EUROPEAN JOURNALS OF EMERGING COMPUTER VISION AND NATURAL LANGUAGE PROCESSING

Safeguarding Identity: A Comprehensive Survey of Anonymization Strategies for Behavioral Biometric Data

Shane T. Aldrich
Department of Computer Engineering, Politecnico di Milano, Italy

Kathy R. Smith
School of Computer Science and Applied Mathematics, University of the Witwatersrand, South Africa

V0LUME01 ISSUE01 (2024)
Published Date: 13 December 2024 // Page no.: - 34-59

ABSTRACT

Behavioral biometrics, encompassing unique patterns in human actions like gait, keystroke dynamics, voice, and eye movements, offer powerful tools for authentication and interaction. However, the rich, often sensitive information embedded in this data poses significant privacy risks, as it can inadvertently reveal personal attributes such as gender, age, health conditions, or emotional states. This article presents a comprehensive survey of anonymization techniques specifically designed for behavioral biometric data. It categorizes existing approaches based on modality (voice, gait, keystroke dynamics, ECG, EEG, eye movements, and gesture recognition) and strategy (data transformation, perturbation, and generative models). The discussion highlights the inherent trade-off between achieving strong privacy guarantees and maintaining sufficient data utility for intended applications. By reviewing the state-of-the-art and identifying persistent challenges, this survey aims to inform future research and foster the development of truly privacy-enhanced behavioral biometric systems.

Keywords: Behavioral Biometrics, Anonymization, Privacy Preservation, Data Utility, Re-identification, Voice, Gait, Keystroke Dynamics, ECG, EEG, Eye Tracking, Differential Privacy, Generative Models.

INTRODUCTION

The rapid pace of digital transformation has ushered in an era of unprecedented data collection, particularly concerning human behavior. Advances in sensing technologies, ranging from sophisticated augmented reality (AR) and virtual reality (VR) headsets to ubiquitous smartphones and wearables, enable the capture of detailed biometric and behavioral data at ever-increasing scales and resolutions [2]. This pervasive data acquisition, often occurring seamlessly and without explicit user awareness, includes browsing habits, location services, and interactions within smart environments equipped with voice assistants and cameras [2].

Behavioral biometrics, a subset of biometrics, distinguishes individuals based on their unique patterns of action, such as how they speak, walk, type, or interact with digital interfaces [103, 303]. Unlike physiological biometrics (e.g., fingerprints, facial features) which are based on static physical characteristics, behavioral biometrics leverage dynamic, time-varying traits. These include, but are not limited to, voice, gait, keystroke dynamics, hand motions, eye gaze, heartbeat (ECG), and brain activity (EEG) [103, 303]. Their continuous and implicit nature makes them highly appealing for applications requiring ongoing authentication and personalized user experiences [13, 22, 154]. For instance, behavioral biometrics can enable seamless user verification in financial transactions, adapt gaming experiences to player profiles, or monitor health conditions remotely [22, 205, 283, 289, 328].

However, the very richness and distinctiveness that make behavioral data effective for identification also render it highly sensitive and vulnerable to privacy breaches [6, 145]. Beyond merely identifying an individual, behavioral patterns can inadvertently reveal a multitude of sensitive attributes, including gender [34, 79, 124, 201, 227, 234], age [34, 148, 234], emotional states [301, 286, 142], health conditions (e.g., Parkinson's, Alzheimer's, or even substance abuse) [58, 59, 128, 118, 123, 304, 60, 240], and even cognitive load or personal interests [50, 146, 176, 120]. The strong correlations between behavioral features and these attributes, coupled with the temporal dependencies inherent in time-series data, make it challenging to protect privacy effectively [2]. The risk of re-identification attacks, where seemingly anonymized data is linked back to an individual, is a significant concern, amplified by the increasing sophistication of machine learning and the availability of large-scale datasets [75, 78, 108, 312].

Given these escalating privacy concerns, the development

and deployment of robust anonymization techniques for behavioral biometric data are no longer merely desirable but essential. The overarching goal is to enable the beneficial utility of this data for various applications while simultaneously ensuring that an individual's identity and sensitive personal attributes remain protected from unintended disclosure or inference [281, 256]. This requires a delicate balance between privacy preservation and data utility, often presenting a complex trade-off.

This article provides a comprehensive and systematic survey of existing anonymization techniques for behavioral biometric data. We aim to:

- Systematize the literature by categorizing current approaches based on the specific behavioral trait they address and the underlying data transformation strategies employed.
- Analyze the conceptual operation, advantages, and limitations of these techniques, highlighting their strengths and weaknesses in achieving privacy goals.
- Discuss the evaluation methodologies used in the literature, identifying areas for improvement and advocating for more rigorous assessment against informed adversaries.
- Identify commonalities and differences across various behavioral biometric traits and their respective anonymization solutions, pinpointing neglected areas and promising future research directions.

By consolidating knowledge in this emerging field, this survey seeks to offer a structured overview that can guide researchers, developers, and policymakers in building more privacy-aware and ethically sound behavioral biometric systems.

Background

This section lays the groundwork for understanding the complexities of behavioral data privacy. We begin by defining key terminology used throughout this survey, differentiate it from related concepts, and then critically review existing surveys to highlight the unique contributions of this work. Finally, we detail the systematic methodology employed to gather and analyze the relevant literature.

Terminology

Precise terminology is crucial for a clear understanding of privacy-enhancing technologies (PETs) in the context of behavioral biometrics.

- Privacy Enhancement/Protection: This broad term refers to any measure taken to obfuscate information from adversarial observers, including service providers. It encompasses various strategies such as data access control, encryption, data minimization, and data modification or perturbation [2]. In this survey, our focus is on techniques that control disclosure when untrusted parties gain access to interpretable data, rather than solely on cryptographic methods that hide data from access entirely [2].
- Anonymity: A specific case of privacy where data cannot be linked to the individual to whom it refers [2]. This extends beyond direct identifiers (e.g., Social Security numbers, full names, which are typically removed in an initial sanitization phase) to include indirect identifiers or quasi-identifiers (e.g., gender, age, zip code, behavioral patterns) that, when combined, can uniquely identify an individual [48].
- Pseudonymity: The process of replacing direct identifiers of a person with new, artificial identifiers (pseudonyms) [4]. While pseudonyms prevent direct linkage to real-world identities, they may still allow for tracking an individual across different data instances or services if the pseudonym is consistent or if enough quasi-identifiers remain. True anonymity aims to prevent even this linkage.
- Identity Disclosure: The threat where an attacker uses available data, including behavioral data, to re-identify an individual, often by linking their behavioral patterns to their real-world identity or to other accounts they hold [4]. For example, a VR headset user might use a pseudonym, but their unique motion or eye-tracking data could be used by a server or other users to identify them across different pseudonymous accounts in a federated metaverse [265]. This is a primary target for anonymization.
- Attribute Disclosure: The threat where an attacker infers sensitive personal attributes (e.g., sex, medical conditions, personal interests, emotional state) from behavioral data that the user did not intend to disclose [4]. For instance, EEG data from a brain-computer interface could be used to infer an alcohol problem based on publicly available datasets [134, 203]. This type of disclosure can lead to discrimination or unwanted targeting.
- Utility: This term quantifies the degree of functionality maintained by a service or application despite the implementation of a privacy mechanism that may hide or perturb part of the data [4]. In the context of behavioral biometrics, utility can refer to the accuracy of biometric authentication, the effectiveness of an adaptive interface, the precision of activity recognition, or the naturalness of a synthesized voice [4]. The goal of anonymization is to maximize privacy while minimizing the degradation of utility.
- Privacy-Utility Trade-off: An inherent challenge in PETs, this refers to the inverse relationship between the level of privacy achieved and the retained utility of the data [4]. Enhancing privacy often comes at the cost of reduced data utility, and vice versa. The design and tuning of anonymization techniques aim to optimize this trade-off for specific application scenarios.

- Statistical Disclosure Control (SDC): A field focused on protecting microdata sets (databases with individual-level information) while ensuring their usefulness for research [294]. Concepts from SDC, such as additive noise masking [122] and suppression [122], often inspire anonymization techniques for behavioral data.

Related Surveys

While the importance of behavioral biometrics has led to numerous surveys, most have primarily focused on their uniqueness and suitability for identification, comparing the accuracy of different approaches and their applicability in various contexts. Surveys by Alzubaidi and Kalita [13], Liang et al. [154], Mahfouz et al. [165], and Meng et al. [182] provide overviews of behavioral biometrics for user authentication. More specific reviews delve into gait recognition [290], keystroke dynamics [19, 272], eye gaze [136], and brainwave biometrics [99]. However, these surveys typically only acknowledge the potential for sensitive inferences or identity leaks without providing an in-depth discussion of privacy countermeasures.

A distinct line of research investigates potential privacy attacks on behavioral data, focusing on attribute inferences [18, 34, 124] or user de-identification [75, 78, 108, 312]. Dantcheva et al. [53] offer an extensive overview of "soft biometrics" (e.g., gender, age, ethnicity) that can be inferred from primary biometrics, particularly from image and video data. Ciriani et al. [48] survey k-anonymity, a technique applicable to protecting soft biometrics in tabular data. More recently, Laishram et al. [149] surveyed privacy-preserving face recognition systems, and Golda et al. [95] and Wang et al. [292] explored the privacy implications of large language models (LLMs) and generative machine learning.

Despite this body of work, a comprehensive view of the problem of anonymizing behavioral data, encompassing a wide range of traits, existing solutions, and persistent challenges, has been largely absent. While Ribaric et al. [245] reviewed de-identification techniques for visual and multimedia content, their coverage of behavioral data protection was limited to video, audio, or image captures of a few traits (voice, gait, and gesture), neglecting other sensor types. Similarly, Nhat Tran et al. [281] surveyed biometric template protection generally, without deep dives into the specific anonymization needs of behavioral biometrics. Meden et al. [180] focused on privacy-enhancing faces, and Shopon et al. [256] provided a broader look at biometric de-identification but with a different taxonomic focus.

The critical gap identified in the existing literature is the lack of a systematic review that:

1. Examines a comprehensive set of both traditional and modern behavioral traits for which anonymization solutions have been proposed.
2. Considers diverse data collection sensors and use cases beyond just video or audio.
3. Provides a comparative analysis of evaluation approaches across different behavioral biometrics.

By addressing these shortcomings, this survey aims to offer a more holistic understanding of the field, reveal similarities and differences between anonymization approaches, and pinpoint open research questions that can drive future advancements.

METHODOLOGY

This survey was conducted following the systematic literature review guidelines proposed by Kitchenham [141], ensuring a rigorous and reproducible approach to identifying and analyzing relevant studies on privacy techniques for behavioral data. The procedure is summarized in Figure 1 in the original document.

Our guiding research question for this survey was: "What techniques are applicable to protect behavioral data privacy?" From this central question, our objectives were to understand how these techniques operate, the level of privacy protection they afford, and their inherent limitations and unresolved challenges.

To address this, we initiated our process by thoroughly exploring existing literature on biometrics [5, 13, 53, 98, 165, 182, 221, 303] to compile a comprehensive list of behavioral traits utilized for person identification. This initial exploration yielded a diverse set of traits, including: brain activity (also referred to as cognitive biometric), eye gaze, facial expression, gait, gesture, handwriting, haptic, heartbeat, keystrokes, lip, motion, mouse, thermal, touch, and voice.

Next, we formulated our search strategy by combining these identified behavioral traits with keywords related to privacy. The core search strings included "[Behavioral Trait] AND "Privacy" OR "[Behavioral Trait] AND "Anonymization" OR "[Behavioral Trait] AND "Re-identification". These search terms were applied across major academic databases in computer science, including IEEE Xplore, ACM Digital Library, DBLP, and Google Scholar. No constraints were placed on the publication date during the initial search, resulting in a broad collection of 364 papers published between 2007 and October 2024.

Following the initial search, a duplicate filtering process was performed to remove redundant entries, ensuring each unique publication was considered only once. Subsequently, a pre-screening phase was conducted to refine the scope of the survey. During this phase, we developed a preliminary taxonomy of privacy solutions, which led to the decision to narrow our focus specifically on anonymization techniques aimed at protecting the publication of behavioral data from identity and attribute disclosure attacks. We were particularly interested in approaches that assumed a scenario where data is collected, sanitized, and then published or shared with a

service or application, while still retaining a degree of utility for the intended service.

Based on this refined scope, the following inclusion and exclusion criteria were applied to select the primary studies for in-depth analysis:

Exclusion Criteria:

1. Publication Format: Documents that were not peer-reviewed academic journal articles or conference papers (e.g., preprints, technical reports not formally peer-reviewed, theses, books) were excluded.

2. Availability: Papers that could not be retrieved or accessed through IEEE Xplore, ACM Digital Library, DBLP, or Google Scholar were excluded.

3. Language: Publications not written in English were excluded to ensure consistent analysis.

4. Redundancy: If multiple papers by the same authors addressed the same work, only the most complete and up-to-date version was included, superseding earlier or less comprehensive publications.

5. Scope Mismatch: Privacy protection techniques that did not primarily focus on identity or attribute anonymization with data utility preservation were excluded. This specifically meant excluding approaches solely focused on encryption for confidentiality (where data is never interpretable by untrusted parties), or access control mechanisms without data transformation.

6. Insufficient Detail: Anonymization approaches described at a high level without sufficient technical detail to properly address our guiding research question regarding their mechanisms, protection levels, and limitations were excluded.

The rigorous application of these criteria yielded a final corpus of 142 peer-reviewed works on behavioral data anonymization. These selected studies were then clustered and analyzed according to the specific behavioral trait they aimed to protect: gait, brain activity, heartbeat, eye gaze, voice, and hand motions (which included handwriting, keystrokes, mouse movements, and hand gestures). It is noteworthy that, despite our comprehensive search, we did not find any suitable papers on facial expression, lip, touch, and haptic traits that met all our defined criteria for anonymization techniques.

This systematic methodology ensures that the survey is comprehensive, unbiased in its selection, and provides a robust foundation for analyzing the state-of-the-art in behavioral data anonymization.

3. Behavioral Data Applications and Privacy Concerns

Behavioral data, by its very nature, is a rich tapestry of human expression and interaction. Its collection and analysis enable a myriad of valuable services for individuals and organizations alike. However, this immense utility is inextricably linked to profound privacy implications. This section delves into the characteristics of behavioral biometric data, outlines the typical scenario in which it is processed, explores its diverse applications, defines the concept of utility within this context, and critically examines the inherent privacy concerns and the models of adversaries attempting to exploit this data.

3.1 Behavioral Biometric Data Characteristics

Behavioral biometric data stands as a unique subclass within the broader category of biometric information. Unlike static physiological biometrics, which capture fixed anatomical features (e.g., fingerprint ridges, iris patterns), behavioral biometrics record dynamic patterns of human action over time. This includes, but is not limited to, the nuances of speech, the rhythm of typing, the unique sway of a walk, or the subtle movements of the eyes.

A fundamental distinction of behavioral biometric data, particularly when compared to explicit fields in traditional microdata sets (where sensitive columns like "name" or "address" are clearly defined), is that its privacy-sensitive components are often implicit. The raw data itself may not immediately appear sensitive, but the patterns within it can reveal profound personal details. For instance, the way a person walks (gait) is not just a sequence of movements; it can implicitly indicate a recent injury if the pattern exhibits a limp, even if no explicit "injury" field exists in the data [2].

Moreover, behavioral biometric data is almost universally captured as a time-series, meaning it consists of a sequence of observations recorded over time. This temporal dependency is critical; the current state of a behavioral trait is highly dependent on its preceding states. Beyond temporal links, there are also physiological dependencies – human bodies operate within inherent physical and biological limitations. For example, the motion of a foot is intrinsically linked to the motion of the corresponding leg. These strong, often complex, dependencies between individual data points, both temporally and physiologically, pose significant challenges for anonymization. Simple randomization or perturbation techniques, which might suffice for independent data points, can be ineffective because an intelligent attacker could leverage these underlying dependencies to reconstruct the original, clear data and extract implicit sensitive information from the supposedly anonymized records [2]. The presence of context information and ingrained habits, which manifest as strong, persistent signals within the data, further compounds the difficulty of achieving truly effective anonymization.

3.2 Scenario

In this survey, we operate under a data-publishing scenario, as illustrated in Figure 2 in the original document. This model assumes a workflow where behavioral data is initially captured, subsequently transformed in a privacy-protective manner, and then either published, processed by, or shared with a service

provider or application. This scenario explicitly includes instances of involuntary publication, which can occur through various means, such as the unintentional leakage of biometric templates from an authentication system, or the commercial sale of fitness tracker data to third parties [3].

Within this framework, a critical assumption is that the utility of the protected and modified data is preserved to a sufficient extent. This means that despite the privacy-enhancing transformations, the received service (e.g., a personalized recommendation, a responsive virtual reality game, or a health monitoring alert) remains meaningful and functional for the user [4]. The challenge lies in finding the optimal balance where privacy is maximized without rendering the data useless for its intended purpose.

3.3 Applications

The pervasive nature of human-computer interaction (HCI) means that virtually every input over time constitutes behavioral biometric data. Traditional input modalities like keystroke patterns and mouse movements have long been central to computer systems [324]. However, the landscape is rapidly evolving with the rise of new input methods such as touch, voice, and gestures, which are anticipated to gain even greater prominence [2]. This shift is particularly evident in emerging fields like mixed reality (MR), which integrates multiple input modalities and necessitates continuous monitoring of users to provide immersive experiences.

Beyond general HCI, behavioral data is invaluable across several specialized domains:

- Healthcare and the Quantified Self: Advances in sensor technology and machine learning have revolutionized healthcare, enabling applications for activity recognition, fall detection, and remote health monitoring [59, 212, 226]. These applications are crucial for elder care, supporting individuals with chronic illnesses, and facilitating early diagnosis. Data commonly collected includes gait and motion information from accelerometers and gyroscopes embedded in devices, as well as biosignals like heartbeat (ECG) and brain activity (EEG). This data can be processed to provide real-time health feedback, guide relaxation exercises, or detect and signal cognitive states such as stress [2].

- Biometric Recognition: This is one of the most significant and extensively researched application areas for behavioral data [13, 115, 165, 182]. The inherent uniqueness of an individual's behavior – whether it's their walking style or typing rhythm – allows for robust identity verification. Given that these patterns can often be sensed implicitly as a person interacts with, wears, or carries a device, behavioral biometrics are frequently considered more user-friendly than traditional biometrics like fingerprints [29, 30]. Consequently, they serve as a compelling alternative or complement to password-based authentication. Academic research has demonstrated the feasibility of numerous behavioral traits for user authentication, including keystroke patterns [272], gait [290], touch [273], mouse movement [324], brain activity [99], and even breathing patterns [41, 42]. Several commercial solutions, particularly in the financial sector, already leverage behavioral biometrics to prevent fraud by detecting anomalous behaviors [22, 205, 283, 289].

- Personalization: A substantial portion of behavioral data-driven applications focuses on personalization, aiming to tailor user experiences. This includes adaptive interfaces and services that dynamically adjust their content or appearance based on predicted user preferences derived from their behavior [4]. Examples span various domains:

  - Online Gaming: Behavioral data is used to personalize gaming experiences, such as dynamically adjusting difficulty levels to provide a more satisfactory experience for the player [328].

  - Recommender Systems: User behavior informs recommender systems that suggest online content or advertisements, enhancing relevance and engagement [241].

  - Education: In educational contexts, behavioral data can be used to tailor learning experiences to a student's mental state (e.g., attention level, stress), optimizing pedagogical approaches [130].

The diverse applications underscore the immense potential of behavioral data to enhance convenience, security, and personalization in the digital realm.

3.4 Utility

The concept of utility in the context of behavioral biometric data refers to the effectiveness and functionality retained by the data after anonymization, for its intended application. This measure is highly dependent on the specific service being provided:

- Biometric Authentication: For an authentication application, the primary measure of utility is the system's ability to accurately verify an individual's identity. This is often quantified by metrics such as the Equal Error Rate (EER), False Acceptance Rate (FAR), or False Rejection Rate (FRR) [4].

- Human-Computer Interaction (HCI): If behavioral data serves as an input modality for computer systems (e.g., gestures, keystrokes), its utility is measured by its precision and timeliness in facilitating user interaction [320]. The input must remain reliable and responsive.

- Healthcare Applications: In healthcare, utility might be assessed by the system's accuracy in detecting abnormal behavioral patterns, monitoring specific health aspects (e.g., step counting), or inferring patient preferences for personalized care [4]. For video-based gait analysis, utility could also encompass the naturalness and convincing appearance of the de-identified gait in the

video [125].

- Personalization: For applications focused on personalization, utility is often tied to the accuracy of predicting user preferences or states (e.g., mood, attention level) to deliver tailored content or experiences.

In essence, the utility of the anonymized behavioral data is assessed by how well it performs the specific task for which it was originally collected, despite the privacy-preserving transformations. The challenge lies in minimizing the degradation of this performance while maximizing privacy.

3.5 Privacy Concerns

The significant amount of personal information implicitly embedded within behavioral data-driven applications gives rise to substantial privacy concerns. As established, behavioral data is inherently rich in individuating information, making it a powerful biometric. Consequently, any entity collecting this data, even if for a legitimate purpose, possesses the capability to identify individuals or infer sensitive attributes, irrespective of the service's primary function. This problem is exacerbated by the fact that individuals may often be unaware of the extent to which their behavior is being measured, either due to a lack of transparency in data collection practices, inadequate consent frameworks, or even covert surveillance.

Beyond direct identity, behavioral data carries a wealth of potentially sensitive information that can be misused. For example, traits like voice, eye gaze, gait, or brain responses are known to correlate with various diseases [59, 304], mental states and emotions [269, 301], and involuntary physiological reactions (e.g., pupil dilation) can betray a person's interests or cognitive load [147, 146, 176]. This means that data collected for one purpose (e.g., authentication) could be repurposed to infer highly personal details, potentially leading to discrimination (e.g., increased insurance premiums due to inferred medical conditions) or threatening user autonomy through highly targeted advertising.

Technically, the general process for inferring identity or other information from behavioral data typically involves four sequential steps, as depicted in Figure 3 in the original document:

1. Data Acquisition: The initial step involves recording and digitizing the raw behavioral data from various sensors (e.g., microphones, cameras, accelerometers, EEG headsets).

2. Feature Representation: Relevant features are extracted from the raw data. These features are designed to capture the unique patterns indicative of identity or specific attributes. For instance, Mel-frequency cepstral coefficients (MFCCs) are common features for voice, while gait energy images are used for gait.

3. Dimension Reduction: The extracted feature representation, often high-dimensional, is typically reduced to a lower-dimensional space. This step aims to simplify the data, remove redundancy, and make it more manageable for subsequent analysis, while ideally retaining the discriminative information.

4. Inference: In the final step, the reduced feature representation is fed into machine learning models (e.g., classifiers, regression models) to perform the desired inference. This could be classifying the user's identity (authentication), assigning them to a specific attribute class (e.g., male/female), or estimating a continuous measure (e.g., degree of depression).

Consider a voice-controlled personal assistant: it might use this workflow to authenticate the user commanding to open an email application. However, the same voice features could be exploited to classify the user's mood, leading to the delivery of highly targeted advertisements – a practice that raises significant ethical and privacy concerns.

The proliferation of affordable consumer wearables equipped with numerous sensors (e.g., VR/AR devices with eye-tracking, head pose detection, and EEG sensors) further exacerbates the privacy issue. Once data is collected, even for a legitimate, user-consented functionality like fraud detection based on behavioral anomalies, it can be subsequently exploited to infer private information. This highlights the urgent need for robust techniques to protect behavioral data against unintended identity and attribute disclosure.

3.6 Attacker Model

To effectively design and evaluate anonymization techniques, it is crucial to clearly define the capabilities and goals of the adversary. Our attacker model assumes the following characteristics:

- Access to Behavioral Biometric Data: The adversary has gained access to the behavioral biometric data of one or multiple users. This access can originate from various sources:
  - Service Provider: The adversary might be the service provider itself, who, despite providing a legitimate service, seeks to exploit the collected behavioral data for secondary, privacy-invasive purposes (e.g., inferring sensitive attributes for targeted advertising).
  - Malicious User: The adversary could be another user of the service who has managed to obtain or infer data belonging to other individuals (e.g., face images downloaded from social media, or shared biometric templates).
  - Data Breach/Leak: The data could have been acquired through an unintentional release, a security breach, or a hack of a system storing biometric data [3].
- Full Access and Inference Capabilities: Since the adversary has full access to the behavioral biometric data

(even if it has undergone some form of initial, weak anonymization), they can freely select and apply any available inference technique. This includes state-of-the-art machine learning models, advanced statistical methods, and data analysis algorithms to perform privacy inferences (e.g., re-identification, attribute inference).

- Prior Knowledge: The adversary may possess additional prior knowledge about the target user. This can include:
  - Biometric Templates: Leaked or compromised biometric templates from other systems.
  - Soft Biometrics: Previously known sensitive attributes (e.g., gender, approximate age, known health conditions) that can serve as auxiliary information to aid re-identification or attribute inference.
  - Publicly Available Datasets: Access to public datasets containing similar behavioral data, which can be used to train powerful inference models or to cross-reference with the anonymized data. For example, publicly available EEG datasets of alcoholic and non-alcoholic individuals could be used to build a classifier for newly gathered data [134, 203].

This robust attacker model dictates that anonymization techniques must not only obscure direct identifiers but also resist sophisticated inference attacks that leverage complex data patterns and external knowledge. The evaluation of such techniques must therefore account for these adversarial capabilities, moving beyond simple accuracy metrics on weakly anonymized data.

4. A Taxonomy of Solutions for Behavioral Data Privacy

Based on our comprehensive literature analysis, we identify two primary privacy threats that necessitate anonymization in scenarios where behavioral data is collected or processed by a third party. These threats directly inform the goals of anonymization techniques and can be understood in terms of the capabilities of the attacker model described previously.

4.1 Privacy Threats and Goals

The two main privacy threats are:

- Identity Disclosure: In this scenario, the attacker's primary objective is to use the behavioral data to identify the individual. This threat assumes the attacker can link the target's behavioral data to their real-world identity. The goal might be to link a user's account and data in a professional application to their account in an entertainment application, thereby gaining a more comprehensive profile of their activities [4]. As an illustrative example, consider a user of a VR headset entering a federated Metaverse. Even if the user employs a pseudonym, the server or other users could potentially use transmitted behavioral data (e.g., controller/headset motions, eye-tracking data) to identify the user across different pseudonyms [265]. This is particularly concerning given that behavioral data is sometimes sold to third parties or unintentionally leaked through breaches [3].
  - Anonymization Goal: Identity Protection: This involves transforming the behavioral biometric data of a person such that their identity can no longer be linked to the data. This goes beyond simple pseudonymization, which merely replaces an identifier with a new one, aiming instead to prevent identification altogether, even through indirect means.
- Attribute Disclosure: Here, the attacker's goal is not necessarily to re-identify the user across accounts, but rather to derive sensitive attributes from the available behavioral data that the user did not intend to disclose. These attributes can be long-living (e.g., sex, age, medical conditions) or short-living (e.g., mental state, temporary health conditions, emotional state) [4]. For example, a classifier trained on publicly available electroencephalogram datasets of alcoholic and non-alcoholic persons [134, 203] could be used to determine if newly gathered EEG data from a brain-computer interface application belongs to a user with an alcohol problem.
  - Anonymization Goal: Attribute Protection: This involves transforming the behavioral biometric data in a way that specific private attributes of the person can no longer be inferred from the data. An extreme form of attribute protection is template protection, where the identity verification (e.g., for authentication) remains possible, but all other attributes are protected [4].

Figure 4(a) in the original document visually represents this taxonomy of anonymization techniques based on their privacy goal.

4.2 Taxonomy of Data Transformation Applied

Building upon the analysis of state-of-the-art protection methods, we have developed a classification of anonymization techniques based on the type of transformation applied to the original data to derive the anonymized, protected data. This taxonomy, depicted in Figure 4(b) in the original document, includes only fundamental concepts, though many advanced anonymization techniques combine multiple of these approaches. A defining characteristic shared by all anonymization methods in this taxonomy is their aim to provide irreversible transformations, meaning it is computationally infeasible or impossible to revert the anonymized data back to its original form.

The primary distinction in our taxonomy is between non-deterministic and deterministic techniques:

- Non-deterministic Methods: These methods incorporate randomness into their transformation process, meaning that the same input data can yield different anonymized outputs on different runs. This inherent variability makes it harder for an attacker to link anonymized data back to the original.

- Random Perturbations: This involves a random transformation of the data into a different domain or a random rearrangement of data components. The randomness is applied in a way that significantly alters the original structure while attempting to retain some statistical properties relevant for utility.

- Noise Injection: These methods directly add random noise to the data points. This concept is analogous to additive noise masking in Statistical Disclosure Control (SDC) [122], where random noise is added to original values before release. While noise injection can be combined with deterministic transformations, its core principle relies on obscuring original values through random additions. The challenge lies in adding enough noise for privacy without destroying utility.

- Deterministic Methods: These methods produce the same anonymized output for identical input data every time. While they lack the inherent variability of non-deterministic methods, their transformations are designed to be irreversible and privacy-preserving. Deterministic methods are further categorized into two main types:

- Removal: This method involves eliminating specific data points or features from the dataset such that they no longer influence the anonymized result.

- Coarsening: This refers to reducing the granularity of the data. It can involve removing parts of each data point, making the data more sparse, or aggregating data into broader categories. For instance, reducing the precision of numerical values or grouping fine-grained temporal events into larger time bins.

- Feature Removal: This involves completely eliminating data points associated with a specific feature or set of features. In SDC, this is known as suppression [122]. If a microdata set contains too few records sharing a combination of quasi-identifier values (making re-identification risky), specific values of individual variables are suppressed (replaced with missing values) to expand the number of records conforming to each combination, thereby enhancing privacy.

- Conversion: These methods transform the data points into a new representation. This new representation is typically derived from the original domain but is designed to obscure identifying information.

- Discrete Conversion: The result of the conversion is a discrete value. This can involve mapping original data points to a limited set of predefined categories or symbols.

- Continuous Conversion: The result of the conversion is a continuous value. This often involves complex mathematical transformations, signal processing techniques, or machine learning models that generate a new, continuous data representation that is unlinkable to the original but retains relevant utility.

This taxonomy provides a structured framework for understanding the diverse array of anonymization techniques applied to behavioral biometric data, allowing for systematic comparison and analysis across different traits and approaches.

5. Anonymization Techniques

This section provides a detailed overview of anonymization techniques, organized by the specific behavioral biometric trait they aim to protect. For each trait, we delve into its utility, the associated privacy threat space, a comprehensive analysis of the anonymization techniques employed (categorized according to our taxonomy), and a critical review of the evaluation methodologies used in the literature. Voice is presented first due to its extensive research, followed by gait, hand motions, heartbeat, eye gaze, and brain activity.

5.1 Voice

The human voice is a complex acoustic signal generated by the interplay of the larynx and the vocal tract. The larynx produces a fundamental frequency (F0), perceived as pitch, while the vocal tract acts as a filter, shaping the sound through its unique resonant properties, which are influenced by its length and configuration [35]. These physical characteristics contribute significantly to an individual's unique vocal timbre.

For analysis, speech signals are often transformed into representations that highlight key acoustic features. The log-spectrum, which approximates human perception of sound intensity, is a common starting point. Applying a Fourier Transform (FFT) or cosine transform to the log-spectrum yields the cepstrum (see Figure 5 in the original document), a representation useful for estimating F0 and other vocal tract characteristics [35]. The Mel-frequency cepstral coefficients (MFCCs) are derived from the cepstrum by sampling frequencies on the Mel scale, which mimics the non-linear human auditory perception, providing a compact representation of the signal's macrostructure [35].

Speaker recognition (identifying who is speaking) is a core application that leverages these vocal features. Traditional methods include Gaussian Mixture Models (GMMs), which represent speakers as distributions of their feature vectors (typically MFCCs) [243]. A Universal Background Model (UBM), a GMM trained on a wide variety of non-target speakers, serves as a general reference. Speaker-specific GMMs are then adapted from the UBM using techniques like Maximum A Posteriori (MAP) adaptation [244]. Supervectors, formed by concatenating the means of MAP-adapted GMMs, can be classified using Support Vector Machines (SVMs) [36]. An extension, Total Variability (TV), maps supervectors to a low-dimensional space (i-vectors) that models both speaker and channel variability, becoming a de facto standard for speaker identification [64]. More recently, x-vectors, extracted via

deep neural networks (DNNs), have emerged as robust speaker embeddings [260].

5.1.1 Utility

Voice recordings serve multiple utilities. Primarily, they facilitate human-to-human information transmission, where intelligibility of speech content is paramount. Increasingly, voice also functions as a critical input modality for computer systems [231], where its utility is measured by the accuracy and responsiveness of voice commands. Beyond content, the mere detection of speech in audio samples can be useful for applications like crowd detection [49]. Crucially, voices uniquely identify their speakers, making them suitable for authentication and recognition purposes in biometric systems [247].

5.1.2 Threat Space

The privacy threats associated with human voices are multifaceted, ranging from direct individual identification to the inference of sensitive attributes and even identity theft.

- Identification: The ability to identify individuals by their voice has long been recognized, both by humans and, increasingly, by automated systems.
- Attribute Inference: Voices convey more than just identity; they enable the inference of attributes such as gender [79] or emotional state [301].
- Identity Theft/Cloning: Modern speech synthesis and voice conversion methods allow for the creation of highly realistic fake voice recordings of a target speaker, enabling sophisticated identity theft or circumvention of speaker authentication systems.
- Semantic Content Disclosure: Unlike many other behavioral biometric traits, voice carries semantic meaning (the actual spoken words), which can be highly sensitive and subject to privacy breaches if exposed to unintended listeners.

5.1.3 Additional Privacy Goal: Speech Blurring

Beyond identity and attribute protection, voice data introduces an additional, unique privacy goal: speech blurring. This aims to intentionally destroy the intelligibility of the speech, thereby protecting its semantic content from unauthorized listeners, even if the speaker's identity might still be partially discernible.

5.1.4 Anonymization Techniques

We now present surveyed anonymization techniques for human voices, categorized by their transformation approach.

Random Perturbation

This category involves introducing randomness to disrupt the unique vocal signature.

- Parthasarathi et al. [214] extended their feature removal methods by shuffling voice blocks to add randomness, aiming to obscure speaker identity in diarization tasks.
- Mtibaa et al. [193] proposed a template protection scheme for GMM-UBM speaker identification systems that shuffles feature vectors based on a secret key, making the template cancelable.
- Sharma et al. [254] utilized a self-attention channel combinator to add noise to voice signals, contributing to de-identification.

Noise Injection

These methods directly add random noise to the voice signal to obscure identifying information.

- Tamesue et al. [271] proposed a simple method to make speech unintelligible by playing pink noise within a specific frequency range at various decibel levels.
- Ma et al. [163] developed a technique for smartphone recordings, where two ultrasound waves interfere to create random low-frequency waves that disrupt smartphone microphones, effectively blocking recordings without being audible to humans. They demonstrated effectiveness up to 5 meters depending on the device.
- Hashimoto et al. [107] proposed adding white noise to prevent speaker identification from recordings in physical spaces. Their evaluation showed a reduction in speaker identification (EER from 2% to 17%) while maintaining high speech intelligibility (short-time objective intelligibility [270] from 1 to 0.9).
- Ohshio et al. [208] trained "babble maskers" from pre-recorded speakers by averaging speech segments. These maskers, selected based on the target person's fundamental frequency and pitch, are applied to de-identify recordings.
- Vaidya et al. [285] explored adding random noise to pitch, tempo, pause, and MFCC features, though their description was concise.
- Hamm et al. [103] introduced a differentially private min-max filter, which adds noise either before or after filtering to minimize privacy risk while maximizing utility.
- Han et al. [104] formally defined "voice-indistinguishability" using differential privacy, applying it to x-vectors as speaker representations. They used angular distance as a similarity metric and established an upper limit on this distance for indistinguishability.
- Qian et al. [236] presented a method to sanitize speech for virtual assistants by applying vocal tract length normalization via a compound frequency warping function and adding Laplace noise for robustness, claiming differential privacy. A follow-up work [235] further investigated its security, and Srivastava et al. [263] evaluated it against stronger attackers.

- Shamsabadi et al. [253] aimed for theoretical privacy guarantees in speaker transformation by adding differential privacy to pitch and context features in their latent spaces, though the impact of correlations between speech segments on DP guarantees remains an open question.

Feature Removal

This involves removing or altering specific acoustic features that are highly indicative of speaker identity.

- Parthasarathi et al. [215] proposed three methods for privacy-aware speaker change detection: adaptive filtering (using LP residuals), removing all subbands except specific high-frequency ones (1.5-2.5 kHz and 3.5-4.5 kHz), and using only the spectral slope. In a separate work [213], they found MFCCs more effective than real cepstrum for speaker diarization. Agarwal et al. [7] proposed a similar scheme, transforming segmented speech into the frequency domain, selecting important peaks, and interpolating a new signal.

- Wyatt et al. [297] and Zhang et al. [322] proposed feature removal for speaker segmentation and conversation detection, extracting non-initial maximum autocorrelation peaks, total autocorrelation peaks, relative spectral entropy, and frame energy. However, a privacy evaluation was missing in both.

- Ditthapron et al. [69] investigated removing non-target speakers in speech assessment by extracting speaker representations from MFCCs via an encoder and filtering out non-target speakers in a matching network. A convincing privacy evaluation was noted as missing.

- Nelus et al. [201] trained a DNN using adversarial learning to extract features that allow gender recognition but not speaker identification, showing a drop in identification from 61% to 26% with only a 1% drop in gender recognition. They also applied a similar system to remove speaker identities from urban sound recordings [200].

- Cohen-Hadria et al. [49] used a neural network to extract voices from mixed background and voice noise. They removed attributes by low-pass filtering at 250 Hz or by using the first five MFCC components to create a new voice, reducing identification from 43% to 29%.

Discrete Conversion

These methods transform voice features into discrete, often cancellable, representations.

- Pathak et al. [217] presented a hashing algorithm for voice authentication, transforming speaker supervectors into low-dimensional "buckets" using locality-sensitive hashing, allowing approximate nearest-neighbor comparisons.

- Portelo et al. [229, 230] proposed a template protection scheme based on secure binary embeddings for speaker identification systems using supervectors and i-vectors. These embeddings maintain a proportional relationship between Euclidean distance of original vectors and Hamming distance of hashes, enabling comparison via SVMs with Hamming distance kernels.

- Billeb et al. [27] proposed a template protection scheme based on fuzzy commitment, extracting features from the magnitude spectrum (via FFT) and applying MAP adaptation of a GMM-UBM system. The template is stored as an error-correcting code and hash combination.

Continuous Conversion

This is the most common category, aiming to create an anonymized speech recording that sounds natural but obscures identity.

- Speaker Transformation: This process manipulates a speaker's voice characteristics to sound like a target speaker (natural or synthetic), while preserving linguistic content.

  - Jin et al. [129] evaluated GMM-mapping-based speaker transformation to a synthetic voice (kal-diphone), also testing duration transformation and extrapolated transformation.

  - Pobar et al. [225] combined GMM mapping with a harmonic stochastic model, applying existing transformation functions without needing a parallel corpus for new speakers, reducing identification accuracy from 97% to 9%.

  - Justin et al. [132] investigated intelligibility of transformed speakers using diphone and HMM-based speech synthesis systems, evaluating word error rate with human listeners.

  - Abou-Zleikha et al. [3] explored optimal target speaker selection for de-identification, formulating it as an optimization problem to minimize identification rate while maximizing reconstruction quality.

  - Pribil et al. [234] modified prosodic and spectral features to make speakers sound older/younger or more female/male for de-identification.

  - Bahamanienezhad et al. [17] developed a speaker transformation using a convolutional encoder/decoder network to map spectral features to a target speaker, fusing the result via averaging or gender-based averaging.

  - Fang et al. [82] used x-vectors, calculating a mean x-vector from random unrelated speakers, achieving EER up to 34% for anonymization. Mawalim et al. [177] improved this by scaling F0, increasing utterance length, and using singular value modification for x-vector combination, boosting EER to 54%. Prajapati et al. [232] further enhanced this with a CycleGAN.

  - Cheng et al. [45] proposed a speaker transformation with separate encoders for content and speaker identity, recombining them in a decoder for anonymized utterances.

- Panarielle et al. [210] used neural audio codecs (NAC) for speaker transformation, independently encoding content and identity before combining them with transformer models.
- Miao et al. [186] developed a language-independent speaker transformation based on the VoicePrivacy [279] challenge baseline, showing effectiveness on English and Mandarin.
- Hintz et al. [110] used a GAN to anonymize stuttering speech, preserving the pathology while removing speaker identity.
- Yang et al. [306] focused on low-latency speaker transformation.
- Yao et al. [308] aimed to improve anonymized speaker distinctiveness by scaling formant and pitch information.
- Meyer et al. [185] proposed a speaker anonymization method that preserves the speaker's prosody.
- Nespoli et al. [202] suggested using two sequential speaker transformation systems for enhanced anonymization.
- Chang et al. [40] and Meyer et al. [184] investigated different averaging strategies for target speaker creation.
- Yuan et al. [316] trained an autoencoder for synthetic data generation to create random speakers.
- Lv et al. [162] used autoencoders to obtain latent speaker representations and selected similar ones from a pool using k-means.
- Yao et al. [307] encoded speakers as matrices, decomposed them with SVD, and logarithmically transformed importance values for reconstruction.
- Miao et al. [187] extended their method [186] by removing the speaker pool and using adversarial perturbation for speaker vector transformation.
- Perero-Codosera et al. [219] also proposed an adversarial perturbation approach for anonymizing X-vectors.
- Yao et al. [309] suggested removing random dimensions of an X-vector to create a new speaker identity.

● Adversarial Perturbation: This technique uses machine learning systems trained with dual loss functions: one to minimize the privacy attribute and another to maximize utility.

- Cheng et al. [43] proposed VoiceCloak, a convolutional perturbation injector that takes room impulse response and original voice to output an anonymized voice.
- Deng et al. [65] presented V-Cloak, a convolutional autoencoder trained to minimize identification while preserving timbre and intelligibility, demonstrating real-time anonymization capabilities.
- Chouchane et al. [47] used adversarial training to create speaker embeddings for speaker verification that work for verification but not sex recognition, addressing fairness concerns.
- Xiao et al. [299] developed a microphone module that anonymizes speakers by adding adversarial perturbation to the sound signal, using a genetic algorithm to find the perturbation, resulting in low latency.
- Ravi et al. [240] developed an adversarial perturbation for depression detection in speakers, aiming to enhance accuracy and privacy through speaker disentanglement.
- Ali et al. [9] proposed an autoencoder to anonymize voice at the network edge for voice assistants, extracting privacy-friendly features by training classifiers on latent codes and performing gradient reversal to unlearn identity, gender, and language features.
- Yoo et al. [314] used a CycleGAN with a variational autoencoder as its generator for speaker anonymization, trained against a DNN speaker recognition system as the discriminator.

● Frequency Warping: Similar to speaker transformation, but specifically focuses on altering the frequency spectrum, often to remove vocal tract length normalization (VTLN) characteristics which are speaker-specific.

- Faundez-Zanuy et al. [83] explored phase vocoder and VTLN for gender protection, reducing gender recognition to chance level, though also impacting identity recognition.
- Valdivielso et al. [1] presented a speaker protection approach that transforms pitch and frequency axes.
- Lopez-Otero et al. [161] used frequency warping and amplitude scaling (affine transformation in cepstral domain) for speaker protection in depression detection contexts, increasing EER for speaker identification from 9.7% to 44% while maintaining depression detection utility.
- Magarinos et al. [164] also used frequency and amplitude warping, applying dynamic frequency warping (DFW) to map source spectrum bins to target spectrum, reducing identification from 99% to 4%.
- Aloufi et al. [11] used a CycleGAN to transform emotional speech to neutral speech, hiding emotional state (from >70% to ~20%) and sex (from ~99% to 50% chance level).
- Srivastava et al. [263] evaluated multiple speaker protection methods against informed attackers, showing large differences in EER between ignorant (up to 50%) and

informed (11%) attackers, highlighting the importance of strong attacker models.

- Patino et al. [218] pseudonymized speakers by transforming their McAdams coefficients (LPC coefficients transformed into poles), showing good performance against ignorant attackers (EER from 3% to 26%) but less against informed ones (5% EER). Gupta et al. [100] improved this by modifying both angle and radius of complex poles.

- Mawalim et al. [178] proposed two frequency modifications: segmenting and resampling speech for pitch shifting, and using phase propagation for segment recombination.

- Ganzepoglu et al. [92] modified the VoicePrivacy challenge B1 baseline to produce better anonymized fundamental frequencies using X-vectors and a mask.

Continuous Conversion + Random Perturbation

- Canuto et al. [39] proposed a template protection method where feature vectors are shuffled via a randomized sum based on a secret key, making templates cancellable.

- Prajapati et al. [233, 257] combined voice conversion with speed perturbation by changing sequence length and tempo, showing that speed perturbation strengthens anonymization.

Continuous Conversion + Noise Injection

- Kondo et al. [143, 144] created "babble maskers" by averaging 10-second speech segments (speaker-dependent or gender-based) and applying them to recordings.

- Qian et al. [236] sanitized speech for virtual assistants using vocal tract length normalization via compound frequency warping and adding Laplace noise for differential privacy. A follow-up [235] and work by Srivastava et al. [263] further analyzed its security.

- Shamsabadi et al. [253] aimed for theoretical privacy guarantees by adding differential privacy to pitch and context features in speaker transformation, encoding them via autoencoders and adding Laplace noise.

5.1.5 Evaluations

Most voice anonymization works evaluate de-identification by comparing recognition rates (identity or attributes) on unmodified and de-identified data, using machine learning models or human listeners. Common metrics include Equal Error Rate (EER), False Positive Rate (FPR), False Negative Rate (FNR), recall, precision, and F1 score [4]. Abou-Zleikha et al. [3] also used entropy and Gini index. The prevalence of EER suggests a focus on speaker verification, though speaker identification is often more appropriate for anonymization evaluation.

Utility loss is assessed by measuring naturalness (via Mean Opinion Score from human listeners) and intelligibility (via Word Error Rate, Phoneme Error Rate, or Short-Time Objective Intelligibility [270] from human listeners or machine learning models) [4]. A critical limitation observed is that most evaluations train recognition models on clear data and test on anonymized data, implicitly assuming the attacker is unaware of the anonymization. This leads to an overestimation of anonymization performance.

The VoicePrivacy challenge [279] is a notable initiative improving speaker anonymization methodology. It uses EER and log-likelihood-ratio cost function (Cllr) for speaker verifiability and Word Error Rate for intelligibility. Crucially, it includes retraining speaker verification systems with anonymized speech data to test against an informed attacker, a practice that has become popular since 2020.

Qian et al. [237] proposed a framework with a "p-leak limit" to quantify maximum privacy leakage per speaker. Zhang et al. [321] also provided a theoretical framework for privacy leakage risk and utility loss in speech data publishing.

5.2 Gait

Human gait refers to the characteristic pattern of limb movements during locomotion, encompassing various manners such as walking, running, or trotting [140]. Gait can be systematically broken down into individual gait cycles (see Figure 6 in the original document), which represent the shortest repetitive task within the entire locomotion process [267]. A gait cycle typically extends from a specific event of one foot (e.g., initial contact with the ground) until the same foot reaches that identical event again. Each cycle comprises a stance phase, where the foot is in contact with the ground, and a swing phase, where the foot is airborne. These two phases alternate for each foot, creating a unique and identifiable rhythm.

Gait has long been a subject of research in both computer science and psychology due to its utility as a behavioral biometric trait for individual identification. For example, Yovel et al. [315] demonstrated its importance in human recognition at a distance, and Pollick et al. [227] showed that humans can even infer gender from point-light displays of walkers, highlighting the inherent information encoded in gait patterns.

Gait recognition methods are diverse, adapted to various capture modalities. Wan et al. [290] categorize these methods based on cameras, accelerometers, floor sensors, and radars. Camera-based recognition is often classified as either model-based, which uses specific kinematic models of the walker (e.g., a pendulum model of the legs) to match individuals, or model-free, which processes the entire gait capture without an explicit model (e.g., averaging silhouettes over time to create a "gait energy image"). Accelerometer-based systems typically average gait data into feature representations, either by segmenting gait into cycles or using fixed-size frames.

5.2.1 Utility

Gait recordings offer significant utility across several domains:

- Medical Diagnosis: Gait patterns are crucial for the diagnosis of gait abnormalities and neurological conditions [140].
- Activity Monitoring: More casual applications include tracking daily steps [261] or recognizing specific activities.
- Video Quality/Naturalness: In video recordings, maintaining the naturalness and convincing appearance of the de-identified gait is important to avoid degrading the overall video quality or user experience [125].
- Biometric Authentication: As a behavioral biometric, gait is used for identity verification and continuous authentication.

5.2.2 Threat Space

The omnipresence of human gait in everyday life, coupled with the unintrusive nature of many capturing methods, makes it highly susceptible to privacy threats.

- Ease of Capture: Gait can be easily captured, often as a byproduct of other recordings (e.g., surveillance cameras), without the subject's active participation or awareness.
- Robustness to Obfuscation: Gait recognition has shown remarkable robustness to video quality degradation and minor obfuscation, making it suitable for surveillance systems [290].
- Identification: It is a strong identifier for individuals.
- Attribute Inference: Beyond identity, gait can reveal private attributes like gender [227], age, and various physiological conditions [282].
- Emerging Threats: With advancements in richer capturing methods like LiDAR [87] or affordable motion capture suits, the threat space for gait is expected to expand, allowing for even more detailed and potentially privacy-invasive inferences.

5.2.3 Anonymization Techniques

Here, we present gait anonymization methods found in the literature, organized by our taxonomy.

Random Perturbation

This involves introducing randomness into gait data to obscure identity.

- Hoang et al. [113] proposed a fuzzy commitment scheme based on Bose-Chaudhuri-Hocquenghem (BCH) codes for storing accelerometer gait templates. Reliable bits from binarized accelerometer data are XORed with a BCH-encoded secret key. While promising for false accept rates, the false reject rate needs improvement for user-friendliness.
- Matovu et al. [175] studied the influence of noise injection on accelerometer/gyroscope authentication systems by merging the original time series with a uniformly distributed noise time series.

Noise Injection

This category focuses on adding noise to gait data, often with the goal of obscuring identity while maintaining utility.

- Tieu et al. [276] developed a CNN-based method to mix the gait of a second person (noise gait) into the original gait in videos. Silhouettes are extracted and merged via a shared-weights CNN. They reported identification rates between 20% and 1% depending on view angle.
- In a follow-up, Tieu et al. [277] improved their method by generating the noise gait using a Generative Adversarial Network (GAN) and fusing it with a self-growing and pruning GAN (SP-GAN), achieving identification accuracy between 30% and 10%. They also proposed a method to colorize the resulting black-and-white silhouettes [278].
- Hanisch et al. [105] investigated adding Laplace noise to body positions in motion capture gait data. Their results indicated that effective anonymization was not possible without destroying utility (measured as naturalness via user study).
- Meng et al. [183] also showed that the noise level required for effective anonymization in motion data often destroys its utility.

Coarsening

This involves reducing the granularity or detail of gait data.

- Nair et al. [198] experimented with coarsening the frame rate, positional accuracy, and dimensionality of VR motion data. They found that while these techniques could reduce identification rates for individual motion sequences, they were not effective for anonymizing motion data on a per-session basis.

Feature Removal

This focuses on eliminating specific features from gait data that are highly indicative of identity.

- Jourdan et al. [131] proposed a feature removal approach for privacy-preserving activity recognition using accelerometers. They extracted temporal and frequency features and found that temporal features contributed more to identity recognition, while frequency features contributed more to activity recognition. By removing temporal features, they achieved a good trade-off (activity recognition from 96% to 87%, identification from 90% to 40%).
- Debs et al. [63] performed a similar feature

removal by transforming the signal using a short-time Fourier transformation and then randomly removing 10% to 90% of the data.

- Garofalo et al. [90] proposed a temporal convolutional network as a feature extractor, trained via adversarial training to minimize identity verification while maximizing attribute classification.

- Rougé et al. [246] developed an anonymization technique for accelerometer motion data. They extracted features via a short-time Fourier transform and trained a random forest classifier for action and identity recognition. Features important only for identification were then removed.

- Hanisch et al. [105] also tested removing body parts from gait motion capture data. They found gait data to be highly redundant, with identification success remaining close to 60% even when only head data was retained.

Continuous Conversion

This involves transforming gait data into a new, continuous representation.

- Blurring: This technique de-identifies persons in videos, including their gait.

  - Agrawal et al. [8] proposed exponential blur (treating video as 3D space and blurring via weighted average of neighbors) and Line Integral Convolution (LIC) (mapping bounding box to vector field for pixel calculation).

  - Ivasic-Kos et al. [125] applied a Gaussian filter to blur walker silhouettes, calculating a weighted average of neighboring pixel colors.

- Halder et al. [102] worked on gait anonymization in videos by extracting gait silhouettes, clustering them into key gait poses using k-means, and then using a GAN to generate new video sequences from the closest key pose.

- Moon et al. [191] investigated adversarial training for anonymizing 3D pose data, training models to maximize action recognition while minimizing identification. Their evaluation on ETRI-activity [127] and NTU60 [157] datasets showed high utility for action recognition and identification rates near chance.

- Nair et al. [196] proposed an adversarial approach for VR motion data anonymization using a Siamese architecture, training the model to achieve good action recognition and low identification by adding a random vector to the motion sequence input.

- Thapar et al. [275] considered anonymizing gait in egocentric videos. They learned identities from gallery videos via camera rotation signatures, then applied this signature to target videos to mix identities, increasing EER for person identification from ~20% to ~50%.

Continuous Conversion + Discrete Conversion

- Hirose et al. [112] proposed an approach combining continuous and discrete conversions for walkers in videos. They extracted silhouettes and gait cycles, transformed the silhouette via a deconvolutional neural network encoder into a silhouette code, and then converted this code using a k-same approach (weighted average of k-nearest neighbors). The gait cycle was transformed via a continuous, differentiable, and monotonically increasing function. A new video was generated by feeding the perturbed silhouette code and gait cycle into a convolutional neural network decoder. Their evaluation showed gait recognition dropping from ~100% to 4-29% depending on the model.

5.2.4 Evaluation

Gait de-identification is primarily evaluated using gait recognition systems or human observers, with recognition accuracy as the main metric. Other metrics include F1 score, Equal Error Rate (EER), or False Acceptance Rate (FAR) [4]. To assess utility loss, a wider variety of metrics are used, often quantifying the naturalness of the de-identified gait or the performance of other recognition tasks (e.g., activity recognition). Matovu et al. [175] used a "biometric menagerie" to observe de-identification influence on different user types in biometric authentication systems. A common limitation, similar to voice, is that many evaluations assume a weak attacker unaware of the anonymization.

5.3 Hand Motions and Gestures

The term "hand motions" serves as an umbrella category encompassing a variety of hand-related behavioral biometric factors, including handwriting, keystrokes, mouse movements, and hand gestures. These traits primarily differ in their recording methodology and the specific types of hand movements involved. Handwriting can be captured either offline (only the final written text) or online (real-time capture of hand movements during writing, typically with a digital pen). This survey focuses on the uniqueness of the writing style rather than the linguistic style (stylometry) of the text. In contemporary digital environments, handwriting has largely been supplanted by typing on keyboards, which itself constitutes a significant biometric factor, as individuals can be identified by the precise timings of their key presses. Similarly, the use of computer mice generates unique patterns through trajectories, speeds, and click events, serving as another biometric identifier. Lastly, hand gestures can be directly captured using optical tracking or accelerometer-based techniques.

Hand motion recognition involves various techniques tailored to these diverse capture modalities:

- Handwriting Recognition: For handwriting, the input sequence is often adjusted for its baseline, scaled to a normal writing style, and segmented to meet classifier requirements [223]. Recognition differs between "online

handwriting" (captured during writing) and "offline handwriting" (captured after writing is complete).

- Mouse Movement Recognition: This relies on features such as trajectory, speed, single, and double clicks performed with a mouse.

- Keystroke-based Hand Motion Recognition: Primarily based on timing differences between key-up, key-down, and key-hold events. Features often include differences between successive events (digraphs) or even three successive events (trigraphs) [327].

- Gesture Recognition: Can be categorized into 2D gestures (performed on flat surfaces like smartphones) and 3D gestures (performed in mid-air). Sherman et al. [255] used finger trajectories, resampling them with cubic spline interpolation to remove jitter and employing dynamic time warping for distance calculation between gestures.

5.3.1 Utility

The utility of hand motions is broad and varied:

- Handwriting: The primary utility is the readability of the resulting text by humans or computers. For signatures, the main purpose is identity identification and verification, where the unique style is paramount, and readability of the name is secondary.

- Input Modalities: For other hand motions (keystrokes, mouse movements, gestures), their utility lies in serving as precise and timely input modalities for computer systems [320], enabling effective human-computer interaction.

- Non-verbal Communication: Hand gestures additionally serve as a form of non-verbal communication [250].

5.3.2 Threat Space

The threat space for hand motions is significant due to their ubiquitous nature in daily tasks and digital device interaction.

- Unavoidable Capture: Our hands are constantly in use, and digital devices often record hand motions implicitly, frequently without user awareness.

- Identification: Numerous studies have demonstrated that individuals can be uniquely identified by their handwriting [223], keystroke dynamics [12], mouse movements [242], and gestures [305].

- Semantic Content: Hand motions can convey semantic meaning (e.g., typed text, passwords, private messages), making their content sensitive.

- Attribute Inference: Specific medical conditions, such as hand tremors in Parkinson's patients [128], can manifest in hand motions. Furthermore, hand motions can convey information about emotional states [264].

5.3.3 Anonymization Techniques

We present suitable methods for hand motion anonymization, noting that we did not find specific papers focusing solely on mouse movement anonymization within our defined scope.

Random Perturbation

This involves introducing randomness to disrupt unique hand motion patterns.

- Maiorana et al. [167] proposed a template protection method for online handwriting that segments a handwriting sequence and then randomly mixes the segments before convolution.

- Maiti et al. [168] applied a similar shuffling approach to prevent keystroke inference attacks via wrist-worn accelerometers, though without convolution. This approach was evaluated with only four participants.

- Vassallo et al. [288] also investigated keystroke permutation, focusing on utility reduction rather than privacy evaluation.

- Goubaru et al. [97] proposed a template protection scheme for online handwriting. They extract a pattern ID using a common template, XOR it with a secret encoded by an error-correcting code, and store the result as a template.

Noise Injection

This category involves adding noise to hand motion data.

- Migdal et al. [188] added delays to keystroke timings to anonymize typing patterns.

- Shahid et al. [252] proposed using the Laplace mechanism on the 2D coordinates of handwritten text to achieve local differential privacy.

Coarsening

This involves reducing the granularity or detail of hand motion data.

- Vassallo et al. [288] explored keystroke suppression to preserve typed text content in continuous authentication scenarios.

- Maiti et al. [168] focused on keystroke privacy by proposing two coarsening methods to prevent inference attacks via wrist-worn accelerometers: blocking access to accelerometer data when typing is detected, and reducing the accelerometer's sampling rate.

Discrete Conversion

These techniques transform hand motion data into discrete, often cancellable, representations.

- Sae-Bae et al. [248] proposed an online handwriting template protection scheme that decomposes signatures into histograms for authentication.

- Migdal et al. [189] presented a template protection scheme for multiple modalities, including keystrokes, by

combining IP addresses with keystroke information and computing a biohash.

- Leinonen et al. [151] investigated keystroke timing data anonymization using two rounding approaches that sort timings into buckets, effectively reducing identification from nearly 100% to below 10%.
- Vassallo et al. [288] explored substituting typed keys with random nearby keys to preserve content in continuous authentication.
- Figueiredo et al. [85] developed a modeling language for designing new gestures, allowing recognition on recording hardware without exposing clear data to applications. No privacy evaluation was performed.
- Mukojima et al. [194] designed a privacy-friendly gesture recognition system that illuminates the hand with a random pixel pattern and reconstructs the hand shape from the captured light using machine learning. Privacy protection was not evaluated.

Continuous Conversion

This involves transforming hand motion data into a new, continuous representation.

- Maiorana et al. [167] proposed two continuous conversions for online handwriting templates: a baseline conversion that segments and convolutes handwriting sequences based on a secret key, and a shifting transformation that applies a shift to the initial sequence.
- Malekzadeh et al. [171] proposed using two separate autoencoders for anonymizing gestures captured via IMU sensors. One autoencoder replaces sensitive sequences with generated neutral ones, while the second minimizes mutual information between data and user identity. Their approach reduced identification accuracy from 96% to 7%.
- Fan et al. [81] also proposed a two-encoder system (task encoding, identity encoding) for privacy-preserving motor intent classification from sEMG data. This system is trained adversarially to reduce identity recognition and increase action recognition.
- Saunders et al. [250] worked on sign language motion anonymization, transferring motions from one person to another. Their technique extracts pose features from source video and combines them with target appearance style (from an appearance distribution) to generate new images. Identity identification from hand motions alone was not evaluated.
- Xia et al. [298] proposed a second approach for sign language video anonymization. They estimate motion regions, use optical flow with a confidence map to encode source and driving video motions, and then generate anonymized video via an autoencoder. They used a loss function focused on hand and face motion differences to preserve sign language utility. Identity identification from hand motions was not evaluated.

5.3.4 Evaluation

Hand motion anonymization is primarily evaluated in the context of authentication, using metrics such as False Positive Rate (FPR), False Negative Rate (FNR), and Equal Error Rate (EER) to assess performance [4]. Additionally, recognition approaches are used to evaluate the accuracy of identity, age, gender, and handedness inference. Goubaru et al. [97] uniquely evaluated the randomness of template bits via occurrences and autocorrelation. Similar to other modalities, a common limitation is that EER might overestimate anonymization performance due to assumptions about attacker knowledge. More critical evaluation approaches are needed.

5.4 Eye Gaze

Eye gaze involves two primary types of movements: fixations and saccades. During visual tasks, such as reading (see Figure 7 in the original document), our eyes alternate between these two modes [4]. Fixations refer to periods where the visual focus is maintained on a single stimulus, allowing for information processing. In contrast, saccades are rapid, ballistic eye movements that quickly shift the gaze between fixations to reorient attention. Even during fixations, the eyes are not entirely still; they exhibit constant, involuntary micro-movements known as microsaccades [4].

Eye-tracking technologies are becoming increasingly accessible in both consumer and research markets. The most common tracking principle involves illuminating the eye with non-visible light sources (e.g., infrared) to generate a corneal reflection. Changes in these reflections are then sensed and analyzed to deduce eye rotation and gaze direction [72]. Eye-tracking hardware configurations vary widely, from cameras embedded in computers, smartphones, and virtual reality (VR) headsets to dedicated external devices or mobile eye-wear [72]. These sensors capture not only movement data related to fixations and saccades (e.g., speed, gaze angle, attention spots, scan path) but also additional features like pupil size variations and blink behavior. Combinations of these features provide rich information for developing eye-gaze-driven applications.

5.4.1 Utility

Eye movements have been studied for over a century across diverse research domains, yielding significant utility:

- Medical Field: Gaze patterns provide valuable information about cognitive and visual processing [16, 106], aiding in the diagnosis of various neurological and psychiatric diseases.
- Human-Computer Interaction (HCI): Eye gaze is used as an input modality to enhance accessibility, improve user experience, and enable adaptive system behavior [50, 169, 228]. This includes implicit

authentication, where stable, unique features of eye movement are leveraged to build biometric systems [136]. Behavioral eye biometrics have shown low EERs (e.g., 1.8% [76]) in authentication.

- Application-Specific Utility: The specific utility to be preserved depends on the underlying application, whether it's accurately predicting the next eye movement, diagnosing a mental disorder, detecting the focus of user attention, or recognizing a user for authentication.

5.4.2 Threat Space

Eye movement data is exceptionally rich in information, making it a significant target for malicious entities or curious service providers seeking to uncover sensitive user attributes or directly identify individuals.

- Identity and Attribute Inference: Beyond biometric identification, research has documented strong correlations between eye movement data and multiple disorders and mental conditions, including Alzheimer's [123, 304], schizophrenia [116, 152, 80], Parkinson's [148], bipolar disorder [89], mild cognitive impairment [304], multiple sclerosis [67], autism [31, 291], and psychosis [80].

- Cognitive and Emotional States: Pupil size is a known indicator of a person's interest [109] and a proxy for detecting cognitive load [146, 176].

- Soft Biometrics and Personality: Recent works have demonstrated that eye data can infer gender, age, and even personality traits [24, 147].

- Covert Information Leakage: Martinovic et al. [173] showed that manipulating images presented to users could cause their EEG signals (and by extension, potentially eye movements) to reveal private information such as bank card details, PINs, or living area.

- Increased Availability: The increasing availability of consumer eye-tracking devices and the proliferation of eye-gaze-driven applications create a significant and imminent privacy threat [6]. Hardware manufacturers like Apple have recognized this, restricting third-party access to eye-tracking information in their Vision Pro Headset.

The two primary threats to eye privacy are re-identification and attribute inference.

5.4.3 Anonymization Techniques

We found multiple recent proposals to protect eye movement data privacy, with many employing noise injection to achieve differential privacy (DP).

Random Perturbation

- David-John et al. [55] adapted a task-based marginal model for eye gaze. They built a distribution of values for each feature vector dimension and then randomly sampled new synthetic data from these distributions. The identification accuracy of the generated synthetic data was close to chance level.

Noise Injection

These methods add random noise to eye gaze data, often guided by differential privacy principles.

- Steil et al. [265] proposed a DP-based technique to protect eye movement data collected in a VR setting while users read different documents. The utility goal was accurate document type prediction, and privacy goals were to avoid gender inferences and protect against re-identification. They applied the exponential mechanism [74] to a database of users' eye features. Experiments showed partial utility preservation (~55%-70% document classification accuracy) while reducing gender inference accuracy to random guesses (~50%) at various noise levels.

- Bozkir et al. [33] evaluated two DP-based perturbations: the standard Laplacian Perturbation Algorithm (LPA) [73] and the Fourier Perturbation Algorithm (FPA) [239]. They proposed a modification to FPA that splits eye data into chunks before adding noise to reduce temporal correlations, which often require more noise for privacy. This modification achieved similar document type classification results to Steil et al. [265] while providing better privacy guarantees (more noise).

- Liu et al. [156] presented a DP-based solution to anonymize eye tracking data aggregated as a heatmap (attentional landscape). The privacy goal was to protect individual gaze maps while preserving the utility of the aggregated heatmap. Their experiments with random selection and additive noise (Gaussian, Laplacian) showed Gaussian noise was best for good privacy guarantees without visually distorting hotspots.

- David-John et al. [57] worked on protecting eye tracking data from VR/AR headsets. They proposed Gaussian noise injection, temporal down-sampling, and spatial down-sampling for one interface model. Gaussian noise injection was found most effective in reducing subject identification rates with high variance values. Wilson et al. [295] also proposed adding Gaussian noise to eye tracking data, showing similar results.

- Hu et al. [119] proposed a local differential private mechanism, Otus, for generating synthetic eye movement trajectories. Their technique separates the field of view into tiles, constructs a graph encoding gaze duration and transition probability, perturbs the graph using the Laplacian mechanism, and then sends it to a server. The server averages user graphs and generates new trajectories via random walks.

- Li et al. [153] proposed Kaleido, a plugin system for anonymizing eye gaze trajectories with differential privacy guarantees. They extended geo-indistinguishability [14] and w-event privacy [137] to account for areas of interest, noting protection primarily

against spatial information, not temporal. They defined an adaptive algorithm to allocate privacy budget per time window. Their results showed user identification reduced to near chance level, though utility was also close to chance.

Coarsening

This involves reducing the resolution or detail of eye gaze data.

- The temporal and spatial down-sampling techniques proposed by David-John et al. [57] are both coarsening-based. Temporal down-sampling showed only a small reduction in identification accuracy, while spatial down-sampling had a larger effect but required significant scaling.
- Wilson et al. [295] proposed a spatial down-sampling approach for the eye gaze angle, mapping 180◦ to 2,160 points and coarsening the gaze angle. This appeared more effective than temporal down-sampling.

Continuous Conversion

This involves transforming eye gaze data into a new, continuous representation.

- Wilson et al. [295] proposed smoothing eye gaze using a sliding window approach, showing that a sufficiently large window reduces identification rates.
- David-John et al. [55] applied k-anonymity to eye movements by grouping and averaging user trajectories. They showed significant drops in identification accuracy even with small k values, though high utility was questionable due to separate processing of feature vectors per task.
- In a follow-up, David-John et al. [56] proposed two synthetic data generation approaches for eye gaze: k-same synth (k-anonymity on fitted Gaussian mixture model parameters to generate fixations and saccades) and event-synth-PD (conditional variational autoencoder to generate new data with given characteristics). They showed event-synth-PD achieved plausible deniability and comparable privacy/utility to Kaleido.
- Fuhl et al. [86] performed eye gaze anonymization using an autoencoder combined with reinforcement learning. The autoencoder learns a latent representation, which a manipulation agent modifies to prevent, e.g., gender classification. A classifier then evaluates the manipulation, providing a loss for training the agent.

5.4.4 Evaluation

Proposals by Steil et al. [265] and Bozkir et al. [33] measure anonymization quality for attribute inference protection using classification accuracy for both the main task and the attribute inference task. For re-identification protection, they assume an attacker with prior knowledge (training classifiers on clean data and testing on anonymized data), also using accuracy. They report the privacy loss parameter (epsilon, ϵ) from DP theory, quantifying the maximum difference between data points of two individuals (smaller ϵ means better privacy).

Liu et al. [156] analyzed the privacy-utility trade-off of anonymized heatmaps using the correlation coefficient (CC) and mean square error (MSE) of noisy heatmaps under different privacy levels (ϵ). These metrics, combined with visual representations, help stakeholders decide acceptable noise levels.

Datasets play a crucial role. The largest available is GazeBaseVR [160], with 407 participants performing five tasks across multiple sessions using a VR headset. Steil et al. [265] collected data from 20 participants reading documents in VR, extracting 52 eye movement features related to fixations, saccades, blinks, and pupil diameter. This dataset is publicly available and used by Bozkir et al. [33]. EHTask [120] contains recordings of 30 people performing four eye gaze tasks in VR. DGaze [71] captures 43 people in five scenes. Liu et al. [156] used a synthetic simulated dataset for heatmap anonymization.

Beyond technical analysis, Steil et al. [265] uniquely considered user privacy concerns, conducting a large-scale survey (N=164) to explore willingness to share gaze data for different services. Their findings indicated discomfort with inferences (gender, race, sexual orientation) and objections to sharing data if such attributes could be leaked. Users generally agreed to share with governmental health agencies or for research but objected to sharing with companies. These insights are a vital step towards user-centered privacy design for behavioral data.

5.5 Heartbeat

An electrocardiogram (ECG) is a graphical representation of voltage over time, capturing the intricate electrical activities of the heart's muscle during depolarization and repolarization with each beat. As depicted in Figure 8 in the original document, a normal cardiac cycle in an ECG graph comprises a distinct sequence of waves:

- P-wave: Reflects the atrial depolarization process (contraction of the atria).
- QRS complex: Represents the ventricular depolarization process (contraction of the ventricles), typically the most prominent part of the waveform.
- T-wave: Denotes the ventricular repolarization (relaxation of the ventricles).

Other significant portions of the ECG signal include the PR, ST, and QT intervals, which represent specific time durations between these waves [323].

Like other biometric systems, ECGs are often converted into abstract, compressed representations, known as biometric templates, before being used for identification tasks. ECG biometric template methods are broadly classified based on the features exploited:

- Fiducial-based techniques: Utilize characteristic

points on the ECG signal (e.g., peaks and boundaries of P, QRS, and T-waves) to extract temporal, amplitude, envelope, slope, and area features [207].

- Non-fiducial-based methods: Do not rely on specific characteristic points, instead using techniques like autocorrelation coefficients, Fourier transforms, and wavelet transforms [207].

- Hybrid methods: Combine both fiducial and non-fiducial features.

5.5.1 Utility

ECG data finds critical applications in two primary domains: healthcare and biometric systems.

- Healthcare: ECGs are indispensable for the diagnosis of various heart diseases and monitoring cardiac health [158]. They are often integrated into stand-alone services or comprehensive e-health systems that provide real-time feedback to patients and hospitals, serving as warnings for impending medical emergencies or as monitoring aids during physical activities.

- Biometrics: ECGs are increasingly used for identification and authentication purposes due to their unique physiological characteristics [284].

5.5.2 Threat Space

Despite their utility, ECGs are classified as health data and are inherently sensitive, requiring robust protection under data-protection regulations.

- Health Condition Inference: The most direct privacy threat is the inference of a patient's physiological or pathological condition (e.g., arrhythmias, heart disease) [158, 323, 23, 190]. Such inferences could lead to adverse consequences like increased insurance premiums or discrimination in employment.

- Other Sensitive Inferences: Less commonly known, ECG data can also reveal sensitive information such as cocaine use [118] or psychological stress levels [224]. The ability to derive both desirable (healthcare-related) and sensitive (privacy-invasive) inferences from the same time-series data presents a significant practical dilemma.

5.5.3 Anonymization Techniques

We survey relevant privacy-protection techniques for ECG data.

Feature Removal

This involves extracting and removing specific features from ECG signals to obscure sensitive information.

- Kalai et al. [317] proposed a template protection scheme for ECG data. They compute the Discrete Cosine Transform (DCT) of the ECG signal's autocorrelation coefficients and then remove those with the lowest energy, forming the biometric template. Two keys are derived: one for authentication and a private key from the complete DCT. Zaghouani et al. [318] presented a similar approach using a quantization step after obtaining the DCT-template.

- Mahmoud et al. [166] proposed decomposing the ECG signal into its wavelet transform, eliminating low-frequency coefficients, and reconstructing the signal for release. Only authorized personnel with a secret key can reconstruct the original ECG. Privacy is evaluated via the Percentage Root Mean Square Difference (PRD) [172], which quantifies distortion between original and protected ECG.

- Djelouat et al. [70] proposed an approach based on Compressive Sensing (CS) [38], combining sampling and compression through random projections. This compresses the ECG signal at the time of sensing, reducing the need to store sensitive data on wearable devices. CS allows for good reconstruction of the original signal at the provider side under certain assumptions.

Continuous Conversion

These methods transform ECG data into a new, continuous representation to preserve privacy.

- Bennis et al. [23] proposed a simple k-anonymity scheme for ECG data. They transform the signal into the frequency domain, select the k closest neighbors, aggregate them into a new signal, and then transform it back into the time domain.

- Piacentino et al. [222] used a Generative Adversarial Network (GAN) to generate synthetic ECG data for training epilepsy monitoring systems, though no privacy evaluation of the synthetic data was performed.

- Jafarlou et al. [126] also used a GAN to generate anonymized ECG data samples. Their approach uses the original ECG sequence as GAN input and incorporates identification accuracy into the training loss. Their evaluation showed lower identification accuracies while still allowing arrhythmia detection.

- Nolin-Lapalme et al. [204] used a GAN for ECG anonymization, specifically aiming to generate sex-neutral ECG samples by including sex classification as part of the GAN loss.

- Pascual et al. [216] used a GAN to generate synthetic epileptic brain activity (EEG) data for training epilepsy monitoring systems, focusing on inter-ictal signals. Their results showed synthetic data reached identification rates close to chance level, but this was more of a pseudonymization as synthetic values for a patient could still be linked.

Random Perturbation + Noise Injection

- Chou et al. [46] addressed the vulnerability of CS to information-theoretic attacks [238] by proposing principal component analysis and SVD on a CS scheme where ECG data is encrypted at the wearable sensor with signal-dependent noise. They measured privacy as mutual

information between original and encrypted ECG, showing high classification accuracy while providing privacy beyond computational secrecy.

Discrete Conversion + Noise Injection

- Zare-Mirakabad et al. [319] aimed to publish ECG data with privacy guarantees by converting time series into symbolic representations using Symbolic Aggregate Approximation (SAX). They then built an n-gram model from the symbol string and ensured a minimum frequency of occurrence for each n-gram (similar to k-anonymity) by adding fake n-grams. Experimental results showed minimal information loss for k up to 20.

Continuous Conversion + Random Perturbation

- Chen et al. [44] and Wu et al. [296] addressed the revocability of ECG-based biometric templates, a crucial property for practical use. They proposed generating cancelable templates as random projections of an ECG data block, allowing distinct templates for the same biometrics. The goal is to make recovery of the original biometric infeasible. Re-identification rates of over 95% were reported using the multiple-signal classification algorithm [26].

- Hong et al. [117] proposed a template-free identification system to prevent privacy issues from compromised templates. They convert ECG data into images using spatial and temporal correlations and use deep learning to train a classifier, reporting over 90% identification rates.

Continuous Conversion + Noise Injection

- Sufi et al. [268] proposed building templates of P, QRS, and T waves through cross-correlations and obfuscating them with synthetically generated additive noise in a concatenated fashion (output of one wave obfuscation serves as input for the next). This creates noisy forms of the waves and templates, accessible to authorized personnel with a key for reconstruction, while unauthorized users only see the noisy signal.

- Huang et al. [121] proposed an authentication system protecting ECG templates with differential privacy in an interactive setting. An analyst queries for Legendre polynomial coefficients (used to fit and compress ECG), and Laplace noise is added based on sensitivity. The privacy parameter (ϵ) regulates the trade-off. However, they appeared to miscompute sensitivity, potentially invalidating their results.

- Saleheen et al. [249] investigated inferences from time series data by a dynamic Bayesian network adversary. When sensitive inferences (e.g., conversation, running, smoking) are likely, corresponding data segments are substituted with plausible non-sensitive data. They proposed a variation of differential privacy to bound leaked information. Utility loss was computed as the absolute difference in inference probability for non-sensitive states. While showing small utility loss for $\epsilon \in [0.05, 0.65]$, the solution is limited to dynamic Bayesian network adversaries and assumes pre-available time-series data, precluding real-time application.

5.5.4 Evaluation

ECG anonymization techniques are evaluated by measuring service functionality degradation using common machine learning metrics like precision, recall, and accuracy. Less frequently, Dynamic Time Warping (DTW) and Percentage Root Mean Square Difference (PRD) are used to assess similarity between original and protected time series [4]. Privacy levels are assessed through various notions, including membership inference attack accuracy, the ϵ parameter of differential privacy, mutual information between original and encrypted ECG, the probability of correct inferences on sensitive attributes (with and without protection), and a notion similar to k-anonymity. The MIT-BIH arrhythmia database [190], containing ECG samples from 47 individuals, and the Physikalisch Technische Bundesanstalt (PTB) Database [32] are commonly used datasets.

5.6 Brain Activity

Brainwaves are measurable electrical impulses generated by the complex interactions of billions of neurons within the human brain. Since Hans Berger first recorded the human electroencephalogram (EEG) in 1924 [101], significant advancements have been made in both hardware devices for measuring brain activity and analytical techniques for processing these signals. Brainwave measurement technologies are broadly categorized as invasive and non-invasive methods. Invasive methods involve directly implanting electrodes near the brain's surface to record signals within the cortex [133]. These are typically restricted to critical clinical applications due to their inherent risks. In contrast, non-invasive methods are more frequently used and applicable to a wider range of fields beyond medicine, such as brain-computer interfaces (BCIs). The most portable and common non-invasive technique is EEG, which records electrical activity through sensors placed on the scalp surface.

An EEG signal is a composite of different brainwave types, each oscillating at distinct frequencies and carrying different kinds of information about the brain's current state [10]. Researchers have attempted to associate specific mental states with these brainwave types. Table 1 in the original document provides a summary of the most important wave types, their respective frequencies, originating locations in the brain, and associated mental states:

- Gamma (γ): 30-100 Hz, Somatosensory cortex, associated with active information processing and strong responses to visual stimuli [2].

- Beta (β): 13-30 Hz, Both hemispheres, frontal lobe, associated with increased alertness, anxious thinking, and focused attention.

- Alpha (α): 8-13 Hz, Posterior regions, both hemispheres; high amplitude waves, associated with resting, eyes closed, no attention [139], and is often the most dominant rhythm.

- Theta (θ): 4-8 Hz, No special location, associated with idling, dreaming, imagining, quiet focus, and memory retrieval.

- Delta (δ): 0.5-4 Hz, Frontal regions; high amplitude waves, associated with dreamless and deep sleep, and unconsciousness.

BCI technologies primarily operate on continuous EEG data recordings (time series data). However, many applications also rely on extracting time-locked brain variations that occur in response to external stimuli. These variations, known as event-related potentials (ERPs) [269], are widely used to detect neurological diseases. In both continuous EEG and ERP-based applications, features are computed from the brainwave data, which can belong to the time and/or frequency domain and span one or multiple channels. Commonly used features include Autoregressive coefficients, Fourier transforms, and Wavelet transforms.

5.6.1 Utility

The utility to be preserved when processing brainwave data is highly dependent on the specific application:

- Clinical Applications: In clinical settings (e.g., for diagnosis or controlling brain-controlled prostheses), the raw EEG information might be critically needed for proper diagnosis or safe operation. Regulations like the HIPAA Privacy Rule [111] are typically in place to protect this personal identifiable information.

- Less-Regulated Fields: In applications outside critical medical contexts, the need for full raw EEG data may not be justified. Prominent EEG applications include:

  - User Authentication: Recognizing the user based on their unique brainwave patterns [99].

  - Personalization of Gaming Experiences: Adapting game dynamics based on a player's cognitive state.

  - Brain-Controlled Interfaces: Providing responsive interfaces for various tasks.

In these cases, the preserved utility should be sufficient to ensure a useful application, meaning accurate user recognition, effective personalized options, and responsive interfaces with tolerable error rates that do not compromise the service's security or usability.

5.6.2 Threat Space

Brain activity is an incredibly rich source of information, making it highly susceptible to privacy threats.

- Unique Identification: Brainwaves possess unique characteristics that allow for individual identification, leading to the development of several biometric systems based on brain activity [99].

- Correlation with Sensitive Attributes: The acquisition of EEG signals raises significant privacy concerns because brainwaves correlate with mental states, cognitive abilities, and various medical conditions [269].

- Side-Channel Attacks: Martinovic et al. [173] demonstrated that by manipulating images presented to users, their EEG signals could inadvertently reveal highly private information, such as bank card numbers, PINs, areas of living, or even whether they knew a particular person. This highlights the potential for sophisticated side-channel attacks.

5.6.3 Anonymization Techniques

A growing number of anonymization techniques for brain activity data rely on machine learning methods, with Generative Adversarial Networks (GANs) and adversarial perturbation schemes being particularly dominant. The increasing availability of EEG datasets has spurred research in this area.

Feature Removal

This involves selecting and removing specific features from EEG data to conceal sensitive information.

- Matovu et al. [174] explored reducing private information leakage from EEG user authentication templates, specifically targeting the inference of alcoholism by an unscrupulous database administrator. Their attribute protection mechanism hypothesized that different template designs (features, channels, frequencies) impact the amount of non-authentication information inferred. They demonstrated this by selecting two templates and calculating their predictive capability for authentication and alcohol consumption.

- Yao et al. [310] proposed using GANs [96] to filter sensitive information from EEG data, aiming to reduce alcoholism inference while maintaining utility for detecting mental tasks (e.g., predicting visual stimulus). Their GAN-based filter uses deep neural networks for domain transformation, translating EEGs from a source domain with both desired and privacy-related features to a target domain with only desired features. Their results showed a significant reduction in classifying EEG sequences from alcoholic users as such (from 90.6% to 0.6%), with only a minor drop in mental task classification accuracy (4.2%). However, the initial mental task classifier accuracy was not strong, raising questions about its generalizability.

Continuous Conversion

These methods transform EEG data into a new, continuous representation.

- Pascual et al. [216] used a GAN to generate synthetic EEG data for training epilepsy monitoring systems, addressing the privacy concerns of sharing large medical EEG datasets. Their generator, a convolutional autoencoder, translates a latent code into an ictal (seizure)

sample, which the discriminator compares to real ictal samples. Their results showed synthetic data reaching identification rates close to chance level, even with few patients, though this was more of a pseudonymization as synthetic values for a specific patient could still be linked.

- Bethge et al. [25] proposed "privacy encoders" to remove sensitive information from brain activity data streams before classification. They trained a convolutional neural network as an encoder using Maximum Mean Discrepancy (MMD) between different encoded datasets as a loss function, aiming for domain-invariant representations. Testing on four datasets, they showed classification of data origin dropping from 99% to 52%, while emotion classification only reduced from 51% to 49%. The preservation of subject identity remains an open question.

- Meng et al. [181] proposed a similar approach, learning a perturbation vector added to the EEG signal via an adversarial scheme, using an action classifier for utility and a biometric recognition system for privacy.

- Singh et al. [258] proposed another adversarial approach using an autoencoder for the transformation.

Continuous Conversion + Noise Injection

- Debie et al. [62] also used a GAN to generate synthetic data from original EEG data. Their approach differs by employing differentially private stochastic gradient descent on the discriminator, reducing the influence of individual data points on gradient computations. Evaluated on the Graz dataset A (9 subjects), their results showed good utility preservation, but no additional privacy evaluation was performed.

5.6.4 Evaluation

EEG anonymization works, similar to gait anonymization, evaluate inference protection quality by comparing prediction accuracy for the protected attribute before and after modifying EEG data. Typical machine learning metrics like accuracy, false positive rates, and false negative rates are used. Utility loss is assessed by measuring the reduction in classification accuracy when using original versus anonymized EEG data.

Various EEG datasets are used for evaluations. The largest is the Temple University Hospital EEG data corpus [206] (579 subjects), followed by the BCI2000 dataset [251] (106 subjects). The dataset by Arias et al. [15] (56 people) was specifically recorded for authentication. A unique dataset is the SUNY medical dataset [134, 203], containing EEG data from alcoholic and control subjects viewing visual stimuli. Several smaller datasets also exist [114, 266, 293].

DISCUSSION

The systematic review of anonymization techniques for behavioral biometric data reveals several overarching characteristics, commonalities, differences, and persistent challenges across various traits.

6.1 Commonalities and Differences Across Traits

All reviewed behavioral biometric traits share a fundamental characteristic: they are captured as time-series data, tracking changes in the trait over time. This temporal dimension is crucial and often implicitly encodes identifying and sensitive information.

A significant distinction among traits lies in their observability:

- Overt Traits: Many traits, such as gait, hand motions, voice, and eye gaze, are overt, meaning they can be observed from a distance and do not necessarily require the active participation or explicit consent of the subject. These traits are frequently captured as a byproduct of other recordings, such as video surveillance or audio recordings in public spaces. This inherent observability makes them particularly vulnerable to passive data collection and subsequent privacy invasions.

- Secret Traits: In contrast, traits like EEG and ECG are secret; they can primarily only be recorded by directly attaching specialized sensors to the subject. This typically implies a more active and intentional participation from the individual, and often occurs in controlled environments (e.g., medical settings, research labs). While this reduces the risk of passive, unnoticed collection, the data itself is highly sensitive.

In terms of research attention, we observed a disparity: voice anonymization is an established and extensively researched field, boasting a large body of literature and dedicated initiatives like the VoicePrivacy Challenge. Conversely, other behavioral biometric traits, particularly EEG, have received comparatively less attention in the context of anonymization. For some traits, such as touch, thermal, and lip-facial expressions, we found no anonymization mechanisms that met our rigorous inclusion criteria, indicating significant gaps in current research.

The utility of these traits is highly diverse and often unique to each trait and its specific application. It ranges from the critical need for naturalness in de-identified motion (e.g., for realistic avatars in VR or for video quality) to the paramount importance of intelligibility in anonymized speech (e.g., for voice assistants).

Despite these differences in utility and observability, the threat space across traits exhibits striking similarities due to the pervasive nature of digital capturing devices. Wearables and mobile devices, being continuously attached to or carried by individuals, are of particular concern as they enable continuous capture of behavioral data. As our literature review has consistently shown, all surveyed traits can be exploited for both identity inference and attribute inference. This inherent vulnerability can be abused for a wide variety of privacy threats, including pervasive surveillance, sophisticated identity theft, and the unauthorized inference of private attributes (e.g., health conditions, emotional states, or personal interests).

Consequently, the fundamental privacy goals – identity protection and attribute protection – remain consistent across all behavioral biometric traits. However, voice stands out with an additional privacy goal: making the semantic content of speech unintelligible to protect against content-based privacy breaches.

6.2 Analysis of Anonymization Techniques

Our analysis of the surveyed techniques (summarized in Table 2 and Table 3 in the original document) reveals that continuous conversion methods are the most prevalent, followed by feature removal and noise injection. Random perturbation and discrete conversion (mostly for template protection) are less common, with coarsening being the least represented category.

Several critical observations emerged regarding these categories:

- Removal Methods: While removal methods (coarsening, feature removal) aim to be irreversible, the high redundancy often present in behavioral biometric data poses a challenge. This redundancy means that even if certain features are removed, an intelligent attacker might still be able to reconstruct or infer the missing data from the remaining correlated information. The effectiveness of removal, therefore, needs careful re-evaluation against stronger adversaries.

- Conversion Methods: For conversion techniques, we frequently observed that the parameter space for the anonymizations is often rather small. If an attacker knows the anonymization technique, they might be able to link clear and anonymized data by brute-forcing the limited parameter space. This highlights a need for larger, more complex, or dynamically changing parameter spaces to enhance security. The general reversibility (or rather, the practical irreversibility) of conversion techniques still requires more rigorous evaluation.

- Noise Injection Techniques: A significant challenge for noise injection methods is the strong dependencies, both temporal and physiological, within behavioral biometric data. These dependencies can be exploited by an attacker to filter out the injected noise, thereby recovering the original, sensitive information. This suggests that simple noise addition might be insufficient against sophisticated filtering algorithms.

- Differential Privacy (DP): While several methods claim to provide differential privacy guarantees, we observed a critical limitation: none of them can be used continuously over time without eventually compromising user privacy. This is due to the fundamental property of sequential composition in differential privacy [179], where the privacy budget is necessarily finite and is consumed with each query or release. This contradicts the intended use of many behavioral biometric applications, such as continuous monitoring in healthcare or ongoing authentication services, which are clearly not single-use. More research is urgently needed on how to effectively apply differential privacy to continuous behavioral data streams without exhausting the privacy budget. The use of related privacy notions intended for continuous observations (e.g., w-event differential privacy [137]) may offer promising avenues.

- Temporal Aspect Neglect: A striking observation was that most methods do not explicitly manipulate the temporal aspect of their data. Notable exceptions include Hirose et al. [112] and Maiti et al. [168]. Given that all behavioral traits inherently manifest as time-series data, exploring techniques that perturb the temporal order or alter time differences between events could lead to more general and robust anonymization techniques applicable across multiple traits.

- Intrinsic Attribute Anonymization: For attribute protection, anonymizing intrinsic attributes (e.g., age, sex, health conditions) proves particularly difficult because it is often unclear which specific parts of the behavioral data are relevant for these attributes. In this regard, generative machine learning approaches (e.g., GANs) appear promising, as these models can learn the complex, intrinsic dependencies between data and attributes, potentially allowing for the generation of synthetic data that retains utility but is decoupled from sensitive attributes.

- User Privacy Awareness: We also noticed a significant lack of even a basic understanding of users' privacy awareness and concerns regarding behavioral privacy in the literature. To design truly effective and user-accepted protection techniques, it is necessary to conduct more user studies to understand their needs and requirements.

6.3 Evaluation Methodology Critique

The evaluation methodology employed across different traits and methods in the surveyed literature is remarkably similar, yet often rudimentary. Typically, an inference or recognition system is applied to both the clear (original) and the anonymized data, and the difference in accuracy is reported. Crucially, this is often done without retraining the inference system on the anonymized data. This approach implicitly assumes a weak attacker model – one who is unaware that the data has been anonymized and therefore uses models trained on clean data. This assumption leads to an overestimation of the anonymization performance, as a sophisticated, informed attacker would retrain their models on the anonymized data or even attempt to reverse-engineer the anonymization process.

A notable exception to this trend is the more recent work in voice anonymization, which increasingly relies on the benchmarking framework of the VoicePrivacy Challenge [279]. This initiative explicitly includes retraining speaker verification systems with anonymized speech data to test against an informed attacker, providing a more realistic assessment of privacy guarantees. This demonstrates how

community-driven initiatives can significantly improve the overall evaluation methodology and comparability within a research field.

Furthermore, we found a scarcity of formal approaches [237, 321] to quantify the privacy of behavioral biometric anonymization methods; most evaluations rely on empirical privacy estimations. Another issue is that the evaluation methodology often too closely mirrors the evaluation of recognition systems (which aim to infer persons in large datasets with potentially poor data quality), whereas an anonymization method should ideally also work effectively for smaller group sizes with high data quality, where re-identification risks are different.

The lack of readily available and diverse datasets (as highlighted in Table 4 in the original document) is a significant impediment, particularly for less-researched behavioral biometric traits. Without standardized, large-scale, and publicly accessible datasets, it remains challenging to compare different anonymization techniques rigorously and to foster robust research in these areas.

6.4 Future Directions

Based on the identified gaps and challenges, several promising avenues for future research emerge:

- Robust Differential Privacy for Continuous Data: Developing and refining differentially private mechanisms that can effectively handle continuous streams of behavioral data without rapidly depleting the privacy budget. This might involve exploring alternative DP notions or novel budget allocation strategies.
- Temporal Anonymization: Investigating anonymization techniques that explicitly manipulate the temporal aspects of behavioral time-series data (e.g., altering event timings, shuffling temporal segments) to enhance privacy across multiple traits.
- Advanced Generative Models: Further research into generative machine learning approaches (e.g., advanced GAN architectures, variational autoencoders) for creating high-fidelity, synthetic behavioral data that is statistically representative of the original but provably unlinkable to individuals. This is particularly promising for anonymizing intrinsic attributes.
- Stronger Attacker Models in Evaluation: Standardizing evaluation methodologies to consistently incorporate informed attacker models, including scenarios where the attacker knows the anonymization technique and can retrain their inference models on anonymized data. This would provide a more realistic assessment of privacy guarantees.
- Community-Driven Evaluation Frameworks: Expanding the model of the VoicePrivacy Challenge to other behavioral biometrics. Establishing common benchmarking frameworks, shared datasets (where feasible and ethical), and standardized metrics would significantly increase comparability and rigor in privacy and utility evaluations across the field.
- User-Centered Privacy Design: Conducting more extensive user studies to understand privacy perceptions, concerns, and preferences regarding behavioral data collection and anonymization. This user-centric approach is vital for designing protection techniques that are not only technically sound but also socially acceptable and usable.
- Anonymization for Digital Twins and Mixed Reality: As the concept of digital twins and immersive mixed reality environments gains traction, an open question is whether independently anonymizing individual behavioral traits is sufficient to create privacy-friendly digital twins. Research is needed on holistic anonymization strategies for combined, multi-modal behavioral data in these complex environments.
- Real-time Applicability: Many current anonymization techniques are computationally intensive. Future work should focus on developing low-latency, real-time anonymization solutions, especially crucial for interactive applications like VR/AR and continuous authentication.
- Explainable Anonymization: Developing methods that can explain how and why certain features or data points contribute to identity or attribute leakage, and how the anonymization process mitigates this. This could aid in designing more targeted and effective privacy-preserving transformations.

Concluding Remarks

The increasing ubiquity of sensors and the growing sophistication of data analytics have made behavioral biometrics a powerful yet double-edged sword. While offering unprecedented opportunities for seamless interaction, enhanced security, and personalized services, the inherent richness of behavioral data simultaneously poses significant threats to individual privacy. Our comprehensive literature review has underscored the critical importance of anonymizing behavioral biometric data to protect people's identities and sensitive attributes.

We have identified a diverse array of behavioral traits, each with unique characteristics and vulnerabilities. Our proposed taxonomy, classifying anonymization techniques by the type of data transformation they perform (random perturbation, noise injection, coarsening, feature removal, and continuous/discrete conversion), provides a structured framework for understanding the current landscape. While voice anonymization stands as a relatively mature research field with numerous insights and a robust evaluation framework (e.g., VoicePrivacy Challenge), many other behavioral biometric traits have received comparatively less attention in terms of dedicated anonymization solutions. Their effective protection, therefore, remains a

significant open research question.

A critical finding from our analysis is the widespread use of rudimentary evaluation methodologies, often relying on the assumption of a weak attacker who is unaware of the anonymization. This leads to an overestimation of privacy guarantees. Improving the evaluation methodology by adopting more rigorous, informed attacker models and establishing community-driven benchmarking initiatives is paramount for advancing the field. Furthermore, we observed that the temporal aspect of behavioral data, which is inherently time-series in nature, has largely been neglected in anonymization approaches. Few techniques explicitly perturb temporal order or time differences, representing a missed opportunity for developing more generalizable and robust anonymization strategies across multiple traits.

In conclusion, the journey towards truly privacy-preserving behavioral biometric systems is ongoing and complex. It necessitates continued innovation in anonymization techniques, a shift towards more rigorous and realistic evaluation methodologies, and a deeper understanding of user privacy concerns. By fostering interdisciplinary collaboration and addressing the identified challenges, we can strive to harness the immense potential of behavioral biometrics while safeguarding the fundamental right to privacy in our increasingly data-driven world.

REFERENCES

Alberto Abad, Alfonso Ortega, António Teixeira, Carmen García Mateo, Carlos D. Martínez Hinarejos, Fernando Perdigão, Fernando Batista, and Nuno Mamede (Eds.). 2016. Advances in Speech and Language Technologies for Iberian Languages (Lecture Notes in Computer Science, Vol. 10077). Springer International Publishing. DOI:https://doi.org/10.1007/978-3-319-49169-1

Mohammed Abo-Zahhad, Sabah Mohammed Ahmed, and Sherif Nagib Abbas. 2015. State-of-the-art methods and future perspectives for personal recognition based on electroencephalogram signals. Biometrics 4, 3 (Sept. 2015), 179–190. DOI:https://doi.org/10.1049/iet-bmt.2014.0040

Mohamed Abou-Zleikha, Zheng-Hua Tan, Mads Graesboll Christensen, and Soren Holdt Jensen. 2015. A discriminative approach for speaker selection in speaker de-identification systems. In European Signal Processing Conference (EUSIPCO). IEEE, 2102–2106. DOI:https://doi.org/10.1109/eusipco.2015.7362755

Richard A. Abrams, David E. Meyer, and Sylvan Kornblum. 1989. Speed and accuracy of saccadic eye movements: Characteristics of impulse variability in the oculomotor system. J. Exp. Psychol. Hum. Percept. Perform. 15, 3 (1989), 529. DOI:https://doi.org/10.1037/0096-1523.15.3.529

Christopher Ackad, Andrew Clayphan, Roberto Martinez Maldonado, and Judy Kay. 2012. Seamless and continuous user identification for interactive tabletops using personal device handshaking and body tracking. In Extended Abstracts on Human Factors in Computing Systems. ACM, 1775–1780. DOI:https://doi.org/10.1145/2212776.2223708

Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics emerging: The story of privacy and security perceptions in virtual reality. In Symposium on Usable Privacy and Security. USENIX, 427–442.

Ayush Agarwal, Amitabh Swain, and S. R. Mahadeva Prasanna. 2022. Speaker anonymization for machines using sinusoidal model. In IEEE International Conference on Signal Processing and Communications (SPCOM). 1–5. DOI:https://doi.org/10.1109/SPCOM55316.2022.9840792

Prachi Agrawal and P. J. Narayanan. 2011. Person de-identification in videos. Trans. Circ. Syst. Video Technol. 21, 3 (Mar. 2011), 299–310. DOI:https://doi.org/10.1109/tcsvt.2011.2105551

Hafiz Shehbaz Ali, Fakhar ul Hassan, Siddique Latif, Habib Ullah Manzoor, and Junaid Qadir. 2021. Privacy enhanced speech emotion communication using deep learning aided edge computing. In International Conference on Communications Workshops. IEEE, 1–5. DOI:https://doi.org/10.1109/ICCWorkshops50388.2021.9473669

Abdulaziz Almehmadi and Khalil El-Khatib. 2013. The state of the art in electroencephalogram and access control. In Conference on Communications and Information Technology (ICCIT). IEEE, 49–54. DOI:https://doi.org/10.1109/iccitechnology.2013.6579521

Ranya Aloufi, Hamed Haddadi, and David Boyle. 2020. Privacy-preserving voice analysis via disentangled representations. In Conference on Cloud Computing Security Workshop. ACM, 1–14. DOI:https://doi.org/10.1145/3411495.3421355

Arwa Alsultan and Kevin Warwick. 2013. Keystroke dynamics authentication: A survey of free-text methods. Int. J. Comput. Sci. Issues 10, 4 (2013), 1.

Abdulaziz Alzubaidi and Jugal Kalita. 2016. Authentication of smartphone users using behavioral biometrics. Commun. Surv. Tutor. 18, 3 (2016), 1998–2026. DOI:https://doi.org/10.1109/comst.2016.2537748

Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geoindistinguishability: Differential privacy for location-based systems. In ACM Conference on Computer and Communications Security (CCS). 901–914.

Patricia Arias-Cabarcos, Thilo Habrich, Karen Becker, Christian Becker, and Thorsten Strufe. 2021. Inexpensive brainwave authentication: New techniques and insights on user acceptance. In 30th USENIX Security Symposium

(USENIX Security). 55–72.

A. Terry Bahill, Michael R. Clark, and Lawrence Stark. 1975. The main sequence, a tool for studying human eye movements. Math. Biosci. 24, 3-4 (Jan. 1975), 191–204. DOI:https://doi.org/10.1016/0025-5564(75)90075-9

Fahimeh Bahmaninezhad, Chunlei Zhang, and John Hansen. 2018. Convolutional neural network based speaker deidentification. In Speaker and Language Recognition Workshop. ISCA, 255–260. DOI:https://doi.org/10.21437/odyssey.2018-36

Dustin Bales, Pablo A. Tarazaga, Mary Kasarda, Dhruv Batra, A. G. Woolard, J. D. Poston, and V. V. N. S Malladi. 2016. Gender classification of walkers via underfloor accelerometer measurements. Internet Things J. 3, 6 (Dec. 2016), 1259–1266. DOI:https://doi.org/10.1109/jiot.2016.2582723

Salil Partha Banerjee and Damon Woodard. 2012. Biometric authentication and identification using keystroke dynamics: A survey. J. Pattern Recog. Res. 7, 1 (2012), 116–139. DOI:https://doi.org/10.13176/11.427