ECG De-anonymization: Real-world Risks and a Privacy-by-design Mitigation Strategy

Hamza Aguelal
Computer Science & IT
University College Cork
Cork, Ireland
ORCID: 0009-0002-9409-4206

Paolo Palmieri
Computer Science & IT
University College Cork
Cork, Ireland
ORCID: 0000-0002-9819-4880

***Abstract*—The growing use of patient data in research underscores its value (for instance, in training AI). It also highlights the need for strong anonymization when health datasets are released publicly due to the risk of de-anonymization attacks. Electrocardiograms (ECG) are widely used, and real patient data have been openly released anonymously. However, ECGs are susceptible to linkage attacks, raising concerns around privacy, non-compliance with regulations such as the General Data Protection Regulation (GDPR), and loss of trust in digital healthcare. In this paper, we present a novel lightweight de-anonymization linkage attack on ECGs, and discuss benchmarking routes and an inclusive privacy protection framework that can be used in mitigating de-anonymization risks. The proposed matching attack leverages Convolutional Neural Networks (CNN)-based and ECG-specific features, and was tested on three open datasets: ECG-ID, MIMIC-IV and MIT-BIH. Unlike authentication-focused works, our study evaluates re-identification from an adversarial perspective, quantifying the risk on anonymized datasets based on metrics that establish a benchmarking baseline. Experimental results demonstrate an average matching accuracy of 97.22%, and nearly 100% for the best result, on the MIT-BIH dataset, for which previous results exist in the literature. Our results are substantially higher than the previous best-performing attack, which achieved an 81.9% accuracy. Consistent results on the two other datasets demonstrate the generality of our approach. The attack emphasizes evaluating de-anonymization risks before publicly releasing datasets. Based on our findings, we formalize recommendations into a new privacy-by-design framework resilient against real-world de-anonymization attacks, including inclusive processes to guide stakeholders in assessing requirements and offering insights into privacy metrics and improvement axes.**

***Index Terms*—De-Anonymization Attack, Electrocardiogram (ECG), Anonymity, Risk Assessment, Privacy-by-design.**

I. Introduction

The digitization of healthcare has led to the expansion of sensitive health information, particularly ECG (Electrocardiogram) recordings. ECG signals, characterized by their distinct biometric patterns, are used in diagnostics and biometrics and are essential for monitoring in telehealth systems and wearable devices. However, as healthcare systems become more interconnected to enhance patient care, the reliance on more advanced resources poses privacy challenges and comes with a risk against anonymization as a mandatory requirement to safeguard patient identities against attacks, primarily leveraging ML/DL (Machine and Deep Learning) algorithms.

Despite the adoption of de-identification, healthcare data, particularly ECGs, faces an escalating risk of de-anonymization attacks since traditional anonymization methods often fail to entirely obscure the unique biometric characteristics. This study exposes these vulnerabilities by simulating the adversarial linkage attack aimed at matching individuals by creating fingerprints based on advanced features combined with Convolutional Neural Networks (CNN) embeddings exploiting the sensitive nature of the temporal and morphological features inherent in ECG data. We highlight that such data breaches have ethical, legal, and personal consequences, and regulatory frameworks like the General Data Protection Regulation (GDPR) [10] emphasize anonymization and pseudo-anonymization; however, they lack enforceable methodologies for the robustness of these techniques. This study provides insights through a structured assessment showcasing linkage threats and attack measures, and we aim to influence advancements in patients' data privacy and security in the age of AI-driven healthcare systems.

Research in ECG biometrics has predominantly emphasized authentication [29] and identification [24]. While some studies explore de-anonymization, they are often constrained by narrow contexts or theoretical assumptions. This study extends the focus on the broader implications of data linkage attacks and approaches the challenges of generalizability and dataset diversity, as we opted for different formats and different timeframes. We suggest a transparent, reproducible, and standardized framework encompassing privacy by design protection and risk evaluation. Our contributions include:

- Linkage Attacks: Development of a matching attack that combines CNN with ECG-specific feature extraction.
- Diverse Dataset Application: Evaluation across datasets, covering various configurations and records timeframes.
- Bias-Free Assessment: Strict separation of datasets with

Funded in part by the European Union (EU), Grant Agreement no. 101095717 (SECURED) and by Taighde Éireann – Research Ireland under Grant no. 18/CRT/6222.

non-overlapping records nor fragments, eliminating bias and reflecting realistic adversarial conditions.
- Security Enhancement Framework: We propose an inclusive framework to minimize risks related to de-anonymization for a continuous risk-based strategy for privacy and health data protection.
- Benchmarking and Real-World Simulation: Establishment of benchmarks for de-anonymization assessment, incorporating real-world healthcare data and GDPR-aligned recommendations.

The remainder of the paper is organized as follows: Sec II reviews de-anonymization and ECG background, Sec III describes datasets, feature extraction, and the linkage attack threat model. Results and analysis with interpretation in Sec IV and Sec V outline our mitigation proposals and framework and end with the conclusion in Sec VI.

II. Background & Related works

This section reviews prior work on de-anonymization attacks and risk assessment in healthcare data and ECG. We draw insights from surveys, experimental studies, and emerging literature to highlight the re-identification of ECG data.

A. De-anonymization and Evolution of Attacks

De-anonymization is breaking or reversing the anonymization process to extract or link back data to individuals. Such techniques were explored in early works [34], mainly relying on auxiliary data to demonstrate anonymization vulnerabilities. Researchers in [2] survey de-anonymization techniques on multiple data types and provide a classification, while findings in [17] reveal that successful attacks in various domains such as social networks [4] occurred after 2009, showing a marked increase in attacks in recent years.

Specifically to health data, early surveys such as [9] show that in 2011, approximately 34% of records were successfully re-identified on average (with the highest number of attacks focusing on genomic data). [1] shows that in 2025, re-identification risks from known attacks are substantially higher, ranging from 12.4% for Medical Records to 96% on genomics, and highlights how attacks have been developed for most health data types.

B. ECG Relevance and Vulnerability to de-anonymization

ECG signals are fundamental in clinical cardiology (e.g. diagnosis of arrhythmias, cardiac electrophysiology, myocardial infarction and other conditions). The unique and stable morphological features, including the P-wave, QRS complex, R-R intervals and consistency over time, make them very useful in several use cases, including authentication [11]. However, this makes ECGs vulnerable to de-anonymization attacks and adversarial learning to exploit features. For instance, [13] presents models that, given a patient's ECG, are able to identify the sex of the patient (with 95.13% accuracy for males and 96.59% for females). The same work can also detect certain conditions with high accuracy. [35] also presents a model able to recognize a patient's sex (although with lower accuracy than [13]) and also their age group (75.5% accuracy). Furthermore, they can de-anonymize a patient through a linkage attack with 81.9% accuracy. Our work focuses on this latter attack, achieving a substantially higher accuracy rate.

C. De-anonymization vs Authentication

In a separate line of work, ECGs have been used for biometric identification, where specific biometric markers or data are used to recognize a user and allow them to access a specific resource. While traditional biometrics such as fingerprint or retina are commonly used, recent research has expanded into further data types, including ECGs [7]. While some of the features used in identifying individuals for authentication may be relevant to de-anonymization, results in this area cannot be directly compared with de-anonymization. In particular, success rates cannot be compared because of the intra-user/patient validation, where subjects are drawn from the same recording or session [32], without a strict separation between training and test data. This means that ECG fragments are overlapping [20] (that is, used in both the training and the testing phases of the model generation), leading to potential bias and allowing models to perform exceedingly well by memorizing patterns rather than generalizing to unseen subjects. On the other hand, in the de-anonymization context, it is crucial to maintain the training and testing data entirely separate. We nonetheless mention below, for completeness, the leading results in ECG biometrics for authentication.

Mitchell et al. [24] reviewed accuracies for ECG-based authentication using fudicial and non-fudicial methods reporting 99.3% as highest results, and [7], [36], [3] reaching 81.33%, 91.1% and 96.4%, respectively as max accuracies on MIT-BIH. On the other hand, [27] and [23] demonstrate 96% and 97.75% success rates on ECG-ID, while [5] used ML for gender identification and achieved 99.2% identification but comparing only two persons to the whole same dataset. However, all prior works aimed either for different tasks (authentication/identification) focused on specific conditions or used intra-subject overlap as in [32] to reach nearly 100% on ECG-ID authentication.

D. Deep Learning for ECG Analysis

DL has revolutionized cardiology by enabling the extraction of intricate cardiac features from ECGs [21], as in [22] using CNN or Long Short-Term Memory (LSTM) for [20], techniques that have improved ECG-based models facilitating clinical, diagnostic assessments and monitoring tasks or cardiac abnormalities detections [16]. However, these models are exploited to be used in our de-anonymization attacks. Therefore, standardized benchmarking frameworks are essential to diagnostic accuracy with patient privacy in clinical and health applications.

E. Privacy Risk and Regulation

Health data are inherently sensitive, encompassing monitoring signals like ECGs that serve as digital biomarkers [11] important for decision-making and treatment, yet their

high dimensionality increases the privacy breach risk, and GDPR mandates rigorous data protection practices under its scope of application, or apply strong data anonymization. It defines anonymized data as information that cannot be linked to an identified or identifiable individual in question by any means reasonably likely to be used [10]. However, traditional anonymization can fail when faced with linkage attacks that exploit quasi-identifiers or inferred information, resulting in non-compliance with GDPR. The empirical evidence of these measures is insufficient without integrating advanced Privacy-Enhancing Technologies (PETs) [19] with a focus on risk-based approaches that balance data utility and clinical value with patient confidentiality in digital health ecosystems.

III. Methods

This chapter outlines our approach to assessing de-anonymization risk in ECG data by formulating an actual de-anonymization attack. Our threat model integrates a realistic setup for the attack, the clinical adversarial scenario illustrated in Fig. 1. A patient (Bob) has ECG records collected in a controlled clinical setting (for example, in a hospital for medical diagnosis). These records are included (with Bob's consent) in a dataset alongside those of other patients affected by the same medical condition as Bob. The hospital publicly releases the dataset in an anonymized form to support medical researchers working on the same condition. At a later stage, an attacker gains access to an ECG facility and is able to see ECG data being recorded as well as the patient name (but no information on medical conditions). Bob goes to the hospital where the ECG machine has been compromised, and undergoes an ECG as part of a regular check-up. The attacker sees Bob's record and performs the de-anonymization attack described in the following to determine whether Bob was also a patient in the anonymized dataset, using only the new ECG record for Bob that was leaked and the publicly available dataset. The attacker is able to identify Bob as one of the patients in the dataset and, therefore, link him to the particular condition the dataset relates to. This highlights the feasibility of the de-anonymization attack and the potential sensitive information leakage in healthcare.

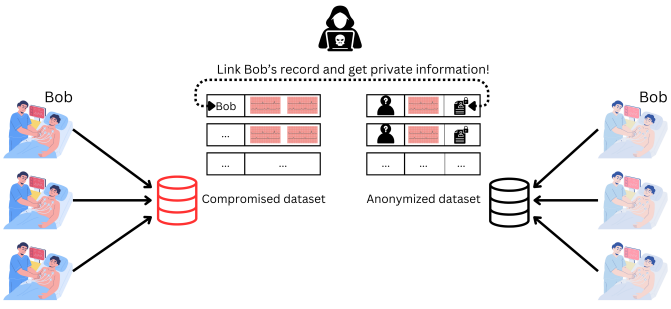

Fig. 1: De-anonymization Adversarial Clinical Scenario

Our methodology is structured in stages that reflect both the technical and clinical imperatives using diverse datasets to evaluate the de-anonymization techniques. The algorithm used to perform the attack is included below as pseudo-code. The Python code is publicly available at https://github.com/ague08/Paper-ECG_De-anonymization_Real-world_Risks_and_Privacy-by-design_Mitigation_Strategy-.git

A. Data and preprocessing

We have targeted different structures' datasets to assess the de-anonymization in different setups: i) MIT-BIH [25] 30-minute long records, captured via dual-lead (a lead represents a specific view of the heart's electrical activity in ECG) setups for 47 subjects. ii) ECG-ID DB [27], consisting of 20-second long records for each of the 90 individuals with multiple recordings per subject (up to 20 records). iii) MIMIC-IV [15] a 10-second recording utilizing 12 leads for 92 patients, also with multiple subject records. Records were captured at various times (clinical visits or episodes of care), on different days, months (6 months for ECG-ID), and even years in the case of MIMIC-IV, reflecting realistic adversarial conditions.

For our setup design, we highlight the following aspects:

- **Limited Auxiliary Information:** Use features and embeddings only. This limits the attacker's power.
- **Time Variability:** Temporal differences in recordings make linkage attacks challenging to ensure that our results reflect de-anonymization risks.
- **Practical Threat Models:** The separation of records emulates real-world healthcare data-sharing scenarios, for instance, between different hospital systems.

To further analyze the ECG signals and extract the unique components, we trigger the noise reduction process and eliminate baseline wander, powerline interference, and motion artefacts to retain the relevant waveforms and complexes (e.g. PQRS complex) by applying a bandpass Butterworth filter [24] using *butter* and *filtfilt* functions to filter in a forward and backward direction to eliminate phase distortion.

For illustration, Tab. I shows original signals with noise alongside filtered signals annotated with clearer R-peak.

TABLE I: ECG Signal: *Noisy* vs *Filtered*

Dataset	Original Signal	Filtered Signal
MIT-BIH	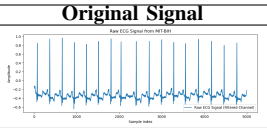	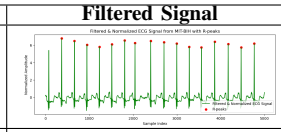
ECG-ID	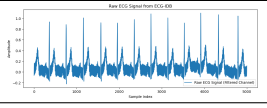	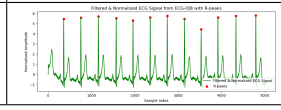

B. Threat model

As described above, our adversarial model simulates a realistic de-anonymization attack. We assume that the attacker has access to a small set of the identified dataset and a larger anonymized one, with limited access to the ECG signals only, with no auxiliary data to meet stricter adversarial conditions. The conception of the setup is highlighted in Fig.2. The scenario is that the adversary de-anonymizes ECG data by linking anonymized records to identified individuals. We aim

for an attacker to have access to a small amount of leaked data compared to the anonymized set, which is much larger to reflect a realistic clinical practice. In practice, this means we split the original datasets (ECG-ID, MIMIC-IV-ECG Demo [15] and MIT-BIH, see Section III-A) between 25% leaked, 75% anonymized data; records are sparse, including independent records; we address cases with fewer than four records by minimizing attacker advantage, and we consider only one record as identified.

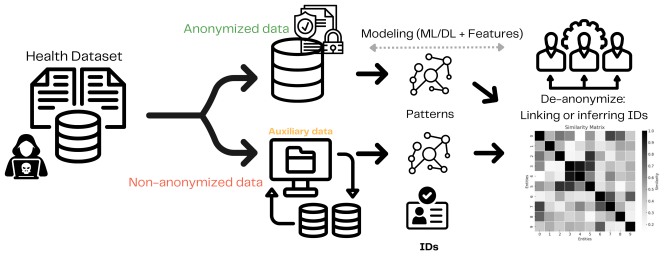

Fig. 2: De-anonymization Matching Attack Process Overview

To create a comprehensive fingerprint for each individual, we combine the traditional features with the CNN embeddings:

- **CNN Embeddings:** Capture temporal dependencies.
- **Advanced Features:** Extract unique features.
- **Normalized and Combined Features:** Normalize and combine embeddings and the features in a fingerprint.

The output representation is a fingerprint stored in a dictionary with person IDs as keys. We outline our attack in Algorithm 1, and the re-identification evaluation as follows:

- **Ranked De-Anonymization:**
  1) Function to compare the fingerprints from the identified dataset against the anonymized subjects.
  2) Uses a similarity-based metric, comparing vectors with Euclidean distance for best results.
  3) Returns a ranked list of potential matches; the top-ranked is the best match.

In addition to the experimental setup, we define a set of evaluation metrics that balance diagnostic accuracy with privacy risk and computational complexity and scalability by analyzing the resource requirements to reflect the feasibility of deploying such attacks in real-world clinical environments.

C. Feature Extraction

De-anonymization relies on extracting invariant features from the ECGs, and our approach utilizes both traditional feature extraction methods and DL-based embeddings to create a comprehensive fingerprint for each individual. The Algorithm 2 details the extraction process after data preprocessing and normalization for better robustness against noise [18]. We focused on the following features to have an extensive representation to reflect the intrinsic cardiac rhythm:

- **Time-domain features:** Extracted from the waveform morphology as R-peaks, computation of intervals (RR-interval) and deviation to characterize rhythm variability (PR, QT) (See Samples in Fig. 3).

Algorithm 1: ECG De-Anonymization Attack

Input : ECG dataset $D = D_{\text{Anon}} \cup D_{\text{De-anon}}$, CNN $\mathcal{M}$, distance $\text{dist}(\cdot,\cdot)$, feature weights $\omega = (\omega_e, \omega_f, \omega_c)$, top-$k$ matches
Output: Mappings from $D_{\text{De-anon}}$ to D_{Anon}, accuracy metrics
Step 1: Extract ECG Fingerprints;
foreach $r_i \in D$ **do**
  Extract fingerprint $\mathbf{F}_i = [\mathbf{f}_t, \mathbf{f}_f, \mathbf{f}_{morph}, \mathbf{f}_{ent}, \mathbf{e}]$ (see Alg. 2);
Step 2: Compute Similarity Matrix;
Compute similarity d_{ij} between each $r_j \in D_{\text{De-anon}}$ and $r_i \in D_{\text{Anon}}$:

$$d_{ij} = \omega_e \text{dist}(\mathbf{e}_j, \mathbf{e}_i) + \omega_f \text{dist}(\mathbf{f}_{morph,j}, \mathbf{f}_{morph,i}) + \omega_c \text{dist}(\mathbf{F}_j, \mathbf{F}_i)$$

Step 3: Identify Top-k Candidates;
foreach $r_j \in D_{De\text{-}anon}$ **do**
  select top-k candidates: $\text{C}_j = \{r_{i1}, r_{i2}, \ldots, r_{ik}\}$ sorted by d_{ij}
Step 4: Predict Identities;
Predict identity by minimum distance match:

$$\hat{r}_j = \arg\min_{r_i \in C_j} d_{ij}$$

Step 5: Evaluate Metrics;
Compute accuracy metrics:

$$\text{Top-1 Acc.} = \frac{|\{\hat{r}_j = r_j^{true}\}|}{|D_{\text{De-anon}}|}, \quad \text{Top-}k\text{ Acc.} = \frac{|\{r_j^{true} \in C_j\}|}{|D_{\text{De-anon}}|}$$

return *Predicted mappings and metrics*

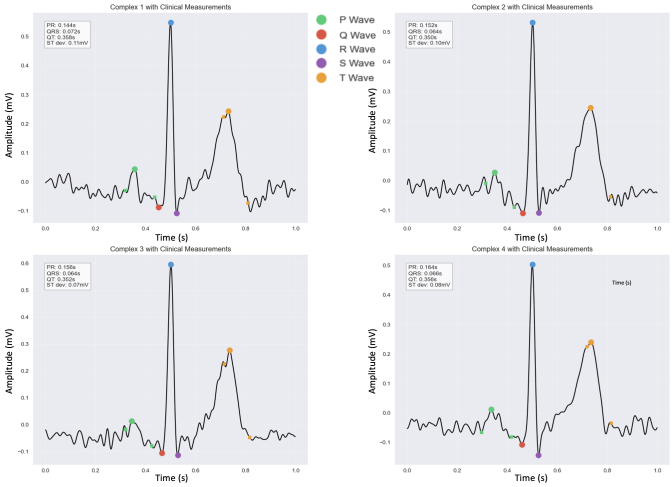

Fig. 3: PQRST Complex Analysis

- **Morphological Features:** Derived from PQRST complex metrics (e.g., amplitudes, PR, QT intervals), capturing the morphological patient differences.
- **Frequency-domain features:** Capture the frequency components of the ECG signal using techniques like Welch's power spectral density (PSD).
- **Non-linear features:** Capture the non-linear dynamics, such as wavelet entropy and fractal dimensions.

In addition to these features, we combine DL-based embeddings. We use a simple CNN (four 1D convolutional layers, each followed by max-pooling for feature extraction, with increasing filter sizes: 128, 256, 512) as an effective model for capturing spatial hierarchies in ECG waveforms.

The process could be summarized in the following steps:

1) **Preprocessing:** The ECG signals are preprocessed to remove noise and artefacts.
2) **Segmentation:** The signals are segmented into individual heartbeats.

3) **Embedding Generation:** High-dimensional vector to capture user-specific patterns and the fusion enhances the discriminative possibility in ECG morphology.

Algorithm 2: Extract Features for De-Anonymization

Input: ECG dataset D, CNN model $\mathcal{M}$, sampling freq. f_s, cutoff f_c
Output: Fingerprints $\{\mathbf{F}_i\}$ for $r_i \in D$
foreach $r_i \in D$ **do**
  Step 1.1: Preprocessing;;
  Filter: $r_i^{filt} = \text{Butterworth}(r_i, 0.5\text{-}45\ Hz, f_s)$;
  Detect peaks: $R = \text{DetectRPeaks}(r_i^{filt})$;
  Step 1.2: Time Features;
  Compute RR intervals/stats: $\mathbf{f}_t = [\mu_{RR}, \sigma_{RR}]$,
  $\mu_{RR} = \frac{1}{n}\sum_k RR_k$, $\sigma_{RR} = \sqrt{\frac{1}{n}\sum_k (RR_k - \mu_{RR})^2}$;
  Step 1.3: Frequency Features;
  Compute PSD and dominant frequency: $\mathbf{f}_f = [f_{dom}]$,
  $f_{dom} = \arg\max_f \text{Welch}(r_i^{filt}, f_s)$;
  Step 1.4: Morphological;
  Extract morphology around R-peaks:
  $\mathbf{f}_{morph} = [P_{amp}, Q_{amp}, R_{amp}, S_{amp}, T_{amp}, PR_{int}, QT_{int}]$;
  Step 1.5: Entropy;
  Compute entropy measures: $\mathbf{f}_{ent} = [H_{sh}(r_i^{filt}), H_{wav}(r_i^{filt})]$;
  Step 1.6: CNN Embeddings;
  Compute average CNN embedding: **foreach** $R_k \in R$ **do**
    $\mathbf{e}_k = \mathcal{M}(r_i^{filt}[R_k - h : R_k + h])$;
  $\mathbf{e} = \frac{1}{|R|}\sum_k \mathbf{e}_k$;
  Step 1.7: Fingerprint;
  Combine features into fingerprint: $\mathbf{F}_i = [\mathbf{f}_t, \mathbf{f}_f, \mathbf{f}_{morph}, \mathbf{f}_{ent}, \mathbf{e}]$;
return *Fingerprints* $\{\mathbf{F}_i\}$

IV. Results and Discussion

A. Results

In this section, we present the attack results on the different non-overlapping splits for each dataset to assess the stability of our architecture under varied temporal conditions.

TABLE II: Attack results (Top-1, Top-3) on three splits of each dataset.

Match	Splits	ECG-ID	MIT-BIH	MIMIC-IV
Top-1	Split 1	90.0%	93.75%	75.32%
	Split 2	92.1%	100%	74.03%
	Split 3	90.0%	97.92%	72.73%
	Mean ± SD	90.7% ± 1.0	97.22% ± 3.4	74.02% ± 1.3
Top-3	Split 1	95.6%	97.92%	85.71%
	Split 2	97.5%	100%	84.42%
	Split 3	95.6%	100%	83.12%
	Mean ± SD	96.2% ± 1.0	99.3% ± 1.2	83.99% ± 0.75

As shown in Table II we achieved a high matching accuracy of 92.1% on ECG-ID and almost similar across splits, with a low ± standard deviation (SD). Similarly, on MIT-BIH, the attack model achieved a high accuracy of 100%. For MIMIC-IV and the purpose of this experiment, we used a subset (demo) of the full MIMIC-IV-ECG dataset and obtained a high accuracy of 75.32% despite the notable changes in the recording sessions revealing the inherent risks related to the sensitive ECG rich in clinical discriminative features such as the PQRST complexes [30] and our cross-validation approach confirms these vulnerabilities. The results show that existing techniques may be insufficient against the linkage attack; thus, there is a need for stronger privacy-preserving techniques [8].

B. Discussion and Key Findings

Improved Matching Accuracy: As highlighted in Tab III, our matching process, to the best of our knowledge, outperforms other de-anonymization works [35] that achieved a high accuracy of 81.9%; also our approach represents more generalizability across diverse clinical scenarios that re-identify specific patterns only or patients' conditions [13] or controlled conditions with the customized dataset, as we align with medical requirements where patients are subject to variability over time. Also, our work differs from the authentication approaches that can tend to have higher accuracies due to the idealized conditions suitable for identity management (already explained in II-C) or the use of complex set-up requiring higher computation. In our prototype, we showcase the importance of the ECG-related morphological features (QRS duration, R-Peak amplitude, PR-interval, QT-interval and P-wave) and how they contribute positively to the matching process. Experimental results indicate that, on average, QRS and RR have a higher impact on the de-anonymization success rate than P-wave for example, in line with previous results [6], [11]. However, determining analytically an exact impact score or rate for each feature remains an open question, likely to require further research on a greater number of datasets.

Accuracy variation: Variation in the accuracies across datasets is explained mainly by the differences in recording conditions (sessions). We achieved the highest accuracy for MIT-BIH as records originate from longer and continuous sessions per individual. Contrariwise, the ECG-ID dataset have slightly lower accuracy since it comprises multiple 20-second records collected over various intervals up to six months. The lower accuracy for the MIMIC-IV-ECG subset is due to higher diversity, shorter records (10 seconds), and more significant temporal gaps (ECGs collected between 2008 - 2019 included as part of MIMIC-IV-ECG [14])

Simplicity of the attack process: Our lightweight, straightforward approach, combining Traditional ML for ECG-Based features (time-domain with the morphological markers) and lightweight CNN, provides a computationally efficient simple model aligning with real-world set-up with no need for sophisticated resources and accessible to adversaries with basic knowledge of ML and signal processing. As an instance, Fig. 4 represents the low computational overhead during training and linking stages across ECG-ID and MIMIC-IV datasets.

Top-3 matching accuracy: Besides perfect matches (Top-1), Top-3 accuracy, in our case, considers the true identity is within the three best-predicted matches (small subset). The attacker can use these close matches to significantly narrow anonymity and as a first step to a multi-layer attack through secondary inferences and the need for auxiliary data, facilitating de-anonymization and increasing privacy risks.

Feature Significance: The biometric properties of ECG signals, such as R-peak amplitudes, QRS durations, and temporal patterns, enabled high accuracy in re-identification attacks.

Variability in Demographics: Datasets with diverse demographic and clinical conditions have a slightly better resistance

TABLE III: De-anonymization Performances on ECG Datasets

Work	Dataset	Individuals	Overlap	Result	Model & Complexity
[35]	Combined datasets (include MIT-BIH)	223	Yes	81.9% accuracy	Transparent ML models, SHAP analysis, moderate complexity.
Our work	ECG-ID DB	90	No	92.1% accuracy	Traditional ML, simple CNN, CPU-based, lightweight.
Our work	MIT-BIH DB	47	No	100% accuracy	Traditional ML, simple CNN, CPU-based, lightweight.
Our work	MIMIC-IV Demo	92	No	75.32% accuracy	Traditional ML, simple CNN, CPU-based, lightweight.

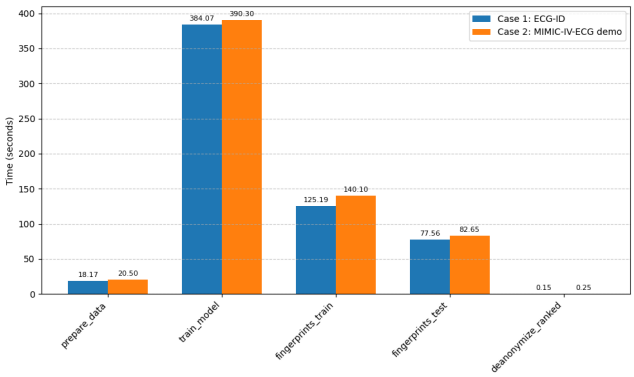

Fig. 4: Execution Times for Pipeline Stages

to attacks but are still insufficient to mitigate the risks. Consequently, we advocate for a risk-based approach to combine clinical diversity, privacy and regulatory requirements.

C. Implications

De-anonymization and specifically matching attacks have far-reaching critical implications, as our experiments highlight vulnerabilities and threats to patients in healthcare systems:

- **Patient Confidentiality:** Re-linking anonymized ECG records to individual patients compromises the confidentiality of Protected Health Information (PHI), potentially disclosing sensitive details and compromising clinical decision-making and trust in institutions.
- **Regulatory Compliance:** Under GDPR, healthcare providers must comply with the framework and its requirements or implement sufficient de-identification. We suggest that current practices may fall short, necessitating elaboration and advanced PETs such as Differential Privacy (DP) [28], synthetic data generation, and federated learning (FL).
- **Data Sharing and Research:** Privacy breaches undermine collaborative research and data-sharing initiatives, slowing the pace of innovation in digital health, remote patient monitoring and personalized medicine.

We summarize our assessment's implications In Tab. IV:

V. Recommendations

Building on our results, we propose a comprehensive framework to provide an inclusive privacy by design and risk-based approach represented in Fig. 5, aligned with GDPR guidelines to mitigate and minimize attacks' threats and to enhance de-anonymization risk evaluation to ensure healthcare protection.

TABLE IV: Privacy, Anonymity, and Trust Implications

Aspect	Implication	Details
Stable biometric patterns	High attack risk	Linking anonymized individuals, threatening confidentiality.
Anonymity Challenges	Ineffective de-identification	Advanced attacks require integration of advanced techniques.
Trust in Healthcare	Weak confidence in systems	Less data-sharing leads to slower advancements and research.
Regulatory Compliance	Non-compliance and legal	Lawfulness, fairness, and ethics imposed by legislations.
Economical	Fines, data and public losses	Financial, reputational and substantial costs (remediation).

It emphasizes proactive integration of security and privacy measures:

- **Advanced Anonymization Techniques:** Integration of PETs to balance the trade-off between data utility and patient privacy. Moreover, re-assessment to refine anonymity and improve trust [33].
- **Risk-based Management:** Adopt a categorized risk-based approach (e.g. high, medium, low) considering data sensitivity (e.g., ECG, genomic) and threats potential: AI-assisted risk detection and monitoring to proactively identify privacy risks and suspicious activities.
- **Continuous Privacy and Security Enhancements:** Continuous improvement cycle that incorporates stakeholder engagement and awareness, improving research and transparency (including negative results reporting), and regular assessment of privacy-preserving mechanisms.
- **Standardized Benchmarks:** Standardized metrics and evaluation protocols that reflect the diverse clinical characteristics and demographics of health data.
- **Legal and Ethical Refinements:** Update regulatory frameworks to explicitly address de-anonymization threats and technical detail and PET related to anonymization thresholds and criteria for acceptable re-identification risk, and implement enhanced Data Protection Impact Assessments (DPIAs) for high-risk scenarios in the face of evolving de-anonymization attacks.

To further enhance the de-anonymization assessment, we detailed an overview of the PETs following the risk-based and use-related cases incorporating privacy-by-design approaches:

1) **In-House use cases for data processing, statistical analysis and ML:** DP preserves patterns essential for training while enhancing anonymization.
   - **Differential Privacy:** Provides strong privacy guarantees by adding noise using Epsilon (ϵ) to retain

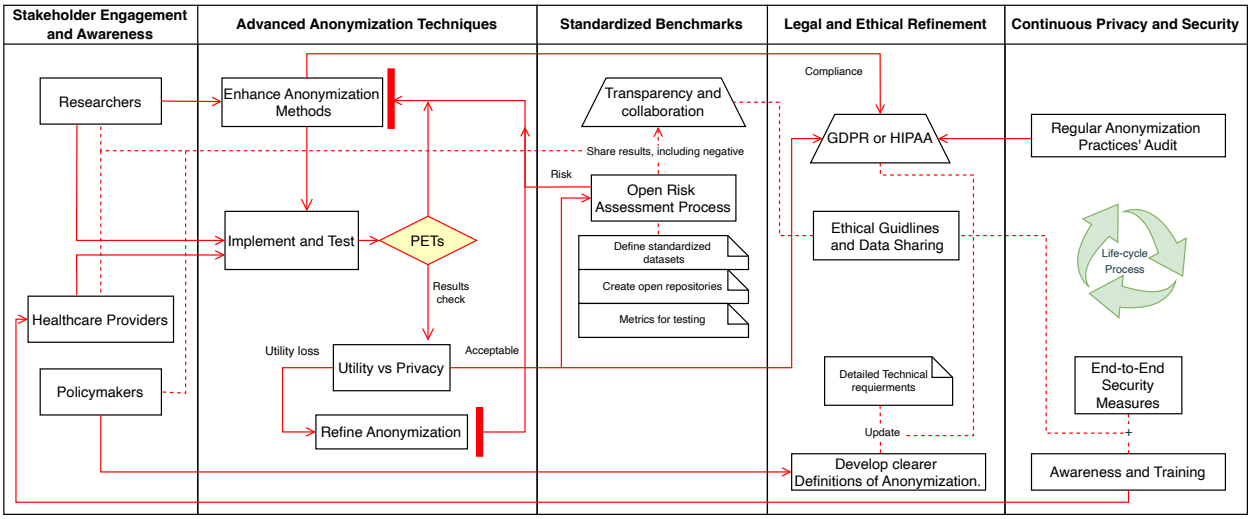

Fig. 5: Framework for Privacy Protection Against De-Anonymization Attacks.

essential utility and privacy trade-offs for specific applications. Implementation considerations:

- Prioritizing DP within in-house data processing environments (research or training models) rather than for open data release, as excessive one will obscure the utility.
- Prioritize highly sensitive attributes and stronger privacy, even at the expense of utility loss.

2) **Data sharing and collaborative research cases:** Data to be shared across institutions and ensuring quality is as important as privacy. DP could fail to maintain high quality in clinical trials or collab research, so alternative PETs should be considered.
   - **Synthetic Data Generation:** Method to generate artificial datasets [26] that mimic the original data while eliminating identifying patterns. Implementation considerations:
     - Security of generators against leaks and vulnerabilities to Membership Inference Attacks.
     - Preserve clinical correlations, such as the temporal dynamics in ECG recordings, to support reliable model training and diagnostics.
   - **Federated Learning:** Allow analysis without transferring raw data by enabling decentralized models' training on datasets. Implementation considerations:
     - Use FL frameworks to retain patient data locally, sharing only encrypted model updates.
     - Secure aggregation protocols, thus meeting regulatory standards.
   - **Advanced Multi-Party Computation (SMPC):** SMPC or Homomorphic Encryption allow collaborative analysis of encrypted sensitive data without exposing raw private information [31].
   - **Hybrid Privacy Approaches:** Combine multiple PETs to offer enhanced protection. For example:
     - Incorporate DP in the model preprocessing step to eliminate persistent patterns so the synthetic outputs do not inherit identifying features.
     - Deploy client-level DP [12] to local model updates before sending them to the central server.

3) **Real-Time Clinical Applications and Edge Analytics:** Real-time processing, as in wearable device monitoring or telemedicine streams, requires privacy by design and minimal latency and computational PETs at the capture before transmission.
   - **Lightweight DP:** Lightweight local DP can be applied directly on edge devices; the calibrated noise can preserve essential patterns (such as ECG signal morphology) while limiting re-identification risks.
   - **Trusted Execution Environments (TEEs):** Isolate real-time processing on edge devices within a dedicated enclave. It can be combined with DP applied to the data outputs before leaving the device.

VI. Conclusion

Health data is highly sensitive and, as such, is considered a prime target for attacks, particularly de-anonymization ones. This work highlights the vulnerability of ECG to de-anonymization by presenting a novel attack with accuracies of 92.1%, 100% and 75.32% on the publicly available and well-studied ECG-ID, MIT-BIH and MIMIC-IV ECG datasets. Based on these results, and given the high risk of de-anonymization identified, we formulate a general framework for the protection of private data. The framework reflects an inclusive risk-based approach, aligned with GDPR requirements. Our recommendations, stemming from the de-anonymization attack, include continuous security monitoring and evaluation of de-anonymization risks (which may vary over time) and advocate for the publication of negative results by researchers working on anonymity, as they can help us understand which

anonymization approaches offer better protections. Differential privacy and advanced privacy-preserving methods like synthetic data generation must be assessed critically using a risk-based perspective to contribute positively to health research and improve trust in systems. In addition, health datasets should be made accessible to anonymity researchers to provide detailed insight into specific datatypes. The current lack of specificity, in fact, allows weak anonymization methods to persist, leaving datasets vulnerable. In the longer term, our results point to a need to include more detailed and actionable requirements in terms of anonymity in regulatory frameworks, to align them with the current advancements in ML/DL and AI-driven innovations. Our work bridges the gap between theoretical risks and practical attacks, emphasizing the need for patients' data protection mechanisms and adopting continuous improvement, we help assess privacy solutions while allowing innovation in healthcare and enhancing privacy-by-design concepts.

REFERENCES

[1] Aguelal, H., Palmieri, P.: De-anonymization of health data: A survey of practical attacks, vulnerabilities and challenges. In: Proceedings of the 11th International Conference on Information Systems Security and Privacy: ICISSP. vol. 2, pp. 595–606. INSTICC (2025). https://doi.org/10.5220/0013274200003899

[2] Al-Azizy, D., Millard, D., Symeonidis, I., O'Hara, K., Shadbolt, N.: A literature survey and classifications on data deanonymisation. In: Risks and Security of Internet and Systems. Springer International (2016)

[3] Belo, D., Bento, N., Silva, H., Fred, A., Gamboa, H.: Ecg biometrics using deep learning and relative score threshold classification. Sensors **20**(15), 4078 (2020)

[4] Bhattacharya, M., Roy, S., Chattopadhyay, S., Das, A.K., Shetty, S.: A comprehensive survey on online social networks security and privacy issues: Threats, machine learning-based solutions, and open challenges. Security and Privacy **6**(1), e275 (2023)

[5] Cabra, J.L., Mendez, D., Trujillo, L.C.: Wide machine learning algorithms evaluation applied to ecg authentication and gender recognition. In: Proceedings of the 2018 2nd International Conference on Biometric Engineering and Applications. pp. 58–64 (2018)

[6] de Chazal, P., O'Dwyer, M., Reilly, R.: Automatic classification of heartbeats using ecg morphology and heartbeat interval features. IEEE Transactions on Biomedical Engineering **51**(7), 1196–1206 (2004). https://doi.org/10.1109/TBME.2004.827359

[7] Deshmane, M., Madhe, S.: Ecg based biometric human identification using convolutional neural network in smart health applications. In: 2018 Fourth International Conference on Computing Communication Control and Automation (ICCUBEA). pp. 1–6. IEEE (2018)

[8] Dwork, C.: Differential privacy: A survey of results. In: Theory and Applications of Models of Computation. pp. 1–19. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)

[9] Emam, K.E., Jonker, E., Arbuckle, L., Malin, B.: A systematic review of re-identification attacks on health data. PLoS ONE **6**(12), e28071 (2011). https://doi.org/10.1371/journal.pone.0028071

[10] European Union: Regulation (eu) 2016/679 general data protection regulation. Official Journal of the European Union (2016), eur-lex.europa.eu/eli/reg/2016/679/oj

[11] Fratini, A., Sansone, M., Bifulco, P., Cesarelli, M.: Individual identification via electrocardiogram analysis. Biomedical engineering online **14**, 1–23 (2015)

[12] Geyer, R.C., Klein, T., Nabi, M.: Differentially private federated learning: A client level perspective. arXiv preprint arXiv:1712.07557 (2017)

[13] Ghazarian, A., Zheng, J., Struppa, D., Rakovski, C.: Assessing the reidentification risks posed by deep learning algorithms applied to ecg data. IEEE Access **10**, 68711–68723 (2022)

[14] Goldberger, A.L., Amaral, L.A.N., Glass, L., Hausdorff, J.M., Ivanov, P.C., Mark, R.G., Mietus, J.E., Moody, G.B., Peng, C.K., Stanley, H.E.: PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. Circulation **101**(23), e215–e220 (2000), https://physionet.org

[15] Gow, B., Pollard, T., Nathanson, L.A., Moody, B., Johnson, A., Moukheiber, D., Greenbaum, N., Berkowitz, S., Eslami, P., et al.: Mimic-iv-ecg demo-diagnostic electrocardiogram matched subset demo

[16] Hannun, A.Y., Rajpurkar, P., Haghpanahi, M., et al.: Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nature Medicine **25**, 65–69 (2019). https://doi.org/10.1038/s41591-018-0268-3

[17] Henriksen-Bulmer, J., Jeary, S.: Re-identification attacks—a systematic literature review. International Journal of Information Management **36**, 1184–1192 (Dec 2016). https://doi.org/10.1016/j.ijinfomgt.2016.08.002

[18] Ingale, M., Cordeiro, R., Thentu, S., Park, Y., Karimian, N.: Ecg biometric authentication: A comparative analysis. IEEE Access **8**, 117853–117866 (2020). https://doi.org/10.1109/ACCESS.2020.3004464

[19] Jordan, S., Fontaine, C., Hendricks-Sturrup, R.: Selecting privacy-enhancing technologies for managing health data use. Frontiers in Public Health **10**, 814163 (2022)

[20] Kim, B.H., Pyun, J.Y.: Ecg identification for personal authentication using lstm-based deep recurrent neural networks. Sensors **20**(11) (2020)

[21] Kiranyaz, S., Ince, T., Gabbouj, M.: Real-time patient-specific ecg classification by 1-d convolutional neural networks. IEEE transactions on biomedical engineering **63**(3), 664–675 (2015)

[22] Labati, R.D., Muñoz, E., Piuri, V., Sassi, R., Scotti, F.: Deep-ecg: Convolutional neural networks for ecg biometric recognition. Pattern Recognition Letters **126**, 78–85 (2019)

[23] Melzi, P., Tolosana, R., Vera-Rodriguez, R.: Ecg biometric recognition: Review, system proposal, and benchmark evaluation. IEEE Access **11**, 15555–15566 (2023). https://doi.org/10.1109/ACCESS.2023.3244651

[24] Mitchell, A.R., Ahlert, D., Brown, C., Birge, M., Gibbs, A.: Electrocardiogram-based biometrics for user identification–using your heartbeat as a digital key. Journal of Electrocardiology **80**, 1–6 (2023)

[25] Moody, G., Mark, R.: The impact of the mit-bih arrhythmia database. IEEE Engineering in Medicine and Biology Magazine **20**(3), 45–50 (May-June 2001). https://doi.org/10.1109/51.932724

[26] Murtaza, H., Ahmed, M., Khan, N.F., Murtaza, G., Zafar, S., Bano, A.: Synthetic data generation: State of the art in health care domain. Computer Science Review **48**, 100546 (2023)

[27] Nemirko, A., Lugovaya, T.: Biometric human identification based on electrocardiogram. In: Proceedings of the XIIIth Russian Conference on Mathematical Methods of Pattern Recognition, Moscow, Russian. pp. 20–26 (2005)

[28] Olatunji, I.E., Rauch, J., Katzensteiner, M., Khosla, M.: A review of anonymization for healthcare data. Big data **12**(6), 538–555 (2024)

[29] Pereira, T.M., Conceição, R.C., Sebastião, R.: Initial study using electrocardiogram for authentication and identification. Sensors **22**(6), 2202 (2022)

[30] Pinto, J.R., Cardoso, J.S., Lourenço, A.: Evolution, current challenges, and future possibilities in ecg biometrics. Ieee Access **6**, 34746 (2018)

[31] Regazzoni, F., Acs, G., Aszalos, A.Z., Avgerinos, C., Bakalos, N., Berral, J.L., Bos, J.W., Brohet, M., Sanz, A.G.C., Davies, G.T., et al.: Secured for health: Scaling up privacy to enable the integration of the european health data space. In: 2024 Design, Automation & Test in Europe Conference & Exhibition (DATE). pp. 1–4. IEEE (2024)

[32] Salloum, R., Kuo, C.C.J.: Ecg-based biometrics using recurrent neural networks. In: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 2062–2066. IEEE (2017)

[33] Seastedt, K.P., Schwab, P., O'Brien, Z., Wakida, E., Herrera, K., Marcelo, P.G.F., Agha-Mir-Salim, L., Frigola, X.B., Ndulue, E.B., Marcelo, A., et al.: Global healthcare fairness: We should be sharing more, not less, data. PLOS Digital Health **1**(10) (2022)

[34] Sweeney, L.: Weaving technology and policy together to maintain confidentiality. The Journal of Law, Medicine & Ethics **25**(2-3), 98–110, 82 (1997). https://doi.org/10.1111/j.1748-720x.1997.tb01885.x

[35] Wang, Z., Kanduri, A., Aqajari, S.A.H., Jafarlou, S., Mousavi, S.R., Liljeberg, P., Malik, S., Rahmani, A.M.: ECG unveiled: Analysis of client re-identification risks in real-world ecg datasets. In: 2024 IEEE 20th IntL. Conf. on Body Sensor Networks (BSN). pp. 1–4. IEEE (2024)

[36] Zhang, Q., Zhou, D., Zeng, X.: Heartid: A multiresolution convolutional neural network for ecg-based biometric human identification in smart health applications. Ieee Access **5**, 11805–11816 (2017)