

How Accurate is the Positioning in VR? Using Motion Capture and Robotics to Compare Positioning Capabilities of Popular VR Headsets

Adam Banaszczyk*

Poznań University of Technology, Poland

Mikołaj Łysakowski

Poznań University of Technology, Poland

Michał R. Nowicki†

Poznań University of Technology, Poland

Piotr Skrzypczyński

Poznań University of Technology, Poland

Sławomir K. Tadeja

University of Cambridge, United Kingdom

ABSTRACT

In this paper, we introduce a new methodology for assessing the positioning accuracy of virtual reality (VR) headsets, utilizing a cooperative industrial robot to simulate user head trajectories in a reproducible manner. We conduct a comprehensive evaluation of two popular VR headsets, i.e., *Meta Quest 2* and *Meta Quest Pro*. Using head movement trajectories captured from realistic VR game scenarios with motion capture, we compared the performance of these headsets in terms of precision and reliability. Our analysis revealed that both devices exhibit high positioning accuracy, with no significant differences between them. These findings may provide insights for developers and researchers seeking to optimize their VR experiences in particular contexts such as manufacturing.

Index Terms: Virtual reality, Positioning accuracy, Motion capture, Robotic trajectories.

1 INTRODUCTION

Virtual reality (VR) technology has rapidly evolved, becoming crucial in fields like education and training [2], engineering design and production [26], and healthcare [27]. Accurate positioning of head-mounted displays (HMDs) is vital for user immersion and interaction in VR [24], ensuring digital objects remain stable and interact correctly with user movements, especially in precision-demanding applications like engineering.

Significant advancements in VR head tracking have focused on improving accuracy, reducing latency, and enhancing user experience [11, 19]. Developments include low-cost lighthouse-type systems [21], outside-in tracking with RGB-D cameras and facial features [1], very accurate visual-inertial tracking [8], and large-area headset localization using the Global Navigation Satellite System (GNSS) [15].

Commercially available VR headsets, like those from Meta [20], use inside-out vision-based positioning, or ego-motion tracking, requiring environmental mapping [11]. These solutions employ visual odometry and simultaneous localization and mapping (SLAM) [28], demanding expensive computation and struggling in environments with repetitive textures, rapid motion, or significant occlusions [9].

Our work presents a new method for evaluating the accuracy of inside-out tracking in novel HMDs, considering possible failure modes. We aim to understand the limits of positioning accuracy in recent VR headsets to guide researchers and developers. Previous research emphasized the importance of precise tracking systems [16, 22], as the perception of body position can influence the feel of presence experienced by VR users [17]. However, compre-



Figure 1: Real-world head movement data acquisition with motion capture setup (OptiTrack) during VR gameplay.

hensive evaluations of the latest VR headsets using standardized testing methods for inside-out visual tracking are lacking.

We initiated a project to evaluate the positioning accuracy of VR headsets using real user movement trajectories acquired during VR gameplay. This approach involves various movement sequences, resulting in diverse headset movement trajectories.

For rigorous evaluation, we used an industrial cooperative robot to precisely repeat these head movement trajectories for different headsets. The trajectories, captured using a state-of-the-art motion capture system, serve as the objective ground truth for our measurements (Fig. 1). This ensures consistency and repeatability, allowing accurate headset performance comparison.

We tested two recent mass-market devices: the Meta Quest 2, aimed at the entertainment market, and the Meta Quest Pro, for professional applications. Meta has been the world’s largest VR equipment vendor in recent years [5].

Our proposed methodology, software, and data, available at Github¹, will benefit the VR development community. Our contributions include: (i) a detailed comparison of the positioning accuracy of Quest Pro and Quest 2 headsets; (ii) an analysis of VR headsets’ tracking precision strengths and weaknesses; and (iii) a reproducible methodology and step-by-step procedure for experiments with various VR headsets using headset-agnostic trajectories, ROS (Robot Operating System) software, and commonly available equipment like a portable motion capture system (OptiTrack) and a cooperative manipulator arm.

*e-mail: adam.a.banaszczyk@gmail.com

†e-mail: michal.nowicki@put.poznan.pl

¹https://github.com/AdamBanaszczyk/head_trajectory

2 STRUCTURE OF THE POSITIONING ACCURACY ASSESSMENT SYSTEM

The proposed system for assessing VR headset positioning accuracy involves several key components working together for precise evaluations (Fig. 1). The system is structured as follows:

- **Trajectory data collection:** this module uses motion capture to collect trajectory data from VR headsets during selected VR applications. It involves calibrating the headset’s coordinate system (used by Unity) with the OptiTrack system and synchronizing timestamps between Unity and OptiTrack. The result is real-world trajectories that can be “played” with different VR goggles mounted on a cooperative robot.
- **Trajectory reproduction:** the collected trajectories are reproduced by a collaborative robot UR5e with a mock-up head as its end-effector, on which the tested VR headset is mounted. Trajectories are logged simultaneously by the ROS controller and the Unity application. This module requires calibrating the transformation between the robot’s end-effector and the Unity coordinate system. Additionally, trajectories are pre-processed and tested in a Gazebo-based simulation to ensure safe playback by the robot.
- **Assessment of positioning accuracy:** this is done using the Absolute Pose Error (APE) from the evo package [12]. The APE in evo with alignment and translational pose relation is equivalent to the Absolute Trajectory Error (ATE) metric defined in [7], and commonly used in robot localization research. The APE aligns both trajectories and measures the distance between the estimated and ground-truth poses, specifically the headset’s coordinate system origin as read from the robot and the same point estimated by the VR headset. This requires time synchronization of headset poses from both sources.




Figure 2: Schematic view of the structure of the proposed positioning accuracy assessment system.

3 ACQUIRING HEADSET TRAJECTORIES IN REAL-WORLD SCENARIOS

3.1 OptiTrack Motion Capture Setup

To accurately track the user’s head motion during HMD usage, we utilized the OptiTrack motion capture system. We started the data capture process by defining a natural playing area for a seated user and surrounding it with six OptiTrack cameras, ensuring that every movement could be accurately recorded (see Fig. 1). Each tested headset required a set of OptiTrack markers attached, ensuring no sensor was obstructed (see Fig 3).

A defined set of markers is later on represented by OptiTrack as a single rigid body corresponding to HMD, with its position defined as the center of mass of used markers. After the calibration procedure, the OptiTrack system tracks the poses of the markers with an expected error smaller than $0.2[mm]$ at $120[Hz]$ [10].

With the experimental setup complete, the operator could enter the area and proceed to use HMD as usual, allowing for simultaneous trajectory data acquisition. Such a non-invasive setup can be easily replicated with any available VR headset without altering it.

3.2 Data Acquisition During VR Gameplay

During the data acquisition, we tested a pool of VR games and experiences based on their compatibility with a defined OptiTrack’s playing area, accessibility, popularity, and user ratings. Each tested app had to be free or available in a playable demo version, suitable for seated gameplay, rated more than 1,000 times, and have a user rating of at least 80% positive. We tested, among others, *Space Pirate Trainer DX*, *Oculus First Contact*, *Bait!*, *Mission: ISS: Quest*, and *Gorilla Tag*.

Game	Score	Ratings
<i>Space Pirate Trainer DX</i>	92%	3498
<i>Oculus First Contact</i>	92%	1451
<i>Bait!</i>	81%	2987
<i>Mission: ISS: Quest</i>	80%	1619
<i>Gorilla Tag</i>	93%	120k+

Table 1: Statistics of tested applications as obtained from <https://vrdb.app/quest/index.us.html>, where Score indicates the average user rating for the application, while Ratings represent the total number of votes cast by users who have purchased or downloaded the app.

While using a headset, we collected 6 degrees of freedom (DoF) positional ground-truth data in the form of Robot Operating System (ROS) messages with the help of OptiTrack. We chose ROS format for its standardized structure and Unix timestamps, guaranteeing uniformity in data comparison across all experiments. However, the OptiTrack software (Motive) does not directly support the ROS format and only measures time increments from the initial entry, which was insufficient for our application. Hence, we created a custom package based on the NatNet client/server networking software development kit (SDK), which allowed us to forward Motive messages to another computer publishing ROS messages. To eliminate the transfer lag that potentially could have disturbed the direct comparison of OptiTrack and HMD data, we connected the computers directly via Ethernet.



Figure 3: Meta Quest Pro with OptiTrack markers attached.

3.3 Setup Calibration: Unity-to-OptiTrack

Having precisely recorded real-world head trajectories in the OptiTrack coordinate system, we needed to determine its exact transformation to the respective trajectory recorded by the HMD in its coordinate system. To obtain such data, we first created a Unity application, which at the frequency of $30[Hz]$ logged the device's timestamps together with 6 DoF quaternion poses of the Center Eye Anchor, provided by the Oculus Software Development Kit (SDK) and the Oculus Integration Package, which manages the VR camera setup and tracks its alignment with the physical movements. Center Eye Anchor references a single point in the VR environment representing the position and orientation of the user's eyes or head. Hence, it was well-suited for our needs.

In the course of all the described experiments with the Meta Quest Pro and Meta Quest 2 headsets, the software version was Meta Quest build 60.0.0, and the OS version was SQ3A.220605.009.A1.

The calibration procedure consisted of a set of specifically prepared relatively slow HMD movements, which aimed at proper saturation of poses, enabling accurate calculation of transformation between two systems. The procedure was performed in the arranged playing area, simultaneously recording data from both introduced systems (OptiTrack and Unity). It was crucial to maintain the identical configuration of OptiTrack markers as utilized during the actual real-world head trajectory recordings since any discrepancies would pollute the transformation.

Determining the precise transformation between Unity's Center Eye Anchor and OptiTrack's rigid body required solving a hand-eye calibration problem [14]. Both systems were rigidly attached to the same actuator, which in this case was the user's head. Therefore, there existed a homogenous transformation X that we aimed to determine that satisfies the equation $AX = XB$, where A denotes the transformation from the Unity base to the Center Eye Anchor, and B represents the transformation between the previous and current poses of the OptiTrack rigid body whenever the system moves. To resolve this problem, we developed a C++ program utilizing the OpenCV [3] and Eigen [13] libraries, which sequentially loaded both trajectories, applied linear interpolation for translation and spherical linear interpolation (slerp) [25] for quaternion rotation, matched their entries by timestamps, and transformed Unity's trajectory from a left-handed system (LHS) into the right-handed system (RHS). The latter was chosen as it matches OptiTrack, ROS, and the majority of robotics applications. Each sequence had its data represented as a transformation from the starting position. Moreover, since OptiTrack's frequency was 4 times greater than Unity's, we diluted its data by that factor to significantly increase the speed of code execution. We thoroughly evaluated the available calibration methods in the OpenCV library, and we have found that the Tsai [29] method sometimes failed to provide accurate solutions, so we selected the Daniilidis method [6] due to its consistent performance.

During hand-eye calibration we discovered a slight time offset between OptiTrack's and Unity's trajectories, which varied for different recordings but was within $0 - 200[ms]$. We tried to eliminate that timestamp mismatch by ensuring that HMD is fully charged and OptiTrack's computer and HMD are connected to the Internet to synchronize internal clocks. However, a slight timestamp offset was still present in some cases. Therefore, we needed to introduce an additional function that could estimate the value of that delay with an accuracy of a few milliseconds. This new program detected axial position peaks for both trajectories and compared the respective timestamps throughout entire recordings. Differences in timestamps at those particular positions were then averaged to calculate the final time offset between Unity and OptiTrack, which was addressed during the hand-eye calibration and any other experiment, allowing for precise time alignment of studied trajectories.

4 MOUNTING HEADSET ON THE ROBOTIC ARM

4.1 Custom Robot Mount



Figure 4: Artificial head with the Meta Quest Pro attached, mounted on the Universal Robots UR5e manipulator tip (up), and the real-world test environment as seen from the perspective of the mounted headset (down).

Accurate replication of human head movements on a robot requires a machine capable of performing demanding movements with high accuracy and precision of repetition, while having a relatively large range and freedom of motion to guarantee the ability to reproduce all sorts of trajectories. Consequently, we decided to use Universal Robots UR5e robotic arm with a maximum reach of $850[mm]$, payload of up to $5[kg]$, six rotating joints, and pose repeatability of $\pm 0.03[mm]$ [30]. Accuracy is understood to be the ability to attain the required point in the workspace. While this parameter is not directly specified by the robot manufacturer (as repeatability is mainly important in industrial tasks), results reported in the literature [23] show that the unidirectional pose accuracy in UR5 is better than $0.1[mm]$, under a combination of two load values and three motion speeds (up to 100% speed). This suggests that the UR5 robot is suitable to replicate the headset's motion in our experiments.

The robot itself was mounted on a dedicated table, resulting in the average height of the Tool Center Point (TCP) within the height of a human head when playing in VR. Although the UR5e robotic arm is very versatile, we needed to propose a customized, secure and robust HMD mounting. Therefore, as all VR headsets are designed to ergonomically rest on the user's head, which is a uniquely shaped body part itself, we decided that the optimal approach was to place the headset on a plastic mannequin bust (see Fig. 4). Furthermore, we designed a custom mounting system to connect such a head to the UR5e manipulator tip, which features threaded holes into which the mounting system was securely fastened. The mounting system consisted of 3D-printed components responsible for precise alignment and connection of all elements and sturdy steel plates providing stability to the base of the mannequin bust (see Fig. 4).

4.2 Setup Calibration: TCP-to-Unity

With HMD properly mounted on the robot, we were able to determine another crucial transformation between the robot’s TCP expressed in the robot’s global frame, with the coordinate system origin at its base and Unity’s Center Eye Anchor recorded with our application running on the headset. Designating accurate TCP poses in the established earlier format, i.e., 6DoF positions with quaternion rotations and timestamps, required setting up communication with the robot’s controller via Real-Time Data Exchange (RTDE) client library, which implements API for UR interface. We eliminated potential transfer delays by connecting an external computer directly to the robot’s controller via Ethernet, which ensured that read timestamps were accurate.

For convenience, we created a ROS publisher for the TCP poses, ensuring the data was in the same format as the OptiTrack output so all the scripts could read it. In addition, the RTDE read TCP’s rotation, which had to be converted from the Rodrigues’ rotation vector (Angle-Axis representation) to the quaternion representation used in typical ROS messages. We set the frequency of the publisher to 30[Hz], matching the value used by the Unity system.

Calibration movement for this task was again designed to properly saturate different headset poses. This time, the process was automated with a script that contained a set of TCP poses the robot traversed with slow and steady joint movements. The commands were uploaded to the controller with the RTDE library, and during the movement, we recorded both Unity and TCP positional data. It is worth mentioning that in order to have the opportunity to fully control the headsets, whenever they were attached to the mannequin head, we used the casting option, which allowed us to stream the HMD display onto an external device, such as a computer or a smartphone. This was crucial for controlling running apps without taking the headset off and disrupting or altering the values of any transformations.

With proper data recorded, we once again solved the hand-eye calibration problem [14] identically to our previous approach. Since the Unity application ran on the HMD, we again needed to address potential time offsets between Unity and the UR5e. This offset was calculated using the peak detection method described earlier.

5 DETERMINING ROBOT MOTIONS

5.1 Robotic Arm Movement Simulator

Having calculated all the required transformations, we could now determine the precise movements of the robot needed to reproduce any HMD’s real-world trajectory that we previously recorded. This necessitated the creation of a simulation environment where various robotic experiments could be safely conducted. We created such a dedicated Docker ROS environment with the following set of essential packages: UniversalRobots/Universal_Robots_ROS_Driver, that was responsible for creating a general interface between UR robots and ROS; ros-industrial/universal_robot, which provided ROS nodes for communication with UR5e controller; MoveIt [4], a robotics motion planning framework.

To test the robot’s movements, we first had to bring up the exact model of UR5e in the Gazebo [18] simulator with a pre-made launch file. Next, we integrated it with MoveIt, which enabled the control of the model and the initiation of motion planning. We also configured UR5e in RViz, a 3D visualizer for the ROS framework, allowing for free manipulation of the simulated robot, which was convenient for setting up different starting poses.

5.2 Calculating TCP Poses and Robot’s Movements

Following developing an accurate and responsive simulator, we calculated the successive TCP poses based on the real-world head

trajectory and verified their sequential executability, ensuring the trajectory’s accurate reproduction.

Firstly, we loaded the head trajectory and diluted the number of individual poses since OptiTrack recorded the data at 120[Hz], which was an excessive amount to reproduce for a robot causing undesired vibration during UR5e’s movement. The value of this dilution parameter was experimentally fitted and resulted in 12 poses per second. Thus, striking a good balance between accurately reproducing the trajectory and minimizing robot oscillations.

Next, we defined the initial TCP pose, as it directly influenced the entire motion sequence, where each pose had to incrementally follow the calculated movement to accurately replicate the movement of the HMD’s Center Eye Anchor. The correct starting pose, expressed in the robot’s RHS, also needed to be determined for each studied trajectory to help prevent self-collision and to ensure the robot did not attempt to reach positions outside its operational range. Consequently, each OptiTrack pose was converted into an increment relative to the first pose of its respective trajectory.

Our approach ensured consistency in HMD movement in its independent system, regardless of the chosen starting pose. With this in mind, in order to calculate a set of UR5e’s TCP poses of interest for each OptiTrack pose of real-world head trajectory, we needed to solve the following problem, depicted in Fig. 5:

$$T_i = \mathbf{B} \cdot {}^{\text{TCP}}\mathbf{T}_{\text{Unity}} \cdot {}^{\text{Unity}}\mathbf{T}_{\text{Opti}} \cdot \Delta O_i \cdot {}^{\text{Opti}}\mathbf{T}_{\text{Unity}} \cdot {}^{\text{Unity}}\mathbf{T}_{\text{TCP}} \quad (1)$$

where \mathbf{B} is an aforementioned initial pose of TCP in the robot’s coordinate system with an origin at its base, each \mathbf{T} refers to a particular transformation matrix and ΔO_i is an increment of a pose at given index in real-world head trajectory recorded with OptiTrack in its respective coordinate system relative to the starting pose, which can be described as:

$$\Delta O_i = O_0^{-1} \cdot O_i \quad (2)$$

The initial increment in the OptiTrack’s head pose is a special case, where the given delta is simply the identity matrix $\Delta O_0 = \mathbf{I}$




Figure 5: Diagram illustrating the relation between systems when determining TCP poses of interest T_i for each OptiTrack pose of the real-world head trajectory O_i . The diagram consists of two fixed elements: $\mathbf{A} = {}^{\text{TCP}}\mathbf{T}_{\text{Unity}} \cdot {}^{\text{Unity}}\mathbf{T}_{\text{Opti}}$, representing the transformation between the TCP and OptiTrack’s Rigid Body corresponding to the VR headset, and \mathbf{B} , which denotes the initial pose of the TCP in its coordinate system. ΔO_i and ΔH_i represent the changes in poses within their respective systems, relative to the starting point, for the trajectory recorded by OptiTrack and the target TCP trajectory. The final TCP poses of interest can be determined with the formula $T_i = \mathbf{B} \cdot \Delta H_i$, where $\Delta H_i = \mathbf{A} \cdot \Delta O_i \cdot \mathbf{A}^{-1}$.

With the target TCP poses calculated, we determined the corresponding joint positions necessary for accurately completing the given trajectory within the MoveIt framework. This process defined the final robot movement path. At this stage, we adjusted the value of the *eef_step* parameter, which refers to the maximum allowable misalignment of TCP pose during trajectory reproduction in a simulated environment, because its base value of *eef_step* = 1[cm] caused the robot to oscillate during some challenging movements.

The most critical factor at this point was to ensure the exact replication of the robot’s movements each time a given sequence was triggered while eliminating any vibrations or similar disturbances. To achieve this, we reduced the precision required for passing through each TCP pose to *eef_step* = 3[cm], which still resulted in the robot performing a movement analogous to the VR user’s head, at the same time guaranteeing the fluidity of every motion. Consequently, in the vast majority of cases, the reproduction of TCP pose was still extremely accurate. Only a few specific cases caused by unfavorable robot joint alignments were problematic and necessitated increasing the value of the *eef_step* parameter during MoveIt motion planning and simulation.

As a result of the simulation, we obtained the robot’s final path together with the corresponding *fraction* value, which describes the extent to which the robot successfully reproduced the TCP poses for the given parameters and the indicated starting pose. Whenever the robot self-collided, tried to reach a point outside its working area, or could not find a direct movement between following trajectory points, the *fraction* value dropped significantly. Thus, the critical task was to find such a configuration in which the *fraction* = 1.0, which meant that all of the defined poses were achieved. Finding such a configuration per single trajectory required several dozen attempts in the simulation, and only in these cases did we accept the robot’s path and proceed with the experiments.

6 ROBOT REPRODUCED REAL-WORLD TRAJECTORIES

The testing proceeding was identical for each VR headset, beginning with mounting the device on the robot in the same manner as described in Section 4.1. This was followed by the calibration procedure detailed in Section 4.2, which provided accurate TCP-to-Unity transformation.

After completing the above activities, we immediately proceeded to replay the robot’s trajectory obtained from the simulation, resulting in a real-world head trajectory in a test environment on the robot. We accomplished this with the assistance of the RTDE library, which directly sent commands containing a set of target joint positions to the robot’s controller.

During such a test, we recorded Unity’s Center Eye Anchor trajectory with the help of our application installed on the headset and UR5e TCP trajectory with our ROS publisher. It was crucial not to disturb the headset’s position between the calibration process and the final test to ensure accurate calculation of the results. Hence, we remotely controlled the headset by enabling the casting option during all the experiments to eliminate any need to touch or adjust the headset.

7 HEADSETS POSITIONING AND ACCURACY

For the final evaluation and comparison of trajectories recorded with different headsets, we used evo [12], which is a Python package specifically designed for such tasks.

To use this package, every studied trajectory needed to be depicted in a compatible format, therefore we saved them as text files, where each row contains 8 entries: *timestamp x y z q_x q_y q_z q_w* corresponding to timestamps, position, and quaternion orientation, all separated by a space. This trajectory format, first introduced in [7], is known as TUM.

Before saving trajectories in the final TUM-formatted text files, they all needed to be converted to a common coordinate system, and their timestamps had to correspond. We started with converting Unity’s trajectory to the RHS by flipping the *Y* – axis, which required negating *y, q_y, q_w* components.

Subsequently, as the ground truth during the final tests was the TCP trajectory in the robot’s coordinate system, accurately recorded by its controller, we needed to transform each pose from the HMD’s Unity trajectory into that system. Essential to this process was the TCP-to-Unity transformation established at the beginning of the final test. We also needed to ensure that all evaluated trajectories had an identical starting pose so that we could align them with the origin of the robot’s coordinate system. The definitive form of compared trajectories *GT, P*, saved to TUM format, could be described with the following equations:

$$GT_i = \mathbf{X}^{-1} \cdot \text{TCP}_i \quad (3)$$

where GT_i is a final ground truth pose at index i , \mathbf{X} is a transformation from the origin of the robot’s system to the first pose of a TCP trajectory, and TCP_i is a pose from TCP trajectory at index i .

$$P_i = {}^{\text{Unity}}\mathbf{T}_{\text{TCP}} \cdot \mathbf{Y}^{-1} \cdot \text{Unity}_i \cdot {}^{\text{TCP}}\mathbf{T}_{\text{Unity}} \quad (4)$$

where P_i is a pose from the final HMD Unity’s trajectory transformed to the robot’s system aligned with its origin, \mathbf{Y} is a transformation from the origin of the Unity’s system to the first position of its trajectory, and ${}^{\text{Unity}}\mathbf{T}_{\text{TCP}}$ is a TCP-to-Unity transformation, while ${}^{\text{TCP}}\mathbf{T}_{\text{Unity}}$ represents its inverse. For the timestamps, we employed the peak-detection method to synchronize all trajectories to matching periods.

7.1 Evaluation of Calibration Accuracy

We first needed to check how precise were the trajectory reproductions in the non-experimental setup, such as the robot’s calibration movement described in Section 4.2. Therefore, we compared two separate yet identical executions of the specified robot with different headsets mounted on the robot, namely the Meta Quest 2 and Quest Pro. During each execution, we collected TCP and HMD Unity trajectories. As for the results of their respective TCP-to-Unity hand-eye calibration, we obtained a satisfactory transformation with 9.3 [mm] translation and 0.178° rotation error for Quest 2 and 6.2[mm] translation 0.209° rotation error for Quest Pro.

With the necessary transformations calculated, we could now precisely evaluate the performance of the headsets and UR5e during calibration. In the presented results, we removed sections of the idle trajectory before initiating the robot’s movement and cut the trajectory just before the final brake, which caused the mounted headset to vibrate significantly.

The APE for TCP trajectories, meaning the pose differences recorded by the robot’s controller during those calibration motions averaged at 0.467[mm] with root mean square error (RMSE) of 0.057[mm]. These results prove the extreme repeatability of the UR5e robotic arm together with a reliable maximum error of only 1.5[mm]. The magnitude of this error aligns closely with the accuracy specified by the manufacturer, therefore we regarded these results as acceptable.

In the case of the HMDs accuracy, we obtained 3.17[mm] and 3.23[mm] of RMSE for Quest 2 and Quest Pro, respectively, with max errors of 10.48[mm] and 11.97[mm], as shown in Fig. 6. Given that achieving these results for each headset required the use of pre-calculated transformations, the outcomes can be regarded as accurate within the specified environment, with pose tracking being nearly identical for both VR devices.

It is important to note that the volume of the calibration trajectory was approximately 0.3 m³, whereas the volume of the trajectory for real-world headset movement, as detailed in the following




Figure 6: APE plot (up) and diagram (down) for Quest 2, Quest Pro, and TCP trajectories between two identical executions of the robot’s calibration motion, also comparing TCP poses during these movements in red. The abscissa displays Unix timestamps (i.e. no. of seconds since Jan 1, 1970), represented in scientific notation.

sections, ranged from 0.11 to 0.13 m^3 . The larger volume of the calibration trajectory compared to the experimental trajectories is advantageous, as it minimizes errors across most tested poses. The relatively small volume for the experimental trajectories is due to the limited range along the z-axis, a result of the user’s seated position.

7.2 Evaluation of Real-World Headset Movement

Using the same mounting setup, i.e. with the TCP-to-Unity transformation and their respective errors remaining identical, we executed the robot’s trajectory, resulting in the HMD movements recorded in Section 3. The movement was recorded during one full gameplay of *Space Pirate Trainer DX*, which lasted about two minutes. Finally, we evaluated the trajectories for Quest 2 and Quest Pro and compared the poses of TCP during two executions.

Firstly, we noticed that the accuracy of trajectory reproduction for TCP poses in case of much more complicated and faster trajectories deteriorated. In such a scenario, the RMSE equaled 1.53[mm], which is almost a threefold increase as compared to the calibration sequence. The maximum error also increased significantly to 3.69[mm]. As for Quest 2, the mean absolute pose RMSE amounted to 2.29[mm], and the mean error equaled 2.05[mm] with the maximum error reaching 8.32[mm]. In the case of Quest Pro, RMSE was equal to 3.49[mm], and the mean error leveled at 3.16[mm] with the maximum error reaching 8.83[mm] (see Fig. 7).

The final positioning accuracy results for these two headsets were comparable, with a slight indication in favor of the Quest 2, which outperformed Quest Pro in every studied metric except for minimal error.

However, the differences between the headsets were within the limits of calibration accuracy. Therefore, it is difficult to assess unambiguously whether Quest 2 is more accurate. Instead, we can conclude that both headsets maintain high accuracy throughout the

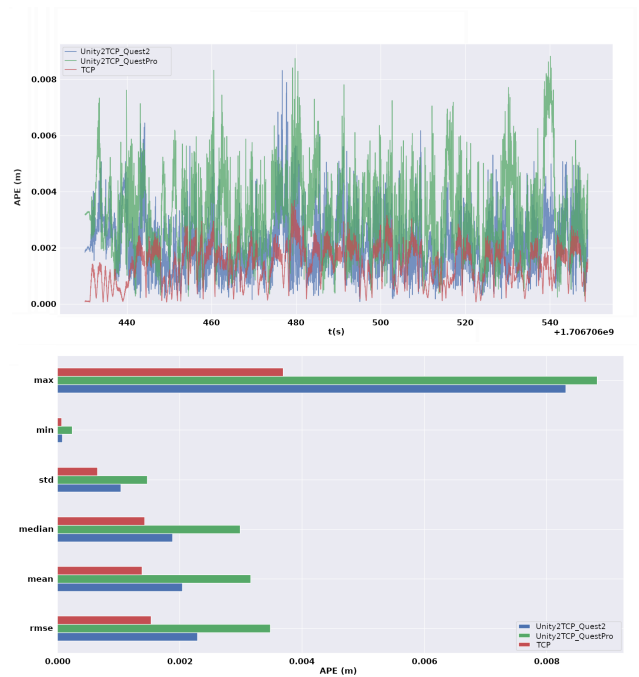


Figure 7: APE plot (up) and diagram (down) for Quest 2, Quest Pro, and TCP trajectories between two identical executions of the robot’s predefined motion plan resulting in Real-World headset movements, with Unix timestamps on the plot abscissa.




Figure 8: The APE plot for a real-world Quest Pro headset trajectory reproduced on a robot. Errors are very small and only visible in the close-up due to the scale.

real-world trajectory without a single error of 9.0[mm].

We repeated the experiment to verify the repeatability of the systems’ performance, yielding the results shown in Tab. 2. As we can observe, the behaviors and relationships between the results are consistent, exhibiting only slightly elevated error values.

Trajectory	RMSE	Mean	Std	Max
Meta Quest 2	3.19	2.85	1.46	9.51
Meta Quest Pro	5.29	4.74	2.34	13.48
TCP	2.40	2.16	1.05	5.11

Table 2: Results for repeated experiment with the real-head trajectory on the robot provided in [mm].

In the analyzed trajectory, we conducted separate tests in which we obstructed two and three out of five RGB cameras on the Quest Pro, each responsible for capturing the surrounding environment. Obscured were successively RGB cameras with id [0, 3] and [0, 2, 3]. The results, presented in Tab. 3, indicate a slight decrease

in tracking accuracy when cameras were obscured. However, covering 2 and 3 cameras yielded comparable results. During this experiment, the hand-eye calibration error was $2.8[mm]$ for translation and 0.179° for rotation.

Trajectory	RMSE	Mean	Std	Max
0 cameras obstructed	3.05	2.69	1.44	8.73
2 cameras obstructed	6.08	5.49	2.62	13.60
3 cameras obstructed	5.56	5.06	2.39	13.59

Table 3: Evaluation of the same trajectory with different numbers of obstructed cameras.

7.3 Evaluation of Long-Period Movement Trajectory

We also examined the accuracy of the headset over a more extended period in real-world conditions. With an exemplary head trajectory recorded during gameplay, we decided to loop it for two hours, simulating an extended gaming session with the Quest Pro.




Figure 9: The APE plot during a two-hour-long looped trajectory with corresponding pieces of axial trajectory from the beginning and the end of the experiment with Unix timestamps displayed on the abscissa.

This test was conducted on a different date than the previous ones, requiring us to first perform another calibration procedure. Its results this time were not so accurate, with a $42.74[mm]$ translation error and 3.96° rotation error. This deterioration was likely caused by an accidental movement of one of the system components. However, for this experiment, the primary objective was to examine the error variation over time. Therefore, we deemed the calibration sufficient for this purpose and proceeded, bearing in mind that values on the ordinate were considerably high.

The visible periodic spike in the pose error is caused by the robot's brake, which occurred during a short moment of reloading the looped trajectory to the controller at its end.

Throughout the experiment, the time offset between the robot and Unity data increased, directly deteriorating the results seen on the APE plot shown in Fig. 9. This also confirms the necessity of our peak detection method, which helps level up this phenomenon. This demonstrates that evaluating longer trajectories recorded in Unity with external ground truth is particularly challenging due to the progressively increasing divergence of their timestamps.

8 CONCLUSION

The presented system for evaluating the accuracy of VR headset positioning, which integrates motion capture and robotic trajectory reproduction, demonstrated high precision in both the collection and reproduction of headset movements. Utilizing advanced spatial calibration and temporal synchronization techniques, we achieved reliable reproduction of real-world VR trajectories.

Our findings indicate that the Meta Quest 2 and Quest Pro headsets exhibit similar levels of accuracy, with the Quest 2 showing a slight advantage in some metrics. The fact that the more affordable Quest 2 matches the accuracy of the more expensive Quest Pro can provide important insight for developers, including those in industrial VR applications. The extended evaluation revealed challenges in time synchronization between the robot and Unity data, underscoring the need for advanced synchronization methods in long-term VR accuracy assessments.

Our approach provides a comprehensive and precise method for assessing VR device's positioning accuracy, which is crucial for enhancing VR experiences and developing more accurate immersive applications. It appears to be the first of its kind to test the accuracy of different headsets on real pre-recorded trajectories and can be easily reproduced from commodity components. Further research can use a motion capture system to independently assess the accuracy of the trajectories replicated by the industrial robot. This shall minimize the influence of any intrinsic calibration issues in the robot arm on the results, but at the expense of a more complicated experimental setup. Moreover, conducting experiments for headsets that are provided by vendors other than Meta and use localization based on different hardware and visual odometry algorithms is an obvious extension.

ACKNOWLEDGMENTS

This work was done at the Poznań University of Technology Centre for Artificial Intelligence and Cybersecurity. The study has been supported by funding provided through an unrestricted gift by Meta.

REFERENCES

- [1] A. Amamra. Smooth head tracking for virtual reality applications. *Signal, Image and Video Processing*, 11(3):479–486, 2017.
- [2] P. Bartyzel, M. Igras Cybulska, S. Tadeja, M. Majdak, G. Łukawski, and D. Hekiert. Voice-Responsive Virtual Reality Training Environment for Occupational and Non-Professional Voice Users. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, 2024. doi: 10.1109/VRW62533.2024.00042
- [3] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.
- [4] D. Coleman, I. Sukan, S. Chitta, and N. Correll. Reducing the barrier to entry of complex robotic software: a moveit! case study, 2014.
- [5] Counterpoint-Research. Extended reality (XR) headset vendor shipment share worldwide from 2020 to 2023, by quarter [Graph]. [Online], <https://www.statista.com/statistics/1222146/xr-headset-shipment-share-worldwide-by-brand/>.
- [6] K. Daniilidis. Hand-eye calibration using dual quaternions. *The International Journal of Robotics Research*, 18(3):286–298, 1999. doi: 10.1177/02783649922066213
- [7] F. Endres, J. Hess, N. Engelhard, J. Sturm, D. Cremers, and W. Burgard. An evaluation of the RGB-D SLAM system. In *IEEE International Conference on Robotics and Automation*, pp. 1691–1696, 2012.
- [8] W. Fang, L. Zheng, H. Deng, and H. Zhang. Real-time motion tracking for mobile augmented/virtual reality using adaptive visual-inertial fusion. *Sensors*, 17(5), 2017.
- [9] Y. Fu, S. Wang, Y. Zhai, and X. Zhan. Visual odometry errors and fault distinction for integrity monitoring. *Aero. Sys.*, 3:1–10, 09 2020.
- [10] J. S. Furtado, H. H. T. Liu, G. Lai, H. Lacheray, and J. Desouza-Coelho. Comparative analysis of optitrack motion capture systems. In F. Janabi-Sharifi and W. Melek, eds., *Advances in Motion Sensing and*

- Control for Robotic Applications*, pp. 15–31. Springer International Publishing, Cham, 2019.
- [11] M. J. Gourlay and R. T. Held. Head-mounted-display tracking for augmented and virtual reality. *Information Display*, 33(1):6–10, 2017.
- [12] M. Grupp. evo: Python package for the evaluation of odometry and slam. <https://github.com/MichaelGrupp/evo>, 2017.
- [13] G. Guennebaud, B. Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
- [14] R. Horaud and F. Dornaika. Hand-eye calibration. *The International Journal of Robotics Research*, 14(3):195–210, June 1995. doi: 10.1177/027836499501400301
- [15] T. E. Humphreys, R. X. T. Kor, P. A. Iannucci, and J. E. Yoder. Open-World Virtual Reality Headset Tracking. In *The 33rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+)*, pp. 2931–2947, September 2020.
- [16] M. S. Ikbal, V. Ramadoss, and M. Zoppi. Dynamic pose tracking performance evaluation of HTC Vive virtual reality system. *IEEE Access*, 9:3798–3815, 2021.
- [17] A. Kim, M. Kwon, M. Chang, S. Kim, D. Jung, and K. Lee. The influence of body position on presence when playing a virtual reality game. In *IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1020–1021, 2019.
- [18] N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, pp. 2149–2154, 2004.
- [19] S. M. LaValle, A. Yershova, M. Katsev, and M. Antonov. Head tracking for the oculus rift. In *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 187–194, 2014.
- [20] Meta. Meta quest 2 virtual reality headset, 2023.
- [21] A. K. T. Ng, L. K. Y. Chan, and H. Y. K. Lau. A low-cost lighthouse-based virtual reality head tracking system. In *International Conference on 3D Immersion (IC3D)*, pp. 1–5, 2017.
- [22] D. C. Niehorster, L. Li, and M. Lappe. The accuracy and precision of position and orientation tracking in the HTC Vive virtual reality system for scientific research. *I-Perception*, 8(3), 2017.
- [23] M. Pollák, M. Kočíško, D. Paulišin, and P. Baron. Measurement of unidirectional pose accuracy and repeatability of the collaborative robot ur5. *Advances in Mechanical Engineering*, 12(12):1687814020972893, 2020.
- [24] D. Reimer, I. Podkosova, D. Scherzer, and H. Kaufmann. Colocation for SLAM-Tracked VR Headsets with Hand Tracking. *Computers*, 10(5), 2021.
- [25] K. Shoemake. Animating rotation with quaternion curves. *SIG-GRAPH*, 19(3):245–254, 1985. doi: 10.1145/325165.325242
- [26] S. Tadeja, P. Seshadri, and P. Kristensson. AeroVR: An immersive visualisation system for aerospace design and digital twinning in virtual reality. *The Aeronautical Journal*, 124(1280):1615–1635, 2020. doi: 10.1017/aer.2020.49
- [27] S. K. Tadeja, T. Bohné, K. Godula, A. Cybulski, and M. M. Woźniak. Immersive presentations of real-world medical equipment through interactive VR environment populated with the high-fidelity 3D model of mobile MRI unit. *Computers & Graphics*, 120:103919, May 2024. doi: 10.1016/j.cag.2024.103919
- [28] C. Theodorou, V. Velisavljevic, V. Dyo, and F. Nonyelu. Visual slam algorithms and their application for ar, mapping, localization and wayfinding. *Array*, 15:100222, 2022.
- [29] R. Tsai and R. Lenz. A new technique for fully autonomous and efficient 3d robotics hand/eye calibration. *IEEE Transactions on Robotics and Automation*, 5(3):345–358, 1989. doi: 10.1109/70.34770
- [30] Universal Robots. UR5e fact sheet. <https://www.universal-robots.com/media/1807465/ur5e-rgb-fact-sheet-landscape-a4.pdf>. Accessed: 2024-07-14.