International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802390

AI TEACH Innovators: A Privacy-First Conversational Business Intelligence Platform using Hybrid RAG, Knowledge Graphs, and Agentic AI with Local LLM Inference

Dr. S. Selvarajan[1], Mr. B. Harish[2], Mr. Ganpat Kumar[3], Mr. K. Dinesh[4], Mr. S. Barath[5]

DEAN & Professor, Department of Computer Science and Engineering, Gnanamani College of Technology, Namakkal, Tamil Nadu, India[1]

UG Scholar, Department of Computer Science and Engineering, Gnanamani College of Technology, Namakkal, Tamil Nadu, India[2-5]

Publication History: Received: 25.02.2026; Revised: 20.03.2026; Accepted: 25.03. 2026; Published: 28.03.2026.

ABSTRACT: Modern enterprises face growing complexity in extracting actionable intelligence from distributed, heterogeneous data sources while simultaneously facing regulatory pressure to maintain data privacy and sovereignty. Traditional cloud-based AI solutions require transmitting sensitive business data to external servers, making them fundamentally incompatible with GDPR, HIPAA, and sector-specific compliance frameworks. This paper presents AI TEACH Innovators, a production-grade, fully offline conversational business intelligence (BI) platform that integrates Hybrid Retrieval-Augmented Generation (RAG), Knowledge Graphs, and a three-node Agentic AI pipeline to enable natural-language querying of enterprise data entirely on-device. The system leverages Ollama as a fully local Large Language Model (LLM) inference engine routed through a LiteLLM proxy, guaranteeing that no data leaves the user's premises at any stage. A five-layer guardrail engine enforces input safety, document scoping, semantic relevance filtering, Personally Identifiable Information (PII) redaction, and hallucination detection. Knowledge Graphs powered by Neo4j provide causal reasoning over Key Performance Indicator (KPI) relationships, while ChromaDB and sentence-transformer embeddings enable dense semantic retrieval. Evaluation results demonstrate sub-15-second end-to-end response times, 100% offline operation, and robust hallucination mitigation with an 87.5% detection rate in the self-critique layer.

KEYWORDS: Retrieval-Augmented Generation, Knowledge Graph, Agentic AI, LangGraph, Local LLM, Ollama, Conversational Business Intelligence, LLM Guardrails, Hallucination Mitigation, Privacy-Preserving AI, ChromaDB, Neo4j, LiteLLM

I. INTRODUCTION

The rapid proliferation of Large Language Models (LLMs) has catalysed a fundamental paradigm shift in how organizations interact with their data assets. Rather than relying on structured query languages, pre-built dashboard tools, or specialist data analysts, decision-makers increasingly expect to pose questions in plain natural language and receive grounded, contextually accurate, and actionable responses in near-real time. This expectation has given rise to a new class of enterprise software commonly referred to as Conversational Business Intelligence (CBI) platforms.

Despite the considerable promise of LLM-powered CBI, three persistent and interrelated challenges have prevented widespread production adoption. First, data privacy and sovereignty: the dominant deployment model for state-of-the-art LLMs relies on cloud-hosted APIs provided by third parties. Submitting financial reports, personnel records, customer databases, and confidential strategic plans to external inference endpoints is directly incompatible with GDPR, HIPAA, and numerous sector-specific regulations in banking and critical infrastructure. Second, factual reliability: standard autoregressive LLMs generate text by predicting statistically probable token sequences, systematically producing hallucinations that can lead to costly or legally consequential business decisions. Third,

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802390

domain specificity: general-purpose LLMs lack the organizational context and proprietary data required for trustworthy enterprise decision support.

Retrieval-Augmented Generation (RAG) has emerged as the principal evidence-based solution to the hallucination problem, grounding model outputs in retrieved, authoritative source documents. Knowledge Graphs (KGs) address the relational reasoning gap by representing causal and relational structures as directed graphs, enabling traversal-based reasoning that complements dense vector retrieval. This paper presents AI TEACH Innovators, a production-grade CBI platform unifying Hybrid RAG, KG-enriched retrieval, a three-node LangGraph agentic reasoning pipeline, and a five-layer guardrail engine, all executing entirely on-device without any cloud dependency.

II. RELATED WORK

A. Retrieval-Augmented Generation

Pan et al. [1] provide a foundational roadmap for unifying large language models and knowledge graphs, demonstrating that structured relational knowledge dramatically improves factual grounding of LLM outputs, identifying three integration paradigms and showing that combining graph-structured knowledge with generative models enables causal and multi-hop reasoning. Procko and Ochoa [2] survey Graph RAG architectures at IEEE CAI 2024, cataloguing how knowledge graph traversal serves as a retrieval backbone that reduces hallucination and improves answer traceability in open-domain QA systems.

B. Knowledge Graphs and Enterprise LLM Systems

Kumar et al. [3] demonstrate that LLM-powered knowledge graph frameworks can unify disconnected enterprise data silos into coherent activity-centric graphs, enabling advanced cross-silo analytics that standard single-stage RAG cannot support. Fayazi et al. [4] present SAGE-QA, an enterprise-level domain-specific Graph RAG system at IEEE BigData 2025, showing that routing queries through a domain knowledge graph before LLM inference yields significantly higher answer precision. Kutlu et al. [5] empirically measure the impact of retrieval strategy on downstream LLM response quality across enterprise document corpora.

C. Agentic AI and Multi-Step Reasoning

Rasheed et al. [8] provide a comprehensive survey of agentic AI in IEEE Access, reviewing architectures, applications, and open challenges in deploying LLM agents across industries, identifying LangGraph-style directed graph pipelines as the leading production approach for complex, verifiable reasoning chains. Li et al. [9] present a technical framework for LLM-powered multi-agent systems at IEEE ICAAI 2025, demonstrating that separating retrieval, analysis, and verification responsibilities across specialized agents achieves measurably higher accuracy than monolithic single-agent inference.

D. LLM Guardrails and Hallucination Mitigation

Devino et al. [6] present a practical engineering study on designing LLM guardrail components in production environments at IEEE/ACM CAIN 2025, identifying five critical guardrail categories that directly inform the five-layer architecture adopted in this platform. Cheung [7] demonstrates at IEEE ICIT 2025 that small language models can serve as effective hallucination detectors in a self-critique pipeline, motivating the Verifier Agent design in the proposed system.

E. Privacy-Preserving LLM Deployment

Baldini et al. [10] evaluate a fully local, privacy-preserving LLM chatbot confirming that on-device inference with no external API calls satisfies GDPR data residency requirements, validating the offline-first approach in this platform. Chen et al. [11] demonstrate that entity-level local differential privacy perturbation can prevent the reconstruction of sensitive source documents from RAG model outputs. Yan et al. [12] survey attack vectors including membership inference and prompt extraction, recommending local inference combined with output guardrails as the most robust defence for regulated industry deployments.

III. EXISTING METHODOLOGIES AND LIMITATIONS

A. Cloud-Based LLM APIs

The prevailing commercial approach relies on cloud-hosted LLM APIs such as OpenAI GPT-4o, Anthropic Claude 3.5, and Google Gemini 1.5 Pro. While these services deliver state-of-the-art inference quality, all input context including proprietary business documents, financial data, and customer records must be transmitted to third-party servers, directly

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802390

incompatible with GDPR Article 46, HIPAA SS164.312, and PCI-DSS requirements. Per-token pricing at scale renders high-volume enterprise deployments economically unsustainable for many mid-market organizations, and these systems provide no built-in guardrails against PII leakage or hallucination.

B. Fine-Tuned Domain-Specific Models

Fine-tuning a pre-trained foundation model requires access to large GPU compute clusters with training cycles measured in days to weeks per domain, making it prohibitively expensive. Fine-tuned models suffer from knowledge staleness, cannot incorporate newly uploaded documents without costly retraining, do not natively support multi-step agentic reasoning, and typically require transmitting proprietary training data to cloud infrastructure.

C. Standard Single-Stage RAG Pipelines

Standard single-stage RAG pipelines exhibit several structural limitations in enterprise BI contexts: they lack causal relational reasoning across KPI hierarchies; they do not support multi-step agent workflows; they provide no systematic guardrail enforcement; and retrieved chunk ordering is determined solely by vector similarity without cross-encoder re-ranking for precision. Most implementations continue to rely on cloud LLM APIs, inheriting the associated privacy limitations.

IV. PROPOSED SYSTEM ARCHITECTURE

A. Four-Tier Platform Overview

The AI TEACH Innovators platform is organized as a four-tier architecture executing locally on the deployment host. Tier 1 (React 18 Frontend) provides user-facing chat, dashboard, upload, and report interfaces. Tier 2 (FastAPI Backend) exposes a structured REST API with four primary endpoints. Tier 3 (LangGraph Agentic Pipeline) implements the three-node reasoning DAG with integrated guardrail enforcement. Tier 4 (Data Infrastructure) comprises ChromaDB for vector storage, Neo4j for the KPI knowledge graph, and Ollama with LiteLLM for fully local LLM inference. No data crosses any network boundary to external services at any point during operation.

B. Data Infrastructure Layer

ChromaDB serves as the persistent vector database storing dense embeddings of all ingested enterprise documents, parsed using Unstructured.io supporting PDF, DOCX, XLSX, and plain-text formats. Each document is segmented into 800-character chunks with 100-character overlaps and encoded using the all-MiniLM-L6-v2 sentence-transformer model, producing 384-dimensional dense vectors. Neo4j serves as the property graph database hosting the KPI influence graph modelled as a directed signed graph where each node represents a business KPI and each INFLUENCES relationship carries a polarity label and weight. Ollama provides local LLM serving with llama3.1:8b quantised to 4-bit precision, requiring approximately 5.5 GB VRAM and delivering 32-38 tokens per second on an RTX 4060.

C. Agentic Pipeline: LangGraph Three-Node DAG

The pipeline state is represented as a typed AgentState TypedDict with explicit fields for query, retrieved passages, KG context, analyst response, verification result, confidence score, and final output. The Retrieval Agent performs dense vector search over ChromaDB, retrieving the top-10 candidate chunks, then cross-encoder re-ranking retains the top-5 most contextually precise passages. In parallel, Neo4j Cypher traversal retrieves related KPI nodes and causal context. The Analyst Agent serialises enriched context into TOON pipe-delimited format, reducing token consumption by 30-60% compared to JSON. The Verifier Agent issues an LLM self-critique call, returning a structured JSON object containing a confidence score (0-100), ungrounded claims list, and final response.

D. Five-Layer Guardrail Engine

Every document RAG query passes through all five guards in strict order; failure at any layer immediately terminates the request. Guard 1 (Input Safety) applies regex-based pattern matching against known prompt injection templates and jailbreak sequences from the OWASP LLM Top 10 taxonomy. Guard 2 (Document Scope) enforces session-level isolation by verifying a valid session_id and ingested documents. Guard 3 (Relevance Gate) computes cosine distance between query and retrieved chunk embeddings; distances greater than 0.75 are blocked. Guard 4 (PII Redaction) scans LLM output to detect and redact email addresses, phone numbers, credit card patterns, and SSN formats. Guard 5 (Hallucination Check) invokes a dedicated LLM self-critique call to flag any ungrounded assertions.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802390

V. TECHNICAL IMPLEMENTATION

A. Document Ingestion and Embedding Pipeline

The pipeline begins when a user uploads files through the React frontend to the /upload FastAPI endpoint. Unstructured.io performs layout-aware parsing identifying headings, paragraphs, tables, list items, and captions as metadata annotations. The sliding window chunker ensures semantically complete sentences are never split across chunk boundaries. The all-MiniLM-L6-v2 model encodes each chunk with sub-100 ms latency enabling approximately 50 chunks per second ingestion throughput. Each vector is upserted into ChromaDB with metadata fields including session_id, source_filename, chunk_index, page_number, and upload_timestamp.

B. Hybrid Retrieval with Cross-Encoder Re-Ranking

The two-stage hybrid retrieval strategy combines the recall breadth of dense vector search with the precision of cross-encoder re-ranking. In the first stage, ChromaDB performs approximate nearest-neighbour search using HNSW indexing to retrieve the top-10 candidate chunks filtered by session_id. In the second stage, a cross-encoder model jointly encodes each query-chunk pair in a single transformer forward pass, producing a relevance score that captures fine-grained semantic interaction beyond what bi-encoder embeddings can represent. The top-5 candidates by cross-encoder score are selected as the final retrieval context.

C. KPI Prediction and Forecasting Engine

The /predict endpoint applies linear regression to compute trend slope and intercept from historical KPI observations. For each future period t, the forecast value is computed as: y-hat(t) = alpha + beta * t +/- 0.1|y-hat(t)|, where alpha is the intercept, beta is the slope, and the confidence interval is 10%. The KPI entity is resolved in Neo4j to retrieve influence relationships, which are serialised and passed to llm_generate() for a contextual business insight interpreting the forecast direction. More sophisticated models such as Facebook Prophet or ARIMA can be substituted as drop-in replacements without modifying the API contract.

VI. RESULTS AND EVALUATION

A. Experimental Setup

All experiments were conducted on a workstation with an NVIDIA GeForce RTX 4060 GPU (8 GB GDDR6 VRAM), Intel Core i7-13700K CPU, 32 GB DDR5 RAM, and 1 TB NVMe SSD, running Ubuntu 22.04 LTS with CUDA 12.2. Ollama served the llama3.1:8b model quantised to 4-bit GGUF precision. The evaluation covered three operational modes: (1) conversational business analytics via the agentic pipeline, (2) document Q&A with all five guardrail layers active, and (3) KPI time-series forecasting. A dedicated pytest suite comprising 21 structured test cases was executed across all five guard layers.

B. Guardrail Evaluation Results

Guard 1 achieved a 100% block rate against all 12 adversarial patterns including seven prompt injection templates, three jailbreak sequences, and two harmful content keyword tests. Guard 2 achieved 100% session isolation across all multi-user scenarios. Guard 3 correctly blocked all four out-of-scope queries exceeding cosine distance 0.75. Guard 4 correctly redacted all eight injected PII instances (two emails, two phone numbers, two credit card numbers, two SSNs) with zero false positives. Guard 5 correctly flagged hallucinated numeric claims in seven of eight test cases; the single false negative occurred when a fabricated figure closely matched a legitimate number in a non-relevant section of the source document.

C. System Performance Metrics

End-to-end query response time averaged 12.4 seconds, within the 15-second target, encompassing all guardrail evaluations, ChromaDB vector search, cross-encoder re-ranking, Neo4j Cypher traversal, and two sequential LLM inference calls. KPI prediction achieved an R-squared of 0.81 on a held-out quarterly financial time-series test set. Document ingestion throughput averaged approximately 50 chunks per second. The three-node LangGraph pipeline achieved 100% grounded responses across 10 evaluated business queries compared to 60% for the single-call LLM baseline. The TOON serialisation format reduced average prompt token count by 43% compared to equivalent JSON encoding. The platform maintained 100% offline operation with zero external API calls throughout all evaluation scenarios.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802390

VII. CONCLUSION

This paper presented AI TEACH Innovators, a production-grade, privacy-preserving conversational business intelligence platform unifying Hybrid RAG, Knowledge Graph causal reasoning, and Agentic AI within a fully offline, zero-cloud-dependency architecture. The five-layer sequential guardrail engine provides deterministic, auditable protection against prompt injection (G1), session scope leakage (G2), semantic irrelevance (G3), PII exposure (G4), and factual hallucination (G5). Evaluation results confirm sub-15-second response times, 100% offline operation, and 100% detection rates across Guards 1-4, with 87.5% hallucination detection on Guard 5. The three-node LangGraph agentic pipeline improved grounded response rates from 60% to 100% across ten evaluated business queries.

VIII. FUTURE WORK

Future work will explore three principal directions. First, replacing the linear regression forecasting engine with locally-deployable time-series models (Facebook Prophet or LSTM) to improve prediction accuracy for seasonal and non-linear KPI trajectories. Second, implementing a federated KG enrichment mechanism to enable collaborative KPI graph construction across organizational departments without transmitting raw data. Third, extending the guardrail suite with a dedicated bias and fairness detection layer suitable for HR and customer-facing AI applications, and introducing an adaptive threshold mechanism that tunes guardrail thresholds based on observed false-positive and false-negative rates during deployment.

REFERENCES

1. S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu, "Unifying large language models and knowledge graphs: A roadmap," IEEE Trans. Knowl. Data Eng., vol. 36, no. 7, pp. 3580-3599, Jul. 2024. DOI: 10.1109/TKDE.2024.3352100
2. T. Procko and O. Ochoa, "Graph retrieval-augmented generation for large language models: A survey," in Proc. IEEE Conf. Artif. Intell. (IEEE CAI), Singapore, 2024. DOI: 10.1109/CAI59869.2024.00205
3. R. Kumar, K. Ishan, H. Kumar, and A. Sharma, "LLM-powered knowledge graphs for enterprise intelligence and analytics," in Proc. IEEE Int. Conf. (ICSPCS), 2025. DOI: 10.1109/ICSPCS64044.2025.11088890
4. Fayazi, C. Wang, M. Bettini, and A. Trezza, "From data to decisions: Enterprise-level domain-specific graph retrieval-augmented generation systems for advanced question answering," in Proc. IEEE Int. Conf. Big Data, 2025. DOI: 10.1109/BigData62323.2024.10936909
5. E. Kutlu, S. Gul, and Y. T. Ic, "Retrieval augmented generation (RAG) and LLM integration," in Proc. IEEE Int. Conf. Comput. Sci. Eng. (UBMK), 2025. DOI: 10.1109/UBMK63289.2024.10845308
6. M. Devino, E. Ju, and P. M. Caldeira Junior, "Designing and implementing LLM guardrails components in production environments," in Proc. IEEE/ACM 4th Int. Conf. AI Eng. (CAIN), 2025, pp. 12-17. DOI: 10.1109/CAIN62612.2025.11029999
7. M. Cheung, "Hallucination detection with small language models," in Proc. IEEE Int. Conf. Inf. Technol. (ICIT), 2025. DOI: 10.1109/ICIT62883.2025.11108162
8. C.Nagarajan and M.Madheswaran - 'Stability Analysis of Series Parallel Resonant Converter with Fuzzy Logic Controller Using State Space Techniques'- Taylor &Francis, Electric Power Components and Systems, Vol.39 (8), pp.780-793, May 2011. DOI: 10.1080/15325008.2010.541746
9. C.Nagarajan and M.Madheswaran - 'Experimental verification and stability state space analysis of CLL-T Series Parallel Resonant Converter' - Journal of Electrical Engineering, Vol.63 (6), pp.365-372, Dec.2012. DOI: 10.2478/v10187-012-0054-2
10. C.Nagarajan and M.Madheswaran - 'Performance Analysis of LCL-T Resonant Converter with Fuzzy/PID Using State Space Analysis'- Springer, Electrical Engineering, Vol.93 (3), pp.167-178, September 2011. DOI 10.1007/s00202-011-0203-9
11. S.Tamilselvi, R.Prakash, C.Nagarajan,"Solar System Integrated Smart Grid Utilizing Hybrid Coot-Genetic Algorithm Optimized ANN Controller" Iranian Journal Of Science And Technology-Transactions Of Electrical Engineering, DOI10.1007/s40998-025-00917-z,2025
12. S.Tamilselvi, R.Prakash, C.Nagarajan," Adaptive sliding mode control of multilevel grid-connected inverters using reinforcement learning for enhanced LVRT performance" Electric Power Systems Research 253 (2026) 112428, doi.org/10.1016/j.epsr.2025.112428
13. S.Thirunavukkarasu, C. Nagarajan, 2024, "Performance Investigation on OCF and SCF study in BLDC machine using FTANN Controller," Journal of Electrical Engineering And Technology, Volume 20, pages 2675–2688, (2025), doi.org/10.1007/s42835-024-02126-w

14. C. Nagarajan, M.Madheswaran and D.Ramasubramanian- 'Development of DSP based Robust Control Method for General Resonant Converter Topologies using Transfer Function Model'- Acta Electrotechnica et Informatica Journal , Vol.13 (2), pp.18-31,April-June.2013, DOI: 10.2478/aeei-2013-0025.
15. C.Nagarajan and M.Madheswaran - 'DSP Based Fuzzy Controller for Series Parallel Resonant converter'- Springer, Frontiers of Electrical and Electronic Engineering, Vol. 7(4), pp. 438-446, Dec.12. DOI 10.1007/s11460-012-0212-0.
16. C.Nagarajan and M.Madheswaran - 'Experimental Study and steady state stability analysis of CLL-T Series Parallel Resonant Converter with Fuzzy controller using State Space Analysis'- Iranian Journal of Electrical & Electronic Engineering, Vol.8 (3), pp.259-267, September 2012.
17. C.Nagarajan and M.Madheswaran, "Analysis and Simulation of LCL Series Resonant Full Bridge Converter Using PWM Technique with Load Independent Operation" has been presented in ICTES'08, a IEEE / IET International Conference organized by M.G.R.University, Chennai.Vol.no.1, pp.190-195, Dec.2007
18. Suganthi Mullainathan, Ramesh Natarajan, "An SPSS and CNN modelling based quality assessment using ceramic materials and membrane filtration techniques", Revista Materia (Rio J.) Vol. 30, 2025, DOI: https://doi.org/10.1590/1517-7076-RMAT-2024-0721
19. M Suganthi, N Ramesh, "Treatment of water using natural zeolite as membrane filter", Journal of Environmental Protection and Ecology, Volume 23, Issue 2, pp: 520-530,2022
20. Rasheed et al., "Agentic AI: A comprehensive survey of technologies, architectures, applications, and challenges," IEEE Access, vol. 13, pp. 43059-43089, 2025. DOI: 10.1109/ACCESS.2025.3549268
21. X. Li et al., "LLM-powered multi-agent systems: A technical framework for collaborative intelligence," in Proc. IEEE Int. Conf. Adv. Artif. Intell. (ICAAI), 2025. DOI: 10.1109/ICAAI64093.2025.11077480
22. M. Baldini, G. Curia, and D. Pedrazzi, "Privacy-preserving LLM-based chatbots for hypertensive patient self-management," Smart Health, vol. 35, p. 100512, Mar. 2025. DOI: 10.1016/j.smhl.2025.100512
23. Y. Chen, H. Li, W. Zhao, and J. Yin, "Mitigating privacy risks in retrieval-augmented generation via locally private entity perturbation," Inf. Process. Manage., vol. 62, no. 4, p. 104138, 2025. DOI: 10.1016/j.ipm.2025.104138
24. L. Yan, C. Tang, Z. Ye, J. Zhao, P. Yu, and Y. Chen, "On protecting the data privacy of large language models (LLMs) and LLM agents: A literature review," Comput. Secur., vol. 152, p. 104413, May 2025. DOI: 10.1016/j.csa.2025.100083
25. Gopinathan, V. R. (2024). Real-Time Fault-Tolerant Multi-Cloud Database Architectures for High Availability Applications. International Journal of Future Innovative Science and Technology (IJFIST), 7(4), 13148.
26. Chandra, S., Rengarajan, A., Sahoo, G. S., & Sharma, S. (2023, December). Identifying Neuronal Damage and Plasticity by Analyzing Changes in Diffusion Tensor Imaging. In International Conference on Data Science, Machine Learning and Applications (pp. 433-438). Singapore: Springer Nature Singapore.
27. Sugumar, R. (2025). Federated AI in Offline-First Mobile Health Architectures for Privacy-Preserving Clinical Intelligence. International Journal of Science, Research and Technology, 8(4), 14589-14600.
28. Murugeshwari, B., Rajalakshmi, S., & Sudharson, K. (2023). Hybrid Approach for Privacy Enhancement in Data Mining Using Arbitrariness and Perturbation. Computer Systems Science & Engineering, 44(3).
29. Pandey, V. K., Mishra, S., Rengarajan, A., Savita, & Roomi, M. M. (2024, March). Enhancing Weather Forecasting with Machine Learning Techniques. In International Conference on Renewable Power (pp. 147-156). Singapore: Springer Nature Singapore.
30. Soundappan, S. J. (2025). Next Generation AI Enabled Holistic Cognitive Platform for Secure Cloud Network Intelligence Enterprise Systems and Digital Trust Optimization. International Journal of Computer Technology and Electronics Communication, 8(5), 11534-11542.
31. Mathew, A. (2022). Leveraging Big Data Analytics to Power AI and ML (Machine Learning) Automation. Educational Research (IJMCER), 4(5), 131-134.
32. Sugumar, R. (2024). AI-Augmented Quality Engineering for Performance Optimization and Test Orchestration in Distributed Systems. International Journal of Science, Research and Technology, 7(5), 12835-12846.
33. Akila, R. (2024). A deep reinforcement learning approach for optimizing inventory management in the agri-food supply chain. J. Electrical Systems, 20(4s), 2238-2247.
34. Mahendran, M., Anbazhagan, K., Pavithran, G., Nivas, A., & Pandey, S. D. (2022). Earthquake Damage Prediction using Machine Learning. Grenze International Journal of Engineering & Technology (GIJET), 8(1).
35. Gopinathan, V. R. (2025). Enterprise AI Frameworks for Financial Data Engineering Behavioural Analytics and Intelligent Cloud Solutions. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 8(4), 12499-12506.
36. Kondalsamy, P., & Kaliappan, K. (2025). An Optimal Prediction of Leaf Disease Based on Hybrid Deep Learnings and Metaheuristic Technique. Traitement du Signal, 42(1), 363.

37. Deivendran, P., Babu, P. S., Malathi, G., Anbazhagan, K., & Kumar, R. S. (2023). Emotion Recognition for Challenged People Facial Appearance in Social using Neural Network. arXiv preprint arXiv:2305.06842.
38. Sugumar, R. (2025). Unified AI Framework for Predictive Data Engineering and Real Time Prescription and Billing Systems. International Journal of Advanced Engineering Science and Information Technology (IJAESIT), 8(5), 17261.
39. Vekariya, V., Kumar, S., & Rengarajan, A. (2024). A distinctive and smart agricultural knowledge-based framework using ontology. In Sustainability in Digital Transformation Era: Driving Innovative & Growth (pp. 207-213). CRC Press.
40. Gopinathan, V. R. (2025). Software engineering practices for AI-driven systems: From development to deployment (MLOps perspective). International Journal of Science, Research and Technology, 8(1), 13493-13500.
41. Mathew, A. R. (2022). Threats and protection on E-sim: a prospective study. Novel Perspectives of Engineering Research, 8, 76-81.
42. Naveena, S., & Kavitha, K. (2025). Gossypium herbaceum: Folium disease identification and classification using Efficient Net-Coordinate Convolutional Neural Network (EcoNet). Engineering Applications of Artificial Intelligence, 152, 110701.
43. Rengarajan, A., Mishra, A., Kulhar, K. S., Shrivastava, V. P., & Alawneh, Y. J. J. (2024, March). Role of Deep Reinforcement Learning in Mitigating Cyber Security Issues: A Review. In International Conference on Renewable Power (pp. 37-48). Singapore: Springer Nature Singapore.
44. Achari, A. P. S. K., & Sugumar, R. (2024, November). Performance analysis and determination of accuracy using machine learning techniques for naive bayes and random forest. In AIP Conference Proceedings (Vol. 3193, No. 1, p. 020199). AIP Publishing LLC.
45. Mathew, A., & Alex, H. (2022). Detect & protect-medical device cybersecurity. Curr. Overview Sci. Technol. Res, 1, 60-68.
46. Sammy, F., Chettier, T., Boyina, V., Shingne, H., Saluja, K., Mali, M., ... & Shobana, A. (2025). Deep Learning-Driven Visual Analytics Framework for Next-Generation Environmental Monitoring. Journal of Applied Science and Technology Trends, 114-122.
47. Anbazhagan, K. (2024). Trustworthy and Adaptive AI Systems for Enterprise Analytics Cybersecurity and Decision Optimization Using API-First and Cloud-Native Architectures. International Journal of Technology, Management and Humanities, 10(03), 65-74.
48. Mathew, A. (2021). Deep reinforcement learning for cybersecurity applications. Int J Comput Sci Mob Compu, 10(12), 32-38.
49. Dhinakaran, D. (2022). Joe Prathap P. M, Selvaraj D, Arul Kumar D and Murugeshwari B," Mining Privacy-Preserving Association Rules based on Parallel Processing in Cloud Computing,". International Journal of Engineering Trends and Technology, 70(3), 284-294.
50. Karthika, K., Anusha, K., Kavitha, K., Harshadha, R., Dharshini, D. S., & Sundhar, N. A. (2025, April). Frequency Reconfigurable Antenna using Advanced Materials: A Study. In 2025 3rd International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA) (pp. 1-6). IEEE.
51. Thavamani, C., & Rengarajan, A. (2024). Clustering related behaviour of users by the use of partitioning and parallel transaction reduction algorithm. International Journal of Advanced Intelligence Paradigms, 29(2-3), 122-132.
52. Sugumar, R. (2025). Unified AI Framework for Predictive Data Engineering and Real Time Prescription and Billing Systems. International Journal of Advanced Engineering Science and Information Technology (IJAESIT), 8(5), 17261.
53. Soundappan, S. J., & Sugumar, R. (2016). Optimal knowledge extraction technique based on hybridisation of improved artificial bee colony algorithm and cuckoo search algorithm. International Journal of Business Intelligence and Data Mining, 11(4), 338-356.
54. SakthiPreetha, A., Kavitha, K., Karthika, K., & Manohari, R. G. (2025, April). A Novel Metasurface-Embedded Antenna for WBAN Communications. In 2025 3rd International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA) (pp. 1-4). IEEE.
55. Murugeshwari, B., Selvaraj, D., Sudharson, K., & Radhika, S. (2023). Data Mining with Privacy Protection Using Precise Elliptical Curve Cryptography. Intelligent Automation & Soft Computing, 35(1).
56. Gopinathan, V. R. (2025). Software engineering practices for AI-driven systems: From development to deployment (MLOps perspective). International Journal of Science, Research and Technology, 8(1), 13493-13500.
57. Anbazhagan, K., Kumar, R., Thilagavathy, R., & Anuradha, D. (2024, March). Shortest Job First with Gateway-based Resource Management Strategy for Fog Enabled Cloud Computing. In 2024 4th International Conference on Data Engineering and Communication Systems (ICDECS) (pp. 1-6). IEEE.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802390

58. Kannadhasan, S., Vasuki, S., Kavitha, K., Karthikeyan, P., & Usha, S. G. A. (Eds.). (2025, April). Preface: Role of Artificial Intelligence and IoT in Engineering, Technology & Science [ICRAETS 2024]. In AIP Conference Proceedings (Vol. 3258, No. 1, p. 010001). AIP Publishing LLC.
59. Dhinakaran, D., Prathap, P. J., Selvaraj, D., Kumar, D. A., & Murugeshwari, B. (2022). Mining privacy-preserving association rules based on parallel processing in cloud computing. International Journal of Engineering Trends and Technology, 70(3), 284-294.