User Identification in Virtual Reality through Behavioral Biometrics and the Influence of Colocated Interactions

Jonathan Liebers
HCI Group
University of Duisburg-Essen
Essen, NRW, Germany
jonathan.liebers@uni-due.de

Frieder Sykora
HCI Group
University of Duisburg-Essen
Essen, NRW, Germany
frieder.sykora@stud.uni-due.de

Niklas Pfützenreuter
HCI Group
University of Duisburg-Essen
Essen, NRW, Germany
niklas.pfuetzenreuter@uni-due.de

Uwe Gruenefeld
GenerIO
Essen, NRW, Germany
uwe@generio.ai

David Goedicke
HCI Group
University of Duisburg-Essen
Essen, NRW, Germany
david.goedicke@uni-due.de

Stefan Schneegass
HCI Group
University of Duisburg-Essen
Essen, NRW, Germany
stefan.schneegass@uni-due.de

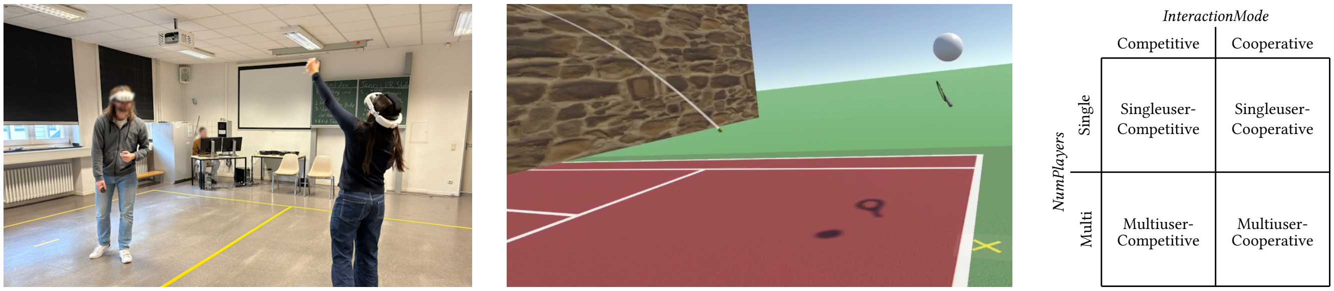

(a) Real Environment.

(b) Shared Virtual Environment.

(c) Our 2 × 2 Study Design.

Figure 1: Impression from our user study in (a). Two participants stand on a tennis court marked on the floor. They interact in a shared virtual environment in a cooperative, Squash-like setting, cf. (b). In our research, we explored the behavioral biometric identifiability of their body movements by manipulating two independent variables through a factorial analysis: *NumUsers* (with levels Singleuser and Multiuser) and *CollaborationMode* (with levels Competitive and Cooperative), as shown in (c).

Abstract

Behavioral Biometrics in Virtual Reality (VR) allow for implicit user identification, as the head- and hand-movements that can be captured from the head-mounted display and the controllers are highly descriptive of the user's true identity. Such body movements have been explored in the past; however, to date, it is unclear how they perform in settings where more than one person interacts in a shared virtual environment. In this work, we explored through a user study (N=40) how behavioral biometrics in VR change when one or more persons interact with each other in a shared virtual environment and whether this is influenced by the nature of the interaction itself. We find that user identification is possible with up to 83.38 % by applying deep learning models, and that particularly cooperative interactions between multiple VR users lead to highly identifiable body movements. Our results help in advancing behavioral biometrics for seamless user identification in VR, as a viable alternative to using PINs and passwords.

CCS Concepts

• **Human-centered computing** → *Human computer interaction (HCI)*; *Virtual reality*; • **Security and privacy** → *Usability in security and privacy*.

Keywords

Authentication, Behavioral Biometrics, Colocated Interactions, User Identification, Virtual Reality

ACM Reference Format:
Jonathan Liebers, Frieder Sykora, Niklas Pfützenreuter, Uwe Gruenefeld, David Goedicke, and Stefan Schneegass. 2025. User Identification in Virtual Reality through Behavioral Biometrics and the Influence of Colocated Interactions. In *31st ACM Symposium on Virtual Reality Software and Technology (VRST '25), November 12–14, 2025, Montreal, QC, Canada.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3756884.3765978

1 Introduction

For most VR applications, it is essential to know *who the user truly is.* Knowing the user's true identity is the central prerequisite to enable both *personalization* and *security.* A personalized VR application has "the ability to provide content and services tailored to individuals", which means that the application can filter and present information in a way that is interesting to the user [15, 68]. Furthermore, a VR application can also load the user's settings

---

VRST '25, Montreal, QC, Canada

ACM ISBN 979-8-4007-2118-2/25/11
https://doi.org/10.1145/3756884.3765978

or profile by knowing the user's identity. Second, having knowledge of the user's identity is important to establish security, as the VR application and system can show personal information, for example, private social media messages, only to legitimate users. By leveraging the information on the user's identity, the system can also refuse to be used by an illegitimate user and thus protect sensitive personal information such as payment data. In general, it is the obligation of the *authentication system* to discover who the user truly is by establishing confidence in the user's identity, and therefore to answer the question, who the user truly is [55, p. 45].

To date, the so-called knowledge-based methods, such as Personal Identification Numbers (PINs) or passwords, are widespread and predominantly used for user authentication. Unfortunately, their usage is associated with a plethora of disadvantages [74]. Users have to perform dozens of authentications per day [37, 38] by entering increasingly complex PINs or passwords which are hard to memorize [77]. In addition, password reuse undermines users' security [64], and at the same time, passwords have been shown to be guessable [20]. In VR, an additional issue appears as a bystander can spy on the hand movements of an immersed VR user to steal their password [9]. Research has shown that users frequently access their devices and, therefore, are frequently bothered by the authentication process [38]. The result is a "great authentication fatigue" [69], leading to PINs and passwords being described as "imperfect authentication" [6] and the question arises whether they are to be considered "dead" [7]. Thus, calls emerge for a shift towards user-centered authentication methods and biometrics [4, 12].

Biometrics can be a solution to these issues. They leverage the user's behavioral or physiological aspects for user authentication [21]. While many biometrics are applicable to VR, there is one form that is native to VR: Kinetic Signatures [33]. Kinetic Signatures are defined to be *traces of users' individual spatiotemporal body behavior that are useful for behavioral biometric user identification* [33]. Essentially, they are the trajectories of the user's hands and head as they can be captured by the Head-Mounted Display (HMD) and hand-held controllers or natural hand-tracking. These behavioral biometric authentication methods have already received attention by the research community. Researchers have explored the movement behavior elicited in games such as Beat Saber [30, 52], Half-Life: Alyx [58], or sports in VR such as ball-throwing [24], archery and bowling [28]. Furthermore, numerous VR activities were explored, such as walking or carrying objects [2, 56], head movements [31, 36, 51], full body kinesiological movements [54], or general movement behavior based on muscle movements throughout a systematic investigation [33].

However, to the best of our knowledge, one important aspect remained unexplored to date and we are the first to investigate it: the effect of more than a single user interacting in a shared virtual environment with respect to their body movement's behavioral biometric identifiability. We hypothesize that a person's behavior changes when interacting with another person in VR. Such a *colocated interaction* can potentially influence the identifiability of their body motion, as the behavior of one person in the virtual environment can depend on the behavior of the other person. In addition, the nature of the interaction can play a role, i.e., whether the colocated interaction is cooperative or competitive. Therefore, we posit our central Research Question (RQ):

RQ How does the number of users in a shared virtual environment and the nature of the colocated interaction influence the biometric identifiability of a Kinetic Signature?

In our work, we explored behavioral biometric user identification in VR during colocated interactions. We conducted a user study with two sessions, where participants participated twice in the same procedure on two different days ($N = 40$). We tested in a 2×2 study design how participants' behavioral biometric identifiability changes when they (a) interact alone or (b) with another person in a shared virtual environment, cf. Figure 1(c). Second, we also tested the nature of participants' interaction, i.e., whether they cooperate during the interaction with an aligned objective or whether they act competitively with a contrary objective. For the interaction itself, our participants engaged in tennis in VR, and we sampled their body movements from the HMD, as tennis has been shown to provide highly identifiable kinetic signatures in previous work [33]. We trained various deep learning models on the behavioral biometric samples, and these models perform with up to 83.38 % accuracy and result in median recall rates of up to 91 % under best conditions. We also captured users' workload and subjective preferences through interviews. Our findings help shape the future of behavioral biometrics, as they deepen the understanding of how a user's body movements can be used for behavioral biometric authentication and what factors influence their identifiability.

> **Contribution Statement [82].** The nature of our contribution is empirical. We present findings on the influence of the number of users involved in a shared virtual environment and the nature of the interaction itself to understand the identifiability of human motion behavior in VR that we explored through a user study (N=40).

2 Foundations and Related Work

Computers, and standalone VR HMDs as such, need to know the true identity of their human user to enable security and personalization. For this reason, users need to perform an "authentication".

2.1 Foudations of Authentication

Authentication is usually implemented through one of three methods: i) knowledge-based methods ("what you know"), such as PINs or passwords, ii) token-based methods ("what you have"), such as a security token or cryptographic secret to possess, or iii) biometrics ("what you are") [55]. The term *authentication* itself is often defined as the "process of confirming an individual's identity, either by verification or by identification" [17]. Thereby, it is an umbrella term. In "verification", the computer verifies an identity that a user claims. For example, when logging into a website, it is usually required that the user provides a 2-tuple of information: (a) a claim-of-identity (e.g., a username or e-mail address that tells which account the user wants to access) and (b) an authenticator (e.g., a corresponding password) [21]. The computer finds the account through the claim-of-identity and correspondingly verifies if the authenticator matches, resulting in an *accept* or *reject* decision [55]. In contrast, in "identification" the authentication takes place only through the authenticator without any claim-of-identity [21]. For example, a fingerprint can be used to directly identify the user, thus the system would return an *identity* or potentially a *reject* [21]. Currently, knowledge-based methods are widespread and the dominant form of authentication, but they have, as aforementioned, a considerable number of associated issues, so calls have emerged to move beyond passwords for future authentication [4].

2.2 Implicit and Continuous Authentication

An important property of biometrics is that they are particularly suited for implicit and continuous authentication. Implicit authentication is "the ability to authenticate [...] users based on actions they would carry out anyway" [22]. It occurs through an implicit interaction which is "an action performed by the user that is not primarily aimed to interact with a computerised system but which such a system understands as input" [70]. Basically, computers, and VR HMDs as such can use the regular behavior of the user that just occurs during natural usage for authentication. The main advantage is that the user is not getting annoyed by having to explicitly deal with an authentication prompt, as the authentication happens transparently, potentially without them noticing it.

A second advantage is that the authentication can happen continuously in the background. Currently, most authentication processes take place before the desired interaction itself. For example, to send an email, one usually needs to first unlock the lock screen of the device before being able to proceed with sending the message. This potentially opens the door to session hijacking, where an adversary takes advantage of the unlocked session. However, with implicit authentication, the computer system can perform continuous authentication [72, 76]. The core principle here is that the user's behavior is implicitly sampled in the background throughout the usage session, and once a strong shift in user behavior is detected, the continuous authentication system can lock the session [72, 76]. Thereby, implicit and continuous authentication relieves the user of the burdens associated with authentication while simultaneously increasing security [72, 76].

2.3 Behavioral Biometric Authentication through Body Movements in VR

Behavioral biometrics in VR gained increasing attention in VR research, where they were used for implicit and continuous user identification. Early work by Kupin et al. introduced the concept of "task-driven biometrics" by analyzing ball-throwing behavior in VR ($N = 12$) [24]. Ajit et al. extended this approach on a larger sample ($N = 33$) and across multiple sessions [1].

Subsequently, Miller et al. introduced a cross-device dataset using Oculus Quest, HTC Vive, and Cosmos, and applied Siamese neural networks [43, 57]. They also introduced real-world constraints into their models [44]. More recently, Li et al. investigated motion forecasting and 3D motion reconstruction from 2D input [26, 27].

Pfeuffer et al. analyzed general movement patterns such as pointing and walking ($N = 22$) and identified inter-user differences based on relative body part positions [56]. Head movement alone proved to be highly distinctive, as it was also confirmed by Mustafa et al. and Sivasamy et al. [36, 51]. Several other approaches explored identification via movement in more complex activities. Abdrabou et al. analyzed tasks like sitting, walking, and object interaction, while Olade et al. focused on object relocation and kinesiological movements [2, 54]. Rack et al. evaluated general-purpose motion representations and published a dataset of 71 users playing "Half-Life: Alyx" [58, 59].

Moore et al. conducted several studies on user identifiability across sessions and domains. They demonstrated that identification rates exceeded 90% within a session but dropped to 40–50% across sessions [48]. Using velocity-based encodings reduced identifiability while retaining task utility [47]. In later work, they compared headset-only models to multi-device configurations and found headset data to better capture physiological uniqueness [46].

Furthermore, Liebers et al. explored the utilization of body normalizations in VR to understand the interplay of behavior and physiology in kinetic signatures [28]. They also explored systematically how different movement components shape and the resulting activities shape identifiability, finding, that particularly highly dynamic, low static activities perform best [33]. Moreover, they conducted an exploration on how behavioral biometrics in VR change over time [30]. Finally, they explored further behavioral modalities, such as the role of head movement behavior during gaze-based authentication and natural hand- and finger behavior [29, 31].

Two large-scale datasets marked exceptions to the smaller sample sizes: Miller et al.'s 360-degree video study with $N = 511$ and Nair et al.'s Beat Saber dataset with over 100,000 users [52, 53]. Finally, with the growing body of research, a few overviews were provided: Stephenson et al. and Garrido et al. presented systematizations of knowledge for authentication and data privacy in VR [49, 73], while Hallal et al. and Giaretta offer literature reviews [14, 16].

Recent studies examined the temporal stability of behavioral biometrics in VR and their dependence on spatiotemporal movement to account for motion learning and changes in behavior. R. Miller et al. reported medium-term changes in motion behavior over 7–18 months [45], while M. R. Miller et al. showed temporal drift in biometric permanence across eight weeks of VR group discussions [42]. Liebers et al. confirmed this trend in a remote field study over eight weeks, where identifiability decreased over time, highlighting the need for retraining of models [30]. These works demonstrate that although spatiotemporal biometrics are distinctive in the short term, their stability diminishes over time, requiring adaptive system designs.

2.4 Multiuser VR & Research Gap

While multiuser interactions and practices in VR already received some attention by researchers [79], such as by Rasch et al. and Weissker et al. conducting experiments on group-based locomotion in VR [62, 80] and undo actions [61], the potential of multiuser interactions for biometric authentication in VR is yet mostly unexplored. Generally, there exists only limited research covering human behavior in colocated interactions and their suitability for user identification, with a few exceptions [65]. Therefore, we explored for the first time the changes imposed on users' biometric kinetic signatures by colocated interactions.

3 Experiment

To explore our research question, "how does the number of users in a shared virtual environment and the nature of the colocated interaction influence the biometric identifiability of a Kinetic Signature?", we conducted a lab-based user-study experiment.

3.1 Hypotheses

In our work, we posit two hypotheses by considering the background of behavioral biometrics and by reasoning on our research question with regard to the foundations of user identification.

> **H1** User behavior in multi-user shared virtual environments yields significantly higher identifiability than behavior recorded in single-user virtual environments.
> **H2** User behavior in cooperative interactions yields significantly higher identifiability than behavior in competitive interactions.

Background. Human behavior needs to fulfill two properties to be suitable for behavioral biometric user identification. The first property is that it is *unique*, which means that it differs from the behavior of other people [21]. Behavioral biometric systems can only distinguish people based on their behavior if the behavior of one person is different from the behavior of another person. For example, a behavior biometric system that distinguishes people by their voice requires that the voices of people be different from one to another. In technical terms, the behavior needs to have a high "inter-class variance", where the term "class" refers to a person's true identity [21]. Additionally, a unique behavior also needs to have a high degree of *stability.* A behavior with a high degree of stability is a behavior that does not change a lot; this property is sometimes synonymously referred to as "permanence" [21]. For example, a voice-based identification system can only reliably identify its user if their voice stays approximately the same – if the voice changes abruptly during the course of a single day, it would likely lead to the voice sample being rejected by the system. In technical terms, this requires a low degree of "intra-class variance" [21].

Reasoning. With regard to our hypotheses, we reason that the behavior of two persons who actively interact with each other in a shared virtual environment tends to mutually adapt to one another. Simply speaking, their behavior is shaped to fit each other, as the behavior of one person is not only formed by themselves but is also formed with respect to the other person that they interact with. We also posit that this might give rise to a joint behavioral signature, where the two involved behaviors are shaped in interplay with one another. The resulting signature takes a *unique* form, as it reflects both each person's intrinsic behavioral patterns and the way these adapt to their counterpart. Such a joint behavioral signature becomes specific to the dyad of users involved: it is shaped by each person's unique and individual intrinsic behavior, but also by the behavior of their counterpart, since a part of the behavior adapts to the other person and becomes specific to them. Conceptually, this can be seen as an analogy to code-switching in linguistics, where speakers adjust their language use depending on the social context and their interaction partner [13]; in our case, however, we assume that their spatiotemporal body behavior adapts. For instance, in a tennis context, this means that the elicited behavior used for biometric identification does not only come from each player in a doubles team but is also influenced by the doubles itself. People in dyadic interactions often adapt to each other [8], on a social level [39], but also on the level of their body behavior [41, 83]. We believe that the behavior tends to be more *stable*, as random changes in behavior are less likely during an interaction with another person. Furthermore, Cooperative interactions are better identifiable than Competitive interactions, as users might develop a stable and aligned strategy to work towards a common goal; whereas surprising the opponent through abruptly changing behavior might be beneficial in Competitive but is detrimental for identifiability and stability. Therefore, Cooperative likely leads to more identifiable interactions.

3.2 Study Design

For our user study, we chose two independent variables to explore throughout our hypotheses and conducted a within-subject, repeated-measures experiment (cf. Figure 1(c)). The first independent variable was *NumUsers* with two Levels: Singleuser and Multiuser. Singleuser denotes that only a single person was in the virtual environment at any time, whereas Multiuser denotes that two persons were present in the virtual environment. Next, the second independent variable was *CollaborationMode* with, again, two levels: Competitive and Cooperative. "Competitive" indicates that the immersed VR users take part in a competitive interaction, and "Cooperative" denotes that they act cooperatively in their interaction. Consequently, we cross the levels of the independent variables and obtain $2 \times 2 = 4$ conditions (cf. Figures 1(c) and 2).

As a dependent variable, we chose the per-participant defined recall rate, which is defined as $\frac{TP}{TP+FN}$, where TP corresponds to true positives and FN to false negatives. Therefore, the recall rate denotes the identifiability of each participant in an interval of $[0, 1]$. In addition, we collected participants' subjective feedback through semi-structured interviews and workload ratings to gain insights into the validity of our experimental conditions.

Our experiment followed a two-session design. This means that participants were invited twice to participate in the user study, so that they participated twice in the procedure. This design allowed us to capture their body movement data on two different days. We used the data obtained in the first study session to train our deep learning models and the data from the second session to determine the recall rate, following a hold-out validation, where the second study session corresponds to the test-set [35].

Finally, our participants performed their interactions in a virtual tennis environment. We chose tennis as an activity, as its motion has a low degree of static movement components in combination with a high degree of dynamic movements, and this combination was shown to be highly identifiable [33]. Since we explored multiuser interactions, i.e., two persons playing tennis, we formed groups of participants in advance. Every participant was assigned one specific partner who did not change during the study, so that they performed their interactions always with the same other person during both study sessions. The assignment was performed mostly randomly without any set rules by the experimenter, with the single exception to account for participants' basic logistics, i.e., that both participants were able to find two corresponding timeslots with both being availble for the study (e.g., a very small number of participants were restricted time-wise by travel or other absence periods and we took care to find a matching partner in such cases so that matching timeslots could be found).

3.3 Apparatus

Our apparatus consists of three components: first, the physical space, second the VR application to play tennis, and third, a logger

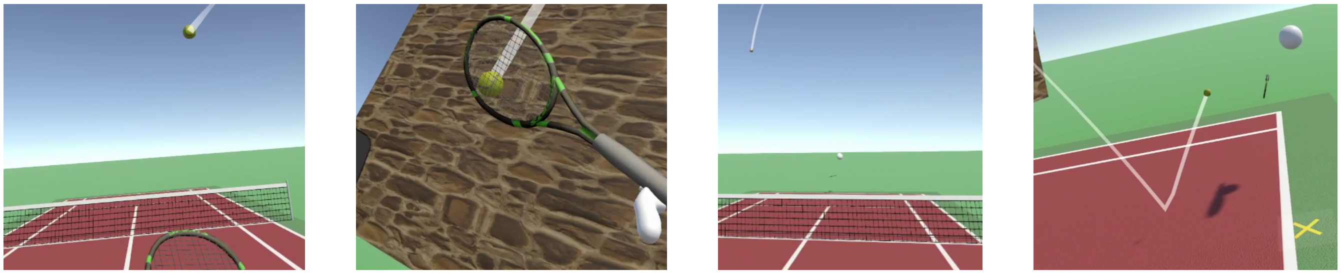
(a) Singleuser-Competitive. (b) Singleuser-Cooperative. (c) Multiuser-Competitive. (d) Multiuser-Cooperative.

Figure 2: Overview of our conditions following our 2×2 design (cf. Figure 1(c)). In (a), a single VR user was immersed in the virtual environment, and they were asked to serve the tennis ball to the opponent's empty side. In (b), the net was replaced with a wall, and a single user was asked to perform a Squash-like interaction, i.e., play the ball repeatedly against the wall and keep it flowing. In (c), two players played tennis against each other in a shared virtual environment. Finally, in (d), two participants cooperated with each other to keep the ball flowing against the wall. The other player in the shared virtual environment was visualized as a sphere that corresponded to their HMD and their hand-held virtual racket.

to capture the player's spatiotemporal body movements (i.e., their Kinetic Signatures). Meta's Quest 3 was used as the standalone HMD together with the handheld controllers. An experimenter was present at all times to observe participants and the apparatus. They were also in charge of running the experiment, as we used Meta's AirLink technology to stream the applications to the respective HMDs. For this, we utilized a WiFi link and two computers with identical specs (Intel i7 9900k, Nvidia Geforce RTX 2080 Super, 32 GB RAM) that were controlled by the experimenter.

3.3.1 Physical Space. We conducted the study under controlled conditions in a secluded seminar room of our institution. We marked a space of 7 m × 3 m (including a safety margin) for each participant on the floor (cf. Figure 1) and the total dimensions of the tennis court were 7 m (width) × 6.50 m (length). The players could move freely within the boundary, as we set the Meta Quest's guardian to the boundaries, in addition to a small safety margin.

3.3.2 Tennis Application in VR. Next, we created a VR application in Unity3D 2022.3.30f1 to play tennis and to implement the experimental conditions of our study design as scenes.

Core Tennis Gameplay Mechanics. The core gameplay mechanics were shared across all study conditions. Participants equipped a virtual racket that was aligned to their controller, with its pivot adjusted to create a natural swing angle. To serve, they could spawn a ball using the X/A button, while the Y/B button triggered the ball machine in the warm-up practice. Players struck balls using their dominant hand, which could be selected at the start of the session. Once hit, a ball followed Unity physics, accelerated at five times the racket velocity but limited to a maximum of 22 m/s. Players were placed in their own instance of the virtual environment in the Singleuser conditions and shared the environment in the Multiuser conditions using Unity's NetCode. Additionally, we implemented a point system. A player scored when the ball was played correctly (e.g., on the empty side of the court or against the wall as required by the condition). If the ball was missed, bounced twice, or not returned properly, the opponent was awarded the point instead. In the Singleuser conditions, the opponent was fictive. In the Multiuser conditions, the opponent was the other participant.

Warm-up Scene. This introductory scene gave participants time to familiarize themselves with the study setup and practice using the controls in a simplified environment. Participants stood on one side of an open tennis court. They could select with which hand they would like to hold the racket and were encouraged to use their dominant hand. On the other side, a red cylinder was placed, which was a simplified representation of the opponent. Participants could either practice their serves by pressing a button on the controller to spawn a ball that they could hit with the racket. Alternatively, using another button on the controllers, they could launch a ball towards them from the other side of the court to perform a counter-shot, similar to reacting to a ball machine. This scene ended when participants felt ready to proceed.

Singleuser-Competitive Scene. In this condition, each participant played alone on a standard virtual tennis court and earned points by serving the ball successfully into the empty opposite field. The virtual environment was not shared; they could not see their human partner, and the ball machine from the warm-up scene was not present. Participants spawned a ball through button presses on the controller and served it into the opposite, empty side of the court. If the ball successfully landed there, they received a point. If they missed, a fictive enemy was awarded the point instead. The points were displayed to participants in all scenes so that they could see their own score and the score of the, in this condition, fictive enemy.

Singleuser-Cooperative Scene. Here, participants again played alone in their own instance of the tennis court, but with a central brick wall replacing the net. They could not see their partners and were alone in the virtual environment. The objective was to serve against the wall and rally with the rebound. Here, instead of the usual net dividing the court, a brick wall was positioned at the center (cf. Figure 2(b)). The setup resembled a wall commonly used in the sport of "Squash" [50]. The objective was to serve the ball and hit it against the wall and each successful hit yielded a point. Players were further encouraged to continue the interaction by attempting to return the ball after it bounced back. If the ball hit the floor twice after contacting the wall, or if a player failed to return it directly to the wall, a fictive enemy was awarded a point.

Multiuser-Competitive Scene. In this Multiuser condition, two participants competed in real time on opposite sides of a shared tennis court. Participants could see each other represented as a simple figure consisting of a sphere for the head and a racket (cf. Figure 2(c)) whose animations were transmitted in real time. In addition, they could see the ball in real time. Then, they played a standard tennis match, aiming to score points by sending the ball onto the opponent's side of the court without it being successfully returned. They took turns serving and returning the ball. Unlike the previous scenes, both participants joined the same virtual environment using Unity3D's Network Manager component. Technically, one participant hosted the game session, while the other joined as a client. The experimenter performed the connection for the participants by having access to the host computers that ran the simulations.

Multiuser-Cooperative Scene. In this cooperative condition, both participants played together on the same side of the tennis court, rallying the ball against a central brick wall while connected over the network. Again, the setup resembled a wall commonly used in the sport of "Squash" [50]. Participants were placed together on the same side of the court and their animations and the visuals of the ball were, again, transmitted in real time so that they could see each other with the same visualization as in the previous scene (cf. Figure 2(d)). They were encouraged to serve the ball in a way that allowed their partner to return it to keep it flowing and were awarded points for doing so. This cooperative setup encouraged collaborative play as participants interacted towards the common goal of keeping the ball in play.

3.3.3 Logger. The final component of our apparatus is the logger. The logger logged the spatiotemporal coordinates of the HMD and both handheld controllers. It operated at 72 Hz and particularly captured the positions ("pos.x", "pos.y", "pos.z") and rotations ("rot.x", "rot.y", "rot.z", and "rot.w") of the HMD, and the left and right controller. In addition, it added metadata such as timestamps and identifiers for the participants, sessions, and conditions. All data was logged to CSV files. Logs were written individually per participant. The logger created the data set that will be used for the biometric identification system after preprocessing. It consists mainly of the spatiotemporal body movements denoted by the positions and rotations.

3.4 Power Analysis

To determine the sample size of our user study, we conducted an a-priori power analysis for a repeated-measures within-factors analysis of variance (RM-ANOVA) using G*Power 3.1.9.7. Given $\alpha = .05$, power = .95, effect size $f = .25$, one group (within-subject) and four measurements, we obtained a total suggested sample size of 36 ($\lambda = 18.00, F = 2.69$).

3.5 Participants

We recruited 40 volunteers (10 female, 30 male, mean age = 25.93, SD = 6.34) for our user study using social media and our institution's mailing lists. Out of the 40, 38 persons were right-handed and two were left-handed. We asked our participants for their agreement with the statement "before participating in this study, I used VR frequently" on a 7-point Likert scale (1 = I strongly disagree, 7 = I strongly agree). We obtained a median of 2 (IQR = 2.5). Furthermore, we presented the statement "I have lots of experience with VR" on the same Likert scale and obtained a median of 2 (IQR = 3.5). Therefore, our participants overall had little experience with VR.

3.6 Procedure

At the beginning of the experiment, we welcomed the participants. They received an introduction to the procedure and the involved devices. Next, we fully answered all their questions and assured that they could cancel their participation in the study at any time without any detriments. Due to the sporty nature of our user study, we informed participants of the possibility to take breaks, drink water, and eat a snack, if they wanted to. Finally, we collected their informed and written consent, demographics, and answers to the aforementioned Likert statements of previous VR experience.

Next, we helped participants put on the HMD and placed them in the warm-up scene. After they felt ready, we tested the four conditions of the user study in a randomized, counter-balanced order using a Latin Square: *Singleuser-Competitive*, *Singleuser-Cooperative*, *Multiuser-Competitive*, and *Multiuser-Cooperative*. Participants engaged in each level for eight minutes. After each condition, we also measured participants' workload by asking them the NASA raw Task Load Index (rTLX) to understand how demanding they found each condition in the study and whether workload would correlate with participants' identifiability or if it would affect the comparability of conditions [66, 67]. In total, each study session took approximately 60 minutes, and we finished the session with a semi-structured interview. After finishing the procedure for the first time, all participants and their partners returned for a second session of the user study on another day. There, they experienced the same four conditions in the same order as in the first session, since we conducted the experimental procedure twice. For the conditions involving two participants (dyadic interactions), they were also assigned the same partner in the second session.

3.7 Ethics

Prior to conducting the study, we obtained approval from our institution's ethics committee. All participants provided informed consent and were explicitly informed that they could withdraw from the study at any time without any disadvantages. To safeguard participants' privacy, we implemented a dual-pseudonymization approach: one randomized identifier was assigned to each participant's data and another to their metadata. After completing the study, we permanently deleted the mapping between these identifiers, rendering any re-identification infeasible. Consequently, the data was fully pseudonymized. We also emphasize that the current work investigates an implicit identification mechanism that potentially allows for implicit user identification throughout the user's body movements. Due to the implicit nature, users might not notice the identification process taking place. Should the findings be employed in practical implementations, it is essential that established ethical standards are upheld. For example, users must be explicitly informed that their data is processed by an implicit identification system, understand the specific purposes of such data processing, and must provide informed consent accordingly.

Table 1: Accuracy ratings in percent of all trained deep learning models that were proposed by Fawaz et al. [18]. The CNN has two variants: with "same" padding, denoted by "(S)", and "valid" padding, which is denoted by "(V)". Mean model accuracy denotes the mean accuracy for the four trained models in this condition. The value marked with an asterisk (*) denotes the highest overall accuracy, which is encountered for the FCN.

Condition	CNN (S)	CNN (V)	Encoder	FCN	Inception	MCDCNN	MCNN	MLP	ResNet	TWIESN	t-leNet
Singleuser-Competitive	24.21	45.21	34.79	47.86	45.86	2.50	2.50	16.14	46.93	27.71	2.50
Singleuser-Cooperative	28.86	67.79	44.64	70.21	67.29	2.50	2.50	09.14	65.79	31.21	2.50
Multiuser-Competitive	50.07	82.35	60.15	83.38	80.96	2.50	2.50	17.50	79.78	44.71	2.50
Multiuser-Cooperative	43.43	78.42	57.21	81.43	77.50	2.50	2.50	17.86	76.43	40.86	2.50
Mean Model Accuracy	36.64	68.44	49.20	70.72*	67.90	2.50	2.50	15.16	67.23	36.12	2.50

4 Analysis

We created a deep-learning-based identification system to collect the recall rate as our dependent variable from the participants' kinetic signatures. We trained the respective deep-learning models with the data that we obtained in the first session of the user study and tested with the data elicited in the second session to obtain the recall rates. Finally, we apply quantitative statistics to explore the effects that the independent variables *NumUsers* and *CollaborationMode* had on the data.

4.1 Data Set

Our data set is created by the logger of the apparatus. It mainly consists of each participant's spatiotemporal coordinates logged from the HMD and controllers. Particularly, we collect all positional ("pos.x", "pos.y", and "pos.z") coordinates, in addition to the rotational coordinates ("rot.x", "rot.y", "rot.z", and "rot.w"). Furthermore, we collect metadata such as timestamps, participant identifiers, and identifiers for the condition in which the behavioral sample was collected. The metadata was only used to partition the data set; it was removed for the model training, hence only participants' positional and rotational coordinates were learned by the models.

A sample related to the point system implemented in the tennis game. Any time a point was scored, a new sample was recorded from the continuous stream of coordinates and the previous sample was marked as finished. This way, we collected 11.348 behavioral samples in the first session (2.51 GB) of the user study in a CSV format and 12.040 behavioral samples in the second session (2.43 GB), which is a total of 23.388 behavioral samples (4.94 GB).

Since every participant played the game at a different pace, we obtained a varying number of samples per participant and per condition. Therefore, to obtain a balanced data set, we calculated the minimum number of samples that every participant contributed per condition, which was 35. Therefore, for further analysis, we incorporate only the first 35 behavioral samples that every participant contributed to the data set and discard all later ones. We seek to avoid an imbalanced data set, as it can have unforeseen consequences on the model training and on important metrics such as accuracy [10].

4.2 Data Split

Since our study followed a two-session design, we use the elicited data and split it by study session to perform a holdout-validation [63]. We prefer a hold-out validation over a cross-validation, since cross-validation bears the risk of overfitting [60]; in addition, we can show that our approach works across days, which is also an important aspect, as it mimics the usage of such behavioral biometric identification systems in the real world. All data that is elicited in the first study session is exclusively used for training, and all data from the second session is used to test the deep learning models and determine the recall rate metric. We never mix data between sessions. Furthermore, we split the data set by condition. That means that we train one model per condition with the data from the first session, and test it exclusively with the second session data elicited in the same condition.

4.3 Preprocessing

We applied the same preprocessing to all behavioral samples, elicited in any condition and any session. First, for each behavioral sample, we removed all metadata so that only positional and rotational coordinates remained. Next, we removed positional biases from the data, especially the "pos.y" coordinate of the HMD, which approximately corresponded directly to the wearer's body height. Furthermore, we removed the second positional bias from the data, which corresponded to the user's location, i.e., the "pos.x" and "pos.z" coordinates in Unity3D's coordinate system. To do so, we calculated local vectors that spanned from the HMD to the left and right controller by subtracting the global coordinates of the controllers from the global coordinates of the HMD, following previous work [28, 56]. Then, we replaced the global coordinates with the newly obtained local coordinates, which were invariant to changes in the global position, so that we did not identify users solely on their global position in the tracking space.

Furthermore, we applied a window slicing algorithm [25]. We selected a window stride of 36 frames, which corresponded to an overlap of 0.5 seconds, as the Meta Quest 3 sampled at 72 Hz. In addition, we chose a window size of 72 frames, which corresponds to exactly one second of sampled data. Through the window slicing, we transformed every participant's biometric samples into a number of new samples. We trained the model on the sliced samples and applied a majority-voting to reconstruct a decision of the original sample from its slices, like previous work [28, 33].

4.4 Deep Learning Models & Training

Finally, we trained eleven different deep learning model architectures on the data from the pre-processing pipeline. We trained one model per condition on the data elicited in the first session of the study and tested it with the data from the second session that was elicited in the same condition. The training results are listed in

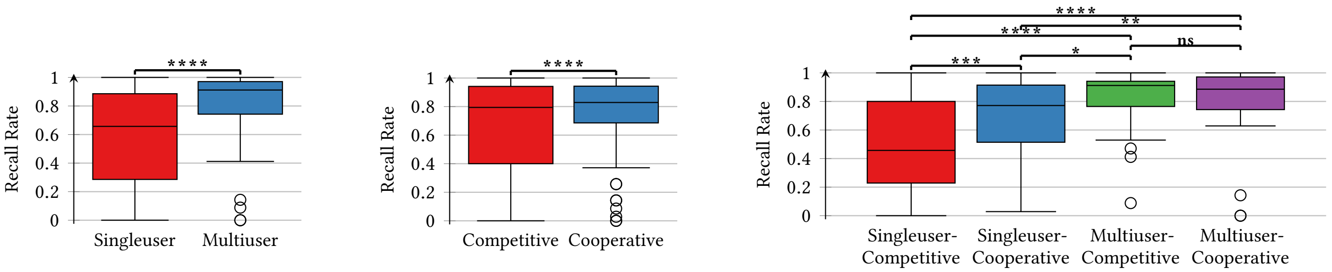

(a) Main Effect 1 for *NumUsers*.

(b) Main Effect 2 for *CollaborationMode*.

(c) Pairwise comparisons of all study conditions.

Figure 3: In (a) and (b), boxplots are provided for the first and second main effect, associated with *NumUsers* and *CollaborationMode*, respectively. In (c), all pairwise comparisons are displayed. Annotated asterisks denote significance levels: **** = p < .0001, *** = p < .001, ** = p < .01, * = p < .05, ns = not significant.

Table 1. We trained all of their proposed models, because we could not foresee which architecture would perform best for our data set. The model architectures are proposed by Fawaz et al. [18] and were previously used for user identification in VR [29, 33, 40]. The models are: a Time Convolutional Neural Network ("Time-CNN") in two variants, using valid ("V") and same ("S") padding [84], a Multi Layer Perceptron ("MLP") [78], a Fully Convolutional Neural Network ("FCN") [78], a Residual Network "ResNet" [78], "Encoder" [71], a Multi-scale Convolutional Neural Network ("MCNN") [11], Time Le-Net "t-LeNet" [25], a Multi Channel Deep Convolutional Neural Network ("MCDCNN") [85], a Time Warping Invariant Echo State Network ("TWIESN") [75], and InceptionTime ("Inception") [19]. Their proposed models are particularly suited for time-series classification [18]; the behavioral biometric samples elicited in our user study are time-series of user behavior. We used the implementation provided by the useckit library and used the default parameters by Fawaz et al. [18, 32].

4.5 Factorial Recall Rate Analysis

Finally, we conducted a factorial analysis of the influences that the independent variables *NumUsers* and *CollaborationMode* imposed on our dependent variable, the recall rate. To do so, we identified the best-performing deep learning model and applied a two-way repeated-measures within-factors analysis of variance (RM-ANOVA) to the recall rates, which we derived from the identity predictions of the best-performing deep learning model.

5 Results

We selected the best-performing model and calculated an RM-ANOVA on the model's predicted recall rates.

5.1 Model Selection

After training all deep learning models, we found that the FCN yielded the highest overall mean accuracy across all four conditions (cf. Table 1). Given its mean accuracy of 70.72 % over all four conditions, it was the best performing one, as it was 2.28 % better on average than the "CNN (V)", which took second place. The FCN also provided the highest overall encountered accuracy of value of 83.38 %. Notably, for our participants' sample ($N = 40$), the base chance for a correct random guess in the model's prediction was $\frac{1}{N} = \frac{1}{40} = 2.5\,\%$. This value was encountered for the MCDCNN, MCNN, and t-LeNet, which did not converge at all.

5.2 Identifiability in Colocated Interactions

For the RM-ANOVA, we first checked whether the obtained recall rates followed a normal distribution. A Shapiro-Wilk test indicated that they are not-normally distributed (W = .8385, p < .0001). Since they were not normally distributed, we applied the aligned-rank transformation [81].

5.2.1 Main Effects. We found two significant main effects, one for *NumUsers* and another one for *CollaborationMode*. First, for *NumUsers* we found that Multiuser (Med. = .91, IQR = .21) yielded higher recall rates compared to Singleuser (Med. = .67, IQR = .60), F(1, 117) = 71.5018, p < .0001, η_p^2 = .3793. This was confirmed by a subsequent post-hoc test (t(117) = 8.4559, p < .0001), which supported **H1**. Figure 3(a) provides a visualization. Second, for *CollaborationMode*, we also found that Cooperative (Med. = .83, IQR = .26) lead to higher recall rates than Competitive (Med. = .79, IQR = .53), F(1, 117) = 20.0197, p < .0001, η_p^2 = .1461. Again, this was confirmed by a post-hoc test (t(117) = -4.4743, p < .0001), which supported **H2**. Figure 3(b) shows the main effect's aggregated recall rates.

5.2.2 Interaction Effect. Furthermore, we found a significant interaction effect (F(1, 117) = 22.5790, p < .0001, η_p^2 = .1618). Figures 3(c) and 4(b) provide an overview. Five out of six possible pair-wise contrast tests showed significant differences following a p-value adjustment using Holm's method. First, the pair-wise comparison between Multiuser-Competitive (Med. = .91, IQR = .16) and Multiuser-Cooperative (Med. = .90, IQR = .23) showed no significant differences (t(117) = -.3919, p = .6958), hence no support for **H2** was given. Second, for the comparison between Multiuser-Competitive and Singleuser-Competitive (Med. = .46, IQR = .55), we found that Multiuser-Competitive yielded significantly higher recall rates (t(117) = 6.8497, p < .0001) that indicated support for **H1**. Third, the comparison between Multiuser-Competitive and Singleuser-Cooperative (Med. = .79, IQR = .41) highlighted that Multiuser-Competitive lead to significantly higher recall rates (t(117) = 2.8369, p = .0107). While this supported **H1**, it contradicted **H2**. Fourth, Multiuser-Cooperative yielded siginficantly higher recall compared

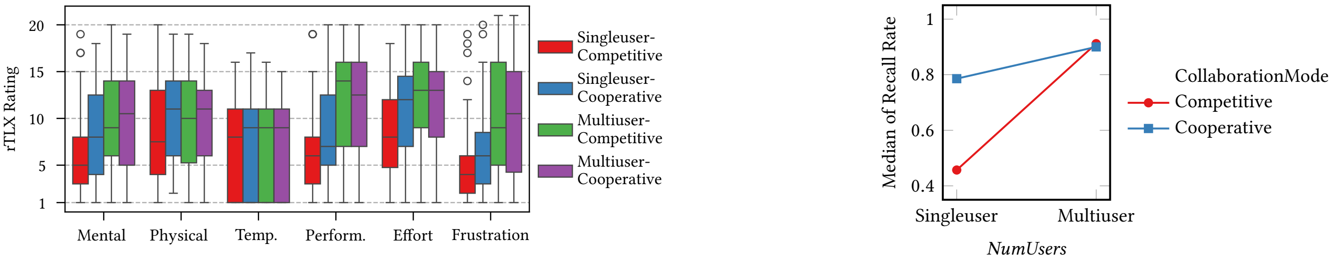

(a) NASA rTLX Scores

(b) Interaction Effect.

Figure 4: In (a), participants' raw NASA TLX scores for each condition and both sessions are shown. The labels are abbreviated and correspond to the six dimensions of the NASA raw TLX [66, 67]. A visualization of the interaction effect is shown in (b).

to Singleuser-Competitive (t(117) = 7.2417, p < .0001), which supported both **H1** and **H2**. Fifth, Multiuser-Cooperative interactions also showed significantly higher recall rates compared to Singleuser-Cooperative interactions (t(117) = 3.2289, p = .0048), which supported **H1**. Finally, we found that Singleuser-Competitive resulted in significantly lower recall rates than Singleuser-Cooperative with t(117) = -4.0128, p = .0004, which supported **H2**.

5.2.3 Summary. To summarize, we found that **H1** which hypothesizes that user behavior in multi-user shared virtual environments yields significantly higher identifiability than behavior in single-user virtual environments was fully supported by its corresponding main effect and all related pair-wise contrast tests. For **H2**, we found a significant main effect and support through two significant comparisons (Multiuser-Cooperative vs. Singleuser-Competitive and Singleuser-Competitive vs. SingleCoop). The comparison between Multiuser-Competitive and Multiuser-Cooperative, however, did not indicate a significant difference, thus providing no support for **H2**. Importantly, the comparison between Multiuser-Competitive and Singleuser-Cooperative contradicted **H2** and supported **H1**. Here, both factors were altered (Multiuser-Competitive vs. Singleuser-Cooperative), whereas the other pairwise comparison that altered both factors (Multiuser-Cooperative vs. Singleuser-Competitive) supported both **H1** and **H2**. This indicates that **H1** outweighs **H2**, which is also reflected in the effect sizes, since the main effect associated with **H1** showed $\eta_p^2 = .3793$, whereas the main effect associated with **H2** showed a lower value of $\eta_p^2 = .1461$.

5.3 NASA rTLX Questionnaire

After every condition, we asked our participants the NASA rTLX questionnaire in both sessions [66, 67]. The results are shown in Figure 4(a). We calculated participants' mean workload score and standard deviation over all TLX dimensions after reversing the performance scale and found for Singleuser-Competitive: M = 8.1035 (SD = 2.4649), Singleuser-Cooperative: M = 9.4910 (SD = 2.5334), Multiuser-Competitive: M = 9.9146 (SD = 2.7868), and Multiuser-Cooperative: M = 9.9903 (SD = 2.8146). The values show that Singleuser-Competitive was associated with the least workload, where all other conditions were in a similar range. Notably, participants' Singleuser-Competitive median rating (cf., Figure 4(a)) for Mental-, Physical, and Temporal Demand, in addition to Effort and Frustation, falls below the median rating of any other condition. In turn, the self-perceived performance falls close to "perfect", which is located at this item's lower end scale. Therefore, Singleuser-Competitive induced the least workload and the highest degree of self-perceived performance. Singleuser-Competitive also lead to kinetic signatures with the lowest degree of identifiability (cf. Figure 3(c)). We observed that the workload associated with Multiuser interactions is mostly equal or higher than the workload associated with Singleuser interactions, with only physical demand and performance being exceptions. However, participants' identifiability for Multiuser interactions outperformed the identifiability of their Singleuser interactions. To understand whether higher workload is associated with higher identifability, we correlated participants' mean workload score with their respective Recall rate using Spearman's correlation coefficient for every condition. We did not find significant results (Singleuser-Competitive: $\rho = -.1112$, $p = .4945$, Singleuser-Cooperative: $\rho = .1518$, $p = .3496$, Multiuser-Competitive: $\rho = .0831$, $p = .6102$, Multiuser-Cooperative: $\rho = .1814$, $p = .2626$). Therefore, we found no evidence that identifiability correlated with perceived workload in any condition.

5.4 Semi-structured Interviews

After participants finished their session, we conducted a brief semi-structured interview. The interview followed three lead questions: i) "what did you think of the study and the tennis interactions in VR", ii) "how did you find the interaction with your partner", and iii) "what did you find easy, and what did you find challenging?". Across both interviews, participants described the VR tennis as "fun" or "enjoyable" 24 times. The controller felt "unnatural" as a racket in eight comments; two participants furthermore reported tracking issues during fast swings. Multiplayer mode was subjectively rated harder than singleplayer in 29 statements, as it was described to be particularly hard to precisely hit the ball. Correspondingly, multiplayer interaction was called "challenging" or "frustrating" 17 times. However, participants also positively described the sensation of impact as in the feeling of striking the virtual ball 7 times, describing it as "satisfying" or "realistic". Overall, singleplayer was labelled "easier" in 24 remarks and thus recommended as an entry point. Cooperative play was preferred to competition in three comments. Finally, five participants stated that seeing their partner in VR increased their motivation. Overall, considering both the quantitative

TLX ratings and the qualitative subjective feedback from the interviews, the study conditions appeared to be perceived as largely comparable, with Multiuser being slightly more demanding.

6 Discussion

In the following, we discuss the ideal identifiability of colocated interactions, overall model accuracies, and limitations related to user-pairing. We conclude with a primer on how the findings of our research can be employed in real VR applications.

6.1 Identifiability of Colocated Interactions

Throughout our exploration, we find strong support for **H1** and can accept that Multiuser interactions yield higher identifiability than Singleuser interactions. This is given by the significant main effect and four out of four significant pair-wise contrast tests, in addition to an effect size of $\eta_p^2 = .3793$. In addition, we also find strong support for **H2** that Cooperative leads to higher identifiability than Competitive, which is given by the significant main effect with an η_p^2 of .1461 with two out of four contrast tests indicating significant differences. However, when the effects of *NumUsers* and *CollaborationMode* interact, the former outweighs the latter. Therefore, the ideal kinetic signature with high identifiability occurs in a Multiuser scenario that is of Cooperative nature.

6.2 Model Accuracy

Overall, our FCN model showed accuracies between 47.86 % in Singleuser-Competitive and up to 83.38 % in Multiuser-Competitive on our balanced data set. This accuracy needs to be interpreted with regard to N, as the base chance for a correct guess is $\frac{1}{N}$, i.e., the identification problem becomes harder the more identities need to be distinguished. In our exploration, we therefore had a base chance of $\frac{1}{40} = 2.50\,\%$ with previous, comparable works reporting accuracies up to 90.91 % [33] (N=24), 69.73 % (N=22) [2], 63.55 % (N=22) [56], or 87.82 % to 98.53 % (N=41). Therefore, our work falls into the approximate range of previous works, but the identification accuracy, of course, primarily depends on the actual behavioral samples and what they were influenced by in the respective user studies, besides the algorithmic components.

6.3 User Pairing Limitations

One central limitation of our work is given by our imposed pairing of participants. Each participant was assigned a fixed partner to interact with in VR that did not change between sessions. Therefore, it is possible that their behavior adapted to each other, i.e., it became specific to the partner player in the virtual environment. We cannot rule out that this impacts generalizability, as it is common in many VR applications that the partner of the interaction changes (e.g., in multiplayer games where the opponents frequently change). A future investigation is needed to understand these effects.

6.4 Utilizing Implicit Identification in VR

While we asked participants to engage in a tennis activity in our user study, playing a VR game is not intended to be a password replacement. To employ our findings, VR designers ideally need to sample users' behavior during a Multiuser scenario of Cooperative nature for implicit and potentially continuous authentication [22, 72, 76]. Also, the user's perspective should be considered; our participants partly perceived the multiuser interactions to be challenging or frustrating and associated them with an overall higher workload. Considering these user-related factors is important for employment in a realistic scenario, and generally speaking, the less the user is bothered, the better. In case that this implicit authentication is not possible, a fallback method for authentication is necessary [3]. Finally, the ethical implications associated with implicit authentication need to be considered (cf. Section 3.7).

6.5 Input Modalities and Future Directions

Our study focused on HMD and controller tracking data as input modalities. While this choice reflects the core tracking capabilities available in most current VR headsets, it also represents a limitation of our work. Recent advances in state-of-the-art biometric sensing and VR hardware have shown that ego-centric cameras, accelerometers, and other multi-modal sensors might provide additional behavioral cues. Exploring whether such modalities alter or enhance identifiability outcomes remains an open question. Similarly, eye-tracking sensors are being increasingly incorporated into HMDs, and they can lead to high identification accuracies [5, 23, 31, 34]. However, for every possible sensor, it is important to consider not only its suitability for implicit authentication as a primary benefit, but also whether it can be meaningfully integrated into the input devices of an HMD in a way that does not interfere with the user experience, to allow for seamless authentication [4].

A related concern in the context of behavioral identification is the potential influence of bias, such as body height. In our preprocessing, we removed the absolute vertical coordinate ("pos.y") to eliminate direct height information by calculating local vectors from the HMD to the controllers. However, we did not normalize participants' body proportions, meaning that distances between headset and controllers could still contain information about overall body height to a very small extent [28]. We note that such structural cues are not available in other sensor modalities, such as accelerometers or egocentric cameras, which do not directly encode anthropometric attributes. Future research may therefore consider explicit normalization strategies to further isolate behavioral patterns from physical traits or apply dedicated normalizations [28].

7 Conclusion

In this work, we explored the identifiability of users in VR through their kinetic signatures and investigated how this is shaped by colocated interactions. We conducted a user study (N=40) in which we varied both the number of immersed VR users and the nature of their interaction through a 2×2 factorial design. Our results show that user identification is feasible in multiuser VR environments, with deep learning models achieving up to 83.38 % accuracy, and that particularly cooperative interactions in shared virtual environments yield highly distinctive body movement patterns. These findings advance our understanding of behavioral biometrics in VR and demonstrate their robustness even in the presence of social interaction, laying a foundation for future biometric authentication systems in VR that are seamless, secure, and user-friendly.

Acknowledgments

We thank the Center for Computational Sciences and Simulation and ZIM of the University of Duisburg-Essen for providing us access to the "amplitUDE" supercomputer.

References

[1] A. Ajit, N. Banerjee, and S. Banerjee. 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)*. IEEE Computer Society, Los Alamitos, CA, USA, 9–97. doi:10.1109/AIVR46125.2019.00012

[2] Yasmeen Abdrabou, Lukas Mecke, Radiah Rivu, Sarah Prange, Quy Dat Nguyen, Vanessa Voigt, Florian Alt, and Ken Pfeuffer. 2023. How Unique do we Move? Understanding the Human Body and Context Factors for User Identification. In *Mensch und Computer 2023*, Markus Stolze, Frieder Loch, Matthias Baldauf, Florian Alt, Christina Schneegass, Thomas Kosch, Teresa Hirzle, Shadan Sadeghian, Fiona Draxler, Kenan Bektas, Katrin Lohan, and Pascal Knierim (Eds.). ACM, New York, NY, USA, 127–137. doi:10.1145/3603555.3603574

[3] Reem AlHusain and Ali Alkhalifah. 2021. Evaluating fallback authentication research: A systematic literature review. *Computers & Security* 111 (2021), 102487. doi:10.1016/j.cose.2021.102487

[4] Florian Alt and Stefan Schneegass. 2022. Beyond Passwords—Challenges and Opportunities of Future Authentication. *IEEE Security & Privacy* 20, 1 (2022), 82–86. doi:10.1109/MSEC.2021.3127459

[5] Sarker Monojit Asish, Arun K. Kulshreshth, and Christoph W. Borst. 2022. User Identification Utilizing Minimal Eye-Gaze Features in Virtual Reality Applications. *Virtual Worlds* 1, 1 (2022), 42–61. doi:10.3390/virtualworlds1010004

[6] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2015. Passwords and the evolution of imperfect authentication. *Communications of the ACM* 58, 7 (2015), 78–87. doi:10.1145/2699390

[7] Bruce Potter. 2005. Are passwords dead? *Network Security* 2005, 9 (2005), 7–8. doi:10.1016/S1353-4858(05)70280-4

[8] Judee K. Burgoon, Lesa A. Stern, and Leesa Dillman. 1995. *Interpersonal Adaptation: Dyadic Interaction Patterns*. Cambridge University Press, Cambridge, United Kingdom.

[9] Ceenu George, Mohamed Khamis, Emanuel von Zezschwitz, Henri Schmidt, Marinus Burger, Florian Alt, and Heinrich Hussmann. 2017. Seamless and Secure VR: Adapting and Evaluating Established Authentication Systems for Virtual Reality. In *Proceedings 2017 Workshop on Usable Security*. Internet Society, San Diego, CA, USA, 12 pages. doi:10.14722/usec.2017.23028

[10] Davide Chicco. 2017. Ten quick tips for machine learning in computational biology. *BioData mining* 10 (2017), 35. doi:10.1186/s13040-017-0155-3

[11] Zhicheng Cui, Wenlin Chen, and Yixin Chen. 2016. Multi-Scale Convolutional Neural Networks for Time Series Classification. arXiv:1603.06995 [cs.CV]

[12] David Ferbrache. 2016. Passwords are broken – the future shape of biometrics. *Biometric Technology Today* 2016, 3 (2016), 5–7. doi:10.1016/S0969-4765(16)30049-2

[13] Penelope Gardner-Chloros. 2009. *Code-switching*. Cambridge university press, Cambridge, United Kingdom.

[14] Alberto Giaretta. 2022. Security and Privacy in Virtual Reality – A Literature Survey. http://arxiv.org/pdf/2205.00208v2

[15] Paul Hagen, Harley Manning, and Randy Souza. 1999. Smart Personalization. *Forrester Research, Cambridge, MA* (1999).

[16] Louisa Hallal, Jason Rhinelander, and Ramesh Venkat. 2024. Recent Trends of Authentication Methods in Extended Reality: A Survey. *Applied System Innovation* 7, 3 (2024), 45. doi:10.3390/asi7030045

[17] International Organization for Standardization (ISO). 2023. ISO/IEC 19092:2023 Financial services – Biometrics – Security framework.

[18] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery* 33, 4 (2019), 917–963. doi:10.1007/s10618-019-00619-1

[19] Hassan Ismail Fawaz, Benjamin Lucas, Germain Forestier, Charlotte Pelletier, Daniel F. Schmidt, Jonathan Weber, Geoffrey I. Webb, Lhassane Idoumghar, Pierre-Alain Muller, and François Petitjean. 2020. InceptionTime: Finding AlexNet for time series classification. *Data Mining and Knowledge Discovery* 34, 6 (01 11 2020), 1936–1962. doi:10.1007/s10618-020-00710-y

[20] J. Bonneau. 2012. The Science of Guessing: Analyzing an Anonymized Corpus of 70 Million Passwords. In *2012 IEEE Symposium on Security and Privacy*. IEEE, San Francisco, CA, USA, 538–552. doi:10.1109/SP.2012.49

[21] Anil K. Jain, Arun A. Ross, Karthik Nandakumar, and Thomas Swearingen. 2025. *Introduction to Biometrics* (2nd ed. 2025 ed.). Springer International Publishing and Imprint Springer, Cham. doi:10.1007/978-3-031-61675-4

[22] Markus Jakobsson, Elaine Shi, Philippe Golle, and Richard Chow. 2009. Implicit Authentication for Mobile Devices. In *Proceedings of the 4th USENIX Conference on Hot Topics in Security (HotSec'09)*. USENIX Association, USA, 9.

[23] Christina Katsini, Yasmeen Abdrabou, George E. Raptis, Mohamed Khamis, and Florian Alt. 2020. The Role of Eye Gaze in Security and Privacy Applications: Survey and Future HCI Research Directions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Regina Bernhaupt, Florian 'Floyd' Mueller, David Verweij, Josh Andres, Joanna McGrenere, Andy Cockburn, Ignacio Avellino, Alix Goguey, Pernille Bjørn, Shengdong Zhao, Briane Paul Samson, and Rafal Kocielnik (Eds.). ACM, New York, NY, USA, 1–21. doi:10.1145/3313831.3376840

[24] Alexander Kupin, Benjamin Moeller, Yijun Jiang, Natasha Kholgade Banerjee, and Sean Banerjee. 2019. Task-Driven Biometric Authentication of Users in Virtual Reality (VR) Environments. In *MultiMedia Modeling*, Ioannis Kompatsiaris, Benoit Huet, Vasileios Mezaris, Cathal Gurrin, Wen-Huang Cheng, and Stefanos Vrochidis (Eds.). Lecture Notes in Computer Science, Vol. 11295. Springer International Publishing, Cham, 55–67. doi:10.1007/978-3-030-05710-7_5

[25] Arthur Le Guennec, Simon Malinowski, and Romain Tavenard. 2016. Data Augmentation for Time Series Classification using Convolutional Neural Networks. In *ECML/PKDD Workshop on Advanced Analytics and Learning on Temporal Data*. Riva Del Garda, Italy, 8 pages. https://halshs.archives-ouvertes.fr/halshs-01357973

[26] Mingjun Li, Natasha Kholgade Banerjee, and Sean Banerjee. 2024. Using Motion Forecasting for Behavior-Based Virtual Reality (VR) Authentication. In *2024 IEEE International Conference on Artificial Intelligence and eXtended and Virtual Reality (AIxVR)*. IEEE, Los Angeles, CA, USA, 31–40. doi:10.1109/AIxVR59861.2024.00012

[27] Mingjun Li, Natasha Kholgade Banerjee, and Sean Banerjee. 2025. Predicting 3D Motion from 2D Video for Behavior-Based VR Biometrics. In *2025 IEEE International Conference on Artificial Intelligence and eXtended and Virtual Reality (AIxVR)*. IEEE, Lisbon, Portugal, 223–230. doi:10.1109/AIxVR63409.2025.00043

[28] Jonathan Liebers, Mark Abdelaziz, Lukas Mecke, Alia Saad, Jonas Auda, Uwe Gruenefeld, Florian Alt, and Stefan Schneegass. 2021. Understanding User Identification in Virtual Reality Through Behavioral Biometrics and the Effect of Body Normalization. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3411764.3445528

[29] Jonathan Liebers, Sascha Brockel, Uwe Gruenefeld, and Stefan Schneegass. 2022. Identifying Users by Their Hand Tracking Data in Augmented and Virtual Reality. *International Journal of Human–Computer Interaction* (2022), 28 pages. doi:10.1080/10447318.2022.2120845

[30] Jonathan Liebers, Christian Burschik, Uwe Gruenefeld, and Stefan Schneegass. 2023. Exploring the Stability of Behavioral Biometrics in Virtual Reality in a Remote Field Study: Towards Implicit and Continuous User Identification through Body Movements. In *29th ACM Symposium on Virtual Reality Software and Technology*. ACM, New York, NY, USA, 1–12. doi:10.1145/3611659.3615696

[31] Jonathan Liebers, Patrick Horn, Christian Burschik, Uwe Gruenefeld, and Stefan Schneegass. 2021. Using Gaze Behavior and Head Orientation for Implicit Identification in Virtual Reality. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology (VRST '21)*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3489849.3489880

[32] Jonathan Liebers, Tristan Kley, Carina Liebers, Uwe Gruenefeld, and Stefan Schneegass. 2024. Useckit: An Open-Source Deep-Learning Toolkit Bundling State-Of-The-Art Algorithms for Evaluating Behavioral Biometrics. In *2024 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, Buffalo, NY, USA, 1–10. doi:10.1109/IJCB62174.2024.10744488

[33] Jonathan Liebers, Patrick Laskowski, Florian Rademaker, Leon Sabel, Jordan Hoppen, Uwe Gruenefeld, and Stefan Schneegass. 2024. Kinetic Signatures: A Systematic Investigation of Movement-Based User Identification in Virtual Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, Florian Floyd Mueller, Penny Kyburz, Julie R. Williamson, Corina Sas, Max L. Wilson, Phoebe Toups Dugas, and Irina Shklovski (Eds.). ACM, New York, NY, USA, 1–19. doi:10.1145/3613904.3642471

[34] Jonathan Liebers and Stefan Schneegass. 2020. Gaze-based Authentication in Virtual Reality. In *Symposium on Eye Tracking Research and Applications*, Andreas Bulling, Anke Huckauf, Eakta Jain, Ralph Radach, and Daniel Weiskopf (Eds.). ACM, New York, NY, USA, 1–2. doi:10.1145/3379157.3391421

[35] Ling Liu and M. Tamer Özsu. 2009. *Encyclopedia of Database Systems: Hold-Out Validation*. Springer, New York. p. 534.

[36] M. Sivasamy, V. N. Sastry, and N. P. Gopalan. 2020. VRCAuth: Continuous Authentication of Users in Virtual Reality Environment Using Head-Movement. In *2020 5th International Conference on Communication and Electronics Systems (ICCES)*. IEEE, Coimbatore, India, 518–523. doi:10.1109/ICCES48766.2020.9137914

[37] Shrirang Mare, Mary Baker, and Jeremy Gummeson. 2016. A Study of Authentication in Daily Life. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 189–206. https://www.usenix.org/conference/soups2016/technical-sessions/presentation/mare

[38] Marian Harbach, Emanuel von Zezschwitz, Andreas Fichtner, Alexander De Luca, and Matthew Smith. 2014. It's a Hard Lock Life: A Field Study of Smartphone (Un)Locking Behavior and Risk Perception. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. USENIX Association, Menlo Park, CA, 213–230. https:

//www.usenix.org/conference/soups2014/proceedings/presentation/harbach

[39] Patrick M. Markey, David C. Funder, and Daniel J. Ozer. 2003. Complementarity of Interpersonal Behaviors in Dyadic Interactions. *Personality and Social Psychology Bulletin* 29, 9 (2003), 1082–1090. doi:10.1177/0146167203253474 arXiv:https://doi.org/10.1177/0146167203253474 PMID: 15189605.

[40] Florian Mathis, John Williamson, Kami Vaniea, and Mohamed Khamis. 2020. RubikAuth: Fast and Secure Authentication in Virtual Reality. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3334480.3382827

[41] Oded Mayo and Ilanit Gordon. 2020. In and out of synchrony—Behavioral and physiological dynamics of dyadic interpersonal coordination. *Psychophysiology* 57, 6 (2020), e13574. doi:10.1111/psyp.13574 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/psyp.13574

[42] Mark Roman Miller, Eugy Han, Cyan DeVeaux, Eliot Jones, Ryan Chen, and Jeremy N. Bailenson. 2023. A Large-Scale Study of Personal Identifiability of Virtual Reality Motion Over Time. doi:10.48550/arXiv.2303.01430

[43] R. Miller, N. K. Banerjee, and S. Banerjee. 2021. Using Siamese Neural Networks to Perform Cross-System Behavioral Authentication in Virtual Reality. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, Lisboa, Portugal, 140–149. doi:10.1109/VR50410.2021.00035

[44] Robert Miller, Natasha Kholgade Banerjee, and Sean Banerjee. 2022. Combining Real-World Constraints on User Behavior with Deep Neural Networks for Virtual Reality (VR) Biometrics. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, Christchurch, New Zealand, 409–418. doi:10.1109/VR51125.2022.00060

[45] R. Miller, N. K. Banerjee, and S. Banerjee. 2022. Temporal Effects in Motion Behavior for Virtual Reality (VR) Biometrics. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, Christchurch, New Zealand, 563–572. doi:10.1109/VR51125.2022.00076

[46] Alec G. Moore, Tiffany D. Do, Nicholas Ruozzi, and Ryan P. McMahan. 2023. Identifying Virtual Reality Users Across Domain-Specific Tasks: A Systematic Investigation of Tracked Features for Assembly. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, Sydney, Australia, 396–404. doi:10.1109/ISMAR59233.2023.00054

[47] A. G. Moore, R. P. McMahan, H. Dong, and N. Ruozzi. 2021. Personal Identifiability and Obfuscation of User Tracking Data From VR Training Sessions. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, Bari, Italy, 221–228. doi:10.1109/ISMAR52148.2021.00037

[48] Alec G. Moore, Ryan P. McMahan, Hailiang Dong, and Nicholas Ruozzi. 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)*. IEEE, Lisbon, Portugal, 556–557. doi:10.1109/VRW52623.2021.00160

[49] Gonzalo Munilla Garrido, Vivek Nair, and Dawn Song. 2024. SoK: Data Privacy in Virtual Reality. *Proceedings on Privacy Enhancing Technologies* 2024, 1 (2024), 21–40. doi:10.56553/popets-2024-0003

[50] Stafford Murray. 2016. *Science of Sport*. Crowood, Marlborough. https://ebookcentral.proquest.com/lib/kxp/detail.action?docID=5286167

[51] Tahrima Mustafa, Richard Matovu, Abdul Serwadda, and Nicholas Muirhead. 2018. Unsure How to Authenticate on Your VR Headset? Come on, Use Your Head!. In *Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics (IWSPA '18)*. Association for Computing Machinery, New York, NY, USA, 23–30. doi:10.1145/3180445.3180450

[52] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2023. Unique Identification of 50,000+ Virtual Reality Users from Head & Hand Motion Data. http://arxiv.org/pdf/2302.08927v1

[53] Vivek Nair, Wenbo Guo, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2023. Berkeley Open Extended Reality Recordings 2023 (BOXRR-23): 4.7 Million Motion Capture Recordings from 105,852 Extended Reality Device Users. doi:10.48550/arXiv.2310.00430

[54] Ilesanmi Olade, Charles Fleming, and Hai-Ning Liang. 2020. BioMove: Biometric User Identification from Human Kinesiological Movements for Virtual Reality Systems. *Sensors* 20, 10 (2020), 2944. doi:10.3390/s20102944

[55] Paul A Grassi, Michael E Garcia, and James L Fenton. 2017. Digital Identity Guidelines: revision 3. doi:10.6028/nist.sp.800-63-3

[56] Ken Pfeuffer, Matthias J. Geiger, Sarah Prange, Lukas Mecke, Daniel Buschek, and Florian Alt. 2019. Behavioural Biometrics in VR: Identifying People from Body Motion and Relations in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300340

[57] R. Miller, N. K. Banerjee, and S. Banerjee. 2020. Within-System and Cross-System Behavior-Based Biometric Authentication in Virtual Reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, Atlanta, Georgia, USA, 311–316. doi:10.1109/VRW50115.2020.00070

[58] Christian Rack, Tamara Fernando, Murat Yalcin, Andreas Hotho, and Marc Erich Latoschik. 2023. Who is Alyx? A new behavioral biometric dataset for user identification in XR. *Frontiers in Virtual Reality* 4 (2023), 11 pages. doi:10.3389/frvir.2023.1272234

[59] Christian Rack, Konstantin Kobs, Tamara Fernando, Andreas Hotho, and Marc Erich Latoschik. 2023. Extensible Motion-based Identification of XR Users using Non-Specific Motion Data. http://arxiv.org/pdf/2302.07517v4

[60] R. Bharat Rao, Glenn Fung, and Romer Rosales. 2008. On the Dangers of Cross-Validation. An Experimental Evaluation. In *Proceedings of the 2008 SIAM International Conference on Data Mining*, Chid Apte, Haesun Park, Ke Wang, and Mohammad J. Zaki (Eds.). Society for Industrial and Applied Mathematics, Philadelphia, PA, 588–596. doi:10.1137/1.9781611972788.54

[61] Julian Rasch, Florian Perzl, Yannick Weiss, and Florian Müller. 2024. Just Undo It: Exploring Undo Mechanics in Multi-User Virtual Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, Florian Floyd Mueller, Penny Kyburz, Julie R. Williamson, Corina Sas, Max L. Wilson, Phoebe Toups Dugas, and Irina Shklovski (Eds.). ACM, New York, NY, USA, 1–14. doi:10.1145/3613904.3642864

[62] Julian Rasch, Vladislav Dmitrievic Rusakov, Martin Schmitz, and Florian Müller. 2023. Going, Going, Gone: Exploring Intention Communication for Multi-User Locomotion in Virtual Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, Albrecht Schmidt, Kaisa Väänänen, Tesh Goyal, Per Ola Kristensson, Anicia Peters, Stefanie Mueller, Julie R. Williamson, and Max L. Wilson (Eds.). ACM, New York, NY, USA, 1–13. doi:10.1145/3544548.3581259

[63] Refaeilzadeh, Payam and Tang, Lei and Liu, Huan. 2009. Cross-Validation. In *Encyclopedia of Database Systems*, LIU, LING and ÖZSU, M. TAMER (Ed.). Springer US, Boston, MA, 532–538. doi:10.1007/978-0-387-39940-9_565

[64] Rick Wash, Emilee Rader, Ruthie Berman, and Zac Wellmer. 2016. Understanding Password Choices: How Frequently Entered Passwords Are Re-used across Websites. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 175–188. https://www.usenix.org/conference/soups2016/technical-sessions/presentation/wash

[65] Alia Saad, Nick Wittig, Uwe Gruenefeld, and Stefan Schneegass. 2022. A Systematic Analysis of External Factors Affecting Gait Identification. In *IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, Abu Dhabi, United Arab Emirates, 9 pages. doi:10.1109/IJCB54206.2022.10007994

[66] Sandra G. Hart. 2006. Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50, 9 (2006), 904–908. doi:10.1177/154193120605000909

[67] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Advances in Psychology*. Elsevier, Amsterdam, The Netherlands, 139–183. doi:10.1016/s0166-4115(08)62386-9

[68] Sarah Theres Völkel, Ramona Schödel, Daniel Buschek, Clemens Stachl, Quay Au, Bernd Bischl, Markus Bühner, and Heinrich Hussmann. 2019. 2. Opportunities and challenges of utilizing personality traits for personalization in HCI. In *Personalized human-computer interaction*, Mirjam Augstein, Eelco Herder, and Wolfgang Wörndl (Eds.). De Gruyter, Berlin and Boston, 31–64. doi:10.1515/9783110552485-002

[69] M. Angela Sasse, Michelle Steves, Kat Krol, and Dana Chisnell. 2014. The Great Authentication Fatigue – And How to Overcome It. In *Cross-cultural design*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Alfred Kobsa, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Demetri Terzopoulos, Doug Tygar, Gerhard Weikum, P. L. Patrick Rau, and P. PatrickL. Rau (Eds.). Lecture notes in computer science Information systems and application, incl. Internet/web and HCI, Vol. 8528. Springer, Cham, 228–239. doi:10.1007/978-3-319-07308-8_23

[70] Albrecht Schmidt. 2000. Implicit human computer interaction through context. *Personal Technologies* 4, 2-3 (2000), 191–199. doi:10.1007/BF01324126

[71] Joan Serrà, Santiago Pascual, and Alexandros Karatzoglou. 2018. Towards a universal neural network encoder for time series. arXiv:1805.03908 [cs.LG] https://arxiv.org/pdf/1805.03908.pdf

[72] Max Smith-Creasey. 2024. *Continuous Biometric Authentication Systems: An Overview* (1st ed. 2024 ed.). Springer International Publishing and Imprint Springer, Cham. doi:10.1007/978-3-031-49071-2

[73] Sophie Stephenson, Bijeeta Pal, Stephen Fan, Earlence Fernandes, Yuhang Zhao, and Rahul Chatterjee. 2022. SoK: Authentication in Augmented and Virtual Reality. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 267–284. doi:10.1109/SP46214.2022.9833742

[74] Viktor Taneski, Marjan Hericko, and Bostjan Brumen. 2014. Password security — No change in 35 years?. In *2014 37th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. IEEE, Opatija, Croatia, 1360–1365. doi:10.1109/MIPRO.2014.6859779

[75] Pattreeya Tanisaro and Gunther Heidemann. 2016. Time Series Classification Using Time Warping Invariant Echo State Networks. In *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, Anaheim, CA, USA, 831–836. doi:10.1109/ICMLA.2016.0149

[76] Issa Traoré and Ahmed Awad E. Ahmed. 2012. Introduction to Continuous Authentication. In *Continuous Authentication Using Biometrics: Data, Models, and Metrics*, Issa Traore and Ahmed Awad E. Ahmed (Eds.). IGI Global, Hershey, PA, USA, 1–22. doi:10.4018/978-1-61350-129-0.ch001

[77] Kim-Phuong L. Vu, Robert W. Proctor, Abhilasha Bhargav-Spantzel, Bik-Lam Tai, Joshua Cook, and E. Eugene Schultz. 2007. Improving password security and memorability to protect personal and organizational information. *International Journal of Human-Computer Studies* 65, 8 (2007), 744–757. doi:10.1016/j.ijhcs.2007.03.007

[78] Zhiguang Wang, Weizhong Yan, and Tim Oates. 2017. Time series classification from scratch with deep neural networks: A strong baseline. In *2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, Anchorage, AK, USA, 1578–1585. doi:10.1109/IJCNN.2017.7966039

[79] Frank Wehrmann and Raphael Zender. 2024. A Systematic Review of State-of-the-Art Practices in Multi-user Virtual Reality Learning. In *Technology Enhanced Learning for Inclusive and Equitable Quality Education*, Rafael Ferreira Mello, Nikol Rummel, Ioana Jivet, Gerti Pishtari, and José A. Ruipérez Valiente (Eds.). Lecture Notes in Computer Science, Vol. 15159. Springer Nature Switzerland, Cham, 483–497. doi:10.1007/978-3-031-72315-5_33

[80] Tim Weissker, Pauline Bimberg, and Bernd Froehlich. 2021. An Overview of Group Navigation in Multi-User Virtual Reality. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, Lisbon, Portugal, 363–369. doi:10.1109/VRW52623.2021.00073

[81] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 143–146. doi:10.1145/1978942.1978963

[82] Jacob O. Wobbrock and Julie A. Kientz. 2016. Research contributions in human-computer interaction. *Interactions* 23, 3 (2016), 38–44. doi:10.1145/2907069

[83] Zhaojun Yang, Angeliki Metallinou, and Shrikanth Narayanan. 2014. Analysis and Predictive Modeling of Body Language Behavior in Dyadic Interactions From Multimodal Interlocutor Cues. *IEEE Transactions on Multimedia* 16, 6 (2014), 1766–1778. doi:10.1109/TMM.2014.2328311

[84] Bendong Zhao, Huanzhang Lu, Shangfeng Chen, Junliang Liu, and Dongya Wu. 2017. Convolutional neural networks for time series classification. *Journal of Systems Engineering and Electronics* 28, 1 (2017), 162–169. doi:10.21629/JSEE.2017.01.18

[85] Yi Zheng, Qi Liu, Enhong Chen, Yong Ge, and J. Leon Zhao. 2014. Time Series Classification Using Multi-Channels Deep Convolutional Neural Networks. In *Web-Age Information Management*, Feifei Li, Guoliang Li, Seung-won Hwang, Bin Yao, and Zhenjie Zhang (Eds.). Springer International Publishing, Cham, 298–310.