Movement Guidance using a Mixed Reality Mirror

Qiushi Zhou
The University of Melbourne
Melbourne, Australia
qiushi.zhou@unimelb.edu.au

Andrew Irlitti
The University of Melbourne
Melbourne, Australia
andrew.irlitti@unimelb.edu.au

Difeng Yu
The University of Melbourne
Melbourne, Australia
difengy@student.unimelb.edu.au

Jorge Goncalves
The University of Melbourne
Melbourne, Australia
jorge.goncalves@unimelb.edu.au

Eduardo Velloso
The University of Melbourne
Melbourne, Australia
eduardo.velloso@unimelb.edu.au

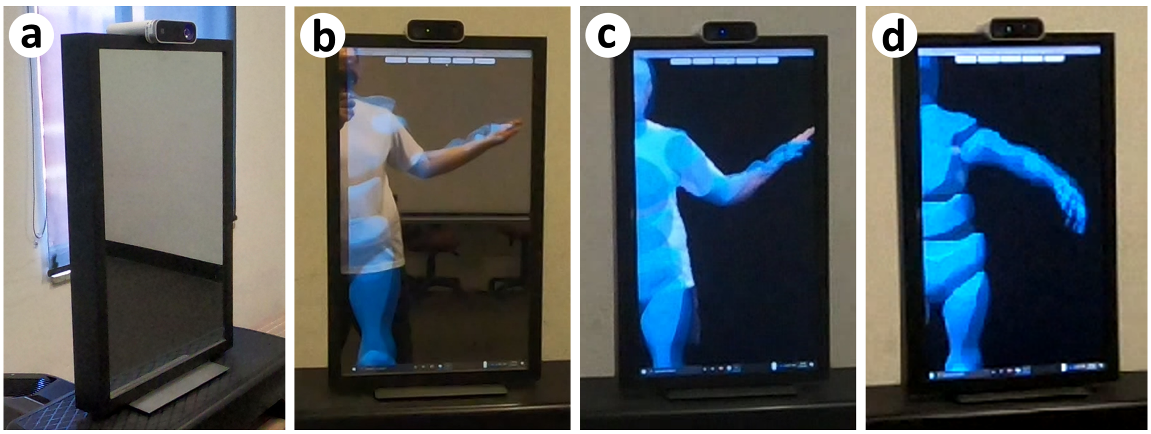

Figure 1: a: Mixed reality mirror consisting of a two-way glass overlaid on a monitor, and a Microsoft Azure Kinect depth sensor; b: A user performing movement acquisition following the guidance from the virtual instructor displayed on the mixed reality mirror. The user's viewing perspective on the virtual instructor is calibrated to their eye-positions following their head movement tracked by the Kinect sensor. The user is able to follow the virtual instructor's arm movement by visually matching it with their own reflection; c: A user performing the same task using a virtual mirror. With the mirror glass removed, the user follows the virtual instructor using their virtual point cloud reflection displayed on the monitor; d: Same task without seeing either the real or virtual reflection of the user.

ABSTRACT

Mirror reflections offer an intuitive and realistic Mixed Reality (MR) experience comparable to other MR interfaces. Their high visual fidelity, and the sensorimotor contingency from the reflected moving body, make the mirror an ideal instrument for MR movement guidance. The translucent two-way mirror display enables users to follow a virtual humanoid instructor's movement accurately by visually matching it with their reflections. In this work, we conduct the first formal evaluation of movement acquisition performance with simple motor tasks, using visual guidance from an MR mirror and a humanoid virtual instructor. Our results of performance and subjective ratings indicate that, comparing with simulated virtual mirror and with traditional screen-based movement guidance, the real MR mirror yields better acquisition performance and stronger sense of embodiment with the reflection, for upper-body movement. But the benefits diminish with larger-range head movements. We provide design guidelines for future mirror movement guidance interfaces and MR mirror experiences at large.

DIS '22, June 13–17, 2022, Virtual Event, USA

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**; **Empirical studies in HCI**; **Mixed / augmented reality**.

KEYWORDS

mirror interface, mixed reality, movement guidance

ACM Reference Format:

Qiushi Zhou, Andrew Irlitti, Difeng Yu, Jorge Goncalves, and Eduardo Velloso. 2022. Movement Guidance using a Mixed Reality Mirror. In *DIS '22: Designing Interactive Systems Conference, June 13–17, 2022, Virtual Event, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

When we see our reflections in a mirror, we experience a most natural and intuitive experience that shares many characteristics of mixed reality (MR). The natural optical reflection makes the mirror a screen that displays a “virtual" scene with lossless visual fidelity and with lighting conditions that automatically blend in with our surrounding. At the same time, an identical representation of the user appears in the mirror, with perfect sensorimotor contingency as the reflection “avatar" flawlessly matches their movements and expressions. The missing component for this scenario to be a complete mixed reality experience is the ability to render additional virtual content. The analogy between a mirror reflection and a virtual avatar in MR is now possible thanks to advances in research related to virtual body ownership illusions. It is now well-understood that realistic visual fidelity and accurate sensorimotor contingency are the two most important factors for the sense of embodiment over virtual avatars [7, 26, 41, 42, 44]. Because these are two strengths of mirror reflections, many researchers have compared them with MR avatars for virtual body ownership illusions. They have found that mirror reflections play a special role in our perception, such that in addition to feeling a strong sense of ownership for the reflection [3, 22, 27], we also perceive the reflected objects in the mirror close to our bodies with an egocentric perspective. This means that we see the mirror reflections from a mental first-person perspective (1PP) instead of a third-person perspective (3PP), despite them being physically external [15, 35, 37].

The benefits afforded by mirrors have been leveraged in many application domains, most notably in motor learning. Concurrent feedback is crucial in the acquisition phase of the motor learning process for learners to correct errors while recreating the movements demonstrated by an instructor [43]. Because concurrent visual feedback requires simultaneous awareness of both the movement guidance and one's own movement, the mirror is the natural interface from which we obtain the latter. For example, in group dancing and fitness classes, instructors and learners are often collocated in front of large mirrors, so that learners are able to observe the movement guidance from the instructor and correct their own movement while they follow along [31, 47]. Despite being adopted in many practical scenarios, this arrangement still presents a challenge to learners: they must balance watching the instructor and attending to their own movement. Further, the disparity between the viewing angle of the instructor and of themselves in the mirror presents an additional cognitive challenge to be overcome [4, 49].

To address this challenge, previous works have explored Mixed Reality (MR) applications to place instructors and learners in the same virtual or augmented environment. Capitalising on the freedom to manipulate space in MR, they have attempted to render the instructor and the learner avatars as superimposed with each other for intuitive movement guidance [8, 16, 40]. While a few works took advantage of the immersive 1PP enabled by Head-Mounted Display (HMD) [16, 18, 57], other works preserved the 3PP for its comprehensive view of the full-body movement [6]. These systems can only adopt either 1PP or 3PP, but cannot achieve both at the same time. To utilise both viewing perspectives, researchers have sought to adopt the mirror metaphor in virtual environments by designing virtual mirrors with superimposition of instructor and

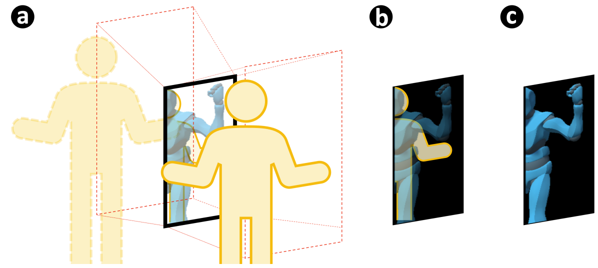

Figure 2: Conceptual diagram illustrating the differences between a real MR mirror, a simulated virtual mirror, and an ordinary screen in the context of this paper. The blue humanoid avatar represents the virtual instructor displayed on the screen, and the yellow figure represents the user. (a) The real MR mirror consists of a two-way mirror overlay on top of a screen. Because the two-way mirror becomes more reflective when less light is projecting through from behind, the black background on the screen becomes fully reflective. Only the virtual instructor is visible from the screen, and it overlaps with the user's real mirror reflection. (b) The virtual mirror overlays the user's RGB point cloud with the virtual trainer, but both images are displayed on the same single-layered screen with a solid black background. (c) An ordinary screen displaying the virtual instructor only.

learner [19, 32, 53, 54]. However, in comparison to real mirrors, virtual mirrors still present latency [53, 54] and limited visual fidelity (See Figure 2 for difference between real and virtual MR mirrors).

Despite the *Holoflector* demonstration by Microsoft already in early 2012 [1], few works have explored or formally evaluated real augmented mirrors as MR interfaces and studied how they can support movement guidance. Those pioneering works all used primitive lines and shapes for guidance through superimposition with the learner, in the form of skeletal figures [1] and body contours [5, 33]. Consequently, those interfaces lose the benefit of leveraging a more intuitive movement interpretation by learning from a realistic humanoid instructor, and could not be evaluated for the users' sense of collocation with the virtual instructor through embodying the reflected virtual space. Whereas the latest commercial home training product *Mirror* [2] simply displays the video recording of an instructor movement on a mirror display, it does so without any consideration of the mapping between the instructor image and the reflection of its user. Despite featuring real mirrors, none of these works exploit the full potential of the mirror as a visually realistic MR interface with perfect sensorimotor contingency.

In this work, we contribute **the first formal evaluation of user performance regarding movement acquisition using a mirror display as an MR interface**. We investigate how the movement acquisition performance is affected by different characteristics of mirror reflection, when it is featured in an MR interface. In this study, we operationalise the characteristics of the mirror reflection as two main factors: FIDELITY encapsulating the visual details of the mirror reflection compared to video displays and the temporal

[1] https://venturebeat.com/2012/02/28/microsoft-holoflector/
[2] https://www.mirror.co/

immediacy with which the reflection updates following user movement, and PERSPECTIVE representing how the optical reflection with the mirror automatically updates our perspective on the reflected scene following head movement. Our results indicate that the real MR mirror provides better depth information of the movement and a better sense of agency and presence, comparing with simulated virtual mirror and with traditional screen-based movement guidance. However, these advantages are reduced for large range of head movements, such as locomotion along the depth axis in front of the mirror. Based on our results, we contribute **design guidelines for future works featuring realistic mirror reflection for movement guidance**.

2 RELATED WORK

In this section, we discuss previous works on movement guidance interfaces in MR, and argue for the potential of MR mirrors as an interface for this purpose. We then give an account of the few interfaces in the literature that features real mirror reflections.

2.1 Movement Guidance in MR with a Virtual Instructor

Augmented feedback is a promising feature in motor learning support systems [43]. Because motor learning often requires seeing a human expert demonstration or its simulation concurrently with the learner's representation, MR environments have been extensively used to place the representations of the learner and the instructor in the same space regardless of physical restrictions. One of the early works in this domain explored the use of immersive HMD VR and evaluated different spatial layouts of the instructors and the learner's avatars [8]. They compared user performance and preference under different layouts and different rendering styles of instructor and learner avatars, including superimposing them with each other. Though they found no significant difference in movement acquisition performance between the conditions, their specific choice of a task involving slow movements—Tai-Chi—and the now dated technology makes it difficult to generalise to today's applications. However, the approach of superimposing the instructor avatar and the learner avatar in MR has inspired a series of later works [8]. These works include early explorations in multimodal movement guidance with projection visualisation [40] and immersive HMD MR applications experimenting with 1PP and 3PP views of the superimposed instructor and learner skeletal figure body representations for static posture guidance [18].

Camporesi and Kallmann evaluated the effects of avatar perspective, stereovision, and display size on the performance of target reaching and of movement reproduction. Using a wall-sized display and a simplistic humanoid avatar, their results highlight the benefits of user perspective, stereovision and the use of avatars [6]. Yu et al. compared user performance of movement acquisition between 3D 1PP view of the arms with corrective visual guidance, 2D mirror view of a video capture of a human instructor, and a 3D 3PP view of a realistic virtual avatar of the user. Although their results indicate that the 1PP view yielded the best performance, the 2D mirror view featured in that work could not match the real mirror reflection in terms of visual fidelity, especially regarding depth information [57]. Overall, despite featuring simulations of real-world mirror reflections in spatial layouts of the instructors and learners, none of those works offered any design or evaluation of interfaces with visual and sensorimotor fidelity matching a real mirror for movement guidance. In this work, we evaluate the user performance and experience with a real mirror screen interface, with which they are able to follow the movement of a virtual instructor superimposed with their reflection in the mirror.

2.2 MR Mirror Interfaces

The use of a mirror as an MR interface has been showcased in an early demo of Microsoft's *Holoflector*, which tracked users' bodies using the first generation of their Kinect sensor. Despite the effective combination of computer graphics and the natural reflection of the mirror, few works have taken this idea of blending mirror reflection with MR rendering further, especially compared to the uptake of other MR platforms such as immersive VR HMDs and see-through AR HMDs. Other demo applications of the Kinect have also featured real and virtual mirrors, such as the virtual clothing augmentation mirror applications [13]. Among previous works exploring the mirror as an MR interface, some have attempted to build MR mirror hardware setups, often borrowing the inspiration of view-dependent rendering [55]. In this kind of rendering, the visual content rendered on the mirror display adapts to user's head movement, remaining consistent with the update of the mirror reflection. Sato et al. first described and implemented such an interface using a half-mirror overlay on a large display while using two video cameras to track markers in front of the display as anchors on which to render virtual objects [38]. Hara and Oda designed an MR virtual mirror display using similar hardware setup for an exhibition featuring CG effects to make it appear as if the user's reflected hand directly interacted with the content displayed in the mirror [17]. Martinez et al. explored different arrangements of a similar setup while enabling interaction to happen around the mirror display [29]. Jang et al. used the Kinect sensor for the same setup with half-mirror overlay on a large display, and proposed methods of calibration between the mirror reflection space and the physical space in front of the mirror captured by the sensor [21]. The authors also explored a focus effect to make the rendering on the mirror display look more realistic [20]. For a review of the use of mirrors or their analogies in MR interfaces, see Portales et al. [34].

A smaller number of previous works used the mirror as a tool for body posture and movement guidance. *Mirrorcle* provided basic visual guides including a body outline of the instructor and corrective lines to indicate angular errors in the user's posture in an effort to ease lower back pain [5]. Similarly, Park et al. described a smart mirror fitness interface that displayed a fixed body contour on a large screen overlaid with a half-mirror and provided corrective feedback on user postures as they were captured by motion tracking cameras [33]. *YouMove* is possibly the one mirror movement guidance interface that is closest to the idea of an MR mirror with full-body tracking and visualisation. It features a wall-sized two-way mirror with back projection. The user sees the 2D skeletal overlay of the instructor on top of their own body, and sees corrective visual feedback indicating mistakes in their posture where their movement deviates from that of the skeletal overlay [1]. However, *YouMove* nevertheless shares similar limitations as the other interfaces we discussed. The most important limitation is that the primitive skeletal overlay does not match the high visual fidelity of the real mirror reflection of the learner. This limitation may prevent the learners from accurately perceiving the movement guidance, such as its depth dimension [11, 30]. Consequently, movement guidance from a humanoid avatar should be more intuitively perceived by users than the guidance from a skeletal overlay [11]. Additionally, because humanoid figures are more familiar to people's daily experience than skeletal figures consisting of straight lines, there are potential benefits from social and psychological perspectives [12]. Also, an MR interface provides better sense of embodiment such as ownership over the self-representation and sense of presence within the reflected MR mirror space, which are potential benefits for movement acquisition performance [35]. Finally, a thorough evaluation of a movement guidance MR interface built upon mirror reflections has yet to be conducted. We present the first of such evaluation by operationalising the two traits of the mirror as an MR interface, namely the high visual fidelity and the sensorimotor contingency it offers with the synchronised movement and the viewing angle of the reflected body following user movement.

2.3 Embodying the Mirror Reflection

Recent research in neuroscience has shown how the sense of ownership of an external body can be induced by a mirror reflection of the whole or part of the body, in a way similar to virtual avatars in MR [35]. The body ownership illusion is a well-studied phenomenon in which healthy people locate themselves in more than one body at the same time while being exposed to humanoid body reduplication, such as virtual avatars [2]. Previous works have found that the anthropomorphic features of the virtual body, its visual fidelity, and the richness of sensory feedback, all have an impact on the feeling of body ownership and agency [42]. Virtual body parts that have realistic personalised visual features increase body ownership and spatial presence [26]. In virtual environments, the visual fidelity of the user body representation modulates the sense of presence and the perception of interaction affordances [41]. Considering movement, other works have found that a realistic sensorimotor mapping between the virtual and the real body helps induce a sense of embodiment [7] and induce a sense of presence in virtual environments [44].

The mirror, as the natural visual interface that offers perfect visual fidelity and sensorimotor contingency through the reflection of its user has been recognised and used by researchers to evaluate its effect on virtual body ownership. They have found that the mirror holds a special place in our perception, such that in addition to feeling a strong sense of ownership for the reflection [3, 22, 27], we also perceive objects reflected in the mirror close to our bodies with an egocentric perspective. Those works suggest that we see mirror reflections from a mental 1PP, and that our mind processes the objects reflected in the mirror as if they were next to our physical bodies instead of being on the other side of the mirror surface [15, 35, 37]. Previous works have found that movement acquisition benefits from seeing the self-avatar in 1PP in MR [57], and that the high visual fidelity of the self-avatar also benefits spatial perception [11]. In light of these findings in the literature, we employ a humanoid virtual instructor and investigate if the egocentric perspective and the high visual fidelity of the mirror reflection contribute to a better sense of embodiment within the same space with the virtual instructor, which subsequently benefit movement acquisition performance. In this work, we evaluate such a setup and ask questions regarding sense of presence and agency over mirror reflection, while comparing quantitative performance.

2.4 Summary

Based on our review of the literature, we identify a vast potential for building ubiquitous movement guidance applications by augmenting real mirrors as realistic MR interfaces. While previous works demonstrated the benefit of superimposing instructors and learners in MR using simulated virtual mirrors or using primitive skeletal overlay on top of real mirrors [8, 18, 40], **no previous work had been able to evaluate augmented real mirror reflection as an MR experience for its movement acquisition performance and the sense of embodiment over the augmented reflection space.** In this work, we evaluate the mirror as an MR experience by superimposing the learner's real mirror reflection with an humanoid virtual instructor, while featuring view-dependent rendering (Figure 1,4). We hypothesise that the benefits of the MR mirror for movement acquisition are twofold: the translucent two-way mirror display blends the virtual image with the reflection while preserving its premium visual details and sensorimotor contingency; the viewing perspective on the virtual instructor adheres to user's eye positions, hence always enables the correct 1PP matching between user reflection and instructor avatar, while offering a realistic MR experience. With these two characteristics of mirror reflection equipped by our apparatus, **we investigate the effect of visual FIDELITY and the effect of realistic mirror PERSPECTIVE change on movement acquisition performance.** Additionally, because the literature suggests that we experience egocentric spatial perception through mirror reflection [15, 35, 37] with potential benefits for movement acquisition [18, 57], **we collect subjective feedback regarding AGENCY and PRESENCE to complement our objective measures of performance.**

3 METHOD

Informed by previous works, we designed and conducted a user study to evaluate the effect of the MR mirror interface on movement acquisition performance and on the sense of embodiment over their mirror reflection during the movement tasks. Specifically, we are interested in investigating the following research questions:

- RQ1: Does the real MR mirror enable better movement acquisition performance due to its superior visual and sensorimotor fidelity comparing with simulated virtual mirror and traditional screen-based interface;
- RQ2: Does the real MR mirror enable better movement acquisition performance due to its intuitive perspective updating when the users' heads move in front of the mirror;
- RQ3: Does superimposing a virtual humanoid instructor with the real mirror reflections of the users induce a stronger sense of agency and presence than with their virtual reflections?

The study was conducted under the approval of the Human Ethics Committee at The University of Melbourne.

3.1 Study Design

To evaluate the effect of the MR mirror interface on user movement acquisition, we compare the mirror-reflection feedback with traditional video-based feedback that are widely adopted in online training videos and in popular Kinect-enabled training applications, such as *NIKE+ Kinect Training* [3] and *EA SPORTS Active 2.0* [4]. Following previous works, we display the instructor image superimposed on the learner's reflection on the mirror [8, 16, 40]. The translucent optical nature of the mirror enables intuitive movement guidance and correction offered by the superimposition without inducing occlusion. To better understand which properties of the mirror are most important for movement guidance, we operationalise the two traits of mirror reflection compared with non-mirror displays. We encapsulate the mirror's high visual fidelity and the sensorimotor contingency into the independent variable dubbed Fidelity. Another important component of the realistic experience offered by mirror reflections is that the spatial viewing perspective on the reflection automatically updates according to the user's head movement. This perspective adaptation offers potential benefits, including continuously providing the correct viewing angle on the instructor's image and better immersive experience. We employ the second independent variable Perspective in the study to evaluate how the availability of this perspective change enabled by head-tracking affects movement acquisition performance.

We designed the study to include two phases. Phase One features two movement tasks performed while standing in place, and Phase Two features two movement tasks that require vertical or horizontal body movements. We employed a repeated-measures design with the independent variable Fidelity for Phase One, and a 3×2 repeated-measures design for Phase Two with two independent variables: Fidelity and Perspective. We did not employ Perspective in the Phase One tasks because they were performed while standing in place without head movement. We compared user performance across three levels of Fidelity patterns (see Figure 3):

- Real Mirror for when participants perform movement acquisition tasks following a virtual instructor displayed on a screen overlaid with a two-way mirror,
- Virtual Mirror in which the virtual instructor is displayed with participants' virtual reflection (video-captured avatar), instead of the real mirror reflection, on an ordinary screen,
- Baseline for when participants only see the instructor on an ordinary screen without their own avatar.

We also compare the acquisition performance with different Perspective updating mechanisms (see Figure 4):

- Dynamic where the participant's viewing perspective on the virtual instructor is constantly calibrated to their viewing angle on the reflection following their head movement
- Fixed where the perspective is fixed according to the head position of the participant at the start of the trial, without further updating despite subsequent head movement.

We measure the Euler angular offsets of the relevant joint rotations for each movement performed by the instructor and the participants as the performance of movement acquisition. We obtain the rotation angles of the shoulder joint around the x (pitch), y (yaw), and z (roll) axis, and the bending (yaw) angle of the elbow joint, as deviation from a T-pose (Figure 5). We measure the offsets in the rotations which are relevant to the correct reproduction of each movement. For Stretch, we measure Shoulder Offset Y, Shoulder offset Z, and Elbow offset, as it requires the arms to be lifted and stretched back while remaining extended straight. For Float, we measure the offsets from all the rotations because it involves the shoulder movement in all directions and the bending of the elbow. For Squat, we measure Shoulder Offset Y, Shoulder Offset Z, and Elbow offset, as this movement involves maintaining the height and width of the shoulder extension and slightly lifting the forearm during the squatting motion. For Entry, we measure all the rotation offsets as it requires the arm to be in the overall correct posture relative to the torso.

We collect subjective ratings from participants after each trial regarding the task difficulty, using the Single Easement Questionnaire. This questionnaire asks participants to rate how easy the task was on a 7-point Likert scale [39]. We also collect subjective ratings for embodiment over their self-representation during the task, either

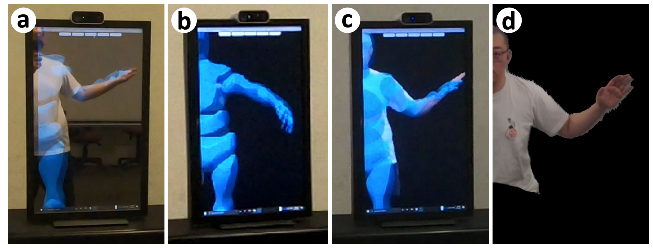

Figure 3: Three levels of Fidelity: (a) Real Mirror for when participants perform movement acquisition tasks following a virtual instructor displayed on a screen overlaid with a two-way mirror; (b) Virtual Mirror in which the virtual instructor is displayed with participants' virtual reflection (video-captured avatar), instead of the real mirror reflection, on an ordinary screen; (c) Baseline for when participants only see the instructor on an ordinary screen without their own avatar. (d) The virtual reflection (for clarity).

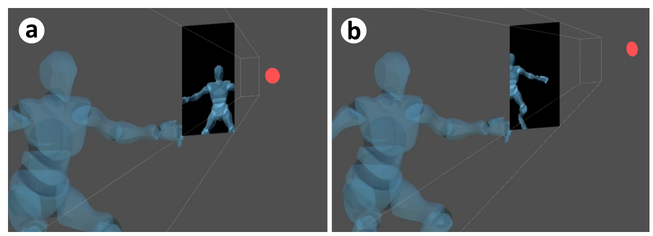

Figure 4: Perspective update following user's eye positions: (a) a casual perspective on the virtual instructor as captured on the screen (black plate) according to the participant's eye positions (red dot); (b) the perspective during the trials, with the participant's head facing the left edge of the screen, which displays the left half of the instructor.

[3] https://news.nike.com/news/introducing-nike-kinect-training
[4] https://www.ea.com/en-au/news/ea-announces-ea-sports-active-2

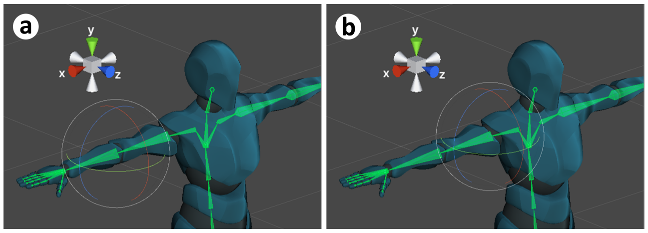

Figure 5: Rotation axis for (a) elbow and (b) shoulder joints.

being the real or the virtual reflection. We adapted the questions regarding "agency and motor control" and "location of the body" in the embodiment questionnaire proposed by Gonzalez-Franco and Peck, and simplified them as two questions regarding the sense of agency and presence, which are appropriate for an MR mirror interface [14, 25]. The ratings were collected at the end of each task, using a questionnaire containing three questions answered with a 7-point Likert scale: *1. Overall, how easy do you think the task was? 2. How much did your reflected/captured arm's actions correspond with your control? 3. To what extent did you feel you were collocated in the same space with the virtual instructor?* Aiming at addressing the research questions RQ1—RQ3, we hypothesised that:

- H1: Movement acquisition errors and perceived task difficulty would be lower with Real Mirror than with Virtual Mirror and Baseline;
- H2: Subjective ratings of the sense of embodiment over the reflections would be higher for Real Mirror than for Virtual Mirror and Baseline;
- H3: For movement tasks with perspective change induced by head movement, movement acquisition errors and perceived difficulty would be lower with Dynamic than with Fixed;
- H4: For movement tasks with perspective change induced by head movement, subjective ratings of the sense of embodiment over the reflections would be higher for Dynamic than for Fixed.

3.2 Apparatus

We assembled the MR mirror setup using a two-way glass mirror (70% reflective, 30% transparent) overlaid in front of a 27 inch screen (brightness: 300 cd/m²) with a custom-built frame under a Microsoft Azure Kinect sensor [5] (Figure 1). The virtual positioning of the Kinect sensor relative to the screen in Unity was calibrated to match their relative positions in reality, on a PC that features an NVIDIA® GeForce® GTX 1080 graphics card, connected with two 27-inch 1080P monitors. The size of the mirror display allows us to evaluate the performance of movement acquisition by measuring and providing visual feedback for a half of the body, which is proven to be appropriate for indicating movement acquisition performance as evidenced by previous work [45]. We implemented the software of the MR mirror in Unity 3D by tracking participants' eye positions with the Kinect and updating the viewing perspective on the instructor accordingly (Figure 4).

We employed a virtual avatar without prominent personal or social traits as the virtual instructor to minimise potential biases. The instructor was scaled to match the body size of each participant in real time, to provide personalised visual feedback [46]. As such, they were able to follow the movements performed by the instructor by visually matching their reflections with the instructor's image in the Real Mirror condition. In the Virtual Mirror condition, the same task could be achieved by matching a video capture (virtual reflection) of the participants with the instructor. We rendered the instructor avatar at 40% opacity such that it did not completely occlude the participant's real or virtual reflection, while still providing enough contrast to make all the arm joints visible (Figure 3). The virtual reflection of the participant was rendered using the RGB point cloud captured by the Kinect sensor at 720p, 30FPS. We measured the delay in the body tracking and rendering of the Kinect sensor using the timecode-view method proposed by [50], and found it to be 165 ms as the mean value of ten attempts (sd = 2.27). We removed the background in the Virtual Mirror condition and only kept the video point cloud avatar of the participants. We removed the background because it better resembles the video-based movement training interfaces, and that it is easier for the participants to recognise their body contour without the distraction induced by the low-res background, as it is at a further distance from the Kinect sensor than the participants (Figure 3).

3.3 Tasks

In this study, we aim to evaluate how an MR mirror helps with movement acquisition, for which we measure the performance as joint angular offsets. Due to the limited size of the interface (27 inch), we focus on the movement of the right arm in this study. Participants were asked to focus on matching their right arm movement but to perform the observed instructor movement with both arms. We employed four movement tasks of increasing complexity, obtained from the online animation library Mixamo [6] (Figure 6):

- Stretch: alternate between a lower backward arm stretch, a higher backward arm stretch, and an idle state;
- Float: a swimming-like arm movement with the two arms paddling forward in alternate phases;
- Squat: a squat movement with the upper-arms remain extended while the forearms gradually lifting as the squat sinks;
- Entry: an entry-to-the-fight movement sequence, with the instructor walking two steps forward, performing the final pose, then returning to the starting point.

While Stretch and Float were performed without lower-body movement, Squat involved vertical body movement, and Entry involved body movement along the depth axis from the display. We only recorded the joint data during key frames for Stretch when the arm was stretching out and pausing for one second, and for Entry when the instructor pauses at the final pose for one second. We recorded the joint data continuously for the other two movements. At the beginning of each trial, the instructor is placed at the same location as the participant by matching their pelvis and both shoulder joints in all tasks but Entry, in which the distance walked by the instructor is determined automatically by its size.

[5]https://azure.microsoft.com/services/kinect-dk/

[6]mixamo.com

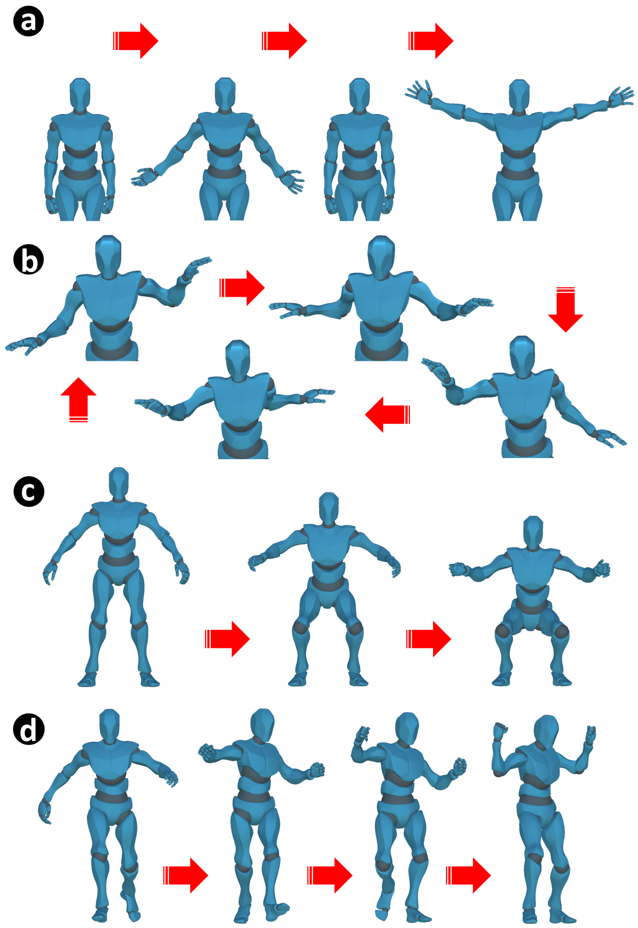

Figure 6: Movement tasks: (a) STRETCH: alternate between a lower and a higher backward arm stretch, and an idle state; (b) FLOAT: a swimming-like movement with two arms paddling forward in alternate phases; (c) SQUAT: a squat movement with the upper-arms remain extending, while the forearms lifting as the squat sinks; (d) ENTRY: an entry-to-the-fight sequence, with the instructor walking two steps forward, performing the final pose, then two steps back.

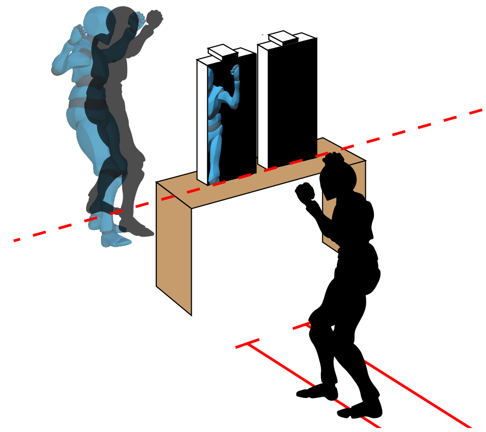

Figure 7: Experimental setup. One MR Mirror display and one ordinary video display, each with a Kinect sensor, were placed side-by-side on a table. Participants (silhouette on the right) stood on the red line facing the left edge of the display used in the current trial. The instructor's image (coloured-figure) is overlaid with the reflection of the participants.

3.4 Participants

We recruited 24 right-handed participants (14 women/9 men/1 non-binary) with a mean age of 24.17 years (Min = 18, Max = 30, SD = 3.48) through university mailing lists. The study lasted one hour on average for each participant, with a $10 gift card compensation.

3.5 Procedures

Upon arrival, participants were informed of the purpose of the study and asked to sign a consent form. We instructed participants to stand in front of the two monitors, one with the mirror overlay and one without. We drew two lines on the floor using red tape to indicate the standing positions. Each line faces the left edge of the active display in the respective condition (the ordinary display for VIRTUAL MIRROR and BASELINE, and the mirror display for REAL MIRROR) (Figure 7). They were allowed to stand on the line between 1.5m to 2m away from the displays, so that they could see the entire right half of the body, with the whole arm visible in the real or virtual reflection as it extended out. We asked the participants to keep the same distance from the displays for the entire study.

To familiarise participants with the interface, we started with a practice movement, in which the instructor stands still and extends the right arm out repeatedly. Before the trials started for each movement, we demonstrated the movement in the Unity editor with a view covering the entire body of the instructor so that the participants could get a holistic view of the movement. We instructed the participants to follow the movements by aiming to correctly reproduce the accuracy and the quality of the arm movement, as if they were learning from a human instructor. We asked the participant to follow the practice movement until they understood the task. Then we started the four movements in the order of increasing complexity, namely STRETCH and FLOAT in PHASE ONE, followed by SQUAT and ENTRY in PHASE TWO. Because ENTRY involved walking back and forth and because the lower-body of the instructor was not visible in the displays, we demonstrated the walking steps to the participants and only commenced after they were able to replicate the sequence. In PHASE ONE, participants performed six repetitions under each of the three FIDELITY levels for each movement. In PHASE TWO, we explained that we added the PERSPECTIVE variable, and informed the participants that they did not have to recognise the difference because they might be too subtle to distinguish, and as a result that they needed to perform twice the repetitions in PHASE TWO. While the order of the four movement tasks were the same for each participant, the order of the conditions within each movement was balanced using a Latin Square. We asked the participants to fill in the questionnaire after each trial, and conducted a semi-structured interview after each movement for the reasons behind their ratings and for their experience.

4 RESULTS

We recorded participant joint angles at 30FPS during key frames, which resulted in 17,199 data points for Stretch, 36,996 for Float, 64,137 for Squat, and 16,817 for Entry. To minimise the effect of the delay between the capturing and the rendering, we time-matched the reference curve of the key movement between the instructor and the participant for each movement. We chose the key movement as the most prominent change that marks the overall rhythm of each task. We used the shoulder z rotation (up and down) for Stretch, the height of the hand for Float, the height of the torso for Squat, and the distance between the torso and the displays for Entry as the key movements for the time matching. To minimise the effect of the trials with task-irrelevant mistakes, we removed outliers (30 frames, 0.2% for Stretch, 62 frames, 0.2% for Float, 126 frames, 0.2% for Squat, 59 frames, 0.4% for Entry) in which any of the offsets between the instructor and the learner was above three standard deviations from the mean ($mean + 3sd$.). We averaged the absolute normalised values of the performance measures for each trial performed by each participant. That resulted in 72 data points (24 participants × 3 Fidelity levels) for each of Stretch and Float, and 144 data points (24 participants × 3 Fidelity levels × 2 Perspective levels) for each of Squat and Entry. The questionnaire after each trial resulted in 216 ratings.

For each performance measure, we applied the Tukey's Ladder of Powers transformation [48] if non-normal distribution of residuals was identified, and subsequently applied the Aligned Rank Transform (ART) if the violation of either residual normality or homoscedasticity persisted (only for Shoulder Offset Z in Squat) [56]. Next, we performed a Repeated-Measures ANOVA on each response variable and post hoc pairwise comparisons with Holm–Bonferroni adjustment to analyse the accuracy of movement acquisition. For Likert scale subjective ratings on task difficulty and on the embodiment questions, we performed non-parametric Friedman tests (with post hoc pairwise comparisons using Conover's test) for Stretch and Float, and transformed the data using ART and performed RM ANOVA for Squat and Entry. We present measures of response variables and all the statistically significant interaction effects in this section. We visualise the results in Figure 8-17, in which the error bars indicate 95% confidence interval.

4.1 Performance Results

4.1.1 Stretch. There were no significant difference in any response variable regarding movement acquisition performance in Stretch.

4.1.2 Float. For Shoulder Offset X, the effect of Fidelity was significant ($F_{2,46} = 3.23, p < .05$), while the offset in Real Mirror ($mean = 16.408, sd = 3.72$) is significantly smaller than in Baseline ($p < .01, mean = 18.190, sd = 3.19$). For Shoulder Offset Z, the effect of Fidelity was significant ($F_{2,46} = 4.11, p < .05$), while the offset in Real Mirror ($p < .05, mean = 10.628, sd = 2.84$) and in Virtual Mirror ($p < .05, mean = 10.508, sd = 2.52$) are significantly smaller than in Baseline ($mean = 12.548, sd = 4.03$). For Elbow Offset, the effect of Fidelity was significant ($F_{2,46} = 12.15, p < .001$), while the offset in Real Mirror ($mean = 28.055, sd = 9.32$) was significantly smaller than in Virtual Mirror ($p < .05, mean = 31.415, sd = 11.03$), which is significantly smaller than in Baseline ($p < 0.05, mean = 34.680, sd = 8.95$) (Figure 9).

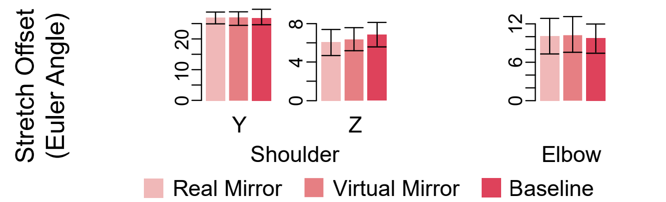

Figure 8: Offsets in joint angles between instructor and participant during Stretch with varying Fidelity.

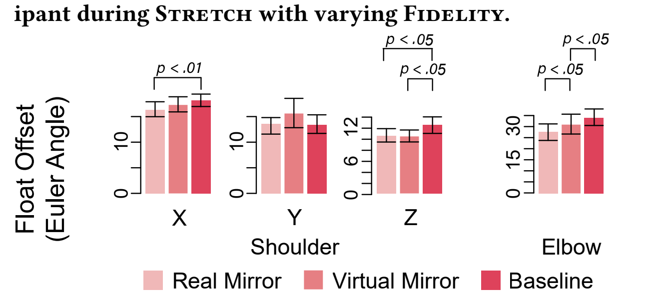

Figure 9: Offsets in joint angles between instructor and participant during Squat with varying Fidelity.

4.1.3 Squat. For Shoulder Offset Z, the effect of Fidelity was significant ($F_{2,115} = 30.83, p < .001$), while the offset in Virtual Mirror ($mean = 13.607, sd = 4.99$) was significantly smaller than in Real Mirror ($p < .05, mean = 15.605, sd = 7.54$), which is significantly smaller than in Baseline ($p < .001, mean = 18.384, sd = 5.32$). For Elbow Offset, the effect of Fidelity was significant ($F_{2,46} = 17.931, p < .001$), while the offset in Real Mirror ($p < .05, mean = 25.308, sd = 8.42$) and in Baseline ($p < .05, mean = 25.131, sd = 8.75$) are significantly smaller than in Virtual Mirror ($mean = 30.092, sd = 8.58$)(Figure 10, 11).

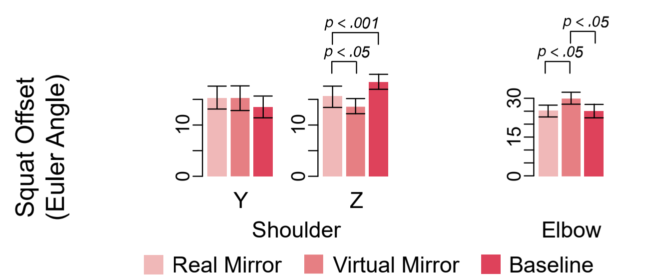

Figure 10: Offsets in joint angles between instructor and participant during Squat with varying Fidelity.

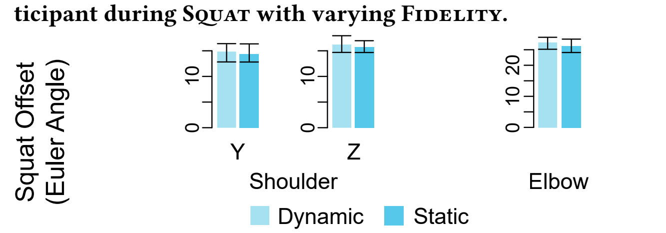

Figure 11: Offsets in joint angles between instructor and participant during Squat with varying Perspective.

4.1.4 Entry. For SHOULDER OFFSET X, the effect of PERSPECTIVE is significant ($F_{2,46} = 9.09, p < .01$), while the offset in TRUE (*mean* = 17.984, *sd* = 9.14) is smaller than in FALSE (*mean* = 20.128, *sd* = 10.26) with insignificant post hoc test results (Figure 12, 13).

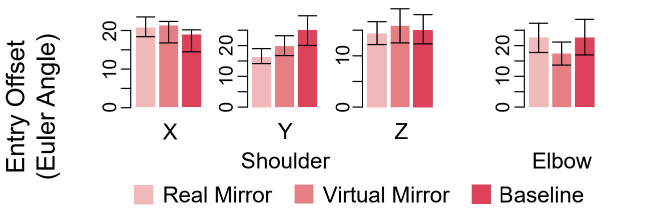

Figure 12: Offsets in joint angles between instructor and participant during ENTRY with varying FIDELITY.

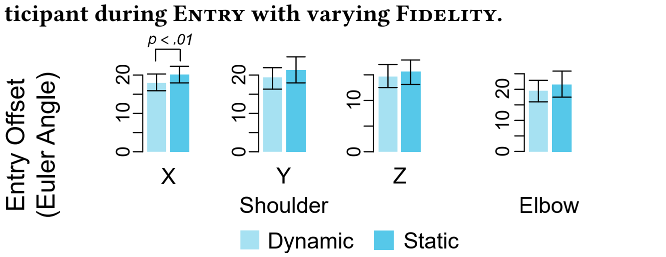

Figure 13: Offsets in joint angles between instructor and participant during ENTRY with varying PERSPECTIVE.

4.2 Subjective Ratings

4.2.1 Stretch. For EASE, the effect of FIDELITY was significant ($\chi^2_2 = 6.79, p < .05$), while the rating in REAL MIRROR (*mean* = 5.917, *sd* = 0.97) is significantly higher than in VIRTUAL MIRROR ($p < .01$, *mean* = 5.500, *sd* = 0.88), which is significantly higher than in BASELINE ($p < .001$, *mean* = 4.917, *sd* = 1.61). For AGENCY, the effect of FIDELITY was significant ($\chi^2_1 = 17.19, p < .001$), while the rating in REAL MIRROR (*mean* = 6.333, *sd* = 0.81) is significantly higher than in VIRTUAL MIRROR (*mean* = 4.375, *sd* = 1.35). For PRESENCE, the effect of FIDELITY was significant ($\chi^2_2 = 7.73, p < .05$), while the rating in REAL MIRROR (*mean* = 5.167, *sd* = 1.31) is significantly higher than in VIRTUAL MIRROR ($p < .01$, *mean* = 4.625, *sd* = 1.10), which is significantly higher than in BASELINE ($p < .001$, *mean* = 3.833, *sd* = 1.63) (Figure 14).

4.2.2 Float. For EASE, the effect of FIDELITY was significant ($\chi^2_2 = 10.93, p < .01$), while the rating in REAL MIRROR (*mean* = 5.208, *sd* = 1.14) is significantly higher than in VIRTUAL MIRROR ($p < .001$, *mean* = 4.042, *sd* = 1.00) and in BASELINE ($p < .001$, *mean* = 4.000, *sd* = 1.79). For AGENCY, the effect of FIDELITY was significant ($\chi^2_1 = 15.70, p < .001$). The rating in REAL MIRROR (*mean* = 5.917, *sd* = 1.02) is significantly higher than in VIRTUAL MIRROR (*mean* = 3.875, *sd* = 1.23). For PRESENCE, the effect of FIDELITY was significant ($\chi^2_2 = 18.69, p < .001$). The rating in REAL MIRROR (*mean* = 5.042, *sd* = 1.43) is significantly higher than in VIRTUAL MIRROR ($p < .001$, *mean* = 4.208, *sd* = 1.25), which is significantly higher than in BASELINE ($p < .001$, *mean* = 3.125, *sd* = 1.78) (Figure 15).

4.2.3 Squat. For AGENCY, the effect of FIDELITY was significant ($F_{1,69} = 19.08, p < .001$). The rating in REAL MIRROR (*mean* = 5.375, *sd* = 1.36) is significantly higher than in VIRTUAL MIRROR (*mean* = 4.354, *sd* = 1.38). For PRESENCE, the effect of FIDELITY was significant ($F_{2,115} = 18.57, p < .001$). The rating in REAL MIRROR ($p < .001$, *mean* = 4.438, *sd* = 1.44) and in VIRTUAL MIRROR ($p < .001$, *mean* = 4.146, *sd* = 1.54) are significantly higher than in BASELINE (*mean* = 3.167, *sd* = 1.37) (Figure 16).

4.2.4 Entry. For EASE, the effect of FIDELITY was significant ($F_{2,115} = 5.01, p < .01$). The rating in VIRTUAL MIRROR ($p < .001$, *mean* = 4.833, *sd* = 1.45) is significantly higher than in BASELINE (*mean* = 4.021, *sd* = 1.86). For PRESENCE, the effect of FIDELITY was significant ($F_{2,115} = 14.51, p < .001$). The rating in REAL MIRROR ($p < .001$, *mean* = 4.458, *sd* = 1.44) and in VIRTUAL MIRROR ($p < .001$, *mean* = 4.521, *sd* = 1.43) are significantly higher than in BASELINE (*mean* = 3.271, *sd* = 1.61) (Figure 17).

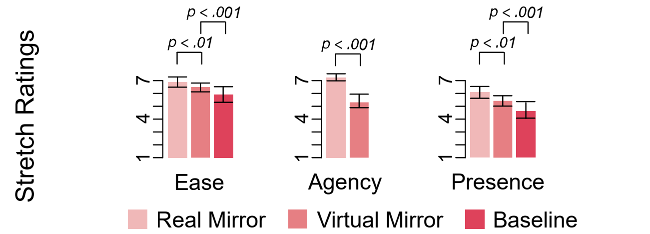

Figure 14: Subjective ratings for STRETCH over FIDELITY.

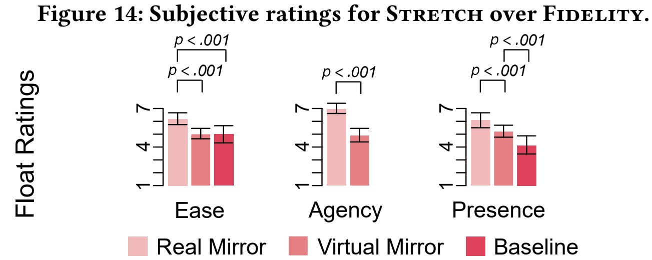

Figure 15: Subjective ratings for FLOAT over FIDELITY.

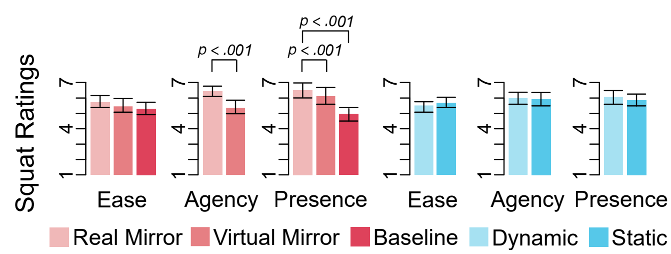

Figure 16: Subjective ratings for SQUAT.

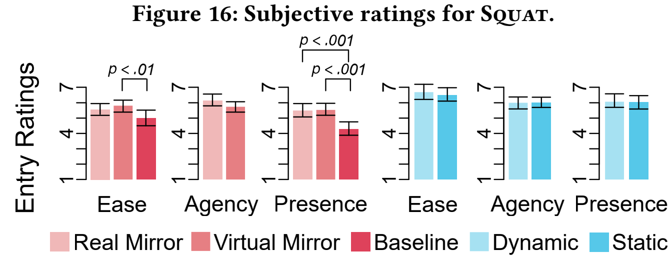

Figure 17: Subjective ratings for ENTRY.

5 DISCUSSION

In this section, we discuss the results in detail with reference to previous works. Overall, our hypothesis H1 was fully supported only by the results in the movement task FLOAT, which features enough complexity but not involving head movement. H2 was supported by the results from all the tasks but ENTRY, because of the large range of head movement along the depth axis in front of the mirror, which disturbed the MR experience. In most cases, H3 and H4 were not supported by our results, mainly because the effect of PERSPECTIVE was not significant.

5.1 Performance Results

5.1.1 Stretch. The STRETCH movement is the most simple of the four tasks. It only involves the arm stretching straight out at a lower and a higher positions (Figure 6 (a)). Hence, it only involves the largest change in the Z axis rotation of the shoulder, while the Y axis rotations were more subtle and less noticeable. We argue that the limited complexity of the movement, especially along the depth axis, may be the reason why we did not find any significant result.

5.1.2 Float. For the FLOAT task, the performance in REAL MIRROR and VIRTUAL MIRROR were better than BASELINE in most measures, while REAL MIRROR yielded significantly better performance than VIRTUAL MIRROR in ELBOW OFFSET. We argue that because this is a more complex movement involving 3 degrees-of-freedom (DoF) rotations of the shoulder and the rotation of the elbow, the benefit of seeing concurrent visual feedback of the real reflection was amplified, hence the better performance over BASELINE in which the participants were not able to see themselves. The increased complexity of the movement was echoed by several participants during the interviews, that the movement was easier for them to follow while seeing their own bodies, due to the increased complexity.

For ELBOW OFFSET, we interpret the significant result as being likely caused by the richer depth visual information in the mirror compared with an ordinary video screen [10]. The majority of changes in the elbow rotations in FLOAT comes from the extension and flexion of the forearm as it reaches forwards and backwards. Because the mirror reflection has superior depth information compared with the virtual reflection due to the fidelity and the stereoscopic image, we argue that the participants were able to correct themselves more easily to match the reaching motions of the instructors. This interpretation is supported by comments such as *"Sometimes (with virtual mirror) it is confusing to see rotating direction front-and-back (P4)."*, *"In the virtual mirror, I couldn't see the whole surface of my arm, which means the sides not directly facing the camera (P6)."*, and *"With real mirror, I had feedback for the depth aspect, but not with the other two conditions (P14)."* Another notable reason given by the participants is that their virtual reflections tend to occlude the instructor more easily, due to the more opaque image compared with the translucent mirror display: *"The occlusion with the joint leads to the fact that I can see how it went backwards in the real mirror, but not in the virtual mirror (P3)."*; *"The virtual reflection has covered the instruction sometimes that I couldn't see it, and my body distracted me. With the real mirror, people are used to it so they do not need any extra energy to focus on the reflected body (P7)."*

Previous works have identified visualising the depth of movement as a pervasive challenge in movement guidance interfaces, and attempted different strategies to address it [1, 6, 45, 57]. *YouMove* and *Physio@Home* added additional video views from different angles to compensate the lack of depth visualisation in those interfaces. Whereas featuring a real mirror, the skeletal overlay in *YouMove* was not able to convey depth information, and the added side view of the same figure was still not intuitive enough to interpret depth movement [1]. Identifying the persisting problem of depth visualisation despite using multiple camera views, the authors of *Physio@Home* proposed an idea of future works featuring real mirror and depth camera [45]. Our MR mirror addresses this challenge and provides a realistic humanoid instructor, instead of using a skeletal overlay. Yu et al. attempted to address the same challenge using a 1PP view inside of an immersive virtual environment. However, their approach is less practical in physical settings outside of immersive environments. In addition, it faced the problem of providing feedback for multiple limbs concurrently, which cannot be solved by head rotation [57]. By using a real mirror, users are still able to benefit from 1PP with their reflection for easier corrective visual feedback [57], while the reflection provides rich depth visualisation. The combination of the humanoid figure and the real mirror, along with the appropriate settings of rendering transparency, enable effective movement depth visualisation on both the instructor avatar and the learner's reflection. We argue that the benefit of this feature in the REAL MIRROR is reflected in the performance of the FLOAT movement.

5.1.3 Squat. For SQUAT, we found that VIRTUAL MIRROR yielded significantly worse performance than REAL MIRROR and BASELINE in ELBOW OFFSET, which measures the arm-bending motion during the squat. On the other hand, VIRTUAL MIRROR yielded the best performance in SHOULDER OFFSET Z, which represents the arm-raising motion during the squat. We infer from these results that the different performances in VIRTUAL MIRROR is induced by its inferior depth visual fidelity compared with REAL MIRROR. As the instructor's arm bends during the squat, the forearm and the hand also extends forward, as (Figure 6 (c)) shows. As a result, participants could recognise this motion and match the instructor more easily with their real reflection than with their virtual reflection. Whereas for SHOULDER OFFSET Z, the better performance could be induced by the clearer visual boundary of the arm than in REAL MIRROR, where the moving background may have induced more cognitive load for participants to match their reflections with the instructor.

There were no significant results for PERSPECTIVE. The feedback from the participants indicates that the benefit of the dynamic perspective change may be diminished by the latency during vertical body movement and the rendered perspective change. Whereas most participants commented that they noticed some latency between the capture and the rendering of their virtual reflection during STRETCH and FLOAT, it was not directly reflected in their performance of those movements. However, the latency appeared to be a larger problem in SQUAT, because the change in the viewing perspective associated with the vertical body movement made the latency more noticeable. Five participants specifically pointed out that the vertical squatting motion made the latency a larger impact on the difficulty of the task: *"Because it involves both arm rotation and body movement, so the delay is more challenging in the virtual* *reflection sometimes because of perspective change in sudden movements. It did not happen much for the previous tasks. (P1).*", *"Because the virtual reflection has a lag, performing the arm movement and the squat at the same time makes the lag more difficult to deal with (P12).*" Whereas the participants could only recognise the latency in Virtual mirror, the latency in the change of viewing angles in Real Mirror would have the same detrimental impact on the performance, but less noticeable comparing with in Virtual mirror where they directly experience it through their virtual reflections.

Previous works in VR have shown that motor performance and simultaneity perception are affected by latency above 75 ms, and that latency above 150 ms is noticeable for video game players [24, 54]. While the latency in Virtual Mirror ($\approx$ 165 ms) is above those marks, its effect is more clearly reflected in Squat than in the other continuous movement, which is Float. We argue that it is due to the perspective update associated with the vertical head movement in Squat, as reflected by the results regarding Perspective.

5.1.4 Entry. The most plausible reason for no significant difference in the performance across Fidelity is that the large range of the walking motion made it more challenging to distinguish subtle differences in the accuracy of the final arm posture (Figure 6 (d)). This interpretation is supported by participants' comments: *"This is keeping moving forward and backward, really hard to catch all the movements compared to the other tasks, because you need to think about the distance when you move. (P2)." "The virtual reflection lag in this even more complicated movement makes it harder (P12)." "The virtual reflection was a bit confusing to follow, as I find the instructor became smaller when it walks back, and it became the same with my body when it walked forward (P23).*" Additionally, the horizontal body movement during walking could make the body move out of the view, while the participants walk straight in front of the display: *"I couldn't see the left side of my body so I didn't know when to stop (P16).*" Previous work has found that the performance of arm movements for target acquisition deteriorates due to stepping movements in straight-walking [59]. We argue that the walking motion in Entry is another source of noise in the performance.

For Perspective, we found that while Dynamic yielded significantly better performance in Shoulder Offset X, the performance was also observably better than Static in all measures for this task. We interpret this result in light of the difference between the two movements, Squat and Entry. Whereas Squat is a continuous movement measured the whole time as the vertical body movement occurred with perspective change, Entry only captured the performance of matching the final posture after the walking motions and the perspective change. From this, we argue that the Dynamic perspective change following head movement is more helpful for performing static posture matching after body movements, than for continuous movement matching during perspective change following concurrent head movement, possibly due to the delay of the video (point cloud) camera between capturing and rendering. P21 gave their comment as *"This movement was easier compared to the squat, because there was a time when I could pause.*"

5.2 Subjective Ratings

5.2.1 Task Difficulty. For the subjective rating of Ease, we found that Real Mirror was consistently rated as being easier than the other Fidelity conditions in the movements Stretch and Float, whereas in Squat and Entry, there was no significant difference between Real Mirror and the other conditions. This is inline with the performance results in Float, where complex arm movements in a static standing pose benefited more from seeing the real mirror reflection. Even though Stretch did not yield significantly better results in Real Mirror, the easy movement may have yielded minimal difference in the performance, which may not be noticed by the participants, as they still rated Real Mirror to be easier. For Squat and Entry, it was more challenging for the participants to subjectively realise the difference in task difficulty between the Fidelity conditions where a larger movement range and a perspective change were involved. This is likely to also be the reason why we did not find any significant difference between any of the subjective ratings for Perspective in these two tasks.

5.2.2 Agency and Presence. For the questions regarding the sense of agency and presence, Real Mirror was rated to be significantly better than Virtual Mirror and Baseline in all movements but Entry. For Agency, Real Mirror received higher ratings as expected due to the instant reaction of the reflected movement and its visual fidelity, which are two traits that directly correlates with the sense of embodiment [7, 26, 42]. One potential reason for that Real Mirror was not rated higher than Virtual Mirror in Entry is that the walking movements induces too much noise in the participants' perception of agency over their real mirror reflection [59]. This is an indication for future works that too much movement involved in the MR mirror interface may induce a lower sense of embodiment even over the mirror reflection of the user, hence negatively impacting their experience in general.

Similarly for Presence, participants rated Real Mirror as not significantly better than Virtual Mirror only in Entry. For the other movements, the sense of presence in Real Mirror is inline with previous findings that the visual fidelity of the user body representation modulates the sense of presence [41], and that realistic sensorimotor mapping between the virtual and the real body induce sense of presence in virtual environments [44]. Apart from the walking motion, the diminished sense of presence in Entry may also be explained as a perceptual mismatch between the 2D instructor avatar and the 3D mirror reflection. Some participants commented that in Entry: *"For the mirror, maybe because of the extra depth information, I sometimes don't know where I am relative to instructor (P11)." "With the virtual reflection, I feel like I was another instructor (P20).*" Indeed, the mismatch between the 2D instructor image and the real mirror reflection may have been amplified when the participant walked and alternated their visual focus between their reflections and the instructor, while the perceptual distance between them changes continuously. Future works should be aware of this issue with locomotion towards the mirror that induce changes in the distance from the mirror to the eyes of the user.

6 DESIGN GUIDELINES

With the findings in our study and previous work on this topic, we offer the following design guidelines for future movement guidance interfaces and MR experience featuring real mirrors in general.

We have found that the mirror has the advantage of providing richer visual information for the depth of movement. **GL1: Future movement guidance interfaces should consider using a real mirror for better depth visualisation of the learner's self-movement, especially for the movements that involve important depth motions.**

We also found that the Real Mirror yielded better performance and was preferred over Virtual Mirror for Stretch and Float, but not for Squat and Entry due to the increasing movement range of the lower body. **GL2: Future MR mirror interfaces should be aware of the potential inferior acquisition outcome when they design for movements that involve large range of head motion, especially when the target movement demands continuous attention from the learners.**

From the performance in Entry, we found that whereas the head movement towards the MR mirror did not yield worse performance than in the continuous Squat movement, the participants' subjective ratings for Agency and Presence reflected a worse embodied experience over their reflections associated with the head movement. **GL3: Future MR mirror interfaces should take precaution with featuring large scale head movement toward the mirror display because it may diminish the embodied MR experience, whereas the perception of static scenes in between head movements may not be affected.**

7 LIMITATIONS AND FUTURE WORK

Due to the fixed size of our prototype, we were only able to evaluate the performance of movement acquisition by measuring and providing visual feedback for one arm. We chose the arm rotation movements following the approach taken by previous works [45], because arm movements have enough DoF to control the complexity of the movement. In the future, we are seeking to implement our prototype using mirrors and displays of larger sizes, and evaluate the effect of the mirror reflection on full-body movement.

Whereas the accuracy of the body tracking feature provided by the Azure Kinect may not be as high as marker-based professional motion tracking systems such as OptiTrack [7], it suited our study where the joint rotation offsets were large enough to be captured by the Kinect sensor. Additionally, our choice of the Kinect sensors enables us and other researchers to build a cost-effective MR mirror interface using portable devices. Future works may explore alternative approaches to feature more accurate body tracking.

The refresh rate, resolution, and latency of the display in our study cannot represent all types of display technologies and devices on the market. Future works could further evaluate the effect of each of those factors on the MR mirror experience and on movement guidance applications. We employed a humanoid virtual avatar with no prominent personal traits to minimise any potential effect of cognitive or social biases. However, there may be more potentially positive effects than biases in using realistic human instructor avatars such as RGB point cloud recordings. Future works could explore the effect of human interaction on movement guidance using MR mirror. In addition, future works could explore the potential of integrating other features, such as contextual information and adaptive feedback, to provide a richer set of guidance [9, 23].

Finally, we evaluated the performance of movement acquisition but not retention, which is more difficult to measure in a short period of time. Future works could explore potential effect of MR mirror on movement retention and other relevant measures for motor learning. Future works on utilising MR mirrors for movement training and motor learning can explore the possibilities of incorporating it in existing multimodal training systems to provide richer feedback for users' performance data, such as in the context of weight lifting [28, 51, 52]. Apart from motor learning, seeing the overlay of the virtual instructor in the mirror with self-reflection also has the potential for inspiring creative movement-making such as in dance and choreography [58]. Future works could explore the use of different visual styles to represent the virtual instructor or other types of visual instructional cues for those purposes [36].

8 CONCLUSION

In this work, we present the first formal evaluation of an MR mirror movement guidance interface. We designed and developed a prototype MR mirror using a two-way mirror overlaid on a desktop monitor. Participants were able to follow the movement performed by a virtual humanoid instructor on the MR mirror by matching their reflections with the instructor movement. We compared participant movement acquisition performance and their subjective rating on difficulty, agency and presence, while they performed four movement tasks in front of the MR mirror display, an ordinary display with their video point cloud avatar, and a baseline condition without seeing themselves. We also evaluated the effect of the real-time updates of the viewing perspective on the virtual instructor according to the change in participants' viewing angle of the mirror reflection. Our results indicate that the MR mirror provides better depth information of the movement and better sense of agency and presence, but those advantages are reduced for large range of continuous head movements, especially along the depth axis in front of the mirror. We contribute towards understanding how to effectively use an MR mirror for movement guidance, and provide design guidelines for future works on MR mirror experience and movement guidance interfaces that feature real mirrors.

ACKNOWLEDGMENTS

This research is supported by an Australian Government Research Training Program (RTP) Scholarship.

REFERENCES

[1] Fraser Anderson, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2013. YouMove: Enhancing Movement Training with an Augmented Reality Mirror. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) *(UIST '13)*. Association for Computing Machinery, New York, NY, USA, 311–320. https://doi.org/10.1145/2501988.2502045

[2] Laura Aymerich-Franch, Damien Petit, Gowrishankar Ganesh, and Abderrahmane Kheddar. 2016. The second me: Seeing the real body during humanoid robot embodiment produces an illusion of bi-location. *Consciousness and Cognition* 46 (2016), 99–109. https://doi.org/10.1016/j.concog.2016.09.017

[3] Marco Bertamini, Nausicaa Berselli, Carole Bode, Rebecca Lawson, and Li Ting Wong. 2011. The rubber hand illusion in a mirror. *Consciousness and Cognition* 20, 4 (2011), 1108–1119. https://doi.org/10.1016/j.concog.2011.04.006

[4] Ivana Bianchi, Marco Bertamini, and Ugo Savardi. 2015. Differences between predictions of how a reflection behaves based on the behaviour of an object, and how an object behaves based on the behaviour of its reflection. *Acta Psychologica* 161 (2015), 54–63.

[7] https://optitrack.com/

[5] A. W. Burghoorn, E. R. Dhaeze, J. S. Faber, J. W. H. Wories, L. M. G. Feijs, and A. A. A. Timmermans. 2015. Mirrorcle: Enhanced Visual Feedback to Support Motor Learning in Low Back Pain. In *Proceedings of the 3rd 2015 Workshop on ICTs for Improving Patients Rehabilitation Research Techniques* (Lisbon, Portugal) *(REHAB '15)*. Association for Computing Machinery, New York, NY, USA, 26–29. https://doi.org/10.1145/2838944.2838951

[6] Carlo Camporesi and Marcelo Kallmann. 2016. The Effects of Avatars, Stereo Vision and Display Size on Reaching and Motion Reproduction. *IEEE Transactions on Visualization and Computer Graphics* 22, 5 (2016), 1592–1604. https://doi.org/10.1109/TVCG.2015.2440231

[7] Mark Carey, Laura Crucianelli, Catherine Preston, and Aikaterini Fotopoulou. 2019. The Effect of Visual Capture Towards Subjective Embodiment Within the Full Body Illusion. *Scientific Reports* 9, 1 (Feb. 2019), 2889. https://doi.org/10.1038/s41598-019-39168-4

[8] Philo Tan Chua, R. Crivella, B. Daly, Ning Hu, R. Schaaf, D. Ventura, T. Camill, J. Hodgins, and R. Pausch. 2003. Training for physical tasks in virtual environments: Tai Chi. In *IEEE Virtual Reality, 2003. Proceedings.* 87–94. https://doi.org/10.1109/VR.2003.1191125

[9] Christopher Clarke, Doga Cavdir, Patrick Chiu, Laurent Denoue, and Don Kimber. 2020. *Reactive Video: Adaptive Video Playback Based on User Motion for Supporting Physical Activity.* Association for Computing Machinery, New York, NY, USA, 196–208. https://doi.org/10.1145/3379337.3415591

[10] Fabian Lorenzo Dayrit, Yuta Nakashima, Tomokazu Sato, and Naokazu Yokoya. 2017. Increasing pose comprehension through augmented reality reenactment. *Multimedia Tools and Applications* 76, 1 (2017), 1291–1312.

[11] E. Ebrahimi, L. S. Hartman, A. Robb, C. C. Pagano, and S. V. Babu. 2018. Investigating the Effects of Anthropomorphic Fidelity of Self-Avatars on Near Field Depth Perception in Immersive Virtual Environments. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 1–8. https://doi.org/10.1109/VR.2018.8446539

[12] Maia Garau. 2003. *The impact of avatar fidelity on social interaction in virtual environments.* Ph.D. Dissertation. UCL (University College London).

[13] Stevie Giovanni, Yeun Chul Choi, Jay Huang, Eng Tat Khoo, and KangKang Yin. 2012. Virtual Try-On Using Kinect and HD Camera. In *Motion in Games*, Marcelo Kallmann and Kostas Bekris (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 55–65.

[14] Mar Gonzalez-Franco and Tabitha C. Peck. 2018. Avatar Embodiment. Towards a Standardized Questionnaire. *Frontiers in Robotics and AI* 5 (2018). https://doi.org/10.3389/frobt.2018.00074

[15] Mar González-Franco, Daniel Pérez-Marcos, Bernhard Spanlang, and Mel Slater. 2010. The contribution of real-time mirror reflections of motor actions on virtual body ownership in an immersive virtual environment. In *2010 IEEE Virtual Reality Conference (VR)*. 111–114. https://doi.org/10.1109/VR.2010.5444805

[16] Ping-Hsuan Han, Kuan-Wen Chen, Chen-Hsin Hsieh, Yu-Jie Huang, and Yi-Ping Hung. 2016. AR-Arm: Augmented Visualization for Guiding Arm Movement in the First-Person Perspective. In *Proceedings of the 7th Augmented Human International Conference 2016* (Geneva, Switzerland) *(AH '16)*. Association for Computing Machinery, New York, NY, USA, Article 31, 4 pages. https://doi.org/10.1145/2875194.2875237

[17] Takenori Hara and Masafumi Oda. 2012. Mixed Reality Mirror Display. In *SIGGRAPH Asia 2012 Emerging Technologies* (Singapore, Singapore) *(SA '12)*. Association for Computing Machinery, New York, NY, USA, 1–3. https://doi.org/10.1145/2407707.2407723

[18] Thuong N. Hoang, Martin Reinoso, Frank Vetere, and Egemen Tanin. 2016. Onebody: Remote Posture Guidance System Using First Person View in Virtual Environment. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction* (Gothenburg, Sweden) *(NordiCHI '16)*. Association for Computing Machinery, New York, NY, USA, Article 25, 10 pages. https://doi.org/10.1145/2971485.2971521

[19] Felix Hülsmann, Cornelia Frank, Irene Senna, Marc O. Ernst, Thomas Schack, and Mario Botsch. 2019. Superimposed Skilled Performance in a Virtual Mirror Improves Motor Performance and Cognitive Representation of a Full Body Motor Action. *Frontiers in Robotics and AI* 6 (2019). https://doi.org/10.3389/frobt.2019.00043

[20] Jae Seok Jang, Soo Ho Choi, Gi Sook Jung, and Soon Ki Jung. 2017. Focused augmented mirror based on human visual perception. *The Visual Computer* 33, 5 (May 2017), 625–636. https://doi.org/10.1007/s00371-016-1212-5

[21] Jae Seok Jang, Gi Sook Jung, Tae Hwan Lee, and Soon Ki Jung. 2014. Two-Phase Calibration for a Mirror Metaphor Augmented Reality System. *Proc. IEEE* 102, 2 (2014), 196–203. https://doi.org/10.1109/JPROC.2013.2294253

[22] Paul M. Jenkinson and Catherine Preston. 2015. New reflections on agency and body ownership: The moving rubber hand illusion in the mirror. *Consciousness and Cognition* 33 (2015), 432–442. https://doi.org/10.1016/j.concog.2015.02.020

[23] Mads Møller Jensen, Majken Kirkegaard Rasmussen, and Kaj Grønbæk. 2014. Design Sensitivities for Interactive Sport-Training Games. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (Vancouver, BC, Canada) *(DIS '14)*. Association for Computing Machinery, New York, NY, USA, 685–694. https://doi.org/10.1145/2598510.2598560

[24] Sophie Jörg, Aline Normoyle, and Alla Safonova. 2012. How Responsiveness Affects Players' Perception in Digital Games. In *Proceedings of the ACM Symposium on Applied Perception* (Los Angeles, California) *(SAP '12)*. Association for Computing Machinery, New York, NY, USA, 33–38. https://doi.org/10.1145/2338676.2338683

[25] Julia M. Juliano, Ryan P. Spicer, Athanasios Vourvopoulos, Stephanie Lefebvre, Kay Jann, Tyler Ard, Emiliano Santarnecchi, David M. Krum, and Sook-Lei Liew. 2020. Embodiment Is Related to Better Performance on a Brain–Computer Interface in Immersive Virtual Reality: A Pilot Study. *Sensors* 20, 4 (2020). https://doi.org/10.3390/s20041204

[26] Sungchul Jung, Gerd Bruder, Pamela J. Wisniewski, Christian Sandor, and Charles E. Hughes. 2018. Over My Hand: Using a Personalized Hand in VR to Improve Object Size Estimation, Body Ownership, and Presence. In *Proceedings of the Symposium on Spatial User Interaction* (Berlin, Germany) *(SUI '18)*. Association for Computing Machinery, New York, NY, USA, 60–68. https://doi.org/10.1145/3267782.3267920

[27] Ioannis Kontaris and Paul E Downing. 2011. Reflections on the Hand: The Use of a Mirror Highlights the Contributions of Interpreted and Retinotopic Representations in the Rubber-Hand Illusion. *Perception* 40, 11 (2011), 1320–1334. https://doi.org/10.1068/p6825 PMID: 22416590.

[28] Yousef Kowsar, Masud Moshtaghi, Eduardo Velloso, Lars Kulik, and Christopher Leckie. 2016. Detecting Unseen Anomalies in Weight Training Exercises. In *Proceedings of the 28th Australian Conference on Computer-Human Interaction* (Launceston, Tasmania, Australia) *(OzCHI '16)*. Association for Computing Machinery, New York, NY, USA, 517–526. https://doi.org/10.1145/3010915.3010941

[29] Diego Martinez Plasencia, Florent Berthaut, Abhijit Karnik, and Sriram Subramanian. 2014. Through the Combining Glass. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) *(UIST '14)*. Association for Computing Machinery, New York, NY, USA, 341–350. https://doi.org/10.1145/2642918.2647351

[30] Betty J Mohler, Sarah H Creem-Regehr, William B Thompson, and Heinrich H Bülthoff. 2010. The Effect of Viewing a Self-Avatar on Distance Judgments in an HMD-Based Virtual Environment. *Presence* 19, 3 (2010), 230–242. https://doi.org/10.1162/pres.19.3.230

[31] Luis Molina-Tanco, Carmen García-Berdonés, and Arcadio Reyes-Lecuona. 2017. The Delay Mirror: A Technological Innovation Specific to the Dance Studio. In *Proceedings of the 4th International Conference on Movement Computing* (London, United Kingdom) *(MOCO '17)*. Association for Computing Machinery, New York, NY, USA, Article 9, 6 pages. https://doi.org/10.1145/3077981.3078033

[32] Kia C. Ng, Tillman Weyde, Oliver Larkin, Kerstin Neubarth, Thijs Koerselman, and Bee Ong. 2007. 3d Augmented Mirror: A Multimodal Interface for String Instrument Learning and Teaching with Gesture Support. In *Proceedings of the 9th International Conference on Multimodal Interfaces* (Nagoya, Aichi, Japan) *(ICMI '07)*. Association for Computing Machinery, New York, NY, USA, 339–345. https://doi.org/10.1145/1322192.1322252

[33] Hye Sun Park, Gun A. Lee, Byung-Kuk Seo, and Mark Billinghurst. 2020. User experience design for a smart-mirror-based personalized training system. *Multimedia Tools and Applications* (Nov. 2020). https://doi.org/10.1007/s11042-020-10148-5

[34] Cristina Portalés, Jesús Gimeno, Sergio Casas, Ricardo Olanda, and Francisco Giner Martínez. 2018. Interacting with augmented reality mirrors. In *Virtual and Augmented Reality: Concepts, Methodologies, Tools, and Applications.* IGI Global, 18–46.

[35] Catherine Preston, Benjamin J. Kuper-Smith, and H. Henrik Ehrsson. 2015. Owning the body in the mirror: The effect of visual perspective and mirror view on the full-body illusion. *Scientific Reports* 5, 1 (Dec. 2015), 18345. https://doi.org/10.1038/srep18345 Number: 1 Publisher: Nature Publishing Group.

[36] Katerina El Raheb, George Tsampounaris, Akrivi Katifori, and Yannis Ioannidis. 2018. Choreomorphy: A Whole-Body Interaction Experience for Dance Improvisation and Visual Experimentation. In *Proceedings of the 2018 International Conference on Advanced Visual Interfaces (AVI '18)*. Association for Computing Machinery, New York, NY, USA, Article 27, 9 pages. https://doi.org/10.1145/3206505.3206507

[37] Giuseppe Riva. 2018. The neuroscience of body memory: From the self through the space to the others. *Cortex* 104 (2018), 241–260. https://doi.org/10.1016/j.cortex.2017.07.013 Special Section: The body and cognition: the relation between body representations and higher level cognitive and social processes.

[38] Hideaki Sato, Itaru Kitahara, and Yuichi Ohta. 2009. MR-Mirror: A Complex of Real and Virtual Mirrors. In *Virtual and Mixed Reality*, Randall Shumaker (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 482–491.

[39] Jeff Sauro and Joseph S. Dumas. 2009. Comparison of Three One-Question, Post-Task Usability Questionnaires. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) *(CHI '09)*. Association for Computing Machinery, New York, NY, USA, 1599–1608. https://doi.org/10.1145/1518701.1518946

[40] Christian Schönauer, Kenichiro Fukushi, Alex Olwal, Hannes Kaufmann, and Ramesh Raskar. 2012. Multimodal Motion Guidance: Techniques for Adaptive and Dynamic Feedback. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction* (Santa Monica, California, USA) *(ICMI '12)*. Association for Computing Machinery, New York, NY, USA, 133–140. https://doi.org/10.1145/2388676.2388706

[41] Sofia Seinfeld, Tiare Feuchtner, Antonella Maselli, and Jörg Müller. 2021. User Representations in Human-Computer Interaction. *Human–Computer Interaction* 36, 5-6 (2021), 400–438. https://doi.org/10.1080/07370024.2020.1724790

[42] Sofia Seinfeld and Jörg Müller. 2020. Impact of visuomotor feedback on the embodiment of virtual hands detached from the body. *Scientific Reports* 10, 1 (Dec. 2020), 22427. https://doi.org/10.1038/s41598-020-79255-5

[43] Roland Sigrist, Georg Rauter, Robert Riener, and Peter Wolf. 2013. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review. *Psychonomic Bulletin & Review* 20, 1 (Feb. 2013), 21–53. https://doi.org/10.3758/s13423-012-0333-8

[44] Mel Slater, Daniel Pérez Marcos, Henrik Ehrsson, and Maria Sanchez-Vives. 2009. Inducing illusory ownership of a virtual body. *Frontiers in Neuroscience* 3 (2009). https://doi.org/10.3389/neuro.01.029.2009

[45] Richard Tang, Xing-Dong Yang, Scott Bateman, Joaquim Jorge, and Anthony Tang. 2015. Physio@Home: Exploring Visual Guidance and Feedback Techniques for Physiotherapy Exercises. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 4123–4132. https://doi.org/10.1145/2702123.2702401

[46] Martijn ten Bhömer and Hanxiao Du. 2018. Designing Personalized Movement-Based Representations to Support Yoga. In *Proceedings of the 2018 ACM Conference Companion Publication on Designing Interactive Systems* (Hong Kong, China) *(DIS '18 Companion)*. Association for Computing Machinery, New York, NY, USA, 283–287. https://doi.org/10.1145/3197391.3205450

[47] Milka Trajkova and Francesco Cafaro. 2018. Takes Tutu to Ballet: Designing Visual and Verbal Feedback for Augmented Mirrors. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1, Article 38 (mar 2018), 30 pages. https://doi.org/10.1145/3191770

[48] John W Tukey et al. 1977. *Exploratory data analysis*. Vol. 2. Reading, Mass.

[49] Laia Turmo Vidal, Elena Márquez Segura, and Annika Waern. 2018. Movement Correction in Instructed Fitness Training: Design Recommendations and Opportunities. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) *(DIS '18)*. Association for Computing Machinery, New York, NY, USA, 1041–1054. https://doi.org/10.1145/3196709.3196789

[50] S. Ubik and J. Pospíšilík. 2021. Video Camera Latency Analysis and Measurement. *IEEE Transactions on Circuits and Systems for Video Technology* 31, 1 (2021), 140–147. https://doi.org/10.1109/TCSVT.2020.2978057

[51] Eduardo Velloso, Andreas Bulling, and Hans Gellersen. 2013. MotionMA: Motion Modelling and Analysis by Demonstration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 1309–1318. https://doi.org/10.1145/2470654.2466171

[52] Eduardo Velloso, Andreas Bulling, Hans Gellersen, Wallace Ugulino, and Hugo Fuks. 2013. Qualitative Activity Recognition of Weight Lifting Exercises. In *Proceedings of the 4th Augmented Human International Conference* (Stuttgart, Germany) *(AH '13)*. Association for Computing Machinery, New York, NY, USA, 116–123. https://doi.org/10.1145/2459236.2459256

[53] Thomas Waltemate, Felix Hülsmann, Thies Pfeiffer, Stefan Kopp, and Mario Botsch. 2015. Realizing a Low-Latency Virtual Reality Environment for Motor Learning. In *Proceedings of the 21st ACM Symposium on Virtual Reality Software and Technology* (Beijing, China) *(VRST '15)*. Association for Computing Machinery, New York, NY, USA, 139–147. https://doi.org/10.1145/2821592.2821607

[54] Thomas Waltemate, Irene Senna, Felix Hülsmann, Marieke Rohde, Stefan Kopp, Marc Ernst, and Mario Botsch. 2016. The Impact of Latency on Perceptual Judgments and Motor Performance in Closed-Loop Interaction in Virtual Reality. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology* (Munich, Germany) *(VRST '16)*. Association for Computing Machinery, New York, NY, USA, 27–35. https://doi.org/10.1145/2993369.2993381

[55] Colin Ware, Kevin Arthur, and Kellogg S. Booth. 1993. Fish Tank Virtual Reality. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (Amsterdam, The Netherlands) *(CHI '93)*. Association for Computing Machinery, New York, NY, USA, 37–42. https://doi.org/10.1145/169059.169066

[56] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. Association for Computing Machinery, New York, NY, USA, 143–146. https://doi.org/10.1145/1978942.1978963

[57] Xingyao Yu, Katrin Angerbauer, Peter Mohr, Denis Kalkofen, and Michael Sedlmair. 2020. Perspective Matters: Design Implications for Motion Guidance in Mixed Reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 577–587. https://doi.org/10.1109/ISMAR50242.2020.00085

[58] Qiushi Zhou, Cheng Cheng Chua, Jarrod Knibbe, Jorge Goncalves, and Eduardo Velloso. 2021. Dance and Choreography in HCI: A Two-Decade Retrospective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 262, 14 pages. https://doi.org/10.1145/3411764.3445804

[59] Qiushi Zhou, Difeng Yu, Martin N Reinoso, Joshua Newn, Jorge Goncalves, and Eduardo Velloso. 2020. Eyes-free Target Acquisition During Walking in Immersive Mixed Reality. *IEEE Transactions on Visualization and Computer Graphics* 26, 12 (2020), 3423–3433. https://doi.org/10.1109/TVCG.2020.3023570