

A Psychophysical Experiment Regarding Components of the Plausibility Illusion

Richard Skarbez *Member, IEEE*, Solène Neyret, Frederick P. Brooks, Jr. *Life Fellow, IEEE*,
Mel Slater, and Mary C. Whitton *Life Senior Member, IEEE*

Abstract—We report on the design and results of an experiment investigating factors influencing Slater’s Plausibility Illusion (Psi) in virtual environments (VEs). Slater proposed Psi and Place Illusion (PI) as orthogonal components of virtual experience which contribute to realistic response in a VE. PI corresponds to the traditional conception of presence as “being there,” so there exists a substantial body of previous research relating to PI, but very little relating to Psi. We developed this experiment to investigate the components of plausibility illusion using subjective matching techniques similar to those used in color science. Twenty-one participants each experienced a scenario with the highest level of coherence (the extent to which a scenario matches user expectations and is internally consistent), then in eight different trials chose transitions from lower-coherence to higher-coherence scenarios with the goal of matching the level of Psi they felt in the highest-coherence scenario. At each transition, participants could change one of the following coherence characteristics: the behavior of the other virtual humans in the environment, the behavior of their own body, the physical behavior of objects, or the appearance of the environment. Participants tended to choose improvements to the virtual body before any other improvements. This indicates that having an accurate and well-behaved representation of oneself in the virtual environment is the most important contributing factor to Psi. This study is the first to our knowledge to focus specifically on coherence factors in virtual environments.

Index Terms—Virtual reality, virtual environments, presence, place illusion, plausibility illusion, immersion, coherence, psychophysics, user studies

1 INTRODUCTION

In 2009, Slater proposed two constructs for evaluating user experience in virtual environments (VEs), those being Place Illusion (PI) and Plausibility Illusion (Psi). Place Illusion is defined as “the illusion of being in a place in spite of the sure knowledge that you are not there.” Plausibility Illusion is defined as “the illusion that what is apparently happening is really happening, in spite of the sure knowledge that it is not” [15]. PI and Psi are conceptually orthogonal components of presence, and it is argued that when one experiences PI and Psi (and therefore presence), one should also respond realistically to stimuli from the virtual scenario. Note that realistic behavior is a sign of presence, but is not itself presence. Following Schubert, we define *presence* as the cognitive feeling of being in a particular scenario. [11]

Traditionally, most presence research has focused on measuring the effects induced by changing aspects of the system hardware or implementation: display modality, field of view, tracking modality, latency, etc. Changing these system characteristics tends to change the sensorimotor contingencies of a virtual environment. (*Sensorimotor contingencies* are the regularities in how sensory stimulation depends on the activity of the perceiver [7]. For example, if one leans toward a particular object, it takes up more of their field of vision.) Using Slater’s definition [18], these factors are all elements of the immersion of a VE system. If “[i]mmersion provides the boundaries within which

PI can occur” [15], then there exists a substantial body of previous work that applies to PI in virtual environments.

In [13], Skarbez introduced the term *coherence* to refer to the aspects of a virtual environment that contribute to Psi, as immersion contributes to PI. In System characteristics that contribute to coherence include behavior of virtual humans in the environment, behavior of physical objects in the environment, and the like. Rovira et al. argue that a virtual environment must satisfy three requirements in order to engender Psi in its users. First, that it must be correlational—the actions of the user elicit responses from the environment. Second, that it must be “self-referential”—there must be elements of the scenario that refer directly to the user. Third, it must be credible—the behavior of the environment must be consistent with the user’s prior knowledge [9]. There has been comparatively little previous work on such coherence factors of virtual environments, and therefore little previous work that applies to the study of Psi.

Here we present the design and results of a study specifically focusing on coherence. The results suggest that when tasked with improving their feeling of Psi, participants tend to first improve the quality of their avatar behavior, then the correspondence between the scenario and the visual appearance of the environment, followed by the behavior of other virtual humans in the environment and then the physical behavior of objects in the environment.

2 BACKGROUND

The constructs of Place Illusion and Plausibility Illusion were introduced to the literature by Slater [15]. Place Illusion (PI) corresponds to the traditional definition of presence as the sense of “being there”. There exists a substantial body of work in the VE research community investigating the connection between immersion and presence—and therefore, between immersion and PI. A recent meta-analysis by Cummings and Bailenson included 83 experiments investigating this connection, and identified more than a hundred other related publications [5].

There are comparatively few studies that investigate the connection between coherence and Psi. Coherence is to Psi as immersion is to PI. Skarbez defines the *coherence* of a virtual scenario as the set of reasonable circumstances that can be demonstrated by the scenario without introducing unreasonable circumstances, and a *reasonable circumstance* as a state of affairs in a virtual scenario that is self-evident given prior knowledge [14].

- Richard Skarbez is with the University of North Carolina at Chapel Hill. E-mail: skarbez@gmail.com.
- Solène Neyret is with the University of Barcelona. E-mail: solene.neyret@ub.edu.
- Frederick P. Brooks, Jr. is with the University of North Carolina at Chapel Hill. E-mail: brooks@cs.unc.edu.
- Mel Slater is with the University of Barcelona, Institució Catalana de Recerca i Estudis Avançats (ICREA), and University College London. E-mail: mel Slater@ub.edu.
- Mary C. Whitton is with the University of North Carolina at Chapel Hill. E-mail: whitton@cs.unc.edu.

Manuscript received 19 Sept. 2016; accepted 10 Jan. 2017.

Date of publication 26 Jan. 2017; date of current version 18 Mar. 2017.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TVCG.2017.2657158

Wang et al. investigated plausibility in the context of shaking hands with a virtual partner. There were three hand behavior conditions: no hand, hand with unreal behavior, and human-driven hand behavior. The authors found significant effects on plausibility between each of the conditions [20].

Slater et al. investigated PI and Psi together in a study [17]. In this experiment, participants were placed in the system configuration with the highest level of immersion and coherence, instructed to remember either their feeling of PI or their feeling of Psi, and then match whichever feeling they were instructed to remember by choosing transitions from lower- to higher-quality configurations. They found that natural sensorimotor contingencies are more important for PI, while correlations between self-actions and events are more important for Psi, illumination realism may be more important for Psi, and the virtual body is important for both PI and Psi. This experimental procedure was also used by Azevedo et al. to investigate the effects of multimodal feedback (vision, 3D sound, haptic, or olfactory) on participants' sense of presence [1] [2].

Yu et al. also considered the plausibility of dynamic events that correlate with user actions in a virtual environment [21]. This study extends previous work from Slater et al., and clarifies that the increased presence observed in the earlier study was due to increased coherence, rather than improved visual quality [16].

Aspects of coherence have been considered by other authors as well under a variety of names. In an experiment performed by Biocca et al., participants performed a task (removing objects from a cadaver) in an ecologically-valid VE where the objects were organs, and in a control VE where the objects were geometric primitives [4]. We consider such scenario-appropriateness of stimuli to be a factor of coherence. The authors do not report on the differences between these conditions; they were collapsed in their analysis because they did not observe differences regarding their variables of interest (reports of visual-to-haptic or visual-to-aural intersensory illusions). Llobera and colleagues considered coherence as it applies to interactive storytelling, identifying narrative coherence and interactivity as two dimensions along which mediated storytelling experiences can be evaluated [6]. Beckhaus and Lindeman coined the term Experiential Fidelity for the extent to which a VE stimulates a user in such a way that "their expectations, attitude, and attention are aligned with the actual VR experience, and that the user's own imagination is stimulated to complete the experience," using techniques such as priming [3]. Parola et al. presented a new definition of presence, "the sense of feeling real", that is on face very similar to Psi. They refer to the presence formation process as an "alignment of external stimuli with an internal set of schemata", which highlights the importance of user expectations and prior experiences in that process [8]. Similarly, in discussing an experiment exploring the neural correlates of breaks in presence, Sjölie and colleagues comment that the "key factor in [maintaining presence] is to avoid anything that 'disproves' it by violating expectations" [12].

3 EXPERIMENT

There has been little previous work exploring Psi or coherence factors in virtual environments. Therefore, this study was designed to focus only on coherence, in order to begin to address said lack of research. This experiment is similar in methodology to the experiment described by Slater et al. in [17]. In that experiment, participants were placed in the system configuration with the highest level of immersion, instructed to remember either their feeling of PI or their feeling of Psi, and then match whichever feeling they were instructed to remember by choosing transitions from lower- to higher-quality configurations. In this experiment, we followed the same method, but were only concerned with Psi, so no participants were instructed to remember their feelings of PI.

In this experiment, we varied four coherence factors: the coherence of virtual human behavior (denoted by VH), the behavior of one's own virtual body (VB), the coherence of physical interactions in the VE (P), and the scenario coherence (S). These are elucidated further below. We refer to each instance of these four factors as a *configuration*, and denote a given configuration by a property vector of the form $C = \{VH,$

$P, S\}$.

We chose the factors in the property vector to ensure reasonable coverage of the different types of coherence (and coherence failures) that can be present in a virtual environment. *Virtual humans* were chosen as one of the factors to represent the coherence of interaction with other characters in the virtual environment. The participant's *virtual body* was chosen as another factor. In the real world, there is a justifiably strong feeling of agency—the sense that I am the one who is causing or generating an action [10]—especially when it comes to the behavior of one's own body. Therefore, we would expect the presence or absence of the virtual body, and the coherence of its behavior if present, to have a strong effect on one's level of plausibility illusion. *Physical interactions* (kicking a football) were chosen to represent the coherence of allowed interactions with the virtual environment. Finally, *scenario coherence* (the matching of the virtual environment to the represented situation) was chosen to represent any other factors, outside of specific interactions with characters or objects in the virtual environment, that may lead one to disbelieve the virtual environment as a whole. As an example, without different priming, participants are likely to expect the virtual environment to behave according to the rules of the real world. Those expectations can be violated in subtle ways by behavior that is technically valid, but feels "wrong." For example, consider a scenario set in the desert at mid-day, where all behavior is technically perfect, except the other virtual characters are wearing winter coats. The quality of the interaction hasn't changed. (Using Rovira et al's terminology, the correlation or self-reference would be the same, but the credibility would be decreased [9].)

VH (Virtual human behavior coherence)

The environment in this experiment is a virtual bar. In all trials, there are three virtual humans in the environment with the participant: a bartender, and two young men having a conversation. After approximately 30 seconds of conversation, one of the men says (some variation of) "I am going to the bathroom," and either does or does not do so, requiring him to cross in front of the participant if he does. The specific behavior of these virtual humans depends on the value of VH.



Fig. 1. Virtual humans in level 2. Note eye contact.

(VH=0) Worst behavior. All virtual humans have only idle animations while talking, and remain in the same place. (That is, neither virtual patron crosses to the bathroom.)

(VH=1) Medium behavior. Virtual humans have realistic conversation and walking animations. When crossing in front of the participant, the crossing VH does not stop or acknowledge the participant.

(VH=2) Best behavior. Virtual humans have realistic conversation and walking animations. When crossing in front of the participant, the crossing VH stops, looks at the participant, and addresses the participant directly about their football playing.

VB (Virtual body behavior coherence)

The appearance and behavior of the participant's avatar could be changed. The different possible levels of the participant's avatar are described below.

(VB=0) Feet-only avatar. In this condition, the participant is represented in the environment only by their feet. (See Figure 2.) This condition, rather than having no visible representation in the environment, was chosen to enable participants to meaningfully perform the task of interacting with the football. The feet are fully tracked as, as they are in the other VB levels.

(VB=1) Static avatar. In this condition, the participant is represented in the environment by a gender-appropriate avatar in a seated T-pose. (See Figure 3.) The avatar's legs move with the participant's; the torso and arms, however, do not move.

(VB=2) Fully-tracked avatar. In this condition, the participant is represented in the environment by a fully body-tracked gender-appropriate avatar. The avatar pose is driven by real-time input from the optical tracking system, as described below.



Fig. 2. Virtual body in level 0 (Only feet).



Fig. 3. Virtual body in level 1. Torso fixed in T-pose visible in the mirror.

P (Physical coherence)

In all trials, participants were directed to play with and control a football between their feet. The behavior of the ball when kicked was

determined by the value of P. The ball was purely virtual; there was no haptic feedback. If the ball was kicked out of reach, the participant could request a new ball to be spawned.

(P=0) Null behavior. When the ball was kicked, the force vector applied to the ball was cancelled out by an opposite force vector. In practice, this meant that the ball could be moved while it was in contact with the foot, but it would never roll or maintain momentum once out of contact with the foot.

(P=1) Semi-normal behavior. When the ball was kicked, it would randomly either behave as if it were in level P=0 or level P=2, with equal likelihood. In practice, this meant that the ball would behave normally 50% of the time, and not move 50% of the time.

(P=2) Normal behavior. When the ball was kicked, the physics engine was used to determine the path of the ball.



Fig. 4. Abstract environment (Scenario level 0).



Fig. 5. Mismatched environment, appearing to be an upscale restaurant (Scenario level 1).

S (Scenario coherence)

For all trials, the participant was in a virtual bar environment of the same physical configuration (tables and chairs in the same positions, mirror hanging on the wall facing the participant, etc.), but the representation of those objects changed depending on the value of S.



Fig. 6. Matched environment, appearing to be a bar (Scenario level 2).

(S=0) Abstract appearance. All models in the environment are replaced with simple geometric primitives. (See Figure 4.)

(S=1) Mismatched appearance. The environment model is of an up-scale restaurant. (See Figure 5.)

(S=2) Matched appearance. The environment model is of a bar. (See Figure 6.)

Discussion of configurations

Altogether, there are 81 possible configurations: 3 physical coherence \times 3 scenario coherence \times 3 virtual human coherence \times 3 virtual body coherence.

It is not absolutely clear at this point which aspects of virtual environments are important to PI and which are important to Psi (that is one of the goals of this study, after all), but one can argue for this choice of factors and levels from existing theory. In VH level 0, the virtual human behavior is neither correlational, referential, nor credible (using the terminology of Rovira et al. [9]). In level 1 it is more credible (the VH does what he says he is going to do), but it is still not correlational nor referential. In level 2, the VH addresses the participant directly, adding correlation and referentiality to the behavior. The virtual body, physical behavior, and scenario factors all offer increasing credibility at higher levels of each factor, while the virtual body and physical behavior also offer increasing levels of correlation and referentiality.

3.1 Participants

Twenty-one participants (10 males, 11 females) were recruited from the local university campus. Their average age was 24 ± 5 (S.D.) years and they were compensated for their time. This experiment was approved by the ethics board of the University of Barcelona.

3.2 Materials

The virtual environment was displayed using an Oculus Rift Development Kit 2 (DK2) head-mounted display (HMD) made by Oculus. The DK2 has a nominal 100° field of view, and a resolution of 960×1080 pixels per eye. It weighs 440 grams.

For head tracking, the internal tracking of the DK2 was used, with an update rate of 1000Hz. For body tracking, participants wore an Optitrack body suit, designed to support real-time whole body tracking of a person. It consists of a black suit with 37 retroreflective markers, which are tracked by 12 infrared cameras. Tracking was handled by the Optitrack Motive software platform.

The experiment was implemented in version 5.2 of the Unity Game Engine. The male and female avatars were created using Mixamo.

3.3 Metrics

There were three types of dependent variables. The first was the configuration $\{VH, VB, P, S\}$ at which a participant declared a matching sensation of reality. The second consisted of the transition set—that is, the specific order of improvements that a user chose to move from one configuration i to another configuration j , combined over all users. The third was the post-experiment questionnaire that was completed by all participants. This included a modified Slater-Usuh-Steed (SUS) presence questionnaire [19], as well as a series of questions asking them to rate the factors in order of which had the most impact on their sense of reality and to explain why.

3.4 Experimental procedures

3.4.1 Pre-experiment

Upon arriving at the lab, participants read an information document, signed an informed-consent form, and completed a demographic questionnaire. Participants were informed both verbally and on paper that they were free to withdraw from the experiment at any time without giving any reasons. After completing this process, participants put on the Optitrack suit and underwent a short calibration procedure, after which they donned the Oculus Rift HMD and began the experiment.

3.4.2 Experiment

Participants were seated wearing the HMD, through which they were able to see the virtual bar when looking around and the virtual body from a first-person perspective when looking down. They were also able to see themselves reflected in a mirror in front of them. The virtual body and the scenario were both at the maximum level during this first exposure. Participants were first instructed to look around the room and describe what they saw. They were instructed to move their arms and legs, and to observe these motions both directly and in the mirror in front of them; this was done to establish a sense of embodiment in the virtual body. After that, they were instructed to play with the ball they had between their feet, which was also at the highest level of coherence. Participants were then shown the highest level of coherence of the conversation between the other virtual humans.

During this exposure to the highest levels of coherence for all four factors (virtual humans, virtual body, physical behavior of the ball, scenario) participants were told to focus on the sensation of reality they were feeling at the moment. They were told that this sensation would be used as a reference for the rest of the experiment and would be referred to as the “optimal sensation of reality”. Then participants were shown all the decreased levels of coherence for each of the elements, in the same order as described above: first, the behavior of the virtual humans, then the behavior of their own virtual body, then the behavior of the ball, and finally the different levels of scenario coherence. After making sure that the participant understood all the improvements they could make to affect their sensation of reality, the experimenter would give the instruction to start the experimental procedure. Participants were told that they would be playing a game in which the goal was to reach the optimal sensation of reality they experienced at the beginning of the experience and that they would earn five points each time they would reach this level of reality. They were also instructed to focus on the elements that were their priority for getting closer to this sensation of reality. Participants started each trial in a random configuration presenting different levels of the factors and were able to improve one factor at a time until they reached the optimal level of reality. Each transition had to be an improvement (i.e. the participant could not choose to transition back to an earlier configuration), and only single-step improvements were allowed (i.e. in order to go from $\{0,0,0,0\}$ to $\{0,2,0,0\}$, the participant first had to choose the improvement $\{0,1,0,0\}$). The improvements were made by telling the experimenter which factor they wanted to improve. Similarly, they identified when the optimal sensation of reality was reached by saying so to the experimenter. Once they had reached what they believed to be the optimal sensation of reality or once all factors of coherence were at their maximum level, the next trial would begin.

There were eight trials in total and an average of six changes per trial. Each participant started from configurations $\{0,0,0,0\}$, $\{1,0,0,0\}$, $\{0,1,0,0\}$, $\{0,0,1,0\}$, and $\{0,0,0,1\}$, and from three configurations randomly chosen from the configurations in which two improvements had already been made. These eight trials were presented in random order. Figure 7 illustrates the configuration space and highlights the possible starting configurations.

3.4.3 Post-experiment

After completing the virtual reality portion of the experiment, participants completed a short post-experiment questionnaire. The whole procedure including information, consent form signing and questionnaires lasted one hour, and the participants were compensated for their participation.

4 RESULTS

4.1 Overview

As was done by Slater et al. in [17], in this work we make the simplifying assumption that the results of the eight trials are statistically independent. They cannot be truly independent, as the same participant carried out each of them, and may have learned from one trial to the next. However, the design of the experiment was such that each trial started from a different initial configuration, and so participants were presented with a different set of possible upgrades to choose from in each trial, and so had to reconsider their priorities each time.

In this section, we report on all three measures: which configurations were identified as matching the optimal sensation of reality, the order of transitions that each participant chose in each trial, and participants' responses to a post-experiment questionnaire.

4.2 Accepted configurations

The participants' task was to improve the various factors until they felt that they had reached the same "level of reality" they had felt in the best possible configuration, $C = \{2,2,2,2\}$. These accepted configurations are shown in Figure 8. (Only configurations that were accepted five or more times are included in the figure, for ease of reading.) Included in that figure are both the percentage of total accepted configurations that a given configuration makes up (yellow lines), and also the probability that that configuration was marked as accepted if it was reached (blue lines). (For example, there were 165 total accepted configurations recorded. Configuration $\{2,2,1,2\}$ was accepted 27 times, so it makes up 16.4% of the total accepted configurations. However, configuration $\{2,2,1,2\}$ was only reached 54 times across all participants. So it was accepted 50% (27/54) of the times it was reached.)

Note that the minimum number of improvements (including improvements which were part of the starting configuration for a trial) for any configuration in Figure 8 is 5, and in fact the average number of improvements for all the accepted configurations included in this figure is 6.90. This can be seen in the Markov Chain in Figure 9.

4.3 Transitions

We constructed a transition probability matrix P from the orders of configurations chosen by each participant. Each participant made approximately six improvements in each of eight trials (that is, six improvements from a starting configuration to a configuration accepted as matching the optimal sensation of reality), for a total of 936 observed transitions. Note that by construction (due to the nature of what transitions are allowed at each step) P is a very sparse matrix, with only 125 non-zero entries in an 81 x 81 matrix. (Recall there are 81 possible configurations, 3 virtual human levels x 3 virtual body levels x 3 physical behavior levels x 3 environment appearance levels; this configuration space is illustrated in Figure 7.) Given P , we can compute the probability distribution over the configurations for any given configuration. If we take as the starting configuration $C = \{0,0,0,0\}$ (the lowest possible levels for each factor), and define s as an 81-vector of all zeros except for a 1 in the element corresponding to C , then sP yields the probability distribution after one improvement has been made, sP^2 the probability distribution after two improvements have been made, and sP^n after n improvements. By construction, configuration $\{2,2,2,2\}$ is

absorbing, so the eighth step adds no information, but we can consider the probability distributions over configurations for the first seven steps.

Figure 10 shows the estimated probability distributions over the functions at each of the transitions (only probabilities greater than 0.01 are shown, for ease of reading). Figure 11 shows the most likely path through the Markov chain. (Note that at transition 4 and at transition 7, there are two approximately equal maximum probabilities, this is reflected by highlighting two nodes in the graph at the 4th and 7th levels.)

A clear majority of users chose to immediately upgrade the virtual body twice, in order to have a fully-tracked virtual body (configuration $\{0,2,0,0\}$). Following that, a majority of users upgraded the environment to level 1, moving out of the abstract environment into the mismatched environment (configuration $\{0,2,0,1\}$). After that, users tended to upgrade either the behavior of the virtual humans or the behavior of the ball to level 1, followed immediately by whichever one of those wasn't chosen first, restoring symmetry at configuration $\{1,2,1,1\}$. Users then tended to upgrade the environment for a second time $\{1,2,1,2\}$, then again were divided over whether to upgrade the virtual humans or the physics behavior to the second level, before finally choosing the other option and reaching configuration $\{2,2,2,2\}$.

4.4 Questionnaires

All participants completed a short post-experiment questionnaire. This included a modified SUS presence questionnaire [19], as well as a series of questions asking them to rate the factors in order of which had the most impact on their sense of reality and to explain why. Comments on these subjective measures are included in the discussion.

5 DISCUSSION

5.1 The virtual body is the most important factor of Psi.

The importance of the virtual body showed in all measures that were collected: matching configurations, transition probabilities, and questionnaires. 99.4% of the matched configurations (when users declared a configuration to be "equally real" as the initial $\{2,2,2,2\}$ configuration) had the virtual body at level 2, and 100% had the body at at least level 1. When it was possible to improve the virtual body (that is, VB=0 or 1), users chose to do so 81.2% (281/346 observations) of the time. When the user had no virtual body (VB=0), that increased to 83.9% (120/143) of the time. Improving the virtual body from VB=0 to VB=1, and then again from VB=1 to VB=2, were the most common first and second improvements to be made. And in our post-experiment questionnaire, 90.9% of participants said that the most important factor to improve was the virtual body. It would seem that having a virtual body that moves with one's own body is extremely powerful for convincing a user that, "This is real."

5.2 Regarding the other factors, it is very important to have them in level 1, but not necessarily in level 2.

85.5% (141/165) of accepted configurations have every factor at level 1 or higher. Only three times (out of 165 total accepted configurations) did a participant accept a configuration where the virtual human behavior was at the lowest level (VH=0), only five times did they accept a configuration where the environment was at level 0, and only 19 times did they accept with the ball behavior at level 0. So participants very much wanted the virtual humans to move, but they did not necessarily have to interact with the participant. Similarly, participants overwhelmingly rejected the abstract environment, but did not necessarily require the matched environment, and the ball needed to move realistically, but not necessarily perfectly.

5.3 The second most important factor seems to be the scenario coherence.

This preference is less strong than the virtual body being most important, but the transition probability distributions show that after improving the virtual body twice, participants then choose to upgrade the environment a slight majority (51.2%) of the time. Also, after reaching configuration $\{1,2,1,1\}$, participants tended to upgrade the environment for a second time a plurality of the time (41.4%), more than

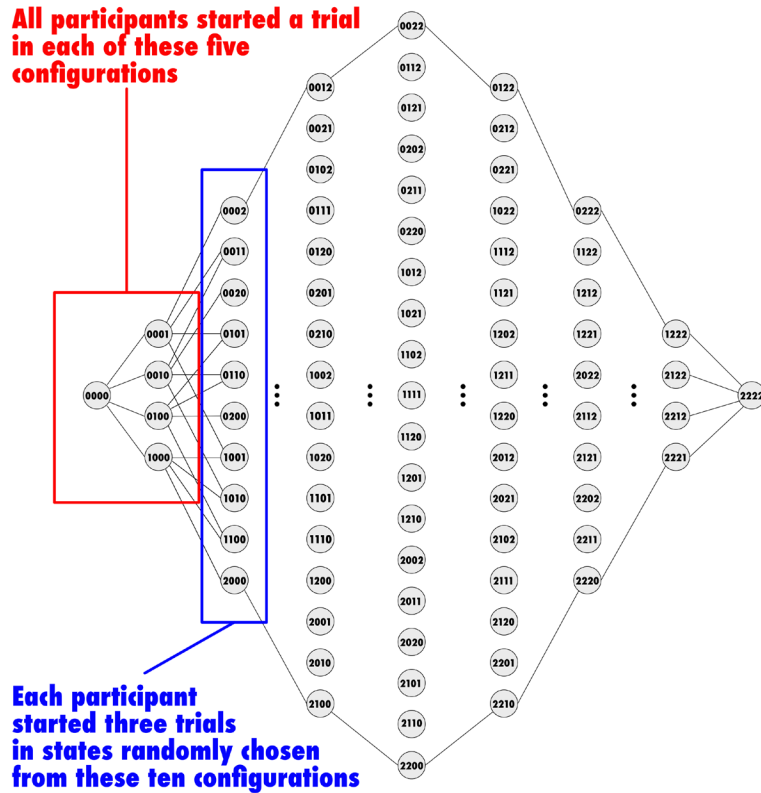


Fig. 7. Markov chain with the possible starting conditions highlighted. Each participant underwent eight trials in eight different starting conditions—the five highlighted by the red box and three sampled from the ten in the blue box. $C=\{VH, VB, P, S\}$.

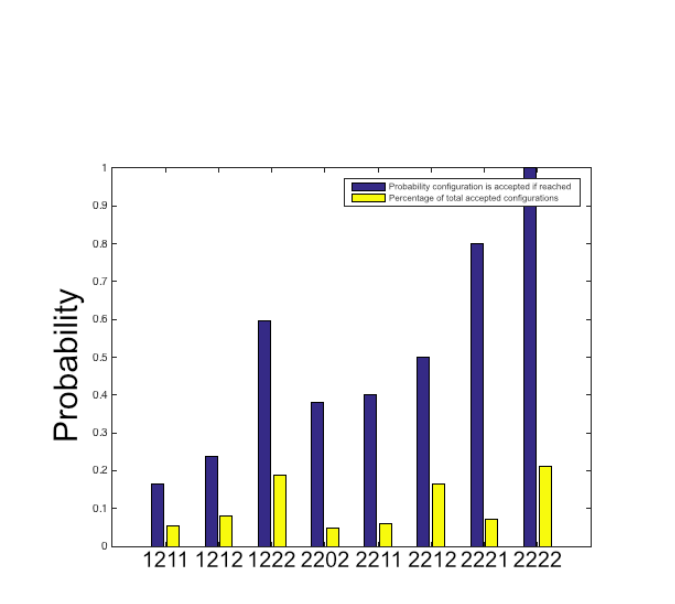


Fig. 8. Accepted configurations and their related probabilities. The blue bars indicate the probability a configuration was accepted if reached, and the yellow bars indicate the percentage of total accepted configurations the given configuration made up. 2222 is an absorbing configuration, as no further transitions are possible. $C=\{VH, VB, P, S\}$.

twice as often as any other option for the sixth improvement (starting from $\{0,0,0,0\}$; if the participant started a trial in any other configuration, one or two improvements had already been chosen for them). Note that this is not just about the visual quality or complexity of the environment, but also the appropriateness of the environment to the scenario presented. Levels 1 and 2 of the environment were designed to be of approximately equal visual quality and complexity. Participant comments such as, “Depending on the conversation that I heard, the environment was important to establish it as real,” “. . . [D]epending on the conversation, it was more clear what kind of bar would fit better,” and “Level 1 was simply too elegant for some football fans with their team shirts,” seem to indicate that this design was at least partially successful, and that participants considered the totality of the scenario and not merely the visual quality.

5.4 Response to the ball was not the same for all participants, but was very important for those participants who interacted with it extensively.

Despite the instructions being the same for all participants, there was a wide range of participant behaviors relating to the ball. Some participants barely looked at it, or touched it once or twice, just to see how realistic its motion was, before ignoring it for the rest of the trial. (This was reflected in participant comments such as, “I considered [the ball’s] presence irrelevant,” or “I had to forget about the other things to focus on the ball because it was at my feet.”) However, some participants interacted with the ball extensively and considered the quality of the ball’s movement to be very important. One participant actually rated it as the single most important factor, saying “. . . even if my body could be deformed or incomplete, the movement of the ball gave me the idea that that place was governed by physical laws similar to those of reality.”

Anecdotally, male participants seemed to interact with the ball more than female participants did. This is reflected in the accepted config-

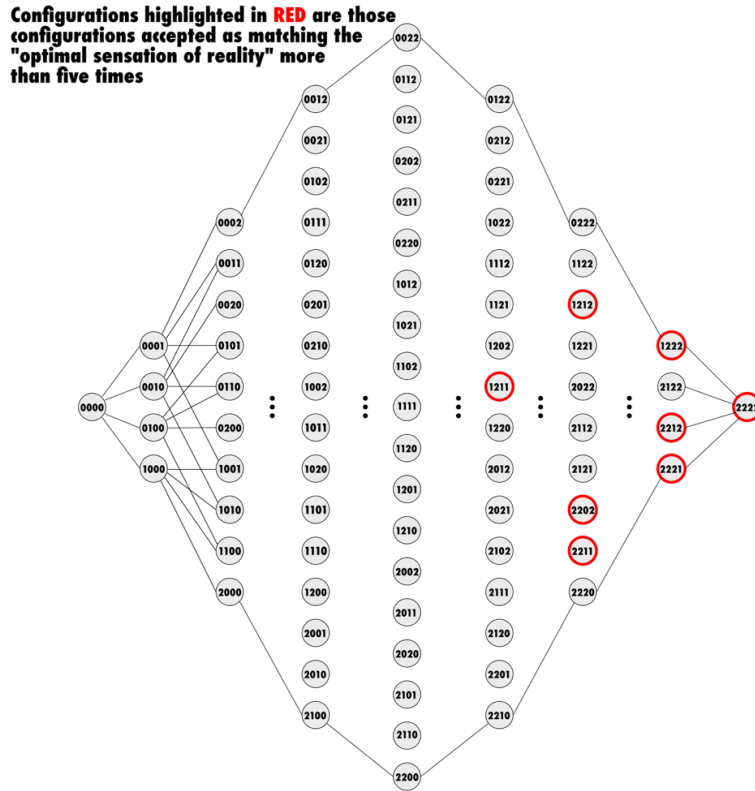


Fig. 9. The most commonly accepted configurations shown on the Markov Chain. $C=\{VH, VB, P, S\}$.

uration data. Males and females accepted the virtual humans, virtual body, and environment at strikingly similar rates. However, males accepted configurations that had the ball's behavior at level 2 63.8% of the time (51/80), while females accepted such configurations only 38.8% of the time (33/85). (Note that this difference between genders was due almost entirely to their difference in behavior at level 1 of the ball's behavior, which females accepted 49.4% of the time, and males accepted only 25% of the time. Both genders accepted the ball at level 0 roughly 11% of the time.) Using Fisher's exact test, this difference is significant, with two-tailed $p=0.0018$. Also, the most commonly accepted configuration among males was $\{1,2,2,2\}$, while the most commonly accepted configuration among females was $\{2,2,1,2\}$. For reference, the accepted configurations for both male and female users are split out in Figure 12.

This difference is also somewhat supported by the questionnaire data, as only two (of eleven) females said that the ball was the first or second most important factor, while five (of ten) males said that the ball was the first or second most important factor.

5.5 Participants who reported lower presence were more likely to accept the ball at level 0 or level 1 than participants who reported higher presence.

As part of the post-experiment questionnaire, participants completed a modified Slater-Usoh-Steed questionnaire with five questions. From this information, we divided the population into those who had low presence—0, 1, or 2 responses of 6 or 7 on the Likert scale (11 participants)—and those who had high presence—3, 4, or 5 responses of 6 or 7 (10 participants). As with the male/female split above, these groups accepted configurations with the ball at level 2 at markedly different rates. The low-presence group accepted configurations in which the ball was at level 0 or level 1 57.5% of the time (50/87), while the high-presence group accepted such configurations only 39.7% of the time (31/78). This difference between the groups is again significant, with two-tailed $p=0.029$.

The design of the experiment does not enable us to say for certain whether these differences are correlation or causation, and if the latter, in what direction. However, we speculate that users who played with the ball more (primarily, but not entirely, male users) interacted more extensively with the scenario and so felt a higher degree of presence. Then, since playing with the ball was important to them, they chose not to accept configurations in which the behavior of the ball was noticeably unrealistic.

5.6 Is the virtual body PI or Psi?

Here and elsewhere in this paper, we present theoretical arguments for the virtual body being a factor that influences Psi. (Objects that move in the virtual environment when and how the user moves provide both correlational and referential feedback to the user; the presence of a body that looks and behaves plausibly adds credibility to the scenario.) Furthermore, our experimental results demonstrate clearly that the virtual body is central to participants' *understanding* of Psi. However, it may be argued that the virtual body is also a component of PI; indeed, previous research suggests that this may be the case [17].

We argue that this apparent dual nature is due to the way the virtual body has been implemented in these studies. In both cases, the movement of the virtual body was controlled by full-body tracking, or not at all. Fully-tracked body motion comprises almost all natural sensorimotor contingencies, and so one would expect it to have a very strong positive influence on PI. However, the virtual body itself is only the means by which this tracking is visualized in the environment. We suspect that if one were to study a virtual scenario in which a virtual body was included, but *not* controlled by full-body tracking (and perhaps not controlled by tracking at all, as in the case of a third-person computer game where a user interacts with the scenario using a controller or a keyboard and mouse), one would find that the behavior of the virtual body would affect Psi primarily or only. The opposite situation (where the virtual body affects PI but not Psi) is difficult or impossible to imagine, because if a user is fully tracked, but their actions do not affect

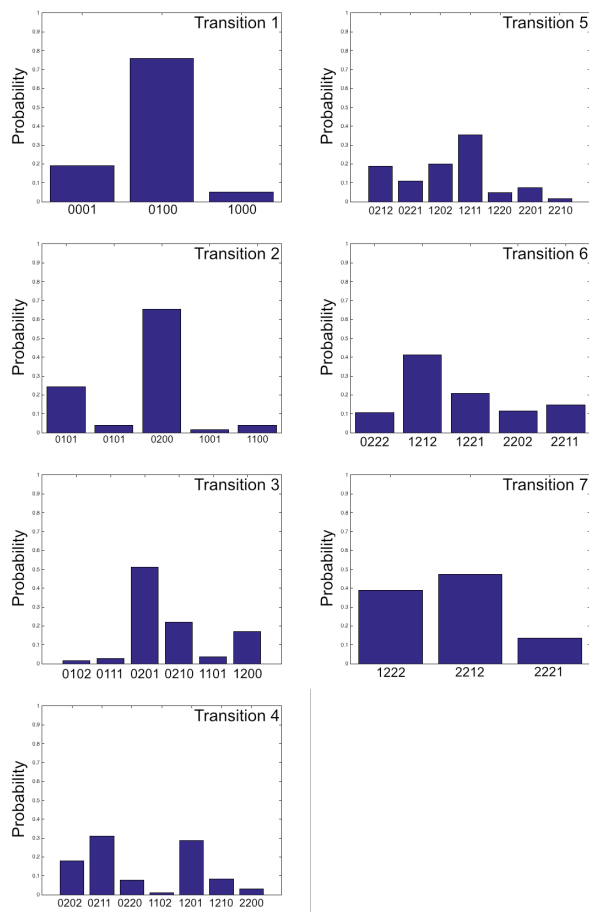


Fig. 10. Transition probability distributions for each step n , $p = sP^n$. $C = \{VH, VB, P, S\}$.

the environment in any way (because that would produce Psi), does the improved tracking actually make any difference to the user?

So, in conclusion, we argue that the presence of and behavior of a virtual body is a factor of the scenario that affects Psi, while the interface that controls that virtual body is a factor that affects PI.

6 CONCLUSION

This study is directly influenced by the study in [17], where Slater and colleagues demonstrated the feasibility of such an experiment, investigating qualia (such as PI and Psi) using matching experiments similar to those used to determine metamers in color science. In that paper, they showed that participants given different instructions (to focus on PI or Psi) chose different matching configurations and made transitions in different orders.

This experiment builds on that one, focusing exclusively on factors thought to influence Psi, namely the behavior of other virtual humans in the scenario, the appearance and behavior of the participant's virtual body, the behavior of other objects in the environment, and the appearance and "scenario-correctness" of the environment itself. Psi has not been previously investigated in this way, nor have the factors contributing to Psi been delineated before.

The results show that the virtual body is the most powerful contributor to Psi of the four factors studied in this experiment, and that result holds across matching configurations, transition probabilities, and post-experiment questionnaires. This suggests that full-body tracking is the technology that can contribute most to Psi, and that the body may indeed be the "focal point where PI and Psi are fused," as claimed in [15]. One thing that is not clear from this experiment, though, is whether the presence of any self-avatar at all is better than none, as participants

almost universally increased the virtual body to the maximum level at the earliest opportunity. It is not clear whether participants here actually thought level 1 of the virtual body was substantially better than level 0, or whether it was merely a necessary step to get to the ultimately desired configuration of having a fully-tracked avatar. More research is needed to answer this question.

This study represents only a first step, intended to further develop our understanding of the factors of coherence and Psi, which have been comparatively understudied compared to immersion and PI. A possible next step is to run a similar experiment identifying and comparing factors of immersion, as this experiment focused only on coherence. More interestingly, one could do an experiment in which participants are prompted to maximize their feelings of PI and Psi, in which the factors include both coherence and immersion factors. This could enable us to make practical suggestions as to whether immersion factors or coherence factors might be more or less important for a given type of task.

ACKNOWLEDGMENTS

Thanks to the Event Lab, University of Barcelona where this experiment was implemented and carried out. This work was partially supported by the European Research Council Advanced Grant TRAVERSE (#227985) to Mel Slater.




Fig. 11. The most likely path taken through the Markov chain, based on the probability distributions shown in Figure 10. $C=\{VH, VB, P, S\}$. (Note that at transition 4 and at transition 7, there are two approximately equal maximum probabilities, this is reflected by highlighting two nodes in the graph at the 4th and 7th levels.)

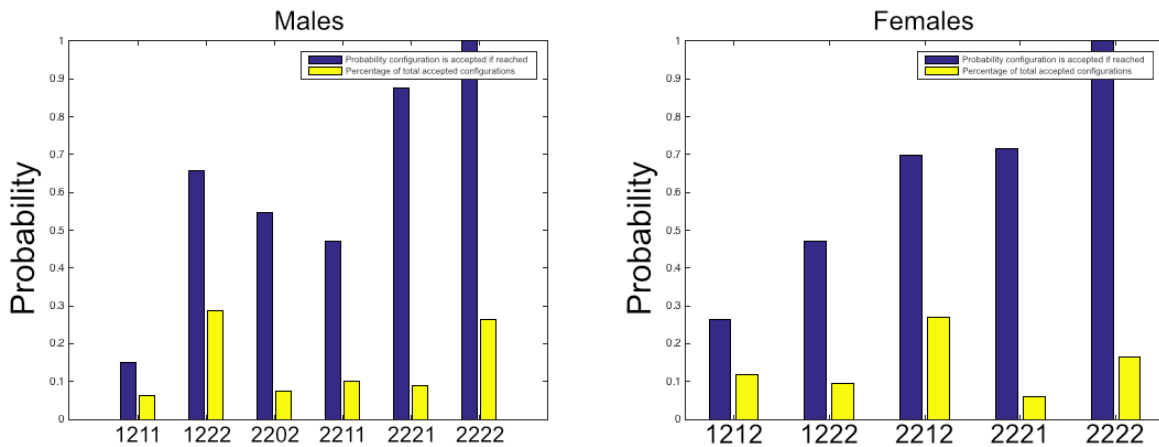


Fig. 12. Accepted configurations and their related probabilities, split by gender. The blue bars indicate the probability a configuration was accepted if reached, and the yellow bars indicate the percentage of total accepted configurations the given configuration made up. $C=\{VH, VB, P, S\}$.

REFERENCES

- [1] A. S. Azevedo. 3-D sound enhanced presence in virtual environments. 2013.
- [2] A. S. Azevedo, J. Jorge, and P. Campos. Combining EEG data with place and plausibility responses as an approach to measuring presence in outdoor virtual Environments. *Presence: Teleoperators and Virtual Environments*, 23(4):354–368, 2016/09/17 2014.
- [3] S. Beckhaus and R. W. Lindeman. *Experiential Fidelity: Leveraging the Mind to Improve the VR Experience*, pp. 39–49. Springer Vienna, Vienna, 2011. doi: 10.1007/978-3-211-99178-7_3
- [4] F. Biocca, J. Kim, and Y. Choi. Visual touch in virtual environments: An exploratory study of presence, multimodal interfaces, and cross-modal sensory illusions. *Presence*, 10(3):247–265, June 2001. doi: 10.1162/105474601300343595
- [5] J. J. Cummings and J. N. Bailenson. How immersive is enough? a meta-analysis of the effect of immersive technology on user presence. *Media Psychology*, 19(2):272–309, 2016. doi: 10.1080/15213269.2015.1015740
- [6] J. Llobera, K. J. Blom, and M. Slater. Telling stories within immersive virtual environments. *Leonardo*, 46(5):471–476, 2016/09/17 2013. doi: 10.1162/LEON_
- [7] J. K. O'Regan and A. Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5):939–973, Oct. 2001.
- [8] M. Parola, S. Johnson, and R. West. Turning presence inside-out: Meta-narratives. *Electronic Imaging*, 2016(4):1–9, February 2016.
- [9] A. Rovira, D. Swapp, B. Spanlang, and M. Slater. The use of virtual reality in the study of people's responses to violent incidents. *Frontiers in Behavioral Neuroscience*, 3(59), 2009. doi: 10.3389/neuro.08.059.2009
- [10] J. Russell. *Agency: its role in mental development*. Essays in Environmental Psychology. Psychology Press, East Sussex, UK, 1996.
- [11] T. W. Schubert. A new conception of spatial presence: Once again, with feeling. *Communication Theory*, 19(2):161–187, 2009. doi: 10.1111/j.1468-2885.2009.01340.x
- [12] D. Sjölie, G. Kalpouzos, and J. Eriksson. Neural correlates of disrupted presence: strange disruptions in a naturalistic virtual environment. In *Proceedings of the 24th International Conference on Artificial Reality and Telexistence and the 19th Eurographics Symposium on Virtual Environments*, pp. 21–28. Eurographics Association, 2014.
- [13] R. Skarbez. A preliminary investigation of place illusion and plausibility illusion. In *IEEE Virtual Reality (VR) Doctoral Consortium*, 2015.
- [14] R. Skarbez. *Plausibility illusion in virtual environments*. PhD thesis, The University of North Carolina at Chapel Hill, 2016.
- [15] M. Slater. Place illusion and plausibility can lead to realistic behavior in immersive virtual environments. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 364:3549–3557, 2009.
- [16] M. Slater, P. Khanna, J. Mortensen, and I. Yu. Visual realism enhances realistic response in an immersive virtual environment. *IEEE Computer Graphics and Applications*, 29:76–84, 2009. doi: 10.1109/MCG.2009.55
- [17] M. Slater, B. Spanlang, and D. Corominas. Simulating virtual environments within virtual environments as the basis for a psychophysics of presence. *ACM Trans. Graph.*, 29:92:1–92:9, July 2010. doi: 10.1145/1778765.1778829
- [18] M. Slater and S. Wilbur. A Framework for Immersive Virtual Environments (FIVE): Speculations on the role of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 6(6):603–616, 2016/02/18 1997. doi: 10.1162/pres.1997.6.6.603
- [19] M. Usoh, E. Catena, S. Arman, and M. Slater. Using presence questionnaires in reality. *Presence: Teleoperators and Virtual Environments*, 9:497–503, 2000. doi: 10.1162/105474600566989
- [20] Z. Wang, J. Lu, A. Peer, and M. Buss. *Influence of Vision and Haptics on Plausibility of Social Interaction in Virtual Reality Scenarios*, pp. 172–177. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. doi: 10.1007/978-3-642-14075-4_25
- [21] I. Yu, J. Mortensen, P. Khanna, B. Spanlang, and M. Slater. Visual realism enhances realistic response in an immersive virtual environment - part 2. *IEEE Computer Graphics and Applications*, 32(6):36–45, Nov 2012. doi: 10.1109/MCG.2012.121