International Journal of Information Security (2026) 25:76
https://doi.org/10.1007/s10207-026-01231-3

REGULAR CONTRIBUTION

Through the looking glass: eye tracking biometrics and the loss of anonymity in extended reality

Antonios Patergianakis[1] · Costas Lambrinoudakis[1]

Received: 20 September 2025 / Accepted: 8 February 2026 / Published online: 1 April 2026

Abstract

The rapid growth of extended reality (XR), encompassing technologies such as virtual reality (VR) and augmented reality (AR), introduces unprecedented challenges in safeguarding personal user privacy. This study explores the extent to which XR eye-tracking biometric technologies can compromise user anonymity and reveal private identities. Leveraging the GazeBaseVR dataset (Lohr et al. in Sci Data 10(1): 177, 2023), which comprises extensive gaze data from hundreds of users, we extracted a comprehensive set of gaze features–including fixation, saccade, and Savitzky–Golay–based velocity statistics, along with position- and task-related measures–that capture the distinctive dynamics of individual eye movements. Using these features, we trained a simple Multi-Layer Perceptron (MLP) classifier to identify individuals based on their unique eye movement patterns. Our model achieved a high identification accuracy of 96.61% on identifying users while watching a video in a VR environment, underscoring the severe privacy risks posed by XR technologies in their collection and processing of biometric data. Beyond presenting these findings, this research highlights the broader implications of XR eye-tracking on user privacy and advocates robust solutions to address these concerns. We urge technologists, policymakers, and privacy advocates to collaborate in establishing comprehensive regulations and privacy-preserving mechanisms to mitigate the potential misuse of XR biometric data. This work aims to inspire further interdisciplinary research to ensure that technological innovation does not come at the expense of fundamental privacy rights.

Keywords Privacy · Eye-Tracking · Extended Reality · AI · User Identification · Biometric Data Protection

1 Introduction

Eye-tracking biometrics within extended reality (XR) environments present a complex landscape of privacy risks that require careful examination [2]. As users navigate these immersive digital spaces, the collection of detailed biometric data - particularly through eye tracking - introduces both promising opportunities and significant privacy challenges. We are currently witnessing a surge in eye tracking technologies [3], as highlighted by industry's recent investments in XR wearables, such as Meta's new AR glasses, the Orion, which feature advanced eye tracking capabilities as their main modus operandi. This increasing use of eye tracking in XR environments undoubtedly increases user interactions but also raises privacy issues, as it enables unregulated, large-scale user identification, it may be used to disclose a wide range of personal information, including sensitive details about individuals [4], and may also lead to potential surveillance infringements [5].

The primary objective of this research is to demonstrate the considerable risks associated with user de-anonymization through the analysis of eye movement data, especially when these data are collected within an XR environment. To achieve this, we employed the comprehensive GazeBaseVR dataset [1], which is well-suited for tasks involving eye movement analysis. Subsequently, we designed and trained a Multi-Layer Perceptron to accurately identify individuals based on the unique patterns of their ocular activity, thereby underscoring the potential privacy risks linked to the collection of such biometric data. Our research demonstrates the potential for significantly enhancing user identification accuracy by combining high-resolution datasets like GazeBaseVR with advanced data processing techniques. For

✉ Antonios Patergianakis
apat@unipi.gr

Costas Lambrinoudakis
clam@unipi.gr

1 Department of Digital Systems, University of Piraeus, Piraeus, Greece

instance, the application of the Savitzky-Golay filter [6] refines velocity calculations, extracting meaningful metrics that improve identification precision. In conjunction with further feature extraction regarding gaze velocity and positional statistics, and with the use of appropriate machine learning classification techniques, we have been particularly successful in user identification. To our knowledge, this is the first study to achieve such a high level of user identification accuracy using eye movement data from a large sample of approximately four hundred participants.

This achievement not only highlights the privacy risks inherent in eye-tracking biometrics but also underscores the urgent need for developing comprehensive strategies and effective privacy-preserving measures to address these risks in extended reality settings. Given the breadth and depth of potential data misuse - ranging from personal identification to profiling - future work should focus on multidisciplinary approaches that integrate technical, regulatory, and ethical considerations. This comprehensive approach will ensure that the integration of eye-tracking technologies into extended reality systems prioritizes user privacy and security while continuing to drive innovation.

Our research contributes to this important discussion by demonstrating a 96.61% identification rate based solely on eye movement data collected during a VR video-watching task, involving over four hundred individuals. These findings highlight the real-world feasibility of eye-tracking as a biometric identifier within extended reality environments. Given the sensitivity of such biometric data, it is imperative to update ethical standards surrounding the collection, storage, and usage of eye-tracking information. Prioritizing informed consent, employing anonymization protocols, and implementing privacy-preserving mechanisms are essential steps toward ensuring that future technological advancements do not compromise individual privacy.

To address the limitations of prior research—which has often been constrained by small cohorts (fewer than 100 subjects) or relied on computationally intensive models that yield lower identification rates—this work bridges a critical gap by validating biometric identification at scale. We leverage the GazeBaseVR dataset to demonstrate that high-accuracy identification (96.61%) is achievable across a diverse population of over 400 users. Central to our approach is a robust feature processing pipeline that employs Savitzky-Golay filtering to refine velocity estimates, from which we extract a comprehensive vector of velocity, position, and fixation statistics. By proving that a lightweight Multi-Layer Perceptron (MLP) can leverage these engineered features to distinguish individuals during passive video-viewing tasks, we highlight a profound practical implication: the barrier to entry for effective surveillance in the Metaverse is dangerously low, requiring neither sophisticated hardware nor active user cooperation.

The paper is organized as follows: Section 2 provides the necessary background by reviewing previous research on eye-tracking biometrics and identifying the gaps that our study addresses. Section 3 outlines our experimental setup, describing how we processed the GazeBaseVR dataset and trained our Multi-Layer Perceptron (MLP) model to recognize users. Section 4 reports our findings, with a specific focus on how identification accuracy varies between passive activities, like watching a video, and active tasks. Section 5 discusses practical ways to protect user privacy, suggesting strategies that balance the need for anonymity with the utility of eye-tracking features. Finally, Section 6 summarizes our main conclusions, addresses the limitations of our current approach, and offers recommendations for future work in the field.

2 Related work

Research on privacy risks associated with eye-tracking biometrics has consistently underscored the threat of user identification in extended reality (XR) environments [2, 4]. In our review, we specifically focus on studies that employed datasets collected through VR-based eye-tracking systems, ensuring that comparisons reflect conditions most relevant to XR applications (Table 1).

Prior work has demonstrated that eye movement features such as saccades, fixations, and task durations can enhance user identification accuracy [7], but these same features raise concerns about the potential misuse of biometric information. Some studies have been constrained by small sample sizes–sometimes involving fewer than one hundred participants [7, 8]–while others have not achieved the high identification accuracy demonstrated in this work [1, 9–11]. For instance, [7] emphasize the role of fixation and saccade durations in differentiating users, while [4] highlight the broader risks posed by large-scale eye-tracking data collection. Collectively, these studies show both the promise of eye-tracking for biometric applications and the privacy risks it entails.

The exploration of eye movements as biometric identifiers has progressed through several distinct phases of maturity, yet prior efforts have frequently grappled with a trade-off between dataset scale and identification accuracy. Early investigations, such as those by Monaco et al. using the EMVIC 2014 dataset, established the theoretical feasibility of the concept but were hindered by low identification rates (39.6%) due to lower-frequency trackers and basic statistical features [12]. While subsequent research by Rigas et al. advanced the state-of-the-art significantly, achieving 88.6% accuracy, these studies were often constrained by small cohort sizes—frequently fewer than 100 participants—limiting the statistical power needed to assert the universality of the privacy threat [11]. Recent advances in the VR domain,

Table 1 Comparison of biometric identification performance on VR-based eye-tracking datasets with more than 100 participants

Study	Features	Classifier	Datasets	Results
Lohr et al.	F, S	STAT	VREM-R1	EER: 9.98%
Lohr et al.	F, S	RBF	VREM-R1	EER: 14.37%
Monaco	G, A	STAT	EMVIC 2014	IR: 39.6%
Rigas et al.	F, S, D	MSF	Own Dataset	IR: 88.6%
Our Study	F, S, SG, P, TP, TD	MLP	GazeBaseVR	IR: 96.61%

Methods: RBF = Radial Basis Function Network, STAT = Statistical test, *MSF = Multi-score fusion, MLP = Multi-Layer Perceptron.*
Features: F = Fixation; S = Saccade; D = Density maps; G = Gaze velocity; A = Acceleration; SG = Savitzky–Golay velocity features; P = Position statistics; TP = Target position statistics; TD = Task duration

such as those by Lohr et al. on the VREM-R1 dataset, introduced ecologically valid constraints but reported error rates in the range of 10–14% [1].

Our work explicitly targets these limitations. By utilizing the GazeBaseVR dataset ($N = 407$) and refining the feature extraction pipeline, we demonstrate a significant leap in performance. The achievement of **96.61% accuracy** in passive video-viewing tasks not only outperforms these prior benchmarks but does so on a dataset with more than **400 subjects**. This comparison underscores a critical evolution in the field: the "uniqueness" of the oculomotor plant is not an artifact of small sample sizes but a robust biometric trait that scales with the population, rendering the privacy threat far more acute than previously understood.

Beyond user identification, eye-tracking data has been shown to reveal a broad range of sensitive personal attributes, including age, gender, health conditions, sexual orientation, interests, and decision-making patterns [13]. These insights are derived from distinct gaze behavior metrics such as gaze velocity, saccade latency, blinking rhythm, and pupil size variation. While eye-tracking biometrics enable natural interaction and immersive XR experiences [14], they also expand the scope of privacy risks [5]. As [5] note, when eye-tracking data is combined with other biometric or behavioral indicators, it can support sophisticated user profiling and unauthorized monitoring. Addressing these threats will require the strict enforcement of privacy standards and the development of privacy-preserving mechanisms that safeguard users while allowing the continued integration of eye-tracking into XR systems.

3 Methodology

In this section, we detail the steps taken to investigate the feasibility of user identification based on eye movement patterns within XR environments. Our primary objective is to assess the extent to which biometric information derived from eye tracking can compromise user anonymity. To achieve this, we first leverage the GazeBaseVR dataset, a large and diverse eye-tracking dataset collected in virtual reality environments. We preprocess this data by applying the Savitzky-Golay filter to enhance the quality and accuracy of velocity estimations. Following preprocessing, we perform extensive feature extraction, deriving velocity and position statistics as well as fixation and saccade metrics that characterize individual gaze behavior. These engineered features are then used to train a Multi-Layer Perceptron (MLP) model designed for user classification tasks. We evaluate the model's performance, focusing particularly on the video-viewing task, to measure the precision and risks associated with biometric identification in realistic XR settings. Throughout this process, we aim to illustrate how advanced data processing and machine learning techniques, when combined with high-resolution XR eye-tracking data, can significantly erode user privacy by enabling highly accurate identification.

3.1 The gazeBaseVR dataset

The GazeBaseVR is a large-scale, longitudinal, binocular eye-tracking (ET) dataset collected at 250 Hz using a VR headset equipped with eye-tracking capabilities. The dataset includes 5,020 binocular recordings from a diverse sample of 407 college-aged participants. Each participant was recorded up to six times over a 26-month period, performing a series of five distinct ET tasks: (1) a vergence task, (2) a horizontal smooth pursuit task, (3) a video-viewing task, (4) a self-paced reading task, and (5) a random oblique saccade task. GazeBaseVR is well-suited for diverse research applications involving ET data in VR, particularly in eye movement biometrics, thanks to its extensive participant base and longitudinal design. Additionally, the dataset includes n participant details to support research on topics such as fairness. [1].

The dataset's high sampling frequency enables reliable and accurate tracking of eye movements, which is critical for identification and behavioral analysis. The variety of visual tasks—ranging from vergence to reading and video-viewing—further enhances its applicability across different research contexts, making it a replicable and generalizable resource. These diverse tasks also allow researchers to explore how eye-tracking data can be utilized for both identification and behavioral analysis under various conditions.

We selected GazeBaseVR specifically because it overcomes the limitations of prior datasets in three critical dimensions. First, with $N = 407$ participants, it provides a population scale significantly larger than previous VR-specific collections–which have often been limited to fewer than 50 subjects. This extensive participant base enables us to capture a wider range of inter-subject variability, ensuring that our identification results are statistically robust and not artifacts of a small cohort. Second, unlike desktop-based trackers (e.g., the original GazeBase) that often operate under restricted viewing conditions, GazeBaseVR captures eye movements within an immersive head-mounted display (HTC Vive Pro Eye). This distinction is crucial, as it ensures high ecological validity by replicating the natural head and eye dynamics users exhibit in authentic XR environments, thereby allowing for a more accurate assessment of privacy risks. Third, the scientific soundness of the data is underpinned by the rigorous technical validation performed by Lohr et al. [1], who confirmed that the system's spatial accuracy and precision are sufficient for fine-grained oculomotor analysis. Crucially, the 250 Hz sampling rate exceeds the Nyquist frequency required to reconstruct the ballistic profiles of saccades, specifically their peak velocities and acceleration onset, which serve as the most discriminative features in our biometric pipeline. By utilizing a dataset that captures these high-frequency dynamics across a longitudinal span of 26 months, we ensure that our identification model relies on stable physiological traits rather than transient sensor artifacts. For our study, GazeBaseVR was an invaluable resource, providing a rich and detailed collection of high quality, diverse data collected in a VR context. It enabled us thus to evaluate identification models under realistic conditions and reinforced the reliability and relevance of our findings. Additionally, the dataset's distinction between visual tasks provided deeper insights into eye-tracking biometrics, further contributing to our understanding of their applications and limitations. At the same time, the longitudinal and multi-session nature of GazeBaseVR introduces challenges in managing and processing such a large corpus. These challenges are not unique to eye-tracking: scalability issues such as random access patterns and computational conflicts are well recognized in large-scale data domains. Prior comparative studies of regression-based approaches for scalability analysis [15] similarly emphasize that the choice of analytical methods has a direct impact on efficiency and reliability when working with massive datasets.

3.2 Data processing techniques

The GazeBaseVR dataset was first processed to ensure that the recorded gaze signals provided a clean and reliable foundation for subsequent biometric analysis. A central component of this stage was the use of the Savitzky–Golay filter, one of the most widely adopted techniques for accurate velocity estimation in eye-movement research [1]. This filter operates by fitting successive windows of adjacent data points with a low-degree polynomial using the method of linear least squares, which has the dual advantage of reducing random noise while preserving the essential structure of the signal. In the context of eye-tracking, this is particularly important because unprocessed velocity estimates are highly sensitive to measurement noise, leading to unreliable detection of saccades, fixations, and other gaze dynamics.

The ability of the Savitzky–Golay filter to smooth data without overly distorting the underlying movement trajectory makes it especially suited for extended reality (XR) settings, where head-mounted displays and complex stimuli often introduce additional variability in gaze measurements. By producing stable velocity profiles, the filter ensures that the downstream extraction of features such as velocity statistics, fixation rates, and saccade characteristics is based on a faithful representation of user behavior. This direct connection between data processing and user differentiation underscores the importance of reliable preprocessing in XR biometrics research. Integrity and accuracy in this stage are not only desirable but fundamental, since any noise or distortion in the input signal would propagate into later steps and degrade identification performance. As highlighted in prior work [1], the careful implementation of filtering standards provides the stability necessary for eye-tracking biometrics to succeed in differentiating between individuals, enabling both precise biometric analysis and accurate identity recognition in XR environments.

To mitigate sensor noise while preserving saccadic onset characteristics, we applied a Savitzky-Golay filter to the raw positional data. We utilized a window length of 7 samples (corresponding to 28 ms at 250 Hz) and a polynomial order of 2. Following smoothing, the gaze velocity was calculated as the Euclidean norm of the horizontal and vertical velocity components:

$$v(t) = \sqrt{v_x(t)^2 + v_y(t)^2}$$

3.3 Feature extraction

Building on the preprocessed signals, we extracted a comprehensive set of features to characterize both the dynamic and spatial dimensions of eye movements. From the smoothed velocity profiles produced by the Savitzky–Golay filter, we calculated descriptive statistics including the mean, standard deviation, and maximum velocity, as well as higher-order measures such as skewness and kurtosis of the velocity distribution. These features capture distinctive patterns in how users perform saccades–rapid, often involuntary eye movements that have been shown to be highly individualistic [16]. The inclusion of skewness and kurtosis is particularly important, as they highlight asymmetries and extreme values in velocity distributions that simpler summary measures might overlook.

In parallel, position statistics were computed, including the mean and standard deviation of gaze coordinates in both the X and Y dimensions. These features capture how long and where individuals tend to fixate, thereby revealing spatial tendencies and dispersion patterns in visual attention. For tasks with defined target positions, we also computed the mean and standard deviation of X and Y target coordinates to assess how closely and consistently a user's gaze aligned with the task demands, which adds another dimension of individual variation.

To further describe gaze behavior, we extracted saccade and fixation metrics: the saccade rate, defined as the proportion of samples above a velocity threshold, and the fixation rate, defined as the proportion below. These measures provide insight into how often and for how long each user engages in exploratory versus steady gaze patterns [1]. Finally, task duration was quantified by the total number of samples collected in each trial, providing a contextual measure of engagement and temporal behavior that, when combined with the other features, contributes to the overall biometric profile.

Together, these features form a multi-faceted representation of eye movement behavior. Velocity statistics capture rapid and unique saccadic characteristics [16], position statistics reflect fixation biases and spread, target-related measures quantify alignment consistency, and fixation/saccade rates describe the rhythm of visual exploration [1]. Task duration then ties these elements together, embedding them within the temporal context of the activity.

To capture the fine-grained temporal evolution of gaze behavior during the task, we employed a high-resolution segmentation strategy. Each recording was divided into $N = 200$ equal temporal portions, ensuring that transient dynamic fluctuations were preserved rather than smoothed out by global averaging. For each individual portion, we extracted a vector of statistical features including the mean, standard deviation, maximum, skewness, and kurtosis of the velocity distribution, as well as the mean and standard deviation of gaze coordinates. The inclusion of higher-order moments is particularly critical, as they quantify the asymmetry and peak characteristics of saccadic profiles unique to each user. These local feature vectors were concatenated with global signal statistics to form the final input vector for the classifier.

Rather than using fixed velocity thresholds (e.g., $30°/s$), we implemented **dynamic per-trial thresholding** to adapt to varying noise levels across sessions. Saccade (T_{sac}) and fixation (T_{fix}) thresholds were defined relative to the global velocity distribution of the specific recording:

$$T_{sac} = \mu_{vel} + 2\sigma_{vel}, \quad T_{fix} = \mu_{vel} - 1\sigma_{vel}$$

where μ_{vel} and σ_{vel} denote the mean and standard deviation of the velocity for the trial.

Importantly, this rich and diverse set of features allowed us to achieve high identification accuracy (96.61%) even when employing a relatively simple Multi-Layer Perceptron (MLP) classifier. The fact that such a straightforward model could leverage these engineered features so effectively underscores the discriminative power embedded in gaze dynamics themselves. This finding illustrates that privacy risks do not only emerge when applying advanced deep learning models, but can already be realized through carefully processed and extracted features, reinforcing the urgency of addressing eye-tracking privacy concerns in XR technologies.

3.4 Model architecture and training process

The extracted features were utilized to train a Multi-Layer Perceptron (MLP) classifier. We selected this architecture because it represents a lightweight, widely used neural network baseline that can be trained efficiently while still capturing non-linear relationships in the data. Our aim was not to design the most complex or specialized model, but to demonstrate that even with a simple and fast classifier, eye movement features are sufficiently distinctive to enable reliable user identification. Several candidate configurations were tested, and the final chosen model offered the best balance between performance and efficiency. This baseline also serves as a foundation for future work, where more sophisticated deep learning models could be employed to push accuracy further.

Model Architecture

The architecture of the MLP classifier (Figure 1) consists of an input layer, three hidden layers, and an output layer, with a skip connection between the input and output layers to enhance learning efficiency. The input layer size corresponds to the number of extracted features, including velocity statistics, position statistics, and other derived metrics as mentioned in the previous section. The hidden layers are structured with 64, 32, and 16 units, respectively, each applying non-linear transformations to the input data. These layers allow the model to learn and distill complex patterns in the data.

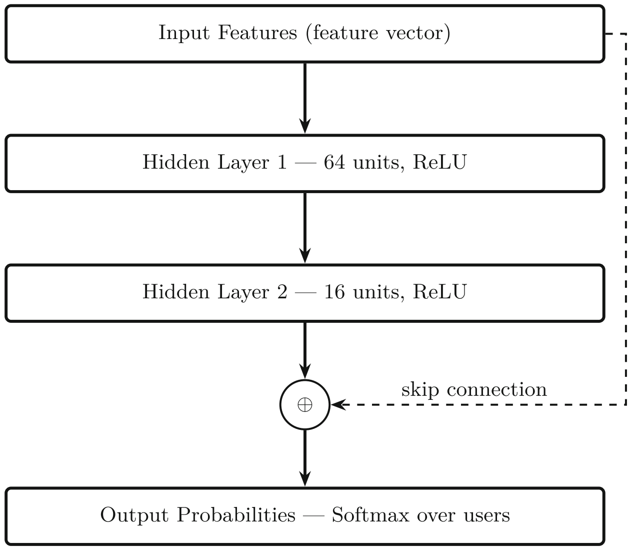

Fig. 1 MLP model architecture

The skip connection links the input layer directly to the output layer, bypassing the hidden layers. This mechanism improves gradient flow, mitigates the vanishing gradient problem, and enhances training efficiency, particularly in architectures with multiple layers. The output layer size corresponds to the number of unique users in the dataset, with each unit representing a specific user. A SoftMax activation function is applied to the output to generate probabilities for classification.

Training Process

The training process was optimized to ensure efficient convergence and robust performance. The model was trained using the Adam optimizer with a learning rate of 0.05, a commonly adopted optimization algorithm known for its ability to manage sparse gradients and noisy data effectively. The cross-entropy loss function was employed to minimize classification error, providing a probabilistic framework for the model's output. The training was conducted over 200 epochs with a batch size of 256, which offered a practical balance between computational efficiency and generalization. These parameters were selected based on reasonable defaults and adjusted empirically to achieve stable convergence. During training, both velocity- and position-based feature vectors were forwarded to the MLP, enabling the model to learn distinct patterns within these metrics.

3.5 Model performance

The MLP classifier achieved a test accuracy of 96.61% when identifying users specifically in the video-viewing task, demonstrating its capability to identify users based on eye movement data. This task is particularly significant, as video consumption is one of the most fundamental and ubiquitous activities in virtual reality (VR). Whereas many XR tasks are context-dependent or application-specific, video consumption constitutes a ubiquitous and unavoidable form of interaction, ensuring that the associated privacy risks apply broadly across the user population. As a result, achieving high identification accuracy in this setting indicates that the privacy risk is not confined to specific use cases, but instead represents a universal concern across XR platforms. This performance highlights the model's ability to process and classify high-dimensional feature vectors effectively, underlining the potential privacy risks associated with eye-tracking biometrics in extended reality platforms.

It is important to clarify that the unit of analysis for this study is the complete recording session. For the video-viewing task, this represents a duration of approximately 38 seconds–a timeframe comparable to a standard online advertisement or a short VR cutscene. The feature vectors described in Section 3.3 were constructed by concatenating the statistics of all 200 temporal segments into a single input vector. This design ensures that our model captures the unique spatiotemporal signature of the user's attention path throughout the stimulus. Therefore, the identification accuracy of 96.61% demonstrates that users can be de-anonymized with near-certainty after less than a minute of passive content consumption. To ensure the transparency and replicability of our results, we implemented a strict evaluation protocol. First, the dataset was partitioned based on recording sessions: for each of the 407 users, the first two sessions were utilized for training, while the third session was held out exclusively for testing. This prevents data leakage and tests the model's ability to recognize users across different points in time. Second, the classification for each trial was determined by the highest output probability (argmax) of the Softmax layer. Finally, the confusion matrix presented in Figure 2 was row-normalized, meaning the diagonal values represent the Recall (true positive rate) for each user, ensuring that the visualization is not skewed by class imbalances. The metrics in Table 2 (Precision, Recall, F1) were computed using macro-averaging, treating all users as equally important regardless of sample size.

As shown in (Figure 2), the confusion matrix demonstrates a strong diagonal dominance with minimal off-diagonal elements, indicating high classification accuracy and low inter-class confusion. This qualitative observation is quantitatively supported by the results presented in (Table 2), where the model achieves a **precision of 0.9883**, **recall of 0.9696**, **F1-score of 0.9746**, and an overall **accuracy of 96.61%**. These metrics confirm the model's strong performance and its ability to generalize effectively across a wide range of classes.

The proposed model provides critical insights into the privacy implications of eye-tracking data. In particular, the high identification accuracy achieved in the video viewing task

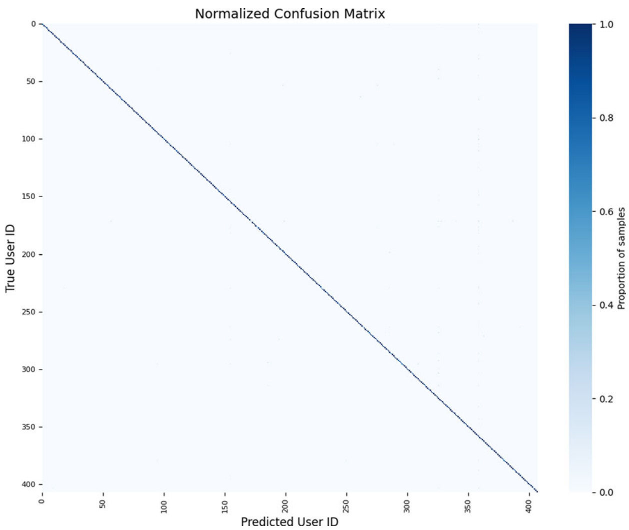

Fig. 2 Confusion matrix demonstrating model performance

Table 2 Performance metrics of our model

Metric	Value
Precision	0.9883
Recall	0.9696
F1 Score	0.9746
Accuracy	0.9661

underscores that these risks are not limited to specialized use cases, but extend to the most common and inevitable forms of XR usage. This work establishes a strong foundation for addressing user identification and privacy threats in extended reality environments, highlighting the urgent need for privacy-preserving measures in such applications.

4 Results

In this study, we successfully demonstrated that user identification based on eye movement data is achievable with remarkably high precision, even among large user groups. Although such high accuracy was attained only during the video viewing task, this outcome aligns with the primary objective of our research—to focus on the most common and ubiquitous activity. By doing so, we aimed to highlight the real-world privacy risks users face during everyday interactions within XR environments.

However, data from the other, more specialized tasks played a crucial role during the feature extraction and model training phases, ultimately contributing to the high level of accuracy achieved in the video-viewing activity. The success of the video-viewing task can largely be attributed to the predictable and uniform nature of its visual stimuli, as all participants were exposed to the same video content. This consistency led to more stable gaze behaviors across users, resulting in clearer and more distinguishable biometric patterns. Additionally, the structured format of the task—where visual attention was typically directed toward a limited set of screen locations—further enhanced the classifier's ability to identify unique user signatures with high precision. These findings not only highlight the video-viewing task as a particularly effective context for biometric identification but also underscore significant privacy implications, given the accuracy with which users can be recognized based solely on their gaze data.

In contrast to the video-viewing task, user identification performance was lower for other tasks, as anticipated. Activities such as horizontal smooth pursuit and random oblique saccades yielded identification accuracies of 68.33% and 72.11%, respectively, highlighting the increased variability in gaze behavior associated with these types of stimuli. Even less consistent results were observed in tasks like self-paced reading (57.37%) and vergence (43.82%), where gaze patterns tend to be more individualized and unpredictable. These lower accuracies are understandable, given the dynamic and less uniform nature of the stimuli, which complicates the extraction of consistent biometric signatures.

Despite the reduced identification performance in these tasks, their inclusion proved valuable. The data collected contributed important features during the feature extraction phase, ultimately enhancing the model's ability to generalize and perform well in more structured and realistic scenarios such as video viewing. By incorporating diverse eye movement data across various tasks, the classifier was better equipped to identify users with higher precision in contexts where gaze patterns are more stable and representative of real-world interactions.

The conditions and outcomes of this study closely mirror real-world scenarios, where XR platforms are capable of collecting diverse eye movement data from a broad user base. Notably, the findings demonstrate that accurate user identification can still be achieved using only a limited subset of data gathered during a video-watching activity. This underscores a critical privacy concern: anonymity within XR environments cannot be guaranteed unless deliberate measures are implemented to mitigate the risks associated with the use of eye-tracking data as a biometric identifier. The study thus highlights the pressing need for privacy-preserving mechanisms in the design and deployment of XR systems.

5 Mitigation strategies: balancing utility and anonymity

The demonstration that eye movement patterns can serve as reliable biometric identifiers necessitates a para-digm shift in how XR systems handle gaze data. As the fidelity of eye tracking increases to meet the demands of foveated rendering and avatar realism, so too does the risk of unauthorized profiling. To resolve this tension between performance and privacy, we must move beyond binary access controls and adopt granular privacy-preserving mechanisms. This section outlines a tiered defense strategy, ranging from cryptographic frameworks to signal obfuscation, designed to decouple the utility of gaze data from its biometric distinctiveness.

5.1 Differential privacy frameworks

Differential Privacy (DP) provides a rigorous mathematical guarantee that the output of an algorithm does not reveal the presence or absence of a specific individual's data. Two specific frameworks have shown promise for eye tracking:

Kaleido: For real-time applications, the Kaleido mechanism implements (ϵ, w, r)-differential privacy [17]. It establishes a "trust boundary" between the eye tracker and the application by applying noise to the gaze stream within a temporal window w and spatial radius r. This satisfies the property of geo-indistinguishability, preventing applications from learning the exact high-precision location of the gaze while still allowing for interaction with larger UI elements.

DPGazeSynth: For scenarios requiring data sharing (e.g., heatmaps), DPGazeSynth employs a generative model based on Markov Chains to synthesize gaze paths. Instead of transmitting raw coordinates, the system generates a synthetic trajectory that mimics the user's general attention but removes the micro-temporal signatures (such as peak saccadic velocity) that encode identity. Studies indicate this can reduce identification rates from $\sim$ 85% to $\sim$ 30% while retaining utility for foveated rendering [18].

5.2 Signal degradation techniques

Biometric identification relies on high-frequency, high-fidelity signal features to distinguish unique physiological traits. Consequently, reducing the signal-to-noise ratio–effectively lowering the quality of the data available to attackers–can effectively thwart these algorithms without necessarily breaking the core interaction mechanics.

Temporal Down-sampling: Our study utilized 250 Hz data, which is sufficient to capture the fine dynamics of saccades. Reducing the sampling rate (e.g., to 30–60 Hz) destroys the high-frequency velocity information required to calculate peak saccadic velocity and acceleration, as these rates fall below the Nyquist frequency needed to reconstruct ballistic eye movements. While effective at reducing re-identification risk [4], this trade-off must be managed carefully, as overly aggressive down-sampling can introduce latency and motion sickness in VR environments, degrading the immersive experience.

Noise Injection and Smoothing: Adding random Gaussian noise to gaze coordinates or applying linear weighted average smoothing can mask fixation stability features [19]. These lightweight privacy mechanisms can prevent the extraction of microsaccadic jitter, a key biometric trait, though excessive noise may result in a "floaty" or unresponsive user experience.

5.3 Oculomotor plant modification

A more advanced defense involves modifying the data stream to simulate a *different* oculomotor plant. By altering the ballistic profile of saccades–such as artificially clamping the peak velocity or modifying the acceleration skewness via percentile matching–a system could theoretically project a "synthetic" biometric identity [20]. This approach targets the underlying physiological "plant" characteristics used for authentication, allowing the user to maintain precise gaze control for interaction while presenting a false biometric signature to the tracking algorithm.

6 Conclusion

User identification through eye movement is an emerging field with profound implications for extended reality (XR) environments. While this study demonstrates high identification accuracy, further advancements can be achieved through improved feature extraction techniques from raw eye movement data and the exploration of alternative machine learning approaches to enhance the effectiveness of identification mechanisms. For instance, recent facial-recognition work shows that multi-scale feature fusion and spatial attention can lift accuracy and robustness under challenging conditions; adapting analogous fusion/attention architectures to gaze sequences may help capture subtle spatio-temporal signatures that current pipelines miss [21].

Furthermore, the findings of this study should be interpreted within the context of specific limitations. First, regarding demographics, the GazeBaseVR dataset is composed primarily of college-aged participants. Since oculomotor characteristics–such as saccadic latency and peak velocity–can shift due to age-related physiological changes, future research must validate these identification risks across broader populations, including children and elderly users, to ensure that privacy protections are universally effective. Second, our model relied on high-fidelity data sampled at 250 Hz, which is sufficient to capture fine-grained micro-saccadic dynamics. In contrast, current mass-market consumer headsets often operate at lower frequencies (60–90 Hz), effectively "smoothing" the gaze signal.

While this signal degradation may currently offer a temporary layer of hardware-based anonymity, this barrier is likely to erode as consumer hardware specifications improve. Finally, the privacy threat is highly context-dependent: while we observed near-perfect identification during passive video viewing (96.61%), accuracy dropped significantly for cognitively active tasks like self-paced reading (57.37%). This indicates that user vulnerability is not static, but fluctuates depending on the nature of the XR interaction, with passive entertainment posing the most acute risk

Ongoing research should place strong emphasis on the development of ethical data usage frameworks and robust privacy protection measures, ensuring that biometric identification in XR is aligned with user rights and data protection standards. Examining the implications of eye-tracking biometrics in the context of extended reality (XR) extends beyond technical innovation and requires a critical focus on privacy and security considerations. While eye-tracking technology enhances our understanding of user attention and interaction patterns within XR environments, it also introduces significant risks related to the collection of sensitive behavioral data—often without explicit user consent. This potential for covert data gathering amplifies user vulnerability to breaches, emphasizing the urgent need for secure system architectures that prevent unauthorized access and misuse of biometric information. Moreover, eye-tracking data may facilitate layered profiling techniques, such as bias detection based on gaze patterns. As noted by Lohr et al. [1], developers and applications may use these insights to infer personal predispositions or preferences, raising additional ethical concerns around manipulation and discrimination. These emerging threats underscore the necessity for comprehensive privacy safeguards and ethical design principles in XR systems. It is imperative to establish rigorous guidelines and protective frameworks that govern the deployment of biometric technologies—ensuring that user rights and autonomy remain central in the evolution of XR applications.

To address the significant privacy concerns associated with eye tracking biometrics, it is imperative that developers implement robust data encryption and anonymization techniques within extended reality environments. Such measures can safeguard sensitive information against potential cyber threats and unauthorized access, ensuring that user data remains protected. Additionally, implementing transparent data usage policies and obtaining explicit consent from users, before collecting eye movement data, can foster trust and compliance with privacy regulations. Moreover, researchers and developers should work collaboratively on establishing industry standards that prioritize ethical data handling, focusing on methods that limit the potential for misuse of biometric information. These standards must emphasize a balance between leveraging the technological benefits of eye tracking for enhanced user experiences and upholding the privacy rights of individuals in virtual spaces.

Various mitigation techniques have been proposed, leveraging differential privacy [5, 22], Gaussian noise addition [4, 18], Temporal down-sampling [4] and k-anonymity [2]. In related biometrics, complementary strategies have emerged; for example, a GAN–blockchain framework generates privacy-preserving synthetic faces while providing an immutable audit layer. Although conceptually distinct from gaze data, referencing such approaches helps situate eye-tracking within the broader privacy-preserving biometrics landscape and highlights domain-specific trade-offs [23]. And while certain researchers have used additional eye movements like blinking to create authentication systems like "BlinKey" [24], others have focused on addressing privacy challenges. For instance, the DPGazeSynth framework offers a promising solution by employing differentially private data synthesis to protect eye-tracking information in virtual reality environments [18]. This approach balances the need for data utility with stringent privacy requirements by utilizing a semi-synthetic method based on the Markov Chain model. The framework's robust privacy guarantees outperform previous solutions, such as Kalεido [17], ensuring that sensitive biometric data remains secure while maintaining functionality.

It remains essential to advance privacy-preserving mechanisms from a technical standpoint to protect user data without impeding innovation in extended reality (XR) environments. Future research should prioritize the design of robust privacy frameworks that integrate advanced cryptographic techniques—such as differential privacy and federated learning—to enable secure data processing without compromising system functionality. At the same time, recent reviews caution that federated learning introduces its own attack surface (e.g., poisoning and malicious clients) and non-IID challenges; coupling FL with IDS-style defenses and careful feature engineering will be essential for secure deployment in XR eye-tracking [25]. The development of standardized privacy protocols for XR platforms could further establish a baseline for user protection, promoting industry-wide accountability and interoperability. Equally critical is the adoption of user-centric approaches, including real-time privacy controls, informed consent mechanisms, and transparency tools, to empower individuals with agency over their eye-tracking data.

Addressing the multifaceted challenges of biometric identification in XR will require sustained collaboration among engineers, designers, ethicists, policy makers and regulators to create systems that are not only technologically sound but also compliant with legal frameworks, ethically responsible and socially acceptable.

Finally, research should also explore the potential risks of adversarial attacks on eye-tracking data and develop countermeasures to mitigate these threats. By fostering a holistic and collaborative approach, the field can ensure that XR technologies continue to evolve in a manner that respects and protects user privacy while unlocking their immense potential for innovation and utility.

Author Contributions A.P. conceived the study, performed the experiments, and wrote the first draft of the manuscript. C.L. supervised the research, contributed to the study design, and revised the manuscript critically for important intellectual content. All authors reviewed and approved the final version of the manuscript.

Funding Open access funding provided by HEAL-Link Greece. The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing Interests The authors have no relevant financial or non-financial interests to disclose.

References

1. Lohr, D., Aziz, S., Friedman, L., Komogortsev, O.V.: GazeBaseVR, a large-scale, longitudinal, binocular eye-tracking dataset collected in virtual reality. Sci. Data **10**(1), 177 (2023). https://doi.org/10.1038/s41597-023-02075-5
2. David-John, B., Butler, K., Jain, E.: For Your Eyes Only: privacy-preserving eye-tracking datasets, in 2022 Symposium on Eye Tracking Research and Applications (ETRA '22) (2022). https://doi.org/10.1145/3517031.3529618
3. Koch, R.: What are you looking at? emerging privacy concerns with eye tracking in virtual reality. Colo. Tech. L.J. **21**(1), 109 (2023)
4. David-John, B., Hosfelt, D., Butler, K., Jain, E.: A privacy-preserving approach to streaming eye-tracking data. IEEE Trans. Visual Comput. Graphics **27**(5), 2555–2565 (2021). https://doi.org/10.1109/tvcg.2021.3067787
5. Bozkir, E., Özdel, S., Wang, M., David-John, B., Gao, H., Butler, K., Jain, E., Kasneci, E.: Eye-tracked virtual reality: a comprehensive survey on methods and privacy challenges (2023). arXiv:2305.14080
6. Schafer, R.W.: What is a savitzky-golay filter? IEEE Signal Process. Mag. **28**(4), 111 (2011). https://doi.org/10.1109/MSP.2011.941097
7. Haria, V., El Abed,A., Maneth, S.: User Identification via Free Roaming Eye Tracking Data (2024). arXiv:2403.09415
8. Asish, S.M., Kulshreshth, A.K., Borst, C.W.: User identification utilizing minimal eye-gaze features in virtual reality applications. Virtual Worlds **10**, 42 (2022). https://doi.org/10.3390/virtualworlds1010004
9. Schröder, C., Al Zaidawi, S.M.K., Prinzler, M.H., Maneth, S., Zachmann, G.: Robustness of eye movement biometrics against varying stimuli and varying trajectory length, in Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (2020), pp. 1–7
10. George, A., Routray, A.: A score level fusion method for eye movement biometrics. Pattern Recogn. Lett. **82**, 207 (2016)
11. Rigas, I., Abdulin, E., Komogortsev, O.: Towards a multi-source fusion approach for eye movement-driven recognition. Inf. Fusion **32**, 13 (2016)
12. Monaco, J.V.: Classification and authentication of one-dimensional behavioral biometrics, in IEEE International Joint Conference on Biometrics, pp. 1–8 (2014)
13. Selinger, E., Altman, E., Foster, S.: Eye-tracking in virtual reality: a visceral notice approach for protecting privacy. Privacy Studies Journal **2**, 1 (2023)
14. Klaib, A.F., Alsrehin, N.O., Melhem, W.Y., Bashtawi, H.O., Magableh, A.A.: Eye tracking algorithms, techniques, tools, and applications with an emphasis on machine learning and Internet of Things technologies. Expert Syst. Appl. **166**, 114037 (2021). https://doi.org/10.1016/j.eswa.2020.114037
15. Latif, N., Arshad, M.S., Javed, S.: A comparative study between regression algorithms to visualize and analyze the scalability of large-scale graphs in big data, in 2021 International Conference on Emerging Technologies (ICET), pp. 1–6 (2021) https://doi.org/10.1109/ICET54505.2021.9689863
16. Harezlak, K., Blasiak, M., Kasprowski, P.: Biometric identification based on eye movement dynamic features. Sensors **21**(18), 6020 (2021). https://doi.org/10.3390/s21186020
17. Li, J., Chowdhury, A.R., Fawaz, K., Kim, Y.: Kaleido: real-time privacy control for eye-tracking systems, in Proceedings of the 30th USENIX Security Symposium, pp. 1793–1810 (2021). https://www.usenix.org/conference/usenixsecurity21/presentation/li-jingjie
18. Ren, X., Fan, J., Xu, N., Wang, S., Dong, C., Wen, Z.: DPGazeSynth: enhancing eye-tracking virtual reality privacy with differentially private data synthesis. Inf. Sci. **675**, 120720 (2024). https://doi.org/10.1016/j.ins.2024.120720
19. Raju, M.H., Komogortsev, O.V.: Real-time lightweight gaze privacy-preservation techniques validated via offline gaze-based interaction simulation (2025). arXiv:2511.09846
20. Griffith, H., Katrychuk, D., Raju, M.H., Aziz, S., Lohr, D.J., Komogortsev, O.V.: The percentile-matching technique for synthetic eye tracking signal degradation: a biometric case study. PLoS ONE **20**(11), 1 (2025). https://doi.org/10.1371/journal.pone.0334605
21. Ghani, M.A.N.U., She, K., Saeed, M.U., Latif, N.: Enhancing facial recognition accuracy through multi-scale feature fusion and spatial attention mechanisms. Electron. Res. Arch. **32**(4), 2267 (2024). https://doi.org/10.3934/era.2024103
22. Steil, J., Hagestedt, I., Huang, M.X., Bulling, A.: Privacy-aware eye tracking using differential privacy, in Proceedings of the 2019 Symposium on Eye Tracking Research and Applications (ETRA '19), pp. 1–9 (2019). https://doi.org/10.1145/3314111.3319915
23. Ghani, M.A.N.U., She, K., Rauf, M.A., Alajmi, M., Ghadi, Y.Y., Algarni, A.: Securing synthetic faces: a GAN-blockchain approach to privacy-enhanced facial recognition. . King Saud Univ. Comput. Inf. Sci. **36**(4), 102036 (2024). https://doi.org/10.1016/j.jksuci.2024.102036
24. Zhu, H., Jin, W., Xiao, M., Murali, S., Li, M.: BlinKey: a two-factor user authentication method for virtual reality devices. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. **4**(4), 164 (2020)

25. Latif, N., Ma, W., Ahmad, H.B.: Advancements in securing federated learning with IDS: a comprehensive review of neural networks and feature engineering techniques for malicious client detection. Artif. Intell. Rev. **58**(3), 91 (2025). https://doi.org/10.1007/s10462-024-11082-w