Veronica Weser*
Dennis R. Proffitt
Department of Psychology
University of Virginia

Making the Visual Tangible: Substituting Lifting Speed Limits for Object Weight in VR

Abstract

We developed a novel interaction technique that allows virtual reality (VR) users to experience "weight" when hefting virtual, weightless objects. With this technique the perception of weight is evoked via constraints on the speed with which objects can be lifted. When hefted, heavier virtual objects move slower than lighter virtual objects. If lifters move faster than the lifted object, the object will fall. This constraint causes lifters to move slowly when lifting heavy objects. In two studies we showed that the size-weight illusion (SWI) is evoked when this technique is employed. The SWI occurs when two items of identical weight and different size are lifted and the smaller item is perceived as heavier than the larger item. The persistence of this illusion in VR indicates that participants bring their real-world knowledge of the relationship between size and weight to their virtual experience, and suggests that our interaction technique succeeds in making the visible tangible.

1 Introduction

The virtue of virtual reality (VR) is its ability to give users the sense that they are "somewhere else." However, this feeling of immersion in a virtual environment (VE) is, for the most part, exclusively dependent on visual and auditory experience. Many modern VR platforms are incorporating fully tracked controllers into their systems, allowing users to interact with their environments. However, this interaction is still largely limited to changing the visual properties of the scene. To truly achieve presence in a VE, the visual must be made tangible. Here we propose a novel interaction technique for use in VR that is based on principles of pseudo-haptic feedback. This technique allows users to experience the weight of lifted virtual objects by manipulating constraints on virtual object kinematics without altering the display of the user's tracked controllers or requiring any other specialized hardware.

By capitalizing on the dominance of vision over the other senses, pseudo-haptic feedback uses visual information to evoke haptic sensations in interactive VEs (Ernest & Banks, 2002). Researchers have used pseudo-haptic feedback in combination with a variety of novel passive haptic interfaces to create compelling paradigms for interaction with virtual objects. For example, Achibet and colleagues (2014) created a "virtual mitten" using a standard hand-exerciser with elastic properties that allowed users to feel resistance when grasping a virtual object. Coupled with a color change that conveyed information about the

Presence, *Vol. 27, No. 1, Winter 2019*, 1–12
doi:10.1162/PRES_a_00319

*Correspondence to vuw3nb@virginia.edu.

success of the grasp, the virtual mitten succeeded in creating a compelling passive haptic experience that allowed users to initiate precise grasping motions and pseudo-haptically perceive differences in grasp effort and effectiveness. Likewise, by attaching standard exercise bands to the wrist and shoulder of a user, passive haptic devices have also been used to simulate reaching out to press against virtual objects, with pseudo-haptic animation changes to convey different levels of virtual object compliance (Achibet, Girard, Talvas, Marchal, & Lécuyer, 2015). A more complex apparatus that involved 4 strings mounted to the corners of the workspace and intersecting at the interactable object was implemented by Paljic and Coquillart (2004). This device simulated virtual surface "stickiness" when the user moved the physical object attached to the strings; a virtual cube could be dragged across virtual surfaces of varying adhesiveness. They found that although users were above chance at discriminating stickiness using only the visual (pseudo-haptic) information of the cube slowing down during sticky interactions, passive haptic feedback increased performance.

Although effective at conveying additional haptic information in virtual environments, passive feedback devices require custom hardware that may be a barrier for use. Pseudo-haptic feedback has been used in isolation to influence how people perceive the haptic properties of virtual objects using a standard input device, like a mouse. For example, Lécuyer and colleagues (2000) succeeded in simulating the experience of friction in the case of using a mouse to drag a block across a surface in a VE. This was achieved by having the observed block's speed be a function of mouse speed multiplied by a gain. A gain greater than 1 caused the block to move faster than the mouse (suggesting low in friction), whereas a gain of less than 1 caused it to move slower than the mouse (suggesting high friction). In other words, the ratio of the user's displacement of the input device and the displacement of the object on the computer screen was manipulated.

This gain manipulation is an example of the control/display (C/D) ratio, or the relationship between the user's sensorimotor input (how fast the hand moves) and its visual result (how fast the cursor moves). This technique is typical of pseudo-haptic feedback in that a visual spatial component is used in place of haptic information and therefore becomes incorporated into the user's mental model of their haptic interaction with the manipulated object in the VE (Lécuyer, 2009). As another example, when users interact with objects via a cursor in a VE, virtual object mass can be simulated by decreasing the C/D ratio for lighter objects so that the object moves faster than the user's hand moves the mouse (Dominjon, Lécuyer, Burkhardt, Richard & Richir, 2005), or in theory by increasing the C/D ratio for heavier objects.

In order to study how virtual mass is perceived, it is important to understand how physical weight judgments are made in real-world conditions. When tasked with comparing the weights of two equally-sized objects, lifters tend to apply the same force to the second object as was required to lift the first (Buckingham & Goodale, 2010b). This principle allowed for the systematic study of pseudo-haptic presentation of virtual object mass by Dominjon et al. (2005), who presented pairs of size-matched virtual objects to be hefted and compared. Lifters in this study interacted with the objects using a force-feedback device (a PHANToM) that applied an identical amount of pressure to the user's hand. All that differed between virtual objects was the C/D ratio with which the user's hand movements applied to the force-feedback device were translated to the 2D display of the VE. If the C/D ratio was less than 1 for the second lifted object, then the application of the same amount of energy by the lifter resulted in a different visual outcome for the second object: the velocity of the second item lifted was greater than the velocity of the first item lifted. Due to visual capture, this conflict between the visual information and the haptic information is resolved in favor of vision: the user's weight estimates were based entirely on the visual C/D ratio, rather than on the (identical) haptic feedback provided by the interface. Thus, objects that moved on screen with a C/D ratio less than 1 were always thought to weigh less than objects with an unmanipulated C/D ratio (Dominjon et al., 2005). Similarly, Taima and colleagues (2014) used mixed-reality to manipulate the early lift phase of object interaction: the user's hand and the lifted object appeared to move more when the object was intended to feel light, and less when the object was intended to feel heavy. Crucially, they compensated for the misalignment between the user's real hand position and the displayed image in the late lift phase of the object. This insured that a sense of ownership over the user's hand was retained. The C/D ratio concept was also used to convey a sense of object weight by altering the animations of fully-tracked, third-person avatars, such that the animation for lifting a heavy object was slowed and for light objects, it was sped up (Jáuregui, Argelaguet, Olivier, Marchal, Multon, & Lecuyer, 2014).

However, if the C/D ratio of a user's fully tracked controllers in an immersive VR environment is altered, the misalignment between real and virtual hand position is likely to result in discomfort, disorientation, and a decrease in feelings of ownership for the controller. Users of commercial VR systems have come to expect accurate tracking of their hands in space, and rely on this feature during object interaction. By focusing only on modifying the behavior of virtual objects, this confoundment can be avoided. As a result, this interaction technique is only appropriate for conveying the weight of virtual objects lifted with tracked controllers, and does not convey a sense of the weight of the controller itself.

In addition to respecting expectations that controllers retain accurate tracking, conventions established in current VR games and other interactive programs dictate that objects with which the user is currently interacting should remain in close visual proximity to the virtual representation of the user's hand for the duration of the interaction. Although effective, Taima et al.'s (2014) method of altering the C/D ratio of the user's hand and lifted object is more computationally expensive than altering the movement parameters of the lifted object alone. However, setting the velocity of an object relative to the user's movements can result in a separation between the object and the user's controller, which would decrease the user's sense of control over the object. One method of using the C/D ratio and retaining a connection between user and lifted object is the implementation of a virtual tether that stretches between the two, as in Ban and Ujitoko (2018). Although a tether was found to be a useful metaphor for conveying pseudo-haptic information about virtual object weight during lifting and dragging interactions on a touch screen (Ban & Ujitoko, 2018), another option is to implement additional constraints on object movement that prevent an accumulation of distance between the controller and the object in VR.

For these reasons, we decided to make the maximum speed a virtual object could move proportionate to its weight, with a very high cap on light-weight objects and an increasingly lower speed limit on heavier objects. Should the distance between object and controller exceed a threshold, the object would drop out of the participant's grasp and fall to the ground with simulated gravity. It is important to note that the virtual controllers were always seen to move accurately. Although forcing a user to drop an object during a lift is inconsistent with real-world lifting behavior, the addition of the dropping criteria reinforced changes in user lifting movement to reflect different virtual weights: users learned to move slowly when lifting heavy objects and that they could move freely while lifting light objects. This manipulation thus parallels the experience of real-world lifting, in which a heaver object, like a bowling ball, must be moved at slower speeds and with more care than a lighter object, like a ping-pong ball. Indeed, when watching the lifting behavior of others in order to judge the weight of an object, artificially lengthening the duration of the time required to lift the object creates a proportionate increase in the weight estimates made by the observer (Hamilton, Joyce, Flanagan, Frith, & Wolpert, 2007).

Given what is known about human weight perception in the physical world, we sought to test whether our method of setting speed caps on hefted objects could successfully translate the feeling of weight to virtual objects. In the real world, numerical weight estimates are best made by comparing an object of unknown weight to one with a weight that is known. In our study, we provided a range of objects with known virtual weights (predetermined speed caps) that spanned the full range of virtual weights users would encounter. Users were tasked with lifting an unknown box and comparing it to the boxes of known weight. The first aim of our study was to assess whether our method of weight instantiation in the virtual world would yield reliably accurate weight estimates.

A further aim was to see whether other factors that influence weight perception in the real world, such as object size, would also alter how the virtual weight was perceived by our users. Thus we always presented two objects of unknown weight simultaneously: a large object and a small object. Unbeknownst to our users, these objects frequently had the same weight. If our method of suggesting object weight through object speed is successful, then weight estimation biases that occur in the real world for large and small objects should also occur in users' experience of the weights of virtual large and small objects. The final aim of our study was therefore to assess whether our method of virtual weight was a good match for real-world weight by testing to see if users experienced the SWI.

In the real world, expectations about the weight of objects formed through past experiences lifting similar objects allow lifters to determine how much grip and load force ought to be applied before initiating the lift. Thus, large items are always lifted with a greater amount of grip and load force than a small objects of the same material (see Johansson & Flanagan 2009, for a detailed review). Following a lift, the weight of the object is then judged relative to one's initial expectations about how heavy the object was expected to be (Ross, 1969). Under contrived laboratory settings in which the weights of a large and small object of the same material have been manipulated to be identical, these assumptions result in a substantial overestimation of the weight of the small object and an underestimation of the weight of the large object. This illusion, known as the SWI, was first documented over 125 years ago (Charpentier, 1891).

Since then, numerous researchers have recreated the SWI under a variety of different circumstances and concluded that the visual experience of the object's size is the driving factor behind estimates of its weight. In one particularly compelling demonstration, participants were shown either a large or a small cube and then were blindfolded before lifting the cube by a small handle. Unbeknownst to the participant, the cube they had just viewed was swapped prior to their lift with a medium-sized cube with an identical handle. Believing they were lifting the object they had just seen, participants reported that the weight of the cube varied as a function of the size of the cube they had just viewed (Buckingham & Goodale, 2010a).

The visual bias of object weight estimates occurs because people expect small objects to be lighter than big ones made of the same material. Prior experiences with similar looking but differently sized objects have invariably had this outcome. However, when lifting these illusory small and large objects in succession, the lifter will find that the larger does not outweigh the smaller object. This contrast between the expected and perceived weight is thought to drive the SWI: the lifter will therefore judge the large object to be lighter than the small one, and in Buckingham and Goodale's (2010a) study, seeing the large object made the standard cube feel lighter than it did after seeing the small cube.

Since the expectations of an object's weight are influenced by its visual appearance, users in our study who associate movement constraints with virtual object weight should then learn to initiate interactions with large objects at slower speeds than small objects. If this method of pseudo-haptically simulating virtual object weight is successful, then a virtual version of the SWI should occur: users should underestimate the weight of large objects, having moved slowly during the interaction, relative to the weight of small objects lifted with faster movements.

2 Experiment 1: Manipulating Virtual Weight

We found that not only were users able to compare the weights of unknown virtual items to known virtual weights and make accurate estimates, but also that the SWI effect applied even in the absence of any actual weight lifting. On the whole, the weights of small virtual objects were overestimated relative to the weights of large virtual objects, as lifters expectations about small things being light and large things being heavy were violated whenever they lifted two virtual objects set to the same weight.

2.1 Methods

2.1.1 Participants. Thirty-four individuals (11 males, 23 females, mean ± SD age: 19.70 ± 1.34 years, range 18–23) with normal or corrected-to-normal vision took part in the study in exchange for credit in an introductory psychology class at the University of Virginia. All participants provided written informed consent. Data from one individual was excluded for failure to follow experimenter instructions.

2.1.2 Materials. Hardware and Equipment. The head-mounted display (HMD) used was an HTC Vive with a resolution of 1200 × 1080 per eye and a 90-Hz display refresh rate. The HTC Vive has an estimated end-to-end system latency of 22 ms (Niehorster, Li, & Lappe, 2017). The virtual environment was programmed using the Unity game engine on a PC with graphics card GeForce GTX970 and made use of "room scale VR" with a tracked environment approximately 2 m × 4 m through which participants could move freely.

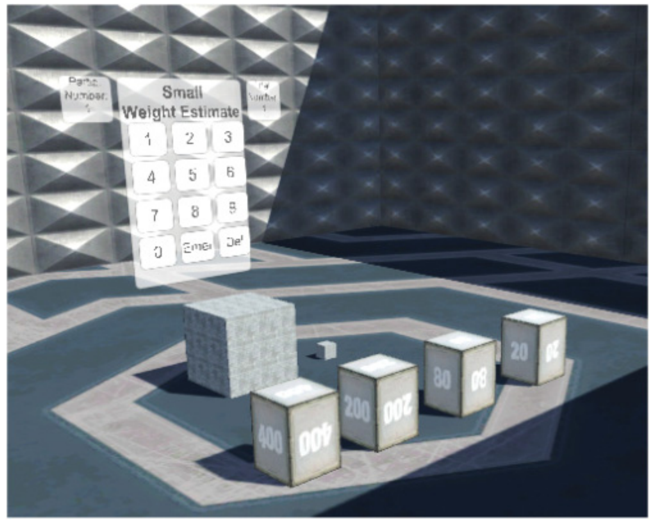

Figure 1. *The virtual environment used in the study.*

Virtual World. The virtual environment was a 10 m × 10 m square room with 5 m high walls and no celling; the default Unity skybox was visible above with the sun as the sole light source. The floor had a realistic stone texture with a geometric pattern achieved by alternating light and dark tiles, and the walls had a realistic diamond plate texture, similar to the metal plating used for industrial ramps (see Figure 1). Participants used the two Vive controllers (one in each hand) to interact with objects in the virtual environment. The fully tracked controllers were visible to participants throughout the course of the study, and participants could pick up objects by inserting the controller inside the object and squeezing the trigger button on the back of the controller.

Virtual "Weight." A user never feels the weight of a virtual object in a typical VR experience, and the speed they can move that object is determined only by how fast they can swing their controller. However, for the purpose of this study, we set upper limits on the speed with which objects could move during interaction in order to simulate the experience of different weights. Object weight was transformed into maximum object speed. Light objects could be lifted quickly, whereas heavy objects could not. If the controller moved at a speed that exceeded the limit set for a given weight, then the object would begin to lag behind the controller. At the point when the controller and the object were visibly separated, the controller would disengage from the object and the object would fall back to the ground. Thus, weight was experienced as how fast an object moved relative to the controller during a successful lift.

This transformation of weight into maximum speed was implemented as follows: During object interaction, the Unity 3D game engine updated the position of the object on a vector in a straight line towards the center of the controller. This value was then multiplied by a speed of 200 m/s and divided by a unitless number representing the object's mass. This resulted in slower speeds for objects with larger masses. For example, the heaviest objects with a mass of 400 moved at a maximum speed of 0.5 m/s. If the distance between the controller and the outer edge of the object exceeded 1 cm, that is, if the participant moved his or her controller faster than the maximum speed allowed for the object, the object and controller became disengaged and the object fell to the ground. By stipulating that the object only fell when it was no longer in direct contact with the controller,

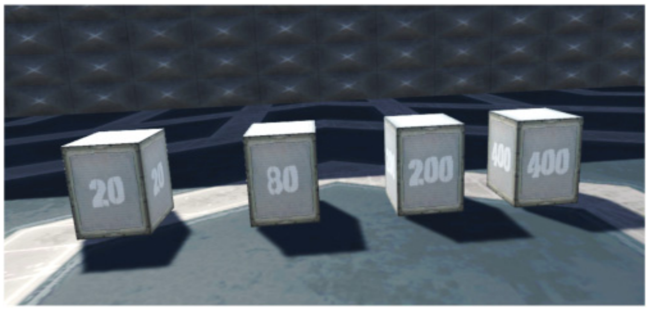

Figure 2. *Weight reference items. The weights were 20, 80, 200, and 400 units, and were programmatically realized by setting the mass of the rigid body of each virtual object to the corresponding amount.*

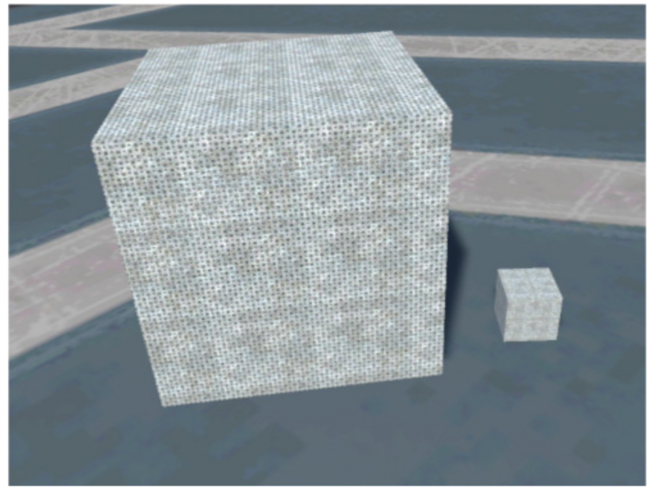

Figure 3. *Weight estimation cubes. The smallest possible size for the small cube was 10 cm^3 (right), and the largest possible size for the large cube was 50 cm^3 (left). The weight of both ranged from 20 to 400 units.*

users were able to adjust their movement speeds in order to better match the speed limits imposed on the objects. Basing the drop criteria on distance between controller and object boundary ensured that objects did not drop while still in contact with the controller, which preserved the user's sense of control over the object during interaction.

Weight References. Users learned the correspondence between maximum speed and weight by lifting 4 virtual reference boxes that appeared to be 40 cm tall and 30 $\times$ 30 cm wide. Each had a different mass that was clearly printed in block letters on the surface of the object (see Figure 2). A separate script attached to the Vive controllers assured that Unity's default gravity would be applied to any object that ceased to be in contact with a Vive controller. Moreover, the physics engine took the programmed object mass into account during object falling, which allowed inertia to effect each object appropriately in accordance with its weight. The masses of the weight references were 20, 80, 200, and 400 units, and were programmatically realized by setting the mass of each virtual object to the corresponding amount. These values were chosen in an informal pilot study in which users were asked to sample a range of weights and pick those best suited to this interaction technique. At the heaviest end, 400 was selected because it was the slowest a user could comfortably move to manipulate an object, while 20 was chosen because it was noticeably different from 0, but still allowed the user to move as fast as desirable during object interaction. The middle values, 80 and 200, were likewise differentiable from 20 and 400, respectively, as well as distinct from one another.

Weight Estimation Cubes. Weight estimation cubes appeared in pairs and varied randomly in size (see Figure 3). The smallest possible size for the small cube was 10 cm^3, and the largest possible size for the large cube was 50 cm^3. On each trial, a value for the dimensions of the small cube was randomly generated between 10 and 30 cm, and then the large cube was generated by adding an amount between 5 and 20 cm to that amount. The mass of the cubes was programmed as above and pseudo-randomly generated so as to sample twice from five equal-sized bins in the range of 20 to 400 units.

2.2 Procedure

Participants were first informed that they would be completing a study examining weight perception in virtual environments. Participants were instructed to lift the weight-reference boxes one at a time with the dominant hand to learn to identify virtual weight values prior to the start of Trial 1. At the beginning of each trial, participants were told to lift each weight estimation cube with the dominant hand and compare the weights of both. The order of lifting (small vs. large) was not prescribed. Thereafter, participants were encouraged to alternate between lifting a reference box and a cube in order to make their estimates as accurate as possible. Aside from the instructions to always lift first with the dominant hand, no other limitations were imposed. We allowed participants to interact freely with the objects to ensure that our procedure would not limit the strategies participants might employ to complete the task. Given our aim of verifying the usefulness of a new interaction technique, we wanted to assess how participants would perceive the virtual weights without imposing procedural constraints. Most participants employed a variety of strategies for comparing the weight estimation cubes and the reference weights. These strategies typically revolved around lifting whatever virtual object was convenient and tended to vary dramatically from trial to trial within a single participant. Many participants chose to lift both a weight estimation cube and a reference box at the same time, one in each hand, to compare the movement of the items. Since no consistent patterns of lifting were observed within or across participants, any findings were deemed unlikely to be the sole product of a lifting order effect.

In general, all participants tended to start each lift slowly and then speed up until it became apparent that the controller was moving away from the boundaries of the object. Participants would then pause and wait for the object to "catch up." Indeed, the time it took for the object to catch up to the controller seemed to provide a metric for the majority of participants who adopted a "shaking" strategy. These participants lifted the objects slowly and then attempted to shake the object as fast as possible while monitoring the speed of the object as it moved back into alignment with the controller. Another strategy frequently employed involved repeatedly lifting and dropping the objects in order to monitor how the physics engine rendered the fall. Several participants commented that the fall assisted in determining whether the object was "light" (<200) or "heavy" (>200), but that watching the way the object trailed behind the controller during movement (such as shaking) provided information that was easier to transform into the numeric scale required for weight estimation.

Figure 4. *Weight input panel. Participants entered their weight estimations directly into the panel by moving a cursor connected to their Vive controllers and squeezing the trigger button when the cursor hovered over a number on the panel.*

After freely interacting with each object to their satisfaction, participants entered their weight estimations directly into a number panel present in the virtual environment (see Figure 4) by moving a cursor connected to their Vive controllers and squeezing the trigger button when the cursor hovered over a number on the panel. The order in which they entered their estimates (small vs. large) was alternated on each trial. Participants completed 10 trials in which they lifted and estimated the weight of 2 cubes per trial. On 8 trials, the mass of both cubes were identical. On 1 trial, the smaller of the cubes had a greater mass than the larger, and on the other mismatched weight trial, it was the larger cube that had the higher mass.

2.3 Results

2.3.1 Weight Estimate Accuracy. Across all trials and cube sizes, participant estimates of virtual object mass and the programmed object mass were strongly positively correlated, $r\,(583) = .60$, $p < .001$, 95% CI [0.55, 0.65], indicating that participants were accurate in their perception of the weight of the virtual objects.

2.3.2 Size-Weight Illusion in VR. Using R (R Core Team, 2017) and the *lmer*( ) function in the lme4 library (Bates, Maechler, Bolker, & Walker, 2014), we fit a linear mixed effect model to predict participant weight estimates from the true weight of each cube, the cube size (whether it was the larger or smaller cube of the pair), a random effect of participant, and a second random effect of trial. The random effects parameters were included to account for repeated measures of the same participant across different trials for cubes of varying size. All main effects and interactions were significant. A significant main effect of true weight (Wald chi-square(1) = 317.08, $p < 0.001$), indicated that the participant estimates of weight increased with programmed representations of weight. Participants estimated the weight of the small cube ($M = 1.20$, $SD = 0.86$) to be significantly heavier than the large cube ($M = 1.16$, $SD = 1.52$) across trials: Wald chi-square(1) = 26.77, $p < 0.001$. Finally, the significant interaction between true weight and cube size (Wald chi-square(1) = 9.78, $p = 0.002$), suggests that the heavier estimates for the small cube of the pair were greatest when the true weights of the pair of cubes were sampled from the heaviest possible range of weights (e.g., 300–400 units) Figure 5 graphs this interaction.

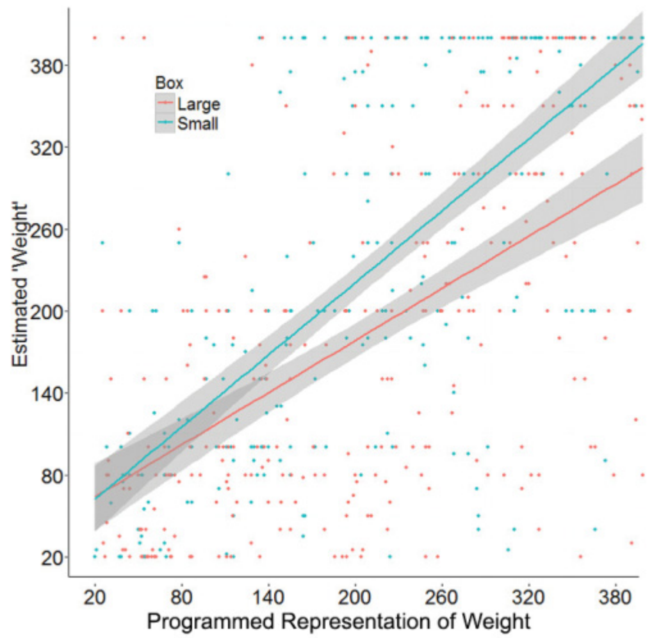

Figure 5. *Participant estimates of weight increased with programmed representations of weight. The small cube was estimated to be significantly heavier than the large cube across trials. Dots represent single cube weight estimates, while shaded areas signify* ± *1* SEM.

3 Experiment 2: Manipulating Virtual Weight

Although users in our study made accurate estimates of the virtual weights of the objects with which they interacted, the quantitative results cannot speak to whether lifting these virtual objects actually "feels" like hefting objects with physical weight. Anecdotally, participants stated that they could feel a difference between heavier and lighter objects, and that the practice of moving slowly to lift something heavy reminded them of real-world object interaction. Moreover, the SWI was replicated using only virtual items, which further suggests that something similar to weight perception is at play. However, the presence of the SWI-like results in Experiment 1 might also be due to the fact that the smallest of the pair of objects was simply more difficult to lift than the larger, given the constraints we imposed programmatically.

In order to produce numerical weight estimates, most participants in Experiment 1 compared the time it took a held cube to catch up to the controller after a small, rapid movement of their grasping hand to the time it took a reference block after a similar movement. The movements made when holding a large cube could be larger, as there was more space for the controller to move within the confines of the large object before the controller and the cube separated. Since objects dropped upon leaving contact with the Vive controller, and the small cube has less volume and surface area, it may have been more challenging for participants to keep their controller in contact with the smaller cube while making use of the prevailing weight estimation strategy described above. This increased challenge could then have been interpreted by participants as an increase in object weight.

Experiment 2 was conducted in order to control for this potential confound specifically as it applies to detecting the presence of a SWI by restricting user interaction to handles. By having users pick up objects by their handles, we were able to ensure that the surface area with which users had to remain in contact was consistent across all objects, regardless of their size. We chose to use handles instead of restricting the interaction to the object center based on the outcome of a pilot study. In our pilot study, participants expressed irritation when the object they were holding fell even when their controller was still in contact with the outer edge of the object. Thus although restricting interaction to the center of the object would have been programmatically easier to execute, the end result was not user friendly. In contrast, the handle is a familiar metaphor for interaction, and participants could clearly tell when they had moved too fast for the object to keep up, allowing a gap to form between handle and controller.

In addition to the increased difficulty of small-object interaction potentially influencing the outcome of the SWI measurement, there was also considerable variability in the weight estimates across participants. One reason for this variability is the inconsistency of object-falling behavior, as larger objects were less likely to fall than smaller object during object shaking. We correctly hypothesized that creating an equalized interaction space for all objects would reduce the noise introduced by unequal-sized objects.

3.1 Methods

3.1.1 Participants. Fifty-five individuals (14 males, 41 females, mean ± SD age: 22.81 ± 3.28 years, range 18–32) with normal or corrected-to-normal vision took part in the study in exchange for $5. All participants provided written informed consent. Participants were recruited using psychology department mailing lists, which included both graduate and undergraduate participants, some of whom were familiar with the paradigm. Therefore, data from four participants who guessed that the study was explicitly designed to replicate the SWI were excluded from analysis. Data from two more individuals were further excluded for failure to follow experimenter instructions, leaving a final sample of 49 participants (13 males, 36 females).

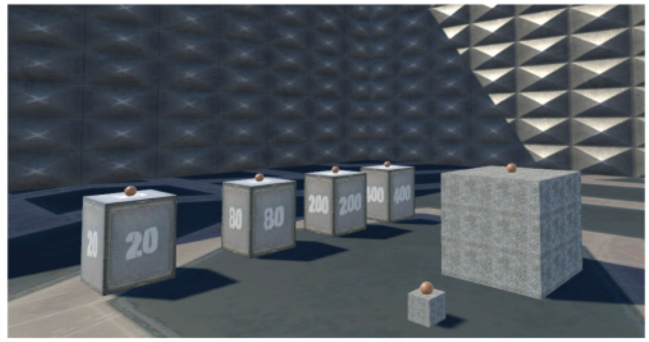

Figure 6. *Weight reference items and estimation cubes fitted with identical handles for Experiment 2.*

3.1.2 Materials. Hardware, Equipment, and Virtual World. The same hardware, physical location, and virtual world used in Experiment 1 were again used in Experiment 2.

Weight References and Estimation Cubes. The same four weight estimation boxes used in Experiment 1 were used in Experiment 2, and the weight estimation cubes followed the same random size and weight parameters as before. The only difference was that in Experiment 2, all cubes could only be lifted by a small, knob-shaped handle centered on the top of each cube (see Figure 6). No part of the reference boxes or weight estimation cubes could be lifted or interacted with, save for the handle.

3.1.3 Procedure. The procedure for Experiment 2 was identical to Experiment 1, save for an additional explanation of the handles, emphasizing that it was only this part of each item that was available for interaction. Following the completion of the experiment, we also asked participants if interacting with the virtual items in the study felt like lifting objects of different weights. If participants said “no,” then we asked them to compare as best they could the experience of the virtual interaction to something they may have experienced in the real world.

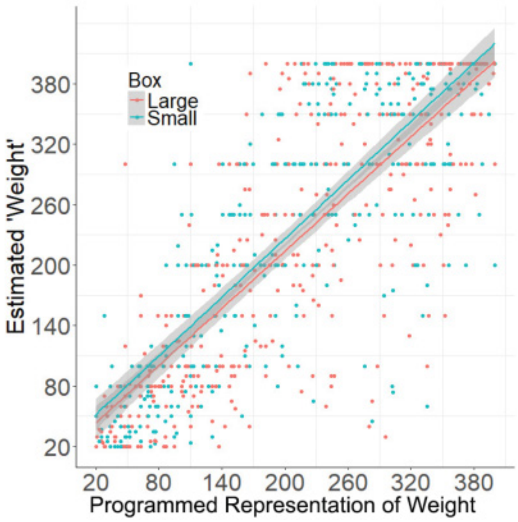

Figure 7. *Participant estimates of weight increased with programmed representations of weight in Experiment 2, but there was no interaction between cube size and weight. Dots represent single cube weight estimates, while shaded areas signify* ± *1* SEM.

3.2 Results

3.2.1 Debrief Question. During the debrief, participants were asked whether interacting with the virtual items in the study felt like lifting items with different weights. Eighty-nine percent of participants said that the experience was most similar to lifting objects of different weights, significantly more often than would be expected by chance, exact binomial p (one-tailed) <0.001 (95% CI [0.78, 0.97]). Of the five participants who said it felt like something other than weight, two participants said it felt like slipperiness of the handles, and one reported the experience was like pulling objects underwater. The final two participants could not articulate how lifting objects compared to anything they had experienced in the real world.

3.2.2 Weight Estimate Accuracy. As in Experiment 1, participant estimates of virtual object mass and the programmed object mass were strongly positively correlated, $r\,(797) = .79$, $p < .001$, 95% CI [0.77, 0.82], indicating that participants were accurate in their perception of the weight of the virtual objects. Figure 7 plots the correlation separately for large and small objects.

3.2.3 Size-Weight Illusion in VR. We followed the same procedure as in Experiment 1. The linear mixed effect model predicted participant weight estimates from the true weight of each cube, the cube size (whether it was the larger or smaller cube of the pair), a random effect of participant, and a second random effect of trial. A significant main effect of true weight (Wald chi-square(1) = 1471.33, $p < 0.001$), indicated that the participant estimates of weight increased with programmed representations of weight. Participants estimated the weight of the small cube ($M = 236$, $SD = 133$) to be significantly heavier than the large cube ($M = 224$, $SD = 132$) across all trials: Wald chi-square(1) = 5.66, $p = 0.017$. The interaction between true weight and cube size did not reach significance ($p > 0.5$).

4 Discussion

Across two experiments we demonstrated that users can learn to accurately identify the "weights" of virtual objects in the absence of any hefted physical mass. By presenting weight visually though the manipulation of maximum object speed, VR users can experience a richer, more nuanced form of object interaction that makes the virtual world tangible. Given that this method requires no additional hardware and can easily be implemented using existing VR systems, we believe it has the potential to easily enhance many current and future VR experiences. The high level of accuracy in weight estimation attained by our users further suggests that it is possible to have tight control over the amount of weight users perceive in their virtual interactions. Heavy and light items can be easily distinguished from one another, and adding a wide variety of weights to interactable objects has the potential to radically enhance the experience of presence in VEs.

Although this interaction technique is not suitable for simulating the weight of specialty controllers or other real-world props instantiated in the virtual environment, it can be readily used to give users a sense of contrast when interacting with a “light” virtual item as opposed to a “heavy” virtual item. In both experiments, most participants attempted to lift the objects as fast as possible and slowed down only when the distance between controller and object center became obvious. In Experiment 1, smaller objects were found to be more difficult to grasp, as they provide less surface area and volume, resulting in less time for the lifter to adjust his or her lifting speed before becoming disengaged from the object. This added difficulty when lifting small objects was especially pronounced for small, heavy objects and likely contributed the interaction between cube size and weight in Experiment 1. The addition of handles in Experiment 2 equalized the differences between large- and small-object interaction, reduced the variation in weight estimates, and eliminated the interaction. As such, Experiment 1 best exemplifies the interaction technique itself, while Experiment 2 validates the results and the SWI finding, but is not ideal for future applications because the addition of handles made all objects *equally difficult* to lift. Therefore, this technique is best suited for objects 10 cm^3 or larger that are engaged by inserting the controller into any part of the object, rather than a small surface like a handle. If this technique is used to add weight to irregularly shaped virtual objects, the amount of time the lifter has to adjust his speed before losing contact with the object will change depending on object orientation of that object. The recommendation that weight added to virtual objects is best used with medium- to large-sized objects with sufficient volume and surface area for easy grasping is consistent with basic intuition about what size objects should weigh the most.

In addition to this size recommendation, items weighing over 200 units tended to be dropped more frequently than the lighter items. Some participants expressed frustration, especially when lifting the heaviest (>300) items, so extreme weights should be used sparingly. On the other hand, the accuracy at identifying the difference between the lightest items and the medium items suggests that valuable weight information can be obtained by users even when the full range of weight possibilities is not sampled. Accuracy of weight estimates can likely be increased by implementing additional features such as altering the color of the lifted object or the controller to visually alert the user when the object and the controller are becoming separated (as in Achibet et al., 2014).

Although in Experiment 1 we did not directly assess the subjective feelings of our users as they interacted with our virtual weighted objects, in both experiments the consistent overestimation of the weight of the smaller cube relative to the weight of the simultaneously presented larger cube harkens to the classic SWI in which the weights of smaller items are always overestimated relative to their identically-weighted larger counterparts (Charpentier, 1891). The replication of the classic SWI suggests that users were experiencing something tangible during their object interaction, and that the virtual experience was tapping into their knowledge of object properties built up over previous interactions with large and small objects manipulated in the real world. That this illusion can be experienced in VR, when no weights are actually lifted speaks both to the power of the illusion and to the immersiveness of VR.

Across two experiments, we have demonstrated that setting a cap on the speed of virtual objects during object manipulation causes the user to interact with the object in a manner consistent with lifting objects of tangible mass. Users are not only accurate in their weight estimates of virtual objects, but they are also susceptible to the SWI, which indicates that our interaction technique is powerful enough to cue users to apply weight hefting strategies used in the physical world to the virtual world. Although VR is already adept at making users feel as if they are somewhere else, true presence cannot be achieved unless users feel that virtual objects are truly present. Novel interactions techniques such as the one proposed here bring us closer to the goal of simulating the real in the virtual.

References

Achibet, M., Girard, A., Talvas, A., Marchal, M., & Lecuyer, A. (2015). Elastic-arm: Human-scale passive haptic feedback for augmenting interaction and perception in virtual environments. *IEEE Virtual Reality*, 63–68.

Achibet, M., Marchal, M., Argelaguet, F., & Lécuyer, A. (2014). The virtual mitten: A novel interaction paradigm for visuo-haptic manipulation of objects using grip force. *IEEE Symposium on 3D User Interfaces*, 59–66.

Ban, Y., & Ujitoko, Y. (2018). Enhancing the pseudo-haptic effect on the touch panel using the virtual string. *IEEE Haptics Symposium*, 278–283.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). lme4: Linear mixed-effects models using Eigen and S4. *R package version, 1*(7), 1–23.

Buckingham, G., & Goodale, M. A. (2010a). Lifting without seeing: The role of vision in perceiving and acting upon the size weight illusion. *PLoS One, 5*(3), e9709.

Buckingham, G., & Goodale, M. A. (2010b). The influence of competing perceptual and motor priors in the context of the size–weight illusion. *Experimental Brain Research, 205*(2), 283–288.

Charpentier, A. (1891). Analyse experimentale quelques elements de la sensation de poids [Experimental study of some aspects of weight perception]. *Archives de Physiologie Normales et Pathologiques, 3*, 122–135.

Dominjon, L., Lécuyer, A., Burkhardt, J. M., Richard, P., & Richir, S. (2005). Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments. *Proceedings of IEEE Virtual Reality*, 19–25.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429–433.

Hamilton, A. D. C., Joyce, D. W., Flanagan, J. R., Frith, C. D., & Wolpert, D. M. (2007). Kinematic cues in perceptual weight judgement and their origins in box lifting. *Psychological Research, 71*(1), 13–21.

Jáuregui, D. A. G., Argelaguet, F., Olivier, A. H., Marchal, M., Multon, F., & Lecuyer, A. (2014). 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.

Johansson, R. S., & Flanagan, J. R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience, 10*(5), 345.

Lécuyer, A. (2009). Simulating haptic feedback using vision: A survey of research and applications of pseudo-haptic feedback. *Presence: Teleoperators and Virtual Environments, 18*(1), 39–53.

Lécuyer, A., Coquillart, S., Kheddar, A., Richard, P., & Coiffet, P. (2000). Pseudo-haptic feedback: Can isometric input devices simulate force feedback?. *Proceedings of IEEE Virtual Reality*, 83–90.

Niehorster, D. C., Li, L., & Lappe, M. (2017). The accuracy and precision of position and orientation tracking in the HTC vive virtual reality system for scientific research. *i-Perception, 8*(3), 2041669517708205.

Paljic, A., & Coquillart, S., 2004. A passive stringed force feedback system for immersive environments. *Proceedings of Eurohaptics Conference*, 82–87.

R Core Team. (2017). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. Retrieved from https://www.R-project.org/

Ross, H. E. (1969). When is a weight not illusory? *The Quarterly Journal of Experimental Psychology, 21*(4), 346–355.

Taima, Y., Ban, Y., Narumi, T., Tanikawa, T., & Hirose, M. (2014). Controlling fatigue while lifting objects using pseudo-haptics in a mixed reality space. *IEEE Haptics Symposium*, 175–180.