



OPEN ACCESS

EDITED BY

David Swapp,
University College London,
United Kingdom

REVIEWED BY

Natasha Banerjee,
Clarkson University, United States
Michael Vallance,
Future University Hakodate, Japan

*CORRESPONDENCE

Christian Rack,
✉ christian.rack@uni-wuerzburg.de

RECEIVED 03 August 2023

ACCEPTED 16 October 2023

PUBLISHED 10 November 2023

CITATION

Rack C, Fernando T, Yalcin M, Hotho A
and Latoschik ME (2023), Who is Alyx? A
new behavioral biometric dataset for user
identification in XR.

Front. Virtual Real. 4:1272234.

doi: 10.3389/frvir.2023.1272234

COPYRIGHT

© 2023 Rack, Fernando, Yalcin, Hotho
and Latoschik. This is an open-access
article distributed under the terms of the
[Creative Commons Attribution License
\(CC BY\)](#). The use, distribution or
reproduction in other forums is
permitted, provided the original author(s)
and the copyright owner(s) are credited
and that the original publication in this
journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Who is Alyx? A new behavioral biometric dataset for user identification in XR

Christian Rack^{1*}, Tamara Fernando¹, Murat Yalcin¹,
Andreas Hotho² and Marc Erich Latoschik¹

¹Human-Computer Interaction (HCI) Group, Informatik, University of Würzburg, Würzburg, Germany,

²Data Science Chair, Informatik, University of Würzburg, Würzburg, Germany

Introduction: This paper addresses the need for reliable user identification in Extended Reality (XR), focusing on the scarcity of public datasets in this area.

Methods: We present a new dataset collected from 71 users who played the game “Half-Life: Alyx” on an HTC Vive Pro for 45 min across two separate sessions. The dataset includes motion and eye-tracking data, along with physiological data from a subset of 31 users. Benchmark performance is established using two state-of-the-art deep learning architectures, Convolutional Neural Networks (CNN) and Gated Recurrent Units (GRU).

Results: The best model achieved a mean accuracy of 95% for user identification within 2 min when trained on the first session and tested on the second.

Discussion: The dataset is freely available and serves as a resource for future research in XR user identification, thereby addressing a significant gap in the field. Its release aims to facilitate advancements in user identification methods and promote reproducibility in XR research.

KEYWORDS

dataset, behaviometric, deep learning, user identification, physiological dataset

1 Introduction

The field of Virtual Reality, Augmented Reality and Mixed Reality [VR, AR and MR; extended reality (XR), in short] matured considerably in the last decade and enables new ways for corporations, institutions and individuals to meet, support and socialize without the need for users to be physically present. Consequently, XR architects are increasingly faced with the challenge of developing solutions that promote authenticity and trust among users, while meeting the high usability requirements of XR applications at the same time. For example, typical password-based verification mechanisms require text input, which is problematic to provide in XR contexts as it comes with notable usability problems (Knierim et al., 2018; Dube and Arif, 2019; Kern et al., 2021; Stephenson et al., 2022). This renders typing-based security mechanisms cumbersome and sub-par.

Recent approaches explore behavioral biometrics as an interesting alternative solution for identification and verification. In addition, such approaches open-up novel features for social XR. For example, they allow continuous monitoring of user movements to provide important cues about avatars' identities, i.e., who is controlling the digital alter egos, which becomes increasingly important when avatars are replicating the appearance of real people (Achenbach et al., 2017; Wenninger et al., 2020). Such approaches do not necessarily need extra hardware, but can work with the data that is

already provided freely by the XR system. Of particular interest in this context is the tracking data that captures the movements of XR users, as this is a constant and readily available data stream that is required for any XR system in order to provide an immersive experience. Analogous to the handwriting of a person, movements in general seem to have very individual characteristics: previous works have demonstrated that identifying information can be extracted from users' movements, so using tracking data provides the foundation for novel approaches to identify and verify XR users.

Following Jain et al. (2011), systems of user recognition can generally be divided into two categories: identification and verification. Identification systems involve the use of biometric inputs from a user to query a database containing the biometric templates of all known users, effectively seeking to answer the question, "Who is this?". Verification systems, on the other hand, compare the biometric input against the templates associated with a specific identity claimed by the user. In this context, "verification" is synonymous with "authentication," akin to password-based verification where a password is verified against a specific user's known credentials, not against all users in the database. This distinction, although often blurred in the literature, is crucial for this article, as these two concepts embody different methodologies, each with its own considerations and applications. This being said, verification and identification approaches are still related on a technical level and can share the same underlying core mechanics (e.g., data preprocessing, machine learning architecture, etc.).

Recent efforts to leverage user motions for identification or verification purposes rely mainly on machine learning models. While these models yield promising results, the works proposing them become susceptible to the ongoing general discussion about a growing reproducibility crisis concerning machine learning methods in research (Kapoor and Narayanan, 2022). One requirement is that researchers must choose the correct methodology and the right evaluation methods to avoid overoptimistic conclusions about their work. Another requirement is reproducibility, as otherwise understanding proposed solutions becomes difficult, errors cannot be found, and new solutions cannot be compared to previous ones. In general, there is a trend that non-replicable publications are cited more often than replicable ones (Serra-Garcia and Gneezy, 2021), which adds to the overall problem. An important component for researchers to make their work reproducible is to use common, publicly available datasets (Pineau et al., 2020).

Unfortunately, there are very few *public* datasets for research concerning XR identification methods: to our knowledge, there are only three datasets featuring tracking data of XR users (Liebers et al., 2021; Miller et al., 2022a; Rack et al., 2022). These datasets are relatively small, or just single-session, which makes them an inadequate choice for the development and the evaluation of larger (deep) machine learning models. This is a severe problem, since the lack of public datasets makes it impossible for researchers 1) to participate in this domain of research if they do not have the time, resources, or access to XR users to create their own dataset, 2) to compare their new methods to current ones, and 3) to reproduce previous solutions. This is in stark contrast to other fields of research, where large public datasets have facilitated the development of impressive systems like BERT (natural language processing) (Devlin et al., 2019), YOLO (object detection) (Wang

et al., 2021) or the recent DALL-E2 (text to image) (Ramesh et al., 2022) and ChatGPT (OpenAI, 2023).

Consequently, we see the lack of public datasets as a serious barrier for mature and practical solutions in this area of research. With our work, we address this issue and publish a new dataset designed for XR user recognition research. The dataset includes the tracking data of 71 VR users who played the game "Half-Life: Alyx" with a HTC Vive Pro for 45 min on two separate days. In addition to recording motion data and eye-tracking from the VR setup, we also collected physiological data for a subset of 31 users, which may be interesting for future work as well. Since the motivation for this dataset is to serve future research for XR user identification and verification, we benchmark two state-of-the-art machine learning architectures that have been shown to find identifying characteristics in user motions.

Recognizing the lack of public datasets as a serious barrier for mature and practical solutions in this area of research, our work provides a new, specially designed dataset for XR user identification research. The dataset includes tracking data from 71 VR users who played the game "Half-Life: Alyx" on an HTC Vive Pro for 45 min on two different days. Alongside the motion and eye-tracking data from the VR setup, we also collected physiological data from a subset of 31 users. While our dataset can potentially be used for both identification and verification research, the focus of this paper is on user identification. To demonstrate the practical value of this dataset, we apply two state-of-the-art machine learning architectures, showing their capacity to identify unique characteristics in user motion data.

Altogether, our contributions are as follows:

1. A new public dataset of 71 VR users, with two sessions per user of 45 min each.
2. Validation of the dataset for the identification task with two state-of-the-art deep learning architectures (CNN and GRU), providing benchmark results for future work.
3. We publish all data and code: <https://github.com/cshell/who-is-alyx>

2 Related work

Our dataset is different to most used by related work, as it is not only freely accessible and relatively large, but also provides a unique setting with VR users performing a wide range of non-specific actions. In the following we discuss the landscape of datasets used by previous works and explain how the "Who Is Alyx?" dataset compares.

2.1 Specific vs. non-specific actions

In this article, we use the terms "specific action" and "non-specific action" to describe the contextualization of input sequences for identification models. *Specific actions* refer to highly contextualized input sequences where the user action is well known and consistent. *Non-specific actions* on the other hand result in less contextualized input sequences that include more arbitrary user motions. We categorize previous work based on the degree of contextualization of input sequences.

TABLE 1 Publicly available XR motion datasets.

Dataset	Liebers et al. (2021)	Rack et al. (2022)	Miller et al. (2022a)	Nair et al. (2023)	Ours
Users	16	34	41 and 16	105,852	71
Footage/user	8 min	60 min	2 min	36 min ^a	90 min
Sessions	2	1	2	Unknown	2

^aMedian value.

Highly contextualized sequences, where user actions are well-defined, are explored by a wide range of works. Here, users are tasked to throw or re-locate virtual (bowling) balls (Ajit et al., 2019; Kupin et al., 2019; Olade et al., 2020; Liebers et al., 2021; Miller et al., 2021), to walk a few steps (Mustafa et al., 2018; Pfeuffer et al., 2019; Shen et al., 2019), to type on a virtual keyboard (Mathis et al., 2020; Bhalla et al., 2021) or to nod to music (Li et al., 2016). In each of these scenarios, the identification model receives short input sequences that include a well defined user action, so most of the variance in the input data should be attributable to the user's personal profile. This allows identification models to specialize on user-specific characteristics within particular motion trajectories. Such scenarios are interesting for authentication use-cases, where the user is prompted to do a specific action and the system can perform the verification based on the same motions each time.

Medium contextualized input sequences originate from scenarios that allow more freedom for users to move and interact. For example, Miller et al. (2020) identified users watching 360° VR videos or answering VR questionnaires for several minutes. Following this work, Moore et al. (2021) explored an e-learning scenario where VR users learned to troubleshoot medical robots in several stages. And recently, Rack et al. (2022) compared different classification architectures using a dataset of participants talking to each other for longer periods. In each of these cases, the general context of the scenario still limits possible user actions (e.g., to "listening" and "talking"), but the exact action, its starting, and its ending point become uncertain. Consequently, the identification task becomes more challenging as models have to deal with more variance not related to the user's identity. Also, it becomes necessary to observe the user for longer periods (i.e., minutes) to make reliable identifications. However, models that can detect user-specific characteristics within such input sequences promise to be applicable to a wider range of scenarios where individual user actions are not clearly specified anymore. This becomes interesting for use cases where XR users cannot be prompted to perform a specific identifying motion and instead need to be monitored passively.

The "Who is Alyx?" dataset is *marginally contextualized*. During recording, the user plays the game for 45 min and executes a wide range of actions along the way. Even though there are several actions that occur repeatedly, like using tools, reloading or using the teleportation mechanism, we consider them non-specific: input sequences for the model could be taken from anywhere in the user recording and may include all, some, or none of these actions. Hence, identification models have to deal with a lot of variance and have to learn user-specific patterns within sequences from a wide range of more or less arbitrary user motions. This makes marginally contextualized datasets like ours an obvious choice to develop identification models that can be applied to a wide range of XR scenarios and use cases.

2.2 Existing datasets

Besides providing a wide range of user actions, the "Who is Alyx?" dataset is also the largest freely available dataset featuring XR user motion data. Although there is already a sizeable body of literature investigating behavior-based identification, to date there are to our knowledge only three *publicly* available datasets viable for the development of methods to identify XR users by their motions. Table 1 lists these datasets and compares them with "Who is Alyx?."

Liebers et al. (2021) published a dataset with 16 VR users, each performing the two specific actions "throwing a virtual bowling ball" and "shooting with bow and arrow." In each of two sessions, users had to complete 12 repetitions of both actions using an Oculus Quest. Each repetition took between 5 and 10 s on average, resulting in a total recording time of 4 min per user per action.

Rack et al. (2022) use the public "Talking With Hands" dataset of Lee et al. (2019) to identify 34 users. The dataset provides full-body motion tracking recordings from participants talking to each other about previously watched movies. After filtering and cleaning the data, the authors used about 60 min per user to train and evaluate their models.

Miller et al. (2022a) recently published a combination of two datasets, featuring 41 and 16 users. These datasets are similar to those of Liebers et al. (2021): they contain recordings of users performing a specific action (i.e., throwing a ball) on two separate days. Additionally, the users repeated the action with three different VR devices, which allows evaluations regarding cross-device generalizability. The datasets contain for each user 10 repetitions of ball-throwing per session and device, each lasting for exactly 3 s, resulting in a total of 1 min of footage per user and device and session.

Each of these datasets has considerable drawbacks for developing and testing new deep learning approaches. The datasets of Liebers et al. (2021) and Miller et al. (2022a) provide only a small amount of data per user and session. This is not just a problem because deep learning models are prone to learn wrong characteristics from individual repetitions, but also because there is too little test data to draw meaningful conclusions about the model's ability to generalize well to unseen data. For example, we replicated the work of Liebers et al. (2021) and followed their proposed methods to train an LSTM model. Here, the test results changed considerably by ± 8 percentage points just by repeating training with differently seeded weights, suggesting that either the training data is too small for robust training or the evaluation data is insufficient to reflect how well the model generalizes. In contrast, the dataset used by Rack et al. (2022) provides much more footage per user, but is only single-session, so it cannot be tested how well a model would recognize the same person on a different day.

Against this backdrop, we created the "Who is Alyx?" dataset with the purpose of not only being publicly available, but to improve upon existing free datasets by providing much more data and two

sessions for each user. In this article we demonstrate that current state of the art identification models can be robustly trained on session one and generalize to session two. This provides future work not only with a large dataset, but also with benchmark scores to compare new solutions with.

Note that just before the submission of this work [Nair et al. \(2023\)](#) published their BOXRR dataset with motion data from over 100,000 online players of the VR games “Beatsaber” and “Tilt Brush.” Despite this exciting development, the “Who is Alyx” dataset continues to be a highly valuable resource for researchers investigating XR user behavior and biometric identification. While our dataset comprises a smaller cohort of 71 users, it still stands out due to several distinctive attributes. Firstly, the data collection process was conducted in a controlled environment, ensuring precise hardware setup and confirming each participant’s active involvement in the XR experience of *Half-Life: Alyx* on two separate days. Notably, the “Who is Alyx” dataset extends beyond mere motion tracking, encompassing not only screen recordings, but additional biometric data such as eye-tracking and physiological signals. This comprehensive approach offers a holistic perspective on users’ physiological responses during XR interactions. Lastly, the two datasets vary in levels of contextualization in the actions performed by users between highly- to medium contextualized (BOXRR) and marginally contextualized (ours) actions, diversifying the data available to the research community. As a result, despite the lower overall user count, the “Who is Alyx” dataset contributes essential and complementary insights to the publicly available XR datasets for researchers investigating XR user identification, biometrics, and behavioral analysis.

3 Dataset

In the following we describe the creation process of the dataset and provide details about the VR game.

3.1 The game

We selected the action game “*Half-Life: Alyx*” ([Valve Corporation, 2020](#)), because it satisfies two of our main objectives for our dataset. First, the game was specifically designed for virtual reality and requires a wide variety of interactions from the players. Second, “*Half-Life: Alyx*” is very entertaining, which made it easier to recruit enough participants that were motivated not only to keep actively engaged throughout one session but also to come back for the second one. In fact, the drop out rate was fairly low with just 5 (out of 76) participants attending only one session.

In the game the player assumes the role of Alyx Vance, who has to rescue her father. Each participant starts the first session in “Chapter 1,” which introduces the player to the core mechanics of the game: navigation, interacting with the environment, and combat. For navigation, the player can either walk around within the limits of the VR setup or use the trackpad of the left controller to teleport to another location nearby ([Figure 1D](#)). Players can grab virtual objects with either hand using the trigger buttons on each controller. Additionally, the game also provides a mechanism to quickly retrieve distant objects using the character’s “gravity gloves”: the player can point at a distant object, press the trigger button

([Figure 1A](#)), and perform a flicking motion with the hand, which repels the object towards the player ([Figure 1B](#)). The player then has to snatch the object out of the air at the right moment using the trigger button again ([Figure 1C](#)). For combat situations the player can select a weapon for the right hand and then shoot with the trigger button ([Figure 1E](#)). To reload a weapon, the player has to press the right controller’s grip button to eject the empty mag, grab behind the shoulders with the left hand to retrieve a new mag from the virtual backpack, insert it, and pull back the rag ([Figure 1F](#)).

After learning these basics, the player ventures on to find Alyx’ father. For this, the game leads through a linear level design and confronts the player with different types of obstacles, such as difficult terrain, bared throughways, or blockading enemies. To solve these challenges, the player has to find the right way, figure out opening mechanisms or engage in combat. There are several mechanisms required to unlock doors or activate machines that reoccur 2 to 4 times each per session: on some occasions, players have to use a hand-scanner to find and fix circuits hidden inside walls ([Figure 1G](#)). To activate upgrade machines, players have to solve 3D puzzles ([Figure 1H](#)). These machines are not mandatory to progress in the game and can be hidden within the level, so the player can ignore them. While the nature of these mechanisms remains the same during each encounter, they become increasingly challenging and require a fair degree of trial and error every time. Additionally, there are several unique challenges that can require a wide range of actions from the player (like [Figure 1I](#)).

3.2 Ethical considerations

The data acquisition study was approved by the institution’s responsible ethics committee. The participants were recruited through the participant recruitment system of our faculty and gave their full consent to publish the collected and anonymized data. Every participant was fully informed about the intents and purposes and the procedure of the data acquisition.

3.3 Setup

The participants were instructed to play the game with an HTC Vive Pro (with 3 lighthouses) within an 10m² area in our lab. An instructor was present all the time and monitored the participant in the real and virtual world, ensuring that the game was played continuously. All participants start their first session in “Chapter 1” and are then free to play the game without any further instructions. After about 45 min, we end the session regardless of where the player was in the game at that time. For the second session, we repeat the same procedure, but change the starting point for each player to “Chapter 3.”

For the first 41 participants we only recorded the tracking data from the head mounted display (HMD) and controllers, along with the eye-tracking data. We noticed an interest within our research group to also collect physiological data, as this makes the dataset viable for other research areas as well. Therefore, we started to have the following 35 participants wear an Empatica E4 wristband and a Polar H10 chest strap. These wearable devices collect physiological data, are easy to setup, and do not restrict the players in their motions—essential requirements as the dataset continues to be



FIGURE 1

Different scenes from the game, illustrating the wide range of different actions required from the players. (A) Gravity grab: player targets object; (B) Gravity grab: hand flip propels object; (C) Gravity grab: player catches object; (D) Player teleports by pressing the trackpad of the left controller; (E) Right hand holds weapons and tools; (F) Player drops old mag to replace it with new one to reload weapon; (G) Player has to find and fix circuits in the wall to unlock barriers; (H) 3D puzzles like this involve both hands and unlock bonuses; (I) Player rolls up a rope and uses a tube to prevent it from unrolling again. Reproduced with permission from Valve Corporation.

intended primarily for research questions around the identification of XR users by their motions.

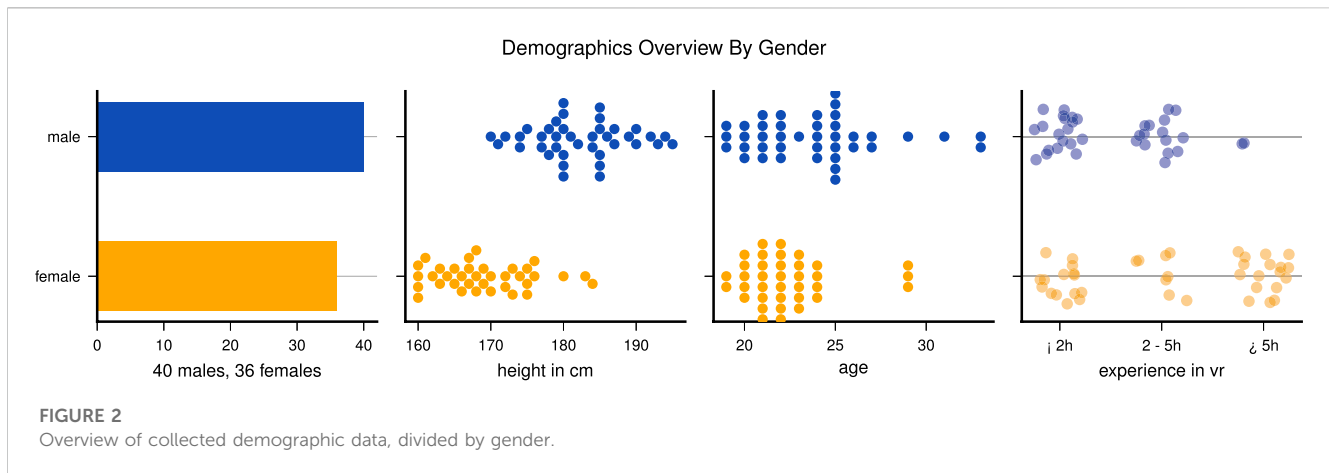
3.4 Recording

Table 2 provides an overview of devices, collected data streams, sampling rates, and the number of recordings per device, as well as the captured physiological features. The sampling rates vary between each data stream, depending on the hard- and software we used for recording. We recorded the tracking data of the first 37 players with an average frequency of 15 frames per second. After a revision of our recording script we increased the sampling rate to 90 frames per second and added tracking of the physiological data for the following 34 players. Since recording the movement data was the main objective, we did not interrupt the session if any of the other data streams failed for any reason, which is why there are some missing eye-tracking and physiological recordings. In Table 2, we count recordings that provide at least 15 min of footage and document any recording faults in the readme file of the data repository.

To retrieve and record the tracking data from the HTC Vive Pro, we use a Python script and the library “openvr” (Bruns, 2022).

TABLE 2 Overview of recorded data.

Device	Sampling rate	1 session	2 sessions
HTC Vive Pro (motions)	Combined	5	71
	15 Hz	4	37
	90 Hz	1	34
Eye-tracking	80 Hz	8	67
Empatica E4 (wristband)		8	27
-Acceleration	32 Hz		
-Heart Rate	1 Hz		
-Electrodermal Activity	4 Hz		
-Photoplethysmography	64 Hz		
-Inter-beat interval	64 Hz		
-Peripheral Body Temp.	4 Hz		
Polar H10 (chest strap)		3	31
-Acceleration	200 Hz		
-Electrocardiogram	130 Hz		



The tracking data includes the positions (x, y, z , in centimeters) and rotations (quaternions: x, y, z, w) for the HMD, left and right controller. The coordinate system is set to Unity defaults, so the “ y ”-axis points upwards. We use HTC’s “SRanipal” to retrieve the eye-tracking data and Unity for calibration and recording.

3.5 Demographics

At the beginning of the first session we collect demographic data and body properties. Altogether there were 36 females and 40 males, all between 21 and 33 years old. [Figure 2](#) summarizes details about age, body height and VR experience. We make all of the collected information available in a summary CSV file in the data repository.

3.6 Insights

In the following, we focus on the tracking data to stay within the scope of this work and leave the analysis of the eye-tracking and physiological data to future work. Altogether, we have recorded 71 participants over the course of two sessions, 5 participants attended just one session, so in sum the dataset includes tracking data of 76 different individuals. Each session is on average 46 min long, with the three shortest sessions lasting between 32 and 40 min.

While screening the dataset, we tried to gain basic insights into how far the participants showed the same or differing behaviour during playing. For basic insights, we plot the travel path of the HMDs on the horizontal axes in [Figure 3](#), which reveals different play styles. Some players tend to stand relatively still without moving much, while others like to walk around or do extensive head movements throughout one session. Interestingly, there was no strong correlation between meters traveled and teleportations made, as indicated by a Pearson correlation coefficient of only 0.25. [Figure 4](#) compares the navigation behavior of the players in both sessions and shows that the players used the teleportation mechanism more frequently in the second session, indicating a learning effect. Also, we noticed that male participants generally used the navigation mechanisms more extensively: males moved 6.9 m and teleported 11 times per minute, while females only moved 5.3 m and teleported 8.5 times per minute on average.

4 User identification deep learning benchmark

We apply two state-of-the-art deep learning models on the dataset to solve the identification task of XR users by their motions. For this, we adopt the methodology of previous works ([Mathis et al., 2020](#); [Rack et al., 2022](#)) as foundation for an experimental setup that evaluates the generalizability of our models in a multi-session setting. With the reported results and subsequent analyses of our model, the results can serve as a baseline for future work that use our dataset as benchmark.

4.1 Methodology

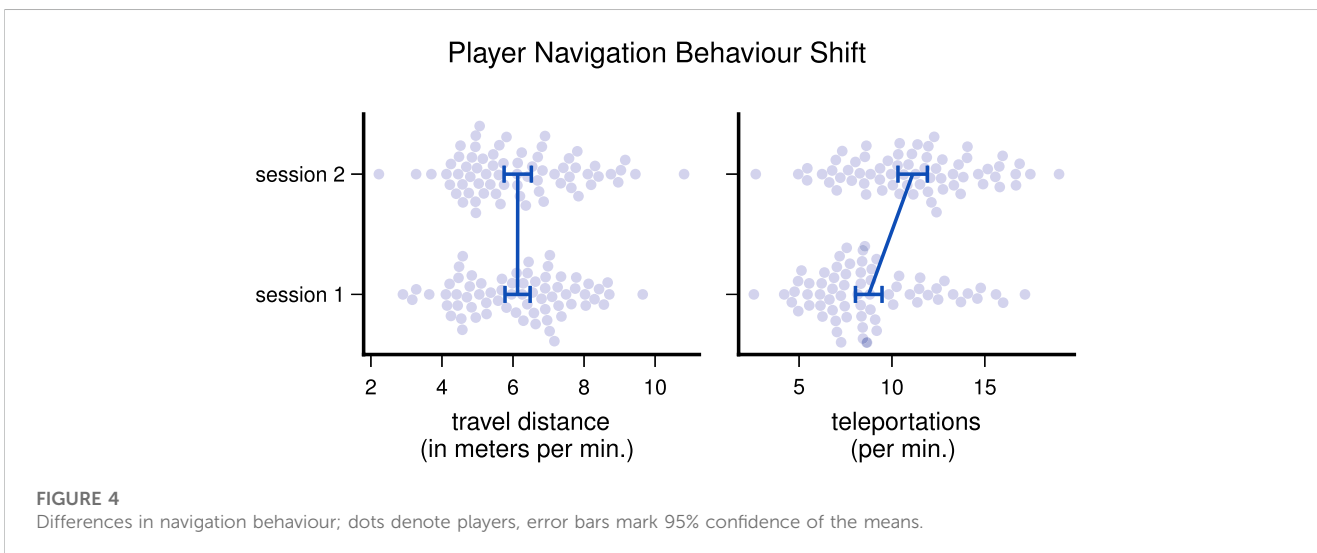
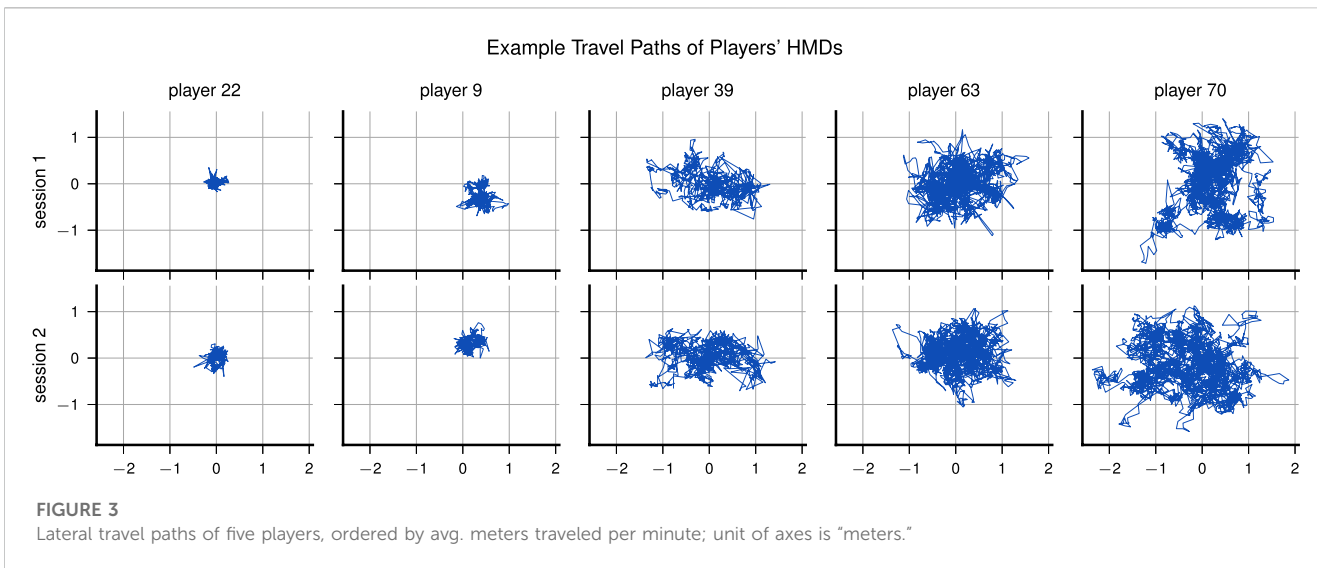
4.1.1 Architecture

For the benchmark, we use a recurrent neural network (RNN), as proposed by [Liebers et al. \(2021\)](#) and [Rack et al. \(2022\)](#), and a convolutional neural network (CNN), as proposed by [Mathis et al. \(2020\)](#) and [Miller et al. \(2021\)](#). Both architectures are able to work directly with sequential data and do not require dimensionality reduction along the time-axis during data preprocessing, as would be the case with multilayer perceptrons, random forests, or the like.

Our choice of the specific RNN architecture is guided by [Rack et al. \(2022\)](#), who compared different types of RNNs. We select the GRU architecture, as it worked equally well as the more commonly used LSTMs, but is a bit faster to train due to a slightly smaller architecture. For the CNN we select an architecture similar to the “FCNN” proposed by [Mathis et al. \(2020\)](#). Additionally, we append a dropout unit to each layer, as we found that this helps to counteract overfitting ([Srivastava et al., 2014](#)). We do not copy the hyperparameter settings used by [Rack et al. \(2022\)](#) or [Mathis et al. \(2020\)](#), but instead perform a hyperparameter search to evaluate a range of different settings (see [Section 4.2.3](#)) and report the final configurations for both architectures in [Table 4](#).

4.1.2 Data encoding

We follow the methods investigated by [Rack et al. \(2022\)](#) for encoding tracking data. As the authors demonstrate, it is necessary to encode the raw tracking data to remove information that does not contain straightforward user characteristics. In our work, we compare the following three different data encodings.



Raw tracking data returned by the VR setup is referenced to some origin relative to the VR scene, which Rack et al. (2022) refer to as scene-relative (SR) encoded data. SR data encode a lot of information not directly tied to the identity of the user, such as user position or orientation, which makes it difficult for machine learning models to learn actual user-specific characteristics. To remove any information about the exact whereabouts of the user, we transform the coordinate system from SR to *body-relative* (BR): we select the HMD as the frame of reference and recompute positions and orientations of the hand controllers accordingly. Since the HMD is always the origin in this coordinate system, its positional features and rotation around the up axis become obsolete and are removed. This only leaves one quaternion encoding the head's rotation around the horizontal axes. Subsequently, there are 18 features for each frame: (pos-x, pos-y, pos-z, rot-x, rot-y, rot-z, rot-w) × (controller-left, controller-right) + (rot-x, rot-y, rot-z, rot-w) of the HMD, all given with respect to the HMD as frame of reference.

Following this, Rack et al. (2022) found that computing the velocities for each feature from BR data yields even better results, since this removes more noise from the data that otherwise confuses the

models. This *body-relative velocity* (BRV) data is computed differently for positions and orientations: for positions, we compute the difference between frames for each feature individually. For rotations, we rotate each quaternion by the corresponding inverted rotation of the previous frame, which yields the delta rotation. This procedure does not change the number of features, so BRV data also consist of 18 features.

Since the BRV encoding produced the best results for Rack et al., we take this train of thought one step further and evaluate the *body-relative acceleration* (BRA) data encoding. For the conversion from BRV to BRA we apply the same steps as previously performed for the conversion from BR to BRV. Like BRV, BRA data consists of 18 features.

4.1.3 Identification procedure

We train the GRU and CNN to map input motion sequences to one of the 71 users. For this we use 300 consecutive frames (about 20 s) of tracking data of one recording (i.e., one user) as input for the model. Since recording lengths during testing consist of much more than 300 frames, we employ a majority voting scheme to produce a prediction for the whole sequence: we retrieve individual

subsequences from the whole sequence by moving a sliding window frame by frame, apply the model on each subsequence and in the end predict the user who was most commonly inferred.

4.1.4 Implementation and training

We use the machine learning framework PyTorch Lightning for implementing both network architectures. Both models are implemented and trained as classification networks with a categorical cross-entropy loss function and the Adam optimizer. We train each model for a minimum of 30 epochs until the monitored minimum accuracy on the validation set deteriorates or stagnates. We save a snapshot of each model at its validation highpoint for later evaluation, since performance can already start declining towards the end of the 30 epochs. We use PyTorch Lightning (Falcon et al., 2020) for implementation. Each job runs with 8 CPU units, one GPU (either NVIDIA GTX 1080Ti, RTX 2070 Ti or RTX 2080Ti), and 20 Gb RAM.

4.2 Experimental setup

In our setup we consider a scenario where our system observes VR users playing the game Half-Life: Alyx for some time on 1 day, and then reidentifies them when they play the game again on another day. We follow the terminology used by the literature (Rogers et al., 2015; Miller et al., 2022a, Miller et al., 2022b), which describes the data collected during on-boarding of users as *enrollment* and the data collected for later identification as *use-time* data. In our analyses, we take care to simulate a real-world user identification scenario: the amount for both, enrollment and use-time data that can be captured per user can be very different depending on the individual use case. Therefore, we limit the available amount of enrollment and use-time data to different lengths and vary this limit in our experiments (e.g., identifying a user within 5 min of use-time data based on 10 min of enrollment data). More specifically, we evaluate enrollment and use-time data of 1, 5, 10, 15, 20, and 25 min per user, as well as all available data. This provides us with evidence of how well the evaluated models perform in different use cases. In the following, we describe different aspects of our experimental setup in more detail.

4.2.1 Data preprocessing

We chose a sequence length of 300 frames with a frequency of 15 frames per second (fps) as input for the neural networks, which resembles a duration of $\frac{300\text{frames}}{15\text{fps}} = 20$ s. This is based on the following considerations. First, we prefer shorter sequences, since the training time considerably increases with larger sequence lengths, especially for the GRU architecture. Second, Rack et al. (2022) demonstrate that it is preferable for samples to cover a longer duration with a lower frequency than a shorter duration with a higher frequency. Lastly, 15 fps matches the frequency of the first half of the dataset (see Section 3.4). Consequently, we resample the tracking data to a constant 15 fps for all users and then sample individual sequences of 300 frames as model inputs.

From each session, we remove the first and the last minute to remove potential artifacts that could have been caused by participants still getting ready or the recording stopping just after exiting the game. Then, we divide the dataset into training,

TABLE 3 Hyperparameter search spaces.

	Hyperparameter	Search space
GRU	Hidden size	(20, 450) neurons
	Number of layers	(1, 8)
	Dropout	(0, 0.6)
	Learning rate	(0.0001, 0.01)
CNN	Kernel sizes	(2, 9)
	Number of layers	(1, 8)
	Channel sizes	(10, 500)
	Dropout	(0, 0.6)
	Learning rate	(0.0001, 0.01)

validation and testing sets. In our case we consider session one for enrollment, hence we divide that session into training and validation split: we select the last 5 min from session one for validation, and the rest for training. The use-time data serves for the final evaluation, so we use entire session two for the test set.

4.2.2 Evaluation metric

For evaluation, we consider the *macro averaged accuracy*, which reports the ratio of correct predictions to the overall predictions made. We use macro-averaging, which takes the accuracy of each class (i.e., user) individually and then computes the average score over all classes. This way we account for eventual class imbalances, even though the footage for each user is fairly balanced.

During our analysis, we compare different lengths of enrollment data t_{enr} per user with different lengths of use-time data t_{use} per user. To report the accuracy score for the various combinations, we use the format $Acc_{\text{use}: t_{\text{use}}}^{\text{enr}: t_{\text{enr}}}$. For example, if a model correctly identified 77% of all 5 min sequences from the test (i.e., use-time) data after being trained on 10 min of enrollment data per user, we write $Acc_{\text{use}: 5}^{\text{enr}: 10} = 0.77$.

For the hyperparameter search, we employ the *minimum accuracy* as key metric to compare models. We do not use the macro averaged accuracy here, because we noticed during preliminary runs that models with similar validation accuracies can still be different in how consistently they identify each class. Some models identified some classes very well, but scored very low for a few, while others performed well for all classes. We think that the latter is preferable, since this allows to eventually identify each class reliably given enough use-time data. The minimum accuracy considers this by taking the individual accuracies of each class and returning the score of the class with the lowest accuracy. This way we prefer models that generalize well for all classes, and dismiss models that neglect individual classes.

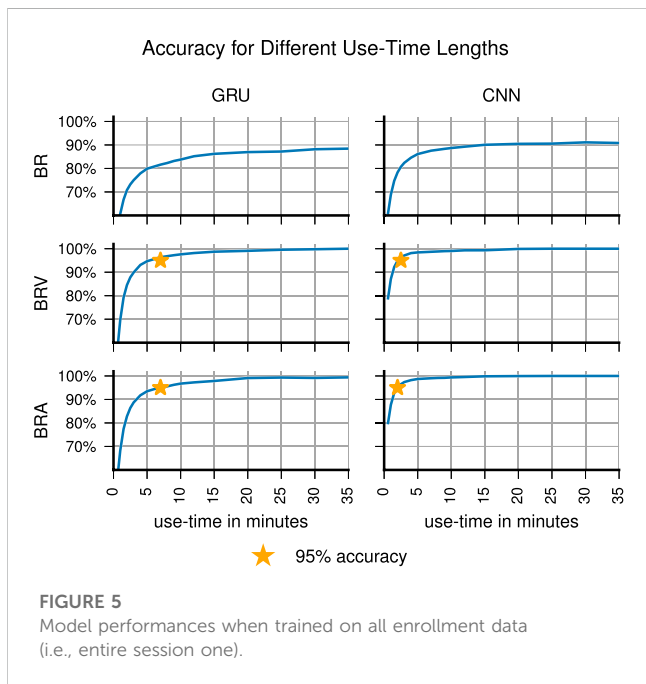
4.2.3 Hyperparameter search

We perform a hyperparameter search for the architectures in combination with each data encoding, which results in six separate searches [i.e., (GRU, CNN) \times (BR, BRV, BRA)]. Table 3 lists the parameters we considered for both architectures.

We use the online service Weights and Biases (Biewald, 2020). The service monitors the training procedure and coordinates the

TABLE 4 Final model configurations after hyperparameter search.

	GRU				CNN				
	Learning rate	Layers	Hidden size	Dropout	Learning rate	Layers	Kernel size	Channel sizes	Dropout
BR	0.0005	3	300	0.32	0.002	4	3	400, 580, 841,1219	0.46
BRV	0.0007	4	250	0.32	0.002	3	3	400; 480; 576	0.42
BRA	0.0005	4	400	0.10	0.002	4	3	600; 900; 1350; 2025	0.44

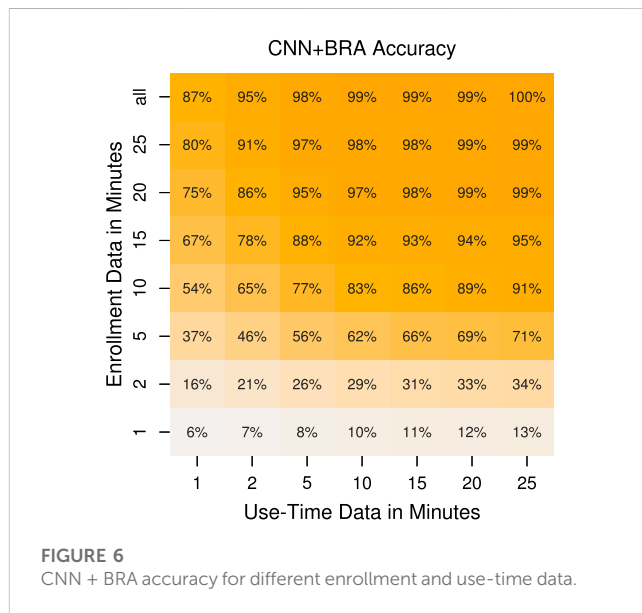


search. Here, we select the Bayesian search method, which observes a target metric during the search to systematically propose new configurations. As target metric we use the minimum accuracy.

4.3 Results

We performed about 100 individual training runs for each architecture and data encoding combination, totaling 600 runs. For each combination we pick the configuration that achieved the highest minimum accuracy score on the validation data. The final configurations are listed in Table 4.

First, we compare the six models on the test data and pick the best model for further analysis. As can be expected, identification accuracy improves when more use-time data is considered for a prediction. This becomes visible in Figure 5 which shows the correlation between amount of use-time data and identification accuracy for each model. Overall, CNNs worked better than GRUs, and BRA worked better than BRV and BR. The CNN + BRA combination performed best and achieved 95% accuracy within 2 min ($Acc_{use:2}^{enr:all}$) and 99% accuracy within 7 min ($Acc_{use:7}^{enr:all}$) of use-time data. The CNN architecture not only outperformed the GRU architecture, but also trained much faster: the CNN + BRA model was trained within 2 h, while the GRU + BRA model



required 11.5 h. Since the BRA + CNN model worked best, we select it for further analysis.

Up to this point we used entire session one for enrollment, which is about 45 min of training data per user. Depending on the target use case, it might not always be feasible to capture that much footage for each user. Therefore, we re-train the model on shorter lengths of enrollment data, ranging from 1 to 25 min. The starting point of the enrollment data within the recording of each user was selected randomly, so we repeat the training for each enrollment length five times and report the mean values in Figure 6.

Having a larger enrollment set appears to be more important than having a larger use-time set. The accuracy for 1 min enrollment and 10 min of use-time achieves only $Acc_{use:10}^{enr:1} = 10\%$, while 10 min enrollment and 1 min of use-time achieves $Acc_{use:1}^{enr:10} = 54\%$. Doubling the amount of enrollment data from 1 to 2 min already improves the accuracy significantly to $Acc_{use:10}^{enr:2} = 29\%$. This implies that if there is too little enrollment data, the model cannot make accurate predictions, even with more use-time data. For instance, there is little difference between 1 and 25 min of use-time data, as long as there is only 1 min of enrollment data ($Acc_{use:1}^{enr:1} = 6\%$ and $Acc_{use:25}^{enr:1} = 13\%$). However, after reaching a certain amount of enrollment data, lower prediction accuracies can be compensated with more use-time data, e.g., $Acc_{use:1}^{enr:5} = 37\%$ compared to $Acc_{use:25}^{enr:5} = 71\%$.

Random factors, such as the random initialization of the network and the GPU architecture, can impact the outcome of

the training process. To test its robustness against these factors, the CNN + BRA model was retrained 10 times with different seeds for random initialization. The identification accuracy of individual input sequences, (i.e., $Acc_{use: 20s}^{enr: all}$), varied by only 1.7 percentage points, ranging from 76.6% to 78.3%. In comparison, a similar test during the investigation of the results of Liebers et al. (2021) resulted in variations of up to 15 percentage points.

5 Discussion and conclusion

With this article, we introduce a new dataset with biometric and behavioral data from VR users to address a critical issue in this field. The lack of such datasets has hindered the development of larger and more accurate machine learning models, making it difficult to compare new methods to current ones or reproduce previous solutions. We believe that publishing our dataset is an important step towards improving the reproducibility and reliability of machine learning methods in XR user identification and authentication research.

Both tested state of the art deep learning models were able to identify users in session two after training on session one. All models performed much better than a random classifier would which validates that the tracking data provides identifying user-specific characteristics. The best model was a CNN architecture that could identify 95% of all 2 min sequences from session two correctly when trained on all data from session one. The CNN architecture also was trained fastest and generally only needed about half the time to reach peak performance compared to the RNN. This performance deviates just slightly (i.e., below ± 1 percentage point of accuracy) if the model is trained with different seeds, which indicates that models can be trained and evaluated in a robust way with the used dataset and methods.

The results also reveal that the acceleration encoding (BRA) yields superior performance compared to BR and BRV. This finding supports the hypothesis put forth by Rack et al. (2022) that abstracting non-motion-related information can enhance the generalization capabilities of the model. We hypothesize that more sophisticated encoding techniques could potentially lead to even better identification models. Thus, this remains a promising area for future research.

It is crucial for the scientific community to have access to diverse and well-documented datasets in order to facilitate progress and advance research. Therefore, we strongly urge authors of future work to make their datasets publicly available, along with any relevant information about data collection and processing. This will enable other researchers to reproduce the results, conduct further analysis, and build on previous work, ultimately leading to a deeper understanding of the field.

A limitation of the dataset is that it contains only marginally contextualized motion sequences of users being passively watched while they engage with the VR application. While this is ideal to explore identification models that can work by passively watching XR users, it is less suited to explore active verification use cases, as these usually expect the user to actively provide a short input sample. However, the preliminary work of Rack et al. (2023) suggests that modern embedding-based models are able to generalize once trained on the “Who Is Alyx?” dataset to a different dataset exhibiting different motions. Consequently, one or more large datasets, like “Who Is Alyx?,” could potentially be used to train powerful multi-

purpose deep learning models, which can then be used not only for identification, but also for verification of users.

This article introduces the “Who Is Alyx?” dataset, specifically designed for investigating user identification based on motion in XR environments. Using two state-of-the-art deep learning architectures, we have established an initial benchmark that future research can build upon. The breadth of our dataset extends beyond the scope of this article, as “Who Is Alyx?” encompasses a wide array of unexplored features that we believe hold promise for further exploration. We invite the scientific community to utilize this dataset in their research and welcome inquiries, feedback, and suggestions from researchers working with our dataset. Furthermore, we have plans to continue extending “Who Is Alyx?” and collect data from other scenarios as well. In doing so, we hope to facilitate the creation of additional datasets encompassing a wider range of XR contexts, pushing towards novel and reliable user recognition systems to improve security and usability in XR.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://github.com/cshell/who-is-alyx>.

Ethics statement

The studies involving humans were approved by the Research Ethics Committee of the Mensch-Computer-Medien faculty of human sciences of the Julius-Maximilians-University Würzburg. The studies were conducted in accordance with the local legislation and institutional requirements. The participants provided their written informed consent to participate in this study.

Author contributions

CR: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Project administration, Software, Validation, Visualization, Writing—original draft, Writing—review and editing. TF: Investigation, Software, Writing—review and editing. MY: Data curation, Writing—review and editing. AH: Resources, Supervision, Writing—review and editing. ML: Resources, Supervision, Writing—review and editing.

Funding

The author(s) declare that no financial support was received for the research, authorship, and/or publication of this article.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Achenbach, J., Waltemate, T., Latoschik, M. E., and Botsch, M. (2017). "Fast generation of realistic virtual humans," in *Proceedings of the ACM symposium on virtual reality software and technology, VRST (ACM)*, F1319. doi:10.1145/3139131.3139154
- Ajit, A., Banerjee, N. K., and Banerjee, S. (2019). "Combining pairwise feature matches from device trajectories for biometric authentication in virtual reality environments," in 2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR), San Diego, CA, USA, 09-11 December 2019 (IEEE), 9–16. doi:10.1109/AIVR46125.2019.00012
- Bhalla, A., Sluganovic, I., Krawiecka, K., and Martinovic, I. (2021). "MoveAR: continuous biometric authentication for augmented reality headsets," in *CPSS 2021 - proceedings of the 7th ACM cyber-physical system security workshop (ACM)*, 41–52. doi:10.1145/3457339.3457983
- Biewald, L. (2020). Experiment tracking with weights and biases. Available at: <https://www.wandb.com/>.
- Bruns, C. (2022). Pyopenvr. Available at: <https://github.com/cnbruns/pyopenvr>.
- Devlin, J., Chang, M. W., Lee, K., and Toutanova, K. (2019). *BERT: pre-training of deep bidirectional transformers for language understanding*. arXiv.
- Dube, T. J., and Arif, A. S. (2019). "Text entry in virtual reality: a comprehensive review of the literature," in *HCI 2019: human-computer interaction. Recognition and interaction technologies (Cham: Springer)*, 419–437. doi:10.1007/978-3-030-22643-5_33/TABLES/5
- Falcon, W., Borovec, J., Wälchli, A., Eggert, N., Schock, J., Jordan, J., et al. (2020). *PyTorchLightning/pytorch-lightning*. Genève, Switzerland: Zenodo. doi:10.5281/ZENODO.3828935
- Jain, A. K., Ross, A. A., and Nandakumar, K. (2011). *Introduction to biometrics*. US: Springer. doi:10.1007/978-0-387-77326-1
- Kapoor, S., and Narayanan, A. (2022). Leakage and the reproducibility crisis in ML-based science. arXiv.
- Kern, F., Kullmann, P., Ganal, E., Korwisi, K., Stingl, R., Niebling, F., et al. (2021). Off-the-shelf stylus: using XR devices for handwriting and sketching on physically aligned virtual surfaces. *Front. Virtual Real.* 2, 684498. doi:10.3389/frvir.2021.684498
- Knierim, P., Schwind, V., Feit, A. M., Nieuwenhuizen, F., and Henze, N. (2018). "Physical keyboards in Virtual reality: analysis of typing performance and effects of avatar hands," in *Conference on human factors in computing systems - proceedings (ACM)*. doi:10.1145/3173574.3173919
- Kupin, A., Moeller, B., Jiang, Y., Banerjee, N. K., and Banerjee, S. (2019). *Task-driven biometric authentication of users in virtual reality (VR) environments*. Cham: Springer International Publishing.
- Lee, G., Deng, Z., Ma, S., Shiratori, T., Srinivasa, S., and Sheikh, Y. (2019). "Talking with hands 16.2M: a large-scale dataset of synchronized body-finger motion and audio for conversational motion analysis and synthesis," in 2019 IEEE/CVF International Conference on Computer Vision (ICCV), Seoul, Korea (South), 27 October 2019 - 02 November 2019 (IEEE), 763–772. doi:10.1109/ICCV.2019.00085
- Li, S., Ashok, A., Zhang, Y., Xu, C., Lindqvist, J., and Gruteser, M. (2016). "Whose move is it anyway? Authenticating smart wearable devices using unique head movement patterns," in 2016 IEEE International Conference on Pervasive Computing and Communications, PerCom, Sydney, NSW, Australia, 14-19 March 2016 (IEEE), 1–9. doi:10.1109/PERCOM.2016.7456514
- Liebers, J., Abdelaziz, M., and Mecke, L. (2021). "Understanding user identification in virtual reality through behavioral biometrics and the effect of body normalization," in *Conference on human factors in computing systems - proceedings (ACM)*, 1–11. doi:10.1145/3411764.3445528
- Mathis, F., Fawaz, H. I., and Khamis, M. (2020). "Knowledge-driven biometric authentication in virtual reality," in *Extended abstracts of the 2020 CHI conference on human factors in computing systems (ACM)*. doi:10.1145/3334480.3382799
- Miller, R., Banerjee, N. K., and Banerjee, S. (2020). "Within-system and cross-system behavior-based biometric authentication in virtual reality," in *Proceedings - 2020 IEEE Conference on Virtual Reality and 3D User Interfaces, VRW, Atlanta, GA, USA, 22-26 March 2020 (IEEE)*, 311–316. doi:10.1109/VRW50115.2020.00070
- Miller, R., Banerjee, N. K., and Banerjee, S. (2021). "Using siamese neural networks to perform cross-system behavioral authentication in virtual reality," in 2021 IEEE Virtual Reality and 3D User Interfaces (VR), Lisboa, Portugal, 27 March 2021 - 01 April 2021 (IEEE), 140–149. doi:10.1109/VR50410.2021.00035
- Miller, R., Banerjee, N. K., and Banerjee, S. (2022a). "Combining real-world constraints on user behavior with deep neural networks for virtual reality (VR) biometrics," in *Proceedings - 2022 IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2022, Christchurch, New Zealand, 12-16 March 2022 (IEEE)*, 409–418. doi:10.1109/VR51125.2022.00060
- Miller, R., Banerjee, N. K., and Banerjee, S. (2022b). "Temporal effects in motion behavior for virtual reality (VR) biometrics," in *Proceedings - 2022 IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2022, Christchurch, New Zealand, 12-16 March 2022 (IEEE)*, 563–572. doi:10.1109/VR51125.2022.00076
- Moore, A. G., McMahan, R. P., Dong, H., and Ruozzi, N. (2021). "Personal identifiability of user tracking data during VR training," in 2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW), Lisboa, Portugal, 27 March 2021 - 01 April 2021 (IEEE), 556–557. doi:10.1109/VRW52623.2021.00160
- Mustafa, T., Matovu, R., Serwadda, A., and Muirhead, N. (2018). "Unsure how to authenticate on your VR headset? Come on, use your head," in *Iwspa 2018 - proceedings of the 4th ACM international workshop on security and privacy analytics, Co-located with CODASPY 2018 (Lubbock, TX, USA: ACM)*, 23–30. doi:10.1145/3180445.3180450
- Nair, V., Guo, W., Wang, R., O'Brien, J. F., Rosenberg, L., and Song, D. (2023). Berkeley open extended reality recording dataset 2023 (BOXRR-23). arXiv. doi:10.25350/B5NP4V
- Olade, I., Fleming, C., and Liang, H. N. (2020). Biomove: biometric user identification from human kinesiological movements for virtual reality systems. *Sensors Switz.* 20, 2944–3019. doi:10.3390/s20102944
- OpenAI (2023). Chatgpt-4. Available at: openai.com/gpt-4.
- Pfeuffer, K., Geiger, M. J., Prange, S., Mecke, L., Buschek, D., and Alt, F. (2019). "Behavioural biometrics in VR: identifying people from body motion and relations in virtual reality," in *2019 CHI conference on human factors in computing systems - CHI '19 (ACM)*, 1–12.
- Pineau, J., Vincent-Lamarre, P., Sinha, K., Larivière, V., Beygelzimer, A., d'Alché-Buc, F., et al. (2020). Improving reproducibility in machine learning research (A report from the NeurIPS 2019 reproducibility program). arXiv.
- Rack, C., Hotho, A., and Latoschik, M. E. (2022). "Comparison of data encodings and machine learning architectures for user identification on arbitrary motion sequences," in 2022 IEEE International Conference on Artificial Intelligence and Virtual Reality, AIVR 2022, CA, USA, 12-14 December 2022 (IEEE).
- Rack, C., Kobs, K., Fernando, T., Hotho, A., and Latoschik, M. E. (2023). Extensible motion-based identification of XR users using non-specific motion data. arXiv.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical text-conditional image generation with CLIP latents. arXiv.
- Rogers, C. E., Witt, A. W., Solomon, A. D., and Venkatasubramanian, K. K. (2015). "An approach for user identification for head-mounted displays," in *ISWC 2015 - proceedings of the 2015 ACM international symposium on wearable computers (ACM)*, 143–214. doi:10.1145/2802083.2808391
- Serra-García, M., and Gneezy, U. (2021). Nonreplicable publications are cited more than replicable ones. *Sci. Adv.* 7, eabd1705. doi:10.1126/sciadv.abd1705
- Shen, Y., Wen, H., Luo, C., Xu, W., Zhang, T., Hu, W., et al. (2019). GaitLock: protect virtual and augmented reality headsets using gait. *IEEE Trans. Dependable Secure Comput.* 16, 484–497. doi:10.1109/TDSC.2018.2800048
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958. doi:10.5555/2627435.2670313
- Stephenson, S., Pal, B., Fan, S., Fernandes, E., Zhao, Y., and Chatterjee, R. (2022). "SoK: authentication in augmented and virtual reality," in 2022 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA, 22-26 May 2022 (IEEE), 267–284. doi:10.1109/SP46214.2022.9833742
- Valve Corporation (2020). Half-life: Alyx. Available at: <https://half-life.com>.
- Wang, C. Y., Yeh, I. H., and Liao, H. Y. M. (2021). You only learn one representation: unified network for multiple tasks. arXiv. doi:10.48550/arxiv.2105.04206
- Weninger, S., Achenbach, J., Bartl, A., Latoschik, M. E., and Botsch, M. (2020). "Realistic virtual humans from smartphone videos," in *26th ACM symposium on virtual reality software and technology (ACM)*, 1–11. doi:10.1145/3385956.3418940