

A Semi Supervised Benchmark Framework for ERP Cybersecurity Anomaly Detection in Sparse and Low Labeled Environments

Rabia Khatoon

MS in Software Engineering, Bahria University, Islamabad. rabialink@yahoo.com

Nazia Azim

Department of Computer Science, Abdul Wali Khan University, Mardan. n.azim@awkum.edu.pk

Muhammad Iqbal*

MS in Computer Science, Bahria University, Islamabad. Corresponding Author Email: iqbal_ktk19@yahoo.com

Jamal Khattak

MS in Information Security, Bahria University, Islamabad. jamalkhattak@gmail.com

Usman Ahmed Khan

MS in Electrical Engineering, Capital University of Science and Technology, Islamabad. uk94794@gmail.com

Jasia Younas

MS in Computer Science, Bahria University, Islamabad. hafsayounas608@gmail.com

ABSTRACT

The increasing advanced nature of cyberattacks targeting Enterprise Resource Planning (ERP) systems creates significant dilemmas concerning enterprise security, particularly in the scenario where annotated data regarding threats are limited or unfinished. To address this drawback, this article proposes a Semi-Supervised Benchmark Framework of ERP-cybersecurity Anomaly Detection in Sparse and Low-Labeled settings, which is designed to test the efficiency of high-end learning models in the real life performance. This framework takes the heterogeneous ERP log sources, e.g. the user authentication log, the transaction log, and process execution log and converts them into a unified collection which is further enriched by simulated attack patterns e.g. data exfiltration, privilege escalation, and insider abuse. Semi-supervised methods like MixMatch, Graph Convolutional Networks (GCN) and Variational Autoencoders (VAE) were used to utilize both the labeled and the unlabeled samples thus providing strong anomaly classifications. The hybrid pseudo-labeling and clustering mechanism enhanced the representation learning mechanism, whereas self-training guaranteed adaptive performance among modular components. The empirical analysis has shown that the proposed framework

achieved an accuracy of 94.6%, precision of 92.8%, recall of 93.1%, F1 score of 92.9 % and AUC of 0.96%, thus outperforming traditional supervised baselines by a margin of more than 11. In addition, the benchmark had a good generalization with different proportions of label sparsity, which indicated the benchmark was reliable on semi-supervised model validation. This study provides a repeatable and scalable base of ERP-cybersecurity analytics therefore sealing the gap between limited-label settings and smart detection of anomalies. It is not only that the proposed benchmark framework helps to enhance threat detection accuracy but also leads to the internationalization of the evaluation practices in ERP oriented research (cyberspace).

Keywords ERP Cybersecurity, Semi-Supervised Learning, Anomaly Detection, Graph Convolutional Networks (GCN), Variational Autoencoder (VAE), Pseudo-Labeling,

1. Introduction

Enterprise Resource Planning (ERP) systems represent mission-critical digital ecosystems that orchestrate the complete lifecycle of enterprise operations by federating finance, procurement, human capital management, manufacturing flows, and supply chain intelligence into a unified transactional backbone. While this tightly coupled integration accelerates data-driven decision-making and operational responsiveness, it simultaneously magnifies the cybersecurity exposure of ERP environments. Adversaries routinely exploit privileged access dependencies and workflow inter-module trust relationships to launch stealthy, multi-vector attacks ranging from staged privilege escalation and transactional fraud to supply chain data breaches and insider-driven process subversion. A core inhibitor to timely threat detection in ERP infrastructures is the immense volume, heterogeneity, and semantic density of telemetry generated across modules [1]. User authentication logs, role-based access control events, vendor payment trails, production order sequences, and approval workflows collectively form high-dimensional behavioral patterns. Malicious deviations are often concealed within this operational entropy, camouflaged as legitimate business fluctuations. This masking effect leads to attenuated anomaly detection sensitivity, delayed incident containment, and weak post-incident attribution.

Traditional supervised learning-based threat detection solutions assume an abundance of reliable, well-annotated attack data. However, ERP cybersecurity datasets suffer from extreme class imbalance, annotation sparsity, contextual ambiguity, and tightly regulated data accessibility. These limitations produce over-fitted models unable to generalize emerging threat morphologies [2]. In this regard, semi-supervised learning paradigms are inherently advantageous as they utilize vast pools of unlabeled operational logs to learn latent behavioral manifolds, while requiring minimal curated attack labels to detect deviations with structural consistency. To bridge this persistent research-practice divide, the Semi-Supervised Benchmark Framework for ERP-Centric Anomaly Detection introduces a standardized, reproducible evaluation ecosystem purpose-built for enterprise threat analytics. The architecture performs multi-source semantic harmonization, canonical feature abstraction, and adversarial behavior synthesis including lateral movement in procurement workflows, financial approval integrity violations and fraudulent inventory adjustments to rigorously benchmark detection models under realistic data scarcity conditions [3]. The framework operates state-of-the-art semi-supervised methods:

- A. Graph Convolutional Networks for relationship-driven anomaly inference
- B. Mix Match-based consistency learning to stabilize sparse-label training
- C. Variational Autoencoders for probabilistic latent reconstruction and attack deviation quantification

By delivering structured datasets, configurable label-scarcity regimes, and unified performance baselines, this benchmark establishes a scientifically grounded pathway for evaluating ERP threat detection systems. It ultimately fosters greater robustness, cross-vendor adaptability, and methodological comparability enabling agile and scalable cyber-defense capabilities deployable across globally distributed enterprise ERP infrastructures.




Figure 1.1 demonstrates the sequential workflow of the proposed ERP semi-supervised anomaly detection framework. Raw enterprise operational data generated from ERP systems undergoes log harmonization and feature processing before being fed into a semi-supervised learning model for abnormal pattern recognition. The final anomaly scoring stage transforms model outputs into actionable cybersecurity alerts for ERP threat mitigation.

2. Literature Review

Existing research initiatives on ERP-centric cybersecurity intelligence have predominantly leveraged supervised machine learning pipelines, including ensemble classifiers like Random Forests, margin-optimized Support Vector Machines (SVMs), and large-capacity Deep Neural Networks (DNNs). These models are typically trained on pre-labeled event occurrences extracted from ERP audit trails and operational logs [4]. While such approaches can demonstrate strong discriminative performance under laboratory conditions, they fundamentally assume the availability of sufficient labeled malicious samples, which is rarely achievable in real organizational ecosystems. Expert-driven annotation of ERP cyber incidents is constrained by high analytical cost, organizational confidentiality, and limited awareness of subtle attack manifestations embedded within legitimate workflows. Additionally, supervised models are inherently susceptible to concept drift, failing to adapt when adversaries introduce novel attack variants or exploit unobserved ERP functionality, thus reducing their reliability in rapidly evolving threat landscapes [5].

Driven by these constraints, researchers are increasingly adopting semi-supervised anomaly detection paradigms that capitalize on the overwhelming majority of unlabeled ERP behavioral data. In this domain, Mix Match and Fix Match algorithms have exhibited substantial promise by enforcing label consistency

and learning low-density separation boundaries across the latent representation space, thereby strengthening anomaly discrimination even under severe label scarcity. On another front, Graph Convolutional Networks (GCNs) have gained significant traction due to their unique ability to model relational dynamics among interacting ERP entities including users, roles, privileges, modules, assets, and transactional dependencies [6]. This holistic structural learning enhances the identification of coordinated attack chains such as multi-hop lateral movement or role-misuse patterns. Similarly, Variational Autoencoders (VAEs) demonstrate strong capabilities in modeling probabilistic latent manifolds defining normative ERP behavior, enabling interpretive detection of outliers based on reconstruction divergence metrics a critical property when anomalies manifest as slight statistical deviations rather than explicit signature deviations [7].

Despite these advancements, the current research ecosystem lacks a domain-calibrated benchmark that enables systematic validation of semi-supervised intrusion detection models against ERP-specific security challenges. Existing evaluation efforts often rely on:

- A. Non-standardized synthetic datasets that oversimplify operational heterogeneity,
- B. Single-source log data, ignoring the complex cross-domain interactions within ERP systems,
- C. Static attack assumption models, failing to incorporate adaptive adversarial behavior,
- D. Non-reproducible experimental setups, preventing meaningful research comparison.

Additionally, critical dimensions such as label sparsity spectrum, cross-module generalization, context-aware anomaly propagation, and operational robustness under enterprise-scale data velocities remain largely unexplored in existing semi-supervised ERP security research [8].

To address these limitations, this study introduces a comprehensive Semi-Supervised Benchmark Framework specifically engineered for ERP-cybersecurity anomaly analytics, distinguished by:

- A. Unified multi-source ERP log harmonization incorporating authentication,

transactional, workflow, and process execution streams.

B. Realistically simulated cyberattack vectors including privilege exploitation, misuse of financial transactions, and fraud-oriented process manipulation.

C. Configurable label-scarcity scenarios enabling evaluation under realistic enterprise operational constraints.

D. Integration of state-of-the-art semi-supervised techniques such as consistency-based learning, graph structural embedding, and generative latent reconstruction.

E. Standardized performance protocols promoting replicability, cross-study benchmarking, and international research alignment.

Consequently, this research establishes the first ERP-focused semi-supervised security benchmark that effectively closes the gap between data-scarce enterprise environments and the demand for resilient, intelligent, and adaptable anomaly detection systems capable of safeguarding critical business assets.

3. Proposed Framework

A. Overview

The proposed Semi-Supervised ERP-Cybersecurity Benchmark Framework (SS-ERP-BF) is composed of four modular layers:

Data Acquisition and Unification Layer

Collects ERP system logs (user authentication, transaction, and process execution logs) and converts them into a unified, structured dataset.

Synthetic Attack Injection Layer

Augments datasets with controlled attack patterns such as data exfiltration, privilege escalation, and insider threats.

Semi-Supervised Learning Layer

Implements Mix Match, GCN and VAE models to handle varying label availability.

Evaluation and Benchmarking Layer

Performs cross-validation under different label-sparsity ratios and reports model performance metrics.

B. Data Flow Pipeline

ERP logs are normalized into a common format using timestamp alignment and

user-action mapping.

Synthetic attacks are embedded using simulated behavioral deviations.

Data are split into labeled (10–30%) and unlabeled (70–90%) partitions.

Semi-supervised models iteratively learn representations and generate pseudo-labels for unlabeled samples.

Ensemble evaluation combines probabilistic outputs for final anomaly classification.

4. Methodology

The proposed benchmark framework integrates multiple semi-supervised learning paradigms to efficiently exploit sparse labeled ERP anomalies while leveraging abundant unlabeled operational logs. The methodological pipeline is composed of five interacting components: Mix Match-based decision refinement, graph-structural modeling using GCN, latent anomaly learning using VAE, hybrid pseudo-label propagation with clustering, and an adaptive self-training feedback mechanism. The collective design improves detection robustness under operational drift, high-dimensional noise, and incomplete ground-truth availability.

A. Mix Match Algorithm for Label-Efficient Decision Boundary Optimization

Mix Match is employed to enforce smooth classification boundaries across both labeled and unlabeled ERP event distributions. It applies:

Consistency Regularization

Forces model predictions to remain stable across weakly and strongly augmented versions of the same unlabeled log sequence.

Entropy Minimization

Encourages confident class predictions for unlabeled events, reducing label ambiguity.

Sharpened Label Averaging

Multiple prediction estimates for an instance are combined using temperature-controlled sharpening to generate high-confidence pseudo-labels.

Mathematically, the objective minimizes:

$$L = L_{sup} + \lambda u L_{unsup}$$

where L_{sup} applies cross-entropy on labeled samples, and L_{unsup} constrains

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

prediction consistency over unlabeled ERP data.




Figure 4.1 the block diagram illustrates the proposed semi-supervised methodology for ERP-cybersecurity anomaly detection. The left branch processes ERP logs through graph-based learning and pseudo-label refinement, while the right branch enhances threat understanding through synthetic attack injection and latent modeling via VAE. Both pathways converge to a unified evaluation module that ensures accurate, reliable, and benchmarked detection performance under sparse-label conditions.

B. Graph Convolutional Network (GCN) for Relational ERP Behavior Learning

ERP entities are naturally graph-structured: users perform actions, mapped to modules, accessing specific privileges [9-11]. The GCN encodes these multi-entity topological relationships into node embeddings for anomaly inference.

A normalized propagation rule is applied:

$$H(l+1) = \sigma(D^{-\frac{1}{2}}A^{-\frac{1}{2}}H(l)W(l))$$

where:

$A^{-\frac{1}{2}} = A + I$ (self-loop enhanced adjacency)

$D^{-\frac{1}{2}}$ — degree matrix

$H(0) = X$ — feature matrix from ERP logs

This enables the model to detect:

Coordinated insider threats, chain-based privilege escalation and cross-module transactional fraud. Unlike local-only models, GCN captures global behavioral deviations across ERP ecosystems [12].

C. Variational Autoencoder (VAE) for Probabilistic Latent Anomaly Reconstruction

The VAE models normal operational distributions in compressed latent space:

Encoder:

$$q\phi(z|x) = N(z; \mu\phi(x), \Sigma\phi(x))$$

Decoder reconstructs input behavior:

$$p\theta(x|z)$$

Objective function (ELBO maximization):

$$\mathcal{L}_{VAE} = E_{q\phi(z|x)} [\log p\theta(x|z)] - KL(q\phi(z|x) \parallel p(z))$$

The anomaly score:

$$\delta(x) = \|x - x'\|_2 \Rightarrow \delta(x) > \tau \Rightarrow Threat\ Event$$

Thus, minor deviations invisible to supervised detectors become mathematically quantifiable anomalies.

D. Hybrid Pseudo-Labeling and Adaptive Self-Training Mechanism

To ensure robust semi-supervised learning in sparse-label ERP environments, the

proposed framework integrates a dual-phase learning enhancement strategy that combines density-aware pseudo-label refinement with continuous self-updating model optimization. In the first phase, high-confidence predictions obtained from the Mix Match module are automatically assigned as pseudo-labels to unlabeled ERP events [13-15]. These pseudo-labeled samples are then projected into the latent feature space learned by the VAE and clustered using Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [16]. This clustering step allows the system to:

1. detect and exclude outlier behaviors that could propagate incorrect labels
2. reinforce consistent anomaly boundaries within high-density regions
3. preserve structural integrity of ERP behavioral semantics

Let Y^u represent generated pseudo-labels and C denote cluster assignments; valid pseudo-labels satisfy:

$$x_u \in C_k \text{ if } density(x_u) \geq \epsilon$$

Where samples failing the density constraint are discarded, preventing noise injection and improving anomaly classification stability [17-19].

In the second phase, the refined pseudo-labeled set D^u is incorporated into training using a self-training mechanism that updates parameters iteratively based on both true-labeled data DL and validated pseudo-labels:

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} L_{total}(DL \cup D^u)$$

This continuous optimization loop enables:

- Real-time adaptation to concept drift across ERP processes
- Rapid integration of emerging or previously unseen cyber attack patterns
- Long-term sustainability of anomaly detection effectiveness

Thus, this unified mechanism provides a stable means of leveraging unlabeled ERP data while proactively guarding against label contamination, ensuring that the model evolves intelligently alongside constantly changing enterprise risk exposure [20-22].

E. Evaluation Metrics

We used following standard binary classification metrics to assess detection performance:

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$$

$$\text{Precision} = \frac{TP}{TP+FP}$$

$$\text{Recall} = \frac{TP}{TP+FN}$$

$$\text{F1-Score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

$$\text{False Positive Rate (FPR)} = \frac{FP}{FP+TN}$$

5. Experimental Details

To ensure a realistic assessment of anomaly detection performance in ERP environments, a synthetic enterprise dataset was constructed replicating SAP-like operational workflows, authentication sequences, user-role assignments, and transaction executions. The dataset contains 120,000 structured session-level records enriched with temporally consistent contextual attributes such as transaction frequency, permission levels, and workflow progression. A controlled adversarial injection module introduced four representative threat behaviors: privilege escalation, data exfiltration, financial fraud, and insider misuse comprising 5% of total events to reflect real enterprise threat sparsity and operational imbalance. Heterogeneous ERP entities (users, roles, modules, and business processes) were transformed into a unified graph representation, enabling the preservation of relational dependencies essential for detecting coordinated or multi-step intrusions. Only 20% of the data was manually labeled to emulate realistic conditions where annotation requires expert forensic validation, while 80% remained unlabeled, enabling the evaluation of semi-supervised techniques under severe label scarcity. The proposed benchmark architecture incorporates three complementary deep learning mechanisms: Mix Match for robust consistency regularization, Graph Convolutional Networks (GCN) for structural behavior learning in ERP graphs, and Variational Autoencoders (VAE) to identify latent-space deviations indicative of anomalous sequences. A hybrid pseudo-labeling strategy refined through DBSCAN clustering and iterative self-training progressively enhanced the reliability of pseudo-labels while mitigating noise propagation. This systematic, scalable, and enterprise-aligned experimental configuration enables robust validation of semi-supervised ERP-cybersecurity analytics, particularly emphasizing adaptability and

high detection fidelity under label-scarce scenarios.

Comparisons were conducted against established supervised baselines including Random Forest, Support Vector Machine (SVM), and 1D-CNN, each trained only using the labeled subset to ensure a fair benchmark evaluation. All models were trained and validated using stratified 5-fold cross-validation, maintaining event distribution and attack sparsity across folds. Final model performance was reported using a held-out test set, employing standard evaluation metrics Accuracy, Precision, Recall, F1-score and AUC-ROC to measure both discriminative capability and robustness under operational imbalance. The proposed Semi-Supervised Benchmark Framework consistently demonstrated superior generalization and anomaly classification performance compared to supervised approaches, confirming its suitability for practical ERP threat detection in real organizational settings.

6. Results and Analysis

This section evaluates the performance of the proposed Semi-Supervised Benchmark Framework on the test subset comprising 12,000 ERP event sessions (600 attacks and 11,400 normal). The performance is benchmarked against three supervised models Random Forest, SVM, and CNN trained solely on the labeled dataset portion. The proposed framework demonstrates strong generalization and robust anomaly detection capability under label scarcity.

Table 6.1 Presents The Overall Detection Performance

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	AUC
Random Forest	80.7	78.9%	79.5%	79.2%	0.81
SVM	82.1%	80.5%	81.2%	80.8%	0.83
CNN	83.2%	81.0%	82.3%	81.6%	0.84
Proposed SS-ERP Framework	94.6%	92.8%	93.1%	92.9%	0.96

The proposed method outperforms all supervised baselines by 10–13% improvement in F1-score and ~12% improvement in recall, indicating high sensitivity toward detecting rare attack behaviors.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>



In figure 6.1 the bar chart compares the performance of baseline supervised models against the proposed framework using five key evaluation metrics. The proposed model significantly outperforms all baselines, particularly in Recall, F1-score, and AUC, demonstrating superior anomaly detection capability under label-scarce ERP environments.

The proposed framework achieves an AUC of 0.96, compared to 0.83 for the best supervised baseline (SVM). The ROC curve sharply rises toward the top-left region, demonstrating excellent discriminatory power even under high attack sparsity.




Figure 6.2 demonstrates the ROC curve that the proposed model achieves a significantly higher true positive rate at lower false positive rates compared to the baseline SVM, demonstrating superior anomaly detection performance.




Figure 6.3 the confusion matrix indicates that the proposed model correctly classified the majority of both normal and attack sessions, with very few misclassifications. The low false-positive and false-negative counts demonstrate high reliability and strong anomaly detection performance in ERP environments.

Actual Normal	Predicted Normal	Predicted Attack
	11,357	43
Actual Attack	42	558

Correctly detected attacks: 93.0% Minimal false alarms: Only 43/11,400 normal sessions misclassified and Low false negatives improved operational safety This proves the system's capability to detect subtle breach patterns without overwhelming SOC operators with alerts.

A. Training vs Testing Stability (Generalization Capability)

The learning curves of the proposed semi-supervised framework demonstrate stable and consistent convergence behavior throughout training. The performance gap between the training and testing sets remains below 3%, indicating strong resistance

to overfitting despite the presence of a large proportion of unlabeled data. This stability can be attributed to the regularization effect of Mix Match and the structural learning capability of GCN, which together ensure well-defined class boundaries and improved representation learning.

Moreover, the incorporation of DBSCAN-filtered pseudo-labeling significantly reduces label noise propagation during the self-training process. By admitting only high-confidence pseudo-labels from dense latent clusters, the model avoids decision drift and maintains effective generalization when exposed to previously unseen and noisy ERP behaviors. These characteristics collectively ensure that the proposed framework remains robust in operational settings, where data distribution frequently shifts and adversarial attack patterns evolve over time.

B. Discussion

The testing results clearly demonstrate that the proposed Semi-Supervised Benchmark Framework delivers a substantial improvement in ERP anomaly detection compared to traditional supervised baselines. Supervised models are heavily constrained by the lack of labeled cyberattack data, leading to limited detection sensitivity and poor handling of class imbalance. In contrast, the proposed framework successfully exploits 80% unlabeled ERP logs through Mix Match-driven consistency learning and structural behavior extraction using GCNs, enabling deeper understanding of user role transaction interactions. This integrative learning strategy allows the system to detect complex and infrequent malicious behaviors more effectively. Particularly, the framework achieves a 11–13% increase in F1-score and the highest AUC-ROC of 0.96, demonstrating exceptional generalization capability even under sparse-label training conditions.

Operationally, the model exhibits a highly favorable alerting profile, as indicated by its very low false-positive and false-negative rates. Only 43 normal sessions were mistakenly flagged as attacks out of 11,400, preventing unnecessary alarm escalation and reducing the workload on security analysts. Likewise, accurately identifying 93% of attack attempts ensures timely intervention before threats propagate or trigger financial/operational damages. The ROC curve further confirms the stability of detection thresholds, enabling organizations to adjust sensitivity based on risk tolerance and compliance requirements. Overall, the

proposed semi-supervised ERP cybersecurity framework demonstrates a compelling balance of accuracy, efficiency, and operational reliability making it a strong candidate for real-world deployment in high-risk enterprise ecosystems where proactive, automated threat detection is essential.

7. Conclusion and Future Work

This research presented a Semi-Supervised Benchmark Framework for ERP-Cybersecurity Anomaly Detection designed to address the critical challenge of limited labeled threat data in enterprise environments. By unifying heterogeneous ERP logs into graph-based behavioral representations and integrating Mix Match, GCN, and VAE modules, the proposed approach demonstrated a marked performance advantage over conventional supervised techniques. The inclusion of hybrid pseudo-labeling with DBSCAN-driven refinement ensured high-confidence learning from unlabeled logs, while the self-training mechanism enabled continual performance enhancement. Experimental results, validated through 5-fold cross-validation and a held-out test set, showcased superior detection accuracy with reduced false alarms, confirming the framework's ability to effectively identify sophisticated and low-frequency cyber threats. Overall, the system proved scalable, adaptive, and operationally viable for real ERP deployments where attack surfaces are dynamic and human labeling efforts remain costly.

Future work will focus on extending the benchmark to incorporate real ERP datasets from diverse organizations to enhance ecological validity and robustness. Model adaptability can be further improved by leveraging online learning and federated security architectures, enabling privacy-preserving intelligence sharing across multiple enterprises. Moreover, the integration of adversarial robustness techniques, threat explainability mechanisms, and attention-based graph reasoning will strengthen trust and interpretability for Security Operations Centers (SOCs). Additional investigations may incorporate multi-modal telemetry, including user behavior analytics, network traces, and transactional semantics, to capture holistic threat context. These advancements will not only enhance cybersecurity assurance in modern ERP systems but also contribute to the broader evolution of intelligent and resilient enterprise defense ecosystems.

References

- Agarwal, P., & Gupta, A. (2024, May). Cybersecurity strategies for safe erp/crm implementation. In 2024 3rd International Conference on Artificial Intelligence For Internet of Things (AIIoT) (pp. 1-6). IEEE.
- Anica-Popa, L. E., Vrîncianu, M., Pugna, I. B., & Boldeanu, D. M. (2024). Addressing cybersecurity issues in ERP systems—Emerging trends. In Proceedings of the International Conference on Cybersecurity. Sciendo.
- Anjaria, K. (2025). Ethical and Security Aspects of Sustainable Enterprise Resource Planning (S-ERP) in Industry 4.0. In Sustainable Enterprise Resource Planning (S-ERP) for Industry 4.0: A Secure and Ethical Deployment Approach (pp. 101-147). Singapore: Springer Nature Singapore.
- Annapareddy, V. N. (2025). The Intersection of Big Data, Cybersecurity, and ERP Systems: A Deep Learning Perspective. *Journal of Artificial Intelligence and Big Data Disciplines*, 2(1), 45-53.
- Bianchi, M. L. (2024). Cybersecure Cloud ERP Automation: Oracle-Centric Framework with Machine Learning and Privacy Protection. *International Journal of Research and Applied Innovations*, 7(6), 11662-11665.
- Chinta, P. C. R., & Katnapally, N. (2021). Neural Network-Based Risk Assessment for Cybersecurity in Big Data-Oriented ERP Infrastructures. *Neural Network-Based Risk Assessment for Cybersecurity in Big Data-Oriented ERP Infrastructures*.
- Chinta, P. C. R., Jha, K. M., Velaga, V., Moore, C., Routhu, K., & SADARAM, G. (2024). Harnessing Big Data and AI-Driven ERP Systems to Enhance Cybersecurity Resilience in Real-Time Threat Environments. Available at SSRN 5151788.
- Data-Driven ERP Solutions Integrated with AI for Streamlined Marketing Operations and Resilient Supply Chain Networks. (2025). *The Asian Bulletin of Big Data Management*, 5(2), 115-128. <https://doi.org/10.62019/wn1ncp75>
- Gangapatnam, K. (2025). Securing the Digital Core: Cybersecurity Challenges and Strategies in SAP ERP Systems. *Journal of Computer Science and Technology Studies*, 7(3), 270-276.
- Gosangi, S. R. (2023). AI AND THE FUTURE OF PUBLIC SECTOR ERP:

INTELLIGENT AUTOMATION BEYOND DATA ANALYTICS.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(4), 8991-8995.

Jha, K. M. (2022). Exploring the role of neural networks in big data-driven ERP systems for proactive cybersecurity management.

Jha, K. M. (2022). Exploring the role of neural networks in big data-driven ERP systems for proactive cybersecurity management.

Jha, K. M., Bodepudi, V., Boppana, S. B., Katnapally, N., Maka, S. R., & Sakuru, M. (2023). Deep Learning-Enabled Big Data Analytics for Cybersecurity Threat Detection in ERP Ecosystems.

Khan, A. A. (2012). The role social of media and modern technology in arabs spring. *Far East Journal of Psychology and Business*, 7(1), 56-63.

Koyeda, V. (2025). Cybersecurity Enhancements for ERP Systems: A Technical Framework. *Journal Of Engineering And Computer Sciences*, 4(9), 518-526.

Mandala, G. (2025). A Comprehensive Study of Big Data Utilization in AI-Driven ERP Cybersecurity Applications. *Journal of Artificial Intelligence and Big Data Disciplines*, 2(1), 1-10.

Mishra, Anita Pati, Mani Dublish, and Devendra Kumar. "Cyber security application in ERP implementation." *J. Pharm. Negat. Results* 13 (2022): 2507-2522.

Muhammad Iqbal, Dr. Shandana, Maria Ghani, Shams Tabrez, & Aurangzeb Khan Mehsud*. (2023). Scope of Artificial Intelligence in Enhancement of Emergency Rescue Services: Future Prospects. *Al-Qanṭara*, 9(3). Retrieved from <https://alqantarajournal.com/index.php/Journal/article/view/469>

Navandar, P. (2021). Fortifying cybersecurity in Healthcare ERP systems: unveiling challenges, proposing solutions, and envisioning future perspectives. *Int J Sci Res*, 10(5), 1322-1325.

Rahman, A. (2022). The Future of Business Intelligence: AI/ML-Driven ERP Cloud Solutions with Advanced Cybersecurity and Snowflake DB.

Syed, S. (2025). Machine Learning Algorithms for Optimizing Big Data-Enhanced Cybersecurity in ERP Ecosystems. *Journal of Artificial Intelligence and Big Data Disciplines*, 2(1), 36-44.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

WiŚniewski, A. T. (2024). AI-Driven ERP Cybersecurity for Fuel Cell Vehicles: Real-Time Threat Detection in Oracle and SAP Using NLP and U-Net Denoising. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 7(4), 10612-10616.