

MmWave for Extended Reality: Open User Mobility Dataset, Characterisation, and Impact on Link Quality

Alexander Marinšek , Sam De Kunst , Gilles Callebaut , Lieven De Strycker , Liesbet Van der Perre 

Abstract—User mobility in extended reality (XR) can have a major impact on millimeter-wave (mmWave) links and may require dedicated mitigation strategies to ensure reliable connections and avoid outage. The available prior art has predominantly focused on XR applications with constrained user mobility and limited impact on mmWave channels. We have performed dedicated experiments to extend the characterisation of relevant future XR use cases featuring a high degree of user mobility. To this end, we have carried out a tailor-made measurement campaign and conducted a characterisation of the collected tracking data, including the approximation of the data using statistical distributions. Moreover, we have provided an interpretation of the possible impact of the recorded mobility on mmWave technology. The dataset is made publicly accessible to provide a testing ground for wireless system design and to enable further XR mobility modelling.

Index Terms—Extended reality, head-mounted display, mobility, wireless, millimeter-wave

I. INTRODUCTION

Millimeter-wave (mmWave) technology has been coined the great enabler of wireless extended reality (XR) [1]–[6]. This is primarily due to the multi-Gbps data rates that mmWave technology can provide, as a consequence of the large bandwidth availability in the mmWave spectrum (30–300 GHz). Hence, mmWave links are distinguished by their real-time high-throughput streaming capabilities. Illustrated in Fig. 1, mmWave links can connect a head-mounted display (HMD) to remote processing infrastructure in the edge, fog, or cloud, where high-fidelity XR content rendering takes place [2], [5]. Thus, remote rendering and mmWave technology can together enable high video quality without compromising on freedom of movement, such as in the case of tethered connections. Moreover, remote rendering can reduce HMD hardware complexity and, therefore, aid ergonomics and affordability.

MmWave systems are known for being highly directional by design, while the mmWave links are notorious for their relatively low robustness [7], [8]. This can result in high losses when the radiation patterns of two communicating devices are misaligned or the link between them becomes obstructed [9], [10]. Equipping HMDs with mmWave technology further magnifies the problem, owing to the high degree of movement exhibited by HMD users [11]. *It is the understanding of this user mobility that holds a significant part of the key to deploying and configuring mmWave networks with support for XR applications.*

All authors are with ESAT-WaveCore, KU Leuven, 9000 Ghent, Belgium. Corresponding author: alexander.marinsek@kuleuven.be

© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This is the accepted version of the published paper with DOI: [\[coming soon\]](#)

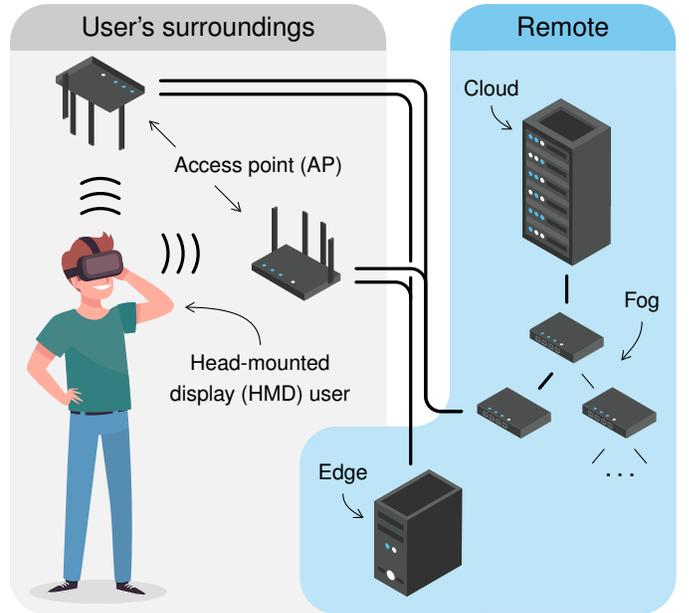


Fig. 1: Wireless remote rendering, which can take place in the network’s edge, fog, or cloud. In a cellular network, the access points (APs) would be substituted by base stations.

User mobility, its impact on mmWave links, and potential mitigation strategies have been the subjects of past research [1], [3], [4], [11]–[13]. However, only a subset of XR use cases, exhibiting limited user mobility, have been considered in the prior art. In this work, we focus on relevant future XR applications featuring a high degree of user mobility that could negatively impact wireless mmWave links. The target applications namely encourage users to explore their environment in all six degrees of freedom (6DoF) and to conduct rapid movements. That is, users both move in a 3-dimensional space and rotate around three rotational axes, as shown in Fig. 2. To facilitate this, we have carried out a tailor-made measurement campaign, allowing us to characterise user movement in high-mobility XR use cases. Furthermore, we elaborate on the characterisation results to assess the possible impact of mobility on mmWave links. *By doing so, we provide the following contributions:*

- A publicly accessible dataset [14], consisting of user tracking data for the head, hands, and body in 6DoF with a total duration of 45 h. This also includes simulator sick-

ness questionnaires (SSQs) and reference measurements for evaluating tracking system accuracy.

- A *characterisation and statistical model-based distribution* approximation of user mobility in a science-fiction role-playing game, an immersive training application, and an action-packed shooter.
- An *evaluation of the possible impact of user mobility on wireless mmWave links*, emphasizing the different adversities that comprise mobility.

The article first describes the measurement campaign and places it alongside prior art in Section II. The collected tracking data are characterised in Section III, while Section IV assesses the possible impact of mobility on mmWave links. Finally, Section V summarises the key findings and outlines future research prospects.

II. MEASUREMENT CAMPAIGN AND MOBILITY DATASET

Several publicly-accessible HMD user mobility datasets and models are readily available at the time of writing [3], [4], [12], [13], [15]–[23], summarised in Table I. These include 360° cinematic experiences [12], [15], [18], [20], virtual museum visits [3], [21], [22], daily activities [16], [17], [19], [23], gaming [13], and pedestrian dynamics [4]. However, the aforementioned feature limited user mobility, which is insufficient for assessing the possible more severe impact of mobility on mmWave links. On the other hand, the prior art that includes relevant high-mobility XR use cases captures exclusively the orientation of the user’s head in 3DoF [13]. The resulting tracking data only permits the analysis of angular head dynamics. To this end, we have conducted a measurement campaign and captured HMD, handheld controller, and body tracker position and orientation in XR applications requiring a high degree of user movement. To the best of our knowledge, this is the first publicly accessible dataset to include 6DoF tracking for multiple body parts in high-mobility scenarios.

A. Measurement setup

Fig. 2 shows the measurement setup and highlights its key components. We collected 6DoF tracking data for the below four body parts:

- The head, using the HTC Vive Pro 2 HMD [25].
- Both hands, with handheld controllers [26].
- The waist, using a Tundra Labs body tracker [27].

The XR devices position themselves using on-board inertial measurement units (IMUs) and external infrared (IR) beacons. The latter sweep the horizontal and vertical angular domain using IR beams with a frequency of 60 Hz to provide precise tracking [28]. We used four IR beacons to ensure a high degree of positioning accuracy and to avoid tracking outage, where the latter could negatively impact tracking accuracy [29]. Tracking data is sampled at a 500 Hz frequency on the HMD and at 250 Hz on the handheld controllers and body tracker. To extract user tracking information, we employed the Brekel OpenVR Recorder [30] and ran it as a background process.

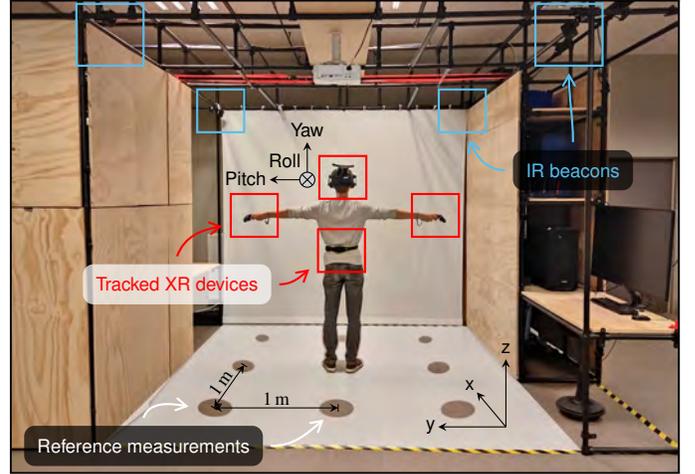


Fig. 2: Measurement environment and a volunteer conducting the T-Pose. The IR beacons are highlighted in blue, while the tracked XR devices are marked in red. The axes left of the HMD show the HMD’s body frame (orientation), with positive rotations determined by the right-hand rule. The axes at the bottom right of the picture show the Cartesian coordinate system (position). The 3×3 m playspace is highlighted white and circles mark the reference measurement locations.

B. Experiment procedure

Fig. 3 shows the procedure applied for each volunteer, with the illustrated steps outlined in the bellow paragraphs.

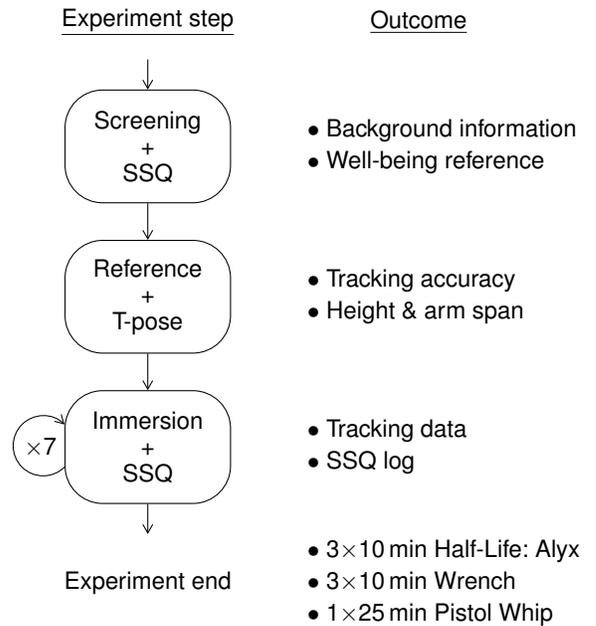


Fig. 3: Experiment procedure per volunteer. The last step (block) is repeated seven times in order to collect the seven tracking data recordings, listed at the bottom right.

1) *Screening and initial SSQ:* A total of 33 volunteers have together contributed to 45 h of tracking data recordings. Volunteers were given a screening questionnaire before the experiment, where they provided informed consent for the further processing of their data, while disclosing their age, gender, height, underlying medical conditions, participation in sports activities, and experience with XR. The 11 female and

TABLE I: An overview of existing mobility datasets and models. The **DoF** column lists the number of degrees of freedom in the collected or generated data, column **Type** indicates whether the tracking data originate from measurements or a mobility model, the **System** column briefly lists the employed measurement hardware or algorithm in the case of mobility models, the **Data** column indicates the provided data types (HMD refers to tracking data for the HMD), the **Use case** column indicates the activity during which data were recorded or for which they were generated, the **Duration** column lists the approximate total duration of the recordings or generated tracking data (model), while the sampling frequency is listed in the last column.

(et al.)	Year	DoF	Type	System	Data	Use case	Duration (h)	Freq (Hz)
Corbillon [12]	2017	3	Data	Razer OSVR	HMD	360° video	6	30
Lo [15]	2017	3	Data	Oculus Rift DK2	HMD, video	360° video	8	30
Faye [16]	2017	3	Data	Nexus 5X phone LG Watch Smartglasses	GPS, steps Heart rate, steps Gaze, Velocity	Daily activities	60	100
Yan [17]	2018	6	Data	Phone IMU	IMU sensors	Phone on body	100	200
Fremery [18]	2018	3	Data	HTC Vive	HMD	360° video	8	200
Chakareski [21]	2020	6	Data	HTC Vive Wireless	HMD	Virtual museum	<1	250
Emery [19]	2021	3/6	Data Model	Custom Oculus Rift HMD + Eye gaze Controllers	HMD (3DoF) Hands (6DoF) Gaze, scenes	Daily activities	4	90
Blandino [3]	2021	3	Model	Extrapolation	HMD	Virtual museum	Variable	250
Struye [22]	2022	3	Model	Deep Learning	HMD	Virtual museum	Variable	17
Dong [13]	2022	3	Data	Oculus Quest 2	HMD, Heart, SSQ	Gaming	13	250
Jin [20]	2022	3	Data	HTC Vive Pro Eye	HMD, Gaze	360° video	54	120
Chukhno [4]	2022	3	Model	Particle statistics	HMD	Pedestrian flow [24]	Variable	Variable
Wei [23]	2023	6	Data	HTC Vive Pro Eye Vive Controllers RGB Camera	HMD, Gaze Hand controls Video	Daily activities	3	50
<i>Ours</i>	2023	6	Data	<i>HTC Vive Pro 2 Vive Controllers Tundra tracker</i>	<i>HMD, SSQ Hand controls Waist tracking</i>	<i>Gaming Training</i>	45	250

22 male volunteers were aged 19–40 years, with the median and mean age, respectively, 25 and 26 years. Following the screening, the volunteers completed an initial SSQ, which serves as a reference of their well-being before the experiment.

2) *Reference measurements and the T-pose:* Reference measurements were made between experiments by placing the four devices on the floor at the positions illustrated in Fig. 2. While the matching against floor markings is not exact due to limited precision when laying the devices on the floor, no position outliers were observed. We calculated that the spread of values around the mean position, the root-mean-square error (RMSE), for each of the 9 reference measurement positions, evaluated 33 times (once per volunteer), is less than 5 mm for all four tracked XR devices in the setup employed in the measurement campaign. This was followed by the *T-pose* recording. Here, the participants stood at the centre of the playspace with their arms extended sideways, as shown in Fig. 2. We then recorded the tracking data for 10 s. This concluded the screening and reference measurement procedures.

3) *Immersion and intermediate SSQs:* Volunteers were immersed in three different applications, spending approximately 20–30 min in each. Breaks were scheduled every 10 min in order for the volunteers to complete an SSQ and to freshen up. The three applications are listed in Table II, with the observed user mobility characteristics listed alongside them. Our aim was to capture mobility profiles that are possibly adverse for wireless mmWave links. This includes *ample movement*, for example, 360° head rotation and movement along the entire playspace. In addition, we were interested in *fast movement*, such as rapid head rotation and swift lateral movement. After reviewing the applications in SteamVR and discussing with

experienced XR users, we foresaw that Wrench would feature ample movement, while Pistol Whip would result in fast movement. Half-Life: Alyx was included since we estimated the in-game puzzles and combat sequences would result in a mixture of both adverse movement types.

C. High-Mobility 6DoF Tracking Dataset

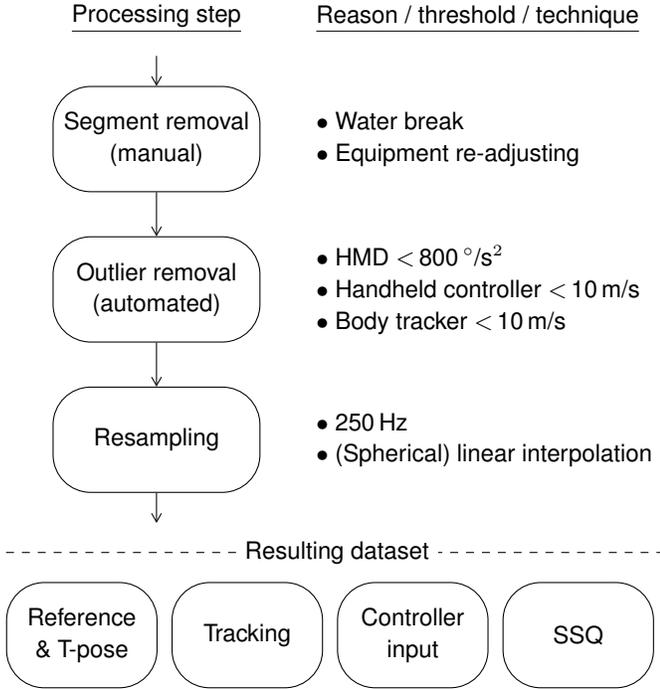
Fig. 4 shows the applied post-processing procedures, outlined in the bellow paragraphs, and lists the four building blocks of the generated dataset, accessible at [14].

1) *Segment and outlier removal:* We have manually removed data segments where the user has stopped to drink water or to re-adjust the XR equipment on their body. Outliers were removed where HMD angular acceleration is greater than 800 rad/s² (peak value during amateur sports [31]), where handheld controller velocity surpasses 10 m/s (martial arts punch velocity [32]), and where the body tracker’s velocity exceeds 10 m/s (athlete sprinting speed [33]). Hence, the removed outliers are physically implausible, while we could not trace their origins to a particular sensor or component.

2) *Resampling:* The published tracking data for the four devices have been synchronised in time and resampled to 250 Hz (controller and body tracker sampling frequency). This consists of subsampling the HMD data from 250 Hz and interpolating the measurement samples for all four devices to bring them onto a common time axis (equal sampling times). Linear interpolation and spherical linear interpolation (SLERP) were used for the position and orientation, respectively. Storing data without further subsampling exceeds the sampling frequency requirements of most user movements. For example, the power-spectral density (PSD) for head orientation during an XR museum visit recorded in [3] showed a drop-

TABLE II: Selected XR applications for the experiment and the corresponding user mobility.

Application	Description	Mobility characteristics
Half-Life: Alyx	A role-playing game set in a post-apocalyptic world. The user traverses three different landscapes while solving puzzles and battling nemeses.	Moderate movement as users explore their surroundings, with fast rotations during occasional high-mobility action sequences.
Wrench	An immersive car mechanic training application. The user is tasked with conducting three different maintenance operations on a car in a mechanic's workshop.	Ample movement along the playspace and ample head rotation when working on a vehicle in a mechanic's garage.
Pistol Whip	An action-packed shooter. The user automatically moves along a linear trajectory in the game while dodging obstacles and shooting or punching villains.	Fast posture changes throughout the heated gunfight, with limited movement due to the automatic in-game forwards trajectory.

**Fig. 4:** Post-processing procedure with the resulting dataset elements at the bottom.

off of more than 20 dB beyond 30 Hz. However, the data sampled at high frequency provide an accurate description of user mobility during fast and accelerated movements, such as running and jumping [31]. Moreover, the high sampling frequency results in high-resolution orientation and position tracking, enabling an accurate evaluation of mmWave link quality. For instance, during a relatively fast rotation at 250 $^\circ/\text{s}$, orientation samples are provided with a 1 $^\circ$ accuracy.

3) *Resulting dataset:* The 6DoF position and orientation data are associated with a 5D grid (axes: user, application, scene, device, and time) and stored in gzip-compressed HDF5 format with a total size of 8 GB. Each sample is represented by seven data variables: three for the position x, y, z (Cartesian notation) and four for the orientation q_x, q_y, q_z, q_w (quaternion notation). We have compiled a separate dataset consisting of controller interaction (button clicks) and published it alongside the tracking dataset. The variables in the controller dataset correspond to individual controller inputs, while it features the same 5D grid as the tracking dataset. Additionally, the

reference and T-pose measurements are accessible in our database at [14], together with the volunteer SSQs. These can, for example, be used to assess the effects of the considered applications on user well-being, which eludes the scope of the work at hand (see [13] for more on the topic). Note that we began the experiment with the HTC Vive wireless adapter (60 GHz mmWave technology [34]) and later switched to a tethered setup due to latency and reliability issues. We noticed high SSQ scores when employing the wireless adapter, with user seven having to stop due to sickness. We changed to a cable and pulley setup (allows a high degree of user movement) starting with user 20. Hence, users 1–19 were using the wireless adapter, while users 20–33 were using the cable instead. This information can be valuable for assessing the effects of current mmWave technology on user sickness and to evaluate the impact of wireless HMDs on user mobility.

As a final remark on the data, we have noticed that there is very little coherence among the available prior art datasets in terms of file format (e.g., CSV and MATLAB files), data types (e.g., quaternions vs. Euler axes), and coordinate conventions (e.g., the directions of the Euler axes). Our approach aims to achieve a self-explanatory dataset structure, to decouple the tracking data from other inputs, and to provide coherency with the coordinate conventions of common Python packages, such as Scipy. While further evaluation in C/C++ and MATLAB would be needed, we note that the current format allows for efficient computation using tools such as Xarray, Dash, and Scipy. Nonetheless, to make the data more approachable, [14] includes tracking data in both the HDF5 format (version 1.1.0) and as individual CSV files (version 0.1.0).

III. MOBILITY CHARACTERISATION

This section characterises user mobility in the three considered applications and describes the results per evaluated mobility variable. With respect to the subsections, these are:

- A: Planar head distance from the centre of the playspace.
- B: Head vertical position, relative to the user's height.
- C: Head pitch angle (looking direction up/down).
- D: Lateral head velocity (combined in 3D Cartesian coord.).
- E: Angular head velocity (combined yaw and pitch rate).
- F: Planar distance between the hands and the head.
- G: Vertical distance between the hands and the head.
- H: Continuous duration that the hands spend above the head.

TABLE III: Mobility characterisation summary, showing the measured data percentiles P_x , where $x \in \{50, 98, 99, 100\}$, and the corresponding fitted distributions. Shades of gray represent the three applications, from darkest to lightest, respectively, Half-Life: Alyx, Wrench, and Pistol Whip. The letters in the ID column show the subsection within Section III where each variable is discussed. The percentiles highlighted with a bold dark blue font bear important information for mmWave link quality, discussed in Section IV. The vertical distance factor shows the complementary percentiles to highlight the lowest relative head position of users. The three distribution parameters, respectively, represent the shape or skewness, location along the abscissa, and the scaling factor. The min/max columns show the boundaries (inclusive, except for 0<), between which the distributions are defined. The rightmost column shows the results of the Kolmogorov-Smirnov goodness of fit evaluation.

	ID	Quantity	Measured percentiles				Fitted distributions						
			P ₅₀	P ₉₈	P ₉₉	P ₁₀₀	Name	Shape	Loc.	Scale	Min	Max	KS
Head	A	Planar distance (m)	0.3	1.1	1.2	2.0	Gamma	2.06	0	0.19	0<	2.0	0.02
			0.4	1.2	1.4	2.3	Gamma	2.54	0	0.19	0<	2.3	0.02
			0.3	0.9	1.0	1.4	Gamma	2.33	0	0.15	0<	1.4	0.02
	B	Vertical distance factor <i>Complementary percentiles</i>	1.0	0.5	0.4	0.3	Skewed Cauchy	-0.54	0.99	0.02	0.3	1.17	0.04
			0.9	0.5	0.3	0.2	Skewed Cauchy	-0.57	0.98	0.04	0.2	1.12	0.04
			1.0	0.5	0.5	0.3	Skewed Cauchy	-0.76	1.00	0.02	0.3	1.18	0.03
	C	Pitch orientation (°)	20	70	75	90	Skewed normal	1.98	2.96	28.23	-80	90	0.03
			32	74	79	90	Skewed normal	-2.96	59.32	41.04	-82	90	0.02
			3	33	44	90	Logistic	/	2.85	5.93	-70	90	0.03
	D	Lateral velocity (m/s)	0.1	0.6	0.8	3.2	Weibull	0.93	0	0.13	0<	3.2	0.04
			0.1	0.7	0.9	2.4	Weibull	0.80	0	0.12	0<	2.4	0.03
			0.2	1.4	1.6	3.2	Weibull	0.85	0	0.29	0<	3.2	0.05
	E	Angular velocity (°/s)	15	180	224	646	Log-normal	1.27	0	15.84	0<	646	0.01
			10	158	197	550	Log-normal	1.38	0	10.20	0<	550	0.01
			21	166	208	549	Log-normal	1.04	0	20.91	0<	549	0.01
Hands	F	Planar distance (m)	0.2	0.6	0.6	1.0	Log-normal	0.32	0	0.33	0<	1	0.04
			0.3	0.5	0.6	1.0	Log-normal	0.22	0	0.48	0<	1	0.01
			0.3	0.6	0.7	1.0	Log-normal	0.29	0	0.44	0<	1	0.03
	G	Vertical distance (m)	-0.4	0.0	0.1	0.6	Weibull	3.91	-1.17	0.82	-1.1	0.6	0.05
			-0.4	0.1	0.2	0.6	Logistic	/	-0.42	0.13	-1.1	0.6	0.01
			-0.4	0.1	0.1	0.6	Weibull	4.45	-1.51	1.20	-1.1	0.6	0.06
	H	Above head duration (s)	0.5	5.4	7.3	22.7	Weibull	0.91	0	0.94	0<	22.7	0.08
			1.3	12.6	16.6	46.7	Weibull	0.82	0	2.14	0<	46.7	0.03
			0.4	2.1	2.6	22.3	Weibull	1.16	0	0.62	0<	22.3	0.04
	I	Lateral velocity (m/s)	0.1	1.3	1.6	9.7	Weibull	0.84	0	0.23	0<	9.7	0.04
			0.1	1.2	1.4	9.1	Weibull	0.80	0	0.20	0<	9.1	0.05
			0.4	3.2	3.9	9.9	Weibull	0.86	0	0.63	0<	9.9	0.02

I: Lateral hand velocity (combined in 3D Cartesian coord.).

Table III summarises the characterisation using the upper percentiles of the measured data and fitted statistical distributions. The fits are derived for multiple distributions by minimizing the negative log-likelihood function. The distribution that best approximates the data is then selected according to the Kolmogorov-Smirnov goodness of fit evaluation.

A. Head planar position

The 3×3 m playspace constrains user movement in the considered XR applications. The rows of Table III corresponding to ID A show that users stay within a single square meter, centred around the middle of the playspace, for more than 50% of the time in all three applications. We notice that users move the furthest in the training XR use case (Wrench), occasionally reaching the edges of the playspace. This is highlighted by the large values of the upper percentiles and the large shape and scale parameters of the Gamma distribution (long tail).

Fig. 5 plots the planar position distributions for the three applications. The three sets of orthogonal lines in the middle represent the principal components, while the probability density functions (PDFs) along the x and y direction are plotted at the top and to the right of the figure, respectively. In addition to the findings in Table III, we notice a covariance in the xy-position in Half-Life: Alyx and Wrench, evident from the angled principal components. The covariance for these originates from car maintenance tasks and puzzles in the two applications. Users appear to have conducted these tasks more often with their right hand, which also resulted in a diagonal head and body movement. Since most volunteers were right-handed, we hypothesise that the biased handedness is the reason for the position covariance. Additionally, we can notice in Fig. 5 that users in Pistol Whip move more along the x axis than they do along the y axis. This is a consequence of the automatic in-game forward movement and of the dodging of obstacles.

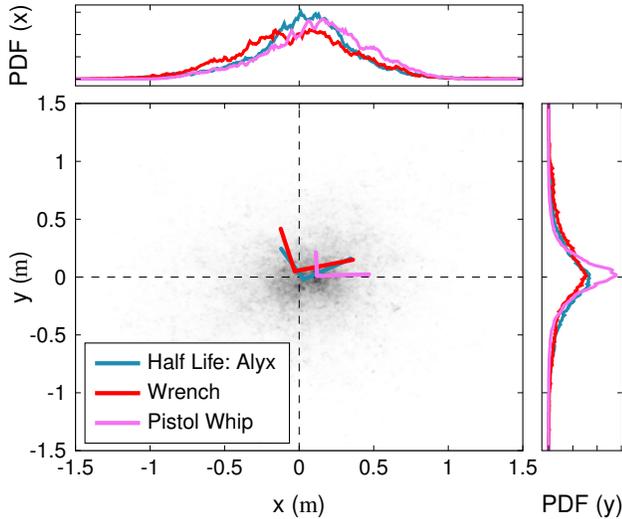


Fig. 5: Planar user position within the playspace. The principal components per application are plotted on top of the position PDF of Half Life: Alyx (a higher probability is darker). The PDF for the x and y axis is shown at the top and right, respectively. Dashed black lines represent zero-values.

B. Head vertical position

Table III, ID B, lists the ratio between the HMD’s vertical position when the user is standing upright, recorded during the preliminary T-pose measurement, and during gameplay. This vertical displacement is presented on a relative scale, since user height can vary considerably, making a comparison in absolute terms less insightful. We notice that users crouch down more in XR training and less in XR games, shown by the 2-times larger scale parameter of the Cauchy distribution for Wrench (medium gray, middle row) in comparison to Half-Life: Alyx and Pistol Whip. The crouching in Wrench appeared mostly when the users were working on a car part close to the floor or even lying underneath the car. On the other hand, Half-Life: Alyx and Pistol Whip require less pronounced crouching.

Fig. 6 plots the PDF of the head’s vertical position to highlight the differences between the applications. Namely, showing the wider distribution of Wrench. Notice that a vertical position factor greater than one was also recorded since users would displace the HMD’s body frame (position) by up to 10 cm when looking up.

C. Head orientation

Fig. 7 shows the two principal components of user head orientation in the yaw-pitch domain. We notice that all applications feature a positive pitch bias (gaze below the horizon). A bias of up to 10° has already been reported in prior art [3], [12], however, the median pitch listed in Table III, ID C, for Half-Life: Alyx and Wrench is considerably larger at 20° and 32° , respectively. Note that the principal components in Fig. 7 start from the mean orientation, which is 23° and 28° , correspondingly. The tendency to look down (positive pitch) in Half-Life: Alyx and Wrench is also visible from the PDFs of the pitch, plotted at the right of Fig. 7, best approximated using skewed normal distributions. Conversely,

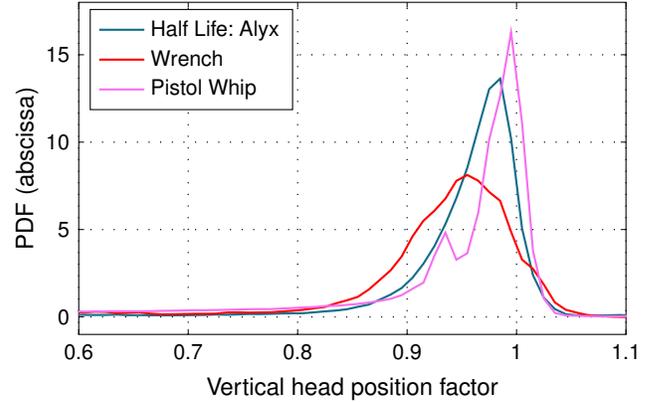


Fig. 6: PDF of the head’s vertical position. Values less than one mean the user’s head was at a lower height than during the T-pose measurement.

the symmetrical shape and heavy tails make Pistol Whip conform with the logistic distribution. Note that the 90° pitch in Table III occurred due to users locating themselves relative to a floor marking at the centre of the playspace.

In terms of yaw, users in Half-Life: Alyx and Wrench explored all orientations, with a slight bias in the forward-facing direction. This is because objects are placed in front of the user when starting the application. For example, the car in Wrench is placed at a yaw of approximately 0° . Moreover, we can observe a negative yaw bias for Wrench, originating from predominant right-hand usage. For example, when reaching for a part under the car’s bonnet, the users looked towards their right hand. Since most volunteers were right-handed, we hypothesise that the biased handedness caused the orientation bias.

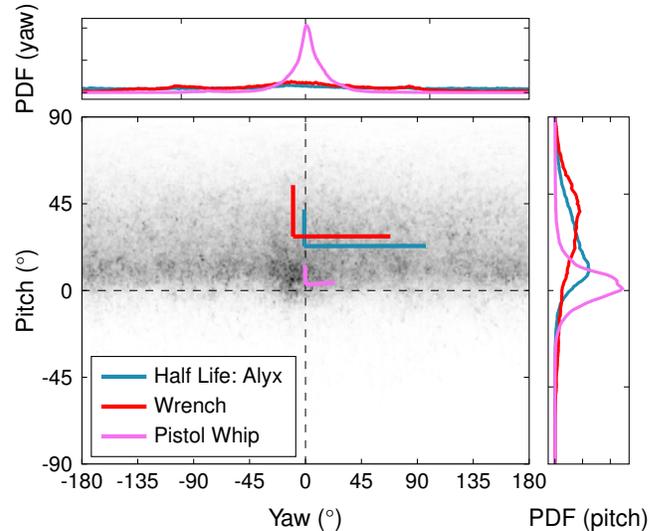


Fig. 7: Head orientation principal components per application on top of the orientation PDF of Half-Life: Alyx (darker corresponds to a higher probability). The PDFs for yaw and pitch are shown at the top and right, respectively. Dashed black lines represent zero-yaw and zero-pitch. Recall that a positive yaw is caused by leftward head rotation, and a positive pitch means the user is looking below the horizon.

D. Head lateral velocity

Table III, ID D, shows the lateral head velocity in each application. That is, the 3-dimensional velocity in a Cartesian coordinate system. We see that users exhibit the fastest lateral movement in Pistol Whip, while the largest value across all applications is 3.2 m/s. Lateral velocity is roughly exponentially distributed, while the Weibull distribution further improves the goodness of fit for the upper percentiles due to its heavier tail. The measured lateral velocity is well-below the velocities observed in other domains, such as vehicular communications or even general sports activities. Instead, XR use cases are characterised by fast rotations.

E. Head angular velocity

Table III, ID E, lists the combined angular velocity along the pitch and yaw axis, which would be observed by a forward-facing antenna. We notice that users can rotate their heads in excess of $360^\circ/\text{s}$, as indicated in prior art [11]. However, for the majority of time, the angular velocity is contained below $100^\circ/\text{s}$ in all three applications. Head angular velocity is best described using the log-normal distribution, since the angular velocity's logarithm is approximately normally distributed. This is also the mobility variable with the best fitting statistical distribution. Fig. 8 confirms that Wrench features the lowest head angular velocity of the three applications. Users exhibit similar angular velocity in Half-Life: Alyx and Pistol Whip, with the former featuring higher peak values. Users in Pistol Whip exhibit higher angular head velocity at lower percentiles, while the highlighted intersection in Fig. 8 indicates the longer tail of Half-Life: Alyx' distribution. In view of the later evaluation of the possible impact on mmWave links, we are interested in understanding whether the high velocity occurs randomly or in bursts.

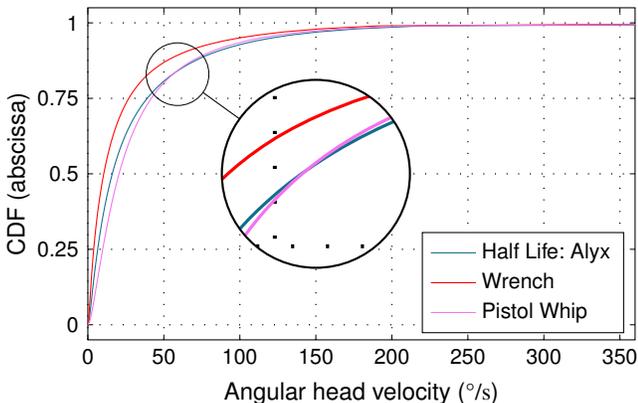


Fig. 8: Cumulative distribution function (CDF) of the head angular velocity.

Fig. 9 shows that the running average of the head's angular velocity during a 0.1 s observation window retains 90% of its short-term value. Thus, multiplying the angular velocity values in Table III, ID E, by the factor 0.90 and further by the window size of 0.1 s allows us to calculate the head rotation during a 0.1 s window. Using Half-Life: Alyx as an example, we can estimate the P_{50} and P_{100} rotation over the course of 0.1 s at approximately 1.5° and 60° , respectively.

The same principle can be applied for other window sizes and applications, where it is evident that the decrease in angular velocity noticeably diverges for longer windows. For instance, the maximal recorded running average of the angular velocity within a window size of 1 s, reduces to about 20% of the maximal short-term velocity (4 ms window size). Conversely, the running average of the median velocity reduces to 60–80% of its short-term value. We notice that the median of the angular velocity's running average reduces the most with in Pistol Whip as window size increases, indicating predominantly shorter-term rotations. A similar trend was observed for other percentiles, which are not plotted in view of clarity. The maximum of the angular velocity's running average rapidly decreases with the window size, reaching less than $360^\circ/\text{s}$ beyond a window size of 0.3 s for all three applications. Note that the 2D angular velocity can be up to a factor of $\sqrt{2}$ larger than the velocity around a single axis of rotation.

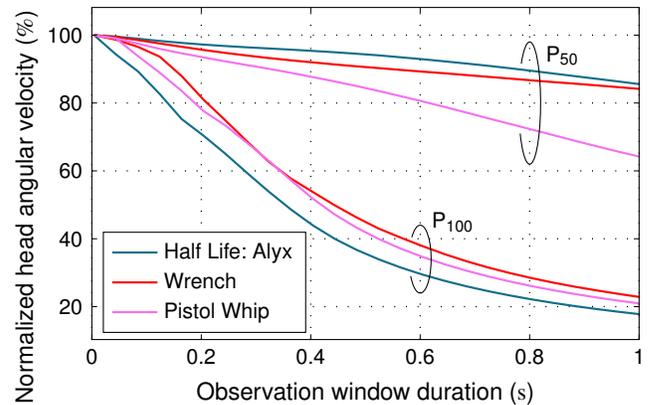


Fig. 9: Running average of the head's angular velocity, dependent on observation window length and normalized relative to the velocity for a 4 ms window. Median and maximal velocity are marked by P_{50} and P_{100} , correspondingly.

F. Planar hand-to-head distance

Table III, ID F, shows that users mostly kept their hands close to their body, indicated by the 0.2–0.3 m median (P_{50}) planar distance (parallel to the floor) between the user's hands and head. The distance's PDF has a sharp cut-in (steep increase at short distances) and, according to P_{98} – P_{100} , a long tail. Given the distance's logarithm is normally distributed and that the planar positions of the hands and HMD rarely coincide, the measurements are best fitted using a log-normal distribution with the shape parameter set to less than one (clear peak). Fig. 10 shows the PDFs of the planar hand distance from the HMD.

G. Vertical hand-to-head distance

Table III, ID G, lists the vertical distance between the users' hands and the HMD. In Wrench, users tend to use their hands in a narrow region with almost equal likelihood of raising or lowering them, represented by the logistic distribution. Half-Life: Alyx and Pistol Whip are better approximated by the longer tails of a Weibull distribution since users often

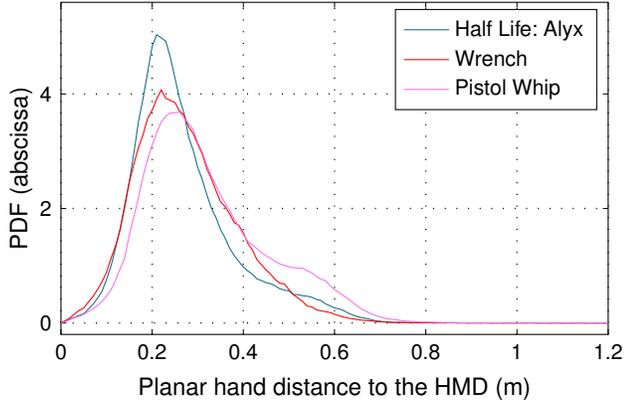


Fig. 10: PDF of the planar distance between the user's hands and the HMD.

lowered their hands when they were not in a shootout. We notice that, in all three applications, the hands spend roughly 3% of time at or above head height. Moreover, considering a 10–15 cm vertical distance between the HMD (eyes) and the user's chin, we see that users keep their hands below the head roughly 90% of the time. Similar to the head's angular velocity, this begs the question: do users place their hands above the head for a long duration, or does the movement resemble a quicker motion, similar to waving?

Fig. 11 shows the vertical position of the hands relative to the HMD. The main difference between the applications in terms of vertical hand distribution is that users tend to use their hands in Pistol Whip at a lower height, compared to Half-Life: Alyx and Wrench. This is because of the game's design, where players have to often lower their hands in order to reload their firearm.

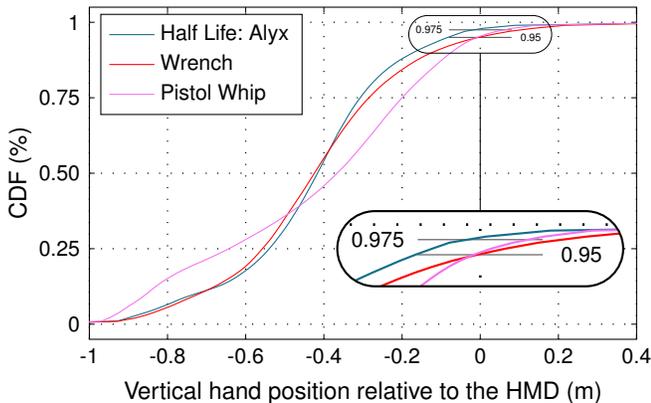


Fig. 11: CDF of the vertical distance between the user's head and hands.

H. Hand persistence above the head

Table III, ID H, shows that users raise their hands above the HMD and lower them back down within the first 0.4–1.3 s on half of the occasions (P_{50}). Furthermore, this duration can grow to more than 20 s in XR games (Half-Life: Alyx and Pistol Whip) and even up to 45 s during the considered XR training (Wrench) – an important finding for mmWave

systems, that has not been reported before. In general, the longest hand persistence above the head was recorded in Wrench, where users also carried out maintenance tasks while the car was lifted above them. Pistol Whip showed the least hand persistence, since users would only quickly aim their firearm and then lower their hand. With a mixture of puzzles and action sequences, Half-Life: Alyx featured moderate hand usage above the head. The slow PDF decay can be well represented using the long tails of a Weibull distribution with the shape parameter set to less than one.

Fig. 12 plots the PDFs of the hand persistence to highlight the differences between the applications. We notice how Wrench features a wider distribution with a larger positive offset, compared to Half-Life: Alyx and Pistol Whip. The latter two have similar PDFs, with Half-Life: Alyx exhibiting somewhat longer hand usage above the head.

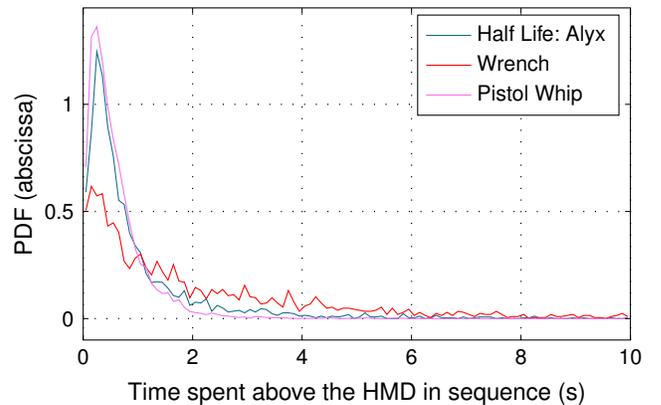


Fig. 12: PDFs showing the likelihood of the user's hands staying above the HMD for a given duration of time.

I. Hand velocity

Table III, ID I, shows that hand velocity is the highest in Pistol Whip, which requires users to punch and aim at moving targets. Users exhibit roughly 65% lower hand velocity in Half-Life: Alyx and Wrench, where there are no tasks which would require fast hand movement. We notice a high maximum hand velocity in all applications, however, since 10 m/s is the punching velocity of martial artists, this could also indicate that the P_{100} values are outliers [32]. For example, due to an underperforming tracking system on the handheld controllers, which are often obstructed by the user and feature fewer IR photodiodes than the HMD. Concluding on the large P_{100} values would require further evaluation of the IR tracking system's accuracy during controller movement ([29] provides a comprehensive assessment for a static HMD). Hand lateral velocity shows a similar distribution to HMD velocity, best approximated using a Weibull distribution.

IV. IMPACT OF MOBILITY ON MMWAVE LINKS

In assessing how HMD mobility negatively impacts mmWave links, two groups of adversities prevail. The first is *shadowing*, where the mmWave link becomes attenuated due to a person or object intercepting it. The second is

TABLE IV: Observations regarding the mobility characterisation and the possible impact on mmWave links.

Mobility	Observation	Possible impact on mmWave links	Adversity
Head planar position	Predominantly static.	Users can obstruct each other's LoS, causing attenuation of up to 15 dB.	External-shadowing
Head vertical position	Can crouch down to 20% of body height.		
Hand planar position	Distance to HMD is about 20–30 cm.	Hands can obstruct a large proportion of the HMD's view, attenuating the LoS by several dB.	Self-shadowing
Hand vertical position	Hands are above the head for 3% of time.		
Hand persistence	Can remain above head for up to 45 s.		
Head yaw orientation	Can rotate for a full 360°.	Losses of several dB or more if insufficient azimuth and elevation coverage are provided.	Antenna misalignment
Head pitch orientation	Bias of up to 32° (gaze below the horizon).		
Head lateral velocity	Slow lateral movement of up to 3 m/s.	Losses of several dB or more, if beam steering at the HMD is not applied frequently enough.	Beam misalignment
Head angular velocity	Fast rotation, possibly exceeding 360 °/s.		

misalignment, where one or more HMD antennas (or arrays) have been moved or rotated by an extent that degrades the mmWave link. The following two sections describe the two adversity groups with regard to general HMD mobility, while Table IV summarizes the observations from Section III and lists the possible impact on the mmWave link.

A. External and self-shadowing

Shadowing occurs when a person or an object obstructs a high-gain propagation path between the HMD and AP. This is typically the line of sight (LoS) component, where a high degree of power is concentrated in mmWave channels [7]. Under the assumption that APs are deployed in a way where there are no objects obstructing the LoS across the entire playspace, there are still two major possible sources of shadowing in XR use cases: other persons and one's own hands [1], [4]. We name the two *external* and *self-shadowing*, respectively.

External shadowing arises when a third person crosses the LoS, usually obstructing it with the torso, neck, or head. Prior art [10] has shown that the resulting mmWave link attenuation can be approximated analytically, for example, using the geometric theory of diffraction (GTD), which depends on the shape of the body part and its distance to the HMD and the AP. According to the GTD and mmWave channel measurements, the obstruction can cause an attenuation of up to 15 dB in practical deployments, where several meters or more separate the HMD, obstructing person, and AP [10]. Similarly, the shadow fading parameter of statistical models for indoor deployments, such as those part of 3GPP 38.901, suggests that attenuation due to shadowing is typically below 10 dB, and that it can potentially reach up to 15 dB. Based on the measurement data, we can assess external obstruction in a potential multi-user scenario. Consider there are several playspaces located next to each other, each with its respective user, and that one or more APs are mounted to the surrounding walls. The highest link obstruction probability is obtained when the centres of two playspaces and an AP align. To determine how often LoS obstruction occurs in such a deployment scenario, we first approximate a person's head and upper body using a vertical cylinder with a radius of 15 cm and 40 cm, respectively. Next, we assess the distance of users from their mean position in the playspace. In particular, we consider Pistol Whip since users exhibit the least amount of lateral movement, resulting in the highest degree of shadowing when the two playspaces and

the AP are aligned. In Pistol Whip, users spend 7% of time within a 7.5 cm distance from the central position and 32% of time within a distance of 20 cm (the corresponding percentiles, P_7 and P_{32} , are not shown in Table III, ID A). Hence, the two numbers respectively represent the worst-case external obstruction probability when link shadowing occurs due to either the head or torso. Determining which of the two applies in a particular scenario further requires information about user body height, user stance, and AP mounting height. The shadowing probability increases if more users are considered.

Self-shadowing happens when the HMD user obstructs the LoS with one or both hands. Although the attenuation caused by a hand or an arm is in general less severe than the attenuation upon obstruction by a person's torso, according to the GTD, it could exceed 5 dB due to the hand's vicinity to the HMD's antennas. According to the measured mobility, self-obstruction, due to raising one's hands above the HMD, has a probability of approximately 3%, making it far less likely than the above-described external shadowing. However, the user's hands can remain above the head for up to 45 s, possibly causing long-term link obstruction. Moreover, the relative proximity of the hands to the HMD, reported in Section III-F, means that a raised hand covers a noteworthy amount of the HMD's angular view. For example, a 7 cm wide palm would obstruct 20° of the HMD's view if placed at the median measured distance from the HMD, i.e., at 20–30 cm.

Providing alternative propagation paths between the HMD and AP can circumvent external and self-shadowing. Therefore, multiple antennas should be placed at different positions on the HMD, while several AP antennas should be distributed in the environment if high link reliability is required. Conversely, consumer deployments might often feature a single AP due to financial reasons, while AP deployment height is usually limited by a low ceiling (e.g., residential buildings). Moreover, mixing users of various body heights in such deployments is likely to lead to some degree of shadowing.

B. Antenna and beam misalignment

MmWave propagation undergoes substantial path loss and is subject to pronounced attenuation during shadowing. To increase gain (mitigate loss), directional antennas with a narrow radiation pattern are employed. Moreover, multiple antennas are assembled into arrays, which yields an even narrower instantaneous radiation pattern, referred to as a beam.

Since the antenna arrays are fixed onto the HMD, which moves relative to a stationary AP, the mobility of the user can lead to so-called *antenna* and *beam misalignment*.

Antenna misalignment occurs when the high-gain regions of the antenna radiation pattern are rotated away from the LoS component. For example, commonly-employed microstrip patch antennas have a directivity gain of several dB, which comes at the price of poor reception from angles exceeding $\pm 90^\circ$ from boresight. Clearly, the HMD requires more than one such antenna to cover the entire horizon (azimuth). However, even an HMD with several antennas, distributed along the azimuth, would observe a varying gain during rotation due to the uneven angular radiation pattern. Furthermore, the large pitch bias observed in this measurement campaign and in prior art suggests that antenna misalignment is likely to also occur in the elevation domain. Certain orientation correction factors, applied during antenna placement on the HMD, could mitigate the impact of this bias on antenna misalignment.

Beam misalignment is the result of antenna array rotation, which is not counteracted using beam steering. MmWave Wi-Fi, that is, the IEEE 802.11ad/ay standard, typically conducts beam steering once per beacon interval, i.e., approximately ten times per second. An HMD employing beams with a 10° half-power beam width and rotating at $100^\circ/\text{s}$ would observe a gain difference of at least 3 dB during a 0.1 s beacon interval, depending on how well the beam was aligned initially. To cope with the fast rotation, the link layer on mmWave HMDs could adapt the beam steering behavior based on user mobility (e.g., using IMU data), which often follows predetermined patterns. For example, the HMD could conduct additional beam steering during periods of elevated mobility or it could proactively steer (or switch) beams based on short-term mobility inference, without occupying airspace with beam probing.

Losses due to antenna and beam misalignment can further reduce the gain of a mmWave link, in addition to the attenuation caused by shadowing. The expected gain reduction of a system with sufficient antennas and proactive beam steering is in the range of several dB, which can increase if insufficient antennas are employed and if inadequate beam steering is provided. Both losses heavily depend on the employed hardware and the mobility characteristics of the XR use case at hand.

V. CONCLUSION AND OUTLOOK

In this paper, we have presented an open 6DoF XR tracking dataset (available at [14]) and the mobility characterisation results, listed in Table III. The tracking data were sampled from two XR games and one XR training application, all of which feature possibly adverse mobility profiles, hence, making the dataset relevant for studies concerning wireless HMDs. The characterisation also proposes statistical distributions that best fit the measured mobility quantities. We have found that user head rotation can surpass $360^\circ/\text{s}$ in all three applications and that users exhibited a considerable head pitch bias, ranging up to 32° (looking down). Moreover, the study revealed that, while the probability of users raising their hands above the head is less than 3%, they can sustain this

posture for durations of up to 45 s. The raised hands severely deteriorate mmWave link quality, although, hands pose a lesser adversity than in mobile telephony. Instead, XR use cases are characterised by the rapid rotation. Our assessment of the possibly adverse effects of the recorded mobility on mmWave links, summarised in Table IV, has shown that mobility poses a substantial challenge in deploying mmWave networks for XR. Namely, due to external shadowing caused by third persons, self-shadowing with one's hands, misalignment of the antenna radiation patterns, and misalignment of array beams.

We began the measurement campaign using a mmWave 60 GHz wireless adaptor on the HMD and switched to a tethered connection after reports of user discomfort during gameplay, indicating that the current mmWave HMDs are not providing the desired user experience. As highlighted, part of the problem may lie in inadequate beam steering to cope with user mobility. The generated database can be used as-is to evaluate novel link management algorithms, such as proactive beam steering (or selection) based on mobility inference. Furthermore, the statistical approximation of the evaluated mobility parameters can be extended to build a mobility model for high-mobility applications, potentially including all 6DoF and user hands in addition to the head. Lastly, the SSQ and tracking measurements could together shed additional light on the connection between user sickness and mobility.

ACKNOWLEDGEMENT

Thanks to our volunteers for making the experiment possible.



This work has received funding from the EU's Horizon 2020 and Horizon Europe programmes under grant agreements No. 861222 (MINTS) and 101096302 (6GTandem).

REFERENCES

- [1] O. Abari, "Enabling High-Quality Untethered Virtual Reality," in *the Proc. of ACM mmNets*, Oct. 2017, pp. 49–49.
- [2] E. Bastug *et al.*, "Toward Interconnected Virtual Reality: Opportunities, Challenges, and Enablers," *IEEE Commun. Mag.*, vol. 55, no. 6, pp. 110–117, 2017.
- [3] S. Blandino *et al.*, "Head Rotation Model for Virtual Reality System Level Simulations," in *the Proc. of IEEE ISM*, Nov. 2021, pp. 43–49.
- [4] O. Chukhno *et al.*, "Interplay of User Behavior, Communication, and Computing in Immersive Reality 6G Applications," *IEEE Communications Magazine*, vol. 60, no. 12, pp. 28–34, Dec. 2022.
- [5] D. G. Morin *et al.*, "Toward the Configuration Implementation of Immersive Augmented Reality Architectures on 5G Networks," *IEEE Commun. Mag.*, vol. 60, no. 2, pp. 46–52, Feb. 2022.
- [6] M. Giordani *et al.*, "Toward 6G Networks: Use Cases and Technologies," *IEEE Commun. Mag.*, vol. 58, no. 3, pp. 55–61, Mar. 2020.
- [7] X. Cai *et al.*, "Dynamic Channel Modeling for Indoor Millimeter-Wave Propagation Channels Based on Measurements," *IEEE Trans. Commun.*, vol. 68, no. 9, pp. 5878–5891, Sep. 2020.
- [8] A. Marinšek *et al.*, "Impact of Array Configuration on Head-Mounted Display Performance at mm Wave Bands," in *the Proc. of IEEE EuCNC/6G Summit*, 2023, pp. 192–197.
- [9] S. K. Saha *et al.*, "Fast and Infuriating: Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware," in *the Proc. of IEEE SECON*, Jun. 2018, pp. 1–9.
- [10] C. Gustafson and F. Tufvesson, "Characterization of 60 GHz shadowing by human bodies and simple phantoms," in *the Proc. of EUCAP (IEEE-supported)*, Mar. 2012, pp. 473–477.
- [11] P. Lincoln, "Low Latency Displays for Augmented Reality," Ph.D. dissertation, Univ. of North Carolina, Chapel Hill, USA, 2017.
- [12] X. Corbillon *et al.*, "360-Degree Video Head Movement Dataset," in *the Proc. of ACM mmSys*, Jun. 2017, pp. 199–204.

- [13] J. Dong *et al.*, “Why VR Games Sickness? An Empirical Study of Capturing and Analyzing VR Games Head Movement Dataset,” *IEEE MultiMedia*, vol. 29, no. 2, pp. 74–82, Apr. 2022.
- [14] S. De Kunst *et al.*, “XR training and gameplay 6DoF mobility dataset”, <https://zenodo.org/records/10836884>, Accessed: 2023-06-15.
- [15] W.-C. Lo *et al.*, “360° Video Viewing Dataset in Head-Mounted Virtual Reality,” in *Proceedings of the 8th ACM on Multimedia Systems Conference*, Jun. 2017, pp. 211–216.
- [16] S. Faye *et al.*, “An Open Dataset for Human Activity Analysis using Smart Devices,” Aug. 2017.
- [17] H. Yan *et al.*, “RIDI: Robust IMU Double Integration,” in *Computer Vision – ECCV 2018*, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds., vol. 11217, 2018, pp. 641–656.
- [18] S. Fremerey *et al.*, “AVtrack360: an open dataset and software recording people’s head rotations watching 360° videos on an HMD,” in *the Proc. of ACM mmSys*, Jun. 2018, pp. 403–408.
- [19] K. J. Emery *et al.*, “OpenNEEDS: A Dataset of Gaze, Head, Hand, and Scene Signals During Exploration in Open-Ended VR Environments,” in *Proc. of ACM ETRA*, ser. ETRA ’21 Short Papers, 2021.
- [20] Y. Jin *et al.*, “Where Are You Looking?: A Large-Scale Dataset of Head and Gaze Behavior for 360-Degree Videos and a Pilot Study,” in *Proceedings of the 30th ACM International Conference on Multimedia*, Oct. 2022, pp. 1025–1034.
- [21] J. Chakareski *et al.*, “6DOF Virtual Reality Dataset and Performance Evaluation of Millimeter Wave vs. Free-Space-Optical Indoor Communications Systems for Lifelike Mobile VR Streaming,” in *Proc. of IEEE Asilomar*, Nov. 2020, pp. 1051–1058.
- [22] J. Struye *et al.*, “Generating Realistic Synthetic Head Rotation Data for Extended Reality using Deep Learning,” in *Proc. of ACM MM, IXR*, Oct. 2022, pp. 19–28.
- [23] Y.-S. Wei *et al.*, “A 6DoF VR Dataset of 3D virtualWorld for Privacy-Preserving Approach and Utility-Privacy Tradeoff,” in *Proc. of ACM MMSys*, Jun. 2023, pp. 444–450.
- [24] F. Farina *et al.*, “Walking Ahead: The Headed Social Force Model,” *PLOS ONE*, vol. 12, no. 1, pp. 1–23, Jan. 2017.
- [25] “HTC Vive Pro 2”, <https://www.vive.com/us/product/vive-pro2/overview/>, Accessed: 2023-11-15.
- [26] “HTC Vive Handheld Controller”, <https://www.vive.com/eu/accessory/controller/>, Accessed: 2023-11-15.
- [27] “Tundra Tracker”, <https://tundra-labs.com/>, Accessed: 2023-11-15.
- [28] “Alan Yates on the Impossible Task of Making Valve’s VR Work”, https://www.youtube.com/watch?v=75ZyteYANTA&t=1121s&ab_channel=HACKADAY, Accessed: 2023-11-15.
- [29] D. C. Niehorster *et al.*, “The Accuracy and Precision of Position and Orientation Tracking in the HTC Vive Virtual Reality System for Scientific Research,” *i-Perception*, vol. 8, no. 3, p. 204166951770820, Jun. 2017.
- [30] “Brekel OpenVR Recorder”, <https://brekel.com/openvr-recorder/>.
- [31] L. E. Miller *et al.*, “An envelope of linear and rotational head motion during everyday activities,” *Biomechanics and Modeling in Mechanobiology*, vol. 19, no. 3, pp. 1003–1014, Jun. 2020.
- [32] K. S. Urbinati *et al.*, “A new virtual instrument for estimating punch velocity in combat sports,” in *the Proc. of IEEE EMBC*, 2013, pp. 571–574.
- [33] T. Haugen *et al.*, “The Training and Development of Elite Sprint Performance: an Integration of Scientific and Best Practice Literature,” *Sports Medicine - Open*, vol. 5, no. 1, p. 44, Nov. 2019.
- [34] “HTC Vive Wireless Adapter”, <https://www.vive.com/eu/accessory/wireless-adapter-full-pack/>, Accessed: 2023-11-15.