

Goran Mladenović<sup>1</sup>  
Aleksandar Stanković  
Arso Vukićević

**Article info:**  
Received 15.01.2026.  
Accepted 01.07.2026.

DOI – 10.24874/IJQR20.03-01



## PRIVACY BY DESIGN IN WEARABLE BIOMETRIC SYSTEMS: QUANTIFYING SIGNAL QUALITY AND PRIVACY RISK TRADE-OFFS

**Abstract:** Wearable biometric systems increasingly support cognitive ergonomics and human-system interaction, forcing design-time trade-offs between physiological signal quality and data-protection requirements. This study presents a system-level framework for jointly assessing signal fidelity and privacy risk in multimodal wearable biometric systems. Using a public affective-physiology dataset of 32 participants, signals were processed into a signal-quality index that captures noise robustness, spectral complexity, artifact prevalence, and temporal stability. Privacy risk was quantified as the worst case across complementary attack vectors, namely re-identification, membership inference, and information-theoretic leakage, so that a single high-risk channel rather than an averaged score governs the estimate.

**Keywords:** Wearable biometrics, Physiological data, Privacy risk, GDPR, Privacy-by-design

### 1. Introduction

Wearable biometric devices are increasingly deployed to support cognitive ergonomics, mental workload assessment, and adaptive human-machine interaction across healthcare, occupational, fitness, and consumer settings, within a rapidly growing global wearable devices market (Grand View Research, 2025). Advances in sensor miniaturization and signal processing enable continuous acquisition of rich physiological data, including electroencephalography (EEG), electro dermal activity (EDA), blood volume pulse (BVP), and respiration. Such systems are typically designed to maximize signal fidelity in order to support reliable inference and decision-making at the system level. At the same time, physiological signals constitute highly sensitive personal data under the EU General Data Protection Regulation (GDPR) and comparable regimes

such as the U.S. HIPAA.

Beyond their raw measurements, they encode stable individual characteristics and can enable re-identification or sensitive inference even without explicit identifiers (Yang & Deravi, 2017; Chikwetu et al., 2023; Lien & Vhaduri, 2023; Blasco et al., 2016). Compliance is therefore not solely a matter of access control or encryption, but a structural property of the sensing and data-processing pipeline itself. Despite this, signal quality and privacy protection have largely been addressed as separate objectives (Ziller et al., 2024; Garcia et al., 2025).

Work in wearable sensing and affective computing focuses on signal-to-noise ratio, artifact robustness, and model performance, treating privacy as an external constraint (Hu et al., 2019; Estreito et al., 2024), whereas privacy-preserving approaches evaluate isolated attack models or anonymization

---

<sup>1</sup> Corresponding author: Goran Mladenović  
Email: [goran.mladenovic@fink.rs](mailto:goran.mladenovic@fink.rs)

techniques without considering their interaction with signal fidelity and utility (Wang et al., 2017; Zhang et al., 2022; Bozorgpanah & Torra, 2024). At the same time, GDPR requires *ex ante*, risk-based assessment, demonstrating that residual risks are proportionate to purpose (Sim et al., 2023; Mone & Shakhlo, 2023).

This separation leaves designers without a quantitative way to reason about privacy-utility trade-offs at design time; in particular, no established, system-level metric reveals whether improvements in signal quality implicitly amplify privacy exposure under GDPR constraint.

A further complication is that wearable systems also enable the inference of cognitive states, introducing risks beyond the protection of neural signals alone (Zhang et al., 2018; Magee et al., 2024). Ienca and Malgieri (Ienca & Malgieri, 2022) argue that mental data should be regarded as personal data whenever combined with other information that allows an individual to be singled out, noting that neural data, such as EEG and fMRI, may, in certain contexts, suffice on their own to identify a person. Building on this, recent scholarship emphasizes regulating mental data and cognitive biometrics, that is, data and inferences capable of revealing mental states, rather than specific sensing modalities (Hallinan et al., 2023; López-Silva et al., 2024; Chandler, 2025; Ardecani et al., 2025). This reinforces the need to evaluate the privacy risks of inferred information, not solely at the signal acquisition stage.

This study addresses that gap with a system-level framework for jointly assessing physiological signal quality and privacy risk, supporting design-time reasoning that ensures analytical performance does not mask unacceptable privacy exposure. Using the DEAP dataset (Koelstra et al., 2011), its contributions are: (i) a multidimensional Signal Quality Index (SQI) characterizing recording fidelity; (ii) a worst-case, multi-vector Privacy Risk Index (PRI) spanning re-identification, membership inference, and

information-theoretical leakage; (iii) their non-compensatory integration into a Privacy-Quality Trade-off Index (PQTI) that prevents high signal quality from masking high privacy risk; and (iv) a counterfactual, data minimization analysis supporting GDPR-aligned design decisions.

The remainder of the paper is organized as follows. Section 2 reviews related work; Section 3 presents the framework and methodology, including the construction of SQI, PRI, and their integration into PQTI; Section 4 reports results on the DEAP dataset and counterfactual analysis; and Section 5 discusses the implications and future work.

## **2. Related works**

Advances in wearable biometric and cognitive sensing are reshaping the privacy-utility landscape of human-system interaction. Modern biometric pipelines no longer serve only identification or authentication; they increasingly enable inference of latent cognitive and affective states, introducing privacy risks tightly coupled to signal-processing and feature-extraction choices (Khan et al., 2020; Hanisch et al., 2021; Prakasha & Sumalatha, 2025). Privacy risk is thus not an external compliance constraint but an intrinsic property of biometric representations. The literature, however, has developed along three largely separate strands: empirical work optimizing signal utility, legal and behavioral work addressing privacy, and methodological work on composite indicators. The present study bridges them.

### **2.1. Inference capacity and the utility focus of prior work**

Physiological signals used in wearables, such as EEG, ECG, and eye-tracking, exhibit high semantic density, allowing models to infer attributes beyond the original sensing objective (Parveen, 2023; Awan et al., 2024; Gahlan & Sethia, 2026). A large body of affective-computing work has accordingly

optimized such signals for utility: on the benchmark DEAP dataset (Koelstra et al., 2011), studies report strong emotion recognition using multi-domain EEG features and classical classifiers (Khateeb et al., 2021), hybrid deep models with explainability (Mouazen et al., 2025), and channel-reduced transformer architectures (Hoffsommer et al., 2025). The same representations can predict health conditions, psychological traits, and behavioral tendencies even when designed for unrelated tasks (Banerjee et al., 2018; Chikwetu et al., 2023), and recent work decodes pre-behavioral and metacognitive states such as confidence and self-monitoring (Bradley et al., 2008; Ohme et al., 2011; Fleur et al., 2021; Polychroni et al., 2022). Critically, none of these utility-focused studies model the privacy exposure these representations create: features optimized for fidelity may simultaneously maximize unintended inference capacity.

## **2.2. Limits of de-identification**

De-identification is structurally inadequate as a safeguard for biometric data. Re-identification accuracy is consistently high, often 86% to 100%, even after explicit identifiers are removed (Chikwetu et al., 2023), and identity can be recovered from seconds- to minutes-long recordings (Zhang et al., 2019). Unlike conventional identifiers, biometric features are persistent and non-revocable: once compromised, they cannot be reset, producing long-term, cumulative exposure (Aguilar et al., 2025; Senthuran et al., 2025; Zhang, 2025). Privacy leakage is therefore not an implementation flaw but a property of biometric feature spaces, motivating quantitative modeling at design time.

## **2.3. Continuous sensing and the limits of consent**

These risks are amplified in IoT-enabled, multimodal ecosystems, where passively

collected signals are fused across devices into persistent "digital shadows" that are externally managed and hard for users to audit (Sarma et al., 2008; Wachter, 2018; Floridi, 2011; Behera et al., 2021; Athreya et al., 2023). Consistent offers limited protection: privacy policies routinely exceed users' reading levels (Banerjee et al., 2018; Van Ooijen & Vrabec, 2019; Zhang et al., 2024), leading to "privacy resignation" and "privacy cynicism," in which data collection is accepted for lack of alternatives. Consent alone, therefore, cannot offset technical privacy risks.

## **2.4. Regulatory gaps and the case for joint, non-compensatory evaluation**

Regulation leaves further gaps. In the U.S., HIPAA covers only designated entities, excluding consumer-grade wearables handled by manufacturers or app providers (Theodos & Sittig, 2020); in the EU, the GDPR's heightened protections remain ambiguous for biometric-derived mental and affective data not tied to health or identification (Ienca & Malgieri, 2022; Solove, 2023), and AI-assisted compliance tooling targets governance rather than signal-level privacy (Hasan & Faruq, 2025). Compliance alone thus does not guarantee privacy; privacy risk must be treated as a system-level performance dimension, jointly optimized with signal quality and utility. Methodologically, composite-indicator theory holds that aggregation should be non-compensatory, so that strength on one dimension cannot mask weakness on another (Mazziotta & Pareto, 2013). The closest prior work, a privacy-utility score for single-lead ECG computed as a geometric mean (Senthuran et al., 2025), demonstrates the value of joint scoring but is limited to a single modality and does not adopt a worst-case formulation. The present study extends this by combining a multidimensional signal-quality index with a worst-case, per-subject privacy risk index for multimodal wearable biometrics.

### 3. Methodology

#### 3.1. Overview of the evaluation framework

The framework integrates signal quality assessment and privacy-risk evaluation into a unified composite indicator to quantify GDPR-aligned, design-time decision-making in a wearable biometrics system. It follows the established principles of composite indicator construction as formalized by the OECD (2008) and combines physiological signal processing with explicit privacy risk assessment.

Privacy risk is quantified through an explicitly modeled adversary within the Privacy Risk Index (PRI), which simulates

realistic worst-case attacks on feature representations and model outputs, including re-identification, membership inference, and information-theoretic leakage. By embedding the adversary as an analytical component, privacy exposure is treated as an observable, measurable system property rather than an abstract or assumed threat. SQI and PRI are constructed independently and then combined into a PQTI through non-compensatory geometric aggregation of SQI and the privacy margin (1-PRI), so that the evaluated privacy risk cannot be offset by high analytical performance, consistent with the preventive logic of GDPR's risk-based approach. Figure 1 illustrates the evaluation architecture supporting design decision-making in wearable biometric systems.




**Figure 1.** Evaluation Architecture for Design Decision Support in Wearable Biometric Systems

#### 3.2. Dataset and experimental scenarios

Physiological data were obtained from the DEAP dataset, which contains recordings from 32 participants exposed to 40 emotional video stimuli.

The dataset provides synchronized multichannel Electroencephalogram (EEG)

together with peripheral bio-signals, including electrodermal activity (EDA/GSR), blood volume pulse (BVP, plethysmograph), and respiration. For each participant and trial, four groups of features were extracted:

- 1) EEG signal-quality metrics, including band-limited signal-to-noise ratio, spectral entropy, artifact

rate, and temporal stability; (2) EDA descriptors (statistical and spectral); (3) simple statistical descriptors of BVP and respiration for downstream privacy analysis; and

- 2) (4) subject self-ratings of valence, arousal, dominance, and liking.

The DEAP dataset serves as a controlled reference environment for evaluating the framework under different data-governance configurations that emulate realistic wearable deployment strategies.

### 3.3. Signal Quality Index (SQI)

The SQI quantifies physiological recording fidelity by integrating indicators that capture noise robustness, signal structure, artifact contamination, and temporal stability. For each participant and trial, signal-to-noise ratio is estimated from a band-pass decomposition of the signal (the in-band component relative to the residual); normalized spectral entropy characterizes signal complexity and information content; artifact prevalence is quantified as the proportion of samples exceeding a robust amplitude threshold; and temporal stability is assessed as the correlation between the first and second halves of each trial.

Each indicator is normalized to  $[0,1]$  using a distance-to-reference transformation relative to empirically defined benchmarks (the 5th and 95th percentiles of the observed distribution), with the direction set so that higher values denote better quality. EEG and EDA sub-indices are first formed by geometrically aggregating their respective indicators, and then combined geometrically into a trial-level SQI; the subject-level SQI is obtained by averaging across trials. Geometric aggregation prevents full compensation between degraded and high-quality dimensions, yielding a compact yet interpretable, sample-relative measure of signal fidelity.

### 3.4. Privacy Risk Index (PRI)

Privacy exposure is quantified through a PRI constructed from complementary attack vectors commonly considered in biometric privacy assessment. The PRI adopts a worst-case-oriented perspective that aligns with GDPR's preventive and harm-avoidance principles, particularly articulated in Articles 25 and 35 (European Parliament, 2016). Biometric re-identification risk is evaluated using a high-capacity Random Forest classifier (Breiman, 2001; Athey et al., 2019; El Haouij et al., 2019), with a Support Vector Machine reported as a weaker baseline (Quitadamo et al., 2017; Zhang et al., 2014; Domínguez-Jiménez et al., 2020), both trained to predict subject identity from extracted physiological features over repeated stratified splits. Membership inference risks are assessed by training a lightweight attacker on the target model's confidence (maximum posterior probability) and evaluating it on a held-out split of members and non-members, avoiding optimistic in-sample estimates. Information-theoretic leakage is estimated as mutual information between feature and subject identity, normalized by the identity entropy,  $\ln K$  (the natural logarithm of the number of subjects  $K$ ; here  $K = 32$ , so  $\ln K \approx 3.47$ ). Each metric lies on a common  $[0,1]$  scale, and the PRI is set to the worst-case (maximum) across vectors, so that a single high-risk channel governs the conservative estimate, and a weak attacker cannot lower it. For comparison, a compensatory geometric aggregation of the same vectors is also reported.

### 3.5. Privacy-Quality Trade-off Index (PQTI) and robustness analysis

The PQTI integrates SQI and PRI to summarize the balance between analytical utility and privacy protection under different design configurations. It is defined as a weighted geometric mean of SQI and the privacy margin  $(1 - \text{PRI})$ ,

$$PQTI = (SQI)^{w_q} \cdot (1 - PRI)^{w_p},$$

with equal baseline weights  $w_q = w_p = 0.5$ , reflecting their comparable importance.

This formulation ensures that PQTI increases only when signal quality is preserved while privacy risk is reduced. To assess robustness, an alternative (compensatory geometric) aggregation is compared with the worst-case PRI, and sensitivity to the weighting is evaluated through Monte Carlo analysis of 10,000 weight draws (Shapiro, 2003; Hammersley, 2013). This provides insight into the stability of PQTI rankings under methodological variation and supports transparent, reproducible evaluation of GDPR-aligned design decisions in wearable biometric systems.

## 4. Results

### 4.1. Processing context and output artifacts

All results were produced in a reproducible Python workflow built on NumPy, SciPy, and Scikit-learn. The pipeline produced (i) modality-specific signal-quality features (EEG, EDA), (ii) trial-level and subject-level Signal Quality Index summaries, and (iii) a privacy-feature matrix for evaluating identity leakage and membership inference. These outputs were then integrated into composite indices quantifying privacy exposure (PRI) and the final privacy-quality trade-off (PQTI). All randomized procedures used fixed seeds, and re-identification and membership inference were evaluated over 20 repeated stratified splits to provide mean estimates with confidence intervals.

### 4.2. Signal quality outcomes and modality contrast

Signal fidelity was quantified by formatting modality-specific sub-indices and integrating them into a trial-level SQI, following the normalization and aggregation (OECD, 2008;

Huang et al., 2009; Xu & Yager, 2009). Across 1,280 trials, the trial-level SQI showed a broad distribution (mean  $\approx 0.352$ , std  $\approx 0.196$ ), spanning the full range from near-zero to high-quality segment and indicating substantial variability in recording reliability across trials and subjects. No single modality collapses the composite, while genuinely degraded segments are still penalized. The variability is practically important for multimodal wearable design: it identifies low-fidelity acquisition conditions that should be mitigated at acquisition time rather than masked by aggregation.

### 4.3. Privacy leakage under re-identification and membership inference

Privacy exposure was evaluated using complementary attack modalities targeting (i) identity re-identification and (ii) training-set membership inference, together with an information-theoretic analysis of identity dependence in the feature space.

**Re-identification (capacity sensitivity).** Subject re-identification showed a strong dependence on attacker capacity. A high-capacity Random Forest attacker achieved a mean accuracy of 0.927 (SD 0.012; 95% CI 0.905-0.945) over 20 repeated stratified splits, indicating that subject identity is highly recoverable by physiological features. In contrast, an SVM baseline achieved 0.289, showing that a weaker attacker extracts substantially less identity information given the same representation. The practical implication is that privacy evaluation must be anchored to realistic (high-capacity) adversaries; otherwise, risk can be substantially underestimated.

**Membership inference.** Membership inference attack (Shokri et al., 2017; Hu et al., 2022), a lightweight logistic-regression attacker (Rahman et al., 2023; Bikmukhamedov & Nadeev, 2019; Hosmer et al., 2013) trained on the one-dimensional maximum posterior probability and evaluated on a held-out split of members and non-

members, achieved AUC = 0.740 (95% CI 0.714-0.770), demonstrating that inclusion of a sample in the training set can be inferred above chance. This reveals a second privacy channel beyond identity classification: even if identity prediction is not the explicit task,

model behavior can leak whether an individual contributed data to training. Table 1 reports these baseline attack results, together with the data-minimization scenarios introduced in Section 3. 2.

**Table 1.** Privacy-attack results on the baseline feature representation and under data-minimization scenarios.

| Scenario                       | Features (n) | RF reidentification | Membership AUC | Normalized MI | PRI (worstcase) |
|--------------------------------|--------------|---------------------|----------------|---------------|-----------------|
| Baseline (full feature set)    | 19           | <b>0.927</b>        | 0.740          | 0.295         | <b>0.927</b>    |
| Data minimisation (EEG only)   | 10           | <b>0.823</b>        | 0.720          | 0.381         | <b>0.823</b>    |
| Anonymized labels (no emotion) | 17           | <b>0.923</b>        | 0.724          | 0.318         | <b>0.923</b>    |
| EEG alpha-band only            | 4            | 0.533               | <b>0.871</b>   | 0.257         | <b>0.871</b>    |

#### 4.4. Information-theoretic identity dependence in feature

To quantify identity leakage independently of any classifier, mutual information (MI) between each feature and subject identity was computed (Guo & Nixon, 2008). The mean normalized MI across features was 0.295, indicating that subject identity is substantially embedded in the extracted representation. However, leakage was broadly distributed rather than concentrated in any single descriptor: values ranged narrowly across the most informative features (maximum 0.414 for alpha-band power variability, with the top nine features all within  $\approx 0.04$  of one another), so no individual feature uniquely identifies a subject. The leading features were EEG band-power descriptors spanning the delta, theta, alpha, and beta bands (both means and variabilities), together with global EEG variance. In contrast, peripheral descriptors such as BVP variability followed closely. This supports a key design-relevant conclusion: identity leakage is not confined to a narrow subset of "biometric" features but is broadly distributed across physiologically meaningful descriptors that are also useful for cognitive/affective inference.

#### 4.5. Composite Privacy Risk Index (PRI)

The Privacy Risk Index (PRI) aggregates

complementary leakage signals into a single conservative score using worst-case (maximum) aggregation across attack vectors. In this implementation, the vectors are re-identification accuracy, held-out membership inference AUC, and normalized mutual-information leakage, all reported in Table 1; a weaker SVM baseline is included for context but, by design, does not lower the worst-case estimate. The resulting composite PRI is 0.927, indicating a structurally high privacy-risk profile under the baseline feature representation.

For comparison, a compensatory geometric aggregation of the same vectors yields only 0.49. Because the value is sensitive to both the inclusion of weak attackers and the normalization of mutual information, the worst-case formulation is preferred as a conservative, GDPR-consistent summary.

A PRI of approximately 0.93 indicates that a single vulnerability does not drive privacy risk. Although a weaker model, such as SVM, shows limited re-identification performance, a high-capacity classifier, information-theoretic dependence, and membership inference each independently reveal substantial leakage. Privacy exposure is therefore structural in nature, arising from several channels rather than from one isolated weakness; a worst-case rather than an averaging aggregation is the appropriate summary.

#### 4.6. PQTI baseline ranking and trade-off behavior

Applying the PQTI (Section 3.5) with equal weights, privacy risk is estimated at the subject level from re-identification recall rather than treated as a single global constant. It ranges from approximately 0.80 to 1.00 across subjects. Consequently, signal quality and privacy exposure jointly determine the ordering, and the subjects with the cleanest signals do not necessarily rank highest if they are also the most identifiable.

In the subject-level summary, s21 achieved the highest PQTI (0.375), followed by s16 (0.373), s27 (0.330), s18 (0.329), and s08 (0.320). This pattern reflects the intended behavior of the index: the highest-ranked subjects combine adequate signal quality with comparatively lower re-identifications exposure, whereas highly identifiable subjects are penalized even when their signal quality is high. For example, s27 has the highest normalized SQI but ranks only third because of its elevated re-identifications risk.

#### 4.7. Counterfactual "what-if" scenario shifts

To assess the interpretability and design sensitivity of the PQTI, counterfactual scenarios (Wachter et al., 2017) were analyzed in which either the relative weighting of privacy or the assumed level of

privacy risk was systematically modified.

Three configurations were considered:

- a) a baseline with equal weighting of signal quality and privacy ( $w_q = w_p = 0.5$ ),
- b) a privacy-weighted scenario emphasizing regulatory strictness ( $w_q = 0.3, w_p = 0.7$ ), and
- c) a better-privacy counterfactual in which the PRI was uniformly reduced by 20% while signal quality remained unchanged.

The results are summarized in Table 2. Increasing the weight assigned to privacy produces a systematic downward shift of PQTI across all subjects, whereas reducing privacy risk shifts all values upward, reflecting the non-compensatory geometric aggregation, in which emphasizing the privacy term penalizes even high-quality signals. For the highest-ranked subject, PQTI fell from 0.375 to 0.285 (s21) and from 0.373 to 0.282 (s16) under the privacy-weighted configuration, and rose to 0.466 and 0.460, respectively, under the -20% risk counterfactual. Critically, although absolute values change, the relative ordering is preserved (Spearman  $\approx 0.99$  between the baseline and privacy-weighted rankings). This confirms that PQTI responds monotonically and predictably to changes in regulatory strictness without introducing ranking instability.

**Table 2.** PQTI under counterfactual weighting and privacy-risk scenarios (top-ranked subjects).

| Subject                                        | SQI_norm | PRI   | Baseline (0.5/0.5) | Privacy-weighted (0.3/0.7) | Better-privacy (PRI - 20%) |
|------------------------------------------------|----------|-------|--------------------|----------------------------|----------------------------|
| <b>s21</b>                                     | 0.703    | 0.800 | 0.375              | 0.285                      | 0.466                      |
| <b>s16</b>                                     | 0.778    | 0.821 | 0.373              | 0.282                      | 0.460                      |
| <b>s27</b>                                     | 0.966    | 0.888 | 0.330              | 0.231                      | 0.428                      |
| <b>s18</b>                                     | 0.720    | 0.850 | 0.329              | 0.243                      | 0.418                      |
| <b>s08</b>                                     | 0.632    | 0.838 | 0.320              | 0.241                      | 0.402                      |
| <b>Rank correlation vs baseline (Spearman)</b> |          |       | <b>1.000</b>       | <b>0.990</b>               | <b>1.000</b>               |

## 5. Results

The results demonstrate that privacy leakage in physiological data is a structural property of the feature space rather than an incidental artifact of modeling choices or individual subjects. A high-capacity re-identification attack recovered subject identity with high accuracy, and the information-theoretic analysis showed that identity information is broadly embedded across features rather than concentrated in any single descriptor (Section 4).

Commonly used physiological representations, particularly EEG band-power and variability features, therefore retain distinctive subject-specific structure even after standard preprocessing and nominal anonymization, the same representations that are central to cognitive and affective analysis.

The pronounced gap between the high-capacity and the weaker attacker highlights the capacity-sensitive nature of biometric privacy risk. While lower-capacity models may appear to offer some protection, stronger nonlinear classifiers reliably exploit higher-order feature interactions to recover identity. This has direct implications for GDPR-oriented assessment: evaluations anchored to a single weak attacker risk substantially underestimating exposure. The PRI addresses this by taking the worst case across attack vectors, so that a weak baseline cannot mask the exposure revealed by a stronger one.

Beyond direct identity inference, the membership inference results reveal an additional privacy channel: even when explicit re-identification is not the objective, model-confidence signals leak information about dataset participation above chance. This is regulatorily salient because GDPR concerns extend to re-identification potential and secondary inferences, not only explicit identity prediction. The convergence of re-identification, membership inference, and information-theoretic leakage provides triangulated evidence that post hoc

anonymization or superficial feature masking is insufficient for protecting physiological data.

Because privacy risk is estimated at the subject level, the Privacy Quality Trade-off Index reflects a genuine interaction between signal quality and identifiability rather than a quality-only ranking. Signal-quality variation arises from electrode placement, motion artifacts, and physiological idiosyncrasies, while per-subject privacy exposure modulates the achievable trade-off. Consequently, the cleanest signal does not guarantee the highest PQTI: a subject with excellent signal quality but high identifiability is penalized relative to one with adequate quality and lower exposure (Section 4). Improvements in signal fidelity alone, therefore, cannot compensate for high structural privacy leakage.

This reinforces the central argument that privacy mechanisms relying solely on downstream transformations are insufficient. The consistently high privacy risk indicates that leakage originates in the core structure of the physiological signals themselves. Effective mitigation, therefore, requires structural data minimization, reducing the number of collected channels, excluding physiologically revealing modalities, or restricting feature representations to task-essential components. The scenario analysis supports this concretely: removing peripheral channels lowered the worst-case privacy risk, whereas removing the affective labels alone left it essentially unchanged (Section 4), indicating that exposure is carried by the signal representation rather than by the labels. These strategies align directly with the GDPR principles of purpose limitation and data minimization by reducing exposure at the point of collection rather than obscuring it afterward.

The counterfactual analyses further validate PQTI as a design-time diagnostic. Increasing the weight assigned to privacy results in a predictable reduction in PQTI, while hypothetical reductions in privacy risk yield systematic increases; in both cases, the

relative ordering of subjects is preserved. Because privacy is modeled at the subject level, these interventions differentiate how individual subjects are affected rather than applying a single, uniform penalty, addressing a limitation of system-level privacy scoring and pointing toward further refinement through attacker-on configuration-specific risk estimates.

From an engineering perspective, the non-compensatory geometric aggregation in PQTI is particularly important: excellence in signal quality cannot mask deficiencies in privacy protection, or vice versa. This aligns with both safety-critical engineering principles and the preventive logic of GDPR. The stability of PQTI rankings across weighting and counterfactual scenarios supports its suitability for engineering decision-making and regulatory assessment, allowing designers to explore alternative sensing configurations and governance priorities without encountering unstable or contradictory evaluations.

More broadly, embedding counterfactual reasoning in PQTI bridges abstract regulatory principles and measurable system properties. By simulating the consequences of hypothetical design choices, the framework operationalizes privacy-by-design and data minimization explicitly and predictively. Rather than treating compliance as a static checklist, PQTI enables scenario-based planning, allowing ethical, regulatory, and functional trade-offs to be evaluated before deployment. This is especially relevant for wearable neuroeconomic and cognitive-ergonomics systems, where strong physiological signals offer high analytical utility but pose substantial, persistent privacy risks.

## **6. Conclusion**

This study introduced the Privacy–Quality Trade-off Index (PQTI) as a GDPR-aligned, system-level framework for quantifying the inherent tension between analytical utility

and privacy risk in wearable biometric systems. By integrating signal-quality metrics with a conservative, multi-vector assessment of privacy exposure, PQTI provides a principled mechanism for evaluating biometric system configurations under realistic adversarial assumptions. The results show that privacy leakage in physiological data is a structural property of commonly used feature representations rather than an incidental byproduct of modeling choices, with strong identity inference persisting even under standard preprocessing.

A key empirical finding is that structural data minimization, implemented through modality selection and channel reduction, is substantially more effective than post hoc anonymization or noise-based transformations.

Improvements in signal quality increase analytical utility but cannot, on their own, offset high structural privacy risk; conversely, a reduction in privacy exposure yields predictable improvements in PQTI while preserving subject rankings, underscoring the value of privacy-by-design strategies that limit data collection at the source. These results position PQTI as a practical diagnostic for design-time decision-making in cognitive ergonomics and wearable IoT systems, enabling engineers to balance performance objectives with regulatory constraints in a transparent, reproducible manner.

The framework also proved robust across weighting schemes and counterfactual scenarios, supporting its use in regulatory assessment and in the comparative evaluation of alternative system designs.

By operationalizing GDPR principles, such as data minimization and privacy-by-design, into a quantitative index, PQTI helps bridge the gap between abstract legal requirements and measurable engineering properties.

Several limitations qualify these findings and motivate future work. The analysis relies on a single controlled dataset (DEAP) recorded with research-grade equipment; generalization to commercial and consumer

wearables, longitudinal recordings, and real-time adaptive systems remains to be established. Future research will extend PQTI in these directions and will refine subject and configuration-specific privacy-risk estimation, including a wider range of adversaries, to enable more granular

counterfactual analysis.

All code, data processing pipelines, and supplementary materials will be released to support reproducibility, benchmarking, and broader community adoption.

## References:

- Aathreya, S., Chaudhary, M., Neal, T., & Canavan, S. (2023). Multimodal context-based continuous authentication. In *2023 IEEE International Joint Conference on Biometrics (IJCB)*, 1-10. <https://doi.org/10.1109/IJCB57857.2023.10448626>.
- Aguilar, D., Martínez-Cruz, A., Ramírez-Gutiérrez, K. A., & Morales-Sandoval, M. (2025). PPG-based biometric authentication: A review on architectures, datasets, attacks and security challenges. *Computational and Structural Biotechnology Journal*, 28, 511-528. <https://doi.org/10.1016/j.csbj.2025.11.006>.
- Ardecani, F. B., Kumar, A., Sabeti, S., & Shoghli, O. (2025). Neural correlates of augmented reality safety warnings: EEG analysis of situational awareness and cognitive performance in roadway work zones. *Safety Science*, 185, 106802. <https://doi.org/10.1016/j.ssci.2025.106802>.
- Athey, S., Tibshirani, J., & Wager, S. (2019). Generalized random forests. *Ann. Statist.*, 47(2), 1148 - 1178, <https://doi.org/10.1214/18-AOS1709>.
- Awan, A. W., Taj, I., Khalid, S., Usman, S. M., Imran, A. S., & Akram, M. U. (2024). Advancing emotional health assessments: a hybrid deep learning approach using physiological signals for robust emotion recognition. *IEEE Access*, 12, 141890-141904, <https://doi.org/10.1109/ACCESS.2024.3463746>.
- Banerjee, S., Hemphill, T., & Longstreet, P. (2018). Wearable devices and healthcare: Data sharing and privacy. *The Information Society*, 34(1), 49-57. <https://doi.org/10.1080/01972243.2017.1391912>.
- Behera, N. K. S., Behera, T. K., Nappi, M., Bakshi, S., & Sa, P. K. (2021). Futuristic person re-identification over internet of biometrics things (IoBT): Technical potential versus practical reality. *Pattern Recognition Letters*, 151, 163-171. <https://doi.org/10.1016/j.patrec.2021.08.007>.
- Bikmukhamedov, R. F., & Nadeev, A. F. (2019). Lightweight machine learning classifiers of IoT traffic flows. In *2019 Systems of Signal Synchronization, Generating and Processing in Telecommunications (SYNCHROINFO)*, 1-5. IEEE. <https://doi.org/10.1109/SYNCHROINFO.2019.8814156>.
- Blasco, J., Chen, T. M., Tapiador, J., & Peris-Lopez, P. (2016). A survey of wearable biometric recognition systems. *ACM Computing Surveys (CSUR)*, 49(3), 1-35. <https://doi.org/10.1145/29682>.
- Bozorgpanah, A., & Torra, V. (2024). Explainable machine learning models with privacy. *Progress in Artificial Intelligence*, 13(1), 31-50. <https://doi.org/10.1007/s13748-024-00315-2>.
- Bradley, M. M., Miccoli, L., Escrig, M. A., & Lang, P. J. (2008). The pupil as a measure of emotional arousal and autonomic activation. *Psychophysiology*, 45(4), 602-607. <https://doi.org/10.1111/j.1469-8986.2008.00654.x>

- Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5-32. doi: <https://doi.org/10.1023/A:1010933404324>.
- OECD, Statistics Working Paper. (2008). Handbook on constructing composite indicators: methodology and user guide. *OECD publishing*. <https://dx.doi.org/10.1787/533411815016>.
- Chandler, J. A. (2025). Knowing the Mind from Brain Data: The Challenge of Prediction and the Fairness of Relying on Objective Data about the Mind. *Hastings Center Report*, 55(4), 15-23. <https://doi.org/10.1002/hast.70000>.
- Chikwetu, L., Miao, Y., Woldetensae, M. K., Bell, D., Goldenholz, D. M., & Dunn, J. (2023). Does deidentification of data from wearable devices give us a false sense of security? A systematic review. *The Lancet Digital Health*, 5(4), e239-e247, [https://doi.org/10.1016/S2589-7500\(22\)00234-5](https://doi.org/10.1016/S2589-7500(22)00234-5).
- Domínguez-Jiménez, J. A., Campo-Landines, K. C., Martínez-Santos, J. C., Delahoz, E. J., & Contreras-Ortiz, S. H. (2020). A machine learning model for emotion recognition from physiological signals. *Biomedical signal processing and control*, 55, 101646, <https://doi.org/10.1016/j.bspc.2019.101646>.
- El Haouij, N., Poggi, J. M., Ghozi, R., Sevestre-Ghalila, S., & Jaïdane, M. (2019). Random forest-based approach for physiological functional variable selection for driver's stress level classification. *Statistical Methods & Applications*, 28(1), 157-185. <https://doi.org/10.1007/s10260-018-0423-5>.
- Estreito, Z., Le, V., Harris\_Jr, F. C., & Dascalu, S. M. (2024). Evaluating an Elevated Signal-to-Noise Ratio in EEG Emotion Recognition. *International Journal of Software Innovation (IJSI)*, 12(1), 1-15. <https://doi.org/10.4018/IJSI.333161>.
- Fleur, D. S., Bredeweg, B., & van den Bos, W. (2021). Metacognition: ideas and insights from neuro-and educational sciences. *npj Science of Learning*, 6(1), 13. <https://doi.org/10.1038/s41539-021-00089-5>.
- Floridi, L. (2011). The informational nature of personal identity. *Minds and machines*, 21(4), 549-566. <https://doi.org/10.1007/s11023-011-9259-6>.
- Gahlan, N., & Sethia, D. (2026). Emotion recognition from auditory Autonomous Sensory Meridian Response (ASMR) using multi-modal physiological signals. *Physiological Measurement*, 015008. <https://doi.org/10.1088/1361-6579/ae35ca>.
- Garcia, R. D., Ramachandran, S. G., Rothenberg, C. E., Krishnamachari, B., & Ueyama, J. (2025). A Survey of Privacy-Preserving Mechanisms on Quality of Experience in Next-Generation Networks. *Computer Networks*, 111899. <https://doi.org/10.1016/j.comnet.2025.111899>.
- Guo, B., & Nixon, M. S. (2008). Gait feature subset selection by mutual information. *IEEE Transactions on Systems, MAN, and Cybernetics-part a: Systems and Humans*, 39(1), 36-46, <https://doi.org/10.1109/TSMCA.2008.2007977>.
- Hallinan, D., Boehm, F., Külpmann, A., & Elson, M. (2023). Information provision for informed consent procedures in psychological research under the general data protection regulation: A practical guide. *Advances in methods and practices in psychological science*, 6(1), 25152459231151944. <https://doi.org/10.1177/2515245923115194>.
- Hammersley, J. (2013). Monte carlo methods. *Springer Science & Business Media*. <https://doi.org/10.1063/1.3047186>.
- Hanisch, S., Arias-Cabarcos, P., Parra-Arnau, J., & Strufe, T. (2021). Privacy-protecting techniques for behavioral data: A survey. *arXiv preprint arXiv:2109.04120*, 1-43. <https://doi.org/10.48550/arXiv.2109.04120>.

- Hasan, M., & Faruq, M. O. (2025). AI-Augmented Risk Detection in Cybersecurity Compliance: A GRC-Based Evaluation in Healthcare and Financial Systems. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 1(01), 313-342. <https://doi.org/10.63125/49gs6175>.
- Hoffsommer, A., Schneider, H., Pavlitska, S., & Zöllner, M. (2025). DEAP DIVE: Dataset Investigation with Vision transformers for EEG evaluation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 51-60.
- Hosmer Jr, D. W., Lemeshow, S., & Sturdivant, R. X. (2013). Applied logistic regression. *John Wiley & Sons*. <https://doi.org/10.1002/0471722146>.
- Hu, H., Salcic, Z., Sun, L., Dobbie, G., Yu, P. S., & Zhang, X. (2022). Membership inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)*, 54(11s), 1-37. <https://doi.org/10.1145/3523273>.
- Hu, X., Chen, J., Wang, F., & Zhang, D. (2019). Ten challenges for EEG-based affective computing. *Brain Science Advances*, 5(1), 1-20. <https://doi.org/10.1177/2096595819896200>.
- Huang, Y. S., Liao, J. T., & Lin, Z. L. (2009). A study on aggregation of group decisions. *Systems Research and Behavioral Science: The Official Journal of the International Federation for Systems Research*, 26(4), 445-454. <https://doi.org/10.1002/sres.941>
- Ienca, M., & Malgieri, G. (2022). Mental data protection and the GDPR. *Journal of Law and the Biosciences*, 9(1), 1-19. <https://doi.org/10.1093/jlb/lisac006>.
- Khan, S., Parkinson, S., Grant, L., Liu, N., & McGuire, S. (2020). Biometric systems utilising health data from wearable devices: Applications and future challenges in computer security. *ACM Computing Surveys (CSUR)*, 53(4), 1-29. <https://doi.org/10.1145/3400030>.
- Khateeb, M., Anwar, S. M., & Alnowami, M. (2021). Multi-domain feature fusion for emotion classification using DEAP dataset. *IEEE Access*, 9, 12134-12142. <https://doi.org/10.1109/ACCESS.2021.3051281>.
- Koelstra, S., Muhl, C., Soleymani, M., Lee, J. S., Yazdani, A., Ebrahimi, T., ... & Patras, I. (2011). Deap: A database for emotion analysis; using physiological signals. *IEEE transactions on affective computing*, 3(1), 18-31. <https://doi.org/10.1109/T-AFFC.2011.15>.
- Lien, C. W., & Vhaduri, S. (2023). Challenges and opportunities of biometric user authentication in the age of iot: A survey. *ACM Computing Surveys*, 56(1), 1-37, <https://doi.org/10.1145/3603705>.
- López-Silva, P., Wajnerman-Paz, A., & Molnar-Gabor, F. (2024). Neurotechnological applications and the protection of mental privacy: an assessment of risks. *Neuroethics*, 17(2), 31. <https://doi.org/10.1007/s12152-024-09565-2>.
- Magee, P., Ienca, M., & Farahany, N. (2024). Beyond neural data: Cognitive biometrics and mental privacy. *Neuron*, 112(18), 3017-3028. <https://doi.org/10.1016/j.neuron.2024.09.004>.
- Mazziotta, M., & Pareto, A. (2013). Methods for constructing composite indices: one for all or all for one? *Rivista italiana di economia, demografia e statistica*, 67, 67-80.
- Mone, V., & Shakhlo, F. (2023). Health data on the go: navigating privacy concerns with wearable technologies. *Legal Information Management*, 23(3), 179-188. <https://doi.org/10.1017/S1472669623000427>.
- Mouazen, B., Benali, A., Chebchoub, N. T., Abdelwahed, E. H., & De Marco, G. (2025). Enhancing EEG-based emotion detection with hybrid models: Insights from DEAP dataset applications. *Sensors*, 25(6), 1827. <https://doi.org/10.3390/s25061827>.

- Ohme, R., Matukin, M., & Pacula-Lesniak, B. (2011). Biometric measures for interactive advertising research. *Journal of interactive advertising*, 11(2), 60-72. <https://doi.org/10.1080/15252019.2011.10722185>.
- Parveen, S. H. (2023). Interweaving artificial intelligence and bio-signals in mental fatigue: unveiling dynamics and future pathways. In *2023 20th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*, 1-9. IEEE, <https://doi.org/10.1109/ICCWAMTIP60502.2023.10387132>.
- Polychroni, N., Herrojo Ruiz, M., & Terhune, D. B. (2022). Introspection confidence predicts EEG decoding of self-generated thoughts and meta-awareness. *Human Brain Mapping*, 43(7), 2311-2327. <https://doi.org/10.1002/hbm.25789>.
- Prakasha, K. K., & Sumalatha, U. (2025). Privacy-preserving techniques in biometric systems: Approaches and challenges. *IEEE Access*, 13, 32584-32616. <https://doi.org/10.1109/ACCESS.2025.3541649>.
- Quitadamo, L. R., Cavrini, F., Sbernini, L., Riillo, F., Bianchi, L., Seri, S., & Saggio, G. (2017). Support vector machines to detect physiological patterns for EEG and EMG-based human-computer interaction: a review. *Journal of neural engineering*, 14(1), 011001. <https://doi.org/10.1088/1741-2552/14/1/011001>.
- Rahman, K., Aziz, M. A., Usman, N., Kiren, T., Cheema, T. A., Shoukat, H., ... & Sajid, A. (2023). Cognitive Lightweight Logistic Regression-Based IDS for IoT-Enabled FANET to Detect Cyberattacks. *Mobile Information Systems*, 2023(1), 7690322. <https://doi.org/10.1155/2023/7690322>Digital Object Identifier (DOI).
- Grand View Research (2025). Retrieved from <https://www.grandviewresearch.com/industry-analysis/wearable-technology-market>.
- Sarma, A., Matos, A., Girao, J., & Aguiar, R. L. (2008). Virtual identity framework for telecom infrastructures. *Wireless Personal Communications*, 45(4), 521-543. <https://doi.org/10.1007/s11277-008-9475-4>.
- Senthuran, V., Thayasivam, U., Natgunanathan, I., Sood, K., & Xiang, Y. (2025). Balancing privacy and health integrity: A novel framework for ECG signal analysis in immersive environments. *Computers in Biology and Medicine*, 192, 110234. <https://doi.org/10.1016/j.compbiomed.2025.110234>.
- Shapiro, A. (2003). Monte Carlo sampling methods. *Handbooks in operations research and management science*, 10, 353-425. [https://doi.org/10.1016/S0927-0507\(03\)10006-0](https://doi.org/10.1016/S0927-0507(03)10006-0).
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, 3-18. IEEE, <https://doi.org/10.1109/SP.2017.41>.
- Sim, J., Kim, B., Jeon, K., Joo, M., Lim, J., Lee, J., & Choo, K. K. R. (2023). Technical requirements and approaches in personal data control. *ACM Computing Surveys*, 55(9), 1-30, <https://doi.org/10.1145/3558766>.
- Solove, D. J. (2023). Data is what data does: Regulating based on harm and risk instead of sensitive data. *North Western University Review*, 118, 1081.
- Theodos, K., & Sittig, S. (2020). Health information privacy laws in the digital age: HIPAA doesn't apply. *Perspectives in health information management*, 18(1).PMCID:PMC7883355 PMID: 33633522.

- European Parliament (2016, 04 27). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data and repealing Directive 95/46/EC. Retrieved from General Data Protection Regulation, *Official Journal of the European Union*, L 119, 1–88. Retrieved from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.
- Van Ooijen, I., & Vrabec, H. U. (2019). Does the GDPR enhance consumers' control over personal data? An analysis from a behavioural perspective. *Journal of consumer policy*, 42(1), 91-107. <https://doi.org/10.1007/s10603-018-9399-7>.
- Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *arXiv preprint arXiv:1711.00399*. <https://doi.org/10.48550/arXiv.1711.00399>.
- Wachter, S. (2018). Normative challenges of identification in the Internet of Things: Privacy, profiling, discrimination, and the GDPR. *Computer law & security review*, 34(3), 436-449. <https://doi.org/10.1016/j.clsr.2018.02.002>.
- Wang, S. L., Hsiu, M. J., Tsai, Y. C., Ting, I. H., & Hong, T. P. (2017). Analysis of privacy and utility tradeoffs in anonymized mobile context streams. *Intelligent Data Analysis*, 21(1\_suppl), S21-S39.
- Xu, Z., & Yager, R. R. (2009). Power-geometric operators and their use in group decision making. *IEEE Transactions on fuzzy systems*, 18(1), 94-105. <https://doi.org/10.1109/TFUZZ.2009.2036907>.
- Yang, S., & Deravi, F. (2017). On the usability of electroencephalographic signals for biometric recognition: A survey. *IEEE Transactions on Human-Machine Systems*, 47(6), 958-969. <https://doi.org/10.1109/THMS.2017.2682115>.
- Zhang, B. (2025). Security Analysis of ECG-Based Authentication in Wearable Devices. (Doctoral dissertation, RMIT University).
- Zhang, J. H., Koivumäki, T., & Chalmers, D. (2024). Privacy vs convenience: Understanding intention-behavior divergence post-GDPR. *Computers in Human Behavior*, 160, 108382. <https://doi.org/10.1016/j.chb.2024.108382>.
- Zhang, J., Yin, Z., & Wang, R. (2014). Recognition of mental workload levels under complex human-machine collaboration by using physiological features and adaptive support vector machines. *IEEE Transactions on Human-Machine Systems*, 45(2), 200-214. <https://doi.org/10.1109/THMS.2014.2366914>.
- Zhang, R., Yan, B., Tong, L., Shu, J., Song, X., & Zeng, Y. (2019). Identity authentication using portable electroencephalography signals in resting states. *IEEE Access*, 7, 160671-160682. <https://doi.org/10.1109/ACCESS.2019.2950366>.
- Zhang, X., Gu, H., Fan, L., Chen, K., & Yang, Q. (2022). No free lunch theorem for security and utility in federated learning. *ACM Transactions on Intelligent Systems and Technology*, 14(1), 1-35. <https://doi.org/10.1145/3563219>.
- Zhang, Y., Qin, F., Liu, B., Qi, X., Zhao, Y., & Zhang, D. (2018). Wearable neurophysiological recordings in middle-school classroom correlate with students' academic performance. *Frontiers in Human Neuroscience*, 12, 457. <https://doi.org/10.3389/fnhum.2018.00457>.
- Ziller, A., Mueller, T. T., Stieger, S., Feiner, L. F., Brandt, J., Braren, R., ... & Kaissis, G. (2024). Reconciling privacy and accuracy in AI for medical imaging. *Nature Machine Intelligence*, 6(7), 764-774. <https://doi.org/10.1038/s42256-024-00858-y>.

---

**Goran Mladenović**

Faculty of Engineering,  
Kragujevac,  
Serbia,

[goran.mladenovic@fink.rs](mailto:goran.mladenovic@fink.rs)

ORCID 0000-0002-0717-9715

**Aleksandar Stanković**

Singidunum University,  
Belgrade,  
Serbia,

[aleksandar.stankovic.23@singimail.rs](mailto:aleksandar.stankovic.23@singimail.rs)

ORCID 0009-0004-6711-0711

**Arso Vukićević**

Faculty of Engineering,  
Kragujevac,  
Serbia,

[arso.m.vukicevic@kg.ac.rs](mailto:arso.m.vukicevic@kg.ac.rs)

ORCID 0000-0003-4886-373X

---