POSTER: ReAvatar: Virtual Reality De-anonymization Attack Through Correlating Movement Signatures

Brandon Falk, Yan Meng, Yuxia Zhan, Haojin Zhu
Shanghai Jiao Tong University, Shanghai, China
{bfalk95,yan_meng,dabeidouretriever,zhu-hj}@sjtu.edu.cn

ABSTRACT

Virtual reality (VR) is on the precipice of entering mainstream entertainment with devices equipped with a multitude of sensing, tracking, and internet capabilities that can reshape the current infotainment industry such as online gaming or conferences with novel features. With VR techniques, the online gamer or conference attendances could choose to keep their identity anonymous by easily altering their appearances (*i.e.*, avatars). However, in this study, we present ReAvatar, a novel de-anonymization attack that identifies users by their virtual avatar via a correlation in specific recorded movements. Using 3D pose estimation, we train a sophisticated machine learning model with user movement data recorded while performing a set of movements in real life and then again with their avatars. We then map correlations between these two sets of movement data using a bespoke agglomerative clustering algorithm and establish relationship between the user's virtual and real-life identity. ReAvatar achieves 89.60% accuracy in detecting a unique user among multiple avatars. The security and privacy implications of this paper will be foundational for users and researchers alike that explore the realm of virtual reality.

CCS CONCEPTS

• **Security and privacy** → **Usability in security and privacy**; **Distributed systems security**;

KEYWORDS

VR security; De-annoymization attack; Behavioral biometrics

ACM Reference Format:
Brandon Falk, Yan Meng, Yuxia Zhan, Haojin Zhu. 2021. POSTER: ReAvatar: Virtual Reality De-anonymization Attack Through Correlating Movement Signatures. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21), November 15–19, 2021, Virtual Event, Republic of Korea.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3460120.3485345

1 INTRODUCTION

Virtual reality (VR) immerses users in experiencing simulated environments or virtual worlds providing a significance sense of realism. The beneficial applications of VR have far reaching implications affecting industries globally. In the medical field, doctors and students use VR to practice surgeries and procedures without concern for fatal mistakes. Millions of children and adults also use VR to watch videos and play games. The market size of VR was valued at USD 15.81 billion in 2020 and is expected to grow at a compound annual growth rate (CAGR) of 18.0% from 2021 to 2028 [10].

In VR worlds, the concept of avatar, a digital representation of the user is widely adopted. An avatar can be customized to resemble a humanoid model which has the shape of a human but may have extra characteristics such as wings, horns, and tail. For instance, in VR games such as VRChat[4], users transform into these avatars to communicate and interact with each other anonymously in virtual settings. These interactions provide a sense of security as a trove of information about the user is "inaccessible". In the emerging VR applications, it is expected to provide the privacy enhancement features by leveraging avatar transformation. For example, the users could arbitrarily change their appearance to avoid the tracking in the VR world.

In this study, we present a novel de-anonymization attack ReAvatar, which can compromise the target VR user's privacy by correlating his physical world identity with his VR counterpart. Different from the existing de-anonymization schemes in VR, which either require modifying the head-mounted display (HMD) or stealthily recording the sensor tracking data of HMD [3, 7–9], ReAvatar requires neither modification of VR hardware or injection of malicious code in the software. ReAvatar directly uses the movement information of the targeted anonymous victim as the attack input, which is accessible in a VR scenario without requiring any privilege. Considering that VR is expected to play an increasing role in online game and social networks, protecting the privacy of users is more than critical for the future development of VR while de-anonymization will result in physical harms and violations of immersive experience [1].

ReAvatar is motivated from the following insights: 1) The VR device worn by the victim records their movement data which is confirmed to compromise identity through unique movement patterns. 2) Even while using multiple avatars, the user's inherent movement signature remains stable and unique. 3) There exists a dynamic mapping correlation between the avatar's movement animation of the target in the VR scenario and the real world. ReAvatar can successfully implement a novel de-anonymization attack on anonymous victims in a real VR scenario with the accuracy of 89.60%. In this paper, our **main contributions are summarized as below:**

- We developed ReAvatar, the first de-anonymization attack at the avatars in the VR. ReAvatar significantly highlights the potential privacy threats in modern VR applications.

Haojin Zhu (zhu-hj@sjtu.edu.cn) is the corresponding author.

CCS '21, November 15–19, 2021, Virtual Event, Republic of Korea.

ACM ISBN 978-1-4503-8454-4/21/11.
https://doi.org/10.1145/3460120.3485345

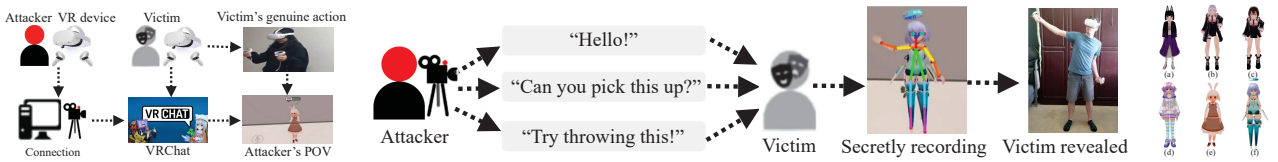

(a) Attack setup. (b) Attack commence. (c) Six avatars.

Figure 1: De-anonymization attack scenario.

Body Part	Left Eye	Right Eye	Left Wrist	Right Wrist	Right Knee	Left Elbow	Right Shoulder
Local X	369.0	359.0	360.0	329.0	340.0	329.0	349.0
Local Y	245.0	245.0	266.0	224.0	349.0	255.0	265.0
Local Z	0.619696	0.632702	0.237653	0.672439	0.343944	0.813862	0.827713

Table 1: A frame capture of the coordinates retrieved

- We proposed a novel computer-vision technique utilizing deep learning to recognize various features of the avatar with a high degree of accuracy and parse the movement data to correlate to their real identity.
- We implement ReAvatar on a popular VR game, VRChat. ReAvatar achieves 89.60% accuracy on correlating VR avatars to real human when employing 6 avatars and 5 users.

2 PROBLEM FORMULATION

The goal of ReAvatar is to reveal the identity of the victim who can hide behind avatars but can't mask their movement signatures.

Attack Assumptions. Figure 1(a) illustrates the beginning of the attack where the adversary uses a VR device to stream and record the content related to the victim to adversary's computer for processing and analysis. Both the adversary and victim will launch the same VR application (*e.g.*, VRChat [4]) and enter the same server. From the adversary's recording point of view (POV) unbeknownst to the victim, all the movement video clips related to the avatar employed by the victim could be collected.

Then, as illustrated in Figure 1(b), the victim will receive a general game command that wouldn't raise any suspicion as they are very common encounters in VR games. The command what is the command usually asks the victim to do some actions (*e.g.*, dancing, picking up things) to enter the next game level. As the victim reciprocates the command innocently, the adversary is capable of extracting movement data.

Input of ReAvatar. In addition to the video movement data from the avatar belonging to the victim, their genuine movement data needs to be collected in order to accomplish correlating the avatar to their identity. In a real situation, parsing through their social media presence and performing social engineering is feasible for obtaining such data. This step is required because without it, only significant clusters will be observed without any association of identities. We also assume the attacker is anonymous to the victim and the victim does not have knowledge of being recorded.

3 METHODOLOGY

ReAvatar consists of three steps: data collection, feature extraction, and identity inference.

Data collection. In this study, we consider the VR device Oculus Quest 2, which supports streaming the VR content directly to a computer. Then, as mentioned in Section 2, when locating and prompting the victim to perform several actions such as waving hello or throwing an object, ReAvatar uses Windows 10 built-in recording software called Xbox Game Bar to record screen with the movements of victim's avatar. The recorded video is processed frame-by-frame for analysis and each video lasts to a maximum of up to 9 seconds to capture the whole progress of an action. This accumulates up to 270 or 540 frames to analyze depending on if the recording is at 30 or 60 frames per second respectfully.

Movement signature extraction. ReAvatar utilizes the concept of part affinity fields (PAF) inspired from OpenPose[2, 6] and both the protofile and the weightfile are sourced from [5]. With this input, ReAvatar is able to detect keywise pairs for poses and extract the local coordinates of (x, y, z) from avatars. To reach this result ReAvatar first needs to map keypoints and assigns pose pairs and for each keypoint find the blob and maxima of each blob. If joint-pair is detected then check every joint within the pair and calculate distance vector. The final step is calculating PAF values to check if interpolated vector is higher than the threshold and add to list.

Table 1 is a snap shot of a singular frame capture for one the avatar 1(b). Only 4 out of 18 parts of interest are tracked in Figure 2(a), considering the realistic and general scenario (*e.g.*, Most VR users have only HMD and hand controllers), only the left wrist, right wrist, and combination of right eye and left eye to be considered by ReAvatar for analysis. Each feature of interest is localized as (X, Y, Z) coordinates. Our avatar is in a 3-dimensional space with the origin point being at the feet. This is why the X and Y values are greater than 200.0 due to being above the above the origin. The Z value reflects the 3-dimensional plane such as the hand being infront of the body. Mapping the coordinates allows for each frame of the video allows us to cluster significance among users.

Identity inference. Figure 2 illustrates output of our enhanced agglomerative k-means algorithm showing a significant difference between the users *i.e.*, user 1 and user 2 from the experiments. In this figure, two users selected 2 of the 6 available avatars as shown in Figure 1(c) and performed the set of actions. It is observed that while the appearance of the users may have been different, the unique movement signatures of the actions remained consistent. This demonstrates users being identified according to their patterns.

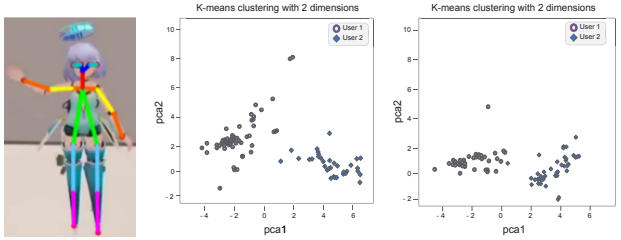

(a) Pose estimation. (b) Avatar A. (c) Avatar B.

Figure 2: Illustration of clusters of significant features when using different avatars.

The patterns are calculated and visualized by our enhanced agglomerative K-means clustering algorithm which utilizes the unsorted (X, Y, Z) data from Table 1 and for each frame analyzes and processes the data. The data is sorted according to their significance allowing for visual observation of each user in a VR environment. The algorithm functions by meshing k-means, principal component analysis (PCA), and agglomerative clustering. ReAvatar will accurately output the identity of the avatar based on the video recordings.

4 EVALUATION

Experimental conditions. We conduct experiments to evaluate the performance of ReAvatar on de-anonymization. We recruit 3 male and 2 female participants between the ages from 20 to 32. 3 participants had prior experience in VR. As shown in Figure 1(a), the participants are instructed to use Oculus Quest 2 device to enter VR environment (*i.e.*, a VR game named VRChat [4] developed for real-time chatting in Unity platform), and we also login in the same environment to record the video of user's avatar at the same time.

In the experiment there is a total of 6 avatars as shown in Figure 1(c). The avatars are humanoid and have been trained upon by the detection algorithm. The participants were then placed in the adversary's POV and asked to perform a set of 3 actions: *wave*, *pick up an object*, *throw an object*. Recording up to 5 sets per participant and change into a new avatar after each set. Each action within a set is recorded for a maximum duration of 9 seconds at 30 frames per second and is regarded as an independent trial totaling for 15 trials per participant. The experiments last 3 days in order to reduce bias in the collected data.

Preliminary results. We used the enhanced agglomerative k-means classifier to conduct the classification. The results are shown in Figure 3. It is observed that the confusion matrix based on the actions performed for accurate de-anonymization of features remains very consistent. ReAvatar categorizes the features of the avatar and the associated (x, y, z) coordinates for each frame yielding and all matrices yield an average of 89.60% accuracy for correctly correlating the feature with associated coordinates. The detection accuracy on "wave", "pick up object" and "throw object" actions are 91.6%, 84.6% and 92.6% respectively. The matrix containing the action to pick up an object is slightly below average compared to its neighbors in terms of accuracy and this is due to the avatar bending down into a new shape that makes it more difficult to detect features of interest.

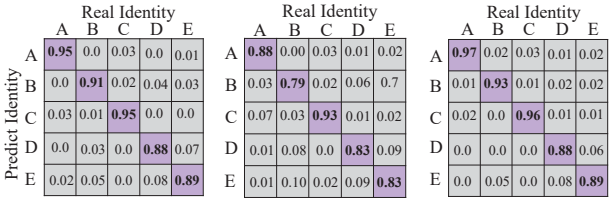

(a) Wave. (b) Pick up object. (c) Throw object.

Figure 3: Confusion matrix among different actions.

5 CONCLUSION & FUTURE WORK

In this poster, we propose an approach to effectively and accurately discern multiple users among multiple virtual avatars and correlate to their real identity. ReAvatar does not require access to either the victim's VR device or inject malicious code in the VR application. We achieve 89.60% accuracy on 5 users and 6 avatars. To further improve the performance of ReAvatar, two directions should be considered. When extracting features from the a video clip, more advanced deep learning implementations other than OpenPose-Caffe model [2] should be explored for future research allowing for improved detection and scalability. Additionally, for the identity inference, developing a more sophisticated clustering algorithm to correlate more complex user data and avatar features is an excellent avenue to pursue.

ACKNOWLEDGEMENT

This study was supported by the National Natural Science Foundation of China under Grant 61972453 and 62132013.

REFERENCES

[1] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics Emerging: the Story of Privacy and Security Perceptions in Virtual Reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 427–442.

[2] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2021. OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 1 (2021), 172–186. https://doi.org/10.1109/TPAMI.2019.2929257

[3] Ceenu George, Mohamed Khamis, Daniel Buschek, and Heinrich Hussmann. 2019. Investigating the Third Dimension for Authentication in Immersive Virtual Reality and in the Real World. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 277–285.

[4] VRChat Inc. 2021. VRChat. https://hello.vrchat.com/.

[5] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*. 675–678.

[6] Sven Kreiss, Lorenzo Bertoni, and Alexandre Alahi. 2019. Pifpaf: Composite fields for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11977–11986.

[7] Mark Roman Miller, Fernanda Herrera, Hanseul Jun, James A Landay, and Jeremy N Bailenson. 2020. Personal identifiability of user tracking data during observation of 360-degree VR video. *Scientific Reports* 10, 1 (2020), 1–10.

[8] 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*. 23–30.

[9] 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*. 1–12.

[10] Grand View Research. 2021. Virtual Reality Market Share & Trends Report, 2021-2028. https://www.grandviewresearch.com/industry-analysis/virtual-reality-vr-market.