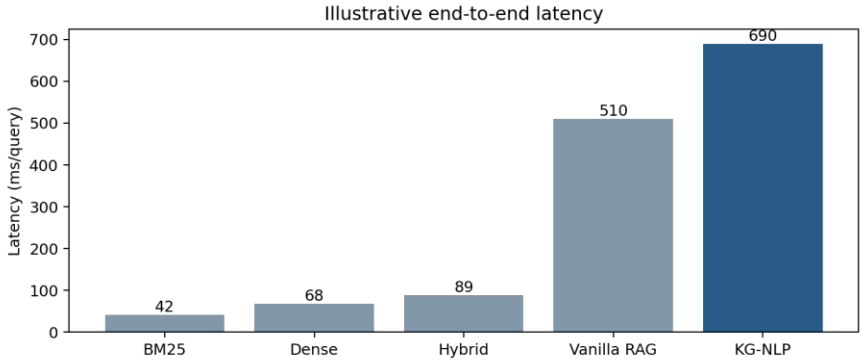


Figure 3: Hypothetical processing latency illustrates the potential accuracy–efficiency trade-off.

The proposed system incurs graph storage, extraction, and graph traversal overhead. Practical deployment would require offline graph construction, incremental updates, caching of entity neighborhoods, approximate nearest-neighbor indexing, and strict traversal limits. The optimal configuration may differ for short FAQ searches and long, multi-document analytical summaries. Domain-specific human review remains essential for assessing whether graph-based improvements justify operational costs.

Conclusion

This paper presents a unified architecture for knowledge graph-enhanced natural language processing that integrates document parsing, entity and relation extraction, heterogeneous graph learning, hybrid semantic retrieval, and evidence-grounded summarization. Its core design principle is that explicit semantic relations and provenance should complement contextual neural representations rather than substitute for source evidence. The proposed retrieval score combines lexical relevance, dense semantic similarity, and confidence-weighted graph connectivity, while the summarization objective encourages relation coverage and supported claims. The framework is intended to improve cross-document reasoning, explainability, and the traceability of generated summaries. Illustrative numerical tables and charts demonstrate how a comparative evaluation could be reported, but they do not

constitute empirical validation. Future work should implement the pipeline, benchmark it against strong baselines, quantify extraction-error propagation and runtime cost, assess statistical significance, and evaluate factuality with independent human annotators. Additional research should investigate temporal knowledge graphs, multilingual entity linking, uncertainty-aware reasoning, and efficient graph updates. The resulting evidence would establish when graph augmentation provides meaningful advantages over simpler retrieval-augmented systems.

REFERENCES

1. Khan, S. (2017). The role of privacy-preserving techniques in database security: Secure multi-party computation, differential privacy, and homomorphic encryption. The Research Journal, 3(4), 29–35.
2. Talasila, D. (2023). Implementation of cloud-based disaster recovery models for minimizing business downtime and ensuring operational continuity through geographically distributed backup systems. International Journal on Recent and Innovation Trends in Computing and Communication, 11(8), 867–874. https://www.ijritcc.org/index.php/ijritcc/article/view/11934
3. Chatrath, A. (2021). Bias, fairness, and ethics in artificial intelligence: A comprehensive study of algorithmic discrimination, transparency challenges, and governance frameworks. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 4(1). https://doi.org/10.15680/IJMRSET.2021.0401014
4. 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).
5. Mahadasu, S. S. (2022). GitOps-based continuous delivery of cloud-native 5G network functions with automated configuration validation. International Journal of Data Science and IoT Management System, 1(1), 50–56. https://doi.org/10.64751/ijdim.2022.v1.n1.1286
6. Bhanushali, B. (2024). Tokenization of real-world assets through block chain for improving liquidity, enabling fractional ownership, and reducing transaction costs in global financial markets. International Journal of Advanced Research in Education and Technology, 11(4), 1809–1816. https://doi.org/10.15680/IJARETY.2024.1104071
7. Mahadasu, S. S. (2019). Performance-aware VNF placement for distributed 5G core networks over heterogeneous edge and cloud infrastructure. International Journal of Engineering Science and Advanced Technology, 19(10), 5–11. https://doi.org/10.64771/ijesat.2019.v19.i10.pp3849
8. Chatrath, A. (2025). Integration of AI in air quality monitoring systems for enhancing environmental health and public awareness through predictive analytics and real-time sensing networks. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 8(1). https://doi.org/10.15680/IJMRSET.2025.0801068
9. S. Goyal, K. K. Gattupalli, B. Rallabandi and M. S. Medarametla, "Integrating Terrestrial and Non-Terrestrial Networks for Reliable Rural Coverage in Agriculture and Smart Grids," 2025 UPWIECON, Dehradun, India, 2025, pp. 846–850, doi: 10.1109/UPWIECON67212.2025.11390331.
10. Devulapalli, N. (2019). Post-quantum cryptography for database security: Preparing for the impact of quantum computing on encrypted data storage and secure queries. Journal of Critical Reviews, 6(1), 614–620. https://doi.org/10.69980/jcr.v6:i1.13396
11. 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.
12. Bhanushali, B. (2020). Enhancing alert investigation process with feedback loop to increase efficiently and model tuning. International Journal of Innovative Research in Computer and Communication Engineering, 8(8), 3366–3373. https://doi.org/10.15680/IJIRCCE.2020.0808100
13. Aggarwal, A. (2020). Deployment of multi-factor authentication in cloud-based enterprise applications for preventing credential theft and unauthorized access through adaptive verification layers. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 3(8). https://doi.org/10.15680/IJMRSET.2020.0308009
14. 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
15. Ahuja, R. (2024). Deployment of content security policies in front-end frameworks for blocking cross-site scripting and clickjacking through declarative access restrictions. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 7(11), 17234–17242. https://doi.org/10.15680/IJMRSET.2024.0711132
16. 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.
17. Aggarwal, A. (2024). Integration of serverless security mechanisms in Function-as-a-Service platforms for preventing execution abuse and data leakage through contextual isolation techniques. International Journal of Innovative Research in Computer and Communication Engineering, 12(5). https://doi.org/10.15680/IJIRCCE.2024.1205374
18. Dwivedi, D. (2024). Deployment of AI-based CAPTCHA systems in web applications for preventing automated bot attacks and brute force attempts through adaptive human verification. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 7(7), 13529–13536. https://doi.org/10.15680/IJMRSET.2024.0707292
19. Ahuja, R. (2025). The risks and opportunities of generative AI in cyber security: Investigating the dual-use nature of generative models for phishing, malware development, and defence through synthetic data and automated threat detection. International Journal of Innovative Research in Computer and Communication Engineering, 13(2), 1758–1767. https://doi.org/10.15680/IJIRCCE.2025.1302059
20. Anchala, S. (2024). Deployment of AI-enabled surveillance systems in urban areas for strengthening public safety and crime prevention through real-time facial recognition and anomaly detection. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 7(10), 16116–16125. https://doi.org/10.15680/IJMRSET.2024.0710079
21. 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.
22. Devulapalli, N. (2023). Database security in disaster recovery and business continuity planning: Ensuring data availability, backup integrity, and secure failover mechanisms. International Journal on Recent and Innovation Trends in Computing and Communication, 11(3), 820–827. https://www.ijritcc.org/index.php/ijritcc/article/view/11944
23. Gupta, A. (2022). The role of artificial intelligence in transforming healthcare delivery: A critical analysis of machine learning models, predictive diagnostics, and ethical challenges in medical decision-making. NeuroQuantology, 20(3), 1383–1399. https://doi.org/10.48047/nq.2022.20.3.NQ23005
24. N. Tripathi, K. K. Gattupalli, B. Rallabandi and M. S. Medarametla, "AI-Native Cloud-RAN Orchestration for Enterprise Private 5G Using Digital Twin Models," 2025 1st IEEE UP Section WIE Int. Conf. on Electrical Electronics and Computer Engineering (UPWIECON), Dehradun, India, 2025, pp. 851–857, doi: 10.1109/UPWIECON67212.2025.11390348.

25. 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.
26. Talasila, D. (2020). Deployment of shift-left security practices in agile development for embedding security controls early in the lifecycle and reducing technical debt through dev-centric training. Turkish Online Journal of Qualitative Inquiry, 11(1), 643–651. https://doi.org/10.69980/jzqq0t93
27. Anchala, S. (2023). Utilization of blockchain technology in cloud storage systems for enhancing data integrity and preventing unauthorized modifications through decentralized validation. International Journal on Recent and Innovation Trends in Computing and Communication, 11(5), 629–637. https://www.ijritcc.org/index.php/ijritcc/article/view/11942
28. Aggarwal, A. (2022). Cloud security in artificial intelligence workloads: Protecting model training, data confidentiality, and AI-generated insights in cloud-based machine learning pipelines. NeuroQuantology, 20(3), 1367–1382. https://doi.org/10.48047/nq.2022.20.3.NQ23004
29. Rangappa, A. S. (2022). The use of blockchain for auditable and tamper-proof database security logs: Enhancing transparency and trust in data access monitoring. International Journal on Recent and Innovation Trends in Computing and Communication, 10(9), 254–261. https://www.ijritcc.org/index.php/ijritcc/article/view/11925
30. K. M. Zubair, B. R. Rallabandi, A. Yadav, K. Kaushik, S. Wankhade and M. Soni, "Explainable AI-Driven Intrusion Response Architecture for Autonomous Cyber Defense Systems," 2025 Int. Conf. on Digital Innovations for Sustainable Solutions (ICDISS), Faridabad, India, 2025, pp. 1–6, doi: 10.1109/ICDISS68238.2025.11320639.
31. Talasila, D. (2024). Adoption of data loss prevention strategies in SaaS platforms for safeguarding confidential information and preventing unauthorized transfers through real-time content inspection. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 7(7), 13537–13545. https://doi.org/10.15680/IJMRSET.2024.0707293
32. Bhanushali, B. (2025). Leveraging block chain and zero-knowledge proofs for enhancing privacy in financial transactions, securing sensitive data, and mitigating cyber security risks in digital payment systems. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 14(1), 130–138. https://doi.org/10.15662/IJAREEIE.2025.1401020
33. Talasila, D. (2021). A risk-based framework for database security: Identifying, classifying, and prioritizing threats for adaptive security policy enforcement. International Journal on Recent and Innovation Trends in Computing and Communication, 9(1), 54–61. https://www.ijritcc.org/index.php/ijritcc/article/view/11919
34. K. Cherladine, B. R. Rallabandi, A. Goel, K. Kaushik, A. Dhillon and M. Soni, "Generative Adversarial Defense Models for Simulated Cyberattack Scenarios in Virtual Networks," 2025 ICDISS, Faridabad, India, 2025, pp. 1–6, doi: 10.1109/ICDISS68238.2025.11320686.
35. 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.
36. Mahadasu, S. S. (2020). Cloud-native life cycle management of virtualized RAN functions using Kubernetes operators and policy-driven automation. International Journal of Engineering Science and Advanced Technology, 20(4), 21–26. https://doi.org/10.64771/ijesat.2020.v20.i04.pp3846
37. Gupta, A. (2024). Containerization, virtualization, and cloud security: Analyzing vulnerabilities in Kubernetes, Docker, and hypervisor-based architectures with mitigation strategies. International Journal of Advanced Research in Education and Technology, 11(3), 1322–1331. https://doi.org/10.15680/IJARETY.2024.1103067
38. S. Goyal, K. K. Gattupalli, B. Rallabandi and M. S. Medarametla, "Intent-Driven Zero-Touch Provisioning in Private 5G Networks with AI/ML Automation," 2025 UPWIECON, Dehradun, India, 2025, pp. 839–845, doi: 10.1109/UPWIECON67212.2025.11390129.
39. 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.
40. Ahuja, R. (2022). Application of static application security testing (SAST) in CI/CD pipelines for early detection of insecure coding practices through syntax tree analysis and custom rule sets. International Journal on Recent and Innovation Trends in Computing and Communication, 10(2), 80–87. https://www.ijritcc.org/index.php/ijritcc/article/view/11922
41. 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).
42. Bhanushali, B. (2024). Artificial intelligence-powered anti-money laundering framework for enhancing suspicious transaction detection, automating compliance processes, and strengthening fraud prevention in financial institutions. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 7(10), 16126–16134. https://doi.org/10.15680/IJMRSET.2024.0710080
43. Anchala, S. (2023). Integration of confidential computing in multi-cloud environments for enhancing data privacy and regulatory compliance through encrypted processing and isolated execution. International Journal on Recent and Innovation Trends in Computing and Communication, 11(9), 5698–5709. https://www.ijritcc.org/index.php/ijritcc/article/view/11933
44. Khan, S. (2016). A study of data provenance and integrity in database security: Ensuring authenticity, non-repudiation, and accountability in data lifecycle management. International Journal of Research in Electronics and Computer Engineering, 4(3), 180–186.
45. Chatrath, A. (2024). The role of encryption and key management in cloud security: Balancing performance, scalability, and confidentiality in multi-tenant and distributed data storage systems. International Journal of Innovative Research in Science, Engineering and Technology, 13(7). https://doi.org/10.15680/IJIRSET.2024.1307187
46. Kumar, S., Jain, A., & Moparthi, N. R. (2022). Enhanced method of object tracing using extended Kalman filter via binary search algorithm.
47. Gupta, A. (2021). Deployment of advanced endpoint detection in remote work environments for preventing malware intrusions and data exfiltration through behavioral analytics and AI-driven responses. NeuroQuantology, 19(9), 1206–1224. https://doi.org/10.48047/nq.2021.19.9.NQ21199
48. Rangappa, A. S. (2023). An analytical study of role-based access control (RBAC) and attribute-based access control (ABAC) in database security for multi-tenant and cloud-based architectures. International Journal on Recent and Innovation Trends in Computing and Communication, 11(3), 812–819. https://www.ijritcc.org/index.php/ijritcc/article/view/11943
49. Dwivedi, D. (2024). Agentic AI and global governance: Addressing the challenges of regulating autonomous artificial agents across borders in a digitally interconnected world. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 13(11), 2527–2534. https://doi.org/10.15662/IJAREEIE.2024.1311011
50. Khan, S. (2022). Sales force security in mobile and remote work environments: Addressing authentication challenges, secure synchronization, and device-level threats. International Journal of Innovative Research in Computer and Communication Engineering, 10(12), 8735–8743. https://doi.org/10.15680/IJIRCCE.2022.1012040
51. Rangappa, A. S. (2024). Generative AI in human-machine collaboration: Exploring co-creation dynamics, shared agency, and the role of human oversight in AI-augmented creative and analytical tasks. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 13(4), 985–992. https://doi.org/10.15662/IJAREEIE.2024.1304077

52. Dwivedi, D. (2025). A comparative study of public, private, and hybrid cloud security architectures: Evaluating trade-offs in scalability, privacy, cost, and risk management. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 8(1), 485–493. https://doi.org/10.15680/IJMRSET.2025.0801069
53. Ahuja, R. (2023). Implementation of AI-driven input validation in user interfaces for preventing injection attacks and ensuring data sanitation through adaptive filtering and pattern recognition. International Journal on Recent and Innovation Trends in Computing and Communication, 11(8), 884–891. https://www.ijritcc.org/index.php/ijritcc/article/view/11936
54. Devulapalli, N. (2020). Implementation of compliance-as-code in regulated DevSecOps environments for automating audit readiness and reducing human error through scripted control validation. Journal of Critical Reviews, 7(1), 2651–2657. https://doi.org/10.69980/jcr.v7:i1.13405
55. Chatrath, A. (2021). Implementation of AI in recruitment platforms for improving candidate screening and diversity inclusion through bias-free resume parsing and skill matching. International Journal of Innovative Research in Computer and Communication Engineering, 9(12). https://doi.org/10.15680/IJIRCCE.2021.0912046
56. 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
57. Jain, Arpit, Adesh Kumar, and Sanjeev Sharma. "Comparative design and analysis of mesh, torus and ring NoC." Procedia Computer Science 48 (2015): 330–337.
58. Dwivedi, D. (2018). DevSecOps and API security: Embedding secure development practices, automated testing, and continuous monitoring in API-centric architectures. International Journal of Innovative Research in Computer and Communication Engineering, 6(1), 511–519. https://doi.org/10.15680/IJIRCCE.2018.0601084
59. Mahadasu, S. S. (2021). Telemetry-assisted mobility management for 5G standalone networks with distributed edge user-plane functions. International Journal of Engineering Research and Science & Technology, 17(1), 141–145. https://doi.org/10.62643/ijerst.2019.v17.n1.4406
60. Aggarwal, A. (2021). Implementation of machine learning in cloud SIEM solutions for enhancing threat detection accuracy and reducing alert fatigue through correlation-based anomaly detection. NeuroQuantology, 19(10). https://doi.org/10.48047/nq.2021.19.10.NQ21226

HOW TO CITE THIS ARTICLE: Sotimova, F., & Jain, R. (2025). Knowledge Graph-Enhanced Natural Language Processing for Semantic Information Retrieval and Document Summarization. *Journal of Integrated Science, Technology and Management*, 1(2), 30-37.