2026 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE) | 979-8-3195-0752-5/26/$31.00 ©2026 IEEE | DOI: 10.1109/CHASE69719.2026.00030

When PPG Meets ECG: Uncovering the Risks of Cross-Modal Biometric Matching

Linfeng Zhu, Kunning Shen, Chenhan Xu, and Huining Li
Department of Computer Science
North Carolina State University
Raleigh, USA
{lzhu28, kshen5, cxu34, hli83}@ncsu.edu

***Abstract*—The rapid proliferation of wearable technology has led to the ubiquitous collection of physiological data, particularly Photoplethysmography (PPG), which shares a deep biological link with clinical Electrocardiogram (ECG) signals. While institutions frequently release large-scale ECG datasets for research under the assumption that removing explicit identifiers (e.g., names) ensures anonymity, this paper uncovers a critical and underexplored privacy vulnerability: an adversary can re-identify specific patients in "de-identified" public ECG databases by matching them against PPG data harvested from consumer wearables. To demonstrate this vulnerability, we propose a novel deep learning framework based on a Pseudo-Siamese network architecture. This model independently extracts features from ECG and PPG signals and maps them into a shared biometric latent space, effectively capturing the intrinsic physiological identity consistency between the two modalities. Comprehensive experiments conducted on the BIDMC dataset provide empirical evidence that cross-modal matching is not only feasible but effective. Our findings confirm that current de-identification protocols are insufficient to protect against biometric re-identification, exposing individuals to severe risks such as medical discrimination and privacy loss.**

***Index Terms*—ECG, PPG, Cross-Modal Biometric Matching, Contrastive Learning, Privacy**

I. Introduction

Heart monitoring has become a cornerstone of both reactive and proactive healthcare, as it provides critical insights into cardiovascular health across diverse environments. In professional clinical settings, high-fidelity Electrocardiogram (ECG) data is routinely collected to diagnose arrhythmias and other cardiac pathologies with medical-grade precision. Meanwhile, the rise of wearable technology has shifted monitoring into the sphere of daily life, where billions of data points are captured via Photoplethysmography (PPG) sensors in smartwatches and fitness trackers [1], [2].

To facilitate medical research, institutions and companies frequently release large-scale ECG datasets (e.g., MIMIC-III, PTB-XL) to the public [3], [4]. These datasets generally are claimed to be de-identified, where explicit identifiers (e.g., names, social security numbers) are removed to comply with privacy regulations. However, a critical yet underexplored privacy vulnerability persists: physiological signals themselves contain unique biometric identity patterns [5], [6]. In light of this, we must evaluate a fundamental security challenge: *can an adversary, such as a malicious insider or an unauthorized third-party application, re-identify a specific patient in a "de-identified" public ECG database by matching it against PPG data harvested from common wearable devices?* If such cross-modal biometric matching is feasible, the identity of patients in public medical datasets is no more "anonymous", as the physiological "fingerprint" remains linked to the individual. The consequences extend far beyond a simple loss of privacy, potentially leading to severe real-world harms such as medical discrimination in insurance or employment, invasive targeted advertising based on sensitive health markers, and potential legal or social consequences for the individuals being identified [7].

The feasibility of this cross-modal attack is grounded in the physiological coupling between the heart's electrical activity (ECG) and the resulting mechanical blood volume changes (PPG), suggesting that unique biometric patterns may be mirrored across both signals [8]–[10]. While recent deep learning research has demonstrated that this biological link enables the reconstruction of ECG waveforms from PPG data [11], these studies primarily focus on clinical fidelity rather than biometric preservation. However, this reconstruction capability provides new opportunities: if the reconstructed signals retain sufficient morphological distinctiveness, they could theoretically serve as a bridge for identification. Similar cross-modal linkages have been explored in other biometric domains. For example, prior work aligns identity features between faces and voices, enabling identity inference by matching a person's face to their voice, or vice versa [12], [13].

In this paper, we propose a deep learning framework for ECG-PPG biometric matching. We leverage a Pseudo-Siamese network architecture to independently extract features of ECG and PPG signals. By mapping these cross-modal signals into a biometric latent space, we reveal the intrinsic physiological identity consistency between the two modalities, suggesting the risks of cross-dataset linkage attacks. The contribution of this paper can be summarized as follows:

- We identify the threat of cross-dataset linkage attacks between wearable PPG data and clinical ECG records, highlighting a significant privacy vulnerability.
- We propose a specialized framework utilizing a Pseudo-Siamese network to effectively capture and align the identity-related features across ECG and PPG signals.
- We conduct comprehensive experiments on BIDMC

2832-2975/26/$31.00 ©2026 IEEE
DOI 10.1109/CHASE69719.2026.00030

Authorized licensed use limited to: N.C. State University Libraries - Acquisitions & Discovery S. Downloaded on November 07,2026 at 21:06:13 UTC from IEEE Xplore. Restrictions apply.

dataset, providing empirical evidence that patient privacy can be leaked through cross-dataset linkage.

II. Methodology

Our goal is to verify whether the given PPG and ECG segments originated from the same identity. Existing research provides a signal translation solution (i.e., using GANs to reconstruct ECG from PPG before matching [14]) for this task. However, such generative methods would introduce severe reconstruction errors, which corrupt the biometric patterns essential for identification. To address this problem, we propose an end-to-end metric learning framework that maps signals directly into a common latent space.

The overall architecture of our proposed cross-dataset linkage framework is illustrated in Figure 1. The pipeline consists of three phases: First, raw ECG and PPG signals are preprocessed to filter noise and are segmented into aligned windows. Second, identity features of the two signals are extracted independently using a Pseudo-Siamese network backbone with SE-ResNet encoders, and a contrastive loss function is applied to map these features into a common latent space, minimizing the distance between matched pairs while maximizing the distance between unmatched ones. Finally, by comparing the similarity score of the given ECG-PPG pair, we can tell if the ECG and PPG segments belong to the same identity. The following subsections detail the matching process.

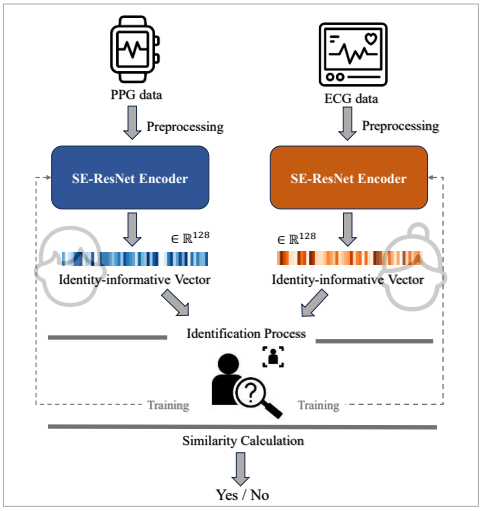


Fig. 1: System overview.

A. Data Preprocessing

Physiological signals are susceptible to environmental noise and motion artifacts. To address this, we apply a systematic preprocessing pipeline to both ECG and PPG signals. Figure 2 illustrates the preprocessing effect on ECG data, demonstrating the transformation from noisy raw signal to clean normalized waveform. The following subsections detail each step.

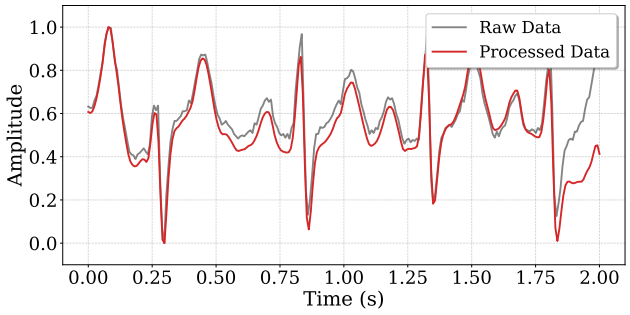


Fig. 2: Visualization of raw and preprocessed ECG signals. The raw Lead II signal (gray) exhibits high-frequency noise, while the processed signal (red) demonstrates the result of resampling, NeuroKit2 cleaning, and normalization.

1) Signal Filtering and Segmentation: First, raw signals are preprocessed using NeuroKit2 library [15] for artifact removal. The continuous PPG and ECG recordings are then resampled to 128 Hz. For segmentation, we employ a sliding window approach with a window size of 4 seconds and a 0.5-overlap, which provides sufficient temporal context for cardiac information while remaining short enough for real-time application. Finally, data normalization is performed on each ECG and PPG segment to eliminate amplitude scaling differences caused by varying sensor contact.

2) Stochastic Data Augmentation: To avoid model overfitting and enhance generalization ability, we employ stochastic data augmentation during the training phase only. The transformations simulate realistic signal variations without altering the underlying personal morphology, which include random amplitude scaling (0.8-1.2), temporal shifting ($\pm$117ms), and additive Gaussian noise with intensity randomly scaled between 0% and 5%.

B. Pseudo-Siamese Network

Standard Siamese networks [16] share weights between branches to find similarities in homogeneous data (e.g., matching two face images). In our case, although both ECG and PPG reflect cardiac activities, the two signals exhibit distinct morphological characteristics. Forcing two branches to share weights would prevent the network from learning modality-specific features, resulting in poor convergence. The weight-sharing architecture (standard Siamese Network) is therefore not suitable.

To address this, we employ a Pseudo-Siamese architecture consisting of two independent encoders with identical structures to extract features of ECG and PPG separately. This design allows each branch to learn the optimal feature extraction procedure for specific modality (ECG or PPG) while the high-level layers are guided by the loss function to project these features into a unified latent space. The core of our framework is a custom 1D SE-ResNet Encoder designed specifically for physiological signal analysis.

Authorized licensed use limited to: N.C. State University Libraries - Acquisitions & Discovery S. Downloaded on November 07,2026 at 21:06:13 UTC from IEEE Xplore. Restrictions apply.

As illustrated in Figure 3, our encoder processes the input signal $\mathbf{X} \in \mathbb{R}^{1\times L}$ through a hierarchical structure to extract abstract semantic features. The architecture mainly consists of three stages. First, a Stem Layer begins with a large-kernel convolutional layer ($k = 7, s = 2$) to capture long-range temporal dependencies, followed by batch normalization, ReLU activation, and a Max Pooling layer ($k = 3, s = 2$). Next, following the Stem Layer, the backbone consists of four sequential 1D residual blocks, each doubling the number of feature channels ($64 \rightarrow 128 \rightarrow 256 \rightarrow 512$). Finally, a dual-pooling module is applied to aggregate temporal information, and a fully connected layer produces a compact embedding bector. The detailed designs of the residual block and the dual-pooling module are detailed below.

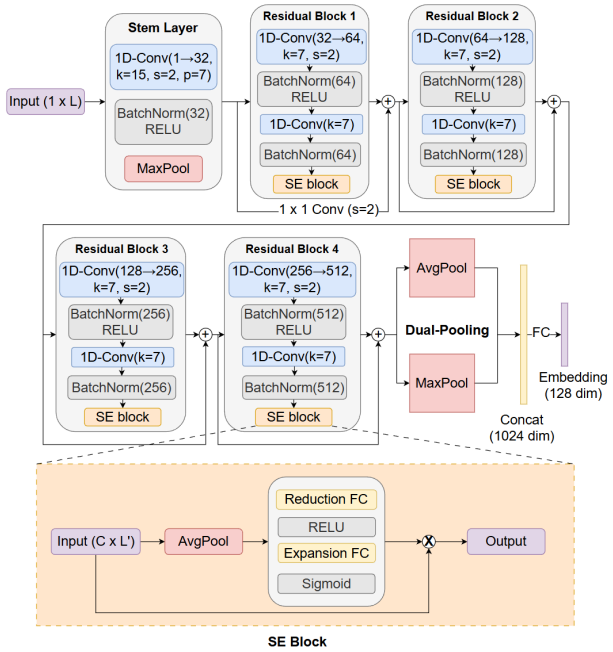


Fig. 3: SE-ResNet Encoder Architecture.

1) 1D Residual Block: The basic building component of our encoder is the 1D Residual Block, designed to capture temporal dependencies while preventing the vanishing gradient problem in deep networks. Unlike the standard ResNet for images which uses 3×3 kernels, our 1D block utilizes a kernel size of k=7. Each residual block consists of two (7x7) convolutional layers. This large receptive field is crucial for physiological signals to capture broader temporal contexts within a single layer. The first convolutional layer applies a stride $s = 2$ followed by batch normalization and a ReLU activation. The second layer performs a convolution with $s = 1$.

To enhance model robustness against motion artifacts common in wearable PPG, we integrate Squeeze-and-Excitation (SE) blocks [17] within the residual units, which can suppress noise-related channels by computing channel attention weights of each channel. The SE block first computes a channel-wise descriptor z through a Global Average Pooling layer, followed by a gating mechanism with two fully connected (FC) layers and a sigmoid activation. The channel attention weights are computed as follows:

$$s = \sigma(W_2 \cdot \delta(W_1 \cdot z)), \tag{1}$$

where δ and σ refer to the ReLU and Sigmoid activation function, respectively. W_1 and W_2 are two FC layers, where $W_1 \in \mathbb{R}^{\frac{C}{r}\times C}$ is the dimensionality-reduction layer with a reduction ratio r (set to 16), and $W_2 \in \mathbb{R}^{C\times\frac{C}{r}}$ is the dimensionality-increasing layer.

2) Dual Pooling: Standard ResNets typically use only Average Pooling, which smooths out signal details. However, physiological signals contain sharp diagnostic features (e.g., R-peaks in ECG) that are critical for identification. Therefore, we employ a Dual Pooling strategy that concatenates the outputs of both Global Average Pooling and Global Max Pooling.

Let $F \in \mathbb{R}^{C\times L'}$ be the feature map from the last residual layer ($C = 512$). We compute both the average and maximum values across the temporal dimension and concatenate them:

$$v = [\text{GAP}(F) \oplus \text{GMP}(F)] \in \mathbb{R}^{2C} \tag{2}$$

The resulting 1024-dimensional vector is then passed through a FC layer to be projected into a 128-dimensional latent space.

C. Contrastive Loss

We aim to map ECG and PPG features into a common latent space where distance metric can reflect identity information. We choose cosine similarity as our distance metric and utilize the Cosine Embedding Loss to optimize the feature alignment. Unlike Euclidean distance, which is sensitive to the magnitude of feature vectors, cosine similarity focuses on the angle between vectors. This is particularly suitable for physiological signals, where the absolute amplitude often fluctuates due to external factors, but the relative morphological shape remains consistent for the same identity. Let $f_e(\cdot)$ and $f_p(\cdot)$ denote the ECG and PPG encoders, respectively. For a given ECG-PPG pair (x_e, x_p) with label y ($y = 1$ for the same identity, $y = -1$ for the different identity), the loss is defined as follows:

$$\mathcal{L} = \begin{cases} 1 - \cos(f_e(x_e),\ f_p(x_p)), & \text{if } y = 1 \\ \max(0,\ \cos(f_e(x_e),\ f_p(x_p)) - m), & \text{if } y = -1 \end{cases} \tag{3}$$

where $\cos(\cdot)$ denotes cosine similarity and m is the margin hyperparameter (set to 0.3). This loss function enforces embeddings of the same identity to be close while those of different identities are far apart, enabling robust cross-modal verification.

III. Experiments

A. Dataset

For evaluation, we utilized the BIDMC PPG and Respiration Dataset [18], which contains waveform signals of 53 8-minute recordings (46 different patients) extracted from the much larger MIMIC II matched waveform Database [19]. To construct the experimental dataset and prevent testing data

Authorized licensed use limited to: N.C. State University Libraries - Acquisitions & Discovery S. Downloaded on November 07,2026 at 21:06:13 UTC from IEEE Xplore. Restrictions apply.

leakage, we performed a 5-fold cross-validation on 46 unique patients, which ensures that the model was always trained and tested on different individuals. In the first four folds, the data were split into a training set of 37 patients and a test set of 9 patients. The final fold consists a training set of 36 patients and a test set of 10 patients. Then, for each patient in the training set, we randomly sampled 1000 ECG-PPG segment pairs, with an equal number of positive (segments from the same identity) and negative (segments from the different identity) samples. For each patient in the unseen test set, 500 segment pairs were sampled for evaluation.

Since BIDMC is derived from MIMIC-II database, it inherits certain technical limitations that pose significant challenges for biometric matching. First, data collection procedures occasionally involved monitors not being reset between patients, leading to single files containing data from different patients. This potential label noise may persist in BIDMC dataset, influencing the matching accuracy. Additionally, the ECG signals are downsampled from high-frequency recordings. Although they are stored at 125 Hz, the actual sampling intervals fluctuate and the effective amplitude resolution is often reduced, which may distort fine-grained signal features essential for high-precision identification.

B. Main Results

TABLE I: Performance evaluation using 5-Fold Cross-Validation on the BIDMC dataset. The AUC and Accuracy (ACC) for each fold and their average.

Fold Index	**AUC**	**ACC (%)**
Fold 1	0.7814	74.68
Fold 2	0.7058	68.84
Fold 3	0.8258	76.76
Fold 4	0.8794	82.29
Fold 5	0.7905	72.38
Average	**0.7966**	**74.99**

To evaluate the risks of privacy leakage via cross-dataset linkage, we formulated the task as a biometric matching problem, utilizing Area Under the ROC Curve (AUC) and accuracy (ACC) as performance metrics. To ensure the robustness of our results, we conducted a 5-fold cross validation experiment.

As the results shown in Table I, our proposed framework achieves an average AUC of 0.7966 and an average accuracy of 74.99%. The model performs better in certain folds (e.g., Fold 3, Fold 4) with AUCs exceeding 0.8 while a performance drop is noticed in Fold 2 (AUC=0.7058). The performance variation could be attributed to the patient diversity of BIDMC dataset. Given the dataset's wide age range from 19 to 90 but relatively small size, it is statistically plausible that certain age ranges present in the test set were not seen in the training set, which causes the performance degradation in certain folds. Despite these fluctuations, the overall performance is enough to confirm that identity-related features are preserved across ECG and PPG signals.

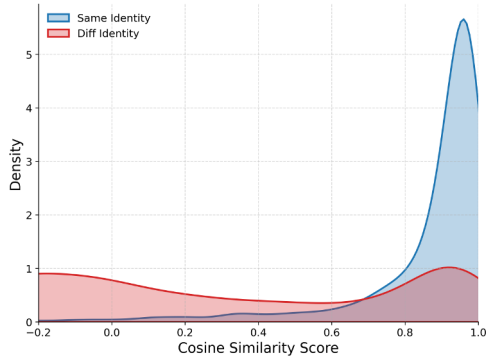


Fig. 4: Distribution of cosine similarity scores for positive pairs and negative pairs on the test set.

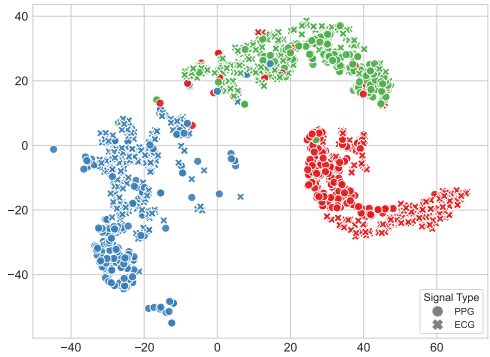


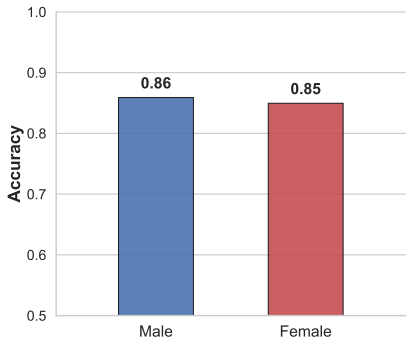
Fig. 5: t-SNE Visualization of the learned feature embeddings in the common latent space, where different colors represent different patient identities

The cosine similarity scores distribution is shown in Figure 4 for positive and negative ECG-PPG pairs. A clear separation is observed between the two distributions: pairs from the same identity typically exhibit high similarity scores (around 0.8-1.0), while pairs from different identities are distributed across lower values. The minimal overlap between these distributions provides intuitive evidence that the learned common latent space successfully aligns the two signals by identity. Despite the distinct morphological differences between the ECG and PPG signals, the intrinsic cardiac identity information is preserved, which implies that an attacker possessing a target's wearable PPG data has a statistically significant chance of identifying corresponding record in a "de-identified" ECG database, confirming the severity of the cross-dataset linkage risks.

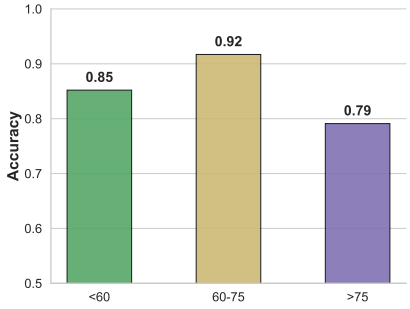
The learned embeddings are visualized using t-SNE, as shown in Figure 5. We randomly selected a subset of subjects from the testing set for clarifying. As observed, the embeddings from distinct clusters corresponding to different identities. Notably, the ECG and PPG embeddings for these

Authorized licensed use limited to: N.C. State University Libraries - Acquisitions & Discovery S. Downloaded on November 07,2026 at 21:06:13 UTC from IEEE Xplore. Restrictions apply.

subjects are tightly mixed within the same identity, indicating that the model successfully aligns the two modalities despite their inherent signal discrepancy. This strong alignment not only demonstrates the effectiveness of the learned representations but also highlights the potential risks of cross-dataset linkage.



(a) Accuracy by Gender



(b) Accuracy by Age

Fig. 6: Demographic bias analysis of the proposed model. (a) Performance between male and female subjects. (b) Performance across different age groups.

C. Bias Analysis

We utilize subjects' demographic information (e.g., sex, age) to analyze potential bias may exist in the proposed model. We focus on positive ECG-PPG pairs for evaluate the model's sensitivity across different groups.

As shown in Figure 6, the model demonstrates high average accuracy for most groups, indicating robust identification capability in recognizing pairs from the same person. Figure 6 (a) shows the performance comparison between the male and female groups. The accuracy is highly consistent between genders, suggesting that the model exhibits negligible gender-based bias. However, as shown in Figure 6 (b), the model achieves optimal performance for age range between 60 and 75, and a noticeable degradation in performance is observed for the elderly group (age > 75), This decline may demonstrate that physiological changes associated with aging, which affects signal quality significantly.

IV. Conclusion

In this paper, we investigated the overlooked privacy risks associated with the proliferation of wearable data. Our experimental results demonstrate that so-called "anonymous" ECG datasets are vulnerable to re-identification attacks via cross-dataset linkage with wearable PPG data. By proposing a Pseudo-Siamese network with SE-ResNet encoders, we successfully aligned the features of these two modalities. Our experiments on the BIDMC dataset demonstrate that the identity information is preserved across modalities. With a matching accuracy nearly 80% and evident identity matching visualization, we validate the feasibility of such cross-dataset linkage attacks. These findings pose a significant security risk and highlight the urgent need for more advanced privacy-preserving mechanisms in the release and sharing of medical research data.

V. Discussion

Our framework demonstrates the feasibility of cross-modal biometric linkage, achieving an average AUC of nearly 0.80. While this does not equate to near-perfect accuracy, it is sufficient to enable practical re-identification when combined with auxiliary information. In realistic threat scenarios, adversaries do not search blindly through millions of uniform records. Instead, they use quasi-identifiers—such as approximate age inferred from PPG morphology, gender, or admission dates—to narrow the candidate pool. Following this initial screening, our cross-modal matcher acts as a high-recall filter, reducing the candidates to a concise shortlist. This makes manual correlation highly feasible and poses a substantive privacy risk.

This vulnerability highlights a structural blind spot in current data-anonymization policies. Regulatory frameworks like HIPAA Safe Harbor implicitly assume that ECGs stripped of metadata are anonymous [20]. However, our findings invalidate this assumption. We demonstrate that raw physiological waveforms carry inherent identity fingerprints, which can be re-linked using data from commodity wearables. Closing this gap requires updated de-identification standards that explicitly address biometric leakage through signal-level transformations, controlled access models, or systematic auditing.

While our findings establish a clear proof-of-concept for cross-modal identity leakage, the exact scope of this vulnerability requires further contextualization. The current evaluation successfully demonstrates biometric linkage within single recording sessions, effectively isolating the fundamental morphological correlation between PPG and ECG. However, real-world wearable data is longitudinal and inherently subject to temporal drift and physiological variations induced by factors such as exercise, stress, or medication [21]. Consequently, investigating the long-term temporal robustness of these cross-modal signatures remains an important direction for future research. Furthermore, expanding this framework to larger, demographically diverse multi-session cohorts will be essential to precisely characterize the attack's effective false-positive rate in the wild, thereby mapping the full landscape of this population-level privacy threat.

Authorized licensed use limited to: N.C. State University Libraries - Acquisitions & Discovery S. Downloaded on November 07,2026 at 21:06:13 UTC from IEEE Xplore. Restrictions apply.

REFERENCES

[1] A. Ometov, V. Shubina, L. Klus, J. Skibińska, S. Saafi, P. Pascacio, L. Flueratoru, D. Q. Gaibor, N. Chukhno, O. Chukhno, A. Ali, A. Channa, E. Svertoka, W. B. Qaim, R. Casanova-Marqués, S. Holcer, J. Torres-Sospedra, S. Casteleyn, G. Ruggeri, G. Araniti, R. Burget, J. Hosek, and E. S. Lohan, "A survey on wearable technology: History, state-of-the-art and current challenges," *Computer Networks*, vol. 193, p. 108074, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128621001651

[2] P. H. Charlton, P. A. Kyriacou, J. Mant, V. Marozas, P. Chowienczyk, and J. Alastruey, "Wearable photoplethysmography for cardiovascular monitoring," *Proceedings of the IEEE*, vol. 110, no. 3, pp. 355–381, 2022.

[3] A. E. Johnson, T. J. Pollard, L. Shen, and et al., "Mimic-iii, a freely accessible critical care database," *Scientific Data*, vol. 3, no. 1, pp. 1–9, 2016.

[4] P. Wagner, N. Strodthoff, R. Bousseljot, and et al., "Ptb-xl, a large publicly available electrocardiography dataset," *Scientific Data*, vol. 7, no. 1, p. 154, 2020.

[5] S. M. Narayan, N. Kohli, and M. M. Martin, "Addressing contemporary threats in anonymised healthcare data using privacy engineering," *npj Digital Medicine*, vol. 8, no. 1, p. 145, 2025.

[6] K. Macierzanka, A. Sau, K. Patlatzoglou, L. Pastika, E. Sieliwonczyk, M. Gurnani, N. S. Peters, J. W. Waks, D. B. Kramer, and F. S. Ng, "Siamese neural network-enhanced electrocardiography can re-identify anonymized healthcare data," *European Heart Journal-Digital Health*, vol. 6, no. 3, pp. 417–426, 2025.

[7] E. Zeng, X. Wu, E. N. Ertmann, L. Huang, D. F. Johnson, A. T. Mehendale, B. T. Tang, K. Zhukoff, M. Adjei-Poku, L. Bauer *et al.*, "Measuring risks to users' health privacy posed by third-party web tracking and targeted advertising," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–26.

[8] A. Jaya Prakash, K. K. Patro, M. Hammad, R. Tadeusiewicz, and P. Pławiak, "Baed: A secured biometric authentication system using ecg signal based on deep learning techniques," *Biocybernetics and Biomedical Engineering*, vol. 42, no. 4, pp. 1081–1093, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0208521622000778

[9] R. Wei, X. Xu, Y. Li, Y. Zhang, J. Wang, and H. Chen, "Pulseid: Multi-scale photoplethysmographic identification using a deep convolutional neural network," *Biomedical Signal Processing and Control*, vol. 88, p. 105609, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S174680942301042X

[10] Y. Ye, G. Xiong, Z. Wan, T. Pan, and Z. Huang, "Ppg-based biometric identification: Discovering and identifying a new user," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, 2021, pp. 1145–1148.

[11] C. Zhong, Z. Guo, A. K. Gehi, C. Xu, and H. Li, "Wearable ppg-to-multi-lead ecg conversion for cardiac monitoring," in *2025 IEEE 21st International Conference on Body Sensor Networks (BSN)*. IEEE, 2025, pp. 1–4.

[12] H. Ning, X. Zheng, X. Lu, and Y. Yuan, "Disentangled representation learning for cross-modal biometric matching," *IEEE Transactions on Multimedia*, vol. 24, pp. 1763–1774, 2022.

[13] A. Nagrani, S. Albanie, and A. Zisserman, "Seeing voices and hearing faces: Cross-modal biometric matching," 04 2018.

[14] R. Kumari, S. Sarkar, D. Dutta, P. Das, and A. Acharyya, "P2e-lgan: Ppg to ecg reconstruction methodology using lstm based generative adversarial network," in *2024 IEEE international symposium on circuits and systems (ISCAS)*. IEEE, 2024, pp. 1–5.

[15] D. Makowski, T. Pham, Z. J. Lau, J. Brammer, F. Lespinasse, C. Schölzel, S. Chen, and H. Pham, "Neurokit2: A python toolbox for neurophysiological signal processing," 07 2020.

[16] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1, 2005, pp. 539–546 vol. 1.

[17] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7132–7141.

[18] M. A. F. Pimentel, A. E. W. Johnson, P. H. Charlton, D. Birrenkott, P. J. Watkinson, L. Tarassenko, and D. A. Clifton, "Toward a robust estimation of respiratory rate from pulse oximeters," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 8, pp. 1914–1923, 2017.

[19] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "Physiobank, physiotoolkit, and physionet," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000. [Online]. Available: https://www.ahajournals.org/doi/abs/10.1161/01.CIR.101.23.e215

[20] U.S. Department of Health and Human Services, Office for Civil Rights, "Guidance regarding methods for de-identification of protected health information in accordance with the health insurance portability and accountability act (hipaa) privacy rule," 2012, https://www.hhs.gov/hipaa/for-professionals/special-topics/de-identification/index.html.

[21] L. Li, C. Chen, L. Pan, L. Y. Zhang, Z. Wang, J. Zhang, and Y. Xiang, "A survey of ppg's application in authentication," *Computers & Security*, vol. 135, p. 103488, 2023.

Authorized licensed use limited to: N.C. State University Libraries - Acquisitions & Discovery S. Downloaded on November 07,2026 at 21:06:13 UTC from IEEE Xplore. Restrictions apply.