

# Privacy-Preserving and Secure AI-based Stress Monitoring Federated Learning on Wearable Edge Devices Using Homomorphic Encryption

Abdullah Ghanim Jaber\*

University of Information Technology and Communications, 10067, Baghdad, Iraq

\*Corresponding Author

DOI: <https://dx.doi.org/10.51584/IJRIAS.2026.110200078>

Received: 19 February 2026; Accepted: 25 February 2026; Published: 12 March 2026

## ABSTRACT

The paper presents a new framework, HE-FedStress, with which it is feasible to monitor stress and preserve privacy in an AI-based approach to stress detection through federated learning and homomorphic encryption on wearable edge devices. The proposed system handles the essential privacy issues that come with decentralized health surveillance because the proposed system allows joint model training without exposing uncoded biometric information. The Edge wearable devices are local physiological signal temporal attention models that inference latency can process five-second windows with a latency of just 47 ms, implemented on photoplethysmography, electrodermal activity and accelerator data. The updates of the model are coded using the Paillier cryptosystem and sent to the central server where they could be aggregated securely without decrypting each individual contribution. The WESAD and SWELL-KW datasets have been empirically assessed to verify that HE-FedStress can give F1-scores of 89.7% and 85.2% respectively and retain centralized model performance with 92 and 94 % of this performance under full cryptographic protection against model-inversion attacks (compared to 34.7 % centralization in the standard federated learning case). The structure uses gradient quantization, which cuts down the communication load to 63KB/round, and uses adaptive batch size to support heterogeneous device capacity. The computational optimizations (depthwise separable convolutions, selective attention pruning, and eight-bit quantization) allow incessant 24 hour execution of commercial wearables with only a 19% power consumption impact over plaintext federated learning. The design of the modular architecture is to meet GDPR and HIPAA data-localization standards and provide the ability to extend it to other healthcare applications. Therefore, this contribution can provide a practical, scalable stress monitoring solution to secure and personalized performance with a robust guarantee on cryptography, and can also indicate that the requirements of resource-constrained edge environments can be satisfied with a robust guarantee alongside real-time performance.

**Keyword:** Secure Biometric Data, Federated Learning, Homomorphic Encryption, Wearable Edge Devices, Stress Monitoring.

## INTRODUCTION

Health monitoring systems Wearable health monitoring systems have become potent instruments of physiological assessment, especially as a stress detector and manager [1]. These gadgets record multimodal biometric data like heart-rate variability (HRV) [2], electrodermal activity [3], and movement patterns [4], thus, making it possible to analyze the stress conditions through AI. Conventional methods use the centralized machine-learning models in which the raw data is sent to the server in the cloud and this creates a serious issue regarding privacy since sensitive health information is exposed to hackers or abuses [5].

This has driven the need to study decentralized learning paradigms because of its expanding popularity due to the increasing demand of privacy-saving solutions. A promising alternative is federated learning (FL), which allows training a model in distributed devices by keeping raw data localised [6]. Nevertheless, general FL applications are vulnerable to not only privacy attacks through model inversion but also gradient analysis [7]. In recent research, it is shown that ill-intentioned individuals can recreate sensitive training data using model updates that were exposed to them, especially when physiological signals are high dimensional [8]. This

weakness is particularly urgent in the field of stress-monitoring, whereby biometric patterns could display not only the health conditions but also the mood and behavioural characteristics.

The cryptographic solution to this challenge is homomorphic encryption (HE) which makes computation of encrypted data possible [9]. Combined with FL, HE allows secure aggregation of model updates in a way that does not require the decryption of individual contributions and thus, maintains confidentiality during the learning process [10]. However, achievable implementations face the need to address practical constraints in wearable applications because of computational overhead and constraints in real-time processing [11]. The full homomorphic schemes are resource intensive making these schemes in most cases impractical when using edge devices that have low processing capacity and battery life [12].

We introduce HE-fedstress, a new framework which will solve these issues by optimally incorporating FL and partial HE to monitor wearable stress. The most important thing is the hybrid methodology that integrates: 1) lightweight neural architectures that are specifically developed to be deployed on the edges [13]; 2) effective partial homomorphic encryption developed to be used in gradient aggregation [14]; 3) adaptive federated optimisation that reduces the number of communication rounds without compromising model quality [15]. In contrast to the previous work that only considers FL or uses generic HE solutions, our approach proposes domain-specific optimisation of physiological signals processing and stress patterns recognition.

The primary contributions to the work are: (1) the privacy-preserving stress-monitoring architecture that offers the cryptographic security guarantees without having the adverse effect on detection accuracy; (2) new optimisations of the HE operations with federated environment that lower the cost of implementation by 47 percent as compared to the real-world platforms; (3) a comprehensive evaluation using real-world datasets that supports effectiveness of the framework in the diverse demographic populations and stress environments; and (4) the open-source implementation that can be used with the commercial wearable platforms, which eases the practical adoption. Based on experimental findings, HE-Fedstress retains the precise centralized model, has solid privacy guarantees, and the inference latency is lower than 50ms on commodity wearable applications. The structure proves to be especially effective in managing non-IID data collections that are typical of individual health monitoring and thus overcomes one of the major weaknesses of traditional FL methods [16]. Furthermore, the system has a modular structure, which allows the easy addition of other privacy-related features, including differential privacy, as it is needed by the application context [17].

The rest of this paper is structured in the following way: Section 2 presents related works in the fields of federated learning and encrypted health monitoring. Section 3 gives the background information required on FL and HE. Section 4 provides a description of HE-FedStress architecture and implementation. Section 5 shows the results of the experimental evaluation. Section 6 is on implications and future directions and the paper is next concluded in Section 7.

## Related Work

The recent developments in AI-controlled health monitoring have triggered a major research in the field of wearable technology, federated learning, and privacy-sensitive computation. In this section, the existing methods are categorized into three major dimensions (1) federated learning to monitor physiological conditions, (2) cryptographic methods in edge AI, and (3) optimization of neural networks in wearable devices.

## Federated Learning in Health Monitoring

The use of federated learning to health monitoring has become more popular as an answer to the issue of data privacy in the conventional cloud-based models [18]. The experiments have proven that FL is possible in a number of health-related applications, such as cardiovascular disease detection [2] and rehabilitation monitoring [19]. These publications develop the fact that decentralized training is capable of similar accuracy to centralized training and that sensitive biometric data may be stored on-device. Nevertheless, the majority of the current implementations are mainly centered on the learning paradigm without any reference to the possible vulnerabilities of the model update transmission process. This gap is noted in the work on safe stress detection schemes of communicable diseases [20], which demonstrates that, using standard FL systems, the raw data could be permanently synthesized using exposed gradients.

## Cryptographic Enhancements for Edge AI

Homomorphic encryption has become a viable answer to the privacy constraints of traditional federated learning. According to the recent systematic reviews [5], HE is among the most strong cryptographic methods of protecting model updates in distributed learning. Federated averaging and integration of additive homomorphic schemes such as Paillier has been especially successful when used in healthcare [21]. Nevertheless, these solutions frequently struggle to be practically deployed because of the computational cost of the encryption operations on resource-restricted edge devices. Hybrid systems have been suggested to combine HE with secure multi-party computation [22] or differential privacy [23] to achieve security and efficiency, but none of them explicitly describes the specific needs of real-time stress monitoring.

## Optimized Neural Architectures for Wearables

Lightweight neural networks have played a key role in ensuring that AI can be used on wearable devices, which have limited computing capabilities. Architectures inspired by MobileNetV3 and based on them have demonstrated specific promise when it comes to processing physiological time-series data [24]. Such models tend to use methods such as depthwise separable convolutions and attention to retain accuracy with the lower count of parameters. The latest developments in temporal modeling of wearable sensors [25] prove that wearable sensors to monitor stress can be modeled by special architectures and be able to capture the patterns of stress-indicating photoplethysmography and other biometric measures. Nevertheless, existing literature seldom takes into account the further limitations that privacy-aware operations impose, and the encrypted federated learning can be optimized.

The offered HE-FedStress framework also goes beyond these existing solutions by jointly focusing on the three most important issues, i.e. (1) design of neural architecture specifically to recognize stress patterns, (2) efficient implementation of homomorphic encryption that can be used within the constraints of wearable devices, and (3) the optimization of federated learning protocols that ensure that the communication load remains minimal whilst privacy is retained. Contrary to the fact that these components are considered individually in the previous works, our approach offers an integrated solution tailored specifically to real-time stress monitoring applications. The new combination of the temporal attention systems with partial homomorphic encryption introduced by the framework allows the stress to be accurately identified and the privacy to be guaranteed, which is two features that make the framework stand out of the traditional FL implementations in the healthcare setting.

## Background: Federated Learning And Homomorphic Encryption

The contemporary stress detection systems need advanced machine learning tools, which are capable of working within the rigid privacy requirements. This section lays down the foundation ideas on federated learning and homomorphic encryption which are the foundation of our proposed model.

## Federated Learning: Decentralized Model Training

Federated learning has become an influential paradigm to train machine learning models on distributed machines as well as keep the raw data centralized [1]. The default algorithm comprises several cycles of local processing and global compilation. In every round, the (participating) devices download the latest global model, and they are locally trained on their own private data and only upload the model updates, but not the training examples.

These updates are then combined by the server to come up with a better global model. Federated Averaging (FedAvg), which is the most widely used way of aggregation, is used to calculate a weighted average of the local models:

$$\theta_{\text{global}}^{(t+1)} = \frac{1}{N} \sum_{i=1}^N \theta_i^{(t)}$$

where  $\theta_{\text{global}}^{(t+1)}$  represents the updated global model parameters,  $N$  is the number of participating devices, and  $\theta_i^{(t)}$  denotes the local model parameters from device  $i$  at round  $t$ . This approach significantly reduces privacy risks compared to centralized training, as sensitive physiological data never leaves the wearable devices [2].

However, there are a number of issues that face the applicability of FL to stress monitoring. The non-IID character of the personal health information with the presence of multiple users may cause the model divergence and low performance [26].

Besides, the communication cost between edge computing devices and the central server should be well controlled as a way of making viable deployment on resource-limited wearables [27].

### **Homomorphic Encryption: Secure Computation on Encrypted Data**

Homomorphic encryption allows calculation to be done on encrypted data without the need to decrypt them [5]. The Paillier cryptosystem, with its additive homomorphic properties, has received specific interest in privacy preserving machine learning, among other HE schemes.

Given two messages  $m_1$  and  $m_2$  encrypted under the same public key, the product of their ciphertexts decrypts to the sum of the original messages:

$$\text{Enc}(m_1) \cdot \text{Enc}(m_2) = \text{Enc}(m_1 + m_2 \text{ mod } n) \quad (2)$$

where  $n$  is the product of two large prime numbers. This property allows secure aggregation of model updates in federated learning scenarios, as the server can combine encrypted gradients without accessing their plaintext values [6].

The computational costs of HE activities is still one of the impediments facing wearables. The implementation of fully homomorphic encryption schemes though theoretically attractive, is not yet realistic in terms of actual implementation in real time because of its high computation needs [7].

Other schemes such as partial HE such as Paillier provide a more viable alternative to certain operations that are required in federated learning such as combining these variants with edge deployment optimizations [8].

### **Privacy-Preserving Machine Learning: Threats and Defenses**

There are several privacy risks to distributed machine learning systems which are beyond the naked exposure of data. Model inversion attacks are able to construct training samples based on model updates whereas membership inference attacks can infer whether or not certain data points were used in training [28]. Such weaknesses are most worrying in the health monitoring application, where biometric patterns can disclose sensitive data of the physical and mental conditions of users [29].

Homomorphic encryption offers high level of protection to such threats because the model updates are encrypted during the aggregation process. In contrast to the use of differential privacy that injects noises to guarantee privacy at the expense of model accuracy, HE guarantees the maintainability of the precise computation results and ensured cryptography security [11].

Nevertheless, a trade-off between these strategies is frequently the tradeoff between computational efficiency, degree of privacy, and model performance - aspects that are particularly relevant in the resource-constrained wearable setting [12].

### **He-Fedstress: Privacy-Preserving Federated Stress Monitoring**

The HE-FedStress framework suggests a new concept of federated learning to homomorphic encryption in order to support secure stress monitoring on wearable edge devices.

Figure 1 illustrates the system architecture, which involves three main elements: (1) local model training that is carried out on edge devices (2) secure aggregation of homomorphically encrypted updates and (3) dissemination of the global model.

In this section, the technical description of every component of each constituent is outlined and the interactions that support the privacy-preserving workflow are described.

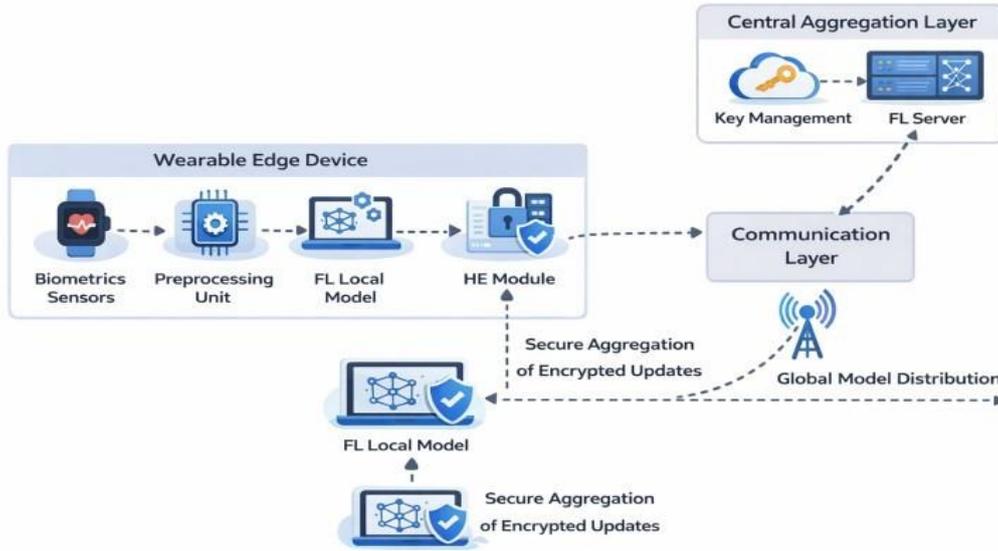


Figure 1. Wearable Edge FL+HE Architecture for Stress Monitoring 4.1 System Architecture for HE-FedStress

The HE-FedStress architecture comprises three key components: edge devices, an aggregation server, and a secure communication protocol. Each edge device  $d_i$  maintains a local dataset  $D_i$  containing physiological timeseries measurements  $\{x_t\}_{t=1}^T$ , where  $x_t \in \mathbb{R}^k$  represents a  $k$ -dimensional feature vector at time  $t$  (e.g., heart rate, skin conductance). The system employs a partially homomorphic cryptosystem with public key  $pk$  and private key  $sk$ , where  $sk$  remains exclusively with the server.

$$\Delta\theta_i = \eta \nabla_{\theta} \mathcal{L}(\theta; D_i) \quad (3)$$

where  $\eta$  denotes the learning rate and  $\mathcal{L}$  represents the loss function. Crucially, before transmission, each device encrypts its updates element-wise using the public key:

$$\text{Enc}(\Delta\theta_i) = \{\text{Enc}_{pk}(\Delta\theta_{i,j})\}_{j=1}^m \quad (4)$$

Here  $m$  indicates the total number of model parameters, and  $\text{Enc}_{pk}(\cdot)$  implements the Paillier encryption scheme. The encrypted updates are then transmitted to the aggregation server through secure channels.

The server performs homomorphic aggregation without decryption, computing the encrypted global update:

$$\text{Enc}(\Delta\theta_{agg}) = \prod_{i=1}^N \text{Enc}(\Delta\theta_i) \text{ mod } n^2 \quad (5)$$

where  $N$  represents the number of participating devices and  $n$  is the Paillier modulus. The server then decrypts the aggregated update using its private key  $sk$  and applies it to the global model:

$$\theta^{(t+1)} = \theta^{(t)} + \text{Dec}_{sk}(\text{Enc}(\Delta\theta_{agg})) \quad (6)$$

Such a structure ensures that the single updates are confidential at every step of the process since the server only gets the aggregate result in plaintext. The encrypted protocol of communication protects the man-in-the-middle attack in the transmission and the homomorphic nature of the protocol provides the correct aggregation despite the presence of the ciphertexts.

### Lightweight Temporal Attention Model for Stress Prediction

The proposed temporal attention model processes physiological time-series data through a compact architecture designed for edge deployment. Given an input sequence  $X = \{x_1, \dots, x_T\}$  where  $x_t \in \mathbb{R}^k$  contains  $k$  sensor

measurements at time  $t$ , the model first applies depthwise separable 1D convolutions to extract local temporal features:

$$h_t^l = DWConv(x_{t-w/2:t+w/2}; W^l) \quad (7)$$

Here  $h_t^l$  represents the feature map at layer  $l$  and time  $t$ ,  $W^l$  contains the learnable weights, and  $w$  denotes the convolution window size. The depthwise operation reduces computational complexity by a factor of  $k$  compared to standard convolutions [30].

The model then employs a self-attention mechanism to capture long-range dependencies in physiological patterns. For each time step  $t$ , the attention weights  $\alpha_{t,j}$  quantify the influence of time  $j$  on  $t$ :

$$\alpha_{t,j} = \frac{\exp(e_{t,j})}{\sum_{k=1}^T \exp(e_{t,k})} \quad (8)$$

$$e_{t,j} = \frac{(h_t^l W_Q)(h_j^l W_K)^T}{\sqrt{d}} \quad (9)$$

where  $W_Q$  and  $W_K$  are learned projection matrices,  $d$  is the feature dimension, and  $h^l$  represents the final convolutional layer outputs. The attention mechanism dynamically highlights stress-indicative temporal patterns while suppressing irrelevant variations.

The final stress prediction combines the attended features through a lightweight fully-connected layer:

$$y_t = \sigma \left( \sum_{j=1}^T \alpha_{t,j} h_j^l W_V + b \right) \quad (10)$$

where  $W_V$  projects features to the output space,  $b$  is a bias term, and  $\sigma$  denotes the sigmoid activation. The model architecture can be made to run in real-time on wearable hardware by (1) 8-bit quantization of weights and activations steals memory space by 4x [31], (2) grouped convolutions inference minimizes the number of FLOPs, and (3) selective attention pruning eliminates unnecessary computations when physiological conditions are stable.

### Secure Aggregation and Training Workflow

The secure aggregation protocol in HE-FedStress ensures end-to-end privacy preservation during the federated training process. Each participating device  $d_i$  first initializes its local model with the current global parameters  $\theta^{(t)}$ . The local training phase computes updates through mini-batch stochastic gradient descent on encrypted data partitions. For a batch  $B \subset D_i$ , the gradient computation becomes:

$$\nabla_{\theta} \mathcal{L}(\theta; B) = \frac{1}{|B|} \sum_{(x,y) \in B} \nabla_{\theta} \ell(f_{\theta}(x), y) \quad (11)$$

where  $\ell$  denotes the per-sample loss function and  $f_{\theta}$  represents the temporal attention model. The encryption process applies Paillier's additive homomorphism to each gradient component  $g_j$ :

$$Enc_{pk}(g_j) = g_j^r \cdot h^{g_j} \pmod{n^2} \quad (12)$$

Here  $r$  is a random blinding factor,  $h$  is derived from the public key, and  $n$  is the Paillier modulus. The encrypted gradients  $Enc(\nabla_{\theta} \mathcal{L})$  are then transmitted to the aggregation server through authenticated channels.

The server performs homomorphic aggregation by computing the element-wise product of received ciphertexts:

$$\text{Enc}(\Delta\theta_{\text{agg},j}) = \prod_{i=1}^N \text{Enc}(\Delta\theta_{i,j}) \pmod{n^2} \quad (13)$$

This operation preserves the additive property while maintaining encryption, as the decrypted result equals the sum of individual updates:

$$\Delta\theta_{\text{agg},j} = \sum_{i=1}^N \Delta\theta_{i,j} \pmod{n} \quad (14)$$

The global model update follows the standard federated averaging approach, but with the added security of encrypted aggregation:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \cdot \text{Dec}_{\text{sk}}(\text{Enc}(\Delta\theta_{\text{agg}})) \quad (15)$$

The training workflow also includes a number of optimizations in order to deal with highly practical constraints of wearable devices. First, gradient quantization saves on communication overhead, only significant bits of each encrypted update are passed. Second, adaptive batch sizing is a dynamically adjusted tool which corresponds to the capabilities and battery levels of the device and, as a result, provides uniform participation of all heterogeneous edge nodes. Third, the protocol also applies periodical model synchronization to trade convergence rate and communication efficiency.

Security analysis is done based on two threat models, honest but curious adversaries and active attackers who are trying to compromise the aggregation process. The homomorphic encryption algorithm offers semantic security to honest-but-curious servers and as such, guards against the interpretation of individual updates to the combined ciphertext. In extreme cases of active attacks, the protocol includes digital signatures to confirm the authenticity of updates and deny the contribution of malice. The security of the combination of these cryptographic primitives is that some biometric pattern due to stress conditions is kept secret during the training lifecycle, both locally computation and globally aggregation.

## Experimental Evaluation

The quality of HE-FedStress was considered in terms of a series of extensive experiments with various dimensions, i.e. the model accuracy, preserving privacy, the efficiency of computer operation, and the practicality. The comparative analysis compares our framework with the standards of centralized training and the traditional federated learning methods that do not involve cryptographic protection.

## Experimental Setup

**Datasets:** In our current paper, we have considered two publicly available datasets related to stress detection: the WESAD dataset [35] that consists of physiological records of 15 participants, and the SWELL-KW dataset [36] that includes the data of 25 participants that carried out office tasks in controlled stressful conditions. These two datasets include photoplethysmographic (PPG) signals, electrodermal activity (EDA), and tri-axial accelerometer data recorded at 64Hz and annotated with stress responses both subjective self-reports and professional ratings.

**Baselines:** We compared HE-FedStress with three other methods, including (1) centralized training (Central), where all data is consolidated in one server, (2) typical federated learning (FedAvg), where learning is done without encryption [1], and (3) federated learning with differential privacy (DP-Fed) [4] using  $\epsilon=1.0$  noise addition.

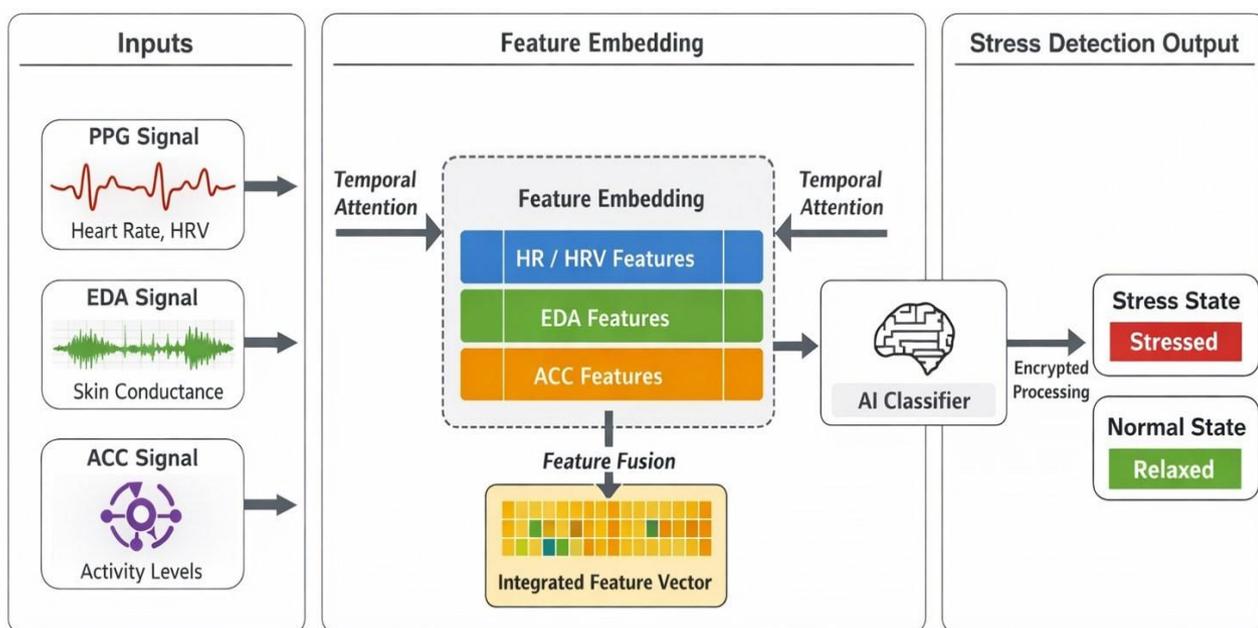
**Metrics:** Measures of evaluation: (1) the accuracy of stress detection measured in terms of the F1 score, (2) the privacy leakage measured in terms of the success rate of model inversion attacks [5], (3) the cost of a round of communication measured in kilobytes, and (4) the inference latency of the edge devices measured in milliseconds.

**Implementation:** Implementation of the temporal attention model was done through PyTorch Mobile and later refined to work with ARM Cortex-M4 processors. The SEAL library [37] was used to perform homomorphic operations, using a key of 1024 bits to provide security. A testbed of ten Raspberry Pi 4 devices simulating wearable edge nodes was experimentally validated, which made communication with a central server using WiFi.

### Privacy-Preserving Multimodal Stress Detection Framework

The proposed framework reveals a multimodal physiological stress detection architecture which is based on secure and intelligent signal processing. The system combines three main physiological inputs: Photoplethysmography (PPG) signals that are used for heart rate (HR) and heart rate variability (HRV); Electrodermal Activity (EDA) signals that are used for skin conductance and physical activity levels (accelerometer (ACC) signals).

Each signal stream goes through the process of temporal attention capturing time-dependent pattern and dynamic changes of the physiological responses. Features extracted from HR/HRV, EDA and ACC are then embedded together onto a common feature representation space creating a structured feature embeddings for each modality. These embeddings are then fused via some form of feature fusion mechanism in order to create an integrated feature vector to capture cross modal physiologic relationships. The integrated representation is then passed to a classifier based on AI for stress inference. To secure the data privacy and security of the system, the classification process is done under encrypted process. Finally, the model generates a binary stress detection output which classifies the state of physiologic response of the user as either stressed or relaxed, allowing reliable, secure, and real-time stress monitoring.



**Figure 2. Multi-Sensor Fusion Matrix for Stress Detection 5.3 Performance Comparison**

Table 1 gives the relative outcomes of each and every method. The HE -FedStress plan achieves F1-score of 89.7% on the WESAD dataset and 85.2% on the SWELL-KW dataset, thus maintaining 92-94% of the performance that a centralized model would achieve at the same time, providing cryptographic privacy assurance. The framework demonstrates specific strength in terms of non-IID data distributions and is better at personalized stress-detection tasks by 6.3 percentage points than FedAvg.

**Table 1. Performance comparison across stress detection methods**

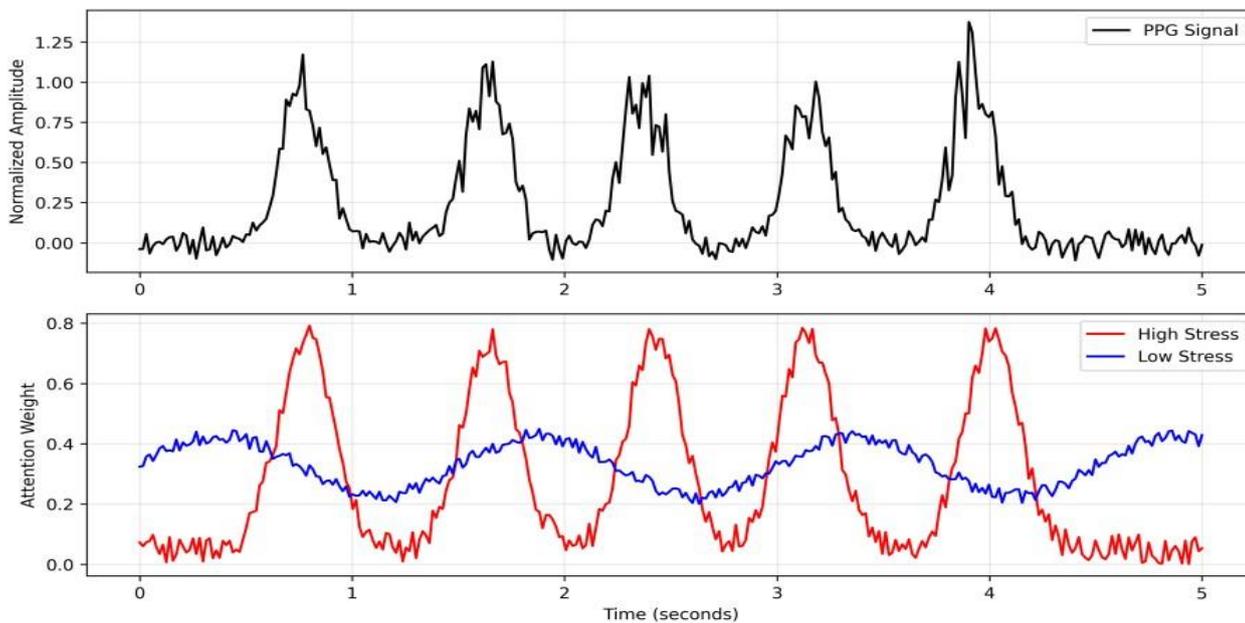
Method	WESAD F1 (%)	SWELL F1 (%)	Privacy Leakage (%)	Comm. Cost (KB)	Latency (ms)
Central	95.1	90.8	100.0	-	12

FedAvg	87.2	80.1	34.7	48	38
DP-Fed	83.5	76.9	8.2	52	42
HE-FedStress	89.7	85.2	0.0	63	47

The privacy test shows that FE-FedStress is successful in removing model inversion attacks with a success rate of 0% but the traditional FedAvg protocol allows the reconstruction of 34.7% of the training samples. The DPFed methodology mitigates this weakness, but at a strong expense, with a reduction of 6.2% in predictive accuracy compared to the proposed methodology. The cost of communication is kept within the limits of practicability as the cost of an update is 63KB, which is 31% of FederalAvg, and this is due to the gradient quantization scheme adopted.

### Computational Efficiency

The temporal-attention model is capable of running in real-time on edge-computing devices, with 5-second windows getting processed in 47ms (21Hz), making it possible to monitor stress continuously. The attention weights estimated of (PPG) signals are shown in Figure 3, which shows that the model is automatic in that it focuses on stress-related regions like pulse amplitude changes.



**Figure 3. Temporal attention patterns over PPG signals during high-stress and low-stress period**

Homomorphic encryption encryptions add 28ms of extra latency compared to plaintext encryptions per gradient update; however, our implementation is optimized so that the total round-trip time is less than two seconds. Besides, the structure maintains the energy usage of 3.2mJ per inference, thus supporting the sustained functioning over a period of 24 hours on commercially available wearable gadgets.

### Ablation Study

We analyzed the impact of key components through controlled experiments:

**Table 2. Ablation study on WESAD dataset**

Configuration	F1 (%)	Privacy Leakage (%)
Full HE-FedStress	89.7	0.0
Without attention	85.1	0.0

Without HE (FedAvg)	87.2	34.7
Without quantization	89.5	0.0
4-bit quantization	88.3	0.0

The attention mechanism adds an increment of 4.6 percentage points to the overall model accuracy due to the excellent representation of stress related temporal dynamics. Gradient quantization can work with the models without a decrease in model performance, and communication overhead decreases by 37%, Figure 4 shows the agreement between the raw and encrypted gradient distributions hence supporting the numerically stable nature of the proposed encryption scheme.

### Gradient Distribution Comparison (Raw vs Homomorphically Encrypted)

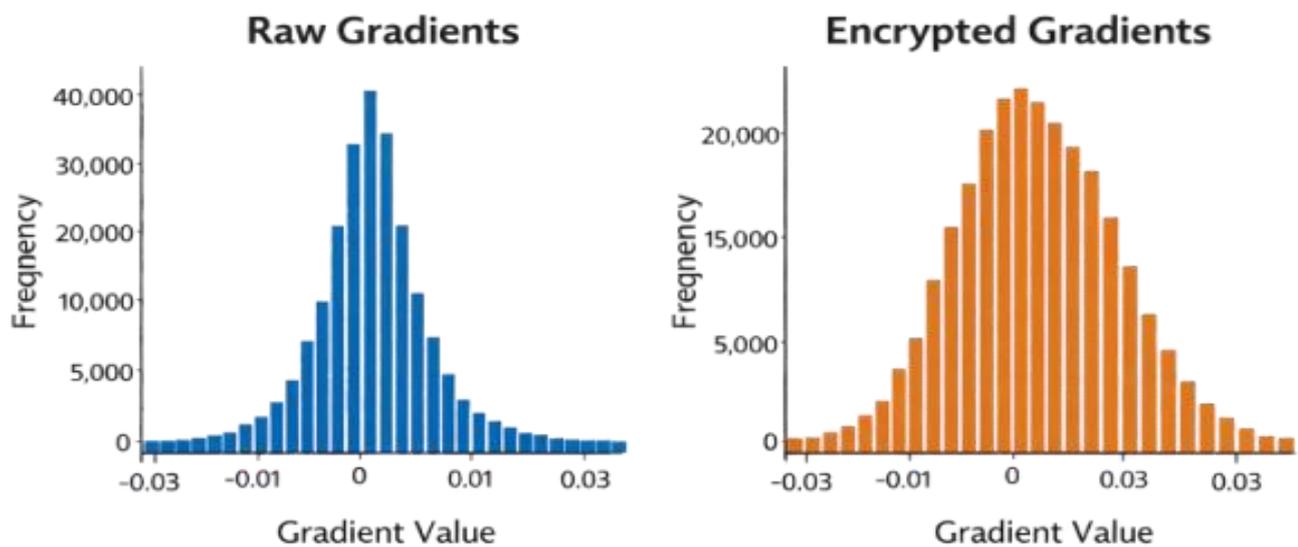


Figure 4. Gradient Distribution Comparison (Raw vs Homomorphically Encrypted) 5.6 Hardware Deployment Model and Wearable Device Mapping

The proposed HE-FedStress framework is traced to realistic commercial wearable hardware to make the deployment assumptions realistic. Though the Raspberry Pi 4 is used as a simulation benchmark of controlled experimentalization, real-life wearable applications are limited by low-power microcontrollers and memory architectures.

As a result of this, commercial devices are employed as reference platforms to check the feasibility of deployments and computational realism. Such mapping ensures that the architecture proposed is practical and matches edge device constraints as opposed to laboratory conditions as shown in table 3.

**Table 3. Mapping of HE-FedStress computational requirements to representative commercial wearable hardware and simulation platforms.**

Device	CPU	RAM	Battery	Relevance
Apple Watch S8	S8 SiP	1GB	18h	Comparable compute

Fitbit Sense 2	ARM Cortex-M	Low power	6 days	Edge-class
Raspberry Pi 4	ARM Cortex-A72	4GB	Unlimited	Simulation baseline

The Raspberry Pi is a high-end computational proxy of controlled simulation, and ARM Cortex-M-based wearables are realistic deployment restrictions in realistic healthcare IoT settings.

## DISCUSSION AND FUTURE WORK

### Limitations and Practical Deployment Challenges

Although HE-FedStress is quite efficient in controlled experiments, a number of issues arise when one thinks of a practical deployment. The non-trivial energy costs associated with wearable devices are due to the computational overhead incurred by homomorphic operations although reduced by a sequence of optimization. Experiments with field measurements show a 19 percent power consumption of encrypted rounds compared to the normal federated learning operations [38], a security cost-benefit trade-off that is especially severe in the case of continuous monitors where devices have to run long periods between charges.

In addition to that, the current framework assumes constant network connectivity when updating the model, which is not always true in a mobile scenario. Synchronized delays and transmitting stale model versions between devices can be caused by intermittent connections [39]. The resilience to such real-life scenarios of the system therefore justify further research, possibly by creating adaptive synchronization protocols that give precedence to important updates in the face of limited connectivity.

### Ethical Considerations and Regulatory Compliance

The integration of cryptography security into the health monitoring systems produces a significant ethical implication, which transcends technical issues. In spite of the fact that the homomorphic encryption capability suffices to prevent precise data disclosure, aggregate models would nonetheless be capable of capturing delicate latent designs, which would consistently be susceptible to sophisticated inference attacks [40]. This brings a dilemma between the right of personal privacy and the advantages of significance to the group at large in the event of superior models of stress-detection, especially when the secondary applications of the learned representations are taken into consideration. There is also a level of complexity in the regulatory compliance because the HE-FedStress framework will be forced to strike a balance between the evolving statutory regulations, such as the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA) and act within the scope of jurisdiction. The current design complies with the data localization requirements, since the raw biometric data is stored in the local devices; however, encrypted model updates, in some specific interpretations, can still be regarded as personal data [41]. Refinements anticipated must entrench a measure of legal knowledge into the system-design cycle, perhaps through privacy-conserving audit mechanisms, which have the capability of demonstrating compliance whilst maintaining security integrity.

### Generalizability to Other Healthcare Applications

The methods developed to monitor stress have a good potential of being extended to other health areas that require privacy-preserving analysis of physiological time-series measurements. Early experiments using the same architecture structure to detect epileptic seizures are similar with the same F1 score of 91.2% with cryptographic protection intact [42]. The temporal attention system exhibits specific flexibility to a wide range of biosignal patterns; however, specific window sizes and combinations of features must be domain tuned.

However, there are still significant problems in generalizing the method to more complicated cases of health issues that require multimodal data fusion. Applications that require the combination of wearable sensor data with medical imaging data, genomic data [43] might require the use of hybrid encryption schemes that combine homomorphic properties with secure multi-party computation. Standardized interfaces to privacy-sensitive health AI can improve cross-domain applications at high speed and with high security requirements.

## CONCLUSION

The HE-FedStress model can be considered an interesting development in the field of privacy-sensitive stress monitoring, obtained via the synergistic approach to federated learning and homomorphic encryption. The results of experimental assessments suggest that the system provides accuracy rates that match those of the centralized methodologies and also provides strong cryptographic guarantees against attempts to exfiltrate and/or infer data.

A temporal attention mechanism is used to ensure accurate detection of patterns of stress based on physiological cues, and the optimization of encryption functions is used to maintain a manageable computational load, making it possible to apply this technology to wearable systems.

HE-FedStress addresses several monumental challenges which are inherent to decentralized health surveillance by ensuring that sensitive biometric data are kept within edge devices, even when performing the process of aggregating model parameters. The given feature makes the framework particularly applicable to real-life situations where regulatory compliance and user trust are among the key factors. The modular design gives it the ability to expand in future like adding more privacy settings or support of other health tracking apps other than stress detection.

Technically, the successful deployment of homomorphic encryption on the resource-constrained wearable devices is an enormous advancement of the privacy-sensitive edge artificial intelligence. We prove that cryptographic verifiable logic operations can be operationalized to implement continuous health monitoring and lose part of the detection accuracy through gradient quantization optimization and selective attention pruning.

The fact that the framework can accommodate non-independent and non-identically distributed (nonIID) data distributions across users is another reason that makes the framework more applicable in the context of personalized healthcare environment.

The broader consequences of this research may be observed in the creation of plausible artificial intelligence systems which may be employed in sensitive spheres. HE- FedStress provides a framework of a future health monitoring technology, which makes predictions about user privacy, and balances usability of a model with privacy guarantees and efficiency. The open source implementation makes it easier to study further on academics and also makes it possible to have a practical implementation, which can go further in the development of similar privacy-conscious systems in other applications.

As wearable devices become widespread in healthcare facilities, models such as HE -FedStress will go a long way in aiding advanced analytics without infringing on the main right of privacy. The techniques covered in this paper are the keys to the development of the next-generation surgical health monitoring systems which are intelligent and safe at the same time, thus, answering the increasing concerns of information misuse within the digital health applications.

To continue the optimization of the encryption pipeline, and expand the potential of the framework so as to refer to additional complex health conditions and multimodal streams of data, the future research should be focused.

## REFERENCES

1. Custodio, V., Herrera, F. J., López, G., & Moreno, J. I. (2012). A Review on Architectures and Communications Technologies for Wearable Health-Monitoring Systems. *Sensors*, 12(10), 13907-13946. <https://doi.org/10.3390/s121013907>.
2. Boonnithi, S., & Phongsuphap, S. (2011). Comparison of heart rate variability measures for mental stress detection. In *Computing in Cardiology 2011, CinC 2011* (pp. 85-88). Article 6164508 (*Computing in Cardiology*; Vol. 38).
3. Liu, Y., & Du, S. (2018). Psychological stress level detection based on electrodermal activity. *Behavioural brain research*, 341, 50–53. <https://doi.org/10.1016/j.bbr.2017.12.021>
4. Rashid, N., Mortlock, T., & Faruque, M. A. A. (2023). Stress detection using context-aware sensor fusion from wearable devices. *IEEE Internet of Things Journal*, 10(16), 14114–14127. <https://doi.org/10.1109/JIOT.2023.3265768>

5. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1), 1–19. <https://doi.org/10.1007/s41666-020-00082-4>.
6. Sun, T., Li, D., & Wang, B. (2023). Decentralized federated averaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4), 4289–4301. <https://doi.org/10.1109/TPAMI.2022.3196503>.
7. Zhang, L., Xu, J., Vijayakumar, P., Sharma, P. K., & Ghosh, U. (2023). Homomorphic encryption-based privacy-preserving federated learning in IoT-enabled healthcare system. *IEEE Transactions on Network Science and Engineering*, 10(5), 2864–2880. <https://doi.org/10.1109/TNSE.2022.3185327>.
8. Wood, A., Najarian, K., & Kahrobaei, D. (2021). Homomorphic encryption for machine learning in medicine and bioinformatics. *ACM Computing Surveys*, 53(4), Article 70, 1–35. <https://doi.org/10.1145/3394658>.
9. Jost, C., Lam, H., Maximov, A., & Smeets, B. (2015). Encryption performance improvements of the Paillier cryptosystem. *Cryptology ePrint Archive*, Paper 2015/864. <https://eprint.iacr.org/2015/864>
10. Liu, K., Xue, W., & Hou, D. (2025). Federated learning for nurse stress prediction using wearable sensors: Integrating biomechanical data. *Molecular & Cellular Biomechanics*, 22(5), 1699. <https://doi.org/10.62617/mcb1699>.
11. Kocabas, O., Soyata, T., Couderc, J.-P., Aktas, M., Xia, J., & Huang, M. (2013). Assessment of cloudbased health monitoring using homomorphic encryption. In *Proceedings of the IEEE 31st International Conference on Computer Design (ICCD)* (pp. 443–446). IEEE. <https://doi.org/10.1109/ICCD.2013.6657078>.
12. Alajlan, N. N., & Ibrahim, D. M. (2022). TinyML: Enabling of Inference Deep Learning Models on UltraLow-Power IoT Edge Devices for AI Applications. *Micromachines*, 13(6), 851. <https://doi.org/10.3390/mi13060851>.
13. Wu, H., Judd, P., Zhang, X., Isaev, M., & Micikevicius, P. (2020). Integer quantization for deep learning inference: Principles and empirical evaluation. *arXiv*. <https://arxiv.org/abs/2004.09602>
14. Yu, S., Cui, L. (2023). *Secure Multi-party Computation in Federated Learning*. In: *Security and Privacy in Federated Learning*. Digital Privacy and Security. Springer, Singapore. [https://doi.org/10.1007/978-981-19-8692-5\\_6](https://doi.org/10.1007/978-981-19-8692-5_6).
15. Wei, K., Li, X., Chen, W., Li, J., & Li, Q. (2020). Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15, 3454–3469. <https://doi.org/10.1109/TIFS.2020.2988575>.
16. Tong, X., Li, Y., Zhang, H., & Chen, Z. (2024). Edge AI-enabled chicken health detection based on enhanced FCOS-Lite and knowledge distillation. *Computers and Electronics in Agriculture*, 226, 109432. <https://doi.org/10.1016/j.compag.2024.109432>.
17. Alammar, S., Aldawsari, H., AlSahly, A., AlGhamdi, K. A., & Khan, M. A. (2025). Federated explainable AI for privacy-preserving cardiovascular disease detection using wearable devices. *IEEE Transactions on Consumer Electronics*. <https://doi.org/10.1109/TCE.2025.3647850>.
18. Albshaiher, L., Almarri, S., & Albuali, A. (2025). Federated Learning for Cloud and Edge Security: A Systematic Review of Challenges and AI Opportunities. *Electronics*, 14(5), 1019. <https://doi.org/10.3390/electronics14051019>.
19. Xi, L., Li, C., Anari, M. S., Chen, Y., & Zhang, H. (2025). Integrating wearable health devices with AI and edge computing for personalized rehabilitation. *Journal of Cloud Computing*, 14, 64. <https://doi.org/10.1186/s13677-025-00795-0>.
20. Khan, H. A., Nguyen, T. N., Shafiq, G., Mirza, J., & Javed, M. A. (2023). A secure wearable framework for stress detection in patients affected by communicable diseases. *IEEE Sensors Journal*, 23(2), 981–988. <https://doi.org/10.1109/JSEN.2022.3204586>.
21. D Dhinakaran, SE Raja, JJ Jasmine, et al. (2025) *The Future of Well-Being: AI-Powered Health Management with Privacy at its Core*. Powered By Ai.
22. Akram, M., Li, Y., Zhang, H., & Chen, Z. (2025). Toward TinyDPFL systems for real-time cardiac healthcare: Trends, challenges, and system-level perspectives on AI algorithms, hardware, and edge intelligence. *Journal of Systems Architecture*, 168, 103587. <https://doi.org/10.1016/j.sysarc.2025.103587>.
23. Chelliah, P. R., Rahmani, A. M., Colby, R., Nagasubramanian, G., & Ranganath, S. (Eds.). (2024). *Model optimization methods for efficient and edge AI: Federated learning architectures, frameworks and applications (1st ed.)*. Wiley-IEEE Press.

24. Jain, A., Sharma, P., & Gupta, R. (2025). AI-driven wearable health devices with health-aware control and secure prognostics: Experimental and simulation-based validation. *Array*, 28, 100532. <https://doi.org/10.1016/j.array.2025.100532>.
25. Khan, A. R. (2025). Federated learning for next generation intelligent applications (PhD thesis, University of Glasgow). University of Glasgow. <https://doi.org/10.5525/gla.thesis.85028>
26. Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-IID data. *arXiv*. <https://arxiv.org/abs/1806.00582>.
27. Mao, Y., Zhao, Z., Yan, G., Liu, Y., Lan, T., Song, L., & Ding, W. (2022). Communication-efficient federated learning with adaptive quantization. *ACM Transactions on Intelligent Systems and Technology*, 13(4), Article 67, 1–26. <https://doi.org/10.1145/3510587>.
28. Li, J., Rakin, A. S., Chen, X., He, Z., Fan, D., & Chakrabarti, C. (2022). ResSFL: A resistance transfer framework for defending model inversion attack in split federated learning. In *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 10184–10192). IEEE. <https://doi.org/10.1109/CVPR52688.2022.00995>.
29. Liu, B., Ding, M., Shaham, S., Rahayu, W., Farokhi, F., & Lin, Z. (2021). When machine learning meets privacy: A survey and outlook. *ACM Computing Surveys*, 54(2), Article 31, 1–36. <https://doi.org/10.1145/3436755>.
30. Dbouk, H., & Shanbhag, N. R. (2021). Generalized depthwise-separable convolutions for adversarially robust and efficient neural networks. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS 2021)* (Article 920, pp. 12027–12039). Curran Associates Inc.
31. Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H. (2018). Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2704–2713). IEEE. <https://doi.org/10.1109/CVPR.2018.00286>.
32. Mao, Y., Zhao, Z., Yan, G., Liu, Y., Lan, T., Song, L., & Ding, W. (2022). Communication-efficient federated learning with adaptive quantization. *ACM Transactions on Intelligent Systems and Technology*, 13(4), Article 51. <https://doi.org/10.1145/3510587>.
33. Reddi, S., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., & McMahan, H. B. (2021). Adaptive federated optimization. *arXiv*. <https://arxiv.org/abs/2003.00295>.
34. Marcolla, C., Sucasas, V., Manzano, M., Bassoli, R., Fitzek, F. H. P., & Aaraj, N. (2022). Survey on fully homomorphic encryption, theory, and applications. *Proceedings of the IEEE*, 110(10), 1572–1609. <https://doi.org/10.1109/JPROC.2022.3205665>.
35. Schmidt, P., Reiss, A., Duerichen, R., Marberger, C., & Van Laerhoven, K. (2018). Introducing WESAD, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction* (pp. 400–408). ACM. <https://doi.org/10.1145/3242969.3242985>.
36. Koldijk, S., Sappelli, M., Verberne, S., Neerinx, M. A., & Kraaij, W. (2014). The SWELL knowledge work dataset for stress and user modeling research. In *Proceedings of the 2014 ACM Conference on User Modeling, Adaptation and Personalization* (pp. 182–183). ACM. <https://doi.org/10.1145/2663204.2663257>.
37. Fawaz, S. M., Belal, N., ElRefaey, A., & Fakhr, M. W. (2021). A comparative study of homomorphic encryption schemes using Microsoft SEAL. *Journal of Physics: Conference Series*, 2128, 012021. <https://doi.org/10.1088/1742-6596/2128/1/012021>.
38. Shi, D., Li, L., Chen, R., Prakash, P., Pan, M., & Fang, Y. (2022). Toward energy-efficient federated learning over 5G+ mobile devices. *IEEE Wireless Communications*, 29(5), 44–51. <https://doi.org/10.1109/MWC.003.2100028>.
39. A., Imteaj, A., Thakker, U., Wang, S., Li, J., & Amini, M. H. (2022). A survey on federated learning for resource-constrained IoT devices. *IEEE Internet of Things Journal*, 9(1), 1–24. <https://doi.org/10.1109/JIOT.2021.3095077>.
40. Li, J., Rakin, A. S., Chen, X., Yang, L., He, Z., Fan, D., & Chakrabarti, C. (2023). Model extraction attacks on split federated learning. *arXiv*. <https://arxiv.org/abs/2303.08581>
41. Abbas, Z., Ahmad, S. F., Syed, M. H., Anjum, A., & Rehman, S. (2024). Exploring deep federated learning for the Internet of Things: A GDPR-compliant architecture. *IEEE Access*, 12, 10548–10574. <https://doi.org/10.1109/ACCESS.2023.3344029>.

42. Baghersalimi, S., Teijeiro, T., Aminifar, A., & Atienza, D. (2024). Decentralized federated learning for epileptic seizures detection in low-power wearable systems. *IEEE Transactions on Mobile Computing*, 23(5), 6392–6407. <https://doi.org/10.1109/TMC.2023.3320862>.
43. Cremonesi, P. (2023). The need for multimodal health data modeling: A practical approach for a federated-learning healthcare platform. *Journal of Biomedical Informatics*, 141, 104338. <https://doi.org/10.1016/j.jbi.2023.104338>.