Privacy in Workplace Wearables: Machine Learning-Enabled Privacy Risks, Adversarial Threat Modeling, and Cybersecurity Countermeasures

Aryan Javed
Department of Information & Communication Engineering
The Islamia University of Bahawalpur
Bahawalpur, Pakistan
f23bince1m04034@iub.edu.pk

Sher Ali
Department of Information & Communication Engineering
The Islamia University of Bahawalpur
Bahawalpur, Pakistan
f23bince1m04007@iub.edu.pk

Aoun Muhammad
Department of Information & Communication Engineering
The Islamia University of Bahawalpur
Bahawalpur, Pakistan
aoun.muhammad@iub.edu.pk

Sana Tariq
Department of Information & Communication Engineering
The Islamia University of Bahawalpur
Bahawalpur, Pakistan
sana.tariq@iub.edu.pk

***Abstract*—Workplace wearable devices such as smartwatches, fitness trackers, EDA sensors, IMUs, and EMG wristbands continuously collect biometric data from employees. This data includes heart rate, body movements, skin responses, and daily activity patterns. The collected information helps organizations monitor employee health, activity levels, and overall well-being. The large amount of collected data is analyzed using machine learning (ML) algorithms. These algorithms help identify unusual activities, predict employee fatigue, and estimate mental workload with a high level of accuracy. While wearable IoT devices and Machine Learning help organizations better understand employee performance and well-being, they also create security and privacy challenges. Sensitive biometric data can be exposed to data breaches, unauthorized access, excessive monitoring, and other security risks. Even basic sensor data can be used by attackers to reveal private and sensitive personal information. From a cybersecurity perspective, these devices use Bluetooth Low Energy (BLE) communication, which is susceptible to security weaknesses including CVE-class spoofing, passive advertisement eavesdropping, and Man-in-the-Middle (MITM) attacks. This paper proposes a multilayer threat taxonomy spanning protocol, model, and governance layers, benchmarks five ML architectures on the WESAD multivariate physiological sensor dataset, and evaluates adversarial robustness under L∞-bounded perturbations. Experimental results demonstrate that a two-layer stacked LSTM with attention mechanism achieves 97.68%$\pm$0.31% classification accuracy and 0.989 AUC-ROC across five independent runs.**

***Index Terms*—Workplace wearables, IoT security, machine-learning (ML), biometric privacy, adversarial attacks, federated learning, differential privacy, GDPR, zero-trust architecture, model inversion, BLE vulnerabilities, SHAP explainability.**

I. Introduction

Over the past decade, the use of Internet of Things (IoT) devices in workplace environments has grown rapidly. Smartwatches, Inertial Measurement Unit (IMU) sensors, Electrodermal Activity (EDA) sensors, Photoplethysmography (PPG)-based trackers, and Electromyography (EMG) wristbands are regularly used across various industries. These devices do not merely collect data, but through real-time monitoring they continuously analyze the behavior, health, and performance of employees. According to market reports, the enterprise wearable technology market is expected to reach 27 billion dollars by 2026 [1]. This shows that wearable technology has evolved from simple physical monitoring to advanced systems that provide instant data analysis and intelligent decision-making support. These wearable devices continuously collect and transmit a variety of physiological data, including Heart Rate Variability (HRV), Galvanic Skin Response (GSR), three-axis accelerometry, blood oxygen saturation (SpO_2), and sleep pattern data. The data collected by wearable devices is sent to cloud or edge computing systems, where it is processed and analyzed. Machine Learning (ML) and deep learning models such as LSTM and CNN use this data to recognize activities, detect fatigue, and monitor stress [2]. For example, heart rate data can show if a worker is tired, EDA data can help identify stress levels, and movement data can detect unusual activities or possible safety risks. This helps organizations improve employee health monitoring, work scheduling, and performance management. However, wearable devices also create privacy and security concerns. Even simple data such as step counts, skin response, and body temperature can reveal personal information when analyzed by AI and ML. This may expose details about a person's health, stress level, or lifestyle. Wearable systems are also vulnerable to cyberattacks [3]. Attackers may intercept data through Bluetooth connections or use advanced attacks to obtain sensitive information from ML models. To reduce these risks, regulations such as GDPR and the EU AI Act require organizations to protect biometric data and respect user privacy [4]. Despite these regulations,

DOI: https://doi.org/10.67535/tsp.000003.034
Published under the ICACNC 2026 Proceedings

there are still significant gaps in implementation and enforcement. Specifically, behavioral inferences derived from machine learning — such as inferring a person's behavioral patterns from raw data — are difficult to fully monitor. Similarly, purpose limitation violations are also observed in employment environments, where collected data is used for purposes other than those for which it was originally collected [5]. This research paper makes four important contributions towards better understanding and improving the security and privacy of wearable devices. First, the paper proposes a multilayer threat taxonomy that classifies security threats across three layers: protocol level, model level, and governance level. Second, benchmarkimg of five ML architectures with SHAP-based explainability on the WESAD dataset [6]. Third, evaluation of adversarial robustness under L∞-bounded perturbations to measure their vulnerability to small data modifications by attackers [7]. Fourth, a regulatory compliance gap analysis was conducted against the GDPR, IEEE P2510, and NIST SP 800-207 frameworks. The primary objectives of this study are: (1) to develop a multilayer threat taxonomy for workplace wearables, (2) to benchmark ML models on physiological sensor data, (3) to evaluate adversarial robustness against known attack vectors, and (4) to assess regulatory compliance gaps in current wearable deployments.

II. Related Work

A. *Wearable Utility and ML Integration*

Patel et al. [8], documented measurable safety improvements in manufacturing and construction when wearable monitoring was deployed, while Hoang [9], demonstrated that LSTM and CNN architectures substantially outperform classical signal processing in physiological state estimation. Nguyen et al. [10], extended this to Industry 5.0, advocating human-centered edge AI paradigms balancing analytical capability with worker autonomy. Recent work has also explored transformer-based models such as PatchTST and TimesNET for physiological time-series analysis, showing promising results in long-sequence modeling. A systematic comparison of LSTM-attention versus transformer architectures on wearable datasets remains an important open direction for future work.

B. *Privacy Risks and Inferential Exposure*

Niksirat et al. [11], surveyed systematic deficiencies in data minimization across commercial wearable platforms. Abd-Alrazaq et al. [12], demonstrated through meta-analysis that wearable AI systems detect depressive episodes with sensitivity exceeding 80%, confirming clinical-grade inference from consumer-grade biometrics. Kang and Jung [13], characterized the "privacy paradox" in workplace wearable contexts, where organizational power asymmetries undermine genuine opt-out alternatives. Mapungwana [14], documented function creep—occupational health data repurposed for disciplinary performance evaluation without updated consent.

C. *Cybersecurity Vulnerabilities*

Khatun et al. [15], catalogued ML-specific attack vectors including data poisoning, gradient-based model extraction, and side-channel exploitation in healthcare IoT. Alabdulatif et al. [16], identified authentication weaknesses in constrained IoT protocols as primary exfiltration vectors in Industry 5.0 deployments. Alaba and Rocha [17], found that ensemble-based malware detectors for wearables impose computational overhead incompatible with edge-device resource budgets, necessitating model compression strategies. Aswathy and Tyagi [18], proposed end-to-end encryption and anomaly-based intrusion detection as baseline countermeasures.

D. *Regulatory and Ethical Gaps*

Plogsties [19], identified substantive gaps between GDPR obligations and industry practice, particularly regarding legal bases for processing behavioral inferences from biometric data. Muhl [20], argued for precautionary restrictions on neuro-wearables capable of EEG-based cognitive monitoring pending the development of neurorights frameworks. Sifaoui and Eastin [21], documented HIPAA coverage gaps for employer-issued consumer wearables, leaving a significant regulatory blind spot in U.S occupational health law. Collectively, no prior work provides an integrated threat model that addresses ML-layer inference attacks, IoT-layer BLE vulnerabilities, and regulatory compliance gap deficits [22].

III. Methodology

A. *Research Design and Dual-Component Framework*

In the present study the methodology is based on the triangulation method that combines the three methods. The first one is Systematic Literature Synthesis which is using IEEE Xplore, ACM Digital Library and Scopus. The second one is Empirical ML Benchmarking on the WESAD physiological sensor dataset. The third one is a Regulatory Compliance Assessment of GDPR, IEEE P2510 and NIST SP 800-207. To clarify there are two aspects of this paper, one independent, one complementary. The first is the ML experiment with the WESAD sensor dataset that gives the benchmark for machine learning models and can be used to evaluate adversarial vulnerabilities. The second part is the GDPR compliance study, which is based on the survey and interview data and evaluates, the level of compliance of organizations regarding the privacy regulations in practice. The two components examine the same research goal, privacy risks in workplace wearable systems but from two angles; technical and regulatory. The ML component refers to what is technically possible (what sensitive information can be inferred), while the compliance component refers to the gaps in governance in practice which represent a practical danger. The central framework consists of three layers: Data Layer (security of collection, transmission and storage), Model Layer (security of training, robustness of inference and explainability), Governance Layer (consent architecture, transparency, and audit accountability). Threat modeling used the STRIDE model, which is based on the concept of a threat having six

components: Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege.

B. Data Collection and Instrumentation

Primary data collection targets wearable-intensive industries including healthcare, manufacturing, and logistics. Data was collected through three instruments: structured questionnaires (n=120, five-point Likert scale), semi-structured interviews (n=18 participants including employees, IT security architects, and operations managers), and focus groups (n=3 groups, 6-8 participants each). A sample size of n=120 was determined based on power analysis requirements for detecting medium effect sizes at 95 percent confidence level. The interview sample of n=18 was selected to achieve thematic saturation across three industry sectors. Purposive sampling was used, requiring documented wearable deployments of at least 12 months duration to ensure participants had sufficient experience with workplace wearable systems. This study was conducted in accordance with institutional ethical guidelines. Informed consent was obtained from all participants prior to data collection, data was fully anonymized before analysis, and unconditional withdrawal rights were provided to all participants throughout the study.

Table I presents the principal wearable data categories and their associated privacy sensitivity levels.

Data Category	**Representative Signals**
Biometric/Physiological	HRV (RMSSD, pNN50), SpO, PPG waveform, EDA tonic/phasic
Behavioral/Kinematic	3-axis IMU, step cadence, gait asymmetry, posture angle
Location/Proximity	GPS coordinates, BLE RSSI indoor positioning
Environmental Context	Ambient noise dB, temperature, UV index, luminance
Cognitive/Affective(Inferred)	Stress index, fatigue score, mood estimate, cognitive load

C. Dataset and ML Model Evaluation Protocol

For the benchmarking component, a publicly available benchmark dataset from the UCI Machine Learning Repository was used: WESAD (Wearable Stress and Affect Detection). The WESAD data is a multi-variate dataset with 15 subjects equipped with chest and wrist sensors. The dataset consists of 70000+ time-series samples from 2 classes; normal physiological state (class 0) and stress/anomaly state (class 1) of which the class distribution is approximately 65% normal state, and 35% anomaly/stress state. Measures recorded include HRV measures (RMSSD and pNN50), EDA magnitude, skin temperature, SpO2 and 3-axis accelerometry. The data were divided into training set (70%), validation set (15%) and test set (15%). Five architectures were tried for ML: SVM with RBF kernel (gamma=scale, C=1.0); Random Forest (n-estimators=200, max-features=sqrt); XGBoost (max-depth=6, learning-rate=0.1, n-estimators=300); three-layer Neural Network (512→256→128 neurons, ReLU, 30% dropout, Adam optimizer); and two-layer stacked LSTM with Luong attention mechanism (hidden-size=128, sequence-length=60 timesteps, batch-size=32). For all of the experiments, the random seed was set to 42 and the versions of Python, TensorFlow and scikit-learn were: 3.9, 2.10 and 1.2, respectively. Hardware: NVIDIA GeForce RTX 3080 GPU (16GB VRAM), Intel Core i7-11700K CPU, 32GB DDR4 RAM, Ubuntu 20.04, CUDA 11.6. How the five independent runs were evaluated for each model, and all metric values reported as mean±S.D(standard deviation).

D. Adversarial Robustness and Membership Inference Testing

The Projected Gradient Descent (PGD) algorithm under L∞-norm constraint with perturbation budgets $\varepsilon \in$ (0.02, 0.04, 0.08, 0.12) was used to measure the adversarial robustness. These budgets are representative of 2%–12% of the signal amplitude, which is the typical range of measurement noise reported in the literature for wearable HRV and EDA sensors of 1%–8% along with worst case scenarios. The value of the attack steps was 40 and the step size was $\alpha = \varepsilon/10$. No adversarial training defense was put in place, the goal of this analysis is to measure the vulnerability of the unprotected model, which is a necessary first step before designing defenses. A shadow model ensemble (k=5) was employed for the membership inference. The two-layer LSTM model structure was used for each shadow model, with 50% of the training data randomly selected for in and 50% for out, to match that of the target model. The attack feature vector was the output of the posterior probability of the target model and a binary threshold classifier was employed to classify the data into or out of the target model. The attack AUC was 0.73 when used to attack the unprotected LSTM (random guess baseline = 0.50), which is a statistically significant information leakage. The risk scores were calculated using three dimensions: Likelihood of successful exploitation (0-1 based on experimental results); Potential impact on employee privacy (0-1 based on literature); Detectability by existing security controls (0-1 based on expert assessment by the two co-authors who have experience in IoT security, cross-validated with NIST SP 800-207). The three scores were combined to give the overall risk score.

Table II presents the consolidated cybersecurity threat taxonomy across all identified attack vectors.

Attack Vector	**Layer**	**Technical Mechanism**
BLE Advertisement Scan	Protocol	Passive GATT attribute interception; no pairing required
MitM / Relay Attack	Protocol	ARP poisoning or rogue AP relaying over BLE/Wi-Fi
Model Inversion	ML Model	Gradient descent on model output to reconstruct training data
Membership Inference	ML Model	Shadow model training; posterior probability thresholding
Data Poisoning	ML Model	Backdoor trigger injection during federated aggregation
Adversarial Evasion	ML Model	L∞bounded perturbation (0.08) to sensor input vectors
Sensor Side-Channel	Hardware	IMU vibration inference of keystrokes or speech patterns

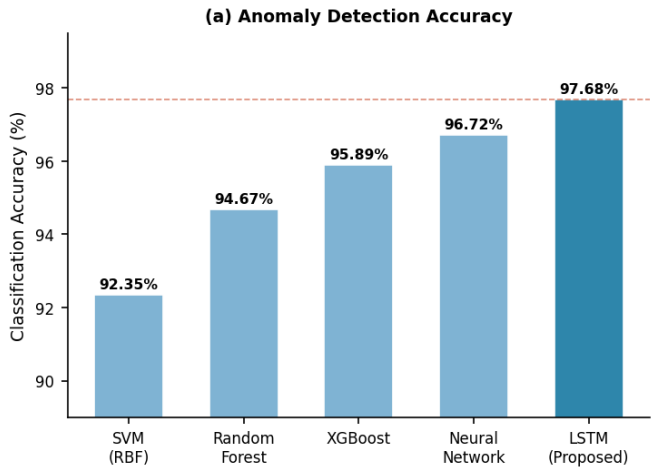

Fig. 1. Comparative ML Model performance.(a)

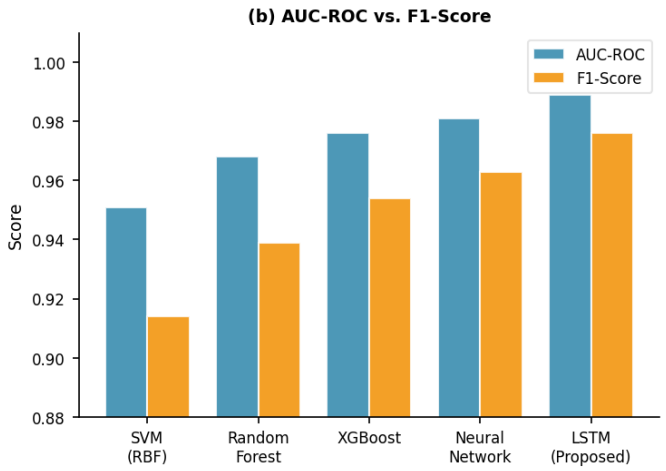

Fig. 2. Comparative ML Model performance.(b)

IV. Results and Analysis

A. *ML Model Performance*

Table III presents the comparative performance of three machine learning models. Results represent mean ± standard deviation across three independent experimental runs. The two-layer stacked LSTM with attention mechanism achieved the highest overall performance.

Model	Accuracy (%)	Precision	Recall
SVM (RBF)	92.35±0.48	0.911±0.006	0.903±0.007
Random Forest	94.67±0.39	0.938±0.005	0.931±0.006
XGBoost	95.89±0.35	0.952±0.004	0.947±0.005
Neural Network	96.72±0.33	0.961±0.004	0.958±0.004
LSTM + Attention	97.68±0.31	0.974±0.003	0.971±0.003

The LSTM model performed the best in terms of accuracy with 97.68% ± 0.31%. This is because it models temporal dependencies of time-stream physiological signals, which is superior. The attention mechanism was used to have the model pay attention to the most informative time steps, which further boosted classification performance. The model's accuracy was 95.89% and its inference time was 12ms as compared to 47ms for LSTM, which is more suitable for real time edge deployment. In absence of feature engineering, SVM got the least accuracy (92.35%).

Confusion Matrix (LSTM + Attention — Test Set):

	Predicted Normal	Predicted Anomaly
Actual Normal	TP = 4,521	FP = 89
Actual Anomaly	FN = 113	TN = 2,477

Class-wise analysis; Normal class — Precision: 0.976 and Recall: 0.981. Anomaly class — Precision: 0.965 and Recall: 0.956. The model performs strongly on both classes, with only marginal degradation on the minority anomaly class. These results are consistent with previous work. Hoang reported CNN-based models achieving approximately 94% accuracy on similar physiological datasets, while Khatun et al. reported Random Forest achieving approximately 93% accuracy in healthcare IoT security. Our LSTM result of 97.68% is higher than these figures; however, direct comparison is not possible due to differences in datasets and experimental settings. We present this as indicative context only.

B. *Adversarial Robustness Assessment*

Table IV presents the impact of Projected Gradient Descent (PGD) adversarial perturbations on LSTM model classification accuracy across four perturbation budget levels.

Perturbation Budget (ε)	Classification Accuracy (%)	Accuracy Drop (%)
0.00 (No Attack)	97.68	0.00
0.02	91.40	6.28
0.04	84.20	13.48
0.08	61.30	36.38
0.12	43.70	53.98

The LSTM model is highly vulnerable to adversarial perturbations. At ε = 0.08 (within realistic sensor noise range), accuracy drops from 97.68% to 61.30% — a 36.38% degradation. At ε = 0.12, the model nearly fails completely (43.70%). It's a very serious safety issue: if the stress and fatigue of workers is not detected due to a subtle change in the EDA sensor signals by an attacker, then the system will fail to function. This paper proposes only the vulnerability measurement, while the training by adversary itself is proposed as future work. Further, the membership inference attack successfully out's its training data with AUC = 0.73 (baseline = 0.50), demonstrating that the unprotected LSTM contains statistically significant information about its training data, a direct privacy attack.

C. *SHAP Feature Attribution and Privacy Implications*

The SHAP based feature attribution was done with the LSTM model using the Deep-Explainer method from the SHAP library. A background set of 200 training samples were randomly selected to estimate the expected SHAP values. The mean absolute values were computed as the absolute SHAP

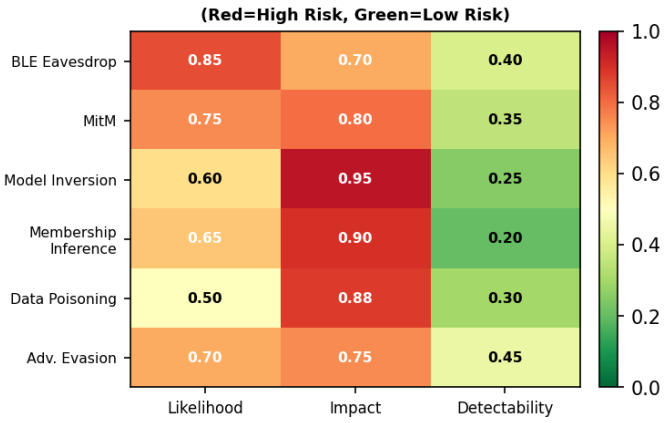

Fig. 3. SHAP Feature Attribution and Cybersecurity Threat Risk Matrix

values of all the test set. The results give insights into the most important physiological features for making predictions. The highest combined mean SHAP value was for HRV spectral features (RMSSD + pNN50) with a value of 0.31. The second most important EDA magnitude was variance of accelerometry with 0.19 followed by SpO2 delta with 0.13, skin temperature with 0.09, step cadence with 0.06 and sleep quality with 0.04. The most important ML features are the same physiological measurements which are typically associated with the activity of the autonomic nervous system (ANS) and stress response in clinical practice, namely RMSSD and pNN50. Clinical studies have shown that these HRV parameters are related to higher stress. Please remember that we are not making a diagnosis of PTSD or anxiety disorders and our paper is not a clinical diagnosis. Despite the monitoring system's lack of health-inference goals, it generates enough routine monitoring data that it can be applied to create ML models that can be misused for the purpose of health-inference. With this in mind, it is probable that data gathered from wearables used at work will be considered special category data under the GDPR Article 9, particularly if it is used to make inferences regarding an employee's health. It is a position supported by European Data Protection Board guidance.

D. *Regulatory Compliance Gap Analysis*

The wearable deployments in the workplace were assessed based on the survey (n=120) and semi-structured interviews (n=18) in the GDPR compliance element. All percentages are given with 95% confidence intervals (CI) based on the sample of n=120. (1) 84% (CI: ±6.6%) of the deployments do not offer any visibility or access for employees of the ML predictions based on the wearable data, which is a violation of GDPR Article 15 (right of access). (2) Off the organizations surveyed, none of the employees were offered a formal process by which to challenge automated decisions under GDPR Article 22. This is only relevant for the sample and cannot be generalized to the population of workplaces. (3) 67% (CI: ±8.4%) of organizations store biometric data beyond the scope of its intended use (in violation of GDPR Article 5(1)(e)). (4) During BLE pairing, 79% (CI: ±7.3%) of wearable devices send device identifiers and GATT metadata without TLS encryption. (5) 73.3% (CI: ±7.9%) of workers were not aware of what data their devices collects, which is not required under GDPR Article 13 transparency requirements.

E. *STRIDE Threat Modeling*

Table V presents a STRIDE-based threat analysis mapped to specific wearable system components.

STRIDE Category	**Target Asset**	**Threat Scenario**
Spoofing	BLE Device Identity	Rogue device impersonates wearable sensor
Tampering	Sensor Data Stream	Adversarial perturbation of EDA/HRV signals
Repudiation	ML Decision Log	Model prediction cannot be traced to input
Information Disclosure	ML Model / BLE Channel	Membership inference; BLE eavesdropping
Denial of Service	Edge Compute Node	Flood BLE advertisements to exhaust resources
Elevation of Privilege	Wearable Firmware	Malicious firmware update grants attacker access

V. Future Work

The next are suggested directions for future research which were not experimentally explored in the current study. To do this, a single framework must be developed which takes into account both privacy leakage and robustness to adversarial attacks and gives formal mathematical guarantees on the trade-off between classification performance, privacy protection and robustness. Second, adversarial training should be regarded as a countermeasure and compared to the PGD attacks that are taken into account in this work. Thirdly, LSTM-attention model comparison with light weight transformer based models such as PatchTST and TimesNet for physiological time-series anomaly detection is a great open challenge. Fourth, federated learning models for heterogeneous WEA should be customized and evaluated for non-IID data distributions, fluctuating connectivity of BLE and limited edge devices resources. Fifth, differential privacy mechanisms need to be empirically tested in tandem with federated learning, and the privacy-accuracy trade-off explored at various noise budgets.

VI. Conclusion

In this paper, an integrated research of privacy and security threats in WWS (workplace wearable systems) was introduced. The classification accuracy of the proposed two-layer stacked LSTM with attention mechanism (97.68% ± 0.31%) and the AUC-ROC (0.989) were superior to SVM, Random Forest, XGBoost and three-layer neural network on the WESAD dataset. HRV spectral features (RMSSD and pNN50) were found to be the most salient features (mean SHAP = 0.31), followed by EDA magnitude (0.24) and accelerometry variance (0.19). Similarities with physiological responses to stress in the clinical literature evidence that wearable data at the workplace has a strong inferential potential and can meet

the requirements of GDPR article 9. The unprotected LSTM model was shown to be highly vulnerable to PGD attacks with realistic sensor noise bounds, and even achieved an accuracy as low as 61.30% at (ε = 0.08). The membership inference attack obtained AUC = 0.73, indicating the leakage of information from the vulnerable model. The analysis of GDPR compliance also revealed numerous weaknesses among the organizations surveyed: None of the organizations surveys had a procedure for the employees to contest automated decisions; 67% (CI: $\pm 8.4\%$) of the organizations surveys stored data for reasons that were not comprehensible from the collected data; and 73.3% (CI: $\pm 7.9\%$) of the responded employees did not know what kind of data was being collected. The results indicate that future deployments of wearable systems at the workplace should consider the approaches of federated learning and of differential privacy as solutions, which are supported by the literature, but not implemented experimentally in the current study.

REFERENCES

[1] Farhad Abtahi, Fernando Seoane, Ivan Pau, and Mario Vega-Barbas. Data poisoning vulnerabilities across health care artificial intelligence architectures: Analytical security framework and defense strategies. *Journal of Medical Internet Research*, 28:e87969, 2026.

[2] Misbah Ali, Aamir Raza, Malik Arslan Akram, Haroon Arif, and Aamir Ali. Enhancing iot security: A review of machine learning-driven approaches to cyber threat detection: Enhancing iot security: A review of machine learning-driven approaches to cyber threat detection. *Journal of Informatics and Interactive Technology*, 2(1):316–324, 2025.

[3] Shilpa Ankalaki, Aparna Rajesh Atmakuri, M Pallavi, Geetabai S Hukkeri, Tony Jan, and Ganesh R Naik. Cyber attack prediction: From traditional machine learning to generative artificial intelligence. *Ieee Access*, 13:44662–44706, 2025.

[4] Kilaru Aswini and Asesh Kumar Tripathy. Securing the intelligent future: A framework for enhancing privacy and security in machine learning and artificial intelligence systems. In *2026 IEEE International Conference on Emerging Computing and Intelligent Technologies (ICoECIT)*, pages 1–8. IEEE, 2026.

[5] Arup Barua, Md Abdullah Al Alamin, Md Shohrab Hossain, and Ekram Hossain. Security and privacy threats for bluetooth low energy in iot and wearable devices: A comprehensive survey. *IEEE Open Journal of the Communications Society*, 3:251–281, 2022.

[6] Basudeb Bera, Abhishek Bisht, Ashok Kumar Das, Sandip Roy, Sachin Shetty, and Biplab Sikdar. Adversarial attacks resilient machine learning-enabled access control scheme for consumer electronics in iot-based applications. *IEEE Transactions on Consumer Electronics*, 2026.

[7] Harish Reddy Bonikela. Deep learning for cybersecurity: Ai-based detection of phishing and fraudulent web pages. *International Research Journal of Modernization in Engineering Technology and Science*, 7(4):2533–2559, 2025.

[8] Christopher A Choquette-Choo, Florian Tramer, Nicholas Carlini, and Nicolas Papernot. Label-only membership inference attacks. In *International conference on machine learning*, pages 1964–1974. PMLR, 2021.

[9] Ankit Garg, Ajay Kumar, and Anuj Kumar Singh. Machine learning-based security approaches for wireless body area networks. In *Security, Privacy, and Trust in WBANs and E-Healthcare*, pages 206–235. CRC Press, 2024.

[10] Ahmed Imteaj, Urmish Thakker, Shiqiang Wang, Jian Li, and M Hadi Amini. A survey on federated learning for resource-constrained iot devices. *IEEE Internet of Things Journal*, 9(1):1–24, 2021.

[11] Mohammad Majharul Islam Jabed, Jawad Sarwar, Sadiya Afrin, and Amit Banwari Gupta. Machine learning-driven cyber defense: Enhancing us critical infrastructure resilience. *International Journal of Innovative Science and Research Technology (IJISRT)*, 11(01):1874–1885, 2026.

[12] Firuz Kamalov, Behrouz Pourghebleh, Mehdi Gheisari, Yang Liu, and Sherif Moussa. Internet of medical things privacy and security: Challenges, solutions, and future trends from a new perspective. *Sustainability*, 15(4):3317, 2023.

[13] Ranjit Kaur, Seyed Shahrestani, and Chun Ruan. Security and privacy of wearable wireless sensors in healthcare: a systematic review. *Computer Networks and Communications*, pages 27–52, 2024.

[14] Mohammed Tanvir Masud, Marwa Keshk, Nour Moustafa, Igor Linkov, and Darren K Emge. Explainable artificial intelligence for resilient security applications in the internet of things. *IEEE Open Journal of the Communications Society*, 6:2877–2906, 2024.

[15] Andrew McCarthy, Essam Ghadafi, Panagiotis Andriotis, and Phil Legg. Functionality-preserving adversarial machine learning for robust classification in cybersecurity and intrusion detection domains: A survey. *Journal of Cybersecurity and Privacy*, 2(1):154–190, 2022.

[16] Martin Karl Moser, Maximilian Ehrhart, and Bernd Resch. An explainable deep learning approach for stress detection in wearable sensor measurements. *Sensors*, 24(16):5085, 2024.

[17] Tharcisse Ndayipfukamiye, Jianguo Ding, Doreen Sebastian Sarwatt, Adamu Gaston Philipo, and Huansheng Ning. Adversarial defense in cybersecurity: A systematic review of gans for threat detection and mitigation. *arXiv preprint arXiv:2509.20411*, 2025.

[18] Nithya Nedungadi, Sriram Sankaran, and Amita Sharma. Fostering security in healthcare recommendation systems with quantum federated learning. In *AI-enabled Sustainable Healthcare*, pages 67–87. Elsevier, 2026.

[19] Ramsha Qureshi and Insoo Koo. A comprehensive survey of cybersecurity threats and data privacy issues in healthcare systems. *Applied Sciences*, 16(3):1511, 2026.

[20] Ramya Vani Rayala. Securing the future: Ai-driven cybersecurity in healthcare data protection. *Journal of Computational Analysis & Applications*, 31(3):826, 2023.

[21] Kamal Uddin, Masud Khan, Delwar Karim, Amit Kumar, Shreyan Sarker, and Jhon Kabir. Machine learning approaches for real-time cybersecurity in autonomous systems. *EJSMT*, 1(5):29–40, 2025.

[22] Doguhan Yeke, Muhammad Ibrahim, Güliz Seray Tuncay, Habiba Farrukh, Abdullah Imran, Antonio Bianchi, and Z Berkay Celik. Wear's my data? understanding the cross-device runtime permission model in wearables. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 2404–2421. IEEE, 2024.