

Improving Virtual Reality Ergonomics Through Reach-Bounded Non-Linear Input Amplification

Johann Wentzel
University of Waterloo
jdwentze@uwaterloo.ca

Greg d'Eon
University of British Columbia
gregdeon@cs.ubc.ca

Daniel Vogel
University of Waterloo
dvogel@uwaterloo.ca

ABSTRACT

Input amplification enables easier movement in virtual reality (VR) for users with mobility issues or in confined spaces. However, current techniques either do not focus on maintaining feelings of body ownership, or are not applicable to general VR tasks. We investigate a general purpose non-linear transfer function that keeps the user's reach within reasonable bounds to maintain body ownership. The technique amplifies smaller movements from a user-definable neutral point into the expected larger movements using a configurable Hermite curve. Two experiments evaluate the approach. The first establishes that the technique has comparable performance to the state-of-the-art, increasing physical comfort while maintaining task performance and body ownership. The second explores the characteristics of the technique over a wide range of amplification levels. Using the combined results, design and implementation recommendations are provided with potential applications to related VR transfer functions.

Author Keywords

interaction techniques; ergonomics; input re-mapping

INTRODUCTION

The large movements typically required in virtual reality (VR) often become fatiguing, cumbersome, or even impractical in constrained environments. For example, large arm swinging or reaching movements in VR when seated at a desk could result in damage to equipment or personal injury. Furthermore, extensive movement in VR may be uncomfortable or even impossible for users with mobility issues.

A typical remedy to this problem is *input amplification*: transforming smaller, more comfortable movements into the larger, more dramatic movements that the user expects. However, many of these techniques come at a cost. Typical amplification techniques allow the user to manipulate objects at distances much further than a typical arm's reach [3, 15, 25]. This unrealistic increase in reach comes at the expense of *body ownership* — the psychological mapping of one's real body to a virtual body [29] — detracting from the user's feeling of presence




Figure 1. The physical controller position (green), relative to a calibrated neutral position, is amplified using a non-linear function so the virtual controller (blue) appears farther away. The transfer function keeps the physical-to-virtual hand offset small near the body, but maximum virtual reach can be achieved with the real controller moving 30% less.

in a virtual environment [13]. Other techniques have been shown to increase comfort while still maintaining body ownership, including Ownershift [7] and Erg-O [19]. However, these techniques remap input based on specific positions of targets in the virtual environment, limiting their applicability for applications with no distinct targets.

We introduce a family of transfer functions for increasing comfort in VR that we call reach-bounded, non-linear (RNL) input amplification. Instead of large movements that extend the user's virtual arm to superhuman levels, which reduces body ownership, the RNL approach allows users to reach within their typical arm length more comfortably and with less strain (Figure 1). The method applies a transfer function to amplify the distance of the hand position relative to a calibrated neutral position near the torso. This creates a separation, or *hand offset*, between the user's virtual hand and real hand, calculated as a percentage of the user's maximum arm reach. The function uses a configurable non-linear Hermite curve, and it can be tuned for different amplification levels. As opposed to other techniques, this method does not require targets in the virtual environment, making it more task-independent and more inter-operable with current VR applications.

We evaluated our approach in two experiments. In the first, we found that two pre-selected levels of RNL curves achieve similar results to the state-of-the-art Ergo-O method [19], reducing ergonomic strain by 6% with no significant reduction in body ownership, while also reducing physical motion by up to 18%. A second experiment explores perception and performance of RNL functions at increasing levels of amplification, from barely noticeable to upper limits. Results show that even small amounts of amplification can provide significant improvement

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-6708-0/20/04 ...\$15.00.

DOI: <https://dx.doi.org/10.1145/3313831.3376687>

to ergonomics, maximum hand offset can reach 20% arm's reach (roughly 14 cm) before body ownership is reduced, 30% (21 cm) before task performance degrades, and 20% before the majority of users perceive amplification.

We make three contributions: (1) an easy-to-implement general purpose technique for improving arm ergonomics while maintaining body ownership; (2) validation of its effectiveness in a user study and meta-comparison; and (3) empirically-informed design guidelines for developing pro-body ownership transfer functions for VR.

BACKGROUND AND RELATED WORK

Amplifying hand motion creates a positional offset between the user's physical hand and its virtual counterpart. Previous work has proposed using offsets for various purposes in VR.

Offsets can be used to exploit visual dominance over proprioceptive cues [4], creating different perceptual effects without loss of body ownership. Rietzler et al. [27] investigate dynamic offsets to create a feeling of virtual object weight, recommending offsets less than 24 cm for optimal immersion. Samad et al. [28] used linearly scaled offsets of up to 40% to create the illusion of object weight. Lloyd et al. [17] investigate the "rubber hand illusion" [2] under increasing translational offsets, recommending hand offsets less than 30 cm to maintain a feeling of body ownership.

Offsets are also used for passive haptics, synthesizing the feel of real objects in virtual space [12]. Examples include haptic retargeting [1], redirected reach toward small physical objects [10, 30], and redirected reach toward targets on a sparse haptic proxy [22]. These techniques set offsets based on physical object positions, then modify the virtual hand position such that the user's real hands are guided toward those objects.

More directly related to our approach are offsets created by amplifying virtual hand positions for the purpose of distant reaching. The Go-Go technique [25] uses a two-stage function where hand movement is not amplified until it exceeds two-thirds of the user's maximum reach. After this point, a quadratic scaling function amplifies the hand position to distances far beyond natural limits. This unbounded reach exceeds the limits set by previous work [2, 27], reducing the user's body ownership as a result. The PRISM technique [8] dynamically computes an offset based on the user's hand speed to switch between distant reaching and up close interaction. However, this technique by design contains motion discontinuities and is not designed with body ownership in mind.

Li et al. [15] tested four amplification methods: zero offset, fixed offset, linearly-scaled offset, and the Go-Go two-stage function. They found that zero offset is best when targets are in reach, and linear offset is best when targets are beyond reach. Their linear function, like Go-Go, modifies a user's reach to be "superhuman," creating too large an offset between the user's real and virtual hands, reducing the related feelings of immersion and body ownership.

Feuchtner and Müller's Ownershift technique [7] explored the interaction between hand amplification offsets and body ownership in VR using a virtual panel GUI. In their experiment,

the participant begins working at an upward reach position, then the system slowly guides the user's physical hands down to a comfortable position while their virtual hands remain high. They evaluated this gradually-scaling dynamic hand offset, showing that gradual adjustment is key to maintaining body ownership.

Erg-O [19] is a reaching technique that focuses on ergonomics and is implicitly designed to maintain immersion and body ownership. It uses a more technical approach for amplifying movement by re-mapping tetrahedral subdivisions of the nearby virtual space corresponding to the user's natural reaching area. A formal evaluation showed Erg-O significantly improved ergonomics by up to 7.2%, with an 11% time reduction in one of three task layouts.

Like our method, Erg-O has a core consideration for ergonomics and is designed to keep virtual arm reach within normal range-of-motion bounds. However, its real-time re-targeting algorithm performs optimizations based on discrete target positions, making it unclear how tasks requiring continuous input and no distinct targets, like drawing, would be amplified. In addition, the implementation and required algorithms (e.g., using simulated annealing with a tuned objective function, requiring complex 3D transformations) may reduce the technique's reproducibility and practical applications for typical VR developers. However, due to similarities in design priorities, we replicate their formal experiment to evaluate our method, and use the results to make a direct meta comparison.

Our approach is inspired by the simplicity and versatility of input amplification functions like Go-Go and Li et al.'s linear function, the design goals of Erg-O to improve ergonomics without making users "superhuman", and Ownershift's principle of gradual adjustment to maintain body ownership.

REACH-BOUNDED NON-LINEAR AMPLIFICATION

The RNL approach is an ergonomics-focused, physically realistic, non-linear input technique with four goals: (1) minimize amplification near the body to maintain precision; (2) minimize physical strain when reaching out by amplifying user input; (3) maintain body ownership with smooth amplification increases and realistic amounts of reach; and (4) achieve the above with a simple, easily-replicable implementation.

Calibration and Amplification

The technique requires a one-time per-user calibration. First, the user selects a comfortable near-body neutral position P_N by pressing a controller button at the desired point. The user then records their maximum reach r_{max} , defining a sphere of reachable points centred at the shoulder position P_S (Figure 2a). Unlike Li et al.'s [15] techniques that rely on the position of the user's head, the reachable sphere and neutral position are coupled to the shoulder, allowing free head movement without affecting the amplification. In Experiment 1, the shoulder position is tracked using a hardware tracker for precision; however, we use inverse kinematics [23] to infer P_S in Experiment 2, showing that this additional hardware is unnecessary.

After calibration, the user's hand position P_H is amplified as follows. First, a ray is cast from P_N through P_H . P_{max} is the




Figure 2. Key geometric points captured in calibration and used for the amplification technique: (a) the user’s maximum reach P_{max} , at distance r_{max} from their shoulder point P_S ; (b) P_H^* is calculated from the user’s hand position P_H .

point where this ray intersects with the edge of the reachable sphere. This point represents the user’s maximum reach in this direction. Then, the current physical offset r is:

$$r = |P_H - P_N| / |P_{max} - P_N|.$$

This physical offset is used to calculate an amplified offset $f(r)$, determining the amplified virtual hand position (P_H^*) as:

$$P_H^* = P_N + f(r) \cdot (P_H - P_N).$$

This process is illustrated in Figure 2b.

Amplification Functions

The core of RNL is a family of amplification functions for $f(r)$. These are designed to maintain precision close to the body by minimizing amplification near the hand’s neutral position P_N , but increase comfort with sufficient amplification when the user extends their arm farther. Amplification behaviour is determined by a Hermite spline with three configurable control points representing normalized units of reach between P_N and P_{max} : CP_0 at the user’s neutral position (P_N), CP_{mid} in the middle of the user’s reach, and CP_{max} at the user’s maximum reach (P_{max}). CP_0 and CP_{mid} determine the amplification’s slope, controlling the intensity of amplification as the user brings their arm forward, while CP_{max} determines the amplification at the user’s maximum reach. CP_0 is fixed at $(0,0)$, meaning that the technique adds no offset at P_N . Figure 3 illustrates the curve configurations used for Experiment 1.

Pilot Study to Guide Curve Shape

We conducted informal, iterative pilot tests with 5 graduate students for feedback on amplification curves with regard to comfort, accuracy, and body ownership. From this we established two design heuristics. First, curves should be kept smooth: discontinuities in the curve or dramatic, instant changes in slope tended to break body ownership almost immediately, as the hand’s position would jump from one point to another rather than with a steady transition. Previous work concurs that gradual amplification is best for body ownership [7].

Second, minimize usage of the area of the curve where the slope is less than one. Early tests revealed that areas where the slope was less than one (meaning the virtual hand moves more slowly than the real hand) prompted unpleasant feelings in users, with a clear loss in body ownership. We mitigate these effects by shifting CP_{max} such that this negative effect

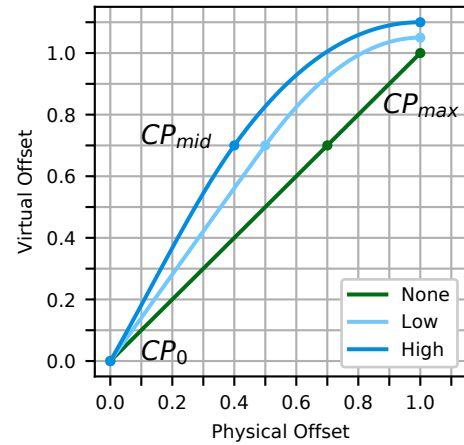


Figure 3. The *Low* and *High* amplification functions used in Experiment 1. These functions modify the relationship between the physical offset r and the virtual offset $f(r)$ from P_N .

only occurs when the virtual controller is beyond the user’s unamplified reach.

Specific Curve Design

Based on this pilot study, we designed two curves with different levels of amplification (Figure 3). These illustrate the general shape of RNL curves, and we use these in Experiment 1. The *Low* amplification curve provides more subtle amplification: $CP_0 = (0,0)$ with slope 1.4, $CP_{mid} = (0.5, 0.7)$ with slope 1.38, and $CP_{max} = (1, 1.05)$ with slope 0. The *High* amplification curve is more aggressive: $CP_0 = (0,0)$ with slope 1.75, $CP_{mid} = (0.4, 0.7)$ with slope 1.4, and $CP_{max} = (1, 1.1)$ with slope 0.

Previous work shows that maintaining body ownership largely depends on the size of the *hand offset*: the separation between the user’s real and virtual hands [17, 27]. Because our design is focused on body ownership, we characterize our non-linear amplification functions by their *maximum hand offset*: the maximum separation between the user’s real hand and virtual hand along their entire range of motion. Because arm lengths differ between people, and for easier reference to our calculation algorithm, we refer to maximum hand offsets by percentage of the user’s arm length. The *High* curve increases a user’s maximum reach by 10% when their arm is fully extended and has a maximum hand offset of 33% of the user’s arm length, around midway through the user’s reach. With a typical arm reach of 61 to 71 cm [24], this represents an increase of maximum reach by 6.1 to 7.1 cm, and a maximum hand offset of 20.1 to 23.4 cm, which is within previously established bounds for maintaining body ownership [17, 27]. Experiment 2 further explores the implications of curve slope and maximum hand offset by testing ten curves covering a range of amplification levels. The details of those curves are provided in that section.

Implementation

The calibration and amplification is implemented in Unity3D to act alongside the SteamVR SDK. The amplification function

control points may be adjusted using Unity3D’s *Animation-Curve* component which natively supports Hermite curves. All code is open source for replication and extension¹.

EXPERIMENT 1 - INITIAL VALIDATION

This experiment aims to show that RNL transfer functions can increase comfort without sacrificing task performance or feelings of body ownership. As such, this experiment is a replication of the study used to evaluate Erg-O [19], considered the state-of-the-art for input amplification techniques that emphasize body ownership. Using a target selection task, we test three amplification functions (none, low, high) with three target layouts (ergonomic, limits of reach, world fixed). Primary measures are trial time, hand movement distance, ergonomics, and user self-reports including body ownership. By replicating the previous experiment, we are able to make a meta-comparison to Erg-O in the discussion to follow.

Participants

We recruited 18 participants (ages 19–29, 13 male, 5 female, 2 left-handed). Participants were recruited by word-of-mouth, and received \$10 for successful completion of the study. 9 participants had at least moderate experience with virtual reality; 11 had at least moderate experience with 3D video games.

Apparatus

Our implementation used a Vive Pro HMD setup powered by an Intel Core i7-7920X CPU and a NVIDIA Titan Xp GPU. The Erg-O experiment calculated ergonomics using a Kinect to track the participant’s shoulders, elbows, hands, and waist. We use Vive Trackers strapped to the participant’s shoulders, elbows, and waist in addition to holding Vive controllers for the same purpose. Participants said they felt the straps holding the trackers in place, but their range of motion was unaffected.

Procedure and Task

Each participant completed a general questionnaire, then we measured their arm span (A) for target placement. In the task, several targets (blue spheres) are presented with two highlighted in green. A trial begins when the participant selects a hand by tapping a virtual indicator placed near their shoulder. The participant then uses the controller in that hand to touch the highlighted spheres in any order. Upon correct selection, the targets return to blue, an audio cue plays, and the next two targets are highlighted. Highlighted targets are randomly chosen, with no pair repeated. The accompanying video demonstrates the task.

Target Layouts

We replicate all three target layouts² used in the Erg-O study (Figure 4). The ERGONOMIC layout places targets in a 5×3 grid, $0.21A$ away from the participant’s torso (where A is the participant’s arm span). The LIMITS layout places 24 targets in a hemisphere $0.44A$ away. The FIXED layout positions targets in world space, independent of the participant. It places targets in two 4×3 grids, spanning across $1.4A \times 0.8A$ with the lowest and highest targets $0.4A$ and $1.2A$ from the floor.

¹<https://github.com/JohannWentzel/RNL-Utilities>

²We use descriptive names for Erg-O study layouts: ERGONOMIC is “Layout 1” in Erg-O, LIMITS is “Layout 2”, FIXED is “Layout 3”.




Figure 4. The ERGONOMIC, LIMITS, and FIXED target layouts used in Experiment 1.

Design

This is a within-subject design with two independent variables: AMPLIFICATION with 3 levels (NONE, LOW, HIGH) and LAYOUT with 3 levels (ERGONOMIC, LIMITS, FIXED). Each participant completed all combinations of LAYOUT and AMPLIFICATION, with the order of AMPLIFICATION determined by a balanced Latin square. The levels of AMPLIFICATION are the two curves described in the previous section, with NONE representing unamplified movement.

Dependent measures are computed from logs. *Time* is the period between touching the first and second highlighted targets³.

Comfort is measured by recording the angles of the participant’s shoulder, elbow, and wrist when selecting a target. These are used to compute a final “Posture Score A” from the RULA ergonomic measurement system [18]. Note that a lower RULA score maps to lower physical effort, meaning higher comfort. *Physical Path Length* (and *Virtual Path Length*) is the ratio of distance travelled by the participant’s physical hand (or virtual hand), divided by the distance between the two highlighted targets.

After each amplification condition, participants self-reported their *comfort*, *ease of reach*, *overstretching*, *sense of control*, and *body ownership* on a scale from 1 to 7. All were in the Erg-O study except body ownership.

In summary: 3 AMPLIFICATIONS \times 3 LAYOUTS \times 30 TRIALS = 270 data points per participant.

Results

For each combination of participant, AMPLIFICATION, and LAYOUT, trials with times, RULA scores, or normalized path lengths more than 3 standard deviations from the mean were excluded as outliers. In total, 197 trials (4.1%) were removed.

In the analysis to follow, we used an AMPLIFICATION \times LAYOUT ANOVA with Holm-Bonferroni corrected post-hoc pairwise t-tests, unless noted otherwise. We verified that sphericity was not violated with any measures.

Time

Time to complete trials was not significantly affected by AMPLIFICATION (Figure 5a), and within each layout, the three amplification levels produced similar times. Residuals for

³Our Time measure is equivalent to TCT in the Erg-O study.

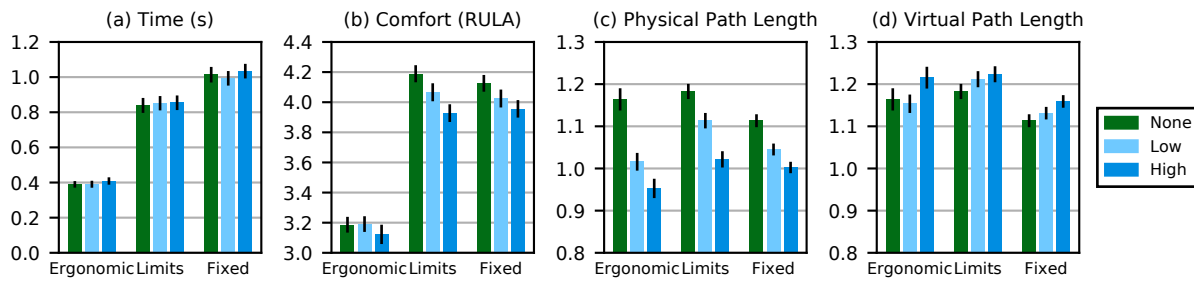


Figure 5. (a) Time, (b) Comfort, (c) Physical Path Length, and (d) Virtual Path Length by AMPLIFICATION and LAYOUT. Error bars are 95% CI.

Time were not normally distributed, so log-transformed values were used for statistical analysis. Although there was a significant main effect of LAYOUT ($F_{2,34} = 273.61$, $p < 0.001$), the more relevant tests for an AMPLIFICATION main effect, or a LAYOUT \times AMPLIFICATION interaction, were not significant.

Comfort based on RULA Score

In the LIMITS and FIXED layouts, both amplification levels had lower RULA scores than the baseline suggesting an increase in comfort (Figure 5b), with HIGH amplification reducing RULA more than LOW. An ANOVA revealed a significant interaction effect of AMPLIFICATION \times LAYOUT ($F_{4,68} = 2.97$, $p < 0.05$), prompting separate post-hoc tests for each LAYOUT. In the LIMITS layout, HIGH and LOW amplifications improved RULA scores by 0.26 and 0.12 over NONE (both $p < 0.01$). In the FIXED layout, HIGH improved by 0.17 ($p < 0.001$) and LOW improved by 0.10 ($p < 0.05$) compared to NONE. We found no significant effects in the ERGONOMIC layout.

Physical and Virtual Path Length

HIGH and LOW amplifications reduced physical path lengths in all layouts (Figure 5c). Pairwise Wilcoxon signed-rank tests revealed significant differences between all AMPLIFICATIONS in every LAYOUT (all $p < .01$). In the ERGONOMIC layout, HIGH (0.95) and LOW (1.02) had shorter paths than NONE (1.16). In the LIMITS layout, HIGH (1.02) and LOW (1.11) had shorter paths than NONE (1.18). In the FIXED layout, HIGH (1.00) and LOW (1.05) were shorter than NONE (1.11). This represents a 10% to 18% decrease in physical path length for HIGH and a 6% to 13% decrease for LOW compared to NONE.

However, HIGH and LOW amplifications also increased virtual path lengths (Figure 5d). Pairwise Wilcoxon signed-rank tests revealed significant differences between all AMPLIFICATIONS in some LAYOUTS. In the ERGONOMIC layout, HIGH (1.22) had longer virtual paths than NONE (1.16) ($p < 0.01$). In the LIMITS layout, HIGH (1.22) and LOW (1.21) had longer virtual paths than NONE (1.18) (both $p < 0.05$). In the FIXED layout, HIGH (1.16) and LOW (1.13) had longer virtual paths than NONE (1.11) (both $p < 0.01$). In the worst cases, these paths are 4.4% and 2.5% longer for HIGH and LOW compared to NONE.

Self-Reports

Questionnaire answers were not strongly affected by AMPLIFICATION. Table 1 shows questionnaire results for each measure. Pairwise Wilcoxon signed-rank tests found no significant differences between these responses for any of the 5 questions.

Question	NONE	LOW	HIGH
Comfort	5.28 \pm 1.49	5.83 \pm 1.15	5.56 \pm 1.10
Ease of Reach	5.44 \pm 1.34	5.83 \pm 0.96	5.78 \pm 1.06
Overstretching	4.44 \pm 1.69	4.22 \pm 1.80	3.72 \pm 1.84
Sense of Control	5.83 \pm 0.71	5.94 \pm 0.94	5.61 \pm 0.98
Body Ownership	6.33 \pm 0.91	5.94 \pm 1.16	5.22 \pm 1.76

Table 1. Mean and stdev for self-report responses by AMPLIFICATION.

Discussion

A primary goal for this experiment was to provide data for a meta-comparison of our technique to Erg-O [19], the state-of-the-art approach. Our results show that RNL performance is generally comparable, providing statistically significant increases in comfort without sacrificing task performance or body ownership.

For Time, our technique did not significantly change trial times. This result is comparable to Erg-O, with one difference: Erg-O's S_R technique significantly improved time in the ERGONOMIC layout. Because our amplification curves aim to provide the user with more precise control in a comfortable range, amplifying only when the user needs to reach outward, in the ERGONOMIC layout the difference between the physical and virtual hands was not large enough for a detectable change in task time. Erg-O provides amplification regardless of arm extension, achieving a significant difference in this task. However, Erg-O's mean task times were higher in every condition (e.g. 700ms in ERGONOMIC versus our 400ms). This may be due to system or interaction differences between Kinect hand tracking and Vive controller tracking.

For Comfort, both RNL variations improved RULA scores in the LIMITS and FIXED layouts, with RULA improvements of 0.25 and 0.17 for high amplification. This is comparable to Erg-O, in which the ergonomic retargeting (E_R) technique improved RULA scores by 0.26 and 0.24 in all three layouts. However, the Erg-O S_R technique only improved RULA in the FIXED layout. Our amplification technique is more similar to Erg-O's S_R technique, so it makes sense that our effects on RULA are similar. One notable exception is the lack of significant difference in the ERGONOMIC layout with our technique: this is likely due to the same design difference discussed above in Time. Note that even LOW amplification achieved statistically significant results, so it is possible to improve comfort with little amplification.

For *Physical Path Length*, both amplification techniques reduced the distance travelled by the user’s hand in all layouts; Erg-O found no such effects. This verifies that RNL reduces large physical movements.

For *Virtual Path Length*, both amplification curves caused participants to move their virtual hands farther than in the unamplified technique. Again, Erg-O found no differences in virtual path length. Our result might suggest that participants overshoot their targets when their motions were amplified, just as mouse pointer acceleration causes overshooting [6]. However, the small increases in virtual path length (only 4.4% in the worst case) combined with positive self-report measures for *ease* and *control* (all above 5.5) suggest participants were able to keep these motions under control.

For *Self-Reports* in general, we came to conclusions similar to those of Erg-O. Responses regarding *comfort*, *ease*, *overstretching*, and *control* trended toward more positive in the HIGH and LOW amplifications than in NONE. There is one borderline case ($p = .08$): the average *body ownership* rating for HIGH amplification appears 1.11 points lower than NONE. This potential negative trend could suggest that the HIGH amplification may cause offsets that are near recommended maximums [17, 27]. Note the average HIGH ownership rating is 5.22, which is still relatively positive.

EXPERIMENT 2 - AMPLIFICATION LEVELS

Experiment 1 shows that RNL can improve comfort without significantly affecting task performance or body ownership. This experiment further explores RNL tuning and side effects by slowly increasing the amount of amplification, and recording any impact on body ownership, comfort, and task performance. Using a more controlled target selection task, ten amplification levels are tested in increasing strength, with maximum positional offsets from 0 to 45% arm’s reach, increasing by 5% at each level. Primary measures are trial time, error, comfort (RULA), and participant self-reports to assess body ownership, strain, and perception by measuring when, if at all, participants notice the amplification.




Figure 6. Levels of AMPLIFICATION used in Experiment 2. These curves modify the relationship between the physical offset r and the virtual offset $f(r)$ from P_N .




Figure 7. The target layout used in Experiment 2, coloured by DISTANCE, from the side (a) and the front (b). Purple targets are CLOSE, blue are MID, green are FAR. Only two targets are visible at a time (c).

Participants

We recruited 18 participants by word-of-mouth (ages 19–33, 11 male, 7 female, 5 left-handed). 11 participants had at least moderate experience with virtual reality; 14 had at least moderate experience with 3D video games. Each received \$15 after completing the study.

Apparatus

We made two hardware changes from Experiment 1. First, we exchanged the Vive Pro HMD for an Oculus Rift S for added visual clarity. Next, we removed the Vive Trackers, which are impractical for real deployment. Instead, an inverse kinematics model infers the shoulder position required by the RNL amplification method, based on HMD and controller positions. We used a Kinect v2 to measure ergonomics during the experiment. The inverse kinematics model was not used to measure ergonomics, and the Kinect was not used for amplification.

Procedure and Task

The initial procedure was the same as Experiment 1, with a general questionnaire, familiarization with the VR equipment, and calibration of maximum reach and near-body neutral position. Each trial displays two 10 cm diameter spherical targets: a green start target at P_N , and a blue target at an experimentally controlled position. A thin green rendered line connects the targets, to visually guide the user and reduce visual search time. The participant selects the targets in sequence with their dominant hand by placing a 1 cm diameter cursor (mounted in the centre of the top face of the virtual controller) inside each target and pressing the controller trigger button. With each selection, a sound indicates success or an error.

We manipulate the control points of the function’s Hermite curve to construct a set of 10 curves that create maximum hand offsets from 0% to 45% of the user’s arm length (Figure 6). These curves amplify user movement over a greater and more fine-grained range of intensities than Experiment 1, from unamplified movement to extremely dramatic amplification ($CP_0 = (0, 0)$ with slope 2.46, $CP_{mid} = (0.28, 0.7)$ with slope 1.44, $CP_{max} = (1, 1.14)$ with slope 0).

Target Positions

We use targets placed at a range of distances from the user to test performance at various offsets from P_N . There are three groups of five targets each: *close*, *mid*, and *far* (Figure 7). The targets are placed directly up, down, left, right, and forward from the user’s waist (for close targets) or dominant shoulder (for mid and far targets). Close targets are $0.3r_{max}$ away from

their central point, while mid and far targets are $0.6r_{max}$ and $0.85r_{max}$ away respectively.

Design

This is a within-subject design with two independent variables: AMPLIFICATION with 10 levels (0 to 9 in increasing order, creating 0% to 45% maximum hand offset in 5% increments) and DISTANCE with 3 levels (CLOSE, MID, FAR). For each AMPLIFICATION level, participants completed 5 blocks of trials. Each block included all 15 targets (5 per DISTANCE) in a random order. There was an initial extra 5 blocks of practice trials using AMPLIFICATION-0.

Note that we did not randomize the order of the amplification levels. With a random order, participants might perceive large amplification changes when jumping between levels with large differences, but none between more similar levels. These strong order effects would measure participants' "relative" sensitivity to the amplification. Instead, we used a fixed, ascending order to measure "absolute" sensitivity, mirroring previous work on mouse pointing control-display gain [9].

Dependent measures are similar to Experiment 1. *Time* is the period between selecting the start target and second target. *Comfort* uses the same RULA calculation as Experiment 1, but at the moment the second target is selected. *Physical Path Length* (and *Virtual Path Length*) is the same ratio of physical (or virtual) hand path over the distance between the two targets. *Error* is a new measure enabled by the Experiment 2 task: the distance from the user's cursor to the second target's center.

At the end of each AMPLIFICATION, participants answered a nine-item questionnaire. The first question (*Affect*) was "compared to when I started, I felt that this was...", from -3 ("much worse") to +3 ("much better"). Questions 2 to 4 were based on the Ownershift study [7], on a scale from -3 ("strongly disagree") to +3 ("strongly agree"): (*Double Hand*) "I felt like I had more than one of the same hand"; (*Part of Body*) "I felt that the virtual hand was part of my body"; and (*Control*) "I felt I could control the virtual hand as if it were my own". The final 5 questions evaluated the physical strain of various upper body areas on the Borg CR-10 scale [11]. Experimenters recorded participant comments regarding their hand movements, as well as their accompanying AMPLIFICATION levels. At the end of the experiment, participants were informally asked if they noticed any input amplification taking place, and if so, at what point during the experiment they became aware.

In summary: 10 AMPLIFICATIONS \times 5 BLOCKS \times 3 DISTANCES \times 5 target positions = 750 data points per participant.

Results

For each combination of participant, DISTANCE, and AMPLIFICATION, we removed outliers by excluding trials with *Time*, *Comfort*, or *Error* more than 3 standard deviations from the mean. 392 trials (2.9%) were removed.

We used the same analysis as Experiment 1, with AMPLIFICATION and DISTANCE as primary factors. When the assumption of sphericity was violated, degrees of freedom were corrected using Greenhouse-Geisser ($\epsilon < 0.75$) corrections.

Learning Effect

We are interested in practised performance, so we examine if earlier blocks took longer and should be removed. For the five blocks of each amplification level, there is a significant main effect for BLOCK on *Time* ($F_{1,17} = 7.70$, $p < .013$, $\eta_G^2 = .01$) and no interaction effects involving BLOCK. Post hoc tests found blocks 1 and 2 significantly slower than block 5 (both $p < .05$). Unless stated otherwise, in the subsequent analysis, blocks 3 through 5 of each AMPLIFICATION are used since they are more representative of practised performance.

Time

Input amplification made selection time noticeably longer in the MID and FAR distances, starting at AMPLIFICATION-6 and AMPLIFICATION-8 respectively (Figure 8a). There was no significant effect of AMPLIFICATION on *Time* for targets at the CLOSE distance, but significant effects on *Time* for both the MID distance ($F_{9,153} = 20.16$, $p < .01$, $\eta_G^2 = .22$) and the FAR distance ($F_{9,153} = 8.94$, $p < .01$, $\eta_G^2 = .09$). Due to significant interaction effect of AMPLIFICATION \times DISTANCE ($F_{18,306} = 12.45$, $p < .01$) we conducted separate post-hoc tests for each DISTANCE. Post-hoc pairwise tests revealed that in the MID distance, participants were unaffected by amplification until AMPLIFICATION-6 (max offset of 30% arm's length), at which point they would become significantly slower than in AMPLIFICATION-0. Post-hoc tests revealed similar behaviour in the FAR distance, with a significant time increase occurring at AMPLIFICATION-8 (max offset of 40% arm's length). This represents time increases of 19% for MID (865 ms to 1031 ms) and 12.6% for FAR (1027 ms to 1157 ms) targets.

Comfort

Input amplification made the selection task more comfortable at all target distances, starting at different AMPLIFICATION levels for each DISTANCE (Figure 8b). There were significant effects of AMPLIFICATION and DISTANCE on *Comfort*. Effects were determined using pairwise Wilcoxon signed-rank tests for each DISTANCE. *Comfort* was unaffected for the CLOSE targets until AMPLIFICATION-5 (max offset of 25% arm's reach), after which point RULA scores significantly decreased (all $p < .01$). There were similar effects for the MID targets starting at AMPLIFICATION-4 (max offset of 20% arm's reach) onward (all $p < .05$), and for the FAR targets at AMPLIFICATION-2 (max offset of 10% arm's reach) onward (all $p < .01$). To illustrate, AMPLIFICATION-6 reduces RULA scores by 6.4% for CLOSE targets (2.96 to 2.78), 7.1% for MID targets (3.80 to 3.53), and 10.0% for FAR targets (4.18 to 3.76).

Error

Generally, selection error was unaffected by AMPLIFICATION or DISTANCE (Figure 8c). Pairwise Wilcoxon signed-rank tests found no effects of AMPLIFICATION or DISTANCE on *Error*, with two exceptions: AMPLIFICATION-2 (max offset of 10% arm's reach) for the MID targets ($p < .05$), and AMPLIFICATION-8 (max offset of 40% arm's reach) for the FAR targets ($p < .05$).

Physical and Virtual Path Length

Input amplification resulted in noticeably less movement at all target distances (Figure 8d). Pairwise Wilcoxon signed-rank tests found that AMPLIFICATION had a significant effect at different levels depending on DISTANCE: AMPLIFICATION-1

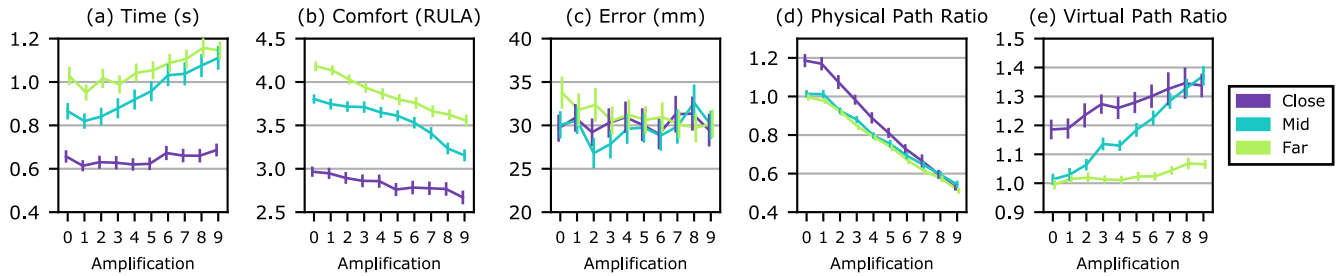


Figure 8. (a) Time, (b) Comfort, (c) Error, (d) Physical Path Length, and (e) Virtual Path Length by AMPLIFICATION. Error bars are 95% CI.

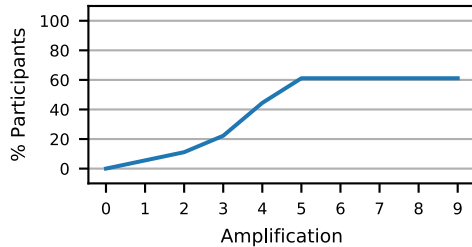


Figure 9. The proportion of participants who noticed that amplification was taking place at each level.

(max offset of 5% arm’s length) for CLOSE and FAR targets (all $p < .05$), and AMPLIFICATION-2 (max offset of 10% arm’s length) for MID targets (all $p < .01$). This effect represents a reduction of physical movement by 9.3% for close targets by applying just a 10% maximum hand offset.

However, input amplification also increased virtual path lengths at all distances (Figure 8e). Pairwise Wilcoxon signed-rank tests found that AMPLIFICATION had a significant effect (compared to AMPLIFICATION-0) at different levels depending on DISTANCE: AMPLIFICATION-3 (max offset of 15% arm’s length) for CLOSE targets (all $p < .01$), AMPLIFICATION-2 (max offset of 10% arm’s length) for MID targets (all $p < .01$), and AMPLIFICATION-5 (max offset of 25% arm’s length) for FAR targets (all $p < .05$). This represents an increase of virtual hand movement by 5% at mid-range targets by applying a max hand offset of 10% max reach.

Perception of Amplification

Participants noticed that amplification was taking place at different points during the experiment (Figure 9). 4 of the 18 participants (22%) noticed it by AMPLIFICATION-3; an additional 7 (39%) noticed by AMPLIFICATION-5; and the remaining 7 (39%) never noticed the amplification. Further, participants with more previous experience with video games noticed the amplification earlier. We found a Spearman’s rank correlation of -0.48 ($p < .05$) between familiarity with video games (on a scale of 1 to 7) and the AMPLIFICATION level at which the participant noticed input amplification taking place.

Questionnaire Responses

Input amplification did not strongly change participants’ affect, body ownership, or strain. Figure 10 shows responses for questions 1 to 4. Pairwise Wilcoxon signed-rank tests showed no significant effects of AMPLIFICATION on self-report answers,

including those for strain and body ownership. The Borg CR-10 questions also showed no significant effect of AMPLIFICATION. Mean values are 2.08 ± 2.07 for the neck, 1.52 ± 1.93 for the forearm, 1.26 ± 1.68 for the hand, 1.98 ± 2.11 for the shoulder, and 1.76 ± 2.05 for the upper arm.

Discussion

This experiment further examines the effects of various properties of RNL functions, to provide a broader understanding of non-linear 3D transfer functions with respect to body ownership, comfort, and task performance. These implications also reveal characteristics that apply to VR input more broadly.

The results for *Time* show that amplification strength and target distance both strongly influence usability. RNL functions are designed to minimize hand offset at positions near the body, reach maximum hand offset midway through the user’s reach, and reduce offset again near maximum reach. The results reflect this design. Close targets saw little change due to hand offset being minimized close to the body. Mid-reach targets were most affected due to being near the area of maximum offset (resulting in the least precise input), and far targets were less affected due to the declining offset at that level of reach.

Error had generally minimal effects of amplification or distance. This is an example of the classic trade-off between speed and accuracy in selection tasks [35]. Participants in this case optimized for speed over accuracy. While there were two small significant differences at single levels of amplification (level 2 for mid, level 8 for far), this was likely a side-effect of participant numbers and not indicative of a larger trend.

Amplification level had a dramatic effect on the user’s *Comfort*. As specified by the RNL function design, amplification benefits grow as targets become further away and less comfortable to reach. User comfort did not improve for close targets until a 25% maximum hand offset, while mid-range and far targets required a 20% and 10% hand offset respectively for comfort to improve. This suggests that if a task requires larger reaching movements, little amplification is necessary to make clear improvements in comfort.

For *Physical Path Length*, even amplifications of as little as 5% of arm’s reach reduced path lengths. This verifies that input amplification is taking place and effectively reducing user movement. Note that even considering the potential for overshooting and correcting motions, the total movement by the participant is still reduced.

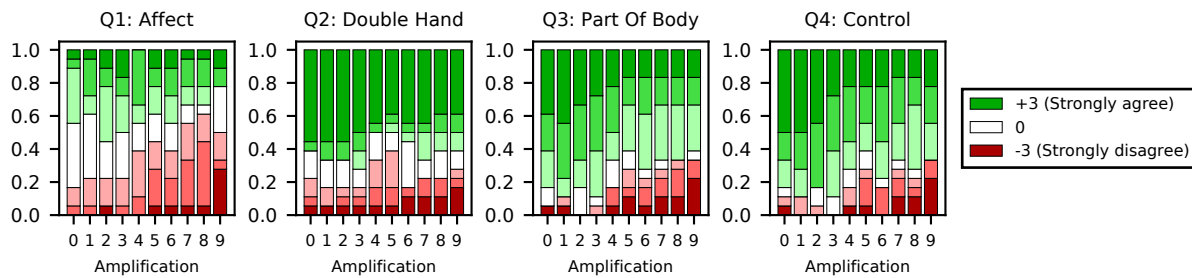


Figure 10. Proportion of questionnaire answers by AMPLIFICATION. Answers were inverted for Q2 for visual comparison.

For *Virtual Path Length*, input amplification caused participants to move their virtual hands more than they would naturally, starting at various levels depending on target distance. Just like with *Time*, mid-range targets were likely less precise due to the increased hand offset at that level of reach.

For *Self-Reports*, while we cannot make strong conclusions without statistically significant results, there may be a possible trend in user preference and body ownership as amplification increases. Mean values for general enjoyment (*affect*) may decline around AMPLIFICATION-4 (max offset of 20% arm's reach), suggesting that the increased amplification made the task less enjoyable. Mean values for body ownership may also decline around AMPLIFICATION-4, suggesting that this level may be the point users begin to lose connection to their virtual hands. Based on an average arm length of 61 to 71 cm [24], a 20% hand offset is 12.2 to 14.2 cm, making this tolerance roughly congruent with body ownership tolerances from previous work [17, 27]. The Borg CR-10 survey also found no significant results, suggesting that RNL functions do not significantly increase the user's physical strain.

We also found a negative correlation between familiarity with 3D video games and the amplification level at which participants noticed amplification taking place. Participants who self-reported higher familiarity with 3D video games noticed amplification taking place earlier than the rest. This could be because of the prevalence of control-display modification in 3D video games, or because video games generally increase awareness of control-display mismatches [16].

Measuring perception of amplification also involves measuring the user's adjustment to changes in amplification levels. In many cases, there was a significant (or borderline not significant) increase in time for the first block of a new amplification level compared to the last block of the previous amplification level. Specifically, levels 4, 5, and 9 were significant (all $p < .05$) and levels 6, 7, and 8 were borderline not significant ($p < .065$). It could be the case that the initial three levels did not require the user to make large enough adjustments to warrant significant learning, and the final level was so far removed from natural movement that learning was much more difficult.

GENERAL DISCUSSION

The two experiments were designed to answer three questions:

1. Does reach-bounded non-linear input amplification perform comparably to the state-of-the-art?

2. How strong can non-linear amplification become before negative effects begin?
3. How do properties like slope, hand offset, and target distance affect the usability of VR transfer functions?

We rephrase the most relevant results as design recommendations, discuss possible limitations in our methods, and discuss future uses for reach-bounded non-linear input amplification.

Design Recommendations

While the RNL family of functions was the main vehicle for testing, our results also provide general design guidelines for VR transfer functions as well as tasks in which they are used.

Consider the Task

Target placement is an important consideration when designing a VR transfer function. Non-linear transfer functions change their control-display ratio dynamically based on the user's reach, requiring users to move at varying rates for equal performance at different distances. However, this means that targets at certain distances may be more difficult to accurately reach than others. We showed that mid-reach selection performance declined the fastest as amplification increased due to these targets being near our function's area of maximum offset. For consistent ease of use, designers should consider the typical reaching distance required in the application task relative to the behaviour of their transfer function.

Another consideration is the task's requirement for user body ownership. Our work concurs with previous work [17, 27], showing that body ownership decreases if hand offset is brought above certain limits. However, this limit (max offset of 20% arm's length) was lower than the limit for task performance (30%). This suggests that tasks that emphasize productivity over body ownership (e.g. 3D modeling) could benefit from further increased comfort if the maximum offset is between 20% and 30%. Designers should consider opportunities for compromise when prioritizing comfort, body ownership, and task productivity.

Consider the Function

Designers should consider three function properties when designing VR 3D input transfer functions: maximum hand offset, slope, and discontinuities.

The maximum hand offset of a transfer function refers to the largest separation between the real and virtual hand when reaching forward. While increasing maximum hand offset will

always increase the user's physical comfort, task performance and body ownership decline when offset surpasses certain levels. Designers should consider that maximum hand offsets larger than 20% (12.2 to 14.4 cm for average arms) may reduce feelings of body ownership and enjoyment, and offsets larger than 30% (18.3 to 21.3 cm) may reduce task performance.

The slope of a transfer function determines the speed of the virtual hand relative to the real hand. Our work shows that a slope of less than one can feel unpleasant, while excessive slope can cause reduced performance and body ownership. For example, task performance in Experiment 2 declined around AMPLIFICATION-6. This RNL function has a maximum slope of 1.92 at $r = 0.14$, meaning that virtual hand speed is nearly doubled at this point, reducing accuracy. Designers implementing non-linear VR transfer functions should consider the benefits and drawbacks of increased or decreased curve slope.

Function discontinuities cause a sudden change in the controller's position or speed, which can immediately break the illusion that the virtual and real hands align. Our approach addressed this with a smooth Hermite curve design that gradually reduces offset when approaching maximum reach. Designers of other functions should consider maintaining body ownership by keeping functions smooth to avoid sudden changes.

Consider the Person

Participants noticed amplification much earlier if they played video games. If a task requires a certain level of body ownership (e.g. a dancing game that relies on kinesthetic feedback) and the typical user has a higher level of experience with video games, the amount of input amplification applied may need to be reduced to maintain body ownership.

Limitations

Continuous Input

A benefit of our approach is that it amplifies movement independently of objects in the VR scene, making it more applicable to general tasks. However, our experimental tasks only cover discrete "point-and-click" input. Future work should focus on determining the usefulness of RNL amplification in situations requiring continuous input, like virtual drawing.

Comparison to Erg-O

Experiment 1 replicates the Erg-O experimental protocol to enable an effective meta-comparison, showing that the RNL approach provides similar results to the state-of-the-art. However, a direct comparison would be possible with an identical implementation and access to Erg-O's participant data.

Perception and Participant Memory

Alongside noting participant comments throughout the experiment, Experiment 2 involved asking if and when they became aware of input amplification. To avoid bias, they were only asked once at the end of the experiment, not after each amplification level. However, this may introduce variance as participants may not immediately comment as they notice a change, and the final perception question may rely on participants' memory of the one-hour experiment session. Further work could explicitly design around this perception question, using tests that require less participant memory.

Controller Models and Ownership

Our experiments use controller models instead of hand models for higher accuracy [5] and ecological validity with current VR applications. This could come at a cost to body ownership [31] however our results demonstrate a good lower bound for body ownership under input amplification.

Future Applications

The RNL method can be easily integrated into existing VR applications or frameworks. However, a typical VR user may not want to spend time calibrating our amplification system for optimal use. We believe a dynamic calibration method could be created, perhaps amplifying user input based on the average amount of reach required in the last few minutes.

Input amplification need not only apply to hand translation. Previous work has studied various linear and non-linear offsets applied to hand rotation [8, 14, 32], head translation [26, 32–34], and head rotation [21, 32]. These techniques could be reproduced or extended using RNL function curves, bounding the function with realistic movement extents similarly to reach-bounded hand amplifications.

Some newer VR headsets track head position without external sensors or beacons by using in-headset cameras and SLAM algorithms [20]. However, hand controllers can only be tracked if they are in view of the headset cameras, and extreme reaches and hand movements can lose tracking and reduce input accuracy. Rather than lose tracking altogether, user input could be amplified to keep the controllers in the headset's tracked area.

CONCLUSION

We described and evaluated an approach we call RNL for amplifying hand movement through easily-configurable Hermite curves. It is simple to implement and potentially applicable to more tasks than the state-of-the-art, but still has comparable impacts on comfort, task performance, and body ownership. Further testing of this technique shows the effects of various aspects of non-linear transfer functions, including target placement, maximum offset, and slope. The insights gained from these tests allow us to make general design recommendations for VR transfer functions.

VR interfaces often overlook users with constrained work spaces, comfort requirements, or mobility issues. As VR matures to include more applications including office productivity, interactive prototyping, or multi-user environments, the challenges of user comfort and small physical spaces will become central to adoption. Our approach allows the easy implementation of input amplification into existing VR applications. The insights gained during this testing process, in conjunction with the development of the RNL system, should help VR become a more comfortable and accessible experience for all.

ACKNOWLEDGEMENTS

This work made possible by NSERC Discovery Grant 2018-05187, the Canada Foundation for Innovation Infrastructure Fund 33151 "Facility for Fully Interactive Physio-digital Spaces," and Ontario Early Researcher Award ER16-12-184.

REFERENCES

- [1] 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 Conference. *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces - ISS '16* (2016), 501–504. DOI: <http://dx.doi.org/10.1145/2992154.2996883>
- [2] Matthew Botvinick and Jonathan Cohen. 1998. Rubber hands 'feel' touch that eyes see. *Nature* 391, 6669 (1998), 756. DOI: <http://dx.doi.org/10.1038/35784>
- [3] Doug A. Bowman and Larry F. Hodges. 2004. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. (2004), 35–38. DOI: <http://dx.doi.org/10.1145/253284.253301>
- [4] Eric Burns, Mary C. Whitton, Sharif Razzaque, M.R. McCallus, A. T. Panter, and Frederick P. Brooks. 2005. The hand is slower than the eye: a quantitative exploration of visual dominance over proprioception. *IEEE Proceedings. VR 2005. Virtual Reality, 2005*. 2005 (2005), 3–10. DOI: <http://dx.doi.org/10.1109/VR.2005.1492747>
- [5] Giuseppe Caggianese, Luigi Gallo, and Pietro Neroni. 2019. *The Vive Controllers vs. Leap Motion for Interactions in Virtual Environments: A Comparative Evaluation*. 24–33. DOI: http://dx.doi.org/10.1007/978-3-319-92231-7_3
- [6] Géry Casiez, Daniel Vogel, Ravin Balakrishnan, and Andy Cockburn. 2008. The impact of control-display gain on user performance in pointing tasks. *Human-Computer Interaction* 23, 3 (2008), 215–250. DOI: <http://dx.doi.org/10.1080/07370020802278163>
- [7] Tiare Feuchtner and Jörg Müller. 2018. Ownershift: Facilitating Overhead Interaction in Virtual Reality with an Ownership-Preserving Hand Space Shift. *The 31st Annual ACM Symposium on User Interface Software and Technology - UIST '18* (2018), 31–43. DOI: <http://dx.doi.org/10.1145/3242587.3242594>
- [8] Scott Frees, G. Drew Kessler, and Edwin Kay. 2007. PRISM interaction for enhancing control in immersive virtual environments. *ACM Transactions on Computer-Human Interaction* 14, 1 (2007), 2–es. DOI: <http://dx.doi.org/10.1145/1229855.1229857>
- [9] Ravin Balakrishnan Géry Casiez, Daniel Vogel and Andy Cockburn. 2008. The Impact of Control-Display Gain on User Performance in Pointing Tasks. *Human-Computer Interaction* 23, 3 (2008), 215–250. DOI: <http://dx.doi.org/10.1080/07370020802278163>
- [10] Dustin T. Han, Mohamed Suhail, and Eric D. Ragan. 2018. Evaluating remapped physical reach for hand interactions with passive haptics in virtual reality. *IEEE Transactions on Visualization and Computer Graphics* 24, 4 (2018), 1467–1476. DOI: <http://dx.doi.org/10.1109/TVCG.2018.2794659>
- [11] Asha Hareendran, Nancy Leidy, Brigitta Monz, Randall Winnette, Karin Becker, and Donald Mahler. 2012. Proposing a standardized method for evaluating patient report of the intensity of dyspnea during exercise testing in COPD. *International journal of chronic obstructive pulmonary disease* 7 (05 2012), 345–55. DOI: <http://dx.doi.org/10.2147/COPD.S29571>
- [12] Brent Insko, Michael J. Meehan, Mary C. Whitton, and Frederick P Brooks Jr. 2001. *Passive Haptics Significantly Enhances Virtual Environments*. Ph.D. Dissertation. University of North Carolina at Chapel Hill.
- [13] Konstantina Kilteni and Raphaela Groten. 2012. The Sense of Embodiment in Virtual Reality. *Presence* 21, 4 (2012), 373–387. DOI: http://dx.doi.org/10.1162/PRES_a_00124
- [14] Joseph J. Laviola and Michael Katzourin. 2007. An Exploration of Non-Isomorphic 3D Rotation in Surround Screen Virtual Environments. In *2007 IEEE Symposium on 3D User Interfaces*. DOI: <http://dx.doi.org/10.1109/3DUI.2007.340774>
- [15] Jialei Li, Isaac Cho, and Zachary Wartell. 2018. Evaluation of Cursor Offset on 3D Selection in VR. (2018), 120–129. DOI: <http://dx.doi.org/10.1145/3267782.3267797>
- [16] Li Li, Rongrong Chen, and Jing Chen. 2016. Playing Action Video Games Improves Visuomotor Control. *Psychological Science* 27, 8 (2016), 1092–1108. DOI: <http://dx.doi.org/10.1177/0956797616650300> PMID: 27485132.
- [17] Donna M. Lloyd. 2007. Spatial limits on referred touch to an alien limb may reflect boundaries of visuo-tactile peripersonal space surrounding the hand. *Brain and Cognition* 64, 1 (2007), 104–109. DOI: <http://dx.doi.org/10.1016/j.bandc.2006.09.013>
- [18] Lynn McAtamney and E Nigel Corlett. 1993. RULA: A survey method for the investigation of work-related upper limb disorders. *Applied Ergonomics* 24, 2 (1993), 91–99. DOI: [http://dx.doi.org/10.1016/0003-6870\(93\)90080-5](http://dx.doi.org/10.1016/0003-6870(93)90080-5)
- [19] Roberto A. Montano Murillo, Sriram Subramanian, and Diego Martinez Plasencia. 2017. Erg-O: Ergonomic Optimization of Immersive Virtual Environments. *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (2017), 759–771. DOI: <http://dx.doi.org/10.1145/3126594.3126605>
- [20] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós. 2015. ORB-SLAM: A Versatile and Accurate Monocular SLAM System. *IEEE Transactions on Robotics* 31, 5 (Oct 2015), 1147–1163. DOI: <http://dx.doi.org/10.1109/TRO.2015.2463671>

- [21] Nahal Norouzi, Luke Bölling, Gerd Bruder, and Greg Welch. 2018. Augmented rotations in virtual reality for users with a reduced range of head movement. *Proceedings of the 12th Annual International Conference on Disability Virtual Reality and Associated Technologies - ICDVRAT '18* (2018), 100–107.
- [22] Eyal Ofek, Christian Holz, Andrew D. Wilson, Hrvoje Benko, and Lung-Pan Cheng. 2017. Sparse Haptic Proxy. (2017), 3718–3728. DOI: <http://dx.doi.org/10.1145/3025453.3025753>
- [23] Mathias Parger, Joerg H. Mueller, Dieter Schmalstieg, and Markus Steinberger. 2018. Human upper-body inverse kinematics for increased embodiment in consumer-grade virtual reality. 1–10. DOI: <http://dx.doi.org/10.1145/3281505.3281529>
- [24] Stephen Pheasant and Christine M. Haslegrave. 2005. *Bodyspace: Anthropometry, Ergonomics and the Design of Work* (3 ed.). Taylor & Francis. <https://books.google.ca/books?id=nwfvSbuyBdMC>
- [25] Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. 1996. The go-go interaction technique: non-linear mapping for direct manipulation in VR. *Proceedings of the 9th Annual ACM symposium on User Interface Software and Technology - UIST '96* (1996), 79–80. DOI: <http://dx.doi.org/10.1145/237091.237102>
- [26] Ivan Poupyrev, Suzanne Weghorst, and Sidney Fels. 2000. Non-isomorphic 3D Rotational Techniques. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. ACM, New York, NY, USA, 540–547. DOI: <http://dx.doi.org/10.1145/332040.332497>
- [27] Michael Rietzler, Florian Geiselhart, Jan Gugenheimer, and Enrico Rukzio. 2018. Breaking the Tracking: Enabling Weight Perception Using Perceivable Tracking Offsets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 128:1–128:12. DOI: <http://dx.doi.org/10.1145/3173574.3173702>
- [28] Majed Samad, Elia Gatti, Anne Hermes, Hrvoje Benko, and Cesare Parise. 2019. Pseudo-Haptic Weight: Changing the Perceived Weight of Virtual Objects By Manipulating Control-Display Ratio. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 320, 13 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300550>
- [29] Mel Slater, Bernhard Spanlang, Maria V. Sanchez-Vives, and Olaf Blanke. 2010. First Person Experience of Body Transfer in Virtual Reality. *PLOS ONE* 5, 5 (05 2010), 1–9. DOI: <http://dx.doi.org/10.1371/journal.pone.0010564>
- [30] 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 2017 - Proceedings*. DOI: <http://dx.doi.org/10.1109/3DUI.2017.7893363>
- [31] Manos Tsakiris, Lewis Carpenter, Dafydd James, and Aikaterini Fotopoulou. 2010. Hands only illusion: multisensory integration elicits sense of ownership for body parts but not for non-corporeal objects. *Experimental Brain Research* 204, 3 (01 Jul 2010), 343–352. DOI: <http://dx.doi.org/10.1007/s00221-009-2039-3>
- [32] WalkinVR. 2019. WalkinVR Driver. (2019). <http://www.walkinvrdriver.com/>
- [33] Betsy Williams, Gayathri Narasimham, Tim P. McNamara, Thomas H. Carr, John J. Rieser, and Bobby Bodenheimer. 2006. Updating Orientation in Large Virtual Environments Using Scaled Translational Gain. In *Proceedings of the 3rd Symposium on Applied Perception in Graphics and Visualization (APGV '06)*. ACM, New York, NY, USA, 21–28. DOI: <http://dx.doi.org/10.1145/1140491.1140495>
- [34] Graham Wilson, Mark McGill, Matthew Jamieson, Julie R. Williamson, and Stephen A. Brewster. 2018. Object Manipulation in Virtual Reality Under Increasing Levels of Translational Gain. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 99, 13 pages. DOI: <http://dx.doi.org/10.1145/3173574.3173673>
- [35] Shumin Zhai, Jing Kong, and Xiangshi Ren. 2004. Speed-accuracy Tradeoff in Fitts' Law Tasks: On the Equivalency of Actual and Nominal Pointing Precision. *Int. J. Hum.-Comput. Stud.* 61, 6 (Dec. 2004), 823–856. DOI: <http://dx.doi.org/10.1016/j.ijhcs.2004.09.007>