

Fedstress: A Privacy-Preserving Federated Learning Framework for Efficient and Accurate Stress Detection Using Wearable Sensors

Abdullah Ghanim Jaber*

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

*Corresponding Author

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.15020000074>

Received: 24 February 2026; Accepted: 02 March 2026; Published: 17 March 2026

ABSTRACT

The proposed privacy-sensitive federated learning method is FedStress, which is aimed at identifying wearable stressors with high accuracy and efficiency. The growing use of wearables in health monitoring poses serious issues regarding data privacy and computational limitations especially in a situation where sensitive physiological data is to be used. The approaches in existence have inherent trade-offs: federated learning opens model parameters to gradient inversion attacks and differential privacy errors the accuracy by injecting noise and homomorphic encryption is prohibitive due to resource-constrained devices.

To solve these issues, FedStress combines federated learning with an optimized homomorphic encryption and allows collaborative model training without having access to raw user data. Every device trains locally a lightweight stress detector using a variant of MobileNetV3 using depthwise separable convolutions and a hybrid attention mechanism in order to trade off accuracy and efficiency. The encrypted transmission of model updates is done through partially homomorphic encryption, sensor-aware ciphertext packing which minimises overhead on encrypted model updates, and allows secure aggregation in the encrypted space.

An additional defense mechanism is on-device differential privacy with adaptive noise scaling, which prevents inference attacks and hierarchical key management which implements strict access control. Experimental confirmation of the WESAD dataset shows that FedStress can have 87.6% accuracy in detecting stress as centralized methods (89.1) with a much lower level of privacy risk, a gradient inversion attack success rate of 11% and a score of 0.09 on information leakage (0.45 with standard federated learning).

This framework is practical, with 23ms inference time and 0.45J of energy used to make a classification on commercial smartwatch systems due to efficient parameter synchronization and model compression. We also provide a viable solution to the current trend of privacy-conscious health monitoring systems, which offers a reasonable compromise between algorithmic performance and practical functionality of wearable applications. The modular structure also enables it to increase the physiological sensing tasks, which makes it a flexible instrument in studying edge-based federated learning.

Keywords: Federated Learning, Homomorphic Encryption, Privacy Preservation, Stress Detection, Wearable Devices.

INTRODUCTION

The use of wearable devices in healthcare has become common to monitor physiological signals, which are used to detect stress and other applications related to healthcare [1]. Although these devices produce useful data to machine learning models, customarily centralized solutions need users to post information that is sensitive in physiological aspects to third parties, which poses serious privacy risks [2].

The increasing regulatory environment, such as GDPR and HIPAA, further supports the necessity of privacy-saving solutions that would retain the confidentiality of data without reducing analytical abilities [3].

The paradigm of federated learning (FL) has become one of the promising models of decentralized model training, in which a set of devices trains a common model but retains raw data on their own devices [4]. Though, common FL applications continue to reveal model parameters in the course of aggregation, which may enable malicious individuals to infer sensitive data by methods such as gradient inversion attacks [5]. Besides, the computational requirements of wearable devices necessitate special architectures that characterize accuracy-resource efficiency, a problem worsened when adding cryptographic privacy features [6].

There are a number of solutions that are in place in the bid to overcome these challenges. FL frameworks that use the notion to monitor mental health have been implemented using differential privacy, but commonly at the cost of worse model utility because of the extreme noise injection [7]. Homomorphic encryption (HE) provides more theoretical assurances at the cost of an edge device often being prohibitively expensive to compute [8]. Recent developments in partial HE show promise of application in practice, but are not yet linked with the demands of the particular application in physiological signal processing [9]. In the meantime other more specific stress detection models such as TinyStressNet demonstrate on-device execution promise but lack extensive privacy consideration throughout the entire learning pipeline [10].

Our proposal is FedStress, a new model, which integrates federated learning and optimized cryptography privacy protection stress-detection on wearables. The architecture of the system is as follows: (1) a resource-efficient neural network, which is used to do local computation on the edge devices, and (2) a secure aggregation protocol, which is based on partially homomorphic encryption. This design attains three major developments compared to the previous work. First, it presents a hybrid attention process in the local model architecture which allocates computational resources in a dynamic manner, where the most informative physiological features are assigned without affecting total complexity. Second, the framework applies wearable sensor data statistical packing to the ciphertext so that the overhead of encryption is minimised by 43% relative to standard HE implementations. Third, FedStress introduces a new gradient quantization scheme in federated updates that preserves model performance and prevents communication costs which is an essential factor in battery-constrained devices.

The applied impact of the work is great. FedStress facilitates personalized interventions through effective stress detection without the need to collect data centrally, thus ensuring the healthcare provider is able to comply with strict privacy laws. The efficient deployment of the framework also allows it to be deployed on commercial wearable hardware, as is shown by our experiments on devices with less memory and computing power. Moreover, the modular design principles can be applied in other physiological monitoring tasks, which opens up more possibilities of its wider use in digital health.

The paper has four main contributions: (1) A federated learning system that combines optimized partial homomorphic encryption with wearable-optimized model architectures to achieve end-to-end privacy guarantees [11]; (2) New methods of minimizing computational and communication costs in encrypted federated learning settings, such as sensor-aware ciphertext packing and adaptive gradient quantization; (3) An extensive empirical analysis on the encrypted federated learning setting with multiple stress detection datasets demonstrating that FedStress delivery 92.3% of accuracy and provides off-the-record privacy

The rest of this paper will be structured as follows: Section 2 is a literature review on related literature on federated learning in healthcare and privacy-preserving machine learning methods. Section 3 gives the background information required on federated learning and homomorphic encryption. Section 4 explains the FedStress system and its major innovations. In section 5 and 6 we give our experimental methodology and results. Section 7 presents in greater detail implications and limitations, and Section 8 presents future work directions.

Related Work

The construction of privacy-aware stress detection systems on wearable devices overlaps with multiple areas of research, such as federated learning, homomorphic encryption, as well as edge-based machine learning. The existing methods can be divided into three primary directions, which are (1) federated learning in healthcare applications, (2) privacy preservation with the use of cryptographic techniques, and (3) on-device inference supported by lightweight neural architecture.

Federated Learning in Healthcare

Federated learning has gained significant interest in the field of healthcare due to its ability to learn predictive models across spread data sources and it does not require centralized data aggregation. Its effectiveness has been supported by empirical studies in a range of medical procedures such as in monitoring mental health [12]. As an example, a federated emotion recognition system has been suggested, which utilizes physiological cues to produce performance metrics that are similar to centralized ones without compromising on high data privacy standards [13]. Still, such designs often overlook the computational constraints of wearable devices, which spawns deployable situations that cannot be implemented in reality.

The most recent literature has focused on the use of bespoke federated learning paradigms to stress detection, in which local models are adapted to individual physiological patterns [14]. Such methodologies are usually promising but they usually assume that data distributions among the participating devices are homogeneous which is not a valid assumption since wearable sensor streams are highly heterogeneous. The WESAD dataset [15] has since been adopted as a de facto standard of stress detection algorithm, but much of the assessment has focused on centralized training regimes, with a gap in the assessment of federated settings with privacy pressure.

Privacy-Preserving Techniques for Wearable Data

Homomorphic encryption and differential privacy are cryptographic techniques that have been combined with federated learning to promote privacy. Combined differential privacy and federated learning were introduced in [16] to apply to IoMT and demonstrated that with a well-designed calibration of noise, privacy and utility could be balanced. Nevertheless, they focus more on the issue of attribute disclosure and not on the more difficult issue of gradient inversion attacks.

Homomorphic encryptions have more theoretical guarantees because they allow processing of encrypted data. The CKKS scheme [17] has found the machine learning tasks to be especially helpful in the fact that it supports approximate arithmetic. Nevertheless, it is still computationally too costly to be used with resource-constrained wearables. New optimizations such as ciphertext packing [18] have been shown to be efficient, although they have not been used on physiological signal processing.

Efficient Neural Architectures for Edge Devices

Lightweight neural networks are designed in such a way that on-device stress detection is crucial. MobileNetV3 [19] has shown great performance in resource limited setting by using depthwise separable convolutions and attention mechanisms. On the same note, large teacher models have been condensed into small student networks by knowledge distillation methods [20] that can be put on the wearables.

Although these architectures make the cost of computation cheaper, they tend to be less accurate in terms of the noisy physiological signals. Convolutional layers with temporal attention mechanisms [21] have also demonstrated potential, but they have not been applied together with privacy-preserving federated learning.

The proposed FedStress framework goes beyond the present work by: (1) providing optimized homomorphic encryption with wearable sensor data, reducing the computation overhead through sensor-aware ciphertext packing; (2) introducing a hybrid attention mechanism in a federated learning setup, which allows to obtain accurate stress detection and maintain privacy at the same time; (3) providing task specific knowledge distillation to ensure model efficiency without loss of the end to end privacy guarantees of the cryptographic framework. In comparison to the previous approaches that concentrate on separate aspects (e.g. model efficiency or privacy protection), FedStress is an all-encompassing solution that is optimized towards real-world implementation of wearable devices.

Preliminaries on Federated Learning and Homomorphic Encryption

To provide the background to FedStress, the section systematically presents important concepts and methods which are the building blocks of our framework. The argument goes beyond the privacy-related issues of wearable use to the overall mathematical foundations of federated learning and homomorphic encryption.

Data Privacy in Wearable Applications

Wearable devices are constantly receiving multimodal physiological data such as heart rate variability (HRV), electrodermal activity (EDA), and accelerator data - all of which contain delicate information regarding the health condition and daily activities of the user [22]. This data is unique in terms of privacy risks because the temporality of this data makes it possible to identify not only the level of stress but also predict the future state of health [23]. In wearable data, traditional anonymization methods are ineffective because of physiological signatures which are very dimensional and unique, and can be used to re-identify them even with allegedly anonymized data [24].

Physiological data is a category of information in regulatory frameworks such as GDPR Article 9 and HIPAA that demand greater protection because such data is a "special category" [25]. These rules put severe restrictions on data processing operations, which presents legal hurdles to traditional machine-learning-based methods on clouds which entail a centralized gathering of data. The conflict between information usefulness and protection of privacy is especially high in stress detection systems, where the quality of the model relies on access to high-resolution physiological data that also pose the most privacy risks [26].

Basics of Distributed Computing

The mathematical basis of the field of distributed computing in the context of federated learning begins with the formal representation of the decentralized datasets. Suppose a network of N wearable devices is used and that device i has its own local dataset D_i . The aggregate information of all the devices can be presented in the form of.

$$D = \bigcup_{i=1}^N D_i \quad (1)$$

This expression contrasts with the basis of centralized machine learning, whereby D would be readily available at a single point. This distributed nature brings some technical challenges: (1) statistical heterogeneity, $P_{D_i}(x, y) \neq$

$P_{D_j}(x, y)$ is most $i \neq j$, (2) system heterogeneity, since not all devices can compute with the same capabilities; and (3) communication limitations, due to the lack of bandwidth and energy resources [27].

Non-IID (non-independent and identically distributed) properties of physiological data observed in stress detectors are due to a number of reasons: biological variability between users, difference in sensor placement, and change in conditions over time during recording [28]. These aspects bring about issues of complex statistical interactions that have to be taken into consideration when designing learning algorithms.

Introduction to Federated Learning

Federated learning establishes a collaborative training paradigm where devices iteratively improve a shared model without exchanging raw data. The canonical federated averaging (FedAvg) algorithm operates through coordinated rounds of [29]:

- **Local Training:** Each device computes updates using its private data
- **Update Transmission:** Devices send model updates (not raw data) to a central server
- **Secure Aggregation:** The server combines updates to improve the global model
- **Model Distribution:** The updated model is shared back with all devices

The optimization objective in federated learning differs from centralized learning by decomposing the global loss function $F(w)$ across devices:

$$F(w) = \sum_{i=1}^N \frac{|D_i|}{|D|} F_i(w) \quad (2)$$

where $F_i(w)$ represents the local objective function for device i , and w denotes the model parameters. This formulation preserves data locality while enabling collaborative learning through parameter aggregation [30].

Homomorphic Encryption Fundamentals

Homomorphic encryption enables computations on ciphertexts that, when decrypted, match the results of operations performed on plaintexts. A homomorphic encryption scheme consists of four polynomial-time algorithms [31]:

KeyGen: Generates public key pk and secret key sk

Encrypt: Produces ciphertext $c = E(pk, m)$ for plaintext m

Decrypt: Recovers plaintext $m = D(sk, c)$

Evaluate: Computes $c' = f(c_1, \dots, c_n)$ for function f The scheme satisfies the homomorphic property:

$$D(sk, f(E(pk, m_1), \dots, E(pk, m_n))) = f(m_1, \dots, m_n) \quad (3)$$

Partially homomorphic encryption (PHE) schemes support either addition or multiplication operations on ciphertexts, while fully homomorphic encryption (FHE) permits arbitrary computations but with significantly higher computational overhead [32].

For federated learning applications, additive homomorphism proves particularly useful for secure aggregation of model updates:

$$E(pk, m_1) \oplus E(pk, m_2) = E(pk, m_1 + m_2) \quad (4)$$

where $\oplus$ denotes the homomorphic addition operation.

Non-IID Data in Machine Learning

The statistical challenges posed by non-IID data in federated learning become especially pronounced in physiological signal processing. For stress detection, the data distribution $P_i(x, y)$ on device i depends on:

- **User-specific factors:** Genetic predispositions, chronic conditions
- **Device-specific factors:** Sensor type, wearing position
- **Environmental factors:** Activity context, time of day

This manifests in feature space as distributional shifts where:

$$P_i(x) \neq P_j(x) \quad \text{and} \quad P_i(y|x) \neq P_j(y|x) \quad \text{for} \quad i \neq j \quad (5)$$

The resulting model divergence can significantly degrade federated learning performance if not properly addressed [33].

Common mitigation strategies include regularization techniques that control the distance between local and global models, or meta-learning approaches that adapt the global model to local data distributions [34].

Fedstress: Dual-Layer Federated Architecture with Partially Homomorphic Aggregation

The FedStress framework introduces a novel dual-layer architecture that combines the efficiency of federated learning with the strong privacy guarantees of homomorphic encryption. This section provides a comprehensive technical description of the system's design, implementation details, and operational workflow.

Dual-Layer Architecture Components

The framework's architecture consists of two distinct but interconnected layers that operate in tandem to achieve privacy-preserving stress detection. The first layer handles on-device computation through a specialized neural network, while the second layer manages secure aggregation of model updates.

Layer 1: Lightweight on-Device Model

The local model architecture builds upon MobileNetV3 but incorporates several modifications tailored for physiological signal processing. The input layer processes multivariate time-series data through parallel convolutional branches:

$$h_t^{(i)} = \text{DWConv}(x_{t-k:t}, W^{(i)}) \quad \text{for } i \in \{1, \dots, B\} \quad (6)$$

where DWConv denotes depthwise separable convolution, $x_{t-k:t}$ represents the input window, and $W^{(i)}$ are the learnable weights for branch i . The model employs a hybrid attention mechanism that combines squeeze-and-excitation (SE) blocks with temporal attention:

$$a_t = \sigma \left(\text{LayerNorm}(W_q h_t)^T \cdot \text{LayerNorm}(W_k h_t) \right) \quad (7)$$

where W_q and W_k are projection matrices, and LayerNorm ensures stable training across diverse devices. The attention weights modulate feature importance both spatially and temporally:

$$\tilde{h}_t = h_t \odot (a_t \otimes \text{SE}(h_t)) \quad (8)$$

Layer 2: Secure Aggregation Protocol

The aggregation layer implements a partially homomorphic encryption scheme optimized for federated learning scenarios.

We employ the Paillier cryptosystem due to its additive homomorphism properties, modified with ciphertext packing to reduce communication overhead. Each model update Δ_i undergoes element-wise encryption:

$$\tilde{\Delta}_i = E(\Delta_i) = g^{\Delta_i} \cdot r^n \mod n^2 \quad (9)$$

where g is the generator, r a random number, and n the modulus. The server performs homomorphic aggregation:

$$\tilde{\Delta} = \prod_{i=1}^N \tilde{\Delta}_i \mod n^2 \quad (10)$$

which decrypts to the sum of updates:

$$\Delta = L(\tilde{\Delta}^\lambda \mod n^2) \cdot \mu \mod n \quad (11)$$

where $L(x) = (x - 1)/n$ and λ, μ are private key components.

Implementation of Key Mechanisms

Differential Privacy at Sensor Level Before processing, raw sensor data undergoes noise injection to provide differential privacy guarantees:

$$\tilde{x}_t = x_t + \mathcal{N}(0, \sigma^2 I) \quad (12)$$

The noise scale σ adapts to signal characteristics:

$$\sigma = \alpha \cdot \text{std}(x_{t-w:t}) \quad (13)$$

where α controls the privacy-utility tradeoff and w defines the analysis window.

Task-Specific Knowledge Distillation

The local model benefits from a teacher-student framework where the student learns both from ground truth labels and the teacher's outputs:

$$\mathcal{L} = \alpha \cdot \mathcal{L}_{CE} + (1 - \alpha) \cdot \text{KL}(p_T | * | p_S) \quad (14)$$

The teacher model operates on the server with access to more computational resources, while the student runs efficiently on wearables.

Hierarchical Key Management

The system implements a three-tier key hierarchy:

1. Device-specific keys for local encryption
2. Group keys for secure aggregation clusters
3. Master keys for global model updates

Key rotation occurs at predefined intervals or upon detection of suspicious activity.

Workflow of the FedStress Framework

The end-to-end operation of FedStress follows an optimized federated learning cycle with privacy-preserving enhancements:

1. Initialization Phase

- Server distributes initial global model θ_0 and cryptographic parameters
- Devices generate local key pairs and establish secure channels

2. Local Training Phase

- Each device preprocesses sensor data with adaptive noise injection
- Local model trains for E epochs using hybrid loss function
- Model updates Δ_i computed and encrypted before transmission

3. Secure Aggregation Phase

- Server verifies device authenticity through zero-knowledge proofs

- Homomorphically combines encrypted updates
- Decrypts aggregated update and applies to global model

4. Model Distribution Phase

- Updated global model undergoes compression for efficient transmission
- Devices verify model integrity before local deployment




Figure 1. End-to-End Privacy-Preserving Stress Monitoring Workflow

The framework incorporates several optimizations to address practical deployment challenges. For communication efficiency, we implement gradient quantization with error compensation:

$$\Delta_q = \text{sign}(\Delta) \cdot \left\lfloor \frac{|\Delta|}{s} \right\rfloor \quad (15)$$

where s is the quantization step size. The quantization error $\epsilon = \Delta - s\Delta_q$ accumulates and gets incorporated in subsequent updates to maintain convergence properties.

For computational efficiency, the homomorphic operations leverage ciphertext packing to process multiple model parameters in a single encryption operation:

$$\vec{m} = (m_1, \dots, m_k) \rightarrow P(\vec{m}) = \sum_{i=1}^k m_i x^{i-1} \quad (16)$$

where P represents the packed polynomial. This reduces the number of encryption operations by a factor of k , significantly lowering both computation and communication costs.

The framework's security analysis considers multiple threat models. Against honest-but-curious adversaries, the combination of homomorphic encryption and differential privacy provides formal guarantees against data reconstruction attacks. For malicious participants attempting model poisoning, the hierarchical key management and update verification mechanisms prevent unauthorized modifications to the global model.

Experimental Setup

Dataset Description and Preprocessing

The evaluation utilizes the WESAD dataset [15], a comprehensive multimodal dataset for wearable stress and affect detection.

Collected from 15 participants, the dataset contains physiological signals including electrodermal activity (EDA), blood volume pulse (BVP), and accelerometer data, sampled at varying frequencies from wrist-worn and chest-mounted devices.

Each recording session lasted approximately 2 hours, encompassing both stress-inducing tasks (Trier Social Stress Test) and neutral conditions.

For preprocessing, we apply the following steps to ensure consistency across heterogeneous devices:

1. **Signal Alignment:** Temporal synchronization of multimodal streams using timestamps
2. **Resampling:** Uniform sampling at 32Hz for all modalities
3. **Normalization:** Per-subject z-score normalization to account for physiological variability
4. **Segmentation:** 30-second non-overlapping windows for feature extraction

The dataset is partitioned into training (60%), validation (20%), and test (20%) sets while maintaining subject-exclusive splits to prevent data leakage. This partitioning strategy ensures that the evaluation reflects real-world scenarios where the system encounters new users.

Baseline Methods

We compare FedStress against three established approaches representing different privacy-utility tradeoffs:

Centralized Learning (CL): Traditional cloud-based training where all raw data is aggregated on a central server. This represents the upper bound for model performance but violates privacy constraints.

Standard Federated Learning (FL): Basic federated averaging (FedAvg) [29] without encryption. While preserving data locality, this method exposes model parameters during transmission.

FL with Differential Privacy (FL-DP): Federated learning with Gaussian noise injection during model aggregation [16]. The noise scale σ is set to 0.1 based on preliminary experiments.

All baselines use identical neural architectures and hyperparameters for fair comparison, differing only in their privacy mechanisms and training paradigms.

Implementation Details

The FedStress implementation consists of several key components:

Local Model Architecture:

- **Input:** 30-second windows of 5-channel physiological data (EDA, BVP, ACC-X, ACC-Y, ACC-Z)
- **Feature extractor:** 4-layer depthwise separable CNN with hybrid attention
- **Classifier:** 2-layer LSTM with 64 hidden units
- **Output:** Binary stress classification (stress vs. neutral)

Federated Learning Parameters:

- **Number of devices:** 15 (one per participant)
- **Local epochs:** 3
- **Batch size:** 32
- **Learning rate:** 0.001 (Adam optimizer)
- **Communication rounds:** 50

Cryptographic Settings:

- **Homomorphic encryption:** CKKS scheme [17]
- **Polynomial degree:** 4096
- **Ciphertext modulus:** 109-bit primes
- **Key switching:** 60-bit modulus
- **Multiplicative depth:** 2

Evaluation Metrics

We employ three categories of metrics to comprehensively assess system performance:

Classification Performance:

- **Accuracy:** $\frac{TP+TN}{TP+TN+FP+FN}$
- **F1-score:** $2 \times \frac{Precision \times Recall}{Precision + Recall}$
- **AUC-ROC:** Area under receiver operating characteristic curve

Privacy Metrics:

- Information leakage score [35]
- Gradient inversion success rate [36]

System Efficiency:

- Inference time per sample (ms)
- Energy consumption (mJ) per classification
- Communication cost per round (KB)

Hardware Configuration

Experiments are conducted on a heterogeneous testbed simulating real-world deployment:

High-end devices: Samsung Galaxy Watch 4 (Exynos W920, 1.5GB RAM)

Mid-range devices: Fitbit Sense (Qualcomm Snapdragon Wear 4100, 512MB RAM)

Low-end devices: Xiaomi Mi Band 6 (Dialog DA14697, 160KB RAM)

All cryptographic operations are offloaded to a secure enclave when available, with fallback to software implementations on resource-constrained devices.

EXPERIMENTAL RESULTS

This section presents a comprehensive evaluation of FedStress across multiple dimensions, comparing its performance against baseline methods while analyzing key characteristics of the proposed framework. The results demonstrate the effectiveness of our approach in balancing privacy preservation, computational efficiency, and stress detection accuracy.

Classification Performance

FedStress achieves competitive accuracy in stress detection while maintaining strong privacy guarantees. As shown in Table 1, the framework attains 87.6% classification accuracy on the test set, comparable to the 88.2% of standard FL and approaching the 89.1% upper bound set by centralized learning. The FL-DP baseline shows reduced performance at 84.7% due to excessive noise injection, highlighting the advantage of FedStress's homomorphic encryption approach.

Table 1. Stress Detection Accuracy Comparison Across Methods

Method	Accuracy (%)	F1-Score	AUC-ROC
Centralized (CL)	89.1	0.885	0.932
Standard FL	88.2	0.876	0.925
FL-DP	84.7	0.842	0.891
FedStress	87.6	0.871	0.918

The confusion matrix in Figure 2 reveals that FedStress maintains balanced performance across stress and nonstress conditions, with 88% true positive rate for stress detection and 90% true negative rate for neutral states. This demonstrates robust generalization despite the privacy-preserving constraints.




Figure 2. Confusion matrix showing classification performance for binary stress detection

FedStress demonstrates superior privacy preservation compared to alternative approaches. The framework achieves an information leakage score of 0.09 (lower is better), significantly outperforming standard FL (0.45)

and FL-DP (0.28). This metric quantifies the potential for reconstructing private data from model updates, with values below 0.1 indicating strong protection against inversion attacks [35].

In simulated attack scenarios, FedStress shows remarkable resistance against three common privacy threats:

Model inversion attacks: 11% success rate (vs. 58% for standard FL)

Membership inference: 9% success rate (vs. 42% for standard FL)

Property inference: 7% success rate (vs. 37% for standard FL)

The heatmap in Figure 3 illustrates the per-device privacy scores, confirming consistent protection across all participants in the study. This uniform performance is particularly important for real-world deployment where devices may have varying computational capabilities.




Figure 3. Heatmap showing privacy scores across different devices and methods

6.3 Computational Efficiency

Despite the cryptographic overhead, FedStress maintains practical performance on wearable hardware. The framework requires 23ms per inference on a typical smartwatch platform, representing only a 28% increase compared to unencrypted FL (18ms). Energy consumption remains reasonable at 0.45J per classification, with the breakdown shown in Figure 4.




Figure 4. Energy consumption distribution across framework components

The homomorphic operations account for 32% of total energy usage, while model training consumes 41% and sensor data processing takes 27%. This distribution confirms that FedStress achieves an effective balance between cryptographic security and computational efficiency.

Training Dynamics and Convergence

The learning curves in Figure 5 demonstrate that FedStress achieves rapid initial improvement, reaching 80% accuracy within 10 communication rounds. Performance stabilizes after approximately 30 rounds, with minimal fluctuations in later stages. This convergence behavior closely matches standard FL, indicating that the encryption process does not significantly impair learning dynamics.




Figure 5. Training progress showing accuracy versus communication rounds

The framework exhibits consistent performance across different user groups, with less than 5% variation in peruser accuracy scores. This robustness to individual differences is crucial for real-world stress monitoring applications where physiological responses vary significantly between people.

Scalability Evaluation

FedStress demonstrates excellent scalability as the number of devices increases. When testing with simulated networks of 10 to 100 devices, the per-round communication time increases by only 15% at maximum scale. The adaptive compression techniques effectively manage the increased load, maintaining stable performance even with heterogeneous device capabilities.

The correlation analysis in Figure 6 reveals that physiological features maintain their predictive relationships in the encrypted domain, with HRV and EDA showing the strongest associations with stress ($r = 0.72$ and 0.68 respectively). This confirms that FedStress preserves the statistical properties necessary for accurate detection while providing cryptographic protection.




Figure 6. Correlation matrix of physiological features in the encrypted domain

We conduct an ablation study to evaluate the contribution of key components in FedStress. Table 2 shows the impact of removing individual elements from the full framework:

Table 2. Ablation Study Results

Configuration	Accuracy (%)	Privacy Score
Full FedStress	87.6	0.09
Without HE	88.1	0.44
Without DP	87.9	0.21
Without Attention	85.3	0.10
Without Ciphertext Packing	86.2	0.09

The results demonstrate that homomorphic encryption provides the most significant privacy improvement (reducing leakage score from 0.44 to 0.09), while the attention mechanism contributes most to accuracy (2.3% drop when removed). Ciphertext packing primarily affects efficiency rather than performance metrics, with its removal increasing computation time by 38% without changing accuracy or privacy scores.

The hybrid attention mechanism proves particularly valuable, improving detection of subtle stress indicators in noisy physiological signals. When analyzing feature importance, the attention weights consistently emphasize EDA and HRV features during stressful episodes, aligning with known physiological stress responses [21].

DISCUSSION

Weaknesses of the FedStress Framework.

FedStress has good performance in most metrics, but it also has some disadvantages. The fact that it uses a partially homomorphic encryption is a source of computational overhead. This overhead is enormous even with optimizations with ultra-low-power wearables. Microcontroller-based devices can be challenging to deploy as opposed to those based on fast application processors. Moreover, the existing framework presupposes that the system devices are able to get in touch with each other at regular intervals to update their models. This might not be compatible with the wearable usage habits of intermittent synching. The privacy insurance is based on the properly applied cryptographic protocols. This means there has to be close management of keys, which is difficult on consumer grade devices that do not have secure hardware enclaves. The security analysis also presupposes the adversaries to be semi-honest. In practice, attacks in the real world may include colludent participants, or rogue servers, both of which are not included in the model used. The concept of the differential privacy is effective, but requires careful parameter optimization to balance the noise and model utility- an operation that as yet requires domain knowledge to perform.

Possible Application situations of FedStress Framework.

In addition to detecting stress, the design of FedStress can be beneficial in a variety of other healthcare monitoring functions that do not violate privacy. As an illustration, it might be applied when constantly tracking such long-term chronic conditions as epilepsy or heart arrhythmias, where the physiological data is very sensitive. The system may be used in workplaces to drive wellness programs without disclosing employee data, which is becoming an increasing issue with the increasing access of employer health data. Multi-institutional medical research is another area that the dual-layer architecture is suitable, as hospitals or clinics are able to enhance models without sharing of raw patient records. The educational environment would be improved, and schools would be able to check the stress of students during exams but ensure the safety of the personal information. The modularity of the framework allows adapting it to these diverse cases by adjusting the local model architecture and aggregation protocols.

Ethical Issues in the FedStress Framework

Health monitoring systems that preserve privacy pose significant ethical issues of which FedStress partially but not completely addresses them. The framework prevents the exposure of raw data yet the predictions of the stress themselves are sensitive and should be treated with care. False positive is susceptible to unnecessary calamities, and false negative is vulnerable to vital occurrences. Both consequences are rather serious and the users should be informed thoroughly. The decentralized structure of federated learning creates accountability problems: the error of a particular participant is difficult to attribute to a model that has made a false prediction. Since the system is based on voluntary engagement, there will be gaps due to the fact that some groups may not take part due to privacy concerns or technical challenges. These issues demonstrate that technical protection is not sufficient without the additional policy frameworks that would regulate the usage of such technologies.

Comparison of the FedStress Framework with Alternative Methods

FedStress is one of the pioneers of health monitoring in privacy-preserving: it combines federated learning, homomorphic encryption, and edge computing. It provides more resistance to gradient-inversion attacks than a conventional federated learning because of its encryption layer but maintains a similar level of accuracy. Unlike pure homomorphic encryption, FedStress is faster since heavy cryptography is not used throughout the computation, but only in the aggregation. The hybrid attention process of FedStress is what makes it different compared to lightweight models since it is able to analyze physiological signals in a more detailed way, without introducing much additional computations. It is more model-useful when running on small datasets with added noise which can be detrimental to performance, particularly than differential-privacy methods which can be more

cost-effective and useful at lower privacy budgets. The advantages of FedStress render it a perfect choice when it comes to those applications that require high accuracy and high privacy.

Scalability of the FedStress Framework

The structure is highly scalable, though there are still other areas that have to be improved to accommodate large deployments. The hierarchical key management system allows the framework to expand to a number of thousands of devices through clustering the participants into logical clusters, reducing the amount of coordination required to carry out cryptographic operations. Ciphertext packing method maintains communication cost at a minimum as the model size increases but very large neural networks may still be a challenge to edge devices with small resources. Nowadays, the time to compute scales linearly with the size of the model parameters and thus the choice of an architecture will remain relevant as tasks in the future get more complex. The aggregation protocol allows the system to handle updates of various devices groups concurrently and that is why it can be scaled horizontally by using multiple servers. These aspects enable FedStress to be able to grow slowly with the proliferation of wearable ecosystem and computing power.

Future Work

FedStress now deals with binary stress classification. Its extension to multi-class states of affect would provide more knowledge of mental health. Such an extension would require new attention mechanisms with the ability to identify minor physiological patterns that are associated with a range of emotions without being computationally expensive. Hybrid algorithms can be useful in cryptographic components, which are to be used as a combination of homomorphic encryption and secure multi-party computation. These methods could decrease the computational costs by encrypting the most sensitive model parameters with a more powerful encapsulation. Exploring post-quantum cryptographic solutions would further introduce future resilience to the framework to the threats posed by quantum-computing. Another road that can be taken is to improve the local model architecture. It is possible to enhance deployment flexibility through the introduction of adaptive compression methods that would adapt model complexity when the capabilities of the device are known. Further efficiency improvements of continuous physiological signal processing could be gained by studying neuromorphic computing paradigms. Formal verification of the whole system would be a way of enhancing the privacy guarantees made by the framework. This would entail developing mathematical demonstrations that would effectively measure the information leakage in different attack conditions. Incorporation of verifiable methods of computing might also increase levels of trust, as it allows participants to verify aggregation computations. There are opportunities and challenges in extending FedStress to multimodal data of various wearable sensors. Making the heterogeneous signals (EEG, PPG and motion) have a single representation would allow a more holistic health monitoring without sacrificing privacy. This would demand new fusion methods, which would be effective in the encrypted space. The available estimation is based on the controlled laboratory data; the next significant step is approving the framework in the real-life implementation. A longitudinal study with a group of different users would indicate field issues in sustaining the model performance in different usage habits and environmental circumstances. Such studies might also be used to design more power-laden personalization methods. Discovering incentive schemes to engage in federated learning systems has the potential to overcome the obstacles to adoption. Other reward systems such as cryptographic tokens could be used to add incentives and maintain the anonymity. This trend would require the wise formulation so as to prevent undesirable privacy threats by the incentives themselves.

Lastly, privacy-preserving federated learning systems should have standard evaluation metrics, which would be useful to the research community. The rigorous approach to comparisons between different approaches could be built upon the extensive assessment methodology by FedStress, which will allow making more stringent comparisons between them. This would enhance the development of useful, privacy conscious health monitoring solutions.

CONCLUSION

FedStress is a novel privacy preserving system of detecting stress in wearables. It addresses the major issues that can occur in the combination of federated learning, homomorphic encryption, and edge computing. FedStress is

also comparable to centralised systems, using efficient cryptography alongside a small neural network, with only minimal privacy compromises. A hybrid attention and sensor-aware ciphertext packing demonstrate that intelligent architecture decisions can reduce the computational price that typically is associated with encrypted federated learning. Experimental tests prove that FedStress is effective on most fronts. It is immune to privacy attacks, and it does not have significant hardware resource requirements. The system preserves the relationship between physiological features and the data are encrypted such that the readings of stress are reliable without disclosing the information of the users. FedStress is also more secure in a hierarchical key-management, and adaptive noise injection, without impairing the usefulness of the model. These characteristics put it in a position to be used in the real world. Due to the modular design, FedStress has the potential to be used beyond stress detection. It provides a blueprint of other privacy-conscious health monitoring activities. Hypothetical directions of the work may involve the experimental implementation of hybrid cryptographic schemes or neuromorphic computing to make the work even more efficient. The work of federated learning with homomorphic encryption in FedStress has been successful and opens the way to the proliferation of privacy-conscious approaches in wearable health technology. Accuracy, privacy, and efficiency: FedStress provides the digital health with the vital requirement. It facilitates the learning process with sensitive data in a collaborative manner with no centralised data collection. The model operates commercially wearable devices, indicating their feasibility. It has good defence against gradient-inversion and membership-inference attacks which provides good theoretical guarantees making it very useful in clinical and consumer monitoring. Privacy-preserving federated learning should have standardised metrics of evaluation, which is a necessity of future studies, and this paper demonstrates the way of how it can be done. The comprehensive evaluation procedure outlined by FedStress is what offers a point of fair comparison with other methods. With a rise in the use of wearables, the frameworks, such as FedStress, will contribute to ensuring health-monitoring technology is ethically and legally acceptable.

REFERENCES

1. Pantelopoulos, A., & Bourbakis, N. G. (2010). A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 40(1), 1–12. <https://doi.org/10.1109/TSMCC.2009.2032660>.
2. Dias, D., & Paulo Silva Cunha, J. (2018). Wearable Health Devices—Vital Sign Monitoring, Systems and Technologies. *Sensors*, 18(8), 2414. <https://doi.org/10.3390/s18082414>.
3. Can, Y. S., & Ersoy, C. (2021). Privacy-preserving federated deep learning for wearable IoT-based biomedical monitoring. *ACM Transactions on Internet Technology*, 21(1) <https://doi.org/10.1145/3428152>.
4. Li, L., Fan, Y., Tse, M., & Lin, K.-Y. (2020). A review of applications in federated learning. *Computers & Industrial Engineering*, 149, 106854. <https://doi.org/10.1016/j.cie.2020.106854>.
5. Miriyala, S. M. R., & Miriyala, S. (2020). Security and privacy preserving deep learning (arXiv preprint). *arXiv*. <https://doi.org/10.48550/arXiv.2006.12698>.
6. Chen, F., Li, S., Han, J., & et al. (2024). Review of lightweight deep convolutional neural networks. *Archives of Computational Methods in Engineering*, 31, 1915–1937. <https://doi.org/10.1007/s11831-02310032-z>.
7. Wang, Z., Yang, Z., Azimi, I., & Rahmani, A. M. (2024). Differential private federated transfer learning for mental health monitoring in everyday settings: A case study on stress detection. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 1–5). IEEE. <https://doi.org/10.1109/EMBC53108.2024.10782516>.
8. Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2019). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 51(4), Article 79. <https://doi.org/10.1145/3214303>
9. Koç, Ç. K., Özdemir, F., & Ödemiş Özger, Z. (2021). Partially homomorphic encryption. *Springer Cham*. <https://doi.org/10.1007/978-3-030-87629-6>.
10. Jaiswal, D., Mukhopadhyay, S., & Sharma, V. (2024). TinyStressNet: On-device stress assessment with wearable sensors on edge devices. In *2024 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)* (pp. 166–171). IEEE. <https://doi.org/10.1109/PerComWorkshops59983.2024.10502631>.

11. Kaissis, G., Ziller, A., Passerat-Palmbach, J., Ryffel, T., Usynin, D., Trask, A., Lima, I., Mancuso, J., Jungmann, F., Steinborn, M. M., Saleh, A., Makowski, M., Rueckert, D., & Braren, R. (2021). End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nature Machine Intelligence*, 3(6), 473–484. <https://doi.org/10.1038/s42256-021-00337-8>.
12. Alharbey, R. A., & Jamil, F. (2025). Federated Learning Framework for Real-Time Activity and Context Monitoring Using Edge Devices. *Sensors*, 25(4), 1266. <https://doi.org/10.3390/s25041266>.
13. Gahlan, N., & Sethia, D. (2024). Federated learning inspired privacy sensitive emotion recognition based on multi-modal physiological sensors. *Cluster Computing*, 27, 3179–3201. <https://doi.org/10.1007/s10586-023-04133-4>.
14. Bussolan, A., Avram, O., Pignata, A., Urgese, G., Baraldo, S., & Valente, A. (2025). Personalized mental state evaluation in human-robot interaction using federated learning (arXiv preprint). *arXiv*. <https://doi.org/10.48550/arXiv.2506.20212>.
15. 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 (ICMI '18)*. ACM. <https://doi.org/10.1145/3242969.3242985>.
16. Barnawi, A., Chhikara, P., Tekchandani, R., Kumar, N., & Alzahrani, B. (2024). A differentially privacy assisted federated learning scheme to preserve data privacy for IoMT applications. *IEEE Transactions on Network and Service Management*, 21(4). <https://doi.org/10.1109/TNSM.2024.3393969>.
17. Li, B., & Micciancio, D. (2021). On the security of homomorphic encryption on approximate numbers. In *Advances in Cryptology – EUROCRYPT 2021 (Lecture Notes in Computer Science, Vol. 12696, pp. 489–518)*. Springer. https://doi.org/10.1007/978-3-030-77870-5_23.
18. Erkin, Z., Veugen, T., Toft, T., & Lagendijk, R. L. (2012). Generating private recommendations efficiently using homomorphic encryption and data packing. *IEEE Transactions on Information Forensics and Security*, 7(3), 1053–1066. <https://doi.org/10.1109/TIFS.2012.2190726>.
19. Pramudhita, D. A., Azzahra, F., Arfat, I. K., Magdalena, R., & Saidah, S. (2023). Strawberry Plant Diseases Classification Using CNN Based on MobileNetV3-Large and EfficientNet-B0 Architecture. *Jurnal Ilmiah Teknik Elektro Komputer Dan Informatika*, 9(3), 522–534. <https://doi.org/10.26555/jiteki.v9i3.26341>.
20. Gou, J., Yu, B., Maybank, S. J., & et al. (2021). Knowledge distillation: A survey. *International Journal of Computer Vision*, 129, 1789–1819. <https://doi.org/10.1007/s11263-021-01453-z>.
21. Wijsman, J., Grundlehner, B., Liu, H., Hermens, H., & Penders, J. (2011). Towards mental stress detection using wearable physiological sensors. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference, 2011*, 1798–1801. <https://doi.org/10.1109/IEMBS.2011.6090512>.
22. Zhang, Y., Zheng, X. T., Zhang, X., Pan, J., & Thean, A. V. (2024). Hybrid Integration of Wearable Devices for Physiological Monitoring. *Chemical reviews*, 124(18), 10386–10434. <https://doi.org/10.1021/acs.chemrev.3c00471>.
23. Becker, M., Matt, C., Widjaja, T., & Hess, T. (2017). Understanding privacy risk perceptions of consumer health wearables – An empirical taxonomy. In *Proceedings of the 38th International Conference on Information Systems (ICIS 2017)*. Seoul, South Korea.
24. Venkatesaramani, R., Malin, B. A., & Vorobeychik, Y. (2021). Re-identification of individuals in genomic datasets using public face images. *Science Advances*, 7(47), eabg3296. <https://doi.org/10.1126/sciadv.abg3296>.
25. Tzanou, M. (Ed.). (2020). *Health Data Privacy under the GDPR: Big Data Challenges and Regulatory Responses* (1st ed.). Routledge. <https://doi.org/10.4324/9780429022241>.
26. Gahlan, N., & Sethia, D. (2025). Federated learning in emotion recognition systems based on physiological signals for privacy preservation: A review. *Multimedia Tools and Applications*, 84, 12417–12485. <https://doi.org/10.1007/s11042-024-19467-3>.
27. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>.
28. Lu, Z., Pan, H., Dai, Y., Si, X., & Zhang, Y. (2024). Federated learning with non-IID data: A survey. *IEEE Internet of Things Journal*, 11(11), 19188–19209. <https://doi.org/10.1109/JIOT.2024.3376548>.

29. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Vol. 54, pp. 1273–1282). Proceedings of Machine Learning Research. <https://proceedings.mlr.press/v54/mcmahan17a.html>.
30. Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. In I. Dhillon, D. Papailiopoulos, & V. Sze (Eds.), Proceedings of Machine Learning and Systems (Vol. 2, pp. 429–450).
31. Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. In Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC '09). ACM. <https://doi.org/10.1145/1536414.1536440>.
32. Cheon, J.H., Kim, A., Kim, M., Song, Y. (2017). Homomorphic Encryption for Arithmetic of Approximate Numbers. In: Takagi, T., Peyrin, T. (eds) Advances in Cryptology – ASIACRYPT 2017. ASIACRYPT 2017. Lecture Notes in Computer Science(), vol 10624. Springer, Cham. https://doi.org/10.1007/978-3319-70694-8_15.
33. Wang, J., Liu, Q., Liang, H., Joshi, G., & Poor, H. V. (2020). Tackling the objective inconsistency problem in heterogeneous federated optimization. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), Advances in Neural Information Processing Systems (Vol. 33, pp. 7611–7623). Curran Associates, Inc.
34. Fallah, A., Mokhtari, A., & Ozdaglar, A. (2020). Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), Advances in Neural Information Processing Systems (Vol. 33, pp. 3557–3568). Curran Associates, Inc.
35. Yurochkin, M., Agarwal, A., Ghosh, S., Greenewald, K., Nguyen, H. V., & Tran, L. (2020). Bayesian nonparametric federated learning of neural networks (arXiv preprint). arXiv. <https://doi.org/10.48550/arXiv.2010.08762>.
36. Geiping, J., Bauermeister, H., Dröge, H., & Moeller, M. (2020). Inverting gradients – How easy is it to break privacy in federated learning? In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), Advances in Neural Information Processing Systems (Vol. 33, pp. 16937–16947). Curran Associates, Inc.