

Privacy-Preserving Machine Learning for Secure Business Intelligence

Otaxanova Umida

Teacher of Finance, Alfraganus University, Tashkent, Uzbekistan

ABSTRACT

Business intelligence (BI) increasingly relies on machine-learning models trained on customer, employee and transaction data, and these models can leak the data they were trained on through their parameters and outputs. Privacy-preserving machine learning offers formal protection, principally through differential privacy (DP), which bounds what any output reveals about any individual record, and through federated learning, which keeps data at its source. Both impose a utility cost. This paper examines the privacy–utility trade-off for a representative BI regression model in a simulation of 4,000 records. Without protection the model attains a hold-out R^2 of 0.69. Centralised training with output-perturbation DP retains R^2 of 0.68 at a privacy budget of $\epsilon = 0.5$ and 0.64 at $\epsilon = 0.25$, falling to 0.38 at $\epsilon = 0.1$; federated training with per-client DP across five sites loses utility faster, retaining 0.62 at $\epsilon = 0.5$ and 0.42 at $\epsilon = 0.25$. The paper explains why local protection costs more than central protection, maps the trade-off to BI use cases by sensitivity and precision requirements, and proposes an architecture and governance for privacy-preserving BI in which the privacy budget is a managed resource. The contribution is a practical account of how organisations can obtain analytical value from sensitive data under formal privacy guarantees.

Keywords: Privacy-preserving machine learning; differential privacy; federated learning; business intelligence; privacy–utility trade-off; data governance.

Journal of Teacher Education and Research (2026). DOI: 10.36268/JTER/21202

INTRODUCTION

Business intelligence has moved from aggregate reporting to predictive and prescriptive models trained on granular data about individuals: customers' transactions, employees' performance, patients' records, users' behaviour. Such models are valuable precisely because they capture individual-level patterns, and for the same reason they are a privacy risk. Model parameters and predictions can reveal whether a particular individual was in the training data and, in some cases, what their record contained [1], [2]. Data protection regulation, contractual obligations and customer trust all constrain how such data may be used.

Privacy-preserving machine learning addresses the risk formally. Differential privacy [3] guarantees that the output of a computation changes little whether or not any single record is included, with a privacy budget ϵ quantifying "little". Federated learning [4] trains models across data holders without pooling the raw data, and can be combined with DP so that the shared model updates are themselves protected. Both techniques trade privacy for utility: DP adds noise that degrades accuracy, and federated training constrains optimisation and amplifies the noise cost when protection is applied locally.

This paper examines the trade-off for a representative BI model, explains its structure, and derives guidance for organisations deciding how to protect analytical uses of sensitive data. It does so in a simulation in which the data-generating process is known so that utility can be measured against a known ceiling.

BACKGROUND

Privacy risks of trained models

Membership-inference attacks determine whether a record was used in training from the model's outputs [1]; model-inversion and reconstruction attacks recover attributes of training records [2].

Corresponding Author: Otaxanova Umida, Teacher of Finance, Alfraganus University, Tashkent, Uzbekistan, e-mail: u.otaxanova@afu.uz

How to cite this article: Umida, O. (2026). Privacy-Preserving Machine Learning for Secure Business Intelligence. *Journal of Teacher Education and Research*, 21(2):17-21.

Source of support: Nil

Conflict of interest: None

Overfitted models and models trained on small datasets are most vulnerable. Aggregation does not by itself protect individuals, since repeated queries or models can be combined to isolate records.

Differential privacy

A mechanism is (ϵ, δ) -differentially private if its output distribution changes by at most a factor of e^ϵ (plus δ) when one record is changed [3]. Smaller ϵ means stronger protection and more noise. DP can be applied to query outputs, to model parameters (output perturbation) or to the training process (gradient perturbation) [5]. The privacy budget is consumed by each release and must be managed across all uses of a dataset.

Federated learning

In federated learning, data remain with their holders, who compute local model updates that a coordinator aggregates [4]. Raw data are never pooled, which addresses data-residency and contractual constraints, but updates can still leak information, so DP is typically applied to them. Applying DP locally, at each holder, protects against an untrusted coordinator but adds noise at each site that does not cancel when sites are few; secure aggregation [6] lets the coordinator see only the sum of updates, which allows noise to be added once, centrally.

BI use cases

BI models differ in the precision they require and the sensitivity of their data. Aggregate dashboards tolerate noise; customer-level scoring does not. Internal analytics on employee data carry high sensitivity; product usage analytics on pseudonymised data carry less. The appropriate privacy mechanism depends on this combination.

Alternatives and their limits

Organisations commonly rely on pseudonymisation, aggregation and access control to protect analytical data. Each is useful but none is a guarantee. Pseudonymised records can be re-identified by linkage; aggregates can be differenced to isolate individuals; access controls protect data at rest but not the models and reports derived from it, which leave the controlled environment. Differential privacy differs in kind: its guarantee is a property of the output, holds regardless of what auxiliary information an adversary has, and composes predictably across releases. It is the only mechanism among these that allows an organisation to state what a published model or report could reveal about any individual.

MODEL AND SIMULATION DESIGN

A BI regression task is simulated with 4,000 records, six standardised features and an outcome generated by a linear model with four informative features and Gaussian noise, so that the attainable R^2 is bounded. Records are split 3,000/1,000 for training and evaluation. Three regimes are compared:

No protection: $\hat{y} = x^T \beta$; Central DP: $\hat{y} = x^T (\beta + \eta)$, $\eta \sim \text{Laplace}(0, \Delta/\epsilon)$; Federated DP: $\hat{y} = x^T (1/m) \sum_c (\beta_c + \eta_c)$. (3)

Central DP fits ridge regression on the pooled training data and perturbs the coefficients with Laplace noise scaled to a sensitivity proxy Δ proportional to the inverse of the training size. Federated DP partitions the training data across $m = 5$ clients of 600 records, fits a local ridge model at each, perturbs each with Laplace noise scaled to the local sensitivity, and averages. Privacy budgets $\epsilon \in \{0.1, 0.25, 0.5, 1, 2, 5, 10\}$ are evaluated, with 20 noise draws per setting. The noise calibration is a design choice intended to represent output-perturbation DP; the harness examines the shape of the utility curve, not the exact ϵ at which a given utility is reached.

RESULTS

Utility versus privacy budget

Table 1 and Figure 1 report hold-out R^2 . Without protection the model attains 0.692. Central DP retains 0.692 at $\epsilon = 10$ and 0.691 at $\epsilon = 2$, 0.688 at $\epsilon = 1$, 0.680 at $\epsilon = 0.5$, 0.638 at $\epsilon = 0.25$ and 0.380 at $\epsilon = 0.1$. Federated DP tracks central DP closely at $\epsilon \geq 2$ but diverges at tighter budgets: 0.675 at $\epsilon = 1$, 0.619 at $\epsilon = 0.5$, 0.421 at $\epsilon = 0.25$, and negative at $\epsilon = 0.1$, where the noise exceeds the signal.

Why federated protection costs more

Under central DP, noise is calibrated to the sensitivity of a model fitted on 3,000 records; under federated DP, each client's noise is calibrated to a model fitted on 600 records, so the per-client noise is five times larger, and averaging five noisy models reduces it only by a factor of $\sqrt{5}$. The net noise is about 2.2 times that of central DP at the same ϵ , which is the gap visible at tight budgets. The gap narrows as clients grow larger or more numerous and vanishes when protection is applied centrally after secure aggregation.

Table 1: Hold-out R^2 of the BI regression model by privacy regime and budget (no-protection $R^2 = 0.692$).

Privacy budget ϵ	Central DP	Federated DP (5 clients)
10	0.692	0.692
5	0.692	0.691
2	0.691	0.686
1	0.688	0.675
0.5	0.680	0.619
0.25	0.638	0.421
0.1	0.380	< 0

Interpretation for BI

For this task, a privacy budget of $\epsilon = 1$ costs less than one percentage point of R^2 under either regime, and $\epsilon = 0.5$ costs one point centrally and seven points federated. Budgets below 0.25 cost a large share of the model's value. The curve's shape, flat at loose budgets and steep at tight ones, is general: there is a region in which strong protection is nearly free and a threshold below which it is expensive.

Data size and the position of the curve

The threshold below which protection becomes expensive moves with the size of the training data, because DP noise is calibrated to sensitivity, which falls as records are added. A dataset ten times larger would shift the curve in Figure 1 leftward, making tight budgets affordable; a dataset ten times smaller would shift it rightward. This dependence has an organisational implication: pooling data across units, where permitted, improves not only model accuracy but the privacy that can be afforded at a given accuracy, which is one argument for central over federated architectures where the governance obstacles to pooling can be overcome.

ARCHITECTURE AND GOVERNANCE FOR PRIVACY-PRESERVING BI

Match mechanism to use case

For aggregate reporting and dashboards, DP on query outputs at modest ϵ is sufficient and inexpensive. For predictive models on pooled sensitive data, central DP during training is the default. For models across organisational or jurisdictional boundaries where pooling is prohibited, federated learning with secure aggregation and central DP on the aggregate is preferred to per-client DP, because it recovers most of the utility gap. Per-client DP is reserved for settings where the coordinator itself is untrusted.




Figure 1: Hold-out R^2 against privacy budget for centralised and federated training with differential privacy.



Manage the privacy budget

The privacy budget is a finite resource consumed by every release from a dataset. Organisations should allocate it across uses, record consumption in a privacy ledger, and refuse releases that would exhaust it. BI teams accustomed to unlimited querying will find this constraining; the ledger makes the constraint visible and forces prioritisation of high-value analyses.

Architecture

A workable architecture places a privacy layer between the data store and the BI and modelling tools. The layer enforces DP on query outputs, applies DP training procedures for models, mediates federated training, maintains the ledger, and exposes utility estimates so that analysts can see the cost of a given budget before committing it. The simulation's utility curve is the kind of estimate the layer should provide.

Governance

Privacy-preserving BI requires decisions that are policy rather than engineering: what ϵ is acceptable for what data, who may spend the budget, and how utility losses are weighed against protection. These decisions belong with data governance bodies, informed by the utility curves and by regulatory and contractual constraints. Formal privacy guarantees also support accountability, since the organisation can state what any release could reveal about any individual.

Communicating the trade-off

Analysts and decision-makers need to see the trade-off in the terms they use. Utility should be expressed as the metric the use case cares about, such as forecast error or classification accuracy, at each candidate privacy budget, and the privacy guarantee should be expressed in terms of what an adversary could learn. Presenting both on the same curve, as in Figure 1, lets governance bodies choose a budget with an understanding of what they are buying and what they are giving up.

Incremental adoption

Organisations can adopt privacy-preserving BI incrementally. Begin with DP on aggregate releases, where the utility cost is negligible and the governance habits can be formed; extend to central DP training for models on pooled sensitive data; and introduce federated learning only where data cannot be pooled. At each stage, measure utility on the organisation's own tasks, since the curve's position varies with data size and model type even though its shape is general.

Regulatory alignment

Data protection regimes increasingly recognise differential privacy as a form of anonymisation or as a safeguard that reduces the regulatory burden of processing. Organisations that adopt it obtain, in addition to the technical guarantee, a documented basis for arguing that analytical outputs do not constitute personal data. The privacy ledger and the recorded ϵ per release supply the evidence; the governance body that sets the budget supplies the accountability.

LIMITATIONS

The simulation uses a linear model, output perturbation and a sensitivity proxy; gradient-perturbation methods and deep models have different utility curves, and the ϵ values reported are not

comparable to those of production DP libraries without careful calibration. Federated training is simulated with equal, independent clients and a single aggregation round. Membership-inference risk is not measured directly. All reported values are simulation outputs.

CONCLUSION

Privacy-preserving machine learning lets organisations extract business intelligence from sensitive data under formal guarantees, at a utility cost that is small for moderate privacy budgets and large for tight ones. In simulation, central differential privacy retained almost all of a BI model's accuracy at budgets near one, while per-client federated protection cost more because noise does not average away across few sites. Organisations should match mechanisms to use cases, treat the privacy budget as a managed resource, and locate the trade-off decisions in data governance.

REFERENCES

- Mahadasu, S. S. (2024). AI-native wireless digital twins for precision agriculture: Edge intelligence and predictive farm operations in 6G ecosystems. *International Journal of Data Science and IoT Management System*, 3(1), 65–72. <https://doi.org/10.64751/ijdim.2024.v3.n1.1287>
- Bhaskara Raju Rallabandi. (2024). DESIGNING NTN–TERRESTRIAL 6G EXPERIMENTS FOR AGRICULTURE AND SMART GRIDS. *American Journal of Management and IOT Medical Computing*, 3(3), 27–32. <https://doi.org/10.64751/ajmimc.2024.v3.n3.pp27-32>
- 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>
- Kaur, M., AlZubi, A. A., Jain, A., Singh, D., Yadav, V., & Alkhayyat, A. (2023). DSCNet: deep skip connections-based dense network for ALL diagnosis using peripheral blood smear images. *Diagnostics*, 13(17), 2752.
- Anchala, S. (2019). Database security in e-governance systems: Protecting citizen records, ensuring transparency, and managing public trust in cloud-based platforms. *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, 2(11), 2268–2276. <https://doi.org/10.15680/IJMRSET.2019.0211014>
- Mahajan, R., Kanabar, K., Dey, R., Jadhav, P. V., Khan, M., & Ghantasala, G. S. P. (2026). Exploring big tech vulnerabilities: Transparent blockchain harvesting. In A. K. Singh, S. Saxena, & A. K. Haghi (Eds.), *Cybernetic shield: Securing the future of machine intelligence* (pp. 167–174). CRC Press. <https://doi.org/10.1201/9781003567585-15>
- Mahadasu, S. S. (2024). Closed-loop sustainability orchestration for energy-aware RAN, edge, and data-center workloads in 5G-Advanced networks. *International Journal of Engineering Research and Science & Technology*, 20(2), 1392–1398. <https://doi.org/10.62643/ijerst.2024.v20.n2.4407>
- Jain, Arpit, Adesh Kumar, and Sanjeev Sharma. "Comparative design and analysis of mesh, torus and ring NoC." *Procedia Computer Science* 48 (2015): 330–337.
- Chatrath, A. (2020). The role of policy-as-code in DevSecOps: Automating compliance, governance, and risk management in cloud-native and hybrid IT environments. *International Journal of Innovative Research in Computer and Communication Engineering*, 8(1). <https://doi.org/10.15680/IJIRCE.2020.0801020>
- Khan, S. (2018). The role of zero-trust models in database security: Eliminating implicit trust and enforcing continuous verification in enterprise data access systems. *International Journal of Research in Electronics and Computer Engineering*, 6(1), 1622–1628.
- 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>
- Jain, Arpit, Nageswara Rao Moparathi, 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.
- Bhanushali, B. (2020). Effective sanctions screening and false positive identification. *International Journal of Innovative Research in Science, Engineering and Technology*, 9(4), 8306–8313. <https://doi.org/10.15680/IJRSET.2020.0904179>
- Jadhav, P. V., Krishnan, S. B., Dey, R., & Chakrabarti, P. (2026). Combining deep learning and statistical models: LSTM-ARIMA ensemble for financial forecasting. *IAENG International Journal of Computer Science*, 53(1), 236–246. https://www.iaeng.org/IJCS/issues_v53/issue_1/IJCS_53_1_24.pdf
- 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>
- Devulapalli, N. (2023). Database security and big data analytics: Protecting large-scale data warehouses, ensuring query privacy, and mitigating emerging threats in real-time analytics systems. *International Journal on Recent and Innovation Trends in Computing and Communication*, 11(7), 622–629. <https://www.ijritcc.org/index.php/ijritcc/article/view/11939>
- Jain, A. (2024, December). Energy-Efficient Signal Processing Architectures for IoT Networks Design. In *2024 International Conference on Emerging Technologies and Innovation for Sustainability (EmergIN)* (pp. 362–367). IEEE.
- Bhanushali, B. (2020). The industry needs a validated model to justify the immense cost of migrating from manual, periodic KYC (know your customer) to automated. *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, 3(11), 1799–1805. <https://doi.org/10.15680/IJRSET.2020.0311009>
- 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>
- Mahadasu, S. S. (2022). Cross-domain service assurance for private 5G networks using streaming telemetry and dependency-aware event correlation. *International Journal of Data Science and IoT Management System*, 1(3), 51–56. <https://doi.org/10.64751/ijdim.2022.v1.n3.1290>
- Aggarwal, A. (2026). Adoption of end-to-end encryption in cloud communication channels for ensuring data confidentiality and trustworthiness through secure protocol enforcement. *Journal of Critical Reviews*, 7(12), 5493–5499. <https://doi.org/10.69980/jcr.v7i12.13407>
- 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.
- 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>
- 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/IJRCE.2020.0808100>
- Anchala, S. (2019). The evolution of DevSecOps in software supply chain security: Addressing dependencies, open-source vulnerabilities, and the prevention of software supply chain attacks. *International Journal of Multidisciplinary and Scientific Emerging Research*, 7(1), 415–422. <https://doi.org/10.15662/IJMSERH.2019.0701047>
- Bhaskara Raju Rallabandi. (2023). CLOSED-LOOP DIGITAL-TWIN MANAGEMENT FOR HYBRID 5G-NTN NETWORKS. *American Journal of Management and IoT Medical Computing*, 2(3), 30–35. <https://doi.org/10.64751/ajmimc.2023.v2.n3.pp30-35>
- 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/IJRSET.2020.0308009>
- Mahadasu, S. S., & Rallabandi, B. R. (2023). Federated learning for privacy-preserving RAN optimization across multi-site enterprise 5G networks. *International Journal of Engineering Science and Advanced Technology*, 23(4), 53–57. <https://doi.org/10.64771/ijesat.2023.v23.i04.pp3847>
- 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/IJRSET.2024.0707293>
- Kanabar, K., Dey, R., Mahajan, R., Jadhav, P. V., Khan, M., & Karimullah, S. (2026). Comprehensive guide to blockchain technology: Ecosystems, implementation, governance, and innovations. In A. K. Singh, S. Saxena, & A. K. Haghi (Eds.), *Cybernetic shield: Securing the future of machine intelligence* (pp. 175–182). CRC Press. <https://doi.org/10.1201/9781003567585-16>
- Dwivedi, D. (2023). The technical and ethical challenges of building safe agentic AI systems: Balancing autonomy, predictability, alignment, and human oversight in next-generation AI models. *International Journal on Recent and Innovation Trends in Computing and Communication*, 11(11), 2024–2031. <https://www.ijritcc.org/index.php/ijritcc/article/view/11928>
- 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.
- 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/IJRSET.2021.0401014>
- Kumar, S., Jain, A., & Moparathi, N. R. (2022). Enhanced method of object tracing using extended Kalman filter via binary search algorithm.
- Gupta, A. (2024). Cloud security in the era of big data: Protecting large-scale storage systems, ensuring query privacy, and safeguarding data lakes against emerging threats. *International Journal of Innovative Research in Science, Engineering and Technology*, 13(11). <https://doi.org/10.15680/IJRSET.2024.131100217>
- Devulapalli, N. (2024). The role of homomorphic encryption in database security: Enabling confidential query processing, secure data sharing, and privacy-preserving analytics. *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, 7(8), 14319–14326. <https://doi.org/10.15680/IJRSET.2024.0708082>
- 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>
- Chatrath, A. (2020). DevSecOps for financial services: An in-depth study of secure CI/CD practices, regulatory compliance, and real-time threat monitoring in cloud-based FinTech applications. *International Journal of Innovative Research in Science, Engineering and Technology*, 9(11). <https://doi.org/10.15680/IJRSET.2020.0911098>
- Bhaskara Raju Rallabandi. (2022). SYNC-ADAPTIVE RESOURCE SCHEDULING IN PRIVATE 5G FOR TIME-CRITICAL APPLICATIONS. *International Journal of Data Science and IoT Management System*, 1(2), 40–45. <https://doi.org/10.64751/ijdim.2022.v1.n2.pp40-45>
- 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>
- 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>
- Aggarwal, A. (2019). Application of secure Virtual Private Cloud (VPC) architectures in public cloud infrastructure for strengthening network isolation and segmenting trust zones through fine-tuned routing. *International Journal of Multidisciplinary and Scientific Emerging Research*, 7(1), 423–431. <https://doi.org/10.15662/IJMSEERH.2019.0701048>
- Jain, J., & Jain, A. (2022). Securing e-Healthcare Images Using an Efficient Image Encryption Model. *Scientific Programming*, 2022(1), 6438331.
- Anchala, S. (2020). Adoption of secure access service edge (SASE) in distributed enterprises for ensuring cloud application protection and network optimization through unified security frameworks. *Turkish Online Journal of Qualitative Inquiry*, 11(4), 2964–2973. <https://doi.org/10.69980/60m26411>
- 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.
- 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.
- 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>
- 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.
- Dwivedi, D. (2023). Integration of secure file upload mechanisms in user-facing applications for preventing malware injections and file-based exploits through content-type validation. *International Journal on Recent and Innovation Trends in Computing and Communication*, 11(6), 768–775. <https://www.ijritcc.org/index.php/ijritcc/article/view/11940>
- Ahuja, R. (2023). The role of artificial intelligence and machine learning in DevSecOps: Leveraging predictive analytics, automated threat detection, and anomaly identification for secure software delivery pipelines. *International Journal on Recent and Innovation Trends in Computing and Communication*, 11(2), 381–389. <https://www.ijritcc.org/index.php/ijritcc/article/view/11946>
- 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>
- Harshitha, G., Kumar, S., Rani, S., & Jain, A. (2021, November). Cotton disease detection based on deep learning techniques. In *4th Smart Cities Symposium (SCS 2021)* (Vol. 2021, pp. 496–501). IET.
- Ahuja, R. (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>