HaptiCoach: Remote Haptic Feedback for Skill Training within Virtual Reality

Jude Afana
Mixed Reality Lab
University of Nottingham
Nottingham, United Kingdom
jude.afana1@nottingham.ac.uk

Joe Marshall
Mixed Reality Lab
University of Nottingham
Nottingham, United Kingdom
joe.marshall@nottingham.ac.uk

Paul Tennent
Mixed Reality Lab
University of Nottingham
Nottingham, United Kingdom
paul.tennent@nottingham.ac.uk

Remote practice of physical skills is now common, but mainstream videoconferencing tools struggle to convey the embodied, touch-based corrections that are routine in face-to-face coaching. In this paper we present HaptiCoach, an exploratory prototype that combines body-tracked virtual reality with a full-body haptic suit, allowing a live human "coach" to deliver remote touch corrections to a student practising alongside a pre-recorded expert instructor. We use this prototype as a research probe to study, qualitatively, how learners perceive and respond to remote touch during physical practice. In a within-subjects study (N=13) using Tai Chi as the practice context, participants interpreted on-body haptic cues as corrective nudges and reported a sense of social presence around the human-triggered virtual coach. Participants also reported that being touched through the suit avoided the awkwardness they would feel if a real person physically adjusted their arm. Our qualitative data is complemented by quantitative measures which we use to provide context for the for participant interview data. We discuss what these patterns might suggest for the design of future remote physical-skill systems.

Remote Communication, Physical Training, Virtual Reality, Haptic feedback, Touch, Remote Learning, Virtual Embodiment

1. INTRODUCTION

Remote learning of physical skills has become common, a shift accelerated by the COVID-19 lockdowns when learners kept training through pre-recorded tutorials and live Zoom or Teams sessions despite appearing in small 2D squares with limited social interaction and no touch (Wang et al. 2024; Zhou et al. 2022; Viverito et al. 2021). That experience showed both the value of remote delivery, in offering flexibility and widening access to classes and teachers, and its limits: a 2D camera view fixes the instructor's perspective, real-time corrective feedback is constrained to spoken language, and the embodied, often tactile, corrections that are common in face-to-face coaching are absent altogether (Iwaanaguchi et al. 2015; Liu et al. 2020; Serhan 2020).

In this paper, we ask a narrower question: when corrective touch is restored to a remote physical practice session, how do learners perceive and respond to it? We approach this in an exploration with a working prototype called HaptiCoach that combines body-tracked virtual reality with a full-body haptic suit so that a live, remote human can reach out and touch a learner's body during practice. We deliberately keep a human in the loop triggering the touch to study the experience of remote touch as a learner-facing phenomenon.

We chose Tai Chi as the practice context because it needs minimal equipment, has a slow pace that makes it manageable to observe, and a local expert was available. We note that the in-place bodily movements of Tai-Chi, whilst not representative of all aspects of physical activities, do have a lot in common with drills used in training for many activities (such as dry-land drills in swimming, some in place drills in football training etc.) (Crowley 2020; Diez-Vega et al. 2014; Rodríguez González et al. 2022), so this work has the potential to support people training in a wide range of sports.

1.1. Contribution

This work contributes an empirical, qualitative account of how learners experience human-mediated remote touch during physical practice, supported by a working prototype that integrates body-tracked VR with full-body haptics. The novelty is in the study of the experience, not in the underlying haptic hardware. We use our prototype as a research probe

© Afana et al. Published by BCS Learning and Development Ltd.

to surface design questions about social presence, attentional capture, learner comfort with corrective touch, and the practicality of such setups, which future work on more scalable, partly or fully automated systems will need to answer.

2. RELATED WORK

2.1. Remote practice of physical skills

Remote practice of physical activities is shaped by two competing forces. On one hand, video-based delivery extends access: a learner can train with an instructor located anywhere, replay tutorials at any time, and avoid the friction of travel and scheduling. Pre-recorded video tutorials offer especially flexible learning, as they are accessible anytime and anywhere, adding more flexibility to the learning process (Parker et al. 2021; Jo et al. 2023). On the other hand, video-mediated practice loses much of what makes co-located coaching effective. Feedback is restricted to a fixed camera angle and spoken language, learners receive no corrective feedback on their own body during practice, and motivation drops compared to in-person sessions (Iwaanaguchi et al. 2015; Liu et al. 2020; Meyer et al. 2021; Wang et al. 2024; Serhan 2020). Trainers have adapted to these constraints (Koerner and Staller 2020), but the gap remains an acknowledged limitation of remote delivery.

2.2. Mixed Reality for skill training

Studies have found that using mixed reality in training has promising results in enhancing performance, especially the use of augmented reality, while virtual reality was considered almost as good as traditional training in industrial and safety contexts (Daling and Schlittmeier 2024; Lawson et al. 2020). For example, work on maintenance training has shown that AR outperformed both traditional and VR training on multi-level tasks due to lower cognitive load (Liu et al. 2022). However, sports training is inherently a challenging application as it requires a high level of freedom of movement and accurate posture evaluation, (Tian et al. 2024) which makes transferring such results to sports difficult. Early studies suggest that VR sports training can improve performance and training efficiency (Tsai et al. 2022; Takala et al. 2020); for example, one study used a multi-sensor system to track full-body movement in VR martial arts (Takala et al. 2020), and another used a Kinect sensor in table tennis training with written feedback for the player to fix their posture (Oagaz et al. 2022).

A common limitation across this body of work is that current MR and VR sports-training systems rarely integrate *adaptive real-time feedback*. Analyses of user movement, posture or performance are typically delivered through visual and verbal cues, and there is little research on adding adaptive real-time feedback in another modality such as touch.

2.3. Digital touch as a missing modality

Touch is the basic form of human interaction with objects and people, and the first sense used to understand the surrounding environment (Field 2010). Jewitt et al. (Jewitt et al. 2020) introduce the term “digital touch” to refer to digital technologies which mediate touch beyond the direct contact of hands or bodies. Tactile systems have been suggested as a way to reduce the social and emotional gap that occurs in remote settings (Heatley Tejada et al. 2020; BBC Radio 4 2020).

For physical skills in particular, touch is not optional. Traditional skills training very often involves physical touch between trainers and students to provide guidance and corrections for movements, such as adjusting a posture, raising an elbow, or steadying a stance (Assandri 2019). Thus touch is a natural candidate for corrective feedback in remote practice. Interestingly virtual touch may also have an advantage over real world touch, in that mediated touch may actually be more socially comfortable than in-person touch for some learners; Direct human touch between people whose relationship is professional rather than personal, or between opposite genders, can feel awkward in face-to-face settings, and it has been argued that a virtual modality bypasses social risk while preserving pedagogical signals(Heslin et al. 1983; Afana et al. 2023; Camps et al. 2013; Oakley et al. 2001).

Several wearable systems already make such mediated touch possible. Ring*U (Choi et al. 2014), for instance, is a wearable ring that sends low-resolution haptic “mini hugs” between intimate partners, illustrating that mediated touch can carry interpersonal meaning even at low physical resolution. With ever improving haptic interaction technology, complex haptic sensations can be digitally controlled, allowing users to feel sensations such as pressure, texture, and movement through wearable devices such as SenseGlove[1], TactSuit Pro[2] and Teslasuit[3]. Whilst digital haptic sensations may not be identical to real-world touch (Martínez et al. 2016), they offer enough resolution to convey location or direction, such as might be used by a remote coach to indicate where on the learner’s body something needs to change.

[1] https://www.senseglove.com/
[2] https://www.bhaptics.com/tactsuit/tactsuit-pro/
[3] https://teslasuit.io

Taken together, these three strands point to a gap. Remote practice loses the corrective channel, mixed reality has not restored it, and digital touch could now supply it. What is unknown is how learners read such a cue when it is delivered.

3. THE HAPTICOACH SYSTEM

HaptiCoach is a multi-user, non-co-located VR practice prototype built around three roles: a live **student**, a pre-recorded **instructor**, and a live **coach**. The student and coach share a virtual environment in which the pre-recorded instructor performs the practice routine. The coach can see the student's full-body avatar and can reach out with the VR controllers to "touch" that avatar, which causes the corresponding region of the student's haptic suit to fire. We use the system as a research probe to investigate the learner's experience of remote touch.

3.1. Design rationale

Two design decisions follow from this framing:

First, haptics are triggered by a human rather than by an automated posture analyser: the coach watches the student and chooses each correction, making this a Wizard-of-Oz implementation. Keeping a human in the loop lets us study the modality of remote touch directly, without being confounded by behaviour of an automated detector. It also preserves what we take to be the point of the arrangement, which is that two people who are not physically together interact through touch in real time. The coach role is deliberately an amateur, trained in the task by the professional instructor, rather than an expert. This is less unusual than it may sound, because tiered expertise is already normal in physical training. Franchised formats such as Les Mills send centrally designed choreography to certified instructors worldwide (Parviainen 2011; Andreasson and Johansson 2016), and a children's beginners' karate class is usually taught by a junior instructor rather than a senior black belt. In both cases the instructor delivers the expert's material without being the expert, as our coach does. This mirrors plausible future arrangements such as teaching assistants in a hybrid class, or automated systems supporting pre-recorded lessons and reflects the fact that the focus is on learners' experience of remote touch, not the coach's level of expertise.

Secondly, the expert instructor is a pre-recorded animated character rather than as a live person. This separates the asynchronous, broadcast-style content (the expert demonstration) from the synchronous, individualised content (the corrective touch). We adopted this arrangement on the expectation that it could allow high-quality training to scale where expert time is scarce.

3.2. Roles

There are three roles in our exploratory study, which are:-

1. **The student**: wears the haptic suit and motion-capture markers, and follows the instructor's movements. They receive haptic corrections on the body part the coach touches.
2. **The instructor**: is a pre-recorded humanoid avatar, animated from motion capture of a real Tai Chi instructor. The instructor is not a live person in any session.
3. **The coach**: is a live remote human, sharing the same virtual scene as the student. The coach is rendered as a floating green head with spherical hands. This was intentionally low-fidelity, so as to be visually distinguishable from the human-form instructor avatar. The role was filled by the lead author, who was trained in the routines by the professional instructor.

3.3. Apparatus and spatial organisation

The student stands inside a 16-camera Vicon motion-capture volume which provides full-body tracking, while the coach is physically not co-located with the student and works from a separate room (Figure 1). Both meet in a shared virtual scene (Figure 2) so they can see each other's avatars and the instructor's character, and can virtually interact in real time. The apparatus is summarised below:

- **Motion capture:** a 16-camera Vicon volume providing full-body tracking of the student.
- **VR headsets:** Meta Quest 3 worn by both the student and the coach, joining a shared Unity scene.
- **Haptic suit:** a Teslasuit jacket worn by the student, delivering electro-muscle-stimulation (EMS) cues across the upper body, calibrated per participant for comfort.

  The Teslasuit's sizing is restrictive (too loose and the EMS does not couple to the skin; too tight and it does not fit). We had two suits, one medium woman and one medium man, used the jacket only with the Vicon mocap trousers worn underneath, and attached reflective markers to the jacket for tracking (Figure 3).
- **Touch trigger:** because the student's avatar is driven by live motion capture, the coach can touch any tracked region with the VR

controllers, firing the corresponding EMS region on the student's body.

- **Software:**
  - Vicon Shogun Live v1.11 (motion capture)
  - Teslasuit firmware v4.5.6 with Control Center v2.5.3 (haptics)
  - Unity 2022.3.33 with the Teslasuit plugin v2.5.1, ViconUnityPlugin v1.3.1, and Meta XR Core SDK (integration)
  - Open Sound Control (OSC) for cross-instance positional data

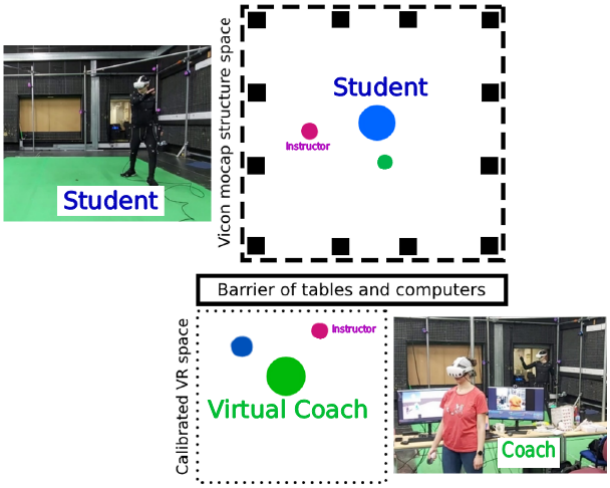

Figure 1: *The student wears the Teslasuit and a VR headset inside the motion-capture volume. The coach is in a separate physical room. Both share the same virtual scene shown in Figure 2.*

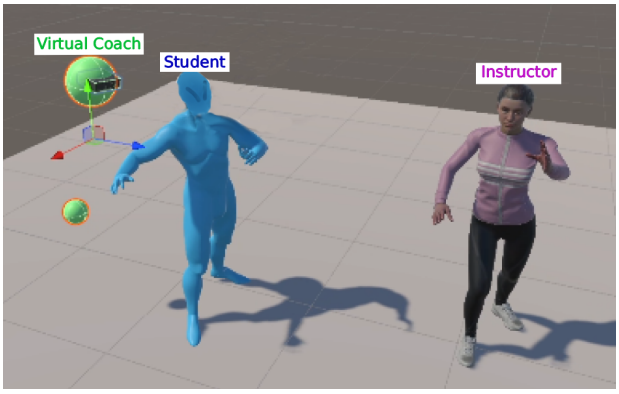

Figure 2: *The shared virtual scene. Left to right: the virtual coach (the green floating head with spherical hands), the student's motion-captured avatar, and the pre-recorded instructor.*

4. USER STUDY

4.1. Research questions and framing

In this exploratory study, our primary interest was qualitative: how do learners perceive, interpret and



Figure 3: *Full gear: Meta Quest headset, Teslasuit jacket with reflective Vicon markers, and Vicon mocap trousers.*

respond to corrective haptic touch delivered by a remote human during a physical practice session? In particular, would learners recognise the coach's intent from a touch (for example, would we be able to touch somebody's back to fix a posture, or move their elbow to a higher angle)? And how would they describe the experience of being touched at a distance? To give users context, we also presented them with sessions with purely pre-recorded 3D capture. We also collected performance scores as a secondary, indicative measure. These scores came from the coach, who was the only person watching the movement as it happened, and sit alongside participants' own descriptions of how the session felt. The understanding of how users receive and experience formative haptic feedback could also be useful in future work that aims to develop AI-driven "virtual coaches" to automate the creation of feedback.

We deliberately excluded verbal feedback from the coach during the session, in order to focus purely on the haptic modality. Pairing every touch with speech would have made it impossible to tell what the touch contributed, and participants were told in advance

that feedback would come only when something needed correcting. We are aware verbal feedback is a necessary part of the feedback process, and a deployable system would clearly combine the two.

4.2. Study design

We used a within-subjects design to compare the effect of adding live haptic feedback to virtual Tai Chi movement training. In a VR scene, the participant would see the instructor's avatar doing pre-recorded movements and they would try to copy it. In a subset of those exercises, there would also be a virtual coach visually present to provide live haptic feedback when necessary. There were two conditions:

- Pre-recorded avatar only
- Pre-recorded avatar in addition to the live virtual coach providing haptic feedback

Each participant learnt 4 poses, two of which used haptic feedback, and two using only pre-recorded visuals. For each pose, participants trained twice to study the learning curve. Ordering of poses was balanced so that 6 participants began with a pose using haptic feedback, and 7 participants began with the pre-recorded-only condition. Ethical approval for the study process was granted by the ethics board of [Anonymised for review].

4.3. Participants

Participants were recruited by convenience sampling. We recruited 13 participants by word-of-mouth and posters. Due to the size of the available Teslasuits, which need to fit their subjects quite specifically, we had to restrict participation only to people who would fit a medium (female) or a medium (male) jacket. Demographics of participants included 7 females and 6 males, all under 30 years old. Only one participant had any previous experience with Tai Chi, from a single school PE lesson.

4.4. Procedure

Once participants arrived at the venue, we explained the study including the consent form and information sheet, then helped them put on the top of the Teslasuit with the Vicon mocap trousers, applied markers for Vicon detection and avatar calibration, calibrated the Teslasuit feedback level to the participant's individual comfort, and checked synchronisation with the Unity Game Engine. Participants were told that there would be 4 poses each repeated twice, that two of the poses would have a "virtual coach" present to provide haptic feedback which would only be triggered if a move or a posture was not correct or needed fixing, and that they were free to move around the expert instructor character and stand wherever felt best for them in the scene. They received printed copies of the information sheet and privacy notice to review, and signed consent forms only after being encouraged to ask questions. As noted above, no Tai Chi training was conducted beforehand and no verbal feedback was given during the session. After the experience, they would change their clothes, we would disinfect the suit, then start the interview. Figure 4 shows the flow through the study.

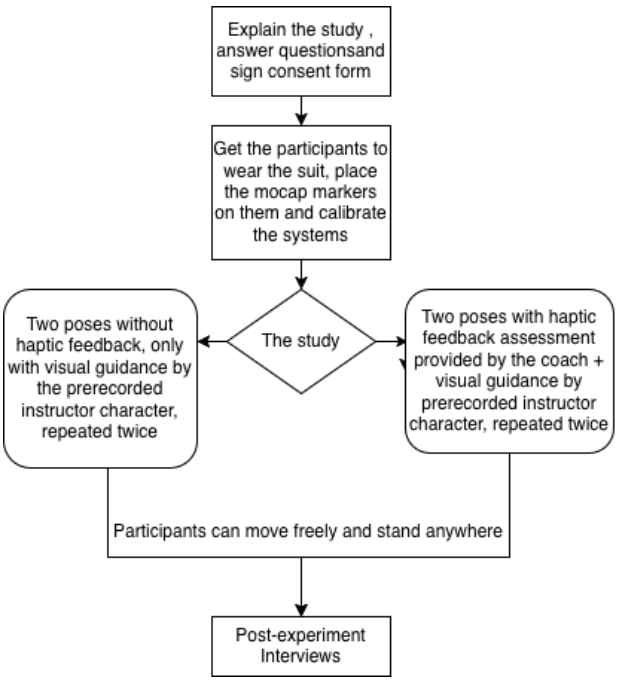

Figure 4: *A flow diagram explaining the procedure of the study.*

4.5. Data collection and analysis

We collected both quantitative and qualitative data during the study. The quantitative data was gathered from the coach's evaluation of the participants' success at performing the various poses, and the qualitative data from interviews after each participant completed all the poses.

Qualitative. After finishing the activity, participants were individually interviewed to gather their perspective throughout the lesson, their thoughts, feelings and experience. Semi-structured interviews (see supplementary material) lasting approximately 30 minutes were performed immediately after the study. We transcribed the interviews and conducted an inductive thematic analysis, following the approach outlined by Braun and Clarke (2006, 2019). We first reviewed the data broadly to achieve familiarisation, then generated initial codes across the whole dataset in NVivo 14, allowing extracts to carry more

than one code. We organised these codes into preliminary categories, then discussed them iteratively and constructed the overarching themes from them.

Quantitative. Performances of the participants were assessed using the assessment criteria provided by the Tai Chi instructor. The criteria sheet (see Supplementary material) gave scores ranging from 1 (lowest) to 5 (best) for body posture, hand movement, the coordination and interconnection of different body parts, and the flow and flexibility of the movement for each part. We scored the recordings ourselves, using the criteria sheet the instructor provided. The reason we used such a manual criteria was our interest in looking at the process from the perspective of the professional human coach. This allowed us to get an observation and overview of how well participants were performing the movements, providing a more ecologically valid model than using an automated system that would only output abstract numbers (Preatoni et al. 2013).

Participants were required to imitate and follow the movement they saw from the 3D model in the scene, and their moves were video recorded to be assessed in comparison to the actual moves and poses. For each participant, we considered the average improvement in performance score between first and second repetitions of poses with and without haptic feedback. We chose to look at within-participant improvement rather than absolute score because the goal of the analysis was not to rank participants against each other but to understand differences in learning rate between conditions over this very short practice arc. We applied repeated-measures ANOVA followed by post-hoc comparisons using Tukey correction for multiple testing, and supplemented this with a Wilcoxon signed-rank test for the repetition effect. Jamovi v2.3.28 was used to apply the statistical tests.

5. FINDINGS

We present the qualitative themes first, as they are the main contribution of this exploratory study, followed by quantitative observations that complement these qualitative reports.

5.1. Qualitative findings

We developed three themes from the interview data: (1) what made the coach socially present, (2) what the touch communicated and how learners responded to it, and (3) the relationship between this experience and current remote-training modalities, including its practicality.

5.1.1. Theme 1: Believed humanity, not visual fidelity, made the coach socially present

Participants were aware that the green floating head represented a real person, and this awareness shaped how they practised. The simplest reading is that low visual fidelity did not prevent, and may even have helped, a social reading of the avatar. P2 described the coach as “a slight presence which was there correcting me, so it was good I didn’t feel like there was somebody actually there and judging me.” The same participant nonetheless reported feeling accountable: “somebody’s actually there and keeping a track of what I’m doing so I have to do better.” P9 echoed this: “the same way you wouldn’t want to disappoint a physical instructor.”

These two readings, the coach as both comfortably abstract and also socially present enough to hold the learner accountable, co-exist rather than contradict each other. Participants seem to derive accountability from the knowledge that the avatar is animated by a real person, not from its visual fidelity. P13 stated this directly: “I literally felt that I am standing in front of a person and trying to match the movements, so the human presence was there.” Conversely, P11, who reported the experience as “a single player game”, is the exception that bounds the claim. When learners disengage from the avatar’s humanity, the social effect attenuates. For these participants the accountability attached to who they believed was behind the avatar rather than to how it looked. We consider what this might mean for design in Section 6.1.

5.1.2. Theme 2: The touch directed attention but left learners to work out the correction

Participants’ reports of the haptic feedback itself were notably varied. The most consistent reading is that the touch functioned as an attentional cue or reminder, directing attention to a body part, but that the interpretive load of figuring out what the cue actually meant fell back on visual reference to the instructor avatar.

An element of surprise came first. P8 said: “I was just worried when I’m going to get the feedback on which part of the body. I knew it wasn’t strong at all. But I was not expecting it. So I think it was all out of the blue when it happened.” This was reported as initial, not persistent, suggesting that participants became less surprised and less distracted over time and more comfortable with the feedback, which is consistent with the quantitative pattern of fewer corrections being needed in the second repetition (see below).

Once participants settled into the cue, many described the touch as a useful reminder. P10 said: “It acted as a reminder that I’m not in the right

position for that arm so I would just like immediately look over at the instructor and try to fix it, it makes you pay attention more." P11 reported the same pattern of immediate corrective action: "When I felt the touch from the suit I immediately corrected my movements, I would look at the 3D model of the instructor and I would correct movements; without the haptic feedback I wouldn't notice that I was doing something wrong." P12 framed it as learning support: "with no haptic feedback I was probably making mistakes without realising it and then getting comfortable in those mistakes, but in the haptic feedback ones they made me more aware that I need to be more thorough with the movements."

A limit of the cue then emerged: confusion. The cue indicated where, but not how. P1 explained: "I got confused by the touch because the direction was missing, the communication was missing, that what am I supposed to do when I feel a touch, like I'm feeling something but how am I supposed to respond to it?" P5 articulated the contrast with in-person coaching: "Whenever I was getting the feedback, I was not sure how exactly to move my hand. In a real scenario, someone would be holding my arm and moving me forward or backward, or telling me to move it, not just pointing at it." These reports converge on a design implication: a pure positional haptic nudge under-specifies the corrective intent. Participants partially closed this gap by re-fixating on the instructor avatar (P1: "I have to constantly look at the instructor like one would in real time and real space"), but this offloads work that an in-person trainer would normally absorb verbally or kinaesthetically. We return to this point in the Discussion.

A particular behavioural strategy under this limit was avoidance. P2 said: "After the first few times I realised that I did not want to feel corrected, was slightly confusing, and then I spent my time trying not to get corrected, so I was following it more thoroughly just so that I could do it better." This is dual-edged. On one hand, it is evidence that the haptic modality is successfully attention-capturing, in that the cue "works" as a behavioural signal. On the other hand, it positions the touch as negative reinforcement, which is not the pedagogical mode an instructor would normally choose. Anxiety about anticipated feedback is not an ideal motivator pedagogically, although a similar dynamic exists in face-to-face training when learners are nervous about being touched or singled out by an instructor. Designers of future systems should be aware that the framing of the cue ("correction" versus "guidance") and its calibrated intensity will materially shape whether learners orient toward the touch or away from it.

When asked whether the suit should be capable of physically moving them into the correct pose, most participants resisted. P1 said: "I don't think I would enjoy if something's controlling my movement, I would instead enjoy the nudging but with a bit of guidance with like auditory guidance or directional guidance." A minority were open to it. P5 liked the idea of the suit being strong enough to move the body correctly, but not having full control: "It would have been a lot better to be strong enough to move it correctly, while I still controlling my body, not giving the suit full control." Taken together, learners preferred to retain agency. They wanted a cue, not a manipulation.

5.1.3. Theme 3: Filling a gap that video cannot, at a hardware cost

Participants spontaneously compared the experience to what they would otherwise use, mainly pre-recorded video and live video calls. The recurring point was that current tools cannot offer corrective signal during practice. P8 said: "In a normal teaching session, verbal feedback will be more than a physical touch. But in the haptic feedback, I could see where exactly I was going wrong without communication. In-person lacks the flexibility that I can get through this. But here, I can just put on the headset whenever I want to do it at my own will, giving me more flexibility."

A related observation concerns self-placement within the shared virtual scene. Participants could either mirror or copy the moves of the instructor, and individual preferences varied. P8 preferred to stand in front and follow directly, whereas P4 stepped behind the instructor to avoid the left-right inversion problem: "I moved a bit behind them to get another view. My left would be their left. And I didn't have to think too much." Notably, P4's position would not be possible in a real class context, but in VR this freedom is available and several participants exploited it to manage their own cognitive load. In a VR setting the coach's and the learner's optimal viewpoints could in principle be decoupled.

P13 felt the gap was filled: "The gap is filled in the setup with this VR and the motion sensor thing, and I could literally see the instructor in front of me and I can follow her, compared to what I would be doing in a video call." Several participants described the setup as filling a gap that pre-recorded video and video calls leave open: it felt like being in a real room, they could see the instructor in front of them and follow her movement, and the haptic feedback helped them perform better.

One participant also raised a social-comfort dimension to the modality. P10 said: "the main thing is the social aspect, because then in the VR headset it's

just you and the avatar. I think that's a good addition because I wouldn't be comfortable with someone coming to physically fix my arm, but this way there's no one touching you but you're still getting the feedback." For this participant, the mediated nature of the touch was itself the affordance, removing the awkwardness of an in-person correction while keeping the corrective signal.

At the same time, practicality was the most consistent negative. The suit was tight, the headset added weight, and calibration was time-consuming. P11 said: "I was aware of the suit on my body and it was a bit tight, so usually in a studio you would wear very comfortable clothes and this was pretty tight, so it did make it a little bit harder." P13 said: "The VR headset is fine. The suit, if it can be designed simpler just to wear it quickly, for me it's the time it took to calibrate, if that is shortened I don't mind it." These reports do not contradict the rest of the findings, but they do constrain where a system like this could plausibly live. A studio or gym setting is more realistic than a domestic one, at least with current hardware.

5.2. Quantitative findings

Unsurprisingly, in all conditions, users performed better in the second repetition of an exercise ($p < .001$). We can also see this reflected in the number of corrective haptic touches the coach needed to deliver. Within the haptic condition, a Wilcoxon signed-rank test showed that participants were corrected significantly less often on the second repetition than on the first ($p = 0.010$), with the first repetition being on average 2.00 units higher.

We can see some supportive evidence for this in the comparison between individual performances (Figure 5), which show slightly lower performances in first repetitions of exercises without the haptic feedback than with, and greater improvement between repetitions in the exercises with it, although a larger study would be required to demonstrate these effects statistically.

6. DISCUSSION

The picture that we draw from the qualitative analysis is that human-triggered remote touch is a workable feedback modality whose value is mostly in directing attention rather than in transmitting full corrective content. Below we draw out three implications:

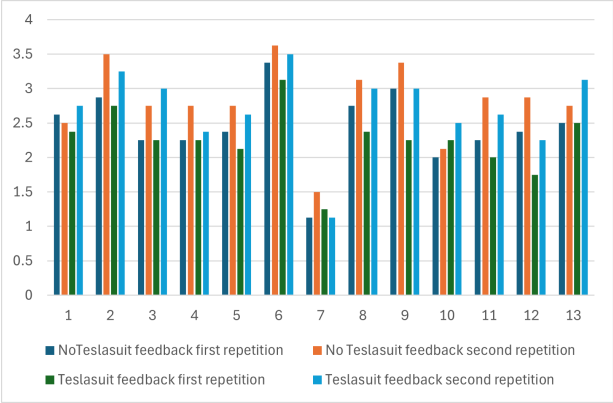

Figure 5: *Per-participant performance scores across condition and repetition.*

6.1. Embodiment of the coach matters less than perceived humanity

The coach was rendered as a low-fidelity floating green head, yet participants reported social presence and accountability, P2, P9 and P13 each describing one or both. That social presence survived such minimal rendering is consistent with prior work finding that non-photorealistic avatars support engagement without uncanny-valley discomfort (Nomura et al. 2025). It also resonates with prior findings on the effect of being observed: trainees who know they are being watched tend to report higher motivation and reach goals more reliably when working with a personal trainer (Lu et al. 2024; McClaran 2003), and although people may feel more anxious about making mistakes when observed, they often perform better in those conditions (Chib et al. 2018). The accountability our participants described (P9: "you wouldn't want to disappoint a physical instructor") suggests that the same effect carries over into a virtual setting, even when the observer is rendered as a green sphere. For designers, this is an opportunity. The high-end work of building photorealistic avatars may not be the limiting factor for socially effective remote coaching. What appears to matter for our participants is the belief that the avatar is animated by a real person; a key question for future work is to consider how this might work with automated feedback, or with a one-to-many coach to student ratio.

6.2. Haptic cues acted as attentional pointers rather than full instructions

Participants treated the touch as a pointer to a body part and then looked at the instructor avatar to figure out what to do. This division of labour, haptics for "where" and visual reference for "what", is workable, and preserves the unique affordance of touch (on-body localisation without disrupting visual attention) while leaving the instruction-rich modality of vision to

the instructor avatar. It is also consistent with prior findings that digital touch is non-trivial to interpret (Martínez et al. 2016). Initial confusion as to the meaning of touch and how to react to it is perhaps partially explained by the lack of experience with the technology; as Martínez et al. (2016) described, longer training time would be needed to allow participants to interpret the touch in a more nuanced manner and to know how to move accordingly.

There is also a learning-transfer hypothesis worth flagging. Participants told us the feedback made them more aware of mistakes they would otherwise have settled into, as P12 put it, and Figure 5 shows greater improvement between repetitions with the feedback. This suggests to us that a learner might wear the suit for initial training and then practise without it, carrying forward what was corrected. This is our reading rather than something participants proposed, but it would suit the gym setting the hardware currently implies. This matches a familiar pattern in physical practice, where drills and aids are used to build a skill that is then performed without them (Diez-Vega et al. 2014; Wen et al. 2016). Our study cannot test this hypothesis as we did not include a transfer condition, but this is a promising area for future skill-training research to explore.

6.3. Practicality is the binding constraint

The most consistent negative feedback was about the hardware: suit fit, calibration time, and the cumulative weight of the headset plus the suit (Laine and Suk 2024; Souchet et al. 2022). None of these are conceptual problems with the modality. They are engineering and consumer-readiness problems with the current generation of equipment. For now they constrain the realistic deployment site of any HaptiCoach-like system to spaces that can absorb the setup cost, such as a studio or gym, rather than a living room. Future, lighter, more forgiving haptic hardware will move this boundary, though participants' accounts suggest they already found the experience worth the setup it currently requires.

7. DESIGN IMPLICATIONS

In this section we discuss a broader design space for haptic-supported virtual training that is ripe for exploration. The implications we draw here are speculative. Our study varied one factor only, the presence or absence of human-triggered corrective touch, so what follows extrapolates from a single point in each dimension below rather than comparing across them. We note that the desired approach to each of these trade-offs may be different depending on the target physical activity, the teaching practices of teachers, and the nature and abilities of the students.

7.1. Live and asynchronous interaction

Training digitally creates new possibilities for asynchronous interaction; however, there is an inherent tension between the extra flexibility and the limits of non-live interaction. Live interactions offer immediacy, spontaneity, and authentic social presence, but they also introduce logistical complexity that limits scalability and flexibility. Asynchronous interactions enhance accessibility, scalability, and convenience by allowing users to engage on their own schedules, but typically sacrifice emotional immediacy and spontaneous interaction. Our system combines the two: an asynchronous instructor with a synchronous coach. In our setup this combination supported both expert demonstration and individualised correction without requiring the expert to be live. Hybrid models, such as an asynchronous instructor who can occasionally be "possessed" by a live expert for brief synchronous moments, could expand this model further.

7.2. Social fidelity and production fidelity

There is often a trade-off between high quality interactivity versus high quality visuals; for example we found it relatively easy to create a high quality 3d representation of the instructor using pre-recorded content. However in order to effectively drive the real-time coach, we used an extremely low quality sphere plus hands avatar for the coach. In our interviews we saw that the coach was experienced as being socially present despite the low visual fidelity. This suggests that visual fidelity is not the only thing which creates a positive training interaction. In this case, the trainee's perception that the simplistic avatar "is" a live human overrides the rudimentary nature of the visual avatar. Designers should not treat photorealism as a prerequisite for social presence. The less-human renderings of the coach even have potential to create advantages in that they may reduce social discomfort around being observed and corrected while preserving accountability. Higher production fidelity also carries the additional risk of triggering uncanny-valley effects in some users (Howard 2017).

7.3. Low-fidelity and high-fidelity touch

Touch in VR spans a wide range, from vibrotactile controllers that buzz the skin without conveying weight (Li et al. 2024), through grounded kinaesthetic rigs that deliver 6-DOF force cues at a single point of perception (Tang et al. 2020), to task-specific props and passive haptics that only work within their intended context (El Saddik 2012; Kim et al. 2023), and lab-only robots whose cost keeps them out of consumer settings (Suzuki et al. 2021). Our findings suggest that for corrective feedback the key requirement may be spatial localisation on the learner's

body (Section 5.1.2), not high force fidelity, which in turn suggests that lighter, simpler wearables (such as distributed vibration motors) may be sufficient for many of the use cases the full Teslasuit is currently used for, or off body approaches such as targeted air blasts (Sodhi et al. 2013).

Touch fidelity is also inherently intertwined with motion capture; in order to accurately give feedback to users, the coach (or automated coaching system) must have knowledge of the physical state of the user's body. Our full-body mocap system meant that the coach could touch the avatar with natural reaching gestures rather than via a menu or interface. This embodied selection of body regions matters, as it preserves the spatial and proprioceptive sense of corrective touch that is central to in-person coaching (Lee-Cultura and Giannakos 2020; Tanaka et al. 2015; Afana et al. 2023). A future, lighter version of such a system could plausibly drop the high-end mocap volume in favour of inside-out body tracking, but we would expect tracking fidelity to bound haptic expressivity, since the coach can only touch what the system can locate.

8. LIMITATIONS AND FUTURE WORK

This was an exploratory, primarily qualitative study; as such the sample size (N=13) is small, and only two repetitions per pose were collected. With this small a participant group we cannot make strong quantitative claims as to the applicability of haptic technology. Larger samples and longer term practice arcs are the natural next step for this work in order to understand how this technology affects performance when used as part of a real training schedule.

A further limitation concerns who assessed performance. We scored the recordings ourselves, and the researcher who did so had run the sessions and acted as the coach; no one running a study can be unaware of it. We partly blinded the scoring: recordings were taken in random order, and although the condition was visible, the researcher scoring them did not know which of a participant's two attempts was the first until all scoring was complete. Ideally an independent rater would have scored the videos, but logistics and practicality made that unworkable.

A further limitation to this work is that the technology used in our study is currently not in a consumer-ready state: the Teslasuit is currently only available at research-grade price points and is complicated to calibrate, the Vicon motion capture system is a fixed location room-scale piece of equipment, and the integration between all elements of the system required multiple PCs and a significant amount of customized interoperability code. Significant maturation of these third-party parts would be needed before such a system could be used outside the lab, particularly in a home setting where most informal physical practice happens (Ainsworth and Hagino 2006; Bårdstu et al. 2020).

Finally, this study focused purely on applying haptic feedback and understanding the user experience of such feedback; in a realistic environment, it would make sense to study a combination of haptic and verbal feedback. This study has demonstrated that remote haptic feedback for training is a feasible option; however for it to become practical more research is needed. We see the following key areas for research.

- Developing a more mature, portable hardware stack, for example by using wireless haptic devices and on-headset or at least on-body body tracking, in order to enable the system to be practically used outside a lab setting.
- Run longer term studies of the system in use in a real training setup to explore the possibilities of using this, and to study the effects on user performance of remote haptic training.
- Explore automatic coaching, or methods to enable semi-automatic coaching or to enable one coach to support many remote users concurrently, in order to allow the technology to scale to meet user demand.

9. CONCLUSION

We explored adding human-triggered digital touch to a remote VR practice session, using a Teslasuit to deliver corrective haptic feedback when learners and instructors are not physically co-present. Our exploratory mixed-methods study suggests three things: the value of remote corrective touch lies less in transmitting full instruction than in directing learner attention; the perceived humanity of the coach matters more than its visual fidelity for our participants; and the binding constraints on deployment are hardware-related, not conceptual.

Beyond the technology itself, what will truly shape whether remote touch becomes a meaningful part of physical practice is the extent to which learners trust the cues they feel and the human intention behind them. As immersive training systems continue to evolve, this sense of human connection and social presence along with trust may become just as important as the feedback itself in supporting effective and engaging remote learning experiences.

REFERENCES

Afana, J., Marshall, J. and Tennent, P. (2023), Designing bodily materialistic representation for effective communication within virtual reality multimedia, CHI'23 Workshop-Body x Materials.

Ainsworth, K. D. and Hagino, C. C. (2006), 'A survey of Ontario chiropractors: their views on maximizing patient compliance to prescribed home exercise', *The Journal of the Canadian Chiropractic Association* **50**(2), 140–155.

Andreasson, J. and Johansson, T. (2016), ''Doing for group exercise what McDonald's did for hamburgers': Les Mills, and the fitness professional as global traveller', *Sport, Education and Society* **21**(2), 148–165.

Assandri, E. P. (2019), 'Is 'touching'essential when teaching ballet?', *Research in Dance Education* **20**(2), 197–207.

BBC Radio 4 (2020), 'Anatomy of touch: Nine things we learned from the world's largest study of touch'. Available from BBC website https://www.bbc.co.uk/.../nine-things-we-learned-from-the-world-s-largest-study-of-touch (May 22, 2025).

Braun, V. and Clarke, V. (2006), 'Using thematic analysis in psychology', *Qualitative Research in Psychology* **3**(2), 77–101.

Braun, V. and Clarke, V. (2019), 'Reflecting on reflexive thematic analysis', *Qualitative Research in Sport, Exercise and Health* **11**(4), 589–597.

Bårdstu, H. B., Andersen, V., Fimland, M. S., Aasdahl, L., Raastad, T., Cumming, K. T. and Sæterbakken, A. H. (2020), 'Effectiveness of a resistance training program on physical function, muscle strength, and body composition in community-dwelling older adults receiving home care: a cluster-randomized controlled trial', *European Review of Aging and Physical Activity* **17**(1), 11.

Camps, Jeroen and, C. T., , Jeroen, S. and Nelissen, J. (2013), 'A situational touch: How touch affects people's decision behavior', *Social Influence* **8**(4), 237–250. Routledge.

Chib, V. S., Adachi, R. and O'Doherty, J. P. (2018), 'Neural substrates of social facilitation effects on incentive-based performance', *Social Cognitive and Affective Neuroscience* **13**(4), 391–403.

Choi, Y., Tewell, J., Morisawa, Y., Pradana, G. A. and Cheok, A. D. (2014), Ring*U: a wearable system for intimate communication using tactile lighting expressions, *in* 'Proceedings of the 11th Conference on Advances in Computer Entertainment Technology', ACE '14, Association for Computing Machinery, New York, NY, USA, pp. 1–4.

Crowley, E. (2020), The specificity of dry-land resistance training exercises to swimming performance, thesis, University of Limerick.

Daling, L. M. and Schlittmeier, S. J. (2024), 'Effects of Augmented Reality-, Virtual Reality-, and Mixed Reality–Based Training on Objective Performance Measures and Subjective Evaluations in Manual Assembly Tasks: A Scoping Review', *Human Factors* **66**(2), 589–626. Publisher: SAGE Publications Inc.

Diez-Vega, I., Molina-Martin, J., Fernandez del Valle, M., Rodriguez Matoso, D. and Rodriguez-Ruiz, D. (2014), *Normalized response speed and jumping-related techniques after training in female volleyball players*, Vol. 9. Cultura, Ciencia y Deporte.

El Saddik, A., ed. (2012), *Haptics Rendering and Applications*, InTech.

Field, T. (2010), 'Touch for socioemotional and physical well-being: A review', *Developmental Review* **30**, 367–383.

Heatley Tejada, A., Dunbar, R. I. M. and Montero, M. (2020), 'Physical Contact and Loneliness: Being Touched Reduces Perceptions of Loneliness', *Adaptive Human Behavior and Physiology* **6**(3), 292–306.

Heslin, R., Nguyen, T. D. and Nguyen, M. L. (1983), 'Meaning of touch: The case of touch from a stranger or same sex person', *Journal of Nonverbal Behavior* **7**(3), 147–157.

Howard, M. C. (2017), 'Investigating the simulation elements of environment and control: Extending the Uncanny Valley Theory to simulations', *Computers & Education* **109**, 216–232.

Iwaanaguchi, T., Shinya, M., Nakajima, S. and Shiraishi, M. (2015), Cyber Tai Chi - CG-based Video Materials for Tai Chi Chuan Self-Study, *in* '2015 International Conference on Cyberworlds (CW)', pp. 365–368.

Jewitt, C., Price, S., Leder Mackley, K., Yiannoutsou, N. and Atkinson, D. (2020), Interdisciplinary Explorations of Digital Touch, *in* 'Interdisciplinary Insights for Digital Touch Communication', Springer International Publishing, Cham, pp. 23–37. Series Title: Human–Computer Interaction Series.

Jo, H.-Y., Seidel, L., Pahud, M., Sinclair, M. and Bianchi, A. (2023), FlowAR: How Different Augmented Reality Visualizations of Online Fitness

Videos Support Flow for At-Home Yoga Exercises, *in* 'Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems', CHI '23, Association for Computing Machinery, New York, NY, USA, pp. 1–17.

Kim, D., Kim, Y. and Jo, D. (2023), 'Exploring the Effect of Virtual Environments on Passive Haptic Perception', *Applied Sciences* **13**(1), 299. Multidisciplinary Digital Publishing Institute.

Koerner, S. and Staller, M. (2020), 'Coaching self-defense under COVID-19: challenges and solutions in the police and civilian domain', *Security Journal* .

Laine, T. H. and Suk, H. J. (2024), 'Investigating User Experience of an Immersive Virtual Reality Simulation Based on a Gesture-Based User Interface', *Applied Sciences* **14**(11), 4935. Multidisciplinary Digital Publishing Institute.

Lawson, G., Roper, T., Shaw, E., Hsieh, M. and Cobb, S. (2020), 'Multimodal virtual environments: an opportunity to improve fire safety training?', *Policy and Practice in Health and Safety* **18**, 1–14.

Lee-Cultura, S. and Giannakos, M. (2020), 'Embodied interaction and spatial skills: A systematic review of empirical studies', *Interacting with Computers* **32**.

Li, Z., Bujić, M., Buruk, O. Bampouni, E., Järvelä, S. and Hamari, J. (2024), 'Haptics-mediated virtual embodiment: Impact of a wearable avatar-controlling system with kinesthetic gloves on embodiment in VR', *Frontiers in Virtual Reality* **5**. Frontiers.

Liu, J., Zheng, Y., Wang, K., Bian, Y., Gai, W. and Gao, D. (2020), 'A Real-time Interactive Tai Chi Learning System Based on VR and Motion Capture Technology', *Procedia Computer Science* **174**, 712–719.

Liu, X.-W., Li, C.-Y., Dang, S., Wang, W., Qu, J., Chen, T. and Wang, Q.-L. (2022), 'Research on Training Effectiveness of Professional Maintenance Personnel Based on Virtual Reality and Augmented Reality Technology', *Sustainability* **14**(21), 14351. Multidisciplinary Digital Publishing Institute.

Lu, Y., Leng, X., Yuan, H., Jin, C., Wang, Q. and Song, Z. (2024), 'Comparing the impact of personal trainer guidance to exercising with others: Determining the optimal approach', *Heliyon* **10**(2), e24625.

Martínez, J., García, A., Oliver, M., Molina, J. P. and González, P. (2016), 'Identifying Virtual 3D Geometric Shapes with a Vibrotactile Glove', *IEEE Computer Graphics and Applications* **36**(1), 42–51.

McClaran, S. R. (2003), 'The Effectiveness of Personal Training on Changing Attitudes Towards Physical Activity', *Journal of Sports Science & Medicine* **2**(1), 10–14.

Meyer, M. J., Molle, A., Judkins, B. N. and Bowman, P. (2021), 'Martial Arts in the Pandemic', *Martial Arts Studies* **0**(11), 7.

Nomura, K., Hatada, Y., Narumi, T. and Kuzuoka, H. (2025), From virtual to physical: Investigating the carryover effects of avatar-mediated communication in intergenerational contexts, *in* 'ACM Symposium on Applied Perception 2025', SAP '25, Association for Computing Machinery, New York, NY, USA.

Oagaz, H., Schoun, B. and Choi, M.-H. (2022), 'Real-time posture feedback for effective motor learning in table tennis in virtual reality', *International Journal of Human-Computer Studies* **158**, 102731.

Oakley, I., Brewster, S. and Gray, P. (2001), 'Can You Feel the Force? An Investigation of Haptic Collaboration in Shared Editors'.

Parker, K., Uddin, R., Ridgers, N. D., Brown, H., Veitch, J., Salmon, J., Timperio, A., Sahlqvist, S., Cassar, S., Toffoletti, K., Maddison, R. and Arundell, L. (2021), 'The Use of Digital Platforms for Adults' and Adolescents' Physical Activity During the COVID-19 Pandemic (Our Life at Home): Survey Study', *Journal of Medical Internet Research* **23**(2), e23389.

Parviainen, J. (2011), 'The standardization process of movement in the fitness industry: The experience design of Les Mills choreographies', *European Journal of Cultural Studies* **14**(5), 526–541.

Preatoni, E., Hamill, J., Harrison, A., Hayes, K., {Van Emmerik}, R., Wilson, C. and Rodano, R. (2013), 'Movement variability and skills monitoring in sports', *Sports Biomechanics* **12**(2), 69–92.

Rodríguez González, L., Melguizo, E., Martín Moya, R. and González Valero, G. (2022), 'Study of strength training on swimming performance. A systematic review', *Science & Sports* **38**.

Serhan, D. (2020), 'Transitioning from Face-to-Face to Remote Learning: Students' Attitudes and Perceptions of using Zoom during COVID-19 Pandemic', *International Journal of Technology in Education and Science* **4**(4), 335–342. Number: 4.

Sodhi, R., Poupyrev, I., Glisson, M. and Israr, A. (2013), 'AIREAL: interactive tactile experiences in

free air', *ACM Trans. Graph.* **32**(4), 134:1–134:10. **URL:** *https://doi.org/10.1145/2461912.2462007*

Souchet, A. D., Philippe, S., Lourdeaux, D. and Leroy, L. (2022), 'Measuring visual fatigue and cognitive load via eye tracking while learning with virtual reality head-mounted displays: A review', *International Journal of Human-Computer Interaction* **38**(9), 801–824. Taylor & Francis.

Suzuki, R., Ofek, E., Sinclair, M., Leithinger, D. and Gonzalez-Franco, M. (2021), 'HapticBots: Distributed Encountered-type Haptics for VR with Multiple Shape-changing Mobile Robots'.

Takala, T. M., Hirao, Y., Morikawa, H. and Kawai, T. (2020), Martial Arts Training in Virtual Reality with Full-body Tracking and Physically Simulated Opponents, *in* '2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)', pp. 858–858.

Tanaka, K., Nakanishi, H. and Ishiguro, H. (2015), 'Physical embodiment can produce robot operator's pseudo presence', *Frontiers in ICT* **2**.

Tang, Y., Liu, S., Deng, Y., Zhang, Y., Yin, L. and Zheng, W. (2020), 'Construction of force haptic reappearance system based on Geomagic Touch haptic device', *Computer Methods and Programs in Biomedicine* **190**, 105344.

Tian, F., Ni, S., Zhang, X., Chen, F., Zhu, Q., Xu, C. and Li, Y. (2024), 'Enhancing Tai Chi Training System: Towards Group-Based and Hyper-Realistic Training Experiences', *IEEE Transactions on Visualization and Computer Graphics* **30**(5), 2713–2723.

Tsai, W.-L., Pan, T.-Y. and Hu, M.-C. (2022), 'Feasibility Study on Virtual Reality Based Basketball Tactic Training', *IEEE transactions on visualization and computer graphics* **28**(8), 2970–2982.

Viverito, N., Flanagan, E., Höchsmann, C., Beyl, R., Altazan, A., Martin, C. and Redman, L. (2021), 'Factors protecting against a decline in physical activity during the covid-19 pandemic', *Medicine Science in Sports Exercise* **53**, 1391–1399.

Wang, L., Gai, W., Jiang, N., Chen, G., Bian, Y., Luan, H., Huang, L. and Yang, C. (2024), 'Effective Motion Self-Learning Genre Using 360° Virtual Reality Content on Mobile Device: A Study Based on Taichi Training Platform', *IEEE Journal of Biomedical and Health Informatics* **28**(6), 3695–3708.

Wen, Y., Peng, Y., Zhao, F. and Zhen, K. (2016), 'The effects of specific drills on the flip turns of freestyle swimmers based on a kinesiology analysis', *Journal of Human Kinetics* **52**, 165–173.

Zhou, T., Zhai, X., Wu, N., Koriyama, S., Wang, D., Jin, Y., Li, W., Sawada, S. and Fan, X. (2022), 'Changes in physical fitness during covid-19 pandemic lockdown among adolescents: A longitudinal study', *Healthcare* **10**, 351.