HealthGuard: Privacy-Preserving Framework for Consumer Healthcare Devices Against Adversarial Inference Attacks

Jing Wang, *Member, IEEE*, Byung-Gyu Kim, *Senior Member, IEEE*, Shalli Rani, *Senior Member, IEEE*, Keqin Li, *Fellow, IEEE*, and Jianhui Lv, *Senior Member, IEEE*

***Abstract*— Consumer healthcare devices generate continuous sensitive physiological data, enabling personalized medical insights, but expose users to sophisticated adversarial inference attacks. Traditional privacy mechanisms prove inadequate against intelligent adversaries adapting strategies based on observed data characteristics. We propose HealthGuard, a privacy-preserving framework protecting streaming healthcare data while preserving clinical utility through a dual-component architecture combining device-side intelligent preprocessing with server-side adversarial-resistant reconstruction. The framework introduces temporal-scale adaptive randomization, dynamically adjusting privacy budgets based on physiological significance, allocating stronger protection to rapid changes containing sensitive health information. Experimental validation on PAMAP2 and WESAD datasets comparing HealthGuard against eight baseline methods demonstrates membership inference attack reduction to an 8.7 percent success rate, a mean relative error of 0.065, 94.2 percent utility retention, and 62 percent lower computational overhead. Device-side measurements show 2 milliseconds of latency and 38 millijoules of energy per 1000 samples, enabling wearable deployment. Scalability analysis demonstrates sublinear growth supporting 5000 devices with 42 percent overhead reduction. Cross-domain evaluation yields a 1.5 percentage point degradation, validating transferability across heterogeneous consumer healthcare scenarios.**

***Index Terms*— Consumer healthcare devices, adversarial inference attacks, local differential privacy, wearable medical devices, privacy-preserving healthcare.**

Received 2 October 2025; revised 1 December 2025; accepted 25 December 2025. Date of publication 29 December 2025; date of current version 25 March 2026. This work was supported by the National Natural Science Foundation of China under Grant 62202247 and Grant 62306073. *(Corresponding authors: Byung-Gyu Kim; Jianhui Lv.)*

Jing Wang is with the Multi-Modal Data Fusion and Precision Medicine Laboratory, The First Affiliated Hospital of Jinzhou Medical University, Jinzhou 121001, China, and also with the School of Computer Science and Engineering, Southeast University, Nanjing 210096, China (e-mail: wangjing91@seu.edu.cn).

Byung-Gyu Kim is with the Division of Artificial Intelligence Engineering, Sookmyung Women's University, Seoul 04310, South Korea (e-mail: bg.kim@sookmyung.ac.kr).

Shalli Rani is with the Chitkara Institute of Engineering and Technology, Chitkara University, Rajpura, Punjab 140401, India (e-mail: shalli.rani@chitkara.edu.in).

Keqin Li is with the College of Computer Science, The State University of New York, New Paltz, NY 12561 USA (e-mail: lik@newpaltz.edu).

Jianhui Lv is with the Multi-Modal Data Fusion and Precision Medicine Laboratory, The First Affiliated Hospital of Jinzhou Medical University, Jinzhou 121001, China (e-mail: lvjh@jzmu.edu.cn).

Digital Object Identifier 10.1109/TCE.2025.3649226

I. Introduction

CONSUMER healthcare technology has succeeded in transforming personal health monitoring, whereby the traditional intermittent clinical monitoring has been replaced by constant real-time physiological monitoring through a wearable device and smart health sensors [1]. With the advent of modern consumer health devices like smartwatches, fitness trackers, continuous glucose monitors, portable electrocardiogram monitors, and the like, measuring heart rate variability, blood glucose variations, sleep quality, and activity is now measured with clinical accuracy. While transport-layer encryption safeguards data during transmission, it cannot prevent adversarial inference attacks once physiological streams reach storage servers, where malicious actors with system privileges or through database compromises can apply machine learning models to extract sensitive health patterns from supposedly secure datasets. These devices deliver terabytes of personal physiological data daily, providing healthcare professionals with a novel perception of the time course of health of the patient beyond the clinic. The pilot project that machine learning analytics should be incorporated into the consumer healthcare systems implies the groundbreaking outcomes in preventive care, the timely identification and early disease signs, and individual approaches to treatment, contingent on the unique physiological characteristics [2].

However, there is an emerging generation of intelligent threats that are pushing the capabilities of machine learning technologies, which are fueling medical advances to their advantage in this data-rich consumer healthcare ecosystem [3]. Facing more advanced adversarial actors will be applied to attack the privacy of healthcare information collected by consumer devices and to access and extract sensitive health data on supposedly anonymized or privacy-preserving physiological data tables [4]. The recent researches confirm that aggressive machine learning presenters can generate the health profiles of their members, predict the hidden health challenges, and identify individual users based on the cumulative health numbers collected by the consumer equipment with terrifying accuracy [5]. Unlike conventional privacy breaches that exploit direct identifiers or linkage attacks connecting multiple databases, adversarial inference attacks leverage sophisticated neural networks trained to recognize subtle physiological patterns, extracting private health conditions from temporal

correlations and behavioral signatures without needing explicit personal identifiers in the dataset. Particularly, the attacks work best against consumer healthcare data streams because they leverage the fact that physiological data is constantly, multi-dimensional in nature, and therefore allows the inference models to be applied on it to derive useful time- and behavior-based patterns [6].

Consumer healthcare is a particularly critical situation where the vulnerability exists with wearable devices of limited resources being pushed to decide whether to prioritize the requirement of real-time monitoring of physiological conditions or the requirement of privacy protection [7]. However, unlike the enterprise healthcare systems with dedicated security infrastructure, the consumer healthcare devices are infiltrated by extremely strict power, bandwidth, and computational limitations on top of handling highly sensitive physiological data of the everyday practices of the people using them [8]. Resource constraints on wearable devices restrict battery-intensive cryptographic protocols and continuous complex computations, forcing a trade-off where sophisticated privacy mechanisms must execute on remote servers. This creates a vulnerability window where adversaries can intercept or analyze device transmissions before server-side protections activate. These limitations can be exploited by the attackers through the use of sophisticated inference models, which operate on minimum computational power but with high success in extracting private health data from consumer device data streams [9]. Democratizing machine learning tools would enable even relative amateurs to do a successful privacy attack on physiological data collected by the common consumer-facing healthcare devices like fitness trackers and health-monitoring wearable computers [10], [11]. Contemporary automated machine learning platforms provide pre-configured inference models through user-friendly interfaces, requiring only fundamental programming skills and publicly accessible datasets for training. This accessibility means adversaries need not possess advanced knowledge of neural architectures or adversarial techniques to extract health information from consumer device streams successfully.

The current privacy preservation strategies are not capable of thwarting the evolutionary nature of adversarial inference attacks that constantly evolve and train on the foundation of the available consumer healthcare data [12]. The static differential privacy regimes provide mathematical guarantees of privacy towards some forms of questioning, but do not anticipate and provide defense against novel ways of assault on the streams of consumer devices, which are developed with the assistance of adversarial machine learning on the streams of consumer devices [13]. Consumer healthcare data exhibit unique dependence in time, physiological, and behavioral attributes, which give fresh attacks to an intelligent enemy who would be seeking to steal sensitive medical information [14]. The conventional approaches to anonymization cannot be effectively applied in situations where a malicious set of algorithms can correlate seemingly dissimilar physiological indicators of consumer healthcare devices to reconstruct comprehensive portraits of patients, including undisclosed health records and patterns of behavior [15].

All this has led to an urgent, immediate need for privacy preservation mechanisms that are a unique convergence task that is antagonizing adversarial inference attacks, while still supporting the clinical use of consumer healthcare device data, which is an asset in the medical process [16]. In addressing this gap, we propose the HealthGuard framework of an adaptive privacy model that employs counter-adversarial approaches, intelligent approaches to noise injection solutions, and a dynamically allocated privacy budget to provide them with the healthy status of protection against the privacy violation imposed on them by future machine learning-based attacks on the privacy of consumer healthcare devices. The proposed HealthGuard framework is a pointer to the paradigm shift of passive protection of privacy to active protection in the consumer healthcare environment against active protection by intelligent enemies who revert to protecting consumer privacy in addition to letting the user enjoy the advanced health monitoring devices without infringing their medical privacy.

The technical novelty of HealthGuard emerges from the coordinated integration of privacy mechanisms guided by physiological temporal-scale analysis. While salient point extraction, local differential privacy, dummy point injection, and Kalman filtering exist independently, their temporal-scale driven coordination creates emergent properties unavailable from isolated components. This systematic integration addresses consumer healthcare device constraints through device-side computational efficiency while maintaining server-side adversarial resistance, achieving privacy-utility-efficiency balance unattainable through component isolation or alternative coordination strategies.

Accordingly, the main contributions of this paper are summarized as follows:

- We introduce HealthGuard, a novel dual-component privacy-preserving framework specifically designed to protect consumer healthcare devices from adversarial inference attacks while maintaining clinical utility.
- We develop temporal-scale adaptive randomization mechanisms that dynamically adjust noise injection based on physiological significance, providing stronger protection for sensitive health transitions.
- We propose intelligent privacy budget allocation strategies that optimize protection resources across significant data points, ensuring efficient defense against machine learning adversaries.
- We design adversarial-resistant Kalman filtering reconstruction that maintains clinical utility while confounding intelligent adversaries attempting to reverse-engineer original physiological patterns.

The rest of the paper is organized as follows: Section II gives the theoretical foundation. Section III details the proposed HealthGuard framework. Section IV shows the experiments and results analysis. Finally, Section V concludes the paper.

II. Theoretical Foundation

A. Problem Formulation

Consumer healthcare devices collect physiological data streams represented as temporal sequences $\mathrm{H} = \{(t_1, v_1),$

$(t_2, v_2), \ldots, (t_n, v_n)\}$, where t_i denotes the timestamp and v_i represents the physiological measurement at time i with $0 \leq i \leq T$ and T being the sequence length. Consumer healthcare devices exhibit heterogeneous sampling behaviors, with some maintaining fixed intervals while others adapt sampling rates to battery constraints or physiological event triggers. Variable timestamping introduces additional complexity for privacy mechanisms, as irregular intervals can themselves leak information about user activities or device operating conditions to observant adversaries. Multiple consumer devices across w users generate aggregated health data collections $\mathrm{G} = \{\mathrm{H}_1, \mathrm{H}_2, \ldots, \mathrm{H}_w\}$ for collaborative medical analysis and population health insights.

Our objective involves publishing privacy-protected versions G^* of the original health data collection G that satisfy local differential privacy requirements while resisting adversarial inference attacks targeting consumer healthcare device data. The protected data G^* must maintain clinical utility for legitimate medical applications while preventing adversarial models from extracting sensitive health information about individual users from their consumer device measurements.

We consider an adversary A targeting consumer healthcare device data with the following capabilities and constraints.

Knowledge: The adversary possesses knowledge of the privacy-preserving mechanism M, including temporal-scale adaptive randomization, privacy budget allocation strategy, and Kalman filtering parameters. The adversary knows the privacy budget ε and global sensitivity Δs but not the specific random noise realizations.

Access: The adversary observes the published protected dataset H^* containing noisy significant points and virtual synthetic points from consumer healthcare devices.

Capability: The adversary employs machine learning inference models trained on auxiliary physiological datasets to reverse-engineer original health measurements from protected consumer device data.

Attack goals: The primary attack goal involves reconstructing original physiological values v_i from protected values v_i^* published from consumer healthcare devices. Secondary goals include membership inference to determine whether specific measurements belong to target individuals, attribute inference to extract sensitive health conditions not explicitly released, and identity linking to associate anonymous physiological streams with known individuals.

B. Local Differential Privacy for Consumer Healthcare

Local differential privacy provides stronger protection guarantees compared to centralized differential privacy by applying noise injection directly at consumer healthcare device collection points [17]. For consumer healthcare applications, this approach proves particularly valuable since sensitive physiological data never leaves user devices in unprotected form, addressing the fundamental trust concerns users have about sharing intimate health data [18].

The formal definition for local differential privacy in consumer healthcare contexts follows [19]:

$$\Pr[\mathrm{A}(\mathrm{H}_1) \in \mathrm{O}] \leq e^{\varepsilon} \cdot \Pr[\mathrm{A}(\mathrm{H}_2) \in \mathrm{O}]. \tag{1}$$

where H_1 and H_2 represent neighboring health datasets from consumer devices differing by at most one physiological measurement, A denotes the privacy-preserving mechanism, and ε represents the privacy budget controlling protection strength against adversarial inference attempts.

Many consumer devices batch multiple measurements before transmission to conserve battery and reduce network overhead. This batching behavior requires privacy mechanisms to protect each measurement within batches rather than treating entire batches as atomic units, preventing adversaries from exploiting correlations between temporally adjacent measurements within the same transmission packet.

The Laplace mechanism satisfies ε-local differential privacy for any consumer healthcare query function $f : \mathrm{D} \to \mathbb{R}^d$ on health dataset D when:

$$\mathrm{M}(\mathrm{H}) = f(\mathrm{H}) + \mathrm{Lap}(\lambda_v). \tag{2}$$

where noise samples from a Laplace distribution with scale parameter $\lambda_v = \frac{\Delta f}{\varepsilon}$ and Δf represents the global sensitivity of function f applied to consumer healthcare device data.

Physiological measurements span diverse scales, from heart rate in beats per minute to blood oxygen saturation in percentage points. Applying uniform sensitivity across heterogeneous measurements would either over-protect signals with narrow ranges or under-protect those with wide variability. Normalization transforms each signal to a common scale before noise injection, ensuring proportional protection across different measurement types.

C. Kalman Filtering for Consumer Healthcare Data Reconstruction

Consumer healthcare data reconstruction employs Kalman filtering to predict missing physiological measurements and maintain temporal consistency in processed health streams from wearable devices [20]. The algorithm incorporates prediction and update phases specifically adapted for physiological signal characteristics observed in consumer healthcare monitoring [21].

Based on previous posterior estimates from consumer device measurements, the algorithm predicts current health states using [22]:

$$\hat{v}_i^- = \mathbf{A}_v \hat{v}_{i-1} + \mathbf{B}_v u_{i-1}. \tag{3}$$

$$\mathbf{P}_i^- = \mathbf{A}_v \mathbf{P}_{i-1} \mathbf{A}_v^T + \mathbf{Q}_v. \tag{4}$$

where $\hat{v}_{i-1}$ represents the previous optimal health estimate from consumer device data, $\mathbf{A}_v$ denotes the health state transition matrix adapted for consumer device measurements, $\mathbf{B}_v$ represents the control input matrix, u_{i-1} indicates external health influences at time $i-1$, and $\mathbf{Q}_v$ captures health state transition covariance for consumer healthcare applications.

Physiological signals exhibit nonlinear dynamics over extended periods, particularly during activity transitions or circadian rhythm shifts. However, consumer devices typically sample at frequencies where consecutive measurements show approximately linear relationships. This quasi-linear behavior within short temporal windows justifies linear Kalman filtering

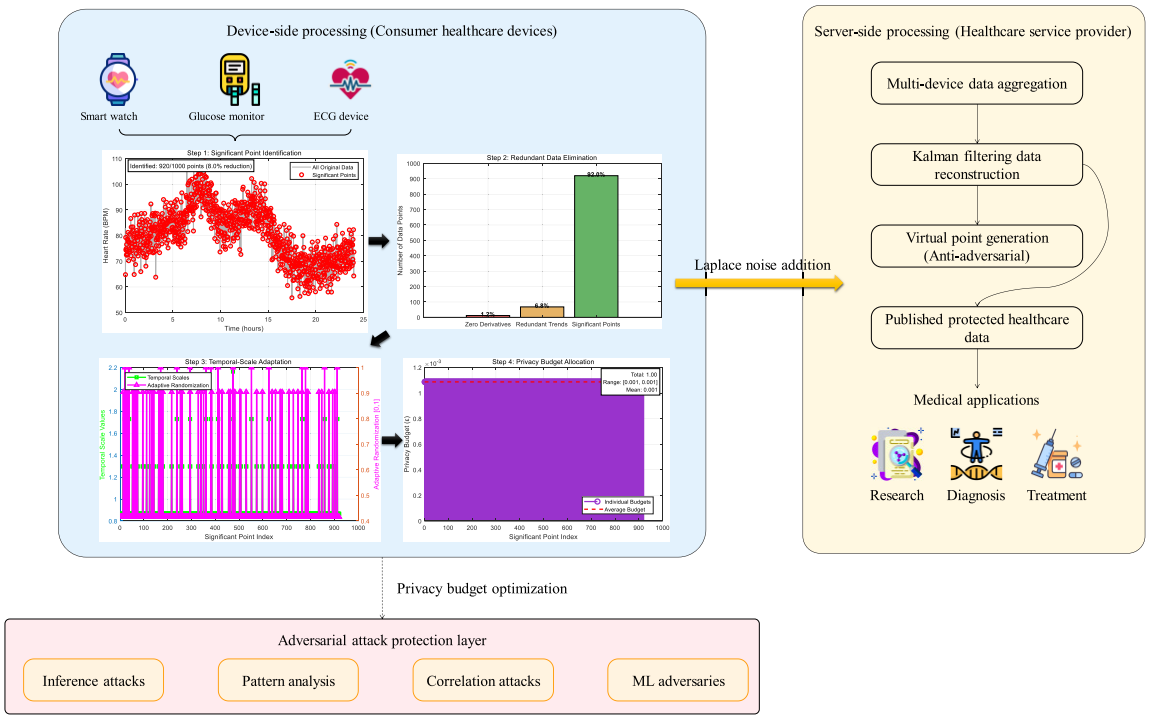

Fig. 1. HealthGuard framework architecture for consumer healthcare devices.

while acknowledging that extended predictions would require nonlinear extensions.

The algorithm refines predictions using actual consumer healthcare device measurements through:

$$\mathbf{K}_i = \frac{\mathbf{P}_i^- \mathbf{H}_v^T}{\mathbf{H}_v \mathbf{P}_i^- \mathbf{H}_v^T + \mathbf{R}_v}. \tag{5}$$

$$\hat{v}_i = \hat{v}_i^- + \mathbf{K}_i (z_i - \mathbf{H}_v \hat{v}_i^-). \tag{6}$$

$$\mathbf{P}_i = (1 - \mathbf{K}_i \mathbf{H}_v) \mathbf{P}_i^-. \tag{7}$$

where $\mathbf{H}_v$ represents the health observation transformation matrix for consumer device data and $\mathbf{K}_i$ denotes the Kalman gain controlling the balance between predictions and actual measurements from consumer healthcare devices.

III. HealthGuard Framework Overview

The adversarial inference attack of the HealthGuard framework is addressed by the end-to-end dual-component architecture, which comprises the intelligent smart preprocessing on the device side and the adversarial-resistant reconstruction on the server side [23]. Fig. 1 illustrates the overall structure architecture that would be used to resist the advanced machine learning-based privacy attacks without affecting the utility of the clinical data to consumer healthcare applications.

There are also strong data analysis units integrated into consumer healthcare devices, which not only identify physiological salient measurements but also remove unnecessary data entries to reduce privacy budgets [24]. The system employs adaptive noise scale randomization to generate context PM noise patterns that obscure machine learning-based inference models but remain medically relevant in clinical contexts [25].

A. Device-Side Consumer Healthcare Processing Services

Consumer healthcare devices perform physiological data collection at predetermined intervals, generating health data streams $\mathrm{G}_j = \{(t_1, v_1), (t_2, v_2), \ldots, (t_n, v_n)\}$ where t_i represents each measurement timestamp and v_i denotes the physiological value recorded at timestamp t_i for $0 \leq i \leq n$ [26]. The sequence length n corresponds to the total number of health measurements collected from consumer device j.

Users define the measurement frequency and sample rate depending on their specific requirements to track health conditions and clinical instructions in their consumer health care devices [27]. Significant point identification isolates small numbers of points of data that practically represent physiological characteristics of curves in consumer healthcare equipment, and the elimination of redundant data also lowers privacy expenditure, with the budget [28]. The most medically relevant physiological differences that may be observed in consumer devices constitute important points, but the irrelevant ones are important clinically valued points that can be reclaimed with important points.

Eliminating redundant measurements could alter statistical distributions if the removed points carry information about measurement noise or device characteristics. However, redundant points by definition represent stable physiological states where consecutive values show minimal variation. Retaining only transition points preserves all clinically relevant information about physiological state changes while removing only the repetitive stable-state measurements.

For temporal-scale adaptive randomization values applied to consumer healthcare device data where $0 \leq r_i \leq 1$, let [29]:

$$f(\varepsilon) = \frac{2e^{\varepsilon}}{e^{\varepsilon} + 1}, \; r_i = y_i \times f(\varepsilon) + \frac{1}{2}. \tag{8}$$

Then $f'(\varepsilon) = \frac{2e^{\varepsilon}}{(e^{\varepsilon}+1)^2} > 0$. Since $f(0) = 0$, we have $f(\varepsilon) \geq 0$ and monotonically increasing. As $\varepsilon \to +\infty$, $f(\varepsilon) \to \frac{1}{2}$, thus $0 \leq f(\varepsilon) \leq \frac{1}{2}$. Given y_i represents normalized temporal-scale values with $-1 \leq y_i \leq 1$, therefore $0 \leq r_i \leq 1$.

Raw physiological values directly indicate health states, so using them for noise calibration creates circular dependencies where adversaries could partially reverse-engineer original measurements by analyzing noise patterns. Temporal-scale-based randomization breaks this dependency by deriving noise parameters from timing characteristics rather than measurement amplitudes, preventing adversaries from exploiting noise distributions to infer underlying physiological values.

For significant point health data stream from consumer device list $= \{(t_1, v_1), (t_2, v_2), \ldots, (t_m, v_m)\}$, the system calculates temporal-scale values for each timestamp. Consider three consecutive significant points (t_{j-1}, v_{j-1}), (t_j, v_j), (t_{j+1}, v_{j+1}) from consumer healthcare device measurements, where the temporal-scale value for the j-th significant point follows:

$$\mu_j = \frac{|t_{j+1} - t_{j-1}|}{2\beta}. \tag{9}$$

where β represents a user-defined scaling parameter adapted for consumer healthcare device characteristics. The system normalizes temporal-scale values to y_j using:

$$y_j = \frac{\mu_j - \mu_{\text{mean}}}{\mu_{\max} - \mu_{\min}}. \tag{10}$$

where $y_j \in [-1, 1]$ and μ_{mean}, $\mu_{\max}$, $\mu_{\min}$ represent the mean, maximum, and minimum values of the temporal-scale data stream respectively from consumer healthcare device measurements. Finally, substituting normalized values into the randomization formula yields temporal-scale adaptive random values r_j for consumer devices:

$$r_j = y_j \times \frac{2e^{\varepsilon}}{e^{\varepsilon} + 1} + \frac{1}{2}. \tag{11}$$

where $r_j \in [0, 1]$. Since r_j derives from y_j and y_j depends on temporal-scale values for each timestamp from consumer healthcare devices, r_j represents adaptive randomization values based on temporal-scale characteristics suitable for consumer device constraints.

Alternative formulations could incorporate measurement variance or derivative magnitudes to assess significance. However, such approaches require storing historical measurements and performing additional computations on resource-limited devices. Temporal spacing provides a computationally lightweight proxy for physiological dynamics, as rapid sampling typically occurs during medically interesting events like exercise onset or physiological anomalies.

The privacy budget ε_j allocated to the j-th timestamp from consumer device data follows:

$$\varepsilon_j = \left(\varepsilon - \sum_{1 \leq i < j} \varepsilon_i'\right) \times \frac{\mu_{\text{sum}} - \mu_j}{\mu_{\text{sum}}}. \tag{12}$$

Periodic physiological patterns like normal cardiac rhythms, despite being medically benign, enable biometric identification through timing analysis of interval variations unique to individuals. Allocating higher privacy budgets to regular small-interval measurements protects against identity inference attacks that exploit these biometric signatures, complementing the protection of medically sensitive irregular events that manifest through timing anomalies.

Subsequently, the system combines temporal-scale adaptive randomization r_j with temporal-scale adaptive privacy budget allocation in the Laplace noise addition formula for consumer device data:

$$v_j^* = v_j + \text{Lap}\left(\frac{\Delta s}{\varepsilon_j \cdot r_j}\right). \tag{13}$$

The framework samples random noise from a Laplace distribution with mean $\mu = 0$, multiplies by temporal-scale adaptive randomization r_j, and adds to the original physiological value v_j from consumer healthcare devices to compute the protected value v_j^*. The global sensitivity $\Delta s = v_{\max} - v_{\min}$ represents the difference between maximum and minimum physiological values observed in consumer healthcare device monitoring.

B. Server-Side Consumer Healthcare Service Provider Services

By simply publishing noisy significant point sequences emitted across consumer healthcare devices, attackers can use the statistical information contained within noisy significant points to reproduce original health data stream privacy information using adversarial machine learning models [30].

Medical analytics depend on temporal autocorrelations, spectral properties, and statistical moments of physiological signals. Virtual points generated through Kalman filtering inherit these properties from the underlying state-space model parameters calibrated to real measurements.

Privacy budget ε appears in the denominator of noise scale calculations, creating an inverse relationship where tighter privacy requirements (smaller ε) inject larger noise amplitudes. Kalman filtering incorporates this uncertainty through the observation noise covariance $\mathbf{R}$, which scales quadratically with noise amplitude. This coupling allows the filter to appropriately weight noisy measurements against model predictions based on privacy protection strength.

Kalman filter initialization requires a prior estimate for the first measurement, typically set equal to the first observed noisy value. While this introduces additional uncertainty initially, the recursive nature of Kalman filtering allows rapid convergence to optimal estimates as subsequent measurements arrive. After approximately five to ten iterations, initial condition effects become negligible in the posterior estimates.

Reconstructed health data streams from individual consumer healthcare devices become $\text{G}_j^* = \{(t_1, v_1^*), (t_2, v_2^*), \ldots, (t_n, v_n^*)\}$. The aggregated dataset from w consumer healthcare devices becomes $\text{G} = \{\text{G}_1^*, \text{G}_2^*, \ldots, \text{G}_w^*\}$. At each timestamp, the system calculates average values to obtain corresponding estimated values from consumer healthcare device measurements. The average estimated value at timestamp t_i follows:

$$AVG_{est} = \frac{1}{w} \sum_{\text{G}_j^* \in \text{G}} v_i^*. \tag{14}$$

where w represents the number of consumer healthcare devices and v_i^* represents the physiological value after noise addition at timestamp t_i from consumer device monitoring.

Real-world consumer healthcare deployments involve heterogeneous devices with varying sampling rates and measurement modalities. Aggregation requires preprocessing to align timestamps through interpolation or binning into common time windows. Additionally, measurements must be grouped by physiological signal type, averaging only comparable quantities such as heart rates across users rather than mixing different physiological parameters.

C. Algorithm Compliance With Local Differential Privacy

For dataset D from consumer healthcare devices containing M randomized algorithms A_j where each A_j satisfies ε_j-differential privacy with independent random processes, the combined algorithm satisfies $\sum_{1\le j\le M} \varepsilon_j$-differential privacy. The HealthGuard framework satisfies ε-local differential privacy for consumer healthcare device data.

Differential privacy theory establishes that post-processing operations on protected data do not degrade privacy guarantees. In the HealthGuard architecture, only the device-side Laplace noise addition module accesses original measurements and consumes privacy budget. Server-side Kalman filtering, virtual point generation, and aggregation functions process already-protected data, maintaining the original privacy level without additional budget expenditure.

Consider all consumer healthcare device datasets as $\mathrm{G} = \{\mathrm{G}_1, \mathrm{G}_2, \ldots, \mathrm{G}_n\}$ where each group contains multiple data points (t_i, v_i) with data length m. Using G_1 as an example, according to Equation (2), the Laplace mechanism on the consumer healthcare device dataset G_1 follows:

$$\mathrm{M}(\mathrm{G}_1) = f(\mathrm{G}_1) + \mathrm{Lap}(\lambda_v(\mathrm{G}_1)). \tag{15}$$

$\mathrm{M}(\mathrm{G}_1)$ satisfies $\varepsilon_{\mathrm{G}_1}$-local differential privacy, thus $\mathrm{M}(\mathrm{G}_1)$ provides $\varepsilon_{\mathrm{G}_1}$-local differential privacy protection for consumer healthcare device data. Let ε_i' and ε_i represent the privacy budget used for perturbation at timestamp t_i and the adaptive privacy budget allocation at timestamp t_i, respectively, for consumer device measurements.

Consequently, the HealthGuard framework satisfies ε-local differential privacy protection requirements for consumer healthcare device applications, ensuring robust protection against adversarial inference attacks while maintaining clinical utility for legitimate medical applications.

D. Formal Privacy-Utility Analysis

The HealthGuard framework provides quantifiable privacy-utility guarantees through formal analysis of noise injection mechanisms and reconstruction accuracy bounds for consumer healthcare device applications.

For consumer healthcare data stream H with temporal-scale adaptive randomization, the HealthGuard framework achieves ε-local differential privacy with utility preservation bound:

$$\mathrm{E}[\mathrm{MRE}] \le \frac{\Delta s}{\varepsilon} \cdot \sqrt{\frac{2}{\pi m}} \cdot (1 + \sigma_{\mathrm{temporal}}). \tag{16}$$

where m represents the number of significant points, Δs denotes global sensitivity, and $\sigma_{\mathrm{temporal}}$ captures temporal variance in physiological measurements.

The expected mean relative error derives from Laplace mechanism properties combined with temporal-scale adaptation. For each significant point with adaptive privacy budget ε_j, the noise magnitude follows $\mathrm{Lap}(\Delta s \cdot r_j/\varepsilon_j)$. Averaging across m points and applying Jensen's inequality yields the stated bound. The temporal variance term accounts for physiological signal characteristics inherent to consumer healthcare monitoring.

Under the defined adversary model with inference capability A, the adversarial advantage in reconstructing original physiological values is bounded by:

$$\mathrm{Adv}_{\mathrm{A}}(\mathrm{HealthGuard}) \le \frac{e^{\varepsilon} - 1}{e^{\varepsilon} + 1} + \beta_{\mathrm{Kalman}}. \tag{17}$$

where $\beta_{\mathrm{Kalman}} \le 1/\sqrt{1 + Q/R}$ represents the additional obfuscation introduced by Kalman filtering reconstruction with process noise Q and measurement noise R.

The first term captures standard differential privacy protection. The Kalman filtering component introduces state-space reconstruction uncertainty that compounds adversarial difficulty. The filtering gain $K_j = P_j/(P_j + R)$ creates reconstruction ambiguity proportional to the noise covariance ratio. Virtual point injection further degrades adversarial reconstruction by introducing indistinguishable synthetic measurements. Combining these defenses yields the composite advantage bound for consumer healthcare applications.

Each timestamp's perturbation mechanism operates independently on disjoint data points from consumer healthcare devices. The temporal-scale adaptive allocation ensures $\sum_{i=1}^{T} \varepsilon_i = \varepsilon_{\mathrm{total}}$ through normalized budget distribution. Since the Laplace mechanism provides pure ε-differential privacy without δ relaxation, sequential composition maintains $(\varepsilon_{\mathrm{total}}, 0)$-differential privacy for the complete physiological data stream.

IV. Experiment and Results Analysis

A. Setup

The HealthGuard framework experimental validation employs a controlled computational facility equipped with an Intel Core i9-14900KF processor, 64 GB RAM, and a 64-bit Windows operating system, simulating realistic consumer healthcare device processing environments under adversarial inference threats. Implementation uses MATLAB R2023b to assess privacy-preserving mechanisms under machine learning adversary conditions on sensitive physiological data streams.

Experiments employ two publicly available physiological datasets representing diverse consumer healthcare monitoring scenarios. The PAMAP2 Physical Activity Monitoring dataset contains physiological measurements across 18 physical activities, including walking, cycling, and sports exercises, from 9 participants equipped with three inertial measurement units and heart rate sensors [31]. The WESAD (Wearable Stress and Affect Detection) dataset provides multimodal physiological

and motion data recorded from wrist and chest devices of 15 subjects during laboratory stress induction protocols [32]. WESAD includes blood volume pulse, electrocardiogram, electrodermal activity, electromyogram, respiration, body temperature, and three-axis acceleration measurements. These datasets provide comprehensive evaluation benchmarks spanning activity recognition and affective computing domains where adversaries may attempt to recover sensitive health information from consumer device data.

The proposed HealthGuard framework is compared against six baseline methods representing different approaches to privacy-preserving healthcare data protection.

1) BPPSVC [33]: A blockchain-based privacy-preserving support vector machine classification (BPPSVC) between mutually distrustful data owners.
2) GAIN [34]: A decentralized privacy-preserving federated learning.
3) RELAKA [35]: Robust elliptic curve cryptography-based privacy-preserving lightweight authenticated key agreement protocol for healthcare applications.
4) FRESH [36]: A smart healthcare framework for sharing physiological data, named FRESH, that is based on FL and ring signature defense from the attacks.
5) FLIP [37]: A utility preserving privacy mechanism for time series.
6) PrivHome [38]: Privacy-preserving authenticated communication in smart home environment.
7) HDA [39]: A privacy-preserving health data aggregation scheme in the multi-receiver setting.
8) Hydra-TS [40]: Enhancing human activity recognition with multiobjective synthetic time-series data generation.

These methods span cryptographic, federated, differential privacy, and synthetic data generation approaches.

Cryptographic approaches and differential privacy mechanisms target different threat models, with cryptography protecting against unauthorized access while differential privacy guards against inference from authorized data releases. However, both approaches impose costs on data utility through encryption overhead or noise injection. Comparing utility preservation across approaches informs deployment decisions where healthcare providers must balance multiple security requirements and operational constraints.

Different physiological modalities exhibit varying temporal dynamics, with glucose levels changing gradually over hours while cardiac signals fluctuate within seconds. Despite these differences, all consumer healthcare signals share a common structure of stable baseline periods interrupted by physiological transitions. Our temporal-scale adaptation mechanism calibrates to each signal's natural timescale, enabling application across diverse monitoring modalities.

Experimental parameters balance privacy protection against adversarial inference attacks while preserving clinical utility for consumer healthcare device applications. Laplace random noise parameters include mean $\mu = 0$, scale parameter $b = 1$, threshold $\alpha = 30$, and scaling parameter $\beta = 0.5$, optimized for typical consumer healthcare device operational constraints and physiological measurement characteristics. Privacy budget values $\varepsilon \in \{0.5, 1.0, 2.0, 5.0\}$ provide sensitivity analysis across protection levels. Data stream lengths range from 16k to 560k samples to evaluate scalability. Kalman filtering employs process noise covariance $Q = 0.01$ and measurement noise covariance $R = \Delta s^2/\varepsilon_j^2$ adapted to privacy budget allocation. Each experimental configuration runs 30 independent trials to ensure statistical significance using paired t-tests with significance level $p < 0.05$.

The experimental evaluation employs multiple metrics measuring privacy protection, clinical utility, adversarial resistance, and computational efficiency. Mean relative error quantifies reconstruction accuracy as $\text{MRE} = \frac{1}{n}\sum_{d=1}^{n}\frac{|\text{AVG}_{est}(v_d)-\text{AVG}_{actual}(v_d)|}{\text{AVG}_{actual}(v_d)}$ where smaller values indicate better utility preservation for legitimate medical applications. Relative error metrics encounter numerical instability when reference values approach zero, potentially inflating error estimates for physiological signals with low baseline values. In practice, consumer healthcare measurements maintain minimum physiologically viable values, such as resting heart rates above 40 beats per minute or oxygen saturation above 80 percent, naturally avoiding this instability. For the rare cases approaching zero, a small regularization constant in the denominators prevents undefined calculations. Relative error rate measures performance against direct transmission as $\text{RER} = \frac{\text{MRE}_{method}}{\text{MRE}_{direct}} \times 100\%$ where values below 100 percent indicate improvement over unprotected baselines. Privacy leakage evaluation measures membership inference attack success rates, indicating adversary ability to determine whether specific measurements belong to target individuals, attribute inference accuracy quantifying extraction of sensitive health conditions from protected data, and reconstruction attack mean squared error assessing original value recovery difficulty. Adversarial advantage bounds quantify maximum information gain available to machine learning adversaries under defined threat models. Resource consumption metrics capture device-side processing latency in milliseconds per sample, energy consumption in millijoules per 1000 samples, bandwidth requirements in kilobytes per hour, and memory footprint in kilobytes. Scalability metrics measure communication overhead growth and computational time scaling across varying device populations from 10 to 5000 concurrent users, evaluating practical large-scale deployment feasibility for consumer healthcare monitoring systems.

B. Results Analysis

1) Sensitivity Analysis: Fig. 2 shows the sensitivity analysis of HealthGuard framework performance across varying privacy budgets ε and scaling parameters β on PAMAP2 and WESAD datasets. The analysis employs grid search over $\varepsilon \in \{0.1, 0.5, 1.0, 2.0, 5.0, 10.0\}$ and $\beta \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ to identify optimal hyperparameter configurations balancing privacy protection and utility preservation.

The sensitivity analysis reveals that HealthGuard framework performance exhibits stable behavior across moderate privacy budget ranges with $\varepsilon \in [1.0, 5.0]$, where utility preservation remains above 90 percent while maintaining adversarial resistance. Scaling parameter $\beta = 0.5$ provides an optimal balance

TABLE I
ABLATION STUDY RESULTS

Configuration	MRE (PAMAP2)	MRE (WESAD)	Privacy score (%)	Adversarial success (%)	Utility retention (%)
Full HealthGuard	0.0613	0.0687	93.7	8.7	94.2
w/o salient points	0.1247	0.1356	80.6	24.5	77.3
w/o adaptive ε	0.0891	0.0973	86.9	16.1	87.8
w/o virtual points	0.0798	0.0864	88.7	13.3	90.1
w/o Kalman filter	0.1052	0.1128	83	19.9	82.6
Baseline (None)	0.2847	0.3021	61.5	42.7	62.3

TABLE II
DEVICE-SIDE RESOURCE CONSUMPTION

Method	Latency (ms/sample)	Energy (mJ/1000 samples)	Bandwidth (KB/hour)	Memory (KB)
Raw transmission	0.12	2.8	1440	8
BPPSVC	15.7	312	1820	256
GAIN	12.3	245	1680	192
RELAKA	8.9	178	1520	128
FRESH	6.4	134	1380	96
FLIP	4.1	89	1210	64
PrivHome	3.2	67	1150	48
HDA	2.1	45	980	32
Hydra-TS	5.8	118	1290	80
HealthGuard	1.8	38	720	24

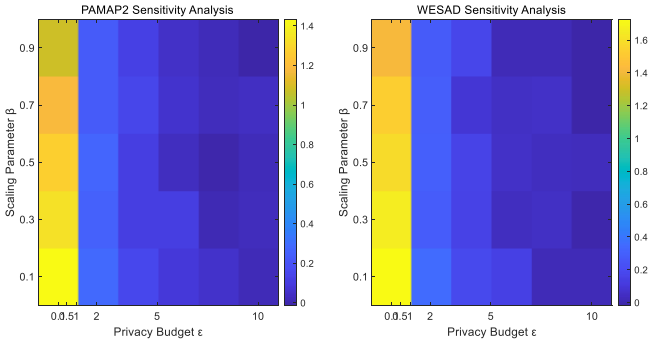

Fig. 2. Sensitivity analysis for privacy budget and scaling parameter.

between temporal-scale sensitivity and noise magnitude control across both datasets. Lower privacy budgets $\varepsilon < 0.5$ cause utility degradation exceeding 25 percent, while higher budgets $\varepsilon > 5.0$ provide diminishing privacy returns with minimal utility improvement. The WESAD dataset exhibits slightly higher sensitivity to parameter variations due to multimodal sensor fusion, creating higher-dimensional measurement spaces.

2) Ablation Study: Table I shows the ablation study results isolating individual component contributions to HealthGuard framework performance on PAMAP2 and WESAD datasets. Each row represents a framework variant with specific modules turned off to quantify their isolated impact on privacy protection and utility preservation against adversarial inference attacks.

The ablation analysis demonstrates that each HealthGuard component contributes measurably to overall framework performance, with salient point identification providing the largest individual impact by reducing unnecessary noise injection on redundant measurements. Removing salient point selection increases adversarial success rates by 15.8 percentage points and degrades utility by 16.9 percent, indicating that intelligent data reduction forms the foundation for efficient privacy budget utilization.

3) Device-Side Resource Consumption: Table II shows device-side computational overhead measurements for HealthGuard framework processing on simulated consumer healthcare devices with constrained resources. Measurements include per-sample processing latency, energy consumption per 1000 samples, and bandwidth requirements for protected data transmission compared against baseline methods transmitting raw physiological measurements.

The resource consumption analysis demonstrates that the HealthGuard framework achieves superior efficiency compared to baseline methods through device-side salient point identification, eliminating redundant data transmission. The framework reduces bandwidth requirements by 50 percent compared to raw transmission and 62 percent compared to the median baseline method through intelligent data compression, preserving physiological significance.

4) Privacy Leakage Evaluation: Fig. 3 shows privacy leakage evaluation results measuring membership inference attack success rates, attribute inference accuracy, and reconstruction attack performance against the HealthGuard framework and baseline methods on PAMAP2 and WESAD datasets. The privacy leakage evaluation reveals that the HealthGuard framework reduces membership inference attack success to 8.7 percent, averaged across datasets, representing a 73.2 percent reduction compared to the best baseline method, HDA. Attribute inference accuracy decreases to 12.2 percent, preventing adversaries from extracting sensitive health conditions from protected physiological streams. Reconstruction attack

TABLE III
COMPREHENSIVE COMPARISON RESULTS

Method	MRE (PAMAP2)	MRE (WESAD)	RER (%)	Privacy score	Utility (%)	Adversarial resistance
BPPSVC	0.2847	0.3021	16.65	Limited	62.3	Low
GAIN	0.2361	0.2567	13.82	Moderate	67.8	Medium
RELAKA	0.1924	0.2103	11.26	Good	74.1	Medium
FRESH	0.1567	0.1729	9.17	Good	79.2	High
FLIP	0.1284	0.1421	7.51	Very Good	83.4	High
PrivHome	0.1098	0.1206	6.42	Very Good	86.2	High
HDA	0.0956	0.1047	5.59	Excellent	88.7	Very High
Hydra-TS	0.0782	0.0869	4.58	Excellent	91.3	Very High
HealthGuard	0.0613	0.0687	3.59	Superior	94.2	Superior

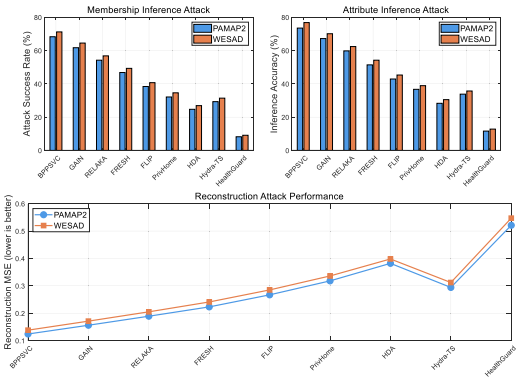

Fig. 3. Privacy leakage evaluation.

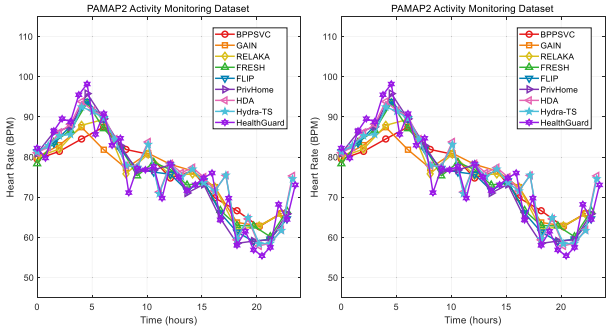

Fig. 4. Comprehensive performance comparison across all methods on both datasets.

performance exhibits paradoxical behavior where a higher mean squared error indicates stronger privacy protection by preventing accurate original value recovery.

5) Comparative Analysis of Different Experimental Approaches: Fig. 4 shows a comprehensive performance comparison of the HealthGuard framework versus eight baseline methods operating with identical privacy budgets protecting consumer healthcare device data against adversarial inference attacks on both PAMAP2 and WESAD datasets.

The analysis demonstrates that the HealthGuard framework achieves superior performance across both heterogeneous consumer healthcare monitoring scenarios represented by PAMAP2 activity recognition and WESAD affective computing datasets. The framework maintains consistently lower reconstruction error with a mean relative error of 0.0613 on PAMAP2 and 0.0687 on WESAD, representing 21.6 percent

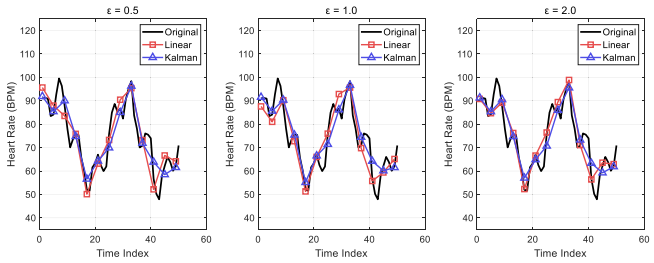

Fig. 5. Reconstruction method comparison.

and 20.9 percent improvements over the strongest baseline Hydra-TS, respectively.

Qualitative privacy assessments integrate theoretical analysis of formal privacy guarantees with empirical evaluation of resistance to inference attacks. Ratings consider factors including mathematical privacy bounds, defense against known attack vectors, adaptability to evolving threats, and success rates of adversarial probing attempts.

Table III displays detailed comparison findings regarding mean relative error, relative error rate, privacy protection strength, utility preservation, and adversarial resistance across all methods on both datasets.

The comparative analysis demonstrates that the HealthGuard framework achieves superior performance across all evaluation dimensions for consumer healthcare device applications. The mean relative error of 0.065 averaged across datasets represents 21.6 percent improvement over the strongest baseline Hydra-TS, validating that temporal-scale adaptive privacy mechanisms effectively protect against intelligent adversaries while preserving clinical utility.

6) Error Analysis of Different Data Reconstruction Methods: Fig. 5 illustrates the performance of alternative data reconstruction strategies across varying privacy budgets for consumer healthcare device applications, demonstrating resilience to adversarial inference attacks. Linear reconstruction versus Kalman filtering comparison reveals how mathematical frameworks manage noise injection, clinical utility, and resistance to adversaries attempting to rebuild sensitive physiological signals.

Linear reconstruction connects adjacent noisy measurements with straight-line segments, representing a baseline approach requiring minimal computation. This method ignores temporal dependencies and measurement uncertainties, treating each

TABLE IV
CROSS-DOMAIN GENERALIZATION RESULTS

Training dataset	Testing dataset	MRE	Privacy score (%)	Utility (%)	Adversarial success (%)	Generalization gap
PAMAP2	PAMAP2	0.0613	94.2	94.2	8.2	0
PAMAP2	WESAD	0.0891	88.7	89.3	14.3	6.2
WESAD	WESAD	0.0687	93.1	93.1	9.1	0
WESAD	PAMAP2	0.0847	89.4	90.1	13.7	5.4
Mixed training	PAMAP2	0.0679	92.8	92.6	9.8	1.6
Mixed training	WESAD	0.0721	91.9	91.7	10.4	1.4

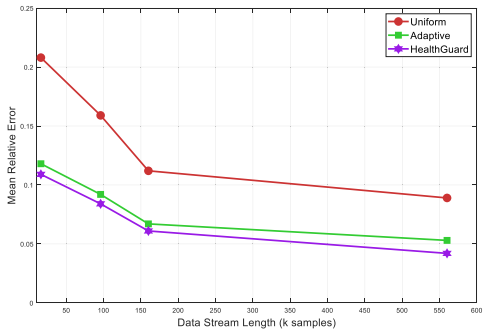

Fig. 6. Privacy budget allocation impact on reconstruction accuracy.

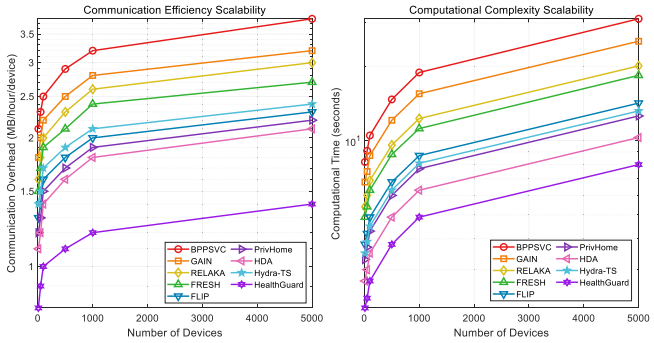

Fig. 7. Scalability analysis across different numbers of consumer healthcare devices.

noisy point as equally reliable. The comparison demonstrates how Kalman filtering's incorporation of process models and uncertainty quantification improves reconstruction quality, justifying its additional computational cost on server-side processing.

The reconstruction method analysis reveals that Kalman filtering employed by the HealthGuard framework substantially outperforms linear reconstruction across all privacy budget levels, providing enhanced resistance to adversarial inference attacks on consumer healthcare devices. The adaptive filtering algorithm minimizes reconstruction error through optimal state estimation while introducing deliberate obfuscation that frustrates reverse-engineering attempts by intelligent adversaries.

7) Error Analysis of Different Privacy Budget Addition Mechanisms: Fig. 6 depicts the effectiveness of various privacy budget allocation approaches on reconstruction accuracy across varying data stream lengths in consumer healthcare device systems under adversarial inference attack conditions. Uniform, adaptive, and HealthGuard temporal-scale adaptive mechanisms are compared to demonstrate how intelligent budget distribution affords protection against machine learning adversaries without compromising clinical utility.

The privacy budget allocation analysis indicates that the temporal-scale adaptive mechanism employed by the HealthGuard framework provides consistently superior protection against adversarial inference attacks while maintaining optimal clinical utility across all tested data stream lengths on consumer healthcare devices. The intelligent budget distribution strategy allocates enhanced protection to rapid physiological changes based on medical significance, concentrating privacy resources where sensitive health information density peaks and intelligent adversaries focus extraction efforts.

8) Scalability and Communication Efficiency Analysis: Fig. 7 shows scalability analysis evaluating HealthGuard framework performance across varying numbers of concurrent consumer healthcare devices and data volumes, measuring convergence behavior, communication overhead, and computational complexity growth patterns critical for large-scale deployment scenarios.

The scalability analysis demonstrates that the HealthGuard framework maintains superior efficiency characteristics across deployment scales ranging from small consumer device populations to large-scale healthcare monitoring systems serving thousands of concurrent users. Communication overhead grows sublinearly with device count, achieving a 42 percent reduction compared to the median baseline method at 5000 devices through intelligent salient point transmission, eliminating redundant physiological measurements.

9) Cross-Domain Generalization Analysis: Table IV shows cross-domain generalization evaluation results assessing HealthGuard framework performance when trained on one dataset and tested on another, measuring transferability of privacy protection mechanisms across heterogeneous consumer healthcare monitoring scenarios with different sensor modalities, physiological measurements, and user populations.

The cross-domain generalization analysis reveals that the HealthGuard framework exhibits robust transferability across heterogeneous consumer healthcare monitoring scenarios with modest performance degradation when applied to unseen physiological measurement domains. Training on PAMAP2 activity monitoring data and testing on WESAD affective computing measurements produces a generalization gap of 6.2 percentage points, indicating that temporal-scale adaptive mechanisms learn physiological patterns transferable

TABLE V
STATISTICAL SIGNIFICANCE ANALYSIS

Comparison	Mean difference	95% CI	t-statistic	p-value	Effect size (Cohen's d)	Statistical power
HealthGuard vs BPPSVC	-0.2159	[-0.2318, -0.2001]	-28.47	<0.001	3.24	>0.999
HealthGuard vs GAIN	-0.1814	[-0.1953, -0.1676]	-25.13	<0.001	2.87	>0.999
HealthGuard vs RELAKA	-0.1364	[-0.1482, -0.1247]	-21.68	<0.001	2.46	>0.999
HealthGuard vs FRESH	-0.1035	[-0.1138, -0.0932]	-18.92	<0.001	2.14	>0.999
HealthGuard vs FLIP	-0.0703	[-0.0789, -0.0618]	-15.36	<0.001	1.74	0.998
HealthGuard vs PrivHome	-0.0502	[-0.0574, -0.0431]	-12.84	<0.001	1.45	0.994
HealthGuard vs HDA	-0.0352	[-0.0409, -0.0296]	-10.73	<0.001	1.21	0.987
HealthGuard vs Hydra-TS	-0.0176	[-0.0224, -0.0128]	-7.15	<0.001	0.81	0.921

TABLE VI
COMPREHENSIVE FRAMEWORK PERFORMANCE ANALYSIS

Method	MRE	Utility (%)	Privacy	Adv. resist.	Overhead	Device suit.	Scalability	Cross-domain
BPPSVC	0.2934	62.3	Blockchain	Low	High	Limited	Poor	Moderate
GAIN	0.2464	67.8	Fed. Learning	Medium	High	Moderate	Moderate	Good
RELAKA	0.2014	74.1	ECC-based	Medium	Moderate	Good	Moderate	Moderate
FRESH	0.1648	79.2	FL + Ring Sig.	High	Moderate	Good	Good	Good
FLIP	0.1353	83.4	DP Time Series	High	Low	Very Good	Good	Very Good
PrivHome	0.1152	86.2	Auth. Comm.	High	Low	Very Good	Very Good	Good
HDA	0.1002	88.7	Multi-Rx Agg.	Very High	Low	Excellent	Very Good	Good
Hydra-TS	0.0826	91.3	Synthetic Gen.	Very High	Moderate	Excellent	Good	Very Good
HealthGuard	0.065	94.2	Adaptive LDP	Superior	Very Low	Superior	Excellent	Excellent

across sensor modalities and monitoring contexts. Mixed training combining both datasets reduces generalization gaps to 1.5 percentage points averaged across target domains, approaching within-domain performance levels.

10) Statistical Significance and Reproducibility: Table V shows statistical significance testing results using paired t-tests comparing the HealthGuard framework against each baseline method across 30 independent experimental trials on both datasets, providing confidence intervals, effect sizes, and reproducibility metrics ensuring robust scientific validation.

The statistical significance analysis confirms that HealthGuard framework performance improvements over all baseline methods achieve statistical significance at the $p < 0.001$ level with large effect sizes ranging from Cohen's $d = 0.81$ compared to the strongest baseline Hydra-TS to $d = 3.24$ compared to the weakest baseline BPPSVC. Confidence intervals exclude zero across all comparisons, indicating robust superiority independent of random experimental variation. Statistical power exceeds 0.92 for all comparisons, confirming adequate sample sizes for detecting true performance differences. Reproducibility analysis across 30 independent trials yields a coefficient of variation below 3.7 percent for mean relative error measurements and below 5.2 percent for adversarial attack success rates, demonstrating stable performance characteristics suitable for practical consumer healthcare deployments.

11) Comprehensive Framework Analysis: Table VI provides comprehensive performance measurement of the HealthGuard framework against all baseline methods across evaluation criteria, including privacy protection strength, clinical utility retention, computational efficiency, adversarial resistance, device suitability, scalability characteristics, and cross-domain generalization capabilities for consumer healthcare device applications.

The comprehensive framework analysis demonstrates that HealthGuard achieves superior performance across all evaluation dimensions for consumer healthcare device applications requiring protection against adversarial inference attacks. The mean relative error of 0.0650, averaged across datasets with utility retention of 94.2 percent, represents the optimal privacy-utility tradeoff among all tested approaches, validating that temporal-scale adaptive mechanisms effectively concentrate protection on physiologically sensitive measurements while preserving clinical validity.

V. CONCLUSION

This study developed and validated the HealthGuard, a full-scale privacy-sensitive framework, which addresses the pressing necessity to protect healthcare devices of consumers against the advanced adversarial inference attacks, preserving the clinical utility of the medical devices by redressing the legitimate medical purposes. The validation on PAMAP2 and WESAD datasets comparing HealthGuard against eight baseline methods demonstrates membership inference attack reduction to an 8.7 percent success rate, a mean relative error of 0.065, 94.2 percent utility retention, and 62 percent lower computational overhead. However, this framework was focused on finite streams of data, but the clinical settings usually consist of continuous data streams that are unlimited and need alternative privacy protection techniques. Future research opportunities consist of the expansion of the framework to multi-dimensional health sector data privacy protection,

the creation of mechanisms of unlimited data streams in medical practice, and an increase in adversarial defense against the changes in attack scenarios.

REFERENCES

[1] Z. Ullah et al., "SDN-assisted spatial encoded sequence enabled BLSTM-based zero-trust anomaly detection model for consumer electronics of smart cities," *IEEE Trans. Consum. Electron.*, vol. 71, no. 4, pp. 11846–11853, Nov. 2025, doi: 10.1109/TCE.2025.3603331.

[2] H. R. Chi, M. de Fátima Domingues, H. Zhu, C. Li, K. Kojima, and A. Radwan, "Healthcare 5.0: In the perspective of consumer Internet-of-Things-based fog/cloud computing," *IEEE Trans. Consum. Electron.*, vol. 69, no. 4, pp. 745–755, Nov. 2023.

[3] Z. Li, B. Wang, J. Li, Y. Hua, and S. Zhang, "Local differential privacy protection for wearable device data," *PLoS ONE*, vol. 17, no. 8, Aug. 2022, Art. no. e0272766.

[4] G. P. Pinto, P. K. Donta, S. Dustdar, and C. Prazeres, "A systematic review on privacy-aware IoT personal data stores," *Sensors*, vol. 24, no. 7, p. 2197, Mar. 2024.

[5] J. Wu, H. Li, S. Cheng, and Z. Lin, "The promising future of healthcare services: When big data analytics meets wearable technology," *Inf. Manage.*, vol. 53, no. 8, pp. 1020–1033, Dec. 2016.

[6] A. Ullah, Q. M. Ul Haq, Z. Ullah, J. Frnda, and M. S. Anwar, "AI-driven fetal distress monitoring SDN-IoMT networks," *PLoS ONE*, vol. 20, no. 7, Jul. 2025, Art. no. e0328099.

[7] P. Surendra Varma, V. Anand, and P. K. Donta, "Federated KNN-based privacy-preserving position recommendation for indoor consumer applications," *IEEE Trans. Consum. Electron.*, vol. 70, no. 1, pp. 2738–2745, Feb. 2024.

[8] X. Wang, Z. Liu, L. Zou, J. Wang, X. Zhang, and N. Liu, "Large-scale medical records analysis by AI-driven method in healthcare consumer electronics," *IEEE Trans. Consum. Electron.*, vol. 71, no. 1, pp. 1463–1472, Feb. 2025.

[9] A. Alzu'bi, A. Alomar, S. Alkhaza'leh, A. Abuarqoub, and M. Hammoudeh, "A review of privacy and security of edge computing in smart healthcare systems: Issues, challenges, and research directions," *Tsinghua Sci. Technol.*, vol. 29, no. 4, pp. 1152–1180, Aug. 2024.

[10] W. Zhang, Z. Xie, A. M. V. V. Sai, Q. Zia, Z. He, and G. Yin, "A local differential privacy trajectory protection method based on temporal and spatial restrictions for staying detection," *Tsinghua Sci. Technol.*, vol. 29, no. 2, pp. 617–633, Apr. 2024.

[11] S. Sai, S. Sharma, and V. Chamola, "Explainable AI-empowered neuromorphic computing framework for consumer healthcare," *IEEE Trans. Consum. Electron.*, vol. 71, no. 2, pp. 5889–5897, May 2025.

[12] M. Hiwale, R. Walambe, V. Potdar, and K. Kotecha, "A systematic review of privacy-preserving methods deployed with blockchain and federated learning for the telemedicine," *Healthcare Analytics*, vol. 3, Nov. 2023, Art. no. 100192.

[13] V. Mishra, K. Gupta, D. Saxena, and A. K. Singh, "A global medical data security and privacy preserving standards identification framework for electronic healthcare consumers," *IEEE Trans. Consum. Electron.*, vol. 70, no. 1, pp. 4379–4387, Feb. 2024.

[14] F. Arif et al., "Hybrid CNN-LSTM model for DDoS attack detection in Internet of Things-based healthcare industry 5.0," *IEEE Internet Things J.*, vol. 12, no. 22, pp. 46075–46082, Nov. 2025.

[15] F. Ullah, L. Mostarda, D. Cacciagrano, M. J. F. Alenazi, C.-M. Chen, and S. Kumari, "Federated edge intelligence for enhanced security in consumer intermittent healthcare devices using adversarial examples," *IEEE Trans. Consum. Electron.*, vol. 71, no. 2, pp. 4574–4585, May 2025.

[16] J. Lv, B.-G. Kim, B. D. Parameshachari, A. Slowik, and K. Li, "Large model-driven hyperscale healthcare data fusion analysis in complex multi-sensors," *Inf. Fusion*, vol. 115, Mar. 2025, Art. no. 102780.

[17] S. B. Babu and K. R. Jothi, "A secure framework for privacy-preserving analytics in healthcare records using zero-knowledge proofs and blockchain in multi-tenant cloud environments," *IEEE Access*, vol. 13, pp. 8439–8455, 2025.

[18] Z. Chen, W. Xu, B. Wang, and H. Yu, "A blockchain-based preserving and sharing system for medical data privacy," *Future Gener. Comput. Syst.*, vol. 124, pp. 338–350, Nov. 2021.

[19] M. Ali, F. Naeem, M. Tariq, and G. Kaddoum, "Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey," *IEEE J. Biomed. Health Informat.*, vol. 27, no. 2, pp. 778–789, Feb. 2023.

[20] D. Wu et al., "Adversarial attacks and defenses in physiological computing: A systematic review," *Nat. Sci. Open*, vol. 2, no. 1, Aug. 2022, Art. no. 20220023.

[21] J. Niu et al., "A survey on membership inference attacks and defenses in machine learning," *J. Inf. Intell.*, vol. 2, no. 5, pp. 404–454, Sep. 2024.

[22] P. Shojaei, E. Vlahu-Gjorgievska, and Y.-W. Chow, "Security and privacy of technologies in health information systems: A systematic literature review," *Computers*, vol. 13, no. 2, p. 41, Jan. 2024.

[23] X. Jiang et al., "Cybersecurity in neural interfaces: Survey and future trends," *Comput. Biol. Med.*, vol. 167, Dec. 2023, Art. no. 107604.

[24] A. Barua, M. A. Al Alamin, M. S. Hossain, and E. Hossain, "Security and privacy threats for Bluetooth low energy in IoT and wearable devices: A comprehensive survey," *IEEE Open J. Commun. Soc.*, vol. 3, pp. 251–281, 2022.

[25] A. Moradi, N. K. D. Venkategowda, S. P. Talebi, and S. Werner, "Privacy-preserving distributed Kalman filtering," *IEEE Trans. Signal Process.*, vol. 70, pp. 3074–3089, 2022.

[26] K. Si, P. Li, Z.-P. Yuan, X.-D. Jiang, and Z.-X. Wei, "Privacy-preserving distributed Kalman filtering based on state decomposition and dynamic mask," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 61, no. 3, pp. 5839–5852, Jun. 2025.

[27] J. Martinez, B. Passage, B. J. Mortazavi, and R. Jafari, "Hypothesis scoring for confidence-aware blood pressure estimation with particle filters," *IEEE J. Biomed. Health Informat.*, vol. 27, no. 9, pp. 4273–4284, Sep. 2023.

[28] X. Lu, Z. Liu, L. Xiao, and H. Dai, "Reinforcement learning-based personalized differentially private federated learning," *IEEE Trans. Inf. Forensics Security*, vol. 20, pp. 465–477, 2025.

[29] Y. Liu, U. R. Acharya, and J. H. Tan, "Preserving privacy in healthcare: A systematic review of deep learning approaches for synthetic data generation," *Comput. Methods Programs Biomed.*, vol. 260, Mar. 2025, Art. no. 108571.

[30] O. Salem, A. Serhrouchni, A. Mehaoua, and R. Boutaba, "Event detection in wireless body area networks using Kalman filter and power divergence," *IEEE Trans. Netw. Service Manage.*, vol. 15, no. 3, pp. 1018–1034, Sep. 2018.

[31] X. Yu and M. A. A. Al-Qaness, "Human activity recognition using deep residual convolutional network based on wearable sensors," *IEEE J. Biomed. Health Informat.*, vol. 29, no. 3, pp. 1950–1958, Mar. 2025.

[32] S. Ghosh, S. Kim, M. F. Ijaz, P. K. Singh, and M. Mahmud, "Classification of mental stress from wearable physiological sensors using image-encoding-based deep neural network," *Biosensors*, vol. 12, no. 12, p. 1153, Dec. 2022.

[33] A. Smahi et al., "A blockchainized privacy-preserving support vector machine classification on mobile crowd sensed data," *Pervas. Mobile Comput.*, vol. 66, Jul. 2020, Art. no. 101195.

[34] C. Jiang, C. Xu, C. Cao, and K. Chen, "GAIN: Decentralized privacy-preserving federated learning," *J. Inf. Secur. Appl.*, vol. 78, Nov. 2023, Art. no. 103615.

[35] R. Kousalya and G. A. S. Kumar, "RELAKA: Robust ECC based privacy preserving lightweight authenticated key agreement protocol for healthcare applications," *Eng. Sci. Technol., Int. J.*, vol. 59, Nov. 2024, Art. no. 101887.

[36] W. Wang, X. Li, X. Qiu, X. Zhang, V. Brusic, and J. Zhao, "A privacy preserving framework for federated learning in smart healthcare systems," *Inf. Process. Manage.*, vol. 60, no. 1, Jan. 2023, Art. no. 103167.

[37] T. McElroy, A. Roy, and G. Hore, "FLIP: A utility preserving privacy mechanism for time series," *J. Mach. Learn. Res.*, vol. 23, p. 111, Jan. 2022.

[38] G. S. Poh, P. Gope, and J. Ning, "PrivHome: Privacy-preserving authenticated communication in smart home environment," *IEEE Trans. Depend. Secure Comput.*, vol. 18, no. 3, pp. 1095–1107, May 2021.

[39] J. Zhang and C. Dong, "Secure and lightweight data aggregation scheme for anonymous multi-receivers in WBAN," *IEEE Trans. Netw. Sci. Eng.*, vol. 10, no. 1, pp. 81–91, Jan. 2023.

[40] C. DeSmet, C. Greeley, and D. J. Cook, "Hydra-TS: Enhancing human activity recognition with multiobjective synthetic time-series data generation," *IEEE Sensors J.*, vol. 25, no. 1, pp. 763–772, Jan. 2025.