A virtual reality platform to evaluate the effects of supernumerary limbs' appearance

Ziyi Jiang[1*], Yanpei Huang[1*], Jonathan Eden[1,2], Ekaterina Ivanova[1,3], Xiaoxiao Cheng[1] and Etienne Burdet[1]

***Abstract*— Supernumerary robot limbs (SL) can expand the ability of users by increasing the number of degrees of freedom that they control. While several SLs have been designed and tested on human participants, the effect of the limb's appearance on the user's acceptance, embodiment and device usage is not yet understood. We developed a virtual reality platform with a three-arm avatar that enabled us to systematically investigate the effect of the supernumerary limb's appearance on their perception and motion control performance. A pilot study with 14 participants exhibited similar performance, workload and preference in human-like or robot-like appearance with a trend of preference for the robotic appearance.**

I. Introduction

Supernumerary limbs (SLs) are robotic devices that act as an extension of the user's body granting them new degrees of freedom beyond that of their natural body [1], [2]. With these new degrees of freedom, one user could perform tasks that would otherwise require teams of individuals to work together, increasing efficiency. Due to these benefits, SLs have shown potential for a range of applications, such as robotic surgery [3] or manufacturing [4].

SL users have been shown to be able to control their devices using their other body parts [5] while working with their natural hands to perform three-handed tasks [6]. With training, this has also resulted in a sense of embodiment over the SL [7]. These studies have however only focused on if the user can control a SL, they have not yet considered if the SL's appearance and presentation alter users' performance and perception.

Virtual reality (VR) provides a platform in which different SL appearances can be readily evaluated without safety issues and the influence of hardware, and a platform for benchmarking before developing wearable supernumerary robotic limbs [8]. Here, current studies have considered SL appearances such as a simple virtual cursor [6], [9], a floating human hand [10] or a connected robotic arm [11]. Furthermore, the design of physical wearable SLs has also considered mechanical robot arm appearances [12] or more human-like appearances [13]. It is known that factors such as target size [14] and realism [15] can impact user performance and perception. However, while it is known that SL appearance can affect embodiment [16], it is not yet known if such factors influence SL operation, where the robot is intended to act as an extension of the user's body and appearance might affect user association with the device.

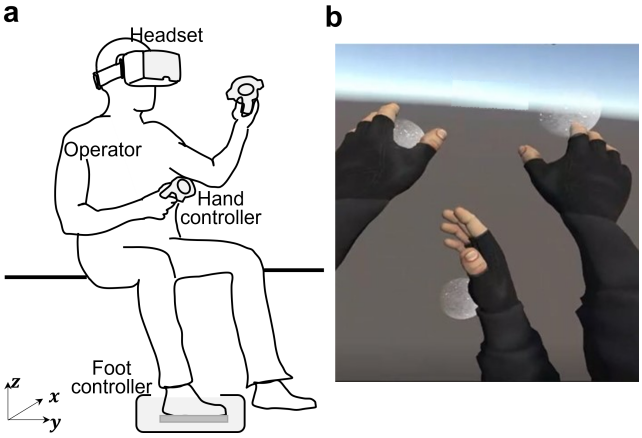

Fig. 1. VR experimental setup (a) and operator view (b). Participants used two hand controllers to move virtual hands 1 and 2, and the foot controller to manipulate virtual hand 3 (supernumerary limb).

In this study, we developed a virtual reality (VR) platform with a module for the SL dynamics and a three-armed avatar with different SL appearances that could be used to evaluate human movement augmentation capability with SLs. An experimental study was then conducted to evaluate the VR platform and investigate how the SL's appearances affect task performance, perceived workload and user preference. For a reaching task, we did not observe an effect of the different appearances in the workload and user preference in both the SL control and the three-hand control modes.

II. Platform

A. VR System

Fig. 1 depicts the developed VR avatar which consists of three virtual arms. When using the platform, the user sat comfortably on a chair, wearing a VR headset (HTC VIVE) which displayed the three-armed virtual avatar in first-person perspective. The user could then control the two virtual natural arms through two hand controllers and a third virtual arm with their dominant foot using a robotic interface. The foot controller was a custom-built haptic passive interface consisting of a pedal connected to an arrangement of eight springs and load cells [17]. Movement of the pedal was

[1] All authors are with the Department of Bioengineering, Imperial College of Science, Technology and Medicine, London, UK. {yanpei.huang, e.burdet}@imperial.ac.uk. [2]Mechanical Engineering Department, the University of Melbourne, Victoria, Australia. [3]School of Electronic Engineering and Computer Science, Queen Mary University of London, London, UK.

*These authors contributed equally. This work was supported by the EC H2020 grants NIMA (FETOPEN 899626) and CONBOTS (ICT 871803).

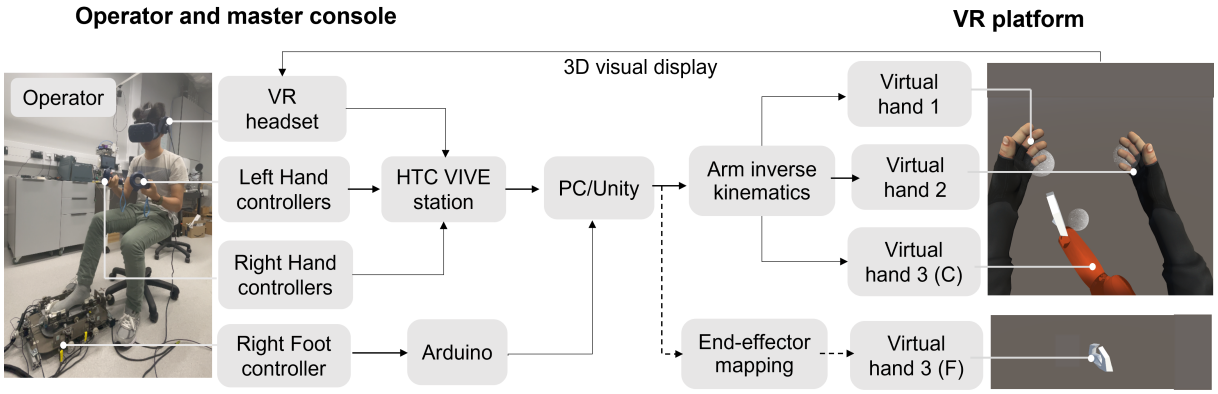

Fig. 2. Platform overview. In the master console on the left, the operator wears a headset, two hand controllers and a foot interface. The hands and foot control signals are collected to control the movement of the virtual arms/hands. Virtual hands 1 & 2 are controlled by the natural left- and right-hand controllers. Virtual hand 3 (supernumerary hand) is controlled by the operator's foot. The controller positions are mapped to the virtual end-effector positions. The joints' angles of the connected virtual arm are then derived based on the inverse kinematics model. The operator can see the virtual limbs in the headset where these arms are extended from the operator's body or float in the air.

measured through a mapping of the spring displacement forces, which could also be felt by the user giving them a sense of the device's position. This interface supports motion in four degrees of freedom (DoF): the pedal's pitch, yaw, and 2-DoF horizontal translations.

Throughout the study, the left and right virtual arms were shown with a human arm appearance. The third virtual hand was then located on the right side of the hip and controlled by the foot controller. The VR platform is a modular system allowing the appearance of the SL to be changed into different forms. In the current appearance library, we have set four SL appearances: Connected human arm and hand (CH); Connected robotic arm and gripper (CR); Floating robotic gripper (FR); and Floating human hand (FH).

For the floating human/robot hand, the hand controller positions in 3-DoF translations are directly mapped to the same positions of the virtual end-effector with a linear scaling. The foot controller also directly controls the third hand's position in the VR scene. Here, the horizontal position of the foot controls the same motion of the SL and the up/down motion was scaled from the pedal's pitch rotation to avoid fatigue from lifting up the leg. For the connected human/robot arm, the joints' position was determined by calculating inverse kinematics.

B. Virtual arm modeling

Using the controlled virtual hand position, the virtual arm was then built using a kinematic model as well as textures that provided the kinematics skeleton with meshes and rendering to enable different appearances and end-effectors.

a) Kinematics modeling: The human arm kinematics was modelled as a 7-DoF serial linkages consisting of a three-DoF should joint, one-DoF elbow joint and three-DoF wrist joint. For the connected hand visualizations the virtual arms' kinematics were calculated using the Forward And Backward Reaching Inverse Kinematics (FABRIK) algorithm [18]. First, the target position of the end-effector in the scene was determined based on the input of the hand/foot controller. The angles and positions of each joint were then calculated by iterating backwards through the wrist, elbow and shoulder positions according to the limb length constraints. As the position of the shoulder should be fixed if the head is not moving, the angle and position of all the joints were calculated again by iterating forward from the starting position of the shoulder, in the order of the elbow and wrist and giving the final position and angle. In addition, the arm's frame was built with respect to the head position and followed the movement of the head.

b) Rendering the arm model: The virtual arm mesh and textures were then built in Blender, where the mapping and binding of joints to the mesh in the inverse kinematic model was completed before it was implemented into Unity. For the human arm (Fig. 3b), the binding of the mesh to all joints was smooth and interlinked, so that each joint change had an effect on the whole mesh, which allowed the arm to behave more naturally and smoothly during movement and twisting. For the robot arm (Fig. 3c), the entire mesh model was split into 32 mesh parts and then individually bound to the corresponding joints. When a joint rotated, only the corresponding mechanical parts were rotated and moved, avoiding distortion of the rigid parts and conforming better to the machine's movement.

The human arm appearance was chosen from TurboSquid [19], while a typical six-joint industrial robot was selected for the robot arm (KUKA KR 6 R900 sixx [20]). Here, this used an alignment of the center of the base, joint 3 and joint 6 of the robotic arm to the joints of shoulder, elbow and wrist of the human arm (Fig. 3c). The robotic arm with gripper was controlled in the same way as the human arm without using its original serial robot kinematics. This was to ensure that the two arms were only different in appearance.

C. Calibration

To fit different users' anatomical properties, we developed a simple calibration procedure for each user to determine their arm lengths and elbow joints position. In this procedure,

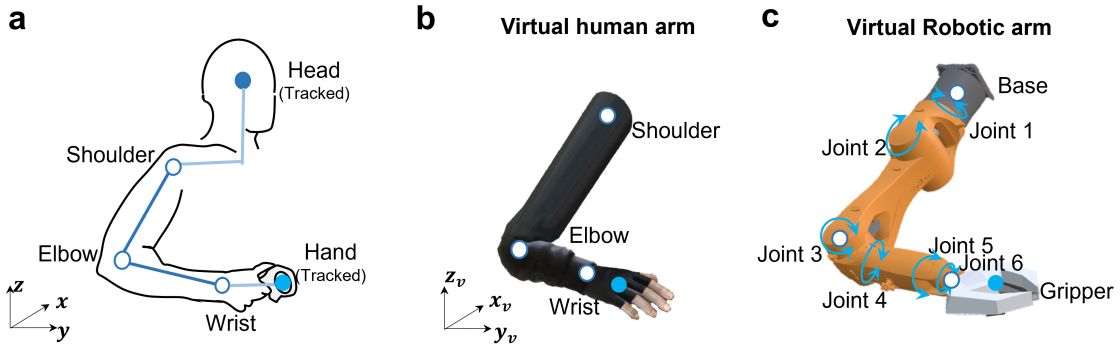

Fig. 3. (a) Human arm kinematics and the visualisations for supernumerary virtual limb in the appearances of (b) human arm and (c) robotic arm.

the operator sat comfortably on the chair and wore the headset while holding the two hand controllers. The head position was tracked by the headset. The shoulder joint was set as a fixed offset from the head position. The operator was instructed to straighten their arms and then move their right hand with the right-hand controller to touch the left arm's elbow joint, and when this occurred to press the controller's trigger button. This position was recorded and identified as the elbow joint. The wrist joint position was directly identified from the hand controller positions. All three virtual arms were set with the same lengths and joints.

III. Pilot Study

As a pilot study of the platform, we evaluated the effect of robotic and human appearances on user performance and perception. The pilot study was approved by the Imperial College London Research Ethics committee (21IC6935). All subjects were informed about the study's purpose and protocol and signed a consent form before starting. 14 subjects without motor impairment (seven male, seven female; mean age = 23 $\pm$1.19 years) participated in the study. Their hand and foot dominance was determined using their Edinburgh Handedness Inventory (EHI) score [21] and the ball-kick dominant leg test [22]. Eleven out of twelve participants were strongly right-handed (Laterality Quotient LQ $>$ 60) while the remaining right-sided participant had mixed dominance (LQ $=$ 54). All participants were right-footed.

A. Procedure

The experiment protocol is depicted in Fig. 4b. The participants first performed the calibration (section II-C) and familiarization session. In the familiarization session, a three-DoF unimanual task was used to familiarize the user with the platform including the use of the hand/foot interfaces and the mapping between the physical limbs and virtual cursors. Then the participants were asked to conduct two test task sessions: a supernumerary-hand task and a three-hand task. Within each task session, the reaching task was performed for 45 trials for each appearance with three blocks for each of the four different appearances (whose order was randomised). After each appearance, participants completed a questionnaire assessing their workload for that given appearance. They also ranked their preference between the four appearances at the end of each session.

B. Task

1) Supernumerary hand task session: In the supernumerary-hand control session, only the SL was displayed (Fig. 4 first row), and the participants were asked to control the SL to reach a target as fast as possible in each trial. In this task, the operator was required to move the virtual hand to reach a series of white transparent target balls that would randomly appear one at a time. If the target was reached, the target ball would disappear and a new ball would appear in a random location. This task also acted as a benchmark to compare the three-hand session against, where there was no interference from the other arms.

2) Three-hand task session: In the three-hand control task all three hands were displayed (Fig. 4 second row), and participants were asked to complete the multi-target reaching task where all three virtual hands were controlled to reach three spatial targets simultaneously. The participant was free to allocate any virtual hand to any target, where once all three targets had been reached, a new set of three targets would appear in new locations. Visual and audio feedback was provided to the participants. When a target was reached, it changed to red color; participants needed to reach all three displayed targets simultaneously, where when the trial was successfully completed a short sound effect was played.

C. Metrics

1) Completion time: The completion time represented the time from the moment that the target set was shown to the moment that it was reached. This only considered the SL in the supernumerary hand session and considered all three targets in the three virtual hands session.

2) Workload and preference: The NASA Task Load Index (TLX) [23] was used to evaluate the workload. Each item was rated for each task within a 100-points range with 5-point steps. Here, the raw NASA TLX was computed as the

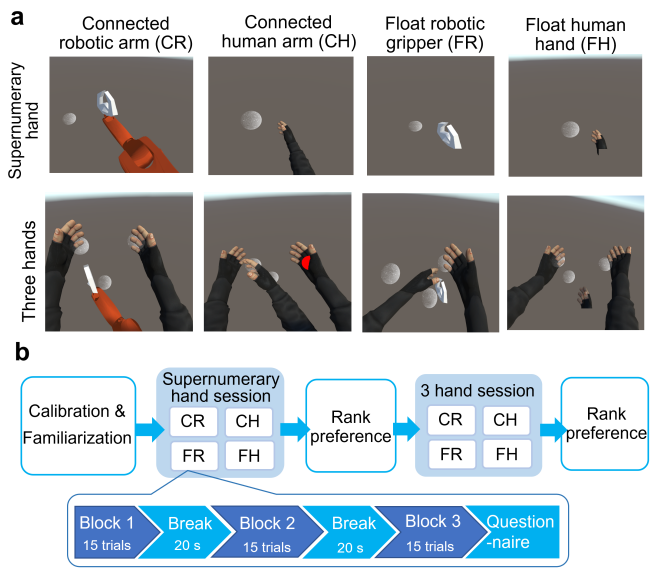

Fig. 4. Experiment conditions (a) and protocol (b).

average score across the six different loads and was used for subsequent analysis. After the supernumerary and three-hand session, we also asked the participants to rank their preference between the four blocks of different appearances (1 – the most preferred, 4 – the least preferred).

D. Statistical analysis

We analyzed the effect of the four different appearances (CR, CH, FR and FH) in the supernumerary-hand session and three-hand session. The normality of the data was tested by the Shapiro-Wilk method. The completion time in the three-hand session and the workload in both the supernumerary-hand and three-hand sessions were normally distributed (all $p > 0.05$). The completion time in the supernumerary-hand session (all $p < 0.013$) and preference ranking were not normally distributed (all $p < 0.04$). For the normally distributed data, a one-way repeated measure analysis of variance (rmANOVA) was used to compare the effect of four appearances. The non-normal distributed data were analysed with the Friedman test. The post-hoc test was carried out for multiple pairwise comparisons.

IV. Results

A. Task completion

The completion time in the supernumerary hand session is illustrated in Fig. 5a. A Friedman test showed that the completion time was affected by the arm visualisation ($\chi^2(3)$= 13.748, p = 0.003). Post-hoc analysis revealed differences in completion time from CH (median = 2.07s) to CR (median = 1.23s) (p = 0.008). The results for the completion time in the three-hand session are shown in Fig. 5b. A rmANOVA showed that there was no effect of the SL appearance on the completion time in the three-hand independent reaching task (F(1.76, 22.89) = 2.82, p = 0.086).

B. Workload

The NASA workload is shown in Fig. 5c. There was no observed difference from using the four appearances on the workload, F(1.78, 23.10) = 1.17, p = 0.32. The overall workloads are about 41.25 ± 15.91, 36.07 ±18.11, 41.19 ± 18.38 and 39.11 ± 17.83 for FH, FR, CH and CR respectively. Similar to the result in the supernumerary-hand session, the overall workload in the three-hand task (Fig. 5d) was not significantly different for four appearances (F(1.61, 20.8) = 0.97, p = 0.38). The average workload was close to but less than 50 (neutral value).

C. Preference

The participant's ranking of the different conditions is depicted in Fig. 5e-f. In the supernumerary-hand session, ten of 14 participants picked the CR or FR as their most preferred appearance (five for each), however, a Friedman test showed no clear difference across the conditions ($\chi^2(3)$=3.43, p=0.330)). For the three-hand session, five of 14 subjects picked CR as the most preferred appearance. Following that, four subjects selected FR as the top appearance, three subjects felt CH was the best and only two participants chose FH as their most preferred appearance. As with the supernumerary hand task, there was no clear impact of the appearance on the ranking ($\chi^2(3)$=3.00, p=0.392)).

V. Discussion

In this paper, we proposed a modular VR platform for evaluating the effect of SL's appearance on performance, perception and preference. The system enables an operator to perform three-handed tasks with a SL of different appearances and will facilitate future studies to evaluate different kinds of SL appearances. We also presented an initial study to evaluate the system and study the effect of the SL appearance on user performance, workload and preference of the SL in controlling the SL alone or in trimanual mode together with the (natural) arm.

Interestingly, the participants could control the SL alone using a foot more efficiently with the connected robotic arm appearance than the connected human hand appearance in the supernumerary hand reaching task, perhaps as the robotic appearance avoids possible confusion with the natural arms. In a task requiring trimanual control thus planning and coordination of multiple hands, we did not find any differences in the completion time when using human or robotic appearances.

Although the appearance of the SL may not affect the performance in three-hand task and the user's workload, it is worth noting there is a trend that the participants preferred the appearance of connected robotic arm among the four appearances. Further investigation over larger populations and different task types will be studied.

The proposed VR system provides multiple further opportunities for future study of human motion augmentation with SLs. First, it may act as a platform to investigate human agency and embodiment of supernumerary limbs in different forms and how that changes with training. Second,

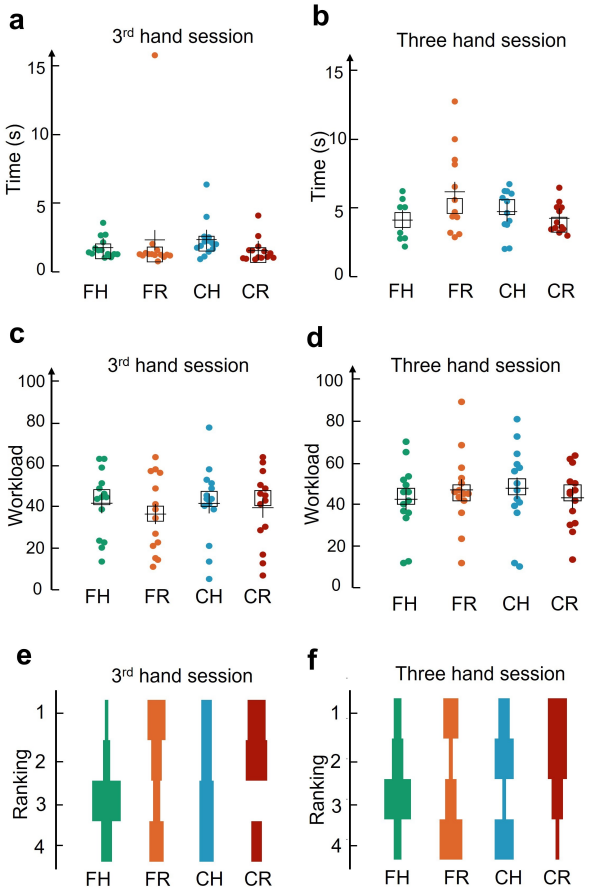

Fig. 5. Completion time in supernumerary-hand (a) and three-hand session (b). Questionnaire result of NASA TLX in supernumerary-hand (c) and three-hand session (d). Subjective ranking of the four appearances in the supernumerary-hand session (e) and three-hand session (f).

the system allows for the benchmarking of different control strategies and inputs, e.g. other hands-free control modes including head motion, eye gaze and voice commands could be compared to the foot-based control considered in this work. Finally, the system can be used to systematically explore the optimal format of the SL (e.g., location, dimensions of the SL) in different trimanipulation tasks.

Acknowledgments

We thank the subjects for taking the time to carry out the experiment.

References

[1] B. Yang, J. Huang, X. Chen, C. Xiong, and Y. Hasegawa, "Supernumerary robotic limbs: A review and future outlook," *IEEE Transactions on Medical Robotics and Bionics*, vol. 3, no. 3, pp. 623–639, 2021.

[2] J. Eden, M. Bräcklein, J. Ibáñez, D. Y. Barsakcioglu, G. Di Pino, D. Farina, E. Burdet, and C. Mehring, "Principles of human movement augmentation and the challenges in making it a reality," *Nature Communications*, vol. 13, no. 1, pp. 1–13, 2022.

[3] Y. Huang, W. Lai, L. Cao, J. Liu, X. Li, E. Burdet, and S. J. Phee, "A three-limb teleoperated robotic system with foot control for flexible endoscopic surgery," *Annals of Biomedical Engineering*, pp. 1–15, 2021.

[4] S.-Z. Ye, P. Jain, A. Walley, Y.-J. Yang, and E. Abdi, "A novel four-degree-of-freedom versus a conventional foot interface for controlling a robotic assistive arm," 2020, pp. 1080–1087.

[5] M. Y. Saraiji, T. Sasaki, K. Kunze, K. Minamizawa, and M. Inami, "Metaarms: Body remapping using feet-controlled artificial arms," in *ACM Symposium on User Interface Software and Technology*, no. 10, 2018, pp. 65–74.

[6] Y. Huang, J. Eden, L. Cao, E. Burdet, and S. J. Phee, "Tri-manipulation: An evaluation of human performance in 3-handed teleoperation," *IEEE Transactions on Medical Robotics and Bionics*, vol. 2, no. 4, pp. 545–548, 2020.

[7] P. Kieliba, D. Clode, R. O. Maimon-Mor, and T. R. Makin, "Robotic hand augmentation drives changes in neural body representation," *Science Robotics*, vol. 6, no. 54, 2021.

[8] F. Aïm, G. Lonjon, D. Hannouche, and R. Nizard, "Effectiveness of virtual reality training in orthopaedic surgery," *Arthroscopy: The Journal of Arthroscopic & Related Surgery*, vol. 32, no. 1, pp. 224–232, 2016.

[9] A. Noccaro, J. Eden, G. Di Pino, D. Formica, and E. Burdet, "Human performance in three-hands tasks," *Scientific Reports*, vol. 11, p. 9511, 2021.

[10] E. Abdi, E. Burdet, M. Bouri, S. Himidan, and H. Bleuler, "In a demanding task, three-handed manipulation is preferred to two-handed manipulation," *Scientific Reports*, vol. 6, p. 21758, 2016.

[11] K. Arai, H. Saito, M. Fukuoka, S. Ueda, M. Sugimoto, M. Kitazaki, and M. Inami, "Embodiment of supernumerary robotic limbs in virtual reality," *Scientific Reports*, vol. 12, p. 9769, 2022.

[12] V. Vatsal and G. Hoffman, "Design and analysis of a wearable robotic forearm," 2018, pp. 5489–5496.

[13] B. L. Bonilla, F. Parietti, and H. H. Asada, "Demonstration-based control of supernumerary robotic limbs," 2012, pp. 3936–3942.

[14] G. Kondraske, "An angular motion fitt's law for human performance modeling and prediction," in *Proceedings of 16th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, 1994, pp. 307–308.

[15] R. P. McMahan, C. Lai, and S. K. Pal, "Interaction fidelity: The uncanny valley of virtual reality interactions," in *Virtual, Augmented and Mixed Reality*, 2016, pp. 59–70.

[16] N. Rosa, R. C. Veltkamp, W. Hürst, A.-M. Brouwer, K. Gijsbertse, I. Cocu, and P. Werkhoven, "Embodiment and performance in the supernumerary hand illusion in augmented reality," *Frontiers in Computer Science*, vol. 3, p. 694916, 2021.

[17] Y. Huang, E. Burdet, L. Cao, P. T. Phan, A. H. T. Meng, and L. Phee, "A subject-specific four-degree-of-freedom foot interface to control a surgical robot," *IEEE/ASME Transactions on Mechatronics*, vol. 25, no. 2, pp. 951–963, 2020.

[18] A. Aristidou and J. Lasenby, "Fabrik: A fast, iterative solver for the inverse kinematics problem," *Graphical Models*, vol. 73, pp. 243–260, 2011.

[19] First Person Arms Pack 3D model, Accessed: May 23, 2022. [Online]. Available: https://www.turbosquid.com/3d-models/unity3d-ar-3d-model-1681559.

[20] Kuka experimental, Accessed: Jun. 10, 2022. [Online]. Available: https://github.com/ros-industrial/kuka_experimental.

[21] R. Oldfield, "The assessment and analysis of handedness: the edinburgh inventory," *Neuropsychologia*, vol. 9, no. 1, pp. 97–113, 1971.

[22] N. van Melick, B. M. Meddeler, T. J. Hoogeboom, M. W. G. Nijhuis-van der Sanden, and R. E. H. van Cingel, "How to determine leg dominance: The agreement between self-reported and observed performance in healthy adults," *PLoS ONE*, vol. 12, p. e0189876, 2017.

[23] S. G. Hart, "NASA-task load index (NASA-TLX); 20 years later," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 50, no. 9, pp. 904–908, 2006.