HeartBreaker: Practical Attacks for Recovering Cryptographic Keys Derived from Biosignals

Froudakis Evangelos[1] and Fabian Monrose[1]

Georgia Institute of Technology, Atlanta, Georgia, USA
froudakis@gatech.edu, fabian@ece.gatech.edu

Abstract. Physiological biosignals are increasingly used as a foundation for security-critical applications, including cryptographic key generation, due to their availability, uniqueness, and perceived resistance to forgery. We show that this premise is fundamentally flawed. Because biosignals arise from shared physiological processes, they are inherently coupled, and leakage from one modality can compromise another. We instantiate this risk through electrocardiograms (ECG), which capture cardiac electrical activity, and photoplethysmograms (PPG), which measure peripheral blood-volume changes. We present HeartBreaker, the first end-to-end attack that reconstructs a victim's ECG from leaked PPG signals and uses it to break state-of-the-art ECG-based key generation schemes. The attack exploits cross-modal physiological coupling using conditional generative models, combining two complementary strategies: a signal-centric approach targeting non-fiducial schemes, and an ECG-centric approach targeting fiducial morphology. Across two clinical datasets, the attack immediately recovers cryptographic keys for 47% of users under a fiducial scheme and 37% under a non-fiducial scheme. For the remaining users, it reduces the key search space by orders of magnitude, enabling practical key recovery within minutes. Experiments on real-world wearable data further confirm that this threat persists under noisy, consumer-grade conditions. Beyond this specific attack, our results highlight a broader principle: biosignals cannot be treated as independent sources of entropy. Security mechanisms built on physiological data must account for cross-modal inference, where compromise of one signal propagates to others. These findings suggest that evaluating biosignal-based security in isolation is insufficient and call for a reassessment of how physiological signals are used in cryptographic systems.

Keywords: ECG-based key generation · biosignal security · biometrics · key generation · side channels

1 Introduction

Physiological biosignals such as electrocardiograms (ECG), which record the heart's electrical rhythm, and photoplethysmograms (PPG), which track pulse and blood oxygen, are rapidly becoming central to digital health and personalized computing. As sensing hardware has become ubiquitous in everyday life, with

increasingly capable sensors embedded in watches, earbuds, and smart rings, biosignal acquisition has shifted from sporadic clinical measurement to continuous, real-world monitoring. This transition has generated volumes of physiological data that are increasingly leveraged beyond health, including for security and authentication [7,3,27,20,9]. Biosignals are appealing in this setting because they are passively collected, user-specific, difficult to imitate, and naturally refreshed over time [27,33]. These properties have motivated their use as a foundation for biometric authentication and cryptographic key generation [7,3,14].

Among the proposed modalities, ECG has emerged as a leading candidate for its stability and strong user specificity [14,20,17]. Early ECG-authentication systems proved brittle to electrode placement, posture, and noise [30], and susceptible to spoofing from partial or noisy observations of a victim's signal [24,10,25,26]. Recent work therefore shifted to *time-invariant cryptographic key generation*, which quantizes stable ECG features into bitstrings processed via fuzzy extractors [17,14,20], reframing the signal as a renewable source of entropy with stronger guarantees than authentication.

Importantly, ECG-based security is no longer confined to academic proposals. Commercial platforms such as B-Secur's HeartKey (an FDA-cleared ECG algorithm suite) and CardioID already deploy ECG-based key generation across wearables, automotive systems, and identity applications [3,7]. As such, the security of biosignal-based key generation is not a speculative concern, but an immediate, real-world issue. This raises a critical question: *are the assumptions underlying these systems robust in modern sensing environments?*

The security of these systems rests on an implicit assumption of acquisition hardness: obtaining a high-fidelity ECG is presumed to require physical contact, making it difficult for an adversary to capture inconspicuously [20,17,14]. This assumption is problematic because unlike ECGs, PPG signals are continuously collected by commodity wearables and are frequently exposed through weak device protections, permissive app permissions, third-party integrations, cloud synchronization, and large-scale breaches such as the 2024 Change Healthcare incident [5,40,2,8,39,37]. Moreover, prior work has shown that PPGs can be spoofed from remotely captured signals [25], reinforcing the concern that PPG should not be viewed as a protected channel. As a result, the fact that PPG signals are abundant and often exposed, creates a powerful, underexplored attack surface. Crucially, we take advantage of the fact that ECG and PPG signals are not independent: they are physiologically coupled measurements of the same cardiac cycle, linked through pulse arrival (PAT) [6]. Although previous studies have exploited this coupling for benign ends such as data augmentation [36,34,28], its security implications remain largely overlooked. In this work, we argue that as long as one modality can be inferred from another, the assumed secrecy of ECG is fundamentally weakened by the widespread availability of PPG.

To demonstrate that threat, we show how the relationship can be directly exploited to break ECG-based cryptographic key generation. We present HeartBreaker (Fig. 1), the first end-to-end attack that reconstructs a victim's ECG from leaked PPG signals and uses it to recover cryptographic keys.

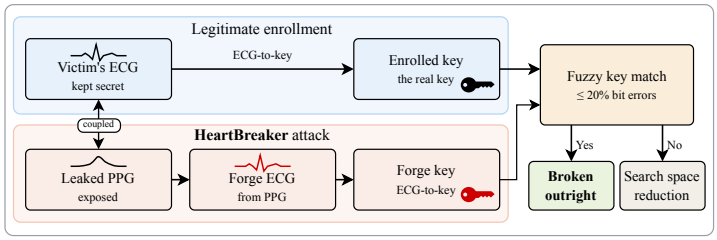

Fig. 1. Overview of HeartBreaker. Because ECG and PPG are coupled measurements of the same heartbeat, a leaked PPG trace can compromise the secret ECG. Whereas legitimate enrollment derives a user's key from a genuine ECG (top), we reconstruct the ECG from leaked PPG and forges a key with the same ECG-to-key algorithm (bottom). This process either recovers the key outright or significantly reduces the residual search space, leaving many keys brute-forceable within minutes.

As depicted in Figure 1, because these signals are coupled measurements of the same heartbeat, an attacker who obtains an exposed PPG can forge a key derived from the secret ECG. HeartBreaker exploits cross-modal physiological coupling using conditional generative models, combining two complementary strategies for attacking the prominent classes of biosignal-based approaches: a signal-centric approach targeting non-fiducial schemes, and an ECG-centric approach targeting fiducial morphology. Together, these strategies enable generalization to unseen users and produce ECG signals that satisfy the tolerance thresholds of real-world ECG-to-key systems. Our *contributions* are as follows:

- *Cross-modal generative attack.* We develop the first end-to-end attack that reconstructs ECG from PPG and uses it to compromise ECG-based cryptographic keys, demonstrating that leakage in one biosignal modality can undermine another.
- *Complementary attack strategies.* We introduce signal-centric and morphology-aware synthesis methods that succeed on complementary subsets of users, enabling robust key recovery across fiducial and non-fiducial schemes.
- *Empirical evaluation.* We evaluate our attacks on clinical and wearable datasets, showing immediate key recovery for a substantial fraction of users and dramatic reductions in residual key search space for the rest.
- *Security implications.* We demonstrate that cross-modal leakage is a practical threat and argue that evaluating biosignals in isolation is no longer tenable in modern sensing ecosystems.

2 Background

Physiological biosignals form a multilayered representation of cardiovascular function, each modality revealing different aspects of the same underlying process. In this section, we briefly describe the cardiac signals relevant to our threat

model, review time-invariant ECG-based key-generation schemes, and summarize prior work on biosignal synthesis and cross-modal translation.

2.1 Heart Signals

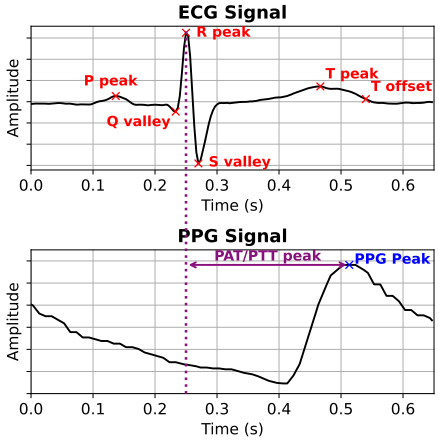

Fig. 2. Example ECG and corresponding PPG waveforms. The ECG's R-peak marks the heartbeat; the PPG rise follows after a short delay (PAT/PTT) as the pressure pulse reaches the peripheral sensor.

The electrocardiogram (ECG) and photoplethysmogram (PPG) are biosignals that capture related manifestations of cardiovascular activity [1]. ECGs measure the heart's electrical polarization cycles via surface electrodes, producing a waveform characterized by landmarks (e.g., QRS complexes) that reflect phases of cardiac excitation [35]. PPGs measure blood volume changes using an optical sensor, capturing the vascular response to each heartbeat (see Fig. 2).

2.2 Time-Invariant Key Generation from ECGs

Time-invariant key generation systems produce cryptographic keys that a user can reliably regenerate from their biosignal, without storing the key directly. These systems aim for high intra-subject reproducibility and low inter-subject similarity, using fuzzy extractors to tolerate physiological variation [17,14,20,9]. Seminal surveys of ECG biometrics distinguish among two families [38]: *fiducial* approaches that extract features from cardiac landmarks such as R-peak timing, QRS width, and interval durations [17]; and *non-fiducial* approaches that characterize the ECG through statistical or transform-domain descriptors without explicit landmark detection [20,14,16].

Because the two families depend on fundamentally different signal characteristics, an attack built for one class may fail against the other. Thus, to be broadly applicable, an adversary must adopt complementary strategies that reproduce morphological features for fiducial schemes, and matches statistical properties for non-fiducial schemes. Table 1 summarizes the distinction that motivates our two attack strategies: fiducial schemes require morphology-aware forgeries, whereas non-fiducial schemes require globally consistent waveform synthesis.

2.3 Generative Models for Biosignal Synthesis

Generative models enable the synthesis of realistic biosignals and translation between related modalities. These models learn a distribution over waveforms and

Table 1. Conceptual distinction between the two main families of ECG-based key-generation schemes and the corresponding attack requirements.

Scheme family	What it extracts	What the forged ECG must match
Fiducial	R peaks, QRS width, amplitudes, intervals	Beat morphology and landmarks
Non-fiducial	Wavelet/statistical descriptors over windows	Global temporal/spectral statistics

can be conditioned on an auxiliary signal—in our case, a victim's PPG—to produce a target output. Two broad families are relevant to our setting. *Adversarial* models (GANs) train a generator against a discriminator to produce realistic waveforms, with architectural choices determining which temporal and spectral properties are enforced [15,24,34]. *Diffusion* models learn to reverse a fixed noising process and have recently been adapted to 1D biosignal synthesis [28,36]. Within both families, we distinguish *signal-centric* architectures, which target global temporal and spectral statistics without cardiac-specific structure, from *ECG-centric* architectures, which focus on morphological features of the signal, such as the QRS complex. This taxonomy directly motivates our dual attack strategy in Section 3; an account of the selected models is given in Appendix B.

3 Methodology

We design an end-to-end attack framework that reconstructs a target's ECG to derive their ECG-based cryptographic key from exfiltrated PPG data. We train generative models on publicly available datasets to learn cross-modal mappings, and at attack time condition them on a leaked PPG trace to synthesize a target-specific ECG, which is fed into a key-generation framework to attempt key regeneration. The attack depicted in Figure 3 work as follows: the adversary conditions the generators on the leaked PPG (Step ①), the generated waveforms are fed to the downstream ECG-to-key algorithm that performs landmark detection (Step ②) and key extraction (Step ③), and the reconstructed key is tested for correctness in Step ④. To comprehensively evaluate resilience across algorithmic classes, we instantiate a *signal-centric* strategy, where a general diffusion model prioritizes high-fidelity waveform reconstruction and targets non-fiducial key systems, and an *ECG-centric* strategy, where models tailored to cardiac morphology target fiducial systems that rely on beat-wise features.

3.1 Threat Model

We consider an adversary whose goal is to recover a victim's cryptographic key generated from ECG using a time-invariant scheme. We model a realistic setting in which a large-scale breach of a consumer fitness platform [37,39] exposes users' PPG recordings, which are routinely collected for heart rate monitoring, sleep analysis, and activity tracking. A subset of these users relies on ECG-based key generation for security-critical tasks such as access control or transaction authorization. Given access to PPG data, the adversary applies the reconstruction techniques described in this work to synthesize victim-specific ECG signals

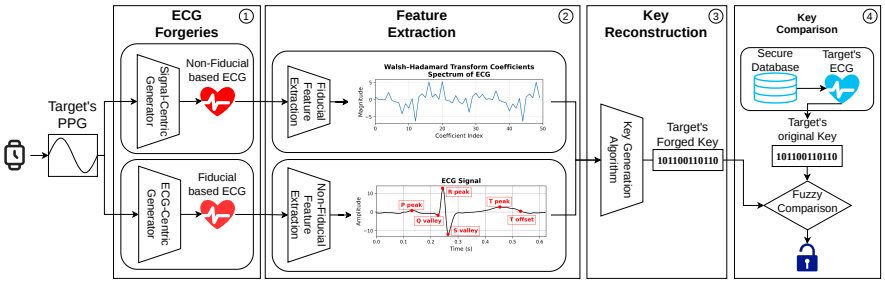

Fig. 3. Overview of the end-to-end PPG-conditioned ECG key-reconstruction attack. (1) A victim's PPG serves as input to two complementary generative models: a *signal-centric* generator focused on global waveform fidelity targeting non-fiducial keying schemes, and an *ECG-centric* generator that reproduces morphological features for fiducial schemes. (2) The synthesized ECGs are processed, and their fiducial/non-fiducial features are extracted. (3) ECG-to-key frameworks focusing on fiducial [17] or non-fiducial [14] features use these to generate a key, (4) which is then compared to the enrolled (true) key using a realistic fuzzy error correction comparison mechanism.

and derive the corresponding cryptographic keys entirely offline, without physical contact, access to ECG sensors, or interaction with the target system. Our threat model makes the following assumptions:

1. **Leaked PPG data.** The adversary has access to one or more PPG traces from the target user. Where available, the adversary may also make use of auxiliary signals co-located with the PPG (e.g., wrist acceleration, ethnicity).
2. **Training resources.** The adversary may obtain publicly available ECG and PPG datasets to train conditional generative models. Crucially, we *do not* require any ECG recordings from the target victim; models are trained across many subjects and conditioned on the victim's PPG at attack time.
3. **Compute and knowledge.** The adversary has modest compute (GPU training resources) and *no* knowledge of the general class of the target key algorithm (e.g., whether it uses fiducial or non-fiducial features), and does not require access to secret helper data (e.g., salts, fuzzy-extractor helper strings), nor to the victim's enrolled ECG templates.

These assumptions mirror real-world conditions under which PPG data are exposed at scale, and capture both strict and more permissive attacker capabilities relevant to evaluating the security of ECG-derived keys. Because ECG-to-key pipelines tolerate intra-subject variability via fuzzy extractors, we treat attack success in terms of whether the resulting key would be accepted after error correction, rather than reporting signal similarity alone.

3.2 Attack Strategy I: Signal-Centric Synthesis

The signal-centric strategy asks whether the global statistical and spectral properties of a victim's ECG can be reproduced over multi-beat windows with suffi-

cient fidelity. Rather than matching individual heartbeat morphology, it targets longer-term statistical consistency, making it well suited for non-fiducial schemes that rely on transform-domain or aggregate features.

Concretely, this strategy is instantiated using a conditional diffusion model [28] adapted to the PPG-to-ECG synthesis task. For training, we select 80% of each dataset's subjects as the training set. Each training example comprises a target ECG segment and the contemporaneous PPG trace as the conditioning input. We train a single global model to learn the conditional distribution $p(\text{ECG} \mid \text{PPG})$. At attack time, the victim's PPG segment is supplied as the conditioning input and the model produces an ECG candidate, which is fed directly to the ECG-to-key algorithm. This strategy produces waveforms with high global fidelity (adequate spectral content and temporal dynamics) and is robust to limited conditioning length and modest misalignment.

3.3 Attack Strategy II: ECG-Centric Synthesis

The ECG-centric strategy asks whether the fine-grained morphology of a victim's individual heartbeats can be reproduced with sufficient fidelity to fool a fiducial key-generation system. By prioritizing beat-level precision rather than long-range statistical consistency, it targets fiducial schemes that derive keys from landmark amplitudes, durations, and inter-peak intervals.

We evaluate this strategy using two adapted models that represent distinct state-of-the-art approaches to ECG-centric synthesis. The first is BioForge [24], a CycleGAN-based model that translates PPG to ECG without requiring paired training data while promoting subject-consistent latent representations. The second is the Region-Disentangled Diffusion Model (RDDM) [36], which focuses denoising on regions most relevant to fiducial key extraction, such as the QRS complex. Both models were modified from their original implementations to incorporate conditioning and task-specific changes required for reliable integration into the ECG-to-key evaluation pipeline; details are provided in Appendix B.

4 Evaluation

We evaluate our attack across two widely used datasets collected under controlled conditions. Our goal is to understand how generative models behave when trained and tested on synchronized ECG–PPG pairs exhibiting clean morphology, stable heart rhythms, and minimal motion artifacts, the operating conditions under which time-invariant ECG-to-key algorithms are typically designed and benchmarked [14,17,20]. We first justify our choice of ECG-to-key algorithms, describe the datasets and evaluation metrics, and outline the training procedures used for each generator. Finally, we present an empirical evaluation of attack success rates. Importantly, our evaluation targets the two dominant architectural designs embodied by deployed commercial ECG security systems [3,7].

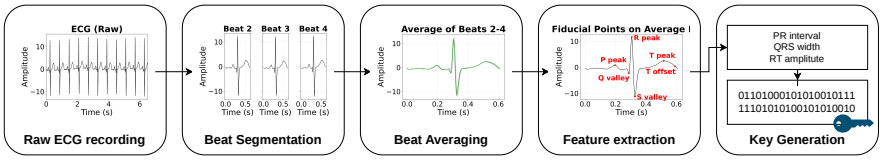

Fig. 4. Overview of fiducial-based ECG key generation. Raw ECG waveforms are segmented into individual beats, aligned, and averaged to obtain a representative waveform. Fiducial landmarks (peaks, valleys, intervals, and others) are then extracted, and a subset of these features is quantized to produce a stable binary key.

4.1 Model and Algorithm Selection

To ensure comprehensive coverage of the ECG-based key generation threat surface, we evaluate against one representative algorithm from each of the two dominant design families. For the fiducial family, we select the algorithm of Hwang *et al.* [17], which provides a complete instantiation of the fiducial pattern (wavelet-based landmark detection, beat-wise normalization, and clustering-based binarization feeding a fuzzy extractor). For the non-fiducial family, we select *Encryption by Heart* (EbH) [14], which instantiates transform-domain feature extraction over fixed-length windows with per-subject quantization, and is the only non-fiducial approach with a readily available implementation. Together, these two algorithms bracket the design space of ECG-based key generation; an attack compromising both directly assesses the threat surface facing ECG-based security systems.

For the generative models, we use one adversarial and two diffusion-based models, covering both attack strategies. BioForge [24] serves as our adversarial ECG-centric model, targeting beat-level morphology suited for fiducial schemes. RDDM [36] provides a diffusion-based ECG-centric perspective, concentrating denoising on QRS neighborhoods. TimeWeaver [28] represents the signal-centric approach, targeting global temporal and spectral consistency aligned with non-fiducial schemes. All three models were adapted from their original implementations for the PPG-to-ECG synthesis task.

Implementation details of the algorithms are provided in the Appendix.

4.2 Analysis of ECG-to-key Algorithms

Fiducial feature-based algorithm The fiducial approach of Hwang *et al.* [17] operates at the level of individual cardiac beats, relying on explicit detection of canonical ECG landmarks such as the onset and offset of the P wave, QRS complex, and T wave. A high-level representation of the pipeline is shown in Figure 4. This framework emphasizes reproducibility and bit-level stability, and its defining property from a security perspective is its sensitivity to per-beat morphology: a synthesized ECG must reproduce the victim's fiducial landmarks with sufficient fidelity in terms of waveform shape, amplitude, and inter-feature geometry to yield a key within the fuzzy extractor's acceptance radius.

Non-fiducial feature–based algorithm The non-fiducial approach of EbH avoids explicit landmark detection, instead characterizing the ECG through statistical and transform-domain descriptors computed over fixed-length windows, quantized into a binary key using per-subject thresholds and a tolerance margin that suppresses unstable components [14]. Because this design does not require beat segmentation or landmarking, it is robust to moderate misalignment, and its defining security property is its dependence on global spectral and statistical consistency across multi-beat windows rather than fine-grained beat morphology.

We verified that our implementations of the [17] and EbH [14] algorithms reproduce intra- and inter-subject key-distance behavior consistent with that reported in the original studies. Across both frameworks, we observe comparable distributions of within-subject and between-subject bit-error rates, confirming that our implementations faithfully replicate the reference designs.

4.3 Datasets

We evaluate on two multi-sensor datasets providing synchronized ECG and PPG recordings under controlled conditions (sufficient population, stable signals, and quasi-stationary cardiac rhythms): **BIDMC** [32] and **CapnoBase** [23].

The **BIDMC** (Beth Israel Deaconess Medical Center) [32] dataset contains 53 adult subjects with synchronized two-lead ECG and fingertip PPG signals, sampled at 125 Hz. The data were collected in resting-state conditions with consistent electrode placement and minimal motion artifacts, providing (high signal stability. We use 80% of the subjects for training and 20% for evaluation.

Likewise, the **CapnoBase** [23] dataset includes synchronized ECG and PPG recordings from 42 pediatric and adult subjects under controlled breathing protocols, sampled at 300 Hz. Compared to BIDMC [32], CapnoBase introduces modest physiological variability while maintaining high signal quality, enabling evaluation under more dynamic yet realistic conditions [23]. We apply the same 80/20 subject-wise split to maintain comparability across datasets.

During verification, the 4 subjects were deemed unfit for use with the fiducial-based algorithm due to unusually high intra-subject variability in the produced keys (bit error rates >10%): `S_0030` in CapnoBase and `bidmc_44`, `bidmc_53`, and `bidmc_24` in BIDMC. These subjects were excluded from our experiments.

4.4 Metrics

Our evaluation quantifies how close a synthesized key is to the enrolled key of a target subject. To do so, we report:

Hamming distance and BER: Let $k \in \{0,1\}^L$ denote the enrolled L-bit key for a subject and let $\hat{k} \in \{0,1\}^L$ be the key derived from a synthesized ECG. The *bit-error rate* (BER) is the normalized Hamming distance:

$$\mathrm{BER}(k, \hat{k}) = \frac{1}{L} \sum_{i=1}^{L} \mathbf{1}\{k[i] \neq \hat{k}[i]\}.$$

In what follows, we use *average Hamming distance* and *average bit-error rate* interchangeably to denote the sample mean of HD (or BER) across trials or subjects. Although BER provides a direct measure of similarity between spoofed and enrolled keys, it does not fully capture an attacker's success. In biometric key-generation systems, security can be overstated if one interprets an L-bit derived key without accounting for the system's built-in tolerance and error-correction mechanisms [18,4].

The security-relevant question is therefore not simply how many bits differ, but whether an attacker can produce a key the system would *accept* under its matching tolerance, and if not, how much brute-force effort remains. In practice, ECG-to-key systems accommodate nonzero intra-subject variability via fuzzy extractors, allowing two keys to match whenever their Hamming distance lies within an acceptance region. Concretely, let $k \in \{0,1\}^L$ denote the enrolled key and $\hat{k} \in \{0,1\}^L$ a candidate key; the system accepts $\hat{k}$ if

$$\mathrm{BER}(k, \hat{k}) \leq \Delta,$$

where $\mathrm{BER}(\cdot,\cdot)$ denotes Bit Error Rate and Δ is the error-correction capacity of the system [17,14,20,16]. Consequently, evaluating attacks solely using raw BER underestimates the effective risk, since any candidate key within the Hamming ball of radius Δ would be accepted as correct.

Δ-Residual Work Factor: To quantify how close an attack is to the system's acceptance threshold, we introduce the *Δ-Residual Work Factor*, which measures the additional effort an adversary must expend to reach the tolerated error level. Intuitively, this metric captures the *residual attacker work*: the remaining distance to key recovery under the system's built-in error tolerance. Formally, let $mean(\mathrm{BER})$ denote the mean bit-error rate achieved by the attacker across trials. We define the Δ-Residual Work Factor as

$$R_\Delta = (mean(\mathrm{BER}) - \Delta) \cdot L$$

where L is the key length and Δ is the system's acceptance threshold expressed as a fraction of bits. Smaller values of R_Δ indicate that the attacker is closer to producing a key that would be accepted by the system, even if the nominal BER remains above the acceptance threshold. In the experiments that follow, we adopt $\Delta = 0.20$, consistent with fuzzy-extractor tolerances reported in prior work [11,20,19,17], and refer to this instantiation as R_{20}.

4.5 Generator Training

To preserve realistic beat morphology and timing required by downstream ECG-to-key algorithms, we *retain each dataset's native sampling frequency* (125 Hz for BIDMC [32], 300 Hz for CapnoBase [23]) and do not perform any downsampling or resampling. The specific training parameters for TimeWeaver [28], BioForge [24], and RDDM [36] are given in Appendix B.

Table 2. Subject-level summary of key-recovery outcomes on the controlled datasets. A subject is *fully recovered* when BER $\leq 20\%$. A subject is considered *near-threshold* when $0 < R_{20} \leq 3$ for the fiducial scheme and $0 < R_{20} \leq 9$ for the non-fiducial scheme. The final column reports the Hamming-ball search space only for near-threshold subjects, computed as $N_{\text{res}}(L, g) = \sum_{i=0}^{g} \binom{L}{i}$ with $g = \lceil R_{20} \rceil$.

Dataset	Scheme	Subjects	Recovered	Near-threshold	Residual search
BIDMC	Fiducial	9	3	3	$\leq 8.2\times10^{2}$
CapnoBase	Fiducial	6	4	2	8.2×10^{2}–1.1×10^{4}
BIDMC	Non-fiducial	12	7	3	5.8×10^{7}–4.5×10^{13}
CapnoBase	Non-fiducial	7	0	4	2.2×10^{9}–9.3×10^{14}

4.6 Attack Success Evaluation

After training, each generator was run in inference mode to synthesize targeted ECGs for every subject in the held-out 20% test partition. We evaluate the success of each generative attack by comparing the keys derived from synthesized ECGs against the enrolled (true) keys of target subjects. Table 2 summarizes the results showing, for each subject, the generator with the lowest mean BER.

Fiducial based ECG-to-key algorithms Figure 5 summarizes the attacks against the fiducial ECG-to-key algorithm inspired by Hwang et al. [17]. On the BIDMC dataset [32], the signal-centric TimeWeaver [28] generator compromises a larger fraction of subjects than either BioForge [24] or RDDM [36]. RDDM [36] in particular never achieves an average BER distance below 25%, indicating limited success in reconstructing usable fiducial features. This likely stems from the lower signal amplitude and higher baseline noise of BIDMC recordings, which weaken the discriminative value of fiducial landmarks [23,32]. Under such conditions, a model that emphasizes global signal quality, rather than localized morphology, produces ECGs whose derived features happen to align more closely with those extracted by the fiducial-based algorithm.

In contrast, the CapnoBase [23] dataset, the ECG-centric models, BioForge [24] and RDDM [36], perform best, achieving the lowest average Hamming distances and, in most cases, reducing bit-error rates well below 20%. This outcome aligns with CapnoBase's higher signal quality and clearer cardiac morphology: since the framework extracts keys from fiducial landmarks such as R-peak timing and QRS width, morphology aware generators outperform signal-centric approaches.

Closer inspection of per-subject outcomes in Figure 5 further illustrates this variability. Subjects `bidmc_20` and `bidmc_22`, are successfully broken by TimeWeaver [28], while both BioForge [24] and RDDM [36] behave poorly. Conversely, `bidmc_30` shows the opposite trend: BioForge [24] achieves lower BER, whereas RDDM [36] and TimeWeaver [28] perform worse than random. These cases exemplify how the robustness of fiducial features varies across individuals. Even within a fiducial framework, some users exhibit low amplitude or noise-susceptible features, making them effectively less fiducial in practice and thus

more vulnerable to signal-centric reconstruction. Others possess strong, stable fiducials that favor morphology-aware generators like BioForge [24].

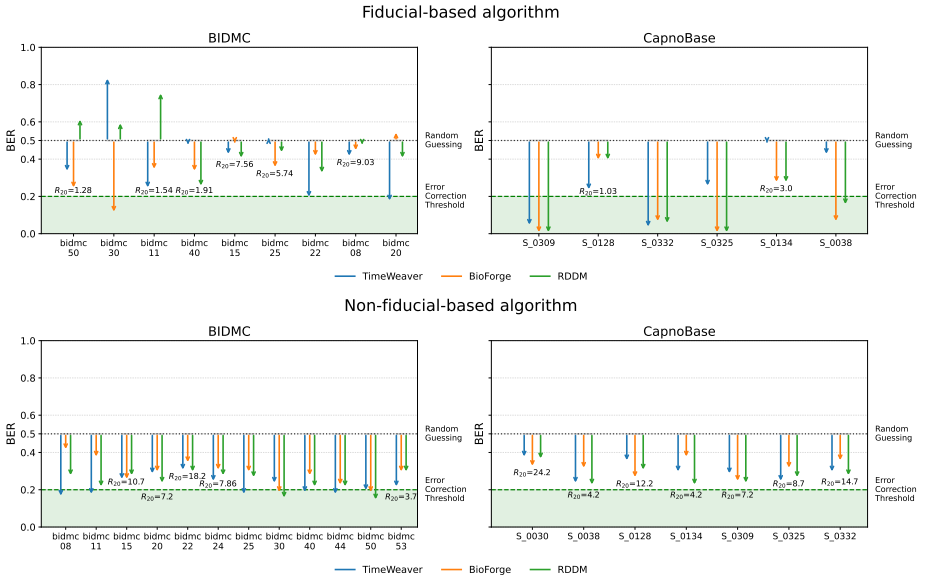

Fig. 5. Attack performance against both ECG-to-key algorithms on BIDMC and CapnoBase. **Top:** fiducial-based algorithm [17], subjects sorted by key length. **Bottom:** non-fiducial algorithm (*Encryption by Heart*) [14], subjects sorted alphabetically (fixed key length of 194 bits). Arrows show the average BER of forgeries generated using TimeWeaver [28], BioForge [24], and RDDM [36].

Following prior guidance that aggregate key-entropy metrics can be misleading [18,4], we interpret "not broken" cases by translating the residual gap to the Δ threshold into the corresponding brute-force search space, namely the size of the remaining Hamming ball. Even when our attacks do not fully recover a user's key, they collapse the residual search space to a dramatically smaller scale.

On BIDMC [32], three of the six unbroken subjects fall within two bits of the $\Delta = 20\%$ threshold, reducing the search space by more than eight orders of magnitude. On CapnoBase [23], the two unbroken subjects lie 1.03 and 3 bits above the threshold, leaving small residual search spaces. In practice, an adversary can search these remaining candidates in milliseconds to seconds.

Takeaway: For fiducial schemes, the attack breaks 7 of 15 subjects. For 5 of the remaining 8, it make key recovery significantly easier by reducing the key search space by eight orders of magnitude.

Non-Fiducial based ECG-to-key algorithms As seen in Figure 5, on the BIDMC dataset [32], all three generators succeed in reducing the search space for the majority of subjects, with TimeWeaver [28] achieving the lowest average Hamming distance overall. This is consistent with the algorithm's reliance on transform-domain and statistical features rather than on morphological landmarks: models that prioritize global signal smoothness and frequency-domain fidelity are better aligned with its feature extraction process. In this setting, TimeWeaver [28] breaks most users, reaching pre-error-correction bit-error rates below the 20% threshold for several subjects, demonstrating that signal-centric synthesis alone is a significant threat to non-fiducial keying schemes.

At the same time, the ECG-centric generators are not redundant. Both BioForge [24] and RDDM [36] reduce the search space for users that TimeWeaver [28] does not compromise. For example, `bidmc_30` exhibits lower BER under the ECG-centric models than under TimeWeaver [28]. Thus, even non-fiducial key derivation depends partly on waveform morphology, and features with fiducial correlations can be spoofed by these generators. Such overlap between morphological and statistical descriptors broadens vulnerabilities in hybrid frameworks.

On CapnoBase [23], attack effectiveness is more limited. None of the models achieves a Hamming distance below 20%, indicating that no user key is fully recovered under standard fuzzy-extractor assumptions. However, both RDDM [36] and TimeWeaver [28] achieve comparable reductions in Hamming distance, with RDDM [36] performing marginally better. This parity likely reflects CapnoBase's higher signal quality and lower inter-beat variability, which constrain the advantage of signal-centric reconstruction while leaving little residual noise for ROI-aware denoising to exploit. GAN, by contrast, performs slightly worse on this dataset, suggesting reduced generalization to CapnoBase's cleaner signals.

As in the fiducial-based case, our attack dramatically reduces the remaining search space even for users whose keys we do not fully recover. The non-fiducial algorithm produces longer keys of 194 bits, so naïve brute force would require $2^{194} \approx 2.5 \times 10^{58}$ guesses. On BIDMC [32], three of the five unbroken users remain only 3.7–7.9 bits above the $\Delta = 20\%$ threshold, and on CapnoBase [23], four of the seven unbroken users lie 4.2–8.7 bits above the threshold. Assuming an offline attacker capable of 10^9 keys per second, these users are brute-forceable within a week, even though their BER never crosses the Δ threshold.

Takeaway: For non-fiducial schemes, the attack breaks 7 of 19 subjects and makes key recovery practical for 7 more by reducing the search space from 10^{58} to 10^{14}.

5 Real-World Feasibility and Robustness Evaluation

To assess the practicality and robustness of our attack under real-world, noisy, sensing conditions, we conduct an validity study using the GalaxyPPG dataset [31], which provides synchronized PPG and ECG recorded during naturalistic activities using consumer wearables (Samsung Galaxy Watch, Empatica E4, and

Polar H10) under conditions that introduce substantial motion artifacts, sensor noise, and physiological variability.

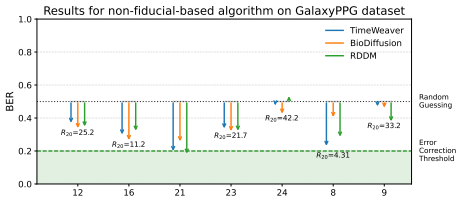

Fig. 6. Performance on the GalaxyPPG dataset against *EbH* [14] using forgeries generated using using TimeWeaver [28], BioForge [24], and RDDM [36].

We extract PPGs from the Empatica E4 as the leaked signal and ECGs from the Polar H10 as ground truth for evaluating key-generation performance. We upsample the PPG signals from 64 Hz to 130 Hz to match the sampling frequency of the ECG, as required by the conditional generative models. As before, we use 80% of participants for training each generative model and evaluate the attack on the remaining, preserving strict cross-subject separation. Unfortunately, we found that because the GalaxyPPG recordings contain strong motion artifacts and highly variable waveform shapes, the fiducial-based key-generation method of Hwang et al. [17] fails even on genuine signals. Most subjects show intra-user BERs close to 50%, far beyond the reproducibility required for time-invariant key generation. Thus, our real-world evaluation focuses on the non-fiducial scheme [14].

Despite the noise, motion artifacts, and variability, our attack remains effective. As seen in Figure 6, out of seven subjects, we achieve *key recovery* for subject 21, breaking the key outright. For subject 8, the attack reduces the R_{20} to 4.31, which means the brute-force effort becomes feasible in *less than one second* at a conservative 10^9 guesses/s. This confirms that cross-modal ECG reconstruction remains viable under the degradation conditions characteristic of commercial wearable sensing, underscoring that once PPG is leaked, an adversary can compromise ECG-derived cryptographic keys even in everyday environments.

Takeaway: The attack remains effective on noisy, consumer-grade wearable data, showing that ECG key compromise from PPG signals is feasible in real-world settings.

6 Discussion

Physiological signals have shifted from clinical measurements to continuously collected data streams embedded in everyday devices. As a result, biosignal-based security has moved from theory into real-world deployment. This shift weakens assumptions that once appeared sound, such as treating ECG as secret because it requires bodily contact. In practice, related modalities like PPG are widely collected and often exposed, creating indirect pathways for compromise.

Our results highlight a broader concern. Biosignals derived from the same physiological process are inherently coupled, and this coupling enables cross-modal inference. For instance, ECG and PPG are not independent sources of

information but different views of the same cardiac activity. When one modality is exposed, it can reveal another. Modern generative models make this practical, allowing attackers to reconstruct signals with enough fidelity to compromise downstream systems. This changes the adversarial model. Compromise no longer requires direct sensor access but can occur through correlated data sources. Thus, treating biosignals as independent sources of entropy is a flawed abstraction.

This perspective has important implications for evaluation. Biosignal-based security can no longer be assessed solely based on the difficulty of acquiring a specific signal. Instead, systems must be evaluated under conditions where related modalities may be observed or inferred. As shown in this paper, cross-modal leakage is not an edge case but a natural consequence of how physiological data is generated and collected. Indeed, our results aptly show that systems that may appear secure in isolation can fail in realistic settings.

These insights also inform future system design. Reducing unnecessary exposure of physiological data through stronger access controls, limited retention, and secure data handling can mitigate risk. At the algorithmic level, reliance on a single feature family is inherently fragile. Combining multiple features may help, but correlated features do not necessarily increase security and can create a false sense of entropy. Designs must therefore explicitly account for such dependencies and be evaluated against adversaries that exploit them. We admit that developing mechanisms that remain robust under cross-modal inference remains an important open problem, but without this shift, such systems risk providing only a false sense of security in modern, data-rich environments.

7 Other Related Work

A substantial body of work studies attacks on ECG biometrics, almost exclusively in the context of *authentication* rather than cryptographic key generation. Prior work has demonstrated presentation attacks on commercial ECG wristbands by synthesizing and replaying counterfeit waveforms, showing that ECG-based authentication is vulnerable to signal injection even under realistic constraints [10,21,13]. Subsequent work explores more systematic spoofing methods, including morphing attacks that interpolate between templates [29,22], dictionary-style attacks that search over template sets [41], and ML–based spoofing using variational autoencoders or adversarial perturbations [24,12,25]. These studies collectively show that ECG authentication systems can be fooled by crafted signals, but stop at breaking decision boundaries and do not examine implications for time-invariant key derivation or key search space.

Our work differs from these studies in three respects. First, we target *time-invariant cryptographic key generation*, not authentication. This shifts the focus from classification error rates to bit-level reproducibility and key search space, requiring different evaluation metrics and threat model. Second, we exploit *ECG–PPG coupling* to reconstruct ECG-derived keys from exfiltrated PPG, whereas previous work primarily evaluates spoofing within and across modalities at the decision level. Third, we analyze two attack paths, signal-centric and

ECG-centric, and show that they compromise disjoint subsets of users and algorithm classes (fiducial vs. non-fiducial keying), revealing structural weaknesses in ECG-based key generation not exposed by prior authentication-focused work.

More importantly, to the best of our knowledge, there is no prior work that mounts an end-to-end, generative, cross-modal attack on time-invariant ECG-based key generation using PPG as the primary leakage channel. Prior studies of ECG-based key generation focus on feature design, error correction, and entropy analysis [17,14,20], typically assuming that attacks would operate at the protocol level or via direct ECG compromise. Our results challenge this assumption by showing that modern generative models can bridge the cross-modal gap from PPG to ECG, materially reducing the key search space and achieving full key recovery for a nontrivial subset of users under realistic error-correction tolerances.

8 Conclusion

This work shows that physiological coupling between ECG and PPG fundamentally weakens the security assumptions underlying ECG-based key generation. By leveraging conditional generative models and two complementary attack strategies, we demonstrate that leaked PPG signals can be used to reconstruct ECG waveforms with sufficient fidelity to compromise cryptographic keys. Across controlled datasets, our attacks recover keys for a substantial fraction of users under both fiducial and non-fiducial schemes, and significantly reduce the remaining search space, in some cases turning an infeasible 2^{194} search problem into one solvable within days or weeks. Experiments on real-world wearable data further confirm that this threat persists outside controlled environments, where key recovery remains feasible despite noise and motion artifacts.

More broadly, our results show that the security of ECG should not be considered in isolation, given that related signals such as PPG are as readily available as they are today. Because biosignal-based cryptography inherits both the structure and the limitations of the physiological processes it relies on, when one signal can be inferred from another, assumptions of secrecy can quickly break down. These findings point to a broader lesson for the design of future systems: without accounting for cross-modal inference, such systems risk providing a false sense of security. Addressing this challenge will require rethinking how physiological signals are used in cryptographic applications and developing designs that remain robust in multi-modal sensing environments.

Acknowledgments.

The authors thank the anonymous reviewers for their detailed feedback that helped improve the paper. This research would not be possible without the generous support of the Julian T. Hightower Chair Professorship.

References

1. Agrafioti, F., Gao, J., Hatzinakos, D., Yang, J.: Heart biometrics: Theory, methods and applications. Biometrics **3**(199-216), 25 (2011)
2. Arias, O., Wurm, J., Hoang, K., Jin, Y.: Privacy and security in internet of things and wearable devices. IEEE transactions on multi-scale computing systems **1**(2), 99–109 (2015)
3. B-Secur: HeartKey® 2.0. `https://www.b-secur.com/solutions/heartkey/` (2024), accessed: April 2026
4. Ballard, L., Kamara, S., Reiter, M.K.: The practical subtleties of biometric key generation. In: USENIX Security Symposium. pp. 61–74 (2008)
5. Bayoumy, K., Gaber, M., Elshafeey, A., Mhaimeed, O., Dineen, E.H., Marvel, F.A., Martin, S.S., Muse, E.D., Turakhia, M.P., Tarakji, K.G., et al.: Smart wearable devices in cardiovascular care: where we are and how to move forward. Nature Reviews Cardiology **18**(8), 581–599 (2021)
6. Block, R.C., Yavarimanesh, M., Natarajan, K., Carek, A., Mousavi, A., Chandrasekhar, A., Kim, C.S., Zhu, J., Schifitto, G., Mestha, L.K., et al.: Conventional pulse transit times as markers of blood pressure changes in humans. Scientific Reports **10**(1), 16373 (2020)
7. CardioID: Next generation ECG biometric authentication. `https://cardioid.digital/` (2022), accessed: April 2026
8. Cilliers, L.: Wearable devices in healthcare: Privacy and information security issues. Health information management journal **49**(2-3), 150–156 (2020)
9. De Oliveira Nunes, I., Rindal, P., Shirvanian, M.: Oblivious extractors and improved security in biometric-based authentication systems. In: European Symposium on Research in Computer Security. pp. 290–312. Springer (2023)
10. Eberz, S., Paoletti, N., Roeschlin, M., Kwiatkowska, M., Martinovic, I., Patané, A.: Broken hearted: How to attack ECG biometrics. In: Network and Distributed System Security Symposium 2017. Internet Society (2017)
11. Garcia-Baleon, H., Alarcon-Aquino, V.: Cryptographic key generation from biometric data using wavelets. In: Electronics, Robotics and Automotive Mechanics Conference. pp. 15–20. IEEE (2009)
12. Garg, A., Karimian, N.: ECG biometric spoofing using adversarial machine learning. In: International Conference on Consumer Electronics. pp. 1–5. IEEE (2021)
13. Giechaskiel, I., Zhang, Y., Rasmussen, K.B.: A framework for evaluating security in the presence of signal injection attacks. In: European Symposium on Research in Computer Security. pp. 512–532. Springer (2019)
14. González-Manzano, L., de Fuentes, J.M., Peris-Lopez, P., Camara, C.: Encryption by heart (ebh)—using ECG for time-invariant symmetric key generation. Future Generation Computer Systems **77**, 136–148 (2017)
15. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Communications of the ACM **63**(11), 139–144 (2020)
16. Hernández-Álvarez, L., Barbierato, E., Caputo, S., de Fuentes, J.M., González-Manzano, L., Encinas, L.H., Mucchi, L.: Keyencoder: A secure and usable EEG-based cryptographic key generation mechanism. Pattern Recognition Letters **173**, 1–9 (2023)
17. Hwang, H.B., Lee, J., Kwon, H., Chung, B., Lee, J., Kim, I.Y.: Preliminary study of novel bio-crypto key generation using clustering-based binarization of ECG features. Sensors **24**(5), 1556 (2024)

18. Jain, A.K., Ross, A., Uludag, U.: Biometric template security: Challenges and solutions. In: European signal processing conference. pp. 1–4. IEEE (2005)
19. Jin, Z., Teoh, A.B.J., Goi, B.M., Tay, Y.H.: Biometric cryptosystems: a new biometric key binding and its implementation for fingerprint minutiae-based representation. Pattern Recognition **56**, 50–62 (2016)
20. Karimian, N., Guo, Z., Tehranipoor, M., Forte, D.: Highly reliable key generation from electrocardiogram (ECG). IEEE Transactions on Biomedical Engineering **64**(6), 1400–1411 (2016)
21. Karimian, N., Woodard, D., Forte, D.: ECG biometric: Spoofing and countermeasures. IEEE Transactions on Biometrics, Behavior, and Identity Science **2**(3), 257–270 (2020)
22. Karimian, N., Woodard, D.L., Forte, D.: On the vulnerability of ECG verification to online presentation attacks. In: 2017 IEEE International Joint Conference on Biometrics (IJCB). pp. 143–151. IEEE (2017)
23. Karlen, W., Turner, M., Cooke, E., Dumont, G., Ansermino, J.M.: Capnobase: Signal database and tools to collect, share and annotate respiratory signals. In: 2010 Annual meeting of the society for technology in anesthesia. p. 27. Society for Technology in Anesthesia (2010)
24. Krish, V., Paoletti, N., Kazemi, M., Smolka, S., Rahmati, A.: Biosignal authentication considered harmful today. In: USENIX Security Symposium. pp. 5521–5536 (2024)
25. Li, L., Chen, C., Pan, L., Tai, Y., Zhang, J., Xiang, Y.: Hiding your signals: A security analysis of PPG-based biometric authentication. In: European Symposium on Research in Computer Security. pp. 183–202. Springer (2023)
26. Li, L., Chen, C., Pan, L., Zhang, J., Xiang, Y.: Video is all you need: Attacking PPG-based biometric authentication. In: Proceedings of the 15th ACM Workshop on Artificial Intelligence and Security. pp. 57–66 (2022)
27. Monrose, F., Reiter, M.K., Li, Q., Wetzel, S.: Cryptographic key generation from voice. In: Proceedings IEEE Symposium on Security and Privacy. S&P 2001. pp. 202–213. IEEE (2001)
28. Narasimhan, S.S., Agarwal, S., Akcin, O., Sanghavi, S., Chinchali, S.: Time weaver: A conditional time series generation model. arXiv preprint arXiv:2403.02682 (2024)
29. Noh, S., Kim, J., Lee, S., Kang, Y., Park, C., Shin, Y.: Broken heart: Privacy leakage analysis on ECG-based authentication schemes. Security and Communication Networks **2022**(1), 7997509 (2022)
30. Odinaka, I., Lai, P.H., Kaplan, A.D., O'Sullivan, J.A., Sirevaag, E.J., Rohrbaugh, J.W.: ECG biometric recognition: A comparative analysis. IEEE Transactions on Information Forensics and Security **7**(6), 1812–1824 (2012)
31. Park, S., Zheng, D., Lee, U.: A PPG signal dataset collected in semi-naturalistic settings using galaxy watch. Scientific Data **12**(1), 892 (2025)
32. Pimentel, M.A.F., Johnson, A.E.W., Charlton, P.H., Birrenkott, D., Watkinson, P.J., Tarassenko, L., Clifton, D.A.: Toward a robust estimation of respiratory rate from pulse oximeters. IEEE Transactions on Biomedical Engineering **64**(8), 1914–1923 (2017). https://doi.org/10.1109/TBME.2016.2613124
33. Rostami, M., Juels, A., Koushanfar, F.: Heart-to-heart (h2h) authentication for implanted medical devices. In: Proceedings of the 2013 ACM SIGSAC conference on Computer & Communications Security. pp. 1099–1112 (2013)
34. Sarkar, P., Etemad, A.: Cardiogan: Attentive generative adversarial network with dual discriminators for synthesis of ECG from PPG. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 488–496 (2021)

35. Serhani, M.A., T. El Kassabi, H., Ismail, H., Nujum Navaz, A.: ECG monitoring systems: Review, architecture, processes, and key challenges. Sensors **20**(6), 1796 (2020)
36. Shome, D., Sarkar, P., Etemad, A.: Region-disentangled diffusion model for high-fidelity PPG-to-ECG translation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 38, pp. 15009–15019 (2024)
37. U.S. Department of Health and Human Services, Office for Civil Rights: Breach portal: Notice to the secretary of HHS breach of unsecured protected health information. `https://ocrportal.hhs.gov/ocr/breach/breach_report_hip.jsf` (2026), accessed: April 2026
38. Uwaechia, A.N., Ramli, D.A.: A comprehensive survey on ECG signals as new biometric modality for human authentication: Recent advances and future challenges. IEEE Access **9**, 97760–97802 (2021)
39. Witty, A.: Responses to questions for the record for Andrew Witty. Tech. rep., U.S. Senate Committee on Finance (2024), full Committee Hearing, May 1, 2024. `https://www.finance.senate.gov/imo/media/doc/responses_for_questions_for_the_record_to_andrew_witty.pdf`
40. Zhang, B., Chen, C., Lee, I., Lee, K., Ong, K.L.: A survey on security and privacy issues in wearable health monitoring devices. Computers & Security p. 104453 (2025)
41. Zhang, B., Li, L., Chen, C., Lee, I., Lee, K., Zhu, T., Ong, K.L.: A novel dictionary attack on ECG authentication system using adversarial optimization and clustering. Knowledge-Based Systems **316**, 113326 (2025)

A ECG-to-Key Algorithm Details

Fiducial Algorithm: Hwang et al. Hwang et al. [17] is a fiducial ECG-to-key scheme that operates on individual beats. After denoising and baseline correction, beats are segmented and fiducial points are detected with a wavelet-based delineator. From each beat, the method extracts 29 fiducial features spanning amplitude, duration, slope, distance, and area. Low-correlation beats are discarded as outliers, and the remaining beats are normalized to a 70 bpm reference to reduce heart-rate variability. Features are then quantized using personalized k-means clustering ($k = 10$), converted to binary codes, concatenated into a subject-specific seed, and passed through a fuzzy extractor to produce the key. Because quantization is personalized, the key length varies by subject.

Non-Fiducial Algorithm: Encryption by Heart (EbH) EbH [14] is a non-fiducial ECG-to-key scheme that uses discrete wavelet transform coefficients and temporal statistics over fixed-length windows. During enrollment, each subject's model stores representative coefficient distributions and magnitude thresholds from baseline windows. At key-generation time, a new ECG segment is transformed with the same basis, and features exceeding subject-specific thresholds are quantized into binary values based on magnitude and stability. A tolerance margin suppresses unstable bits to improve reproducibility. The resulting 194-bit bitstring serves as the subject's cryptographic seed, and the design is robust to moderate misalignment because it does not require beat segmentation.

B Generative Model Details

BioForge [24] is an ECG-centric contrastive CycleGAN for unpaired PPG-to-ECG translation. It uses an encoder–decoder generator with attention-enhanced skip connections and dual time- and frequency-domain discriminators to promote morphological fidelity and realistic spectral content. Its objective combines adversarial, cycle-consistency, identity, and supervised-contrastive losses. The supervised-contrastive term encourages same-subject latent representations to cluster together, biasing generation toward subject-specific morphology even in an unpaired. In our implementation, we follow the original architecture and loss design, modify the code base to avoid downsampling the input signals, and train the model for 24 epochs as in [24].

RDDM [36] is an ECG-centric diffusion model that focuses denoising capacity on morphology-critical signal regions. It modifies the diffusion processes using ROI masks centered on R-peaks to preserve QRS structure. The reverse denoiser uses a high-frequency path for QRS recovery and a complementary path for lower-frequency P/T regions and inter-beat baselines, with PPG conditioning injected at multiple temporal resolutions. ROI masks are derived from the default detector-based proxies, and the model is trained for 1000 epochs with its default objective plus landmark and spectral auxiliary losses [36].

TimeWeaver [28] is a signal-centric conditional diffusion model for multivariate time series. It treats PPG and ECG as channels of a joint signal and uses a dedicated conditioning pathway within a UNet-based denoiser to preserve long-range temporal dependencies and cross-channel alignment. Unlike RDDM, it does not explicitly model QRS regions or other fiducials, instead focusing on global cadence, inter-beat timing, and overall spectral structure under PPG conditioning. We use the default training configuration of [28] and train for 500 epochs. Conditioning channels are dataset-specific: BIDMC [32] uses wrist EDA, wrist temperature, PPG, and three-axis wrist acceleration, while CapnoBase [23] uses age, weight, and PPG. All channels are normalized and processed through TimeWeaver's conditioning encoder, and the model is trained globally across pooled training subjects without victim-specific fine-tuning [28].

C Ethical Considerations

Our study has dual-use implications because it demonstrates attacks on ECG-based key-generation schemes. To minimize harm, we use only previously published datasets [32,23,31], collect no new biosignal data, involve no human subjects, and make no attempt to re-identify individuals. All data are used solely to evaluate the studied constructions under subject-disjoint protocols.

D Open Science

To support open science, we make the datasets, model implementations, and ECG-to-key algorithms used in this study available at:

`https://github.com/EvanFr/HeartBreaker`