Perception of Redirected Pointing Precision in Immersive Virtual Reality

Henrique G Debarba*[1,3] Jad-Nicolas Khoury†[1] Sami Perrin‡[1] Bruno Herbelin§[2] Ronan Boulic¶[1]

[1] Immersive Interaction Group, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
[2] Center for Neuroprosthetics, École Polytechnique Fédérale de Lausanne, Geneva, Switzerland
[3] Artanim Foundation, Geneva, Switzerland

Abstract

We investigate the self-attribution of distorted pointing movements in immersive virtual reality. Participants had to complete a multi-directional pointing task in which the visual feedback of the tapping finger could be deviated in order to increase or decrease the motor size of a target relative to its visual appearance. This manipulation effectively makes the task easier or harder than the visual feedback suggests. Participants were asked whether the seen movement was equivalent to the movement they performed, and whether they have been successful in the task. We show that participants are often unaware of the movement manipulation, even when it requires higher pointing precision than suggested by the visual feedback. Moreover, subjects tend to self-attribute movements that have been modified to make the task easier more often than movements that have not been distorted. We discuss the implications and applications of our results.

Index Terms: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality H.5.2 [Information Interfaces and Presentation]: User Interfaces—Evaluation/methodology

1 Introduction

The premise of Virtual Reality (VR) is to deliver a synthetic world that can be experienced as if it were real. Ideally, VR should mediate all input and output channels of a person to a point where she can no longer detect a discrepancy between the expected and rendered outcome to her actions. This i/o feedback loop is expected to register and interpret the users actions and provide appropriate sensory replacement [8]. But human perception is not a perfect capture of reality, and much of the information we experience as being collected from the external world are the product of brain inference [27]. This is an important enabling factor for VR technologies, as they do not have to match the physical reality and the physiological limits to be effective, and only have to be as good as human perception and expectations.

With the increase in VR popularity and its recent availability in the consumer market, we argue that a better understanding of the limits of human perception can lead to new venues for effective VR interaction. For instance, studies have experimentally manipulated the agreement of sensory signals, showing that vision is generally predominant over other senses [6, 28]. That is, discrepant sensory feedback (e.g. visio-proprioceptive and visuo-vestibular) tends to be solved in favor of vision. In this context we explore the self-attribution of bodily movements that have their visual feedback (as seen in VR) altered by factors that are external to the control of the users. Here we define the term self-attribution as the state where users are more likely than not to acquire the perception that they have complete control over the movements of a virtual hand.

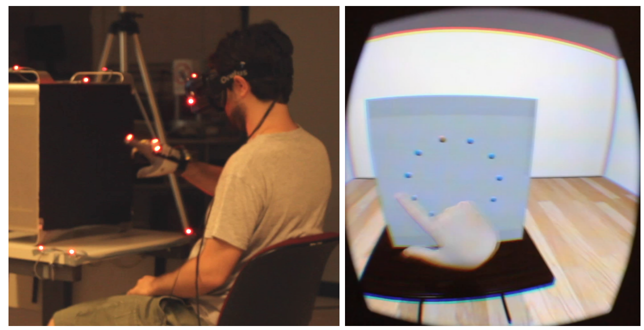

(a) (b)

Figure 1: Overview of the experimental setup. The subject sat in a chair wearing an HMD and a motion tracking glove (a). The task consisted of tapping on a sequence of targets placed in a circular arrangement (b).

We present a distortion function and an experiment to investigate aspects of the self-attribution of distorted pointing movements. In the experiment we manipulate the difficulty to accurately complete a tapping task (Fig. 1). The tapping task could be made easier or harder by changing the mapping from the physical to the virtual hand, i.e. the virtual hand position could diverge from the physical hand position. More specifically, the interaction (motor) area of the target could be made bigger or smaller than its visual size. Thus, we warp the space around the target, being capable of fitting a bigger or smaller physical area (motor) than the virtual feedback (visual) suggests, effectively facilitating or hindering the completion of the task (Fig. 2).

Understanding what are the circumstances leading to the self-attribution of a manipulated movement is relevant from different perspectives. For example, in the context of VR interaction, it can help designing assistive technologies that sustain the sense of accomplishment in applications that involve motor performance, such as physical rehabilitation, or to alleviate the impact of low quality tracking information. Self-attribution is also interesting in the study of mechanisms of bodily control and of mental disorders such as schizophrenia, which has been associated to an impairment – as compared to healthy subjects – in self-attribution [16].

The contribution of this paper are twofold: (i) providing a distortion function that manipulates the effective motor size of a target for VR interaction; and (ii) investigating aspects of self-attribution of distorted pointing movements that interfere with the difficulty of a task.

*e-mail: hgdebarba@gmail.com
†e-mail: jad-nicolas.khoury@epfl.ch
‡e-mail: sami.perrin@epfl.ch
§e-mail: bruno.herbelin@epfl.ch
¶e-mail: ronan.boulic@epfl.ch

2 RELATED WORK

In a pioneering study, Nielsen [26] demonstrated that when the visual feedback of the movement of the hand of a healthy unaware subject is replaced by a similar movement of a second person, the subject might erroneously attribute the seen movement to himself. This was the case even when there was a discrepancy between the performed and seen movements, with subjects reporting the feeling of strangeness and the impression that their hands have been pulled by some external force.

The altered motor control perception reported by Nielsen echoes on cross-modal illusions explored in VR, such as in the notion of pesudo haptics proposed by Lecuyer et al [19,22,23]. Pseudo haptics examines cross-modal perception to create the subjective sensation of haptic interaction with objects of different physical properties through the manipulation of the control-display ratio (CDR) of a mouse. In a related topic, [20] explored the distortion of movements in order to redirect haptic sensations. The goal was to use a passive haptic device as a proxy to a different or more complex virtual object. The authors have aimed at identifying potential performance changes, demonstrating that one can transfer skills obtained in such condition [21]. However, their work has only informally evaluated whether users perceive these manipulations. This topic has also been explored by Ban et al [3, 25] in the context of augmented reality, where haptic feedback of complex symmetric objects was redirected to a proxy physical cylinder. This work was further extended to explore pinching gestures [4]. Finally, Azmandian et al [1] presented a system for tactile interaction with multiple virtual objects using a single passive haptics proxy object.

Burns et al have explored two aspects of visuo-proprioceptive mismatch in order to propose interaction techniques preventing the interpenetration of the virtual hand with the virtual environment [10]. The first concerns the perception of a location mismatch between a physical and a virtual hand, demonstrating that a person may be strikingly unaware of visuo-proprioceptive mismatches which were gradually introduced over a long period of time [12]. The second aspect concerns the perception of movements with spatiotemporal distortions (speed increase/decrease relative to tracked speed) [11].

Distinct from the aforementioned works, we manipulate the difficulty to realize a task in VR, and aim to detect the thresholds to which this manipulation goes unnoticed by the user. Our manipulation is related to those presented in 2D interaction techniques such as Semantic Pointing [5], where the motor space of interactive elements are augmented to facilitate selection, as well as to the manipulation of the index of difficulty of a task [2]. By understanding how users perceive such manipulations in VR we can modulate the difficulty of a task to sustain interest and involvement without interfering with the experience of motor control of the subject.

3 DISTORTION FUNCTION

Our distortion function makes the tapping of a target at position $\mathbf{p}_{tgt}$ easier or harder by manipulating the position mapping of the physical to the virtual hand ($\mathbf{p}_{physical}$ and $\mathbf{p}_{virtual}$ respectively) i.e. the virtual and physical hands positions could diverge.

A distance range (d_{range}) is used so that when the target to physical hand distance ($d_{physical} = ||\mathbf{p}_{physical} - \mathbf{p}_{tgt}||$) is bigger than d_{range} no remapping occurs. The d_{range} is also used to normalize the values to the range $[0,1]$, and then scale this normalized remapping back into the virtual world units. Our dynamic remapping uses properties of exponentiation of values between 0 and 1 so that when the exponent (a in Equation 1) is bigger than 1 the motor radius of the target becomes bigger relative to its visual radius, and when it is less than 1 the motor radius of the target becomes smaller than the visual radius. Thus facilitating or hindering the completion of the goal directed task (Fig. 2).

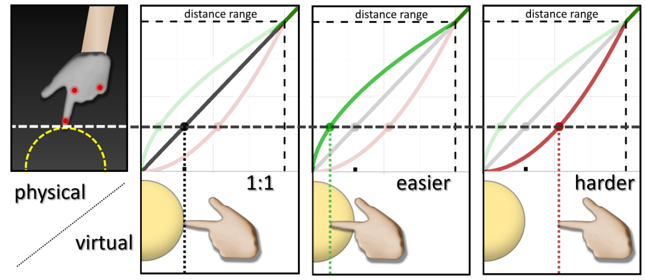

Figure 2: Overview of the pointing distortion function. The vertical axis depicts the physical finger position ($\mathbf{p}_{physical}$), while the horizontal axis depicts the virtual finger position ($\mathbf{p}_{virtual}$). Note that physical and virtual movements happen towards the same direction, and that this graphic represents the 1D case mapping of physical into virtual position for different settings of the distortion function. The lower left corner of each mapping plot (1:1, easier and harder) represents the center of the target ($\mathbf{p}_{tgt}$). The green and red colors represents a facilitating and a hindering distortion respectively.

$$d_{virtual} = \begin{cases} d_{range} \times (\frac{d_{physical}}{d_{range}})^a, & if\ d_{physical} \leq d_{range} \\ d_{physical}, & otherwise \end{cases} \tag{1}$$

Where the exponent a is defined by:

$$a = \log_{\frac{r_{motor}}{d_{range}}}(\frac{r_{visual}}{d_{range}}) \tag{2}$$

Where r_{visual} represents the visual radius of the target, and r_{motor} represents the motor radius of the target. Therefore, when $r_{motor} < r_{visual}$ the task becomes harder than the visual feedback suggests, and when $r_{motor} > r_{visual}$ the task becomes easier. Mind that our distortion model assumes that r_{visual} and r_{motor} are both smaller than d_{range}.

The current position of the virtual hand ($\mathbf{p}_{virtual}$) is then set using:

$$\mathbf{p}_{virtual} = \begin{cases} \mathbf{p}_{tgt} + \frac{\mathbf{p}_{physical} - \mathbf{p}_{tgt}}{d_{physical}} \times d_{virtual}, & if\ d_{physical} \neq 0 \\ \mathbf{p}_{physical}, & otherwise \end{cases} \tag{3}$$

An overview of the variables of the distortion function is presented in Fig. 3.

Based on the definition of index of difficult (ID) used in the Fitts law [30], our distortion function alters the difficulty of a pointing task.

$$ID = \log_2(\frac{D}{W} + 1) \tag{4}$$

Where D represents movement amplitude and W the target width, or diameter in our case. Intuitively, increasing the target diameter (W) results in a lower ID, while reducing the target diameter results in a higher ID, modulating the ID of the task while the D parameter is kept constant.

The proposed distortion alters the motor radius of the target relative to its visual radius. As a result, we have a visual ID that is the expected difficulty, and a motor ID that is the actually experienced difficulty.

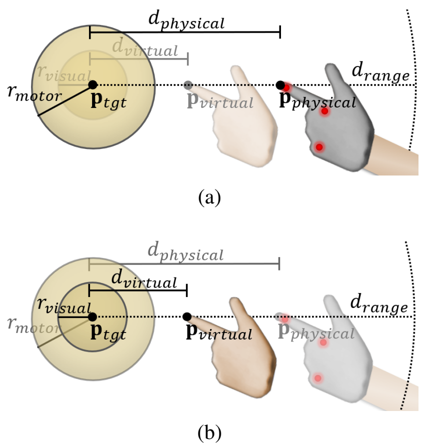

Figure 3: Variables of the distortion function with physical variables and motor radius in evidence in (a) and virtual variables and visual radius in evidence in (b). Notice that the physical and virtual positions of the hand do not overlap as a result of the distortion. Also notice that the difference in positions is exaggerated for illustrative purposes in this figure.

4 EXPERIMENT

4.1 Equipment

The participant wore an Oculus Development Kit 2 head mounted display (HMD) to visualize the virtual scene (960x1080 pixels per eye, 100 degrees field of view, 75Hz). Headtracking was performed using its inertial sensors and corrected for drift around the vertical axis using optical tracking (this yields lower latency than using the optical tracking alone, without damaging the correctness of the tracking).

A PhaseSpace ImpulseX2 system with 18 cameras was used for optical capture of participant and virtual objects. A total of 14 LED markers were used, 4 attached to the HMD, 3 attached to the hand, 3 attached to the table, and 4 attached to the tapping surface. The glove had a marker over the index fingertip, and two over the back of the hand. A rigid and flat stick was positioned between the top of the subject's index finger and the glove in order to prevent the finger from flexing. Thus preserving a congruent index position with the seen virtual hand. The table and the tapping surface were also tracked. Fig. 1 presents an overview of the setup.

The markers on the glove were pre-calibrated. To compensate for small changes in length of the index finger, we calibrate the tapping surface by pointing at three predefined corners of the tapping surface. The plane defined by these points was then used to translate and rotate the virtual representation. We set an activation tolerance of 2 millimeters over the tapping surface. That is, when the fingertip of the virtual hand was within a distance of 2mm from the tapping surface, these were considered as touching. This tolerance was adopted to prevent possible tracking or calibration imprecision from interfering with the detection of a touch. Moreover, subjects were asked to raise the finger when moving from target to target, thus preventing undesired selections. Note that the targets consist of circles over a tapping surface, and thus had a very narrow activation range (the 2mm tolerance). We therefore decided to constrain the distortion to the plane parallel to the taping surface, and use the physical hand to define the virtual hand position in the axis perpendicular to the surface. We did so because a distortion in this axis would simply affect the dynamic of the movement, without any functional role in the difficulty of the task. Moreover, because an activation tolerance was needed to address potential tracking imprecision, the manipulation of position in the axis perpendicular to the tapping surface would also play an unintended role of compressing or expanding the activation tolerance of the target.

The virtual environment was developed using the Unity game engine. It consists of a virtual hand, a chair, a table, and the tapping surface. We assessed a "motion to photon" latency in the range of 30ms to 40ms.

4.2 Task

The subject had to perform a multi-directional pointing task, as described by the annex B of ISO 9241-411 [18]. This task consists of multiple pointing movements – 11 in this experiment – toward targets equally spaced over the borders of a circle (Fig. 4). Subsequent targets are defined as to maximize the distance between them, and the sequence follows a clockwise direction. We remap the distance between the end-effector and the center of the target. The resulting distortion makes the motor radius of the target become bigger or smaller than its visual size, and therefore easier or harder to hit.

After each round of movements the subject was asked two questions:

- Did the virtual hand moved exactly like you?
- Did you miss any target?

The first question is meant to assess the thresholds of movement distortion perception for the proposed manipulation. The second question is meant to verify is subjects may be affected by a success bias, where they self attribute movements that have accomplished the task more often than movements that did not. An overview of a pointing trial is shown in Fig. 4.

4.3 Design

We manipulated two variables; the visual index of difficulty of the task (*visualID*), and the difference in the index of difficulty caused by the distortion (*diffID*), i.e. the motor index of difficulty resulting from a deviation to the visualID. The *visualID* could be set to 4 or 5, we did so while keeping the distance between targets constant ($D = 27cm$) and solving for the required visual diameter of the target using the equation

$$W_{visual} = r_{visual} \times 2 = \frac{D}{2^{visualID} - 1}. \tag{5}$$

The *diffID* could be set to -2.5, -2, -1.5, -1, -0.5, 0 (no distortion), 0.5, 1, 1.5, 2 or 2.5. A positive *diffID* means that the motor size of the target became smaller than the visual feedback suggests (i.e. harder), while a negative *diffID* had the opposite effect (i.e. easier). Similarly to W_{visual}, the motor diameter W_{motor} is computed with the equation:

$$W_{motor} = r_{motor} \times 2 = \frac{D}{2^{visualID+diffID} - 1} \tag{6}$$

The motor diameter values used in the experiment are presented in Table 1.

Table 1: The motor diameter (W_{motor}) values of the target for all combinations of the *visualID* and *diffID* variables and movement amplitude $D = 27cm$ used in the experiment. Diameter values in *cm*.

visualID \ *diffID*	-2.5	-2	-1.5	-1	-.5	0	.5	1	1.5	2	2.5
4	14.77	9	5.8	3.86	2.62	1.8	1.25	.87	.61	.43	.3
5	5.8	3.86	2.62	1.8	1.25	.87	.61	.43	.3	.21	.15

Moreover, the d_{range} in our distortion function was set to $D/2 = 13.5cm$. This allows the past and current pointing targets in a trial to be simultaneously active while ensuring the display of a continuous movement between the two targets. That is, the distortion effect

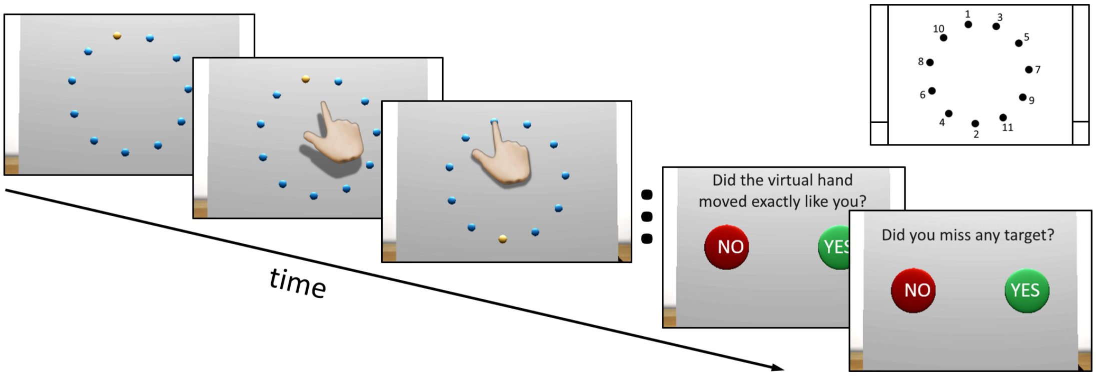

Figure 4: Overview of a trial. The subject had to tap 11 targets in a multi-direction pointing task. The current target is highlighted in orange. The movement could be distorted in the region surrounding each target. Once the participant taps the last target, they are asked whether the seen movement corresponds to the performed movement, and whether they have missed a target. The frame on the top right corner describes the target tapping order used in the trials.

of the past target is damped before the current target can affect the virtual hand position.

The experiment was divided into two blocks. Each block containing three stages, the first and the third stages contained 4 sequences of tapping with no distortion ($diffID = 0$), while the second stage presented each of the 22 combinations of *visualID* and *diffID* three times, with a randomized presentation order.

We designed this experiment to analyze three aspects of redirected interaction:

First, we want to assess the **thresholds of self-attribution for the proposed distortion function , i.e. the thresholds after which the distortion becomes more likely to be perceived than not by users**. To quantify these limits we adopt concepts and procedures from psychophysics. Psychophysics acts on the understanding of how a stimuli affects one's sensation/perceptions, and it is often employed to assess the minimum necessary change ΔI to a stimulus intensity I so that one can perceive a difference between $I+\Delta I$ and I with a high degree of confidence, normally more than 50% of the time. While the ΔI describes an amount relative to a specific stimulus intensity I, the Weber Law states that the ΔI can be defined by a constant proportion k of the stimulus when the stimuli intensity is not extreme [15]. This has been shown to model the stimuli and perception threshold relation considerably well for different sensory modalities and tasks. The Weber constant k can be defined by the ratio:

$$\frac{\Delta I}{I} = k \tag{7}$$

Therefore, we focus on measuring the constant k by estimating ΔI (*diffID* threshold of self-attribution) for different levels of the standard stimulus I (*visualID*). Once we know what is an admissible k for our distortion (one for facilitating and one for hindering) we want to compute a limit *motor* stimulus given a *visual* stimulus.

$$I_{motor} = I_{visual} + \Delta I = I_{visual} + I_{visual} * k = (1+k) * I_{visual} \tag{8}$$

Similar to [11], our study is distinct from regular psychophysics paradigms in that we assess the ΔI across different sensory modalities, i.e. to identify a discrepancy the user may rely on the incongruent visual and proprioceptive feedback of the movement, as well as on the differences between predicted (from motor commands) and realized visual feedback.

Second, we want to know what is the point of maximal self-attribution of movements, and whether it matches the point of no distortion. We hypothesize that **subjects are biased to self-attribute movement distortions that make the task easier**. This hypothesis is based on recent studies on the sense of being the agent of actions [17], and the role of higher order cognitive processes in the self-attribution of actions [31]. We can support this if we obtain a point of maximal agreement with the first question when a facilitating distortion is present, and that presents a difference that is statistically different from no distortion. Another measure that we implement is the relation between the two questions that we ask. That is, if subjects self-attribute movements when they report no missed target ("no" to second question) more often than when they report to have missed a target ("yes" to second question).

Third, when a mapping that makes the task harder is applied, subjects have to interact with greater accuracy in order to hit the target. We want to know whether **subjects change their motor behavior to interact with higher accuracy without noticing the movement discrepancy**. We will assess the pointing accuracy that subjects presented in the range at which they self attribute distorted movements, and compare to the situation of no distortion. Pointing accuracy will be defined as the effective index of difficulty (ID_e), which describes the difficulty of the task completed by the participant – based on the tapping error during the execution of the task – instead of the difficulty defined by the task parameters. The ID_e is defined by $ID_e = log_2(\frac{D}{W_e}+1)$ where $W_e = 4.133 \times SD_x$ and SD_x is the standard deviation of the taping error [18].

4.4 Results

Fifteen subjects participated in the experiment; they had to read and sign a written informed consent form to participate, and were compensated 20 CHF/hour for their participation. Two subjects were excluded from the analysis, one was excluded for not following closely the instructions, and one was excluded due to inconsistent distortion recognition performance that did not allowed the valid fit of the Gaussian function parameters used in our analysis. Pointing movements that took more than 4 seconds or with error above $7.5cm$ have been marked as not valid, and trials with less than 7 valid pointing movements (out of the 10 possible, the first movement was not analyzed) were excluded.

Results in terms of proportional change of difficulty ($\frac{diffID}{visualID}$) are presented in Fig. 5. We fit a Gaussian function to the proportion of positive answers to the self-attribution questions for each subject per level of the *visualID*. With the peak position ($\bar{x}$) and standard deviation (SD) parameters of the distribution we then compute the mean across the levels of *visualID* per subject.

We obtained $\bar{x} = -.062$ (Fig. 5 red dashed line) and $SD = .05$, which presents a statistically significant difference from 0 as ana-

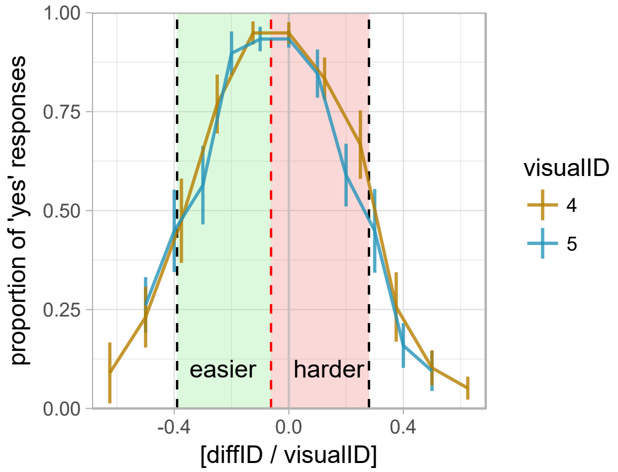

Figure 5: Summary of self-attribution of movements per movement distortion intensity as defined by the proportion *diffID*/*visualID*.The green/red shaded area represents the range where subjects were more likely to self-attribute a movement than not. Subjects tended to self-attribute distorted movements more often when the distortion made the task slightly easier (red dashed line) than the unmodified movement (i.e. $diffID/visualID = 0$). Error bars represent the standard error of the mean.

lyzed with a t-test ($t_{(12)} = -4.4$ $p < .001$). Moreover, considering the relation between the answer to the first and the second question, subjects were less likely to self-attribute a movement when they were aware that at least one target in the trial was missed ($t_{12} = 8.36$ $p < .001$, Fig. 6). These results suggest that **a manipulation from visual to motor space that slightly facilitates the task might be perceived as correct by users more often than when no manipulation is present**, therefore supporting our **second** point of interest.

Furthermore, we define the thresholds – **first** point of interest – as the points where subjects are likely to identify the discrepancy more than 50% of the time. Based on the parameters of the Gaussian distribution, these can approximated using $\bar{x} \pm SD * 1.1775$, which yields a helping threshold of $M = -.39$ $SD = .13$ and a hindering threshold of $M = .28$ $SD = .08$ (shaded area in Fig. 5).

Finally, we found a significant difference when comparing the levels 0 (no manipulation) and 1 (hindering manipulation) of *diffID* for both levels of the *visualID* variable with the t-test ($t_{(12)} = 3$ $p < .02$

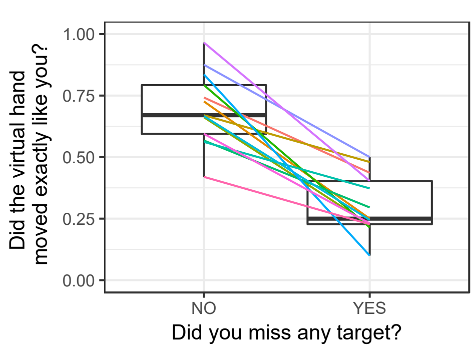

Figure 6: Relation of mean self-attribution agreement by perception of errors. Subjects were more likely to self-attribute a movements when they were not aware of pointing errors. The colored lines illustrate results for individual subjects.

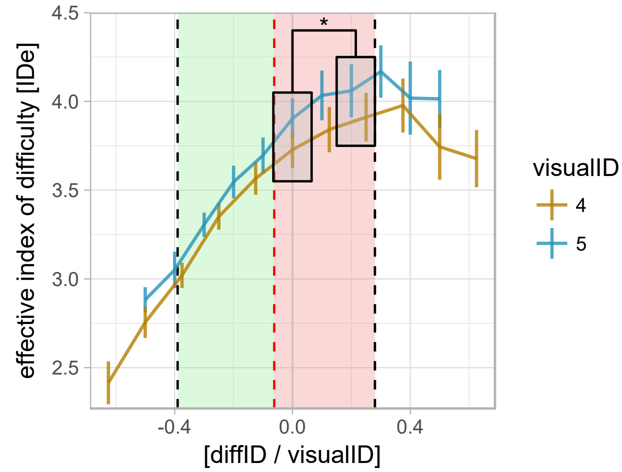

Figure 7: Summary of effective index of difficulty (ID_e). Subjects presented increased ID_e for two hindering distortions (grey shading) among the hindering manipulations that subjects were prone to identify as correct (ref shading).

for *visualID*= 4 and $t_{(12)} = 2.2$ $p < .05$ for *visualID*= 5). However, we failed to reject the difference between levels 0 and 0.5 of *diffID* ($t_{(12)} = 1.7$ $p > .1$ for *visualID*= 4 and $t_{(12)} = 1.5$ $p > .16$ for *visualID*= 5). For *diffID* values above one were likely to be identified by subjects, and therefore were not evaluated here. We conclude that our **third** point of interest is plausible as an increasing trend can be observed in Fig. 7, suggesting that **when subjects are faced with a manipulation that makes the task harder than suggested by the visual feedback, they tend to adjust for this manipulation without being necessarily aware of it** (up to a certain threshold).

5 Discussion and Conclusions

We proposed a distortion function that manipulates the difficulty to accomplish a pointing task. Our results show that subjects perform poorly in detecting discrepancies in spite of the visuo-motor discrepancy that the manipulation introduces. Additionally, we found that subjects are biased toward self-attributing distorted movements that make the task easier.

We note that the bias to self-attribution movements that make the task easier finds support in previous literature (Inoue et al [17]). Additionally, the evidence that subjects self-attribute movements more often when they believe that the task was successfully accomplished indicates a relation with the theory known as self-serving bias [24], which describes the bias to self-attribute successful actions as a mechanism to maintain an enhanced self-esteem.

Moreover, we believe that the findings that we present here could be used in designing more engaging VR interactions. For instance, movement distortion can be used to manipulate the difficulty of tasks, and consequently leverage the challenge of an interactive application so that it matches the skills of the user and promotes the state of flow [9, 13]. This could be the case in physical rehabilitation. For instance, post-stroke patients often experience reduced fine grained control of a limb movement [29]. Recovering from such condition usually requires the repetition of the movement as well as the feedback of completion. By distorting the movement with virtual reality a patient may practice in a more motivating environment, where she is capable of completing the task (an imprecise movement is transformed into a complete and precise movement) in order to maintain a sense of achievement. Recovery could then be accompanied with a progressive reduction of the helping distortion. We note the interest for a similar form of application presented by Dukes et al [14], in their system a short range movement could span a larger portion of the reachable space with the goal to provide a complete view of the action to the patient. We suggest the expansion of such rehabilitation system towards 2 complementary directions: the manipulation of movement precision in addition to movement amplitude; and the observation of self-attribution as a potential mean to manage the sense of achievement of the patient.

In future work, we propose to assess the transfer of training in a redirected situation. That is, does the skill trained in a system that manipulates the visual feedback of actions in order to control the difficulty of completing a task transfers to real activities? We also plan to expand the study of self-attribution for whole body manipulation. For instance, on the subject of movement manipulation for preserving self contact explored by Bovet et al [7].

ACKNOWLEDGMENTS

We wish to thank the anonymous reviewers for their comments. This work has been supported by the SNFS project "Immersive Embodied Interactions in Virtual Environments" with grant 200020-159968.

REFERENCES

[1] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson. Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pp. 1968–1979. ACM, 2016. doi: 10.1145/2858036.2858226

[2] R. Balakrishnan. "beating" fitts' law: Virtual enhancements for pointing facilitation. *Int. J. Hum.-Comput. Stud.*, 61(6):857–874, Dec. 2004. doi: 10.1016/j.ijhcs.2004.09.002

[3] Y. Ban, T. Kajinami, T. Narumi, T. Tanikawa, and M. Hirose. Modifying an identified curved surface shape using pseudo-haptic effect. In *Haptics Symposium (HAPTICS), 2012 IEEE*, pp. 211–216. IEEE, 2012. doi: 10.1109/HAPTIC.2012.6183793

[4] Y. Ban, T. Narumi, T. Tanikawa, and M. Hirose. Displaying shapes with various types of surfaces using visuo-haptic interaction. In *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology*, VRST '14, pp. 191–196. ACM, 2014. doi: 10.1145/2671015.2671028

[5] R. Blanch, Y. Guiard, and M. Beaudouin-Lafon. Semantic pointing: Improving target acquisition with control-display ratio adaptation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pp. 519–526. ACM, New York, NY, USA, 2004. doi: 10.1145/985692.985758

[6] M. Botvinick and J. Cohen. Rubber hands 'feel' touch that eyes see. *Nature*, 391(6669), 1998. doi: 10.1038/35784

[7] S. Bovet, H. G. Debarba, B. Herbelin, E. Molla, and R. Boulic. The critical role of self-contact for embodiment in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 24(4), 2018. doi: 10.1109/TVCG.2018.2794658

[8] D. A. Bowman and R. P. McMahan. Virtual reality: How much immersion is enough? *Computer*, 40(7):36–43, 2007. doi: 10.1109/MC.2007.257

[9] R. Brondi, L. Alem, G. Avveduto, C. Faita, M. Carrozzino, F. Tecchia, and M. Bergamasco. Evaluating the impact of highly immersive technologies and natural interaction on player engagement and flow experience in games. In *Entertainment Computing - ICEC 2015*, pp. 169–181. Springer, 2015. doi: 10.1007/978-3-319-24589-8_13

[10] E. Burns. *Macbeth: Management of Avatar Conflict by Employment of a Technique Hybrid*. PhD thesis, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, 2007. AAI3262622.

[11] E. Burns and F. P. Brooks. Perceptual sensitivity to visual/kinesthetic discrepancy in hand speed, and why we might care. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST '06, pp. 3–8. ACM, New York, NY, USA, 2006. doi: 10.1145/1180495.1180499

[12] E. Burns, S. Razzaque, A. T. Panter, M. C. Whitton, M. R. McCallus, and F. P. Brooks. The hand is more easily fooled than the eye: Users are more sensitive to visual interpenetration than to visual-proprioceptive discrepancy. *Presence*, 15(1):1–15, 2006. doi: 10.1162/pres.2006.15.1.1

[13] M. Csikszentmihalyi and I. S. Csikszentmihalyi. *Optimal experience: Psychological studies of flow in consciousness*. Cambridge university press, 1992.

[14] P. S. Dukes, A. Hayes, L. F. Hodges, and M. Woodbury. Punching ducks for post-stroke neurorehabilitation: System design and initial exploratory feasibility study. In *2013 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 47–54, 2013. doi: 10.1109/3DUI.2013.6550196

[15] G. Fechner. *Elemente der Psychophysik*. Number v. 1 in Elemente der Psychophysik. Breitkopf und Härtel, 1860.

[16] N. Franck, C. Farrer, N. Georgieff, M. Marie-Cardine, J. Daléry, T. d'Amato, and M. Jeannerod. Defective recognition of ones own actions in patients with schizophrenia. *American Journal of Psychiatry*, 158(3):454–459, 2001. doi: 10.1176/appi.ajp.158.3.454

[17] K. Inoue, Y. Takeda, and M. Kimura. Sense of agency in continuous action: Assistance-induced performance improvement is self-attributed even with knowledge of assistance. *Consciousness and Cognition*, 48:246 – 252, 2017. doi: 10.1016/j.concog.2016.12.003

[18] ISO. Ergonomics of human-system interaction – part 411: Evaluation methods for the design of physical input devices. ISO 9241-411:2012, International Organization for Standardization, Switzerland, 2012.

[19] D. A. G. Jauregui, F. Argelaguet, A. H. Olivier, M. Marchal, F. Multon, and A. Lecuyer. Toward "pseudo-haptic avatars": Modifying the visual animation of self-avatar can simulate the perception of weight lifting. *IEEE Transactions on Visualization and Computer Graphics*, 20(4):654–661, 2014. doi: 10.1109/TVCG.2014.45

[20] L. Kohli. Redirected touching: Warping space to remap passive haptics. In *3D User Interfaces (3DUI), 2010 IEEE Symposium on*, pp. 129–130. IEEE, 2010. doi: 10.1109/3DUI.2010.5444703

[21] L. Kohli, M. C. Whitton, and F. P. Brooks. Redirected touching: Training and adaptation in warped virtual spaces. In *3D User Interfaces (3DUI), 2013 IEEE Symposium on*, pp. 79–86. IEEE, 2013. doi: 10.1109/3DUI.2013.6550201

[22] A. Lécuyer, J.-M. Burkhardt, and C.-H. Tan. A study of the modification of the speed and size of the cursor for simulating pseudo-haptic bumps and holes. *ACM Trans. Appl. Percept.*, 5(3):14:1–14:21, Sept. 2008. doi: 10.1145/1402236.1402238

[23] A. Lecuyer, S. Coquillart, A. Kheddar, P. Richard, and P. Coiffet. Pseudo-haptic feedback: can isometric input devices simulate force feedback? In *Virtual Reality, 2000. Proceedings. IEEE*, pp. 83–90. IEEE, 2000. doi: 10.1109/VR.2000.840369

[24] A. H. Mezulis, L. Y. Abramson, J. S. Hyde, and B. L. Hankin. Is there a universal positivity bias in attributions? a meta-analytic review of individual, developmental, and cultural differences in the self-serving attributional bias. *Psychological bulletin*, 130(5):711, 2004. doi: 10.1037/0033-2909.130.5.711

[25] T. Narumi, Y. Ban, T. Kajinami, T. Tanikawa, and M. Hirose. Augmented perception of satiety: Controlling food consumption by changing apparent size of food with augmented reality. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pp. 109–118. ACM, 2012. doi: 10.1145/2207676.2207693

[26] T. I. Nielsen. Volition: A new experimental approach. *Scandinavian Journal of Psychology*, 4(1):225–230, 1963. doi: 10.1111/j.1467-9450.1963.tb01326.x

[27] V. S. Ramachandran, R. L. Gregory, et al. Perceptual filling in of artificially induced scotomas in human vision. *Nature*, 350(6320):699–702, 1991. doi: 10.1038/350699a0

[28] I. Rock and J. Victor. Vision and touch: An experimentally created conflict between the two senses. *Science*, 143(3606):594–596, 1964. doi: 10.1126/science.143.3606.594

[29] B. Rohrer, S. Fasoli, H. I. Krebs, R. Hughes, B. Volpe, W. R. Frontera, J. Stein, and N. Hogan. Movement smoothness changes during stroke recovery. *Journal of Neuroscience*, 22(18):8297–8304, 2002.

[30] R. W. Soukoreff and I. S. MacKenzie. Towards a standard for pointing device evaluation, perspectives on 27 years of fitts law research in hci. *International Journal of Human-Computer Studies*, 61(6):751 – 789, 2004. doi: 10.1016/j.ijhcs.2004.09.001

[31] M. Synofzik, G. Vosgerau, and A. Newen. I move, therefore i am: A new theoretical framework to investigate agency and ownership. *Consciousness and Cognition*, 17(2):411 – 424, 2008. Social Cognition, Emotion, and Self-Consciousness. doi: 10.1016/j.concog.2008.03.008