Sensitivity to Hand Offsets and Related Behavior in Virtual Environments Over Time

KRISTOPHER KOHM, Clemson University, USA
JOHN PORTER III, Clemson University, USA
ANDREW ROBB, Clemson University, USA

This work explored how users' sensitivity to offsets in their avatars' virtual hands changes as they gain exposure to virtual reality. We conducted an experiment using a two-alternative forced choice (2-AFC) design over the course of four weeks, split into four sessions. The trials in each session had a variety of eight offset distances paired with eight offset directions (across a 2D plane). While we did not find evidence that users became more sensitive to the offsets over time, we did find evidence of behavioral changes. Specifically, participants' head-hand coordination and completion time varied significantly as the sessions went on. We discuss the implications of both results and how they could influence our understanding of long-term calibration for perception-action coordination in virtual environments.

CCS Concepts: • **Human-centered computing → Virtual reality**.

Additional Key Words and Phrases: Body awareness, hand offsets, calibration, longitudinal

1 INTRODUCTION

There has been a growing number of studies on body awareness within the virtual reality (VR) community in recent years. This work ranges from exploring how body awareness influences embodiment of virtual avatars (i.e. [11]) to techniques that redirect whole body movement in VR. The application of this redirection includes redirected walking (i.e. [1]) and even redirected jumping (i.e. [21, 28]). Another growing area of this redirection research relates to offsetting the user's virtual hand from their real hand. Like redirected walking, which manipulates head rotation and translation to decouple a user's virtual position from their physical position, hand offsets use translations to decouple the position of the user's virtual hand from their real hand's position. Studying how users respond to this decoupling can give us a broader understanding of perception-action coordination in VR. It can also lead to recommendations for implementing these offsets in VR applications.

Several studies have looked at the practical implications of research on proprioception in virtual environments. One such use is to help reduce strain on the VR user's hand [17]. A publication released in 2005 by Burns et. al looked at users' sensitivity to visual-proprioceptive conflict compared to visual interpenetration [9]. They recommended that users were much less sensitive to visual-proprioceptive conflict, making that a better choice between the two to design with. These offset techniques are readily available to developers. Work by Zenner et. al introduced a toolkit which implemented several of the hand redirection techniques such as body warping and world warping [57] for Unity [1]. Through the use of offsets, reaching in VR applications could require less

[1]https://unity.com/

Authors' addresses: Kristopher Kohm, kckohm@clemson.edu, Clemson University, Clemson, South Carolina, USA, 29643; John Porter III, jjporte@clemson.edu, Clemson University, Clemson, South Carolina, USA, 29643; Andrew Robb, arobb@clemson.edu, Clemson University, Clemson, South Carolina, USA, 29643.

1544-3558/2022/9-ART $15.00
https://doi.org/10.1145/3561055

movement and therefore be easier. The offsets could also change the user's action capabilities, making previously unavailable actions like reaching over extended distances possible.

Perceptual work regarding body awareness has largely focus on measuring the *perceptual thresholds* where the offset between a user's real and virtual hand becomes noticeable. This awareness of the spatial and mechanical status of the musculoskeletal framework is known as proprioception [48]. Proprioception can be sensed independently of the visual system which is important for virtual reality (specifically HMD's). When in an HMD, users have little to no visual information about the outside world, including the position of their own body. Therefore, it is up to these internal mechanisms and visual information in the headset of tracked body movement to give them the location and orientation of things like their limbs.

As a theoretical framework to our study we use the ecological approach to perception, where invariants in the environment lawfully generate information for the organism (or VR user) about their surroundings [7]. The ecological approach, compared to the traditional approach, includes the concept of calibration. The user can calibrate their movements to the invariants, even when the environment is manipulated. As mentioned above, organisms have certain action capabilities available to them which are referred to as affordances [19]. These affordances let the user known how they can interact with their environment and thus drive their perception-action coordination. We expect the user to calibrate over time even when their affordances have been altered. For this study, we manipulate the affordances of their own body by extending or compressing their reaching capability.

Bingham and Romack investigated the effects of calibration on participants' ability to adapt to displacement prisms which skewed their vision while completing a task requiring reaches across several sessions and days [6]. They found that there was no rate of change in calibration during the blocks but that later trials required less calibration because participants were more accurate at the start of subsequent blocks. In other words, participants were able to calibrate despite the visual displacement. In real applications that use body offsets, users would have time to calibrate when an offset is present. The ecological approach posits that calibration does not just occur with explicit feedback about performance but instead happens constantly. Since users are always receiving information about the interactions in their environment (or virtual environment) they can always be calibrating to it. Thus, we expect that their performance and behavior would change over time with this calibration.

The ecological approach also includes the concept of control laws. Control laws relate information to the action system so that organisms can adapt their behavior based on the situation [54]. By using these control laws, organisms (or VR users) can correct for alterations in their actions online (during the task).

To the best of our knowledge, no prior work has considered how user's sensitivity to body offsets might change with experience over time, or how their behavioral responses to such offsets might change. This is important information for researchers and developers as many applications that employ these techniques are likely to see repeated use by the same users. As such, if perceptions and behaviors do evolve over time these applications should adjust with the user or base the design on users with greater experience.

Head-hand coordination is one behavior that may change with calibration to the offsets over time. Eye-head coordination is needed for gaze fixation [15], but head-hand coordination is also needed for some fine motor tasks. Pelz et al. studied eye, head, and hand coordination during a task that required participants to pick up and move blocks [37]. They found that changes in head gaze were often linked to hand movement through short synergies. In ecological perception, synergies reduce the complexity to carry out an action by reducing the degrees of freedom the individual can control [7]. As users become familiar with the virtual environment, it is possible that these synergies will change. This would change their behavior related to the head-hand coordination.

In this work we explored how offsets applied to participants' hands during a block manipulation task impacted their sensitivity to these offsets and behavior during the task. We expected these changes would occur as a result of calibration. However, given there is little related work on perception-action coordination over time we were unsure how their sensitivity or behaviors might change.

Our hypotheses included:

H1: Participants' sensitivity to the hand offsets will change over time.
H2: Participants' control of their virtual hand will rely less on visual perception and more on proprioception over time.
H3: Participants' performance, despite the offsets, will improve over time.

2 RELATED WORKS

2.1 Hand Offsets

Proprioception plays an important role in VR systems, specifically for HMD's when the user's view of the outside world (and their own body) is obscured. A study by Rohde et. al measured the effect of this drift in a non-virtual environment through the use of the Rubber Hand Illusion [43]. This showed how proprioception is conveyed outside of virtual reality and how proprioception is subject to multiple sensory inputs. In VR, these offsets are usually applied to a virtual avatar's orientation. In many of the studies discussed below, these offsets were applied to the virtual position of the user's hand.

There are multiple ways to apply offsets for hand redirection in virtual environments. These offsets can be applied uniformly or non-uniformly. Work by Benda et. al provides estimation thresholds for fixed (uniform) hand offsets in six directions and found that thresholds vary significantly based on the direction of the offset [5]. Specifically, they found that the negative Z ("close") offset direction had the smallest threshold range of 7.83 cm and the negative Y ("down") offset direction had the largest threshold range of 13.37 cm. This implies that magnitude as well as direction can vary offset detection threshold. Like this work, we utilized a 2-AFC method and fixed positional offsets to a virtual hand to measure human perception of offset hand placement. However, our experiment applied offsets in a 2-D plane rather than a 3-D axis. Studies that utilize non-uniform offsets are usually referred to as gain-based interactions. These gain-based interactions are an alternative way of seamlessly redirecting a user's hand in VR (e.g. [14], [58]). Instead of the offset remaining at a set distance from the user's hand it tends to increase as their hand moves further from their body.

Ogawa et. al found that influences other than offset amounts, such as virtual avatar realism, can have significant impacts on the perception of proprioceptive offsets (in the case of this study, the shifts were applied horizontally) [36]. It is interesting that just by changing self-representation in a virtual environment could change a user's affordances. The "avatar" presented in our study is just the participants' virtual hand to avoid any external influences of body awareness.

Azmandian et. al introduced haptic retargeting as a technique that maps an object in the non-virtual world to multiple objects in the virtual world [2]. Their work shows several techniques for aligning physical and virtual objects which included world manipulation, body manipulation, and a hybrid technique. The body warping technique presented in their paper applies a translation offset to reach movements so that the users hand reaches the desired destination. To achieve the desired result of hand redirection, hand offsets are often combined with a variety of other techniques (e.g. [49], [31], [20]).

These studies show how a VR user's action capabilities for reaching can be altered. These alterations can be fixed (as with our study) or be non-uniform depending on how the user interacts with their surroundings. Though these methods differ, participants' sensitivity and behavior have been shown to change when the offsets are presented.

2.2 Long-Term Exposure VR Studies

There are several different kinds of studies related to long-term exposure in VR that have been conducted over several different time frames. Steinicke and Bruder exposed a participant to an IVE for 24 hours straight (with short breaks allowed during the session) [47]. They reported that the participant experienced varying levels of

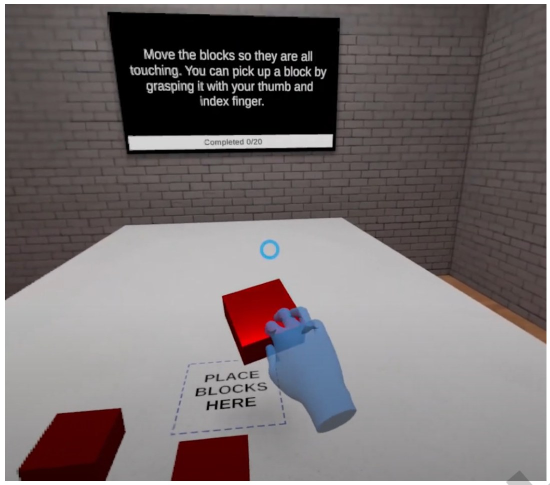

Fig. 1. This image shows the testing environment for the hand offsets. The virtual hand was controlled via the participants' physically tracked hand. The red blocks were picked up through a pinching motion and placed inside of the bounding box.

sickness throughout the course of the experiment based on the SSQ that was not mitigated by time. A similar study by Nordahl et al. exposed two participants to prolonged exposure in HMD's for 12 hours [35]. They found that participants' reported sickness went up and down inconsistently throughout the study.

Zielasko also self-reported sensitivity to simulator sickness, but instead of prolonged exposure to VR in one session they discussed their seven year exposure to VR systems [59]. This work highlights the need for longitudinal studies, specifically for simulator sickness, to examine the complexity of the problem. A survey of studies related to simulator sickness by Dużmańska et al. found that duration of sickness and a threshold time for maximum sickness varies greatly between studies [13]. Porter and Robb found evidence that suggested the way consumers talk about simulator sickness and other factors in VR in online forums may change as they gain experience over prolonged periods of exposure [39]. Baileson and Yee studied users in a collaborative virtual environment across 10 weeks in 15, 45 minute sessions [3]. They found that users reported less sickness as the participants gained experience in the virtual environment, but also that participants looked at each other less over time. Porter et al. also found a small effect related to simulator sickness over time through a study comparing immersive and non-immersive virtual environments in 6, 45 minute sessions over 3 weeks [38].

Many longitudinal studies in VR are system usability studies where the researchers study participants' behavior in the context of particular VR systems over time. This work spans studies related to training (e.g. [42, 46, 56]), therapeutic applications (e.g. [12, 16, 18, 23, 25, 51]), education (e.g. [22, 50]), and computer-supported collaboration (e.g. [24, 33, 52]). The primary focus of these studies is usually on specific experiences in VR and not general changes in interactions as they gain experience.

3 METHODS

3.1 Participants

We recruited 20 participants via an email sent to students enrolled at Clemson University. Participants were required to have normal or corrected-to-normal vision and to have had less than one hour of prior experience using VR. We had a mix of undergraduate and graduate level students participate whose ages ranged from 18-33 years old with a median age of 20 years old. 14 participants identified as male, five participants identified as female, and one participant identified as non-conforming/non-binary. 17 of the participants were in a Computer Science related field and three were in a variety of other fields. 14 of the participants identified as white, three participants identified as Asian, two participants identified as Hispanic, one participant identified as Black or African American, and one participant identified as Native American or Alaska Native. Two participants were left-handed. Participants received a $50 gift card upon enrolling in the study and a $75 gift card at the end of the study. This study was approved by the IRB office at Clemson University.

3.2 Participants' Usage Patterns

Based on the usage data recorded, participants spent an average of 17.13 hours ($sd = 10.58, min = 7.12, max = 45.24$) using the headset during the study. 119 unique applications were launched by participants, though the majority of these were used for less than one hour across all participants. 33 were used for more than one hour, with eleven of these being used for more than five hours. These applications were our experimental application, the Oculus home application, YouTube VR, Virtual Desktop, Oculus Link, Beat Saber, VRChat, Bait!, SUPERHOT VR, Sports Scramble, and SculptrVR. Discounting the experimental application, participants spent an average of 12.81 hours ($sd = 8.56, min = 3.21, max = 35.46$) in VR over the course of the four weeks.

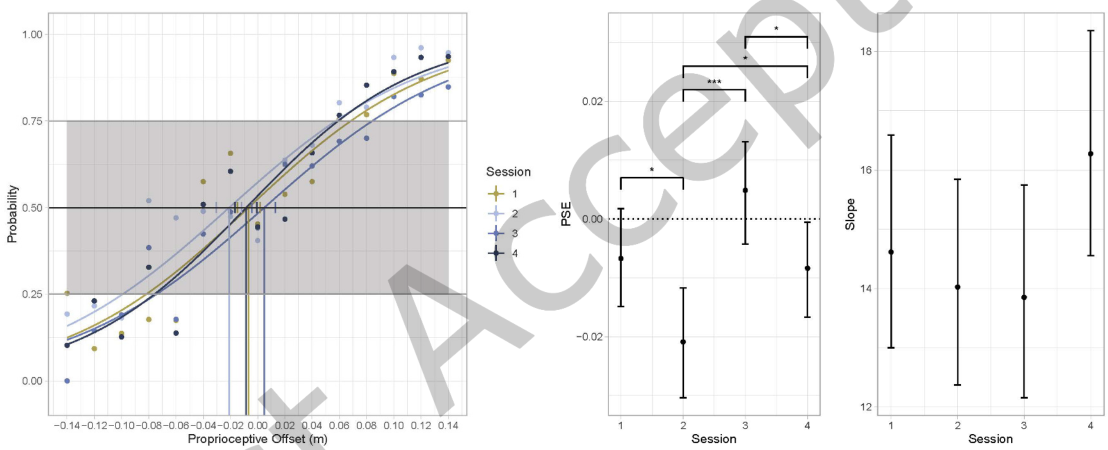

(a) The fitted psychometric functions (b) PSE by session (c) Slope by session

Fig. 2. Though there are significant differences between some of the PSE values, there is no pattern of change that would suggest participants became more or less sensitive to positive offsets over time. The probability on the vertical axis of the figure on the left indicates how likely participants are to detect a positive offset based on the offset along the horizontal axis. Similarly, the slopes of the psychometric curves do not follow a pattern that would indicate sensitivity to offsets changed significantly across sessions.

3.3 Procedure

Each participant was provided an Oculus Quest headset to use for the duration of the four week study. Participants were asked to spend at least 10 hours per week using the headset including time spent completing the study activities and any applications of interest to participants (e.g. games, movies, social applications). This time limit was impractical for most participants as shown by their usage patterns above. In addition to the body awareness activity, participants were also asked to complete two other experimental activities, one dealing with sensitivity to rotational gains and another with navigation through a virtual environment. While the focus of this paper is on the results of the body awareness experiment, we briefly explain the other activities to provide context for everything participants did during the four week study.

Participants completed each of the three activities before repeating them again. The order in which activities were completed was randomly assigned to each participant, but was held constant across the weeks for a given participant. Participants accessed these activities through a custom VR application we developed for the Quest which enforced the order in which activities were completed and reminded participants when each of the activities should be completed each week. Participants also completed two semi-structured interviews during the experiment. The first interview took place after the first or the second week (for logistic reasons half of the participant interviews were scheduled for the first week); the second interview was completed two weeks after the first interview. Participants were given detailed instructions on how to set up the Quest and how to install our custom application using the Sidequest[2] application. They were also instructed how to access data about their activities stored on the headset about their activities.

The software saved the data from each experiment as files on the Quest. Accessing the data for the amount of time participants used the headset required the installation of a second application: Activity Watch [3]. Activity Watch is an open-source time tracking application with clients for several operating systems, including Android (the OS used on the Oculus Quest). Activity Watch allowed us to access information about which applications participants were using, and for how long they were used. Each week, participants were asked to upload logs from our activities and from Activity Watch to a unique Google Drive folder. This allowed us to keep track of their progress and send reminders to participants who fell behind on completing activities.

The "rotation sensitivity and "navigation" activities are outside of the scope of this paper, however we provide a short description of each activity here. In the "rotation sensitivity" activity, participants stood in the center of a virtual garage and performed a series of rotations to the left and the right, with a positive or negative gain applied to one of the rotations. After turning to the left and the right, participants were asked to indicate whether they believed a gain had been applied to the first or the second rotation. In the "navigation" activity, participants used either teleportation or joystick locomotion to explore twenty maze-like levels. In order to reach the end of the level, they had to first find a keycard located in one of the rooms. After finding the keycard, they could use it to open a locked door leading to the exit. Once arriving at the exit, participants were asked to point back in the direction of where they had started and where they had found the keycard.

Figure 1 shows an example of what participants saw during the body awareness activity. A TV was used to give participants instructions at the start of each session and after each trial. A table in front of the participants held the blocks and had a small square on it with text indicating where the participants should place them. There were also buttons that floated slightly below the TV. Participants interacted with these buttons using a gaze pointer. The selection of the target was seen with a raycast line that started in the HMD and followed in the direction of the participants' gaze. The buttons were used to advance though the instructions and to collect responses after each trial. A button was also visible above the TV to recenter the virtual avatar's position in the room.

Participants were asked to be seated throughout the experiment. Prior to starting, participants were also asked to enable hand tracking on their Quests. They could see their tracked hands during the experiment as translucent blue hands floating in front of them. To pick up and drop the virtual blocks they used a pinching motion. Participants completed a total of 64 trials in each session.

[2]https://sidequestvr.com/

[3]https://github.com/ActivityWatch/aw-android

3.4 Apparatus

The body awareness experiment employed a 2-AFC design. For each trial, participants completed two block placing tasks. They used their dominant hand to put three blocks which had been semi-randomly placed on the table into the marked square on the table. Only a virtual copy of their dominant hand was present in the environment. One of these tasks had an offset in one of eight directions applied to it. The eight offset amounts ranged from 0.0 m to 0.14 m in increments of 0.02 m. The eight directions were North, Northeast, East, Southeast, South, Southwest, West, and Northwest. These directions were applied along a 2D plane parallel to the table. Right and left (i.e. East and West) were along the x-axis and forward and backwards (i.e. North and South) were along the z-axis.

After both tasks were completed, participants were asked which of the two tasks they believed had an offset applied to their virtual hand. A different hand offset amount was used in each of the 64 trials within a session. The order that the offsets were given was randomized within each session along with the direction of the offset.

3.5 Analyses

We used the quickpsy [29] package in R [41] to fit a logistic function to the 2-AFC data. We assumed a lapse rate of 0% with the point of subjective equality (PSE) and slope as free parameters. The lapse rate was assumed to be 0% since we had no prior work that suggested a higher rate. We also used a non-parametric bootstrap procedure with 1000 trials to estimate the 95% confidence intervals on the parameters. The fitted functions are shown in Figure 2a and the function parameters are shown in Table 1. Each of the functions fitted the data well as indicated by the non-significant p-values for the goodness-of-fit [40].

We considered all directions towards the participant's body as negative offset and all of the directions away from the participant's body as positive offsets for analysis. Specifically, moving closer towards the center of the participant's body for the x-axis was the negative direction. This was important based on the participant's dominant hand. For all participants, moving closer to their body along the z-axis was considered a negative offset. For the 18 right-handed participants this meant that West, Southwest, South, and Southeast were treated as negative offsets. For the two left-handed participants East, Southeast, South, and Southwest were treated as negative offsets. The direction along the x-axis was switched based on handedness because we considered movement towards the center of the participants' bodies as closer to them. Typically psychometric functions are fit to offsets in positive and negative directions so we transformed the data in this way to get a better fit [40].

Throughout the experiment we recorded the location of participants' head and virtual hands. These locations were based on world space within the environment and recorded in object log files. We used this data to calculate the angle between their virtually offset hand and the forward vector on their HMD. Without access to data about their eye movements, this forward vector can be used to roughly approximate their gaze direction in the virtual scene. This head-hand angle could estimate how much participants were looking at their virtual hand when completing the block stacking task.

The object log files also included time stamps based on in-game time. We annotated the angles with information regarding how the participants were handling the blocks at specific moments in time. For each angle we noted whether they had grasped, held, or ungrasped the block. The grasped and ungrasped events were recorded in the object log files the instant participants picked up or put down a block. The held angles were the average angle between a grasped/ungrasped pair. This approximated the head-hand angle while the participant was holding the block. We also recorded how long it took participants to complete each of the two tasks in each trial across the sessions.

We used linear-mixed models (LMMs) to analyze the head-hand angle and completion time data. The models were created in R first using the 'buildmer' [53] and 'lme4' [4] R packages. Buildmer automatically tests different possible models based on a set of independent variables and uses the likelihood-ratio test and minimum Bayesian information criterion to select the model that best fits the observed data [45]. The models suggested by buildmer was then considered from a theoretical perspective; in our case, all suggested models were theoretically sound. We then checked for violations of assumptions and made any necessary modifications to the suggested model formula based on the results. Once the final model was specified, we fit it to the data using the lmer command provided by lme4. The lmerTest package [27] was used to estimate p-values using the Satterthwhaite degrees of freedom method [44] for the models generated by lmer.

Each of the models was checked for normality of residuals, normality of random effects, a linear relationship, homogeneity of variance, and collinearity (see the 'check_model' function in the 'performance' R package [30]). Regarding collinearity, no predictors in any of the head-hand angle models had a variance inflation factor (VIF) greater than five. In the completion time model only the interaction effect between session and trial had a VIF greater than five. The VIF for this interaction was roughly six suggesting a moderate level of collinearity. Since the VIF for both session and trial individually were low we decided the interaction effect was not problematic and therefore we did not remove it from the completion time model.

All of the participants completed at least three sessions of this activity (two participants were missing one session each), so we did not remove any of the participants' data from the analysis. The 'check_outliers' function of the 'performance' R package [30], with the outlier threshold set to a z-score of three, was used to identify and exclude data points. For the head-hand angle data 730 out of 114,291 data points were removed. For the completion time data 1,044 out of 9,984 data points were removed. Outliers were excluded prior to the construction of the LMMs. Figures for all three sets of data were generated in R using ggplot2 [55].

3.6 Interpreting Linear Mixed Models

We discuss the advantages and disadvantages of LMMs below since they are often not used in VR research. However, there is a precedence for their used in related work such as psychological research [32].

LMMs are useful at analyzing data where conditions or measurements are repeated multiple times [26]. LMMs are also more robust than ANOVAs for assumption violations and missing or imbalanced data. This is particularly important for our study where participants completed the experiments at home. This meant that we had less control over how the data was collected and could end up with missing data. Further, we had less control over when data was collected. LMMs can handle data that is captured on an irregular schedule which was useful for that reason. Given these two variable factors, LMMs are especially appropriate for analyzing longitudinal data [10].

LMMs model both *fixed* and *random* effects [32]. Fixed effects model the influence of predictors on the measured quantities, while random effects model the unexplained variability associated with individual differences between participants. LMMs independently model each participants' response to the data. Random effects can be modeled as *random intercepts*, where the intercept of each participant's individual model is allowed to vary based on the best fit to that participant's data. Random effects can also be modeled as *random slopes*, where the slope of each participant's individual model is also allowed to vary based on the best fit to their particular data. Including random effects in the model improves LMMs' abilities to precisely model the response of fixed effects independently of variability between participants.

The fitted model includes an intercept, which describes the predicted value of the modeled variable when all predictors are set to zero. The model also includes fitted beta values for each fixed effect, which describe how much the modeled variable is predicted to change when the fixed effect unit increases by one. Beta values can vary greatly based on the dependent variables in a model. If the range of the variables are small then the beta will be small and vice versa. Thus, to determine the importance of a beta value we provide standardized beta values (abbreviated std. beta in our tables). These values are scaled based on the distribution of each dependent variable. After scaling, we can more accurately predict the relative contribution each fixed effect has on the independent variable. The standardized beta values tell us the standard deviation change in the dependent variable when the independent variable changes by one standard deviation [8].

We also provide both marginal and conditional r^2 values. Marginal r^2 values (abbreviated m. r^2 in our tables) report the amount of variance explained by the fixed effects alone and conditional r^2 values report the amount of variance explained by both the fixed and random effects [34]. Thus, the conditional r^2 values tell us the explanatory power for the entire model.

4 RESULTS

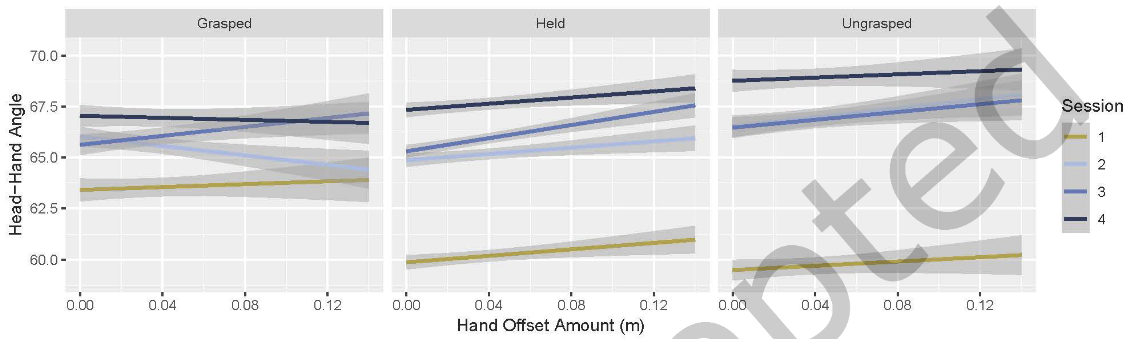

Fig. 3. Head-hand coordination angle increased consistently with session, however an interaction between session and offset amount was only observed for the held and ungrasped events.

4.1 Psychometric Model

To evaluate H1, we explored properties of the psychometric curves fitted to the 2-AFC data. The properties of these curves are presented below.

Table 1. Psychometric Parameters

Session	Goodness of Fit	PSE		Slope	
	p-value	Estimate	95% CI	Estimate	95% CI
1	0.829	-0.007	[-0.015, 0.002]	14.614	[13.001, 16.587]
2	0.838	-0.021	[-0.030, -0.012]	14.024	[12.370, 15.842]
3	0.973	0.005	[-0.004, 0.013]	13.849	[12.158, 15.742]
4	0.887	-0.008	[-0.017, -0.001]	16.271	[14.554, 18.346]

4.1.1 Point of Subjective Equality. While significant differences were observed between sessions for the PSE, a consistent pattern of change did not emerge (see Figure 2b). Two of the observations did not differ significantly from zero, as indicated by the 95% confidence intervals crossing zero, and a third barely demonstrated a significant difference from zero. The final observation suggested that the PSE was roughly 2 cm away from zero offset. Taken as a whole, these results do not suggest that a consistent change in PSE was present over time. Futher, little evidence suggests that the PSE was distinguishable from zero, suggesting participants were equally sensitive to positive and negative offsets.

4.1.2 Slope. No significant differences were observed between sessions for the slope of the psychometric model (see Figure 2c). As such, no evidence was found that participants became more or less sensitive to proprioceptive offsets over time.

4.2 Head-Hand Coordination

To explore H2, we modeled hand-head coordination in three events: when participants began grasping a block, the duration when participants held a block (and were typically moving it to its destination), and when participants placed a block. Figure 3 shows the aggregate data for all participants across all sessions and for each of the events. The fitted models for these three events are presented below.

4.2.1 Head-hand Coordination When a Block was Grasped. This model predicts the head-hand coordination angle at the moment when participants grasped a block. The model included session as a fixed effect and participant ID as a random effect. The model's total explanatory power (conditional r^2) is 0.16 and the explanatory power of the fixed effect alone (marginal r^2) is 0.0046. The parameter values for the model are shown in Table 2.

Based on the model's intercept (when session was zero), participants' head-hand angle was 62.45 degrees when they grasped the blocks. For each unit of session this angle increased by 1.17 degrees. This suggests that participants' head gaze was less focused on the block being grasped over time, which may indicate that participants became more confident acting when an offset was present over time as they felt less of a need to orient their heads towards the position of the visual hand. Notably, offset did not factor into this model. This suggests that participants gazed at the blocks when grasping the same amount regardless of the offset of their hands. This could imply that visuo-motor control was more important for grasping blocks than proprioception.

Table 2. Grasped Head-Hand Angle Parameters

Fixed Effect	beta [95% CI]	Std. beta [95% CI]	t(32283)	p	m. r^2
(Intercept)	62.45 [59.14, 65.76]	–	36.98	<0.001	–
Session	1.17 [1.00, 1.34]	0.07 [0.06, 0.08]	13.14	<0.001	0.0046

4.2.2 Head-hand Coordination While a Block was Held. This model predicts the the head-hand coordination angle while participants held a block. The model included session and offset as fixed effects, and participant ID as a random effect. The model's total explanatory power (conditional r^2) is 0.20 and the explanatory power of the fixed effects alone (marginal r^2) is 0.03. The parameter values for the model are shown in Table 3.

Based on the model's intercept (when session was zero), participants' head-hand angle was 58.38 degrees when they were holding the blocks. For each unit of session this angle increased by 2.31 degrees. As with the grasped model, this suggests that participants' head gaze was less focused on the block being grasped over time which may indicate that participants became more confident acting when an offset was present over time. Unlike the grasped model, offset was observed to also influence the head-hand coordination angle, such that the head-hand angle increased slightly as the absolute offset amount increased. This effect is small in comparison to the effect of session, corresponding to an increase of 1.37° at the largest offset of 0.14 m. This may result from an increased reliance on proprioception. Since participants were already holding the block, they did not need to visually orient towards it to continue holding it. As such, the coordination between their head and hand may have been more strongly influenced by proprioception.

Table 3. Held Head-Hand Angle Parameters

Fixed Effect	beta [95% CI]	Std. beta [95% CI]	t(48810)	p	m. r^2
(Intercept)	58.38 [55.57, 61.19]	–	40.68	<0.001	–
Session	2.31 [2.20, 2.42]	0.17 [0.16, 0.18]	41.07	<0.001	0.0284
Offset Amount (m)	9.77 [7.23, 12.31]	0.03 [0.02, 0.04]	7.54	<0.001	0.0009

4.2.3 Head-hand Coordination When a Block was Ungrasped. This model predicts the head-hand coordination angle at the moment when participants ungrasped a block. The model included session and offset as fixed effects, and participant ID as a random effect. The model's total explanatory power (conditional r^2) is 0.14 and the explanatory power of the fixed effects alone (marginal r^2) is 0.02. The parameter values for the model are shown in Table 4.

Based on the model's intercept (when all predictors were zero), participants' head-hand angle was 58.87 degrees when they ungrasped the blocks. For each unit of session this angle increased by 2.59 degrees. As with both previous models, this may indicate that participants became more confident acting when an offset was present over time. Unlike the grasped model, but similar to the held model, offset also influenced the head-hand coordination angle when ungrasping a block. The head-hand coordination angle increased slightly as the absolute offset amount increased, corresponding to an increase of 1.06° at the largest offset of 0.14 m. This is surprising as placing a block seems more similar to grasping a block than holding a block, as blocks must be placed at specific locations which increases the need for visual orientation. However, this task required less precision when placing a block than when grasping a block, as the blocks could be placed anywhere within a 0.2x0.2 m zone.

Table 4. Ungrasped Head-Hand Angle Parameters

Fixed Effect	beta [95% CI]	Std. beta [95% CI]	t(32454)	p	m. r^2
(Intercept)	58.87 [55.92, 61.81]	–	39.18	<0.001	–
Session	2.59 [2.41, 2.76]	0.15 [0.14, 0.16]	28.60	<0.001	0.0223
Offset Amount (m)	7.595 [3.53, 11.66]	0.02 [0.00874, 0.03]	3.66	<0.001	0.0003

4.3 Completion Time of Trials

To explore H3, we modeled completion time of each task in each trial. The model included session, trial, offset amount, and the interaction between session and trial as fixed effects, and participant ID as a random effect. The model's total explanatory power (conditional r^2) is 0.34 and the explanatory power of the fixed effects alone (marginal r^2) is 0.16. The parameters for the model are shown in Table 5.

Based on the model's intercept (when all predictors were zero), participants' completion time of each task was 5.26 seconds. For each unit of session this time decreased by 0.53 seconds. Each trial decreased the completion time by 0.02 seconds which would be equivalent to 1.28 seconds by the end of 64 trials. Session also moderated the effect of trial as shown by increasing the completion time by 0.004 seconds. Increasing the offset amount made the tasks more difficult and therefore increased the completion time. This increase was 3.815 seconds for every meter of offset amount. For every 0.02 meter change this would be equivalent to a 0.076 second increase. A graph of the completion time data across trials can be found in Figure 4.

Table 5. Completion Time Parameters

Fixed Effect	beta [95% CI]	Std. beta [95% CI]	t(8933)	p	m. r^2
(Intercept)	5.261 [4.99, 5.53]	–	39.18	<0.001	–
Session	-0.53 [-0.57, -0.49]	-0.34 [-0.36, -0.32]	-24.85	<0.001	0.046
Trial	-0.0236 [-0.03, -0.02]	-0.19 [-0.20, -0.17]	-15.05	<0.001	0.0164
Offset Amount (m)	3.815 [3.35, 4.28]	0.14 [0.12, 0.16]	16.19	<0.001	0.0193
Session:Trial	0.004 [0.00294, 0.00512]	0.06 [0.05, 0.08]	7.25	<0.001	0.0038

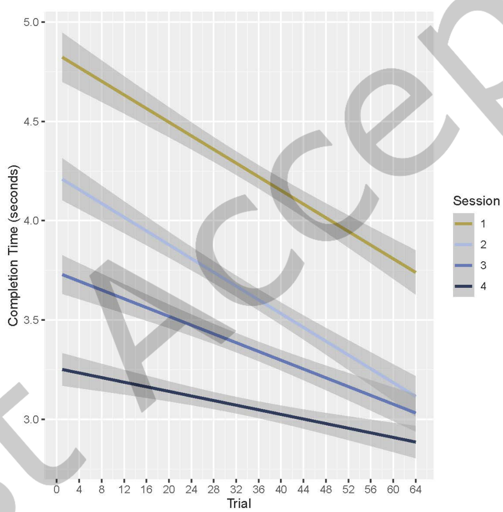

Fig. 4. This figure shows how completion time of tasks within trials changed based on the trial. The data is separated by session which also had a main effect. There was a small interaction effect between session and trial.

5 DISCUSSION

Our analyses of the data did not provide evidence to support H1. While significant differences were observed between sessions for the modeled PSEs, a clear pattern of change was not observed in the data. No significant differences were observed for changes in slope between sessions, indicating that sensitivity of body offsets did not change across the four sessions. Taken together, we found no evidence that participants' sensitivity to body offsets changed with exposure.

Regarding H2, we found evidence that participants' hands drifted away from the center of their head gaze over time. This suggests that participants became more confident manipulating objects using proprioception rather than just visual perception. This was true when grasping a block, when holding and moving a block, and when ungrasping a block (which typically corresponded to placing the block at its destination). They also became more confident in the presence of an offset over time as indicated by the interaction effects between offset and session. Notably, this interaction effect was present when holding or ungrasping a block and not when grasping a block.

While the effect of offset amount was small, this may indicate an increased reliance on proprioception when holding and ungrasping blocks compared to when grasping a block. Since the head-hand angle increased with the offset amount, participants' visual hands were closer to the center of their head gaze when offsets were lower. As such, participants did not strongly reorient their head gaze towards the position of their visual hand when an offset was present. This may indicate the position of their real hand, sensed via proprioception, played a larger role while holding and ungrasping blocks. The lack of an effect of offset amount on head-hand angle during the grasping circumstance may be attributed to the precision required when grasping a block, compared to the other circumstances. Fine motor motions were required to grasp the block which may best be achieved with visual information. Gross motor motions were sufficient to move and place the block since the target location was several times the size of the block.

Since trial was not included as a predictor in any of the head-hand models we can conclude that fatigue did not have a significant effect in each session. If it had, trial would have had a significant effect on each of the angles. Further, the effect of session in these models can show learning effects over time.

Regarding H3, we observed improved performance indicated by task completion time as session and trial increased. We expected an increase in task performance from learning effects, but it is noteworthy that we observed an interaction effect between session offset amount. The interaction between session and offset amount moderated the effect of offset amount. While the main effect of offset amount increased the time required to complete a trial, the interaction effect worked to counteract this. This reduced the time penalty associated with an increased offset amount. This may suggest that participants were able to correct for the hand offset better over time.

The contrasting results of these three hypotheses suggest that while participants' sensitivity to hand offsets may not change their behavior does. Consciously, they did not improve at selecting which task had the offset applied. Unconsciously, they improved at completing the tasks throughout the sessions. This improvement may have also led to their change in behavior.

Given that participants were unable to detect changes in hand offsets more accurately over time it seems unlikely they were able to calibrate to the offsets. This lack of improved accuracy could be explained by the fact that the offsets were presented in a random order. Thus, participants had little time to calibrate before their offset changed. Instead, participants probably experienced an improvement in online control using control laws. Since we found that participants improved at the task and their behaviors change it seems likely that their ability to correct for hand offsets improved between sessions. Therefore, the improvement in online control could have contributed to their change in head behavior, not calibration. This idea could be explored further in future work.

5.1 Limitations

Eye gaze could have more accurately measured the trade-off between vision and proprioception in this experiment better than head gaze. Given the nature of this study eye gaze was not available for use. However, head gaze still provides useful information about the synergies between head-hand coordination such as in [37]. Further, some of the participants used applications that support hand-tracking (i.e. YouTube VR). We do not know if participants used the hand-tracking features. If they did, they could have adjusted to the tasks more quickly across sessions than they would have using hand-tracking in the experiment alone.

6 CONCLUSION

This work highlighted how hand offsets influence VR users over time. We did not find evidence that participants' sensitivity for detecting these offsets changed. We did find evidence that behaviors related to these offsets changed. The difference in head-hand coordination and completion time suggests that participants improved at correcting to the offsets as the experiment went on. These differences in behavior were not large effects but they were present. While we originally expected that participants might have calibrated with repeated exposure to the offsets over time, the randomized nature of the offsets may have prevented calibration. Instead, participants may have learned to have better online control. This could be useful for VR training simulations that may want to specifically prevent calibration from occurring but still improve at the task.

ACKNOWLEDGMENTS

This work was supported in part by the US National Science Foundation (CISE HCC) under Grant #1717937.

REFERENCES

[1] Mahdi Azmandian, Timofey Grechkin, Mark Bolas, and Evan Suma. 2016. The redirected walking toolkit: a unified development platform for exploring large virtual environments. In *2016 IEEE 2nd Workshop on Everyday Virtual Reality (WEVR)*. 9–14. https://doi.org/10.1109/WEVR.2016.7859537

[2] Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. 2016. *Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences*. Association for Computing Machinery, New York, NY, USA, 1968–1979. https://doi.org/10.1145/2858036.2858226

[3] Jeremy N Bailenson and Nick Yee. 2006. A longitudinal study of task performance, head movements, subjective report, simulator sickness, and transformed social interaction in collaborative virtual environments. *Presence: Teleoperators and Virtual Environments* 15, 6 (2006), 699–716.

[4] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67, 1 (2015), 1–48. https://doi.org/10.18637/jss.v067.i01

[5] Brett Benda, Shaghayegh Esmaeili, and Eric D. Ragan. 2020. Determining Detection Thresholds for Fixed Positional Offsets for Virtual Hand Remapping in Virtual Reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 269–278. https://doi.org/10.1109/ISMAR50242.2020.00050

[6] Geoffrey Bingham and Jennifer L Romack. 1999. The rate of adaptation to displacement prisms remains constant despite acquisition of rapid calibration. *Journal of Experimental Psychology: Human Perception and Performance* 25, 5 (1999), 1331.

[7] Julia JC Blau and Jeffrey B Wagman. 2022. Introduction to Ecological Psychology: A lawful approach to perceiving, acting, and cognizing. (2022).

[8] Johan Bring. 1994. How to standardize regression coefficients. *The American Statistician* 48, 3 (1994), 209–213.

[9] E. Burns, S. Razzaque, A.T. Panter, M.C. Whitton, M.R. McCallus, and F.P. Brooks. 2005. The hand is slower than the eye: a quantitative exploration of visual dominance over proprioception. In *IEEE Proceedings. VR 2005. Virtual Reality, 2005.* 3–10. https://doi.org/10.1109/VR.2005.1492747

[10] Avital Cnaan, Nan M Laird, and Peter Slasor. 1997. Using the general linear mixed model to analyse unbalanced repeated measures and longitudinal data. *Statistics in medicine* 16, 20 (1997), 2349–2380.

[11] Diane Dewez, Rebecca Fribourg, Ferran Argelaguet, Ludovic Hoyet, Daniel Mestre, Mel Slater, and Anatole Lécuyer. 2019. Influence of Personality Traits and Body Awareness on the Sense of Embodiment in Virtual Reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 123–134. https://doi.org/10.1109/ISMAR.2019.00-12

[12] Felipe Augusto dos Santos Mendes, José Eduardo Pompeu, Alexandra Modenesi Lobo, Keyte Guedes da Silva, Tatiana de Paula Oliveira, Andrea Peterson Zomignani, and Maria Elisa Pimentel Piemonte. 2012. Motor learning, retention and transfer after virtual-reality-based training in Parkinson's disease–effect of motor and cognitive demands of games: a longitudinal, controlled clinical study. *Physiotherapy* 98, 3 (2012), 217–223.

[13] Natalia Dużmańska, Paweł Strojny, and Agnieszka Strojny. 2018. Can simulator sickness be avoided? A review on temporal aspects of simulator sickness. *Frontiers in psychology* 9 (2018), 2132.

[14] Shaghayegh Esmaeili, Brett Benda, and Eric D. Ragan. 2020. Detection of Scaled Hand Interactions in Virtual Reality: The Effects of Motion Direction and Task Complexity. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 453–462. https://doi.org/10.1109/VR46266.2020.00066

[15] Yu Fang, Ryoichi Nakashima, Kazumichi Matsumiya, Ichiro Kuriki, and Satoshi Shioiri. 2015. Eye-head coordination for visual cognitive processing. *PloS one* 10, 3 (2015), e0121035.

[16] AL Faria, J Couras, MS Cameirão, T Paulino, GM Costa, and S Bermúdez i Badia. 2016. Impact of combined cognitive and motor rehabilitation in a virtual reality task: an on-going longitudinal study in the chronic phase of stroke. *Proceedings of the 11th ICDVRAT* (2016).

[17] Tiare Feuchtner and Jörg Müller. 2018. Ownershift: Facilitating Overhead Interaction in Virtual Reality with an Ownership-Preserving Hand Space Shift. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) *(UIST '18)*. Association for Computing Machinery, New York, NY, USA, 31–43. https://doi.org/10.1145/3242587.3242594

[18] Daniel Freeman, C Thompson, Natasha Vorontsova, G Dunn, L-A Carter, P Garety, E Kuipers, Mel Slater, Angus Antley, Ed Glucksman, et al. 2013. Paranoia and post-traumatic stress disorder in the months after a physical assault: a longitudinal study examining shared and differential predictors. *Psychological medicine* 43, 12 (2013), 2673–2684.

[19] James J Gibson. 1977. The theory of affordances. *Hilldale, USA* 1, 2 (1977), 67–82.

[20] Eric J. Gonzalez and Sean Follmer. 2019. Investigating the Detection of Bimanual Haptic Retargeting in Virtual Reality. In *25th ACM Symposium on Virtual Reality Software and Technology* (Parramatta, NSW, Australia) *(VRST '19)*. Association for Computing Machinery, New York, NY, USA, Article 12, 5 pages. https://doi.org/10.1145/3359996.3364248

[21] Daigo Hayashi, Kazuyuki Fujita, Kazuki Takashima, Robert W. Lindeman, and Yoshifumi Kitamura. 2019. Redirected Jumping: Imperceptibly Manipulating Jump Motions in Virtual Reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 386–394. https://doi.org/10.1109/VR.2019.8797989

[22] Tanya Hill and Hanneke du Preez. 2021. A Longitudinal Study of Students' Perceptions of Immersive Virtual Reality Teaching Interventions. In *2021 7th International Conference of the Immersive Learning Research Network (iLRN)*. IEEE, 1–7.

[23] Xieyining Huang, Kensie M Funsch, Esther C Park, Paul Conway, Joseph C Franklin, and Jessica D Ribeiro. 2021. Longitudinal studies support the safety and ethics of virtual reality suicide as a research method. *Scientific reports* 11, 1 (2021), 1–12.

[24] Negar Khojasteh and Andrea Stevenson Won. 2021. Working Together on Diverse Tasks: A Longitudinal Study on Individual Workload, Presence and Emotional Recognition in Collaborative Virtual Environments. *Frontiers in Virtual Reality* 2 (2021), 53.

[25] Hyun-Jin Kim, Seulki Lee, Dooyoung Jung, Ji-Won Hur, Heon-Jeong Lee, Sungkil Lee, Gerard J Kim, Chung-Yean Cho, Seungmoon Choi, Seung-Moo Lee, et al. 2020. Effectiveness of a participatory and interactive virtual reality intervention in patients With social anxiety disorder: longitudinal questionnaire study. *Journal of medical Internet research* 22, 10 (2020), e23024.

[26] Charlene Krueger and Lili Tian. 2004. A comparison of the general linear mixed model and repeated measures ANOVA using a dataset with multiple missing data points. *Biological research for nursing* 6, 2 (2004), 151–157.

[27] Alexandra Kuznetsova, Per B. Brockhoff, and Rune H. B. Christensen. 2017. lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software* 82, 13 (2017), 1–26. https://doi.org/10.18637/jss.v082.i13

[28] Yi-Jun Li, De-Rong Jin, Miao Wang, Jun-Long Chen, Frank Steinicke, Shi-Min Hu, and Qinping Zhao. 2021. Detection Thresholds with Joint Horizontal and Vertical Gains in Redirected Jumping. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 95–102. https://doi.org/10.1109/VR50410.2021.00030

[29] Daniel Linares and Joan López-Moliner. 2016. quickpsy: An R Package to Fit Psychometric Functions for Multiple Groups. *The R Journal* 8, 1 (2016), 122–131. https://journal.r-project.org/archive/2016-1/linares-na.pdf

[30] Daniel Lüdecke, Mattan S Ben-Shachar, Indrajeet Patil, Philip Waggoner, and Dominique Makowski. 2021. performance: An R package for assessment, comparison and testing of statistical models. *Journal of Open Source Software* 6, 60 (2021).

[31] Brandon J. Matthews, Bruce H. Thomas, Stewart Von Itzstein, and Ross T. Smith. 2019. Remapped Physical-Virtual Interfaces with Bimanual Haptic Retargeting. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 19–27. https://doi.org/10.1109/VR.2019.8797974

[32] Lotte Meteyard and Robert AI Davies. 2020. Best practice guidance for linear mixed-effects models in psychological science. *Journal of Memory and Language* 112 (2020), 104092.

[33] Fares Moustafa and Anthony Steed. 2018. A longitudinal study of small group interaction in social virtual reality. In *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology*. 1–10.

[34] Shinichi Nakagawa, Paul CD Johnson, and Holger Schielzeth. 2017. The coefficient of determination R 2 and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. *Journal of the Royal Society Interface* 14, 134 (2017), 20170213.

[35] Rolf Nordahl, Niels C Nilsson, Ali Adjorlu, Eduardo Magalhaes, Silvin Willemsen, Nikolaj S Andersson, Jonas Wang, and Stefania Serafin. 2019. 12 Hours in Virtual Reality: Two Cases of Long-Term Exposure to Consumer-Grade Virtual Reality. In *Proc. of IEEE VR Workshop on Immersive Sickness Prevention*.

[36] Nami Ogawa, Takuji Narumi, and Michitaka Hirose. 2021. Effect of Avatar Appearance on Detection Thresholds for Remapped Hand Movements. *IEEE Transactions on Visualization and Computer Graphics* 27, 7 (2021), 3182–3197. https://doi.org/10.1109/TVCG.2020.2964758

[37] Jeff Pelz, Mary Hayhoe, and Russ Loeber. 2001. The coordination of eye, head, and hand movements in a natural task. *Experimental brain research* 139, 3 (2001), 266–277.

[38] John Porter III, Matthew Boyer, and Andrew Robb. 2018. Guidelines on successfully porting non-immersive games to virtual reality: a case study in minecraft. In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play*. 405–415.

[39] John Porter III and Andrew Robb. 2019. An analysis of longitudinal trends in consumer thoughts on presence and simulator sickness in VR games. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. 277–285.

[40] Nicolaas Prins et al. 2016. *Psychophysics: a practical introduction*. Academic Press.

[41] R Core Team. 2017. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/

[42] Aylen Ricca, Amine Chellali, and Samir Otrnane. 2021. The influence of hand visualization in tool-based motor-skills training, a longitudinal study. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 103–112.

[43] Marieke Rohde, Massimiliano Di Luca, and Marc O Ernst. 2011. The rubber hand illusion: feeling of ownership and proprioceptive drift do not go hand in hand. *PloS one* 6, 6 (2011), e21659.

[44] Franklin E Satterthwaite. 1946. An approximate distribution of estimates of variance components. *Biometrics bulletin* 2, 6 (1946), 110–114.

[45] Gideon Schwarz. 1978. Estimating the dimension of a model. *The annals of statistics* (1978), 461–464.

[46] Sherrill J Smith, Sharon Farra, Deborah L Ulrich, Eric Hodgson, Stephanie Nicely, and William Matcham. 2016. Learning and retention using virtual reality in a decontamination simulation. *Nursing education perspectives* 37, 4 (2016), 210–214.

[47] Frank Steinicke and Gerd Bruder. 2014. A self-experimentation report about long-term use of fully-immersive technology. In *Proceedings of the 2nd ACM symposium on Spatial user interaction*. 66–69.

[48] Barry C Stillman. 2002. Making sense of proprioception: the meaning of proprioception, kinaesthesia and related terms. *Physiotherapy* 88, 11 (2002), 667–676.

[49] Mohamed Suhail, Shyam Prathish Sargunam, Dustin T. Han, and Eric D. Ragan. 2017. Redirected reach in virtual reality: Enabling natural hand interaction at multiple virtual locations with passive haptics. In *2017 IEEE Symposium on 3D User Interfaces (3DUI)*. 245–246. https://doi.org/10.1109/3DUI.2017.7893363

[50] Tuukka M Takala, Lauri Malmi, Roberto Pugliese, and Tapio Takala. 2016. Empowering students to create better virtual reality applications: A longitudinal study of a VR capstone course. *Informatics in Education* 15, 2 (2016), 287–317.

[51] Ioannis Tarnanas, Winfried Schlee, Magda Tsolaki, René Müri, Urs Mosimann, and Tobias Nef. 2013. Ecological validity of virtual reality daily living activities screening for early dementia: longitudinal study. *JMIR serious games* 1, 1 (2013), e2778.

[52] Viswanath Venkatesh and Philip Johnson. 2002. Telecommuting technology implementations: a within-and between-subjects longitudinal field study. *Personnel psychology* 55, 3 (2002), 661–687.

[53] Cesko C Voeten. 2020. Using 'buildmer'to automatically find & compare maximal (mixed) models.

[54] William H Warren. 2006. The dynamics of perception and action. *Psychological review* 113, 2 (2006), 358.

[55] Hadley Wickham. 2009. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. http://ggplot2.org

[56] Alexander Winkler-Schwartz, Khalid Bajunaid, Muhammad AS Mullah, Ibrahim Marwa, Fahad E Alotaibi, Jawad Fares, Marta Baggiani, Hamed Azarnoush, Gmaan Al Zharni, Sommer Christie, et al. 2016. Bimanual psychomotor performance in neurosurgical resident applicants assessed using NeuroTouch, a virtual reality simulator. *Journal of surgical education* 73, 6 (2016), 942–953.

[57] André Zenner, Hannah Maria Kriegler, and Antonio Krüger. 2021. *HaRT - The Virtual Reality Hand Redirection Toolkit*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3411763.3451814

[58] André Zenner and Antonio Krüger. 2019. Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 47–55. https://doi.org/10.1109/VR.2019.8798143

[59] Daniel Zielasko. 2021. Subject 001-A Detailed Self-Report of Virtual Reality Induced Sickness. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, 165–168.