Visual Realism Enhances Realistic Response in an Immersive Virtual Environment—Part 2

Insu Yu, Jesper Mortensen, and Pankaj Khanna ▪ *University College London*

Bernhard Spanlang ▪ *University of Barcelona*

Mel Slater ▪ *ICREA-University of Barcelona*

Does realistic illumination contribute to the sense of presence in virtual environments, or does it only seem so because realistic lighting includes dynamic effects such as shadows as participants move around the scene? An experiment to answer this question shows that both realistic lighting and dynamic effects are important contributors to presence, but for different reasons.

People experiencing VR might have the illusion of being in the virtual place and, consequently, might carry out actions as if the situation and events depicted were real. These actions can be unconscious and physiological, such as changes in heart rate, and behavioral, such as smiling at a virtual human character. They might also extend to the participant's emotions and thoughts. This sensation of being in the virtual place, operationalized by responding to the virtual environment as if it were real, is called *presence*, a concept with a long history and roots in teleoperator systems.[1] VR practitioners and application builders need to understand how VR system properties influence an average participant's sensation of presence.

In part 1 of this article, we described an experiment that compared presence in a virtual-environment scene illuminated by either ray casting or real-time ray tracing.[2] The fundamental distinction between these methods is that ray casting produces no shadows or reflections as participants move in the environment, whereas ray tracing produces dynamic effects such as real-time shadows and mirror reflections (see the sidebar, "Ray Casting versus Ray Tracing"). We assessed presence for each method, using both a questionnaire and physiological responses. The results indicated that ray tracing was more effective. We suspected that the dynamic nature of the ray-tracing illumination might account for its significantly higher presence, but we couldn't rule out that its overall higher illumination quality might be the cause.

In this article, we settle this issue by extending this experiment to determine whether the overall improved appearance or the dynamic changes produced higher presence for ray tracing.

Controlling for Illumination Quality

As in the earlier experiment, this extension compared presence in a virtual room under two conditions, but this time both conditions included dynamic shadows and reflections. In the first condition, we illuminated the room using standard OpenGL Gouraud-interpolated shading that included dynamically changing shadow umbras and reflections. In the second condition, the room was rendered with global illumination that featured soft shadows and mirror reflections and included other global-illumination effects such as color bleeding. Because both conditions included dynamic changes, any difference in results would result solely from the overall illumination quality.

Because the dynamic aspects were the same, we

 Published by the IEEE Computer Society

expected to find no difference in presence between the two illumination conditions. This would mean that overall illumination quality isn't critical to achieving presence but that dynamic changes to the illumination are—at least those caused by the participant's own movements. (This assumes that the application itself isn't intrinsically bound up with the issue of illumination, as would be the case, for example, in architectural-walkthrough applications).

The VR Evaluation System

We implemented the experimental scene on a five-PC cluster driving a four-wall Trimension Reactor system similar to a Cave Automatic Virtual Environment. (From now on, we'll just call our environment system the Cave.) A dedicated PC rendered each wall; a single master PC collected data and ran the simulation. The Cave has three 3 × 2.2 m back-projected screens—front, left, and right—and a 3 × 3 m front-projection surface on the floor. The cluster computers contained Intel Pentium 3.2-GHz processors with 1 Gbyte of RAM and Nvidia Quadro FX 5600 graphics cards.

We fitted the participants with shutter glasses synchronized with the projectors, delivering active stereo at 45 Hz for each eye. An InterSense IS-900 device attached to the top of the glasses tracked the participant's head.

The Scene

Figure 1 shows the scene we used: a library with furniture, some bookshelves, and a large mirror over a fireplace. Three area light sources illuminated the scene: two on top of the bookshelves, pointing up, and another opposite the mirror on the back wall.

We designed the virtual room to be the same size as the Cave so that its walls coincided with the Cave walls. The mirror was on the Cave's back wall—that is, the one that participants would naturally face when entering the Cave.

We used three types of rendering:

- The *baseline* was simple OpenGL-based interpolated shading.
- *OGL* was the same as the baseline, but with dynamic shadows and reflections (see Figure 2a).
- *VLF* was full global illumination using the virtual light field method described in the next section (see Figures 2b and 2c).

The dynamic shadows and reflections were identical for OGL and VLF. Figure 3 shows a person in the Cave with the scene rendered using VLF.

Participants were represented by male or female avatars (see Figure 2). Participants could see their own real body, so the avatar was invisible except for casting shadows and mirror reflections as if it were visible. The avatar was quite detailed, comprising approximately 10K polygons, textured and hand-rigged for realistic skin deformations, depending on the pose of a skeletal structure embedded in the model. The avatar model was from AXYZ Design (www.axyzdesign.com) and was animated with inverse kinematics (IK) derived only from the head tracker and a wand, each with six degrees of freedom. The IK method also used the relative position and angle between the head and the wand to estimate the elbow position. This partially replicated the participants' real body movements as they held the wand in their right hand; the avatar's right arm moved accordingly and showed up in the avatar's mirror reflection and

Ray Casting versus Ray Tracing

In ray casting, only one ray is traced from the viewpoint into a scene for every pixel (or subpixel). The illumination is determined solely by the direct light's effect on the first surface the ray intersects. So, this method illuminates only local effects, with no shadows or reflections.

In contrast, in ray tracing, the first ray through the pixel follows a path determined by specular surfaces, so that it is reflected from mirror surfaces and includes shadow rays. Thus, ray tracing portrays light interreflections between specularly reflecting surfaces. Real-time ray tracing therefore also portrays highlights and reflections that dynamically change according to the participant's head position and gaze direction.

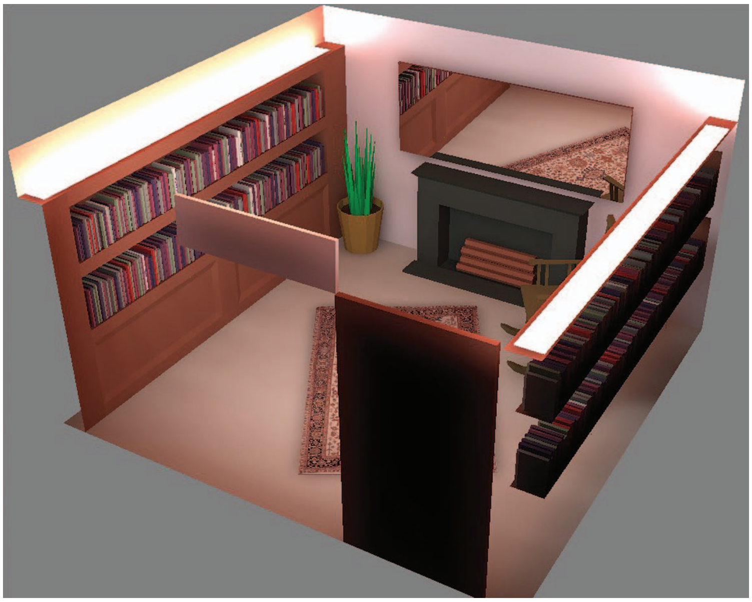

Figure 1. An overview of the scene. We chose the virtual library's size to coincide with the physical size of the system displaying it.

Figure 2. The scene rendered with (a) OGL (Open-GL-based interpolated shading including dynamic shadows and reflections), (b) VLF (virtual light field) showing the male avatar, and (c) VLF showing the female avatar. Dynamic shadows and reflections of the body appear in all the renderings, but Figures 2b and 2c have overall global illumination, and Figure 2a has only local illumination.

shadows. Moreover, the system correctly rendered head turns on the avatar representation and reflected both up and down body movements and whole-body turns in the mirror.

Rendering for Global Illumination

To provide global illumination at real-time frame rates, we developed hybrid rendering. We employ the VLF method to account for global illumination effects between the scene's static elements.[3] This accounts for full L{S|D|G}*E illumination in static environments and so adds color bleeding effects, caustic reflections from surfaces with nondiffuse BRDFs (bidirectional reflectance distribution functions), and soft shadows. The VLF method essentially stores a sampling of the outgoing radiance from all points in the scene in all directions; a rendering application can then interpolate this data to produce renderings from arbitrary vantage points. A GPU can compute a scene's VLF in linear time according to the number of polygons, providing solutions with tens of thousands of polygons with millions of irradiance and radiance samples in minutes. We can then use the converged VLF for real-time rendering, adding only a small constant time cost to a traditional renderer.

Integrating dynamic changes into a global-illumination solution is difficult. However, by breaking up the problem and attacking the light transport modes that contribute most to the image, we can achieve real-time frame rates and still support significant global-illumination effects. We can separate the problem into three main transport modes contributing to the image:

- field radiance scattered off the avatar toward the eye,
- soft shadows cast by the avatar, and
- specular reflections of the avatar.

We then focus on solving those modes. Unfortunately, this doesn't take into account reflections from the avatar and thus doesn't account for color bleeding caused by the avatar. However, the magnitude of illumination that has undergone multiple diffuse reflections is generally low and adds little to the image.

To solve the field radiance incident on the avatar, we must be able to provide the irradiance rapidly at an arbitrary spatial position and direction. To provide irradiance calculations at real-time frame rates, we use an irradiance volume.[4,5] We subdivide the scene's bounding volume into a voxel set, with each voxel storing irradiance retrieved from the VLF projected to a spherical harmonic representation.

We can interpolate irradiance values at arbitrary positions by interpolating neighboring voxels.

Calculating physically correct soft shadows is notoriously difficult. However, the GPU can render perceptually plausible soft shadows in real time. We use the Percent Closer Soft Shadows method[6] to sample a standard shadow map stochastically and provide approximate umbra and penumbra regions of a shadow caused by an area light source.

The VLF already appropriately accounts for reflections (and caustics) for the scene's static parts but not for the dynamic elements. Using a reflection rendering pass,[7] we easily render the visible scene and dynamic objects in real time onto a reflective surface, which is texture mapped to the specular geometry.

Used together, these techniques achieve a striking effect. The dynamic elements merge well with the surrounding scene, featuring impinging color bleeding, caustics, and soft shadows, and are visible in reflective surfaces (see Figures 2 and 3).

Experimental Design

The experiment was a within-groups design with two groups. Both groups experienced the baseline condition first. One group then experienced OGL and subsequently VLF; the other experienced VLF and then OGL. Table 1 summarizes the conditions.

The within-groups design let us compare the results of each participant against himself or herself across the two conditions. We adopted this approach, which differed from the part 1 between-groups experiment,[2] because we believed that the differences between OGL and VLF were subtle enough to keep participants from realizing the experiment's purpose and having this knowledge bias their responses.

On the other hand, their responses to events in the environment could change in subsequent exposures. However, if we consider only the results of their first exposure after the baseline (either OGL or VLF), we can also analyze the data as a between-groups experiment.

Experimental Procedures

The experiment involved 21 participants (six female), but data was missing for one participant, whose results we exclude from this analysis. We recruited the participants from the University College London (UCL) campus and paid them the equivalent of US$10. The UCL Ethics Committee approved the experiment under informed consent. A video of the experiment's main part is at www.youtube.com/watch?v=X2A_LVC--N8.

Figure 3. A photograph of a person in our Cave, with the room rendered with global illumination. The ceiling projector generates the shadow behind the person. The photograph shows a faint virtual shadow of the person's virtual body directed toward the wall and a mirror reflection of the avatar representing the person.

When the participants first entered each environment, we had them look around and report what they saw. This was to ease them into the environment and give us time to check that the equipment was working properly. Soft music played in the background.

To help participants adapt to the environment, we first exposed them to the baseline condition. This condition simply showed the virtual library with no avatar reflections, dynamic shadows, or events. It lasted for 150 seconds, during which the participants simply looked at and moved around the room.

Next, they experienced either VLF or OGL, according to a predetermined random order. Their third experience was in the remaining condition. The second and third trials each lasted 180 seconds and consisted of the same experience. Participants saw a reflection in the mirror of an avatar with upper-body movements that partially followed

Table 1. The experimental conditions.

Condition*	Illumination model	Dynamic shadows and reflections	Actions: books falling and character appearing
Baseline	Interpolated shading	No	No
OGL	Interpolated shading	Yes	Yes
VLF	Global illumination	Yes	Yes

**Both the baseline and OGL were based on OpenGL; VLF employed the virtual light field method.*

their own movements, as we described earlier. We had them examine the book titles, which we told them contained a clue about what was happening in the environment. We did this to encourage them to move and look around the environment.

After one minute into the second and third trials, books started falling from the bookshelves, controlled by a physics simulation. After all the books had fallen, one was displayed on the head of the avatar's mirror reflection. This wasn't one of the experimental issues, but we were curious to see whether participants would react to this—for example, by trying to touch their head to feel the virtual book. Such a reaction would mean the participant passed the *mirror test*, a sign of self-recognition. After another minute passed, a virtual boy suddenly appeared and floated around the room three times and then disappeared. Participants would typically first see this character in the mirror. This event aimed to inject a shock into the environment, useful for measuring differences in physiological responses between the conditions.

At the end of each of the three experiences, we replaced the scene with the questionnaire described in the next section. After participants answered the questions, we instructed them to close their eyes until the next session started.

Overall, the environment looked somewhat strange, with book titles suggesting paranormal events. The books falling and the floating boy's sudden appearance added to this strangeness, as did the background music.

The Presence Questionnaire

After each session, participants scored their responses to seven queries displayed on the Cave walls on a 0 to 100 scale, in which 100 was the most positive response:

Q1. I sometimes felt myself to be in the library much as if I was in a real place.
Q2. When I think about it now, I remember the library more as if it were a place I visited rather than a computer-generated world.
Q3. I sometimes forgot about the real world of the laboratory and reacted as if I were in the library.
Q4. I sometimes thought about sitting on the chair.
Q5. I sometimes thought that the plant might be a real one.
Q6. Sometimes my thoughts, feelings, and actions were as if I were in a real place.
Q7. Even though the person in the mirror did not look like me, I sometimes had the feeling as if I were seeing myself.

Q1, Q2, Q3, and Q6 related to presence. Q4 and Q5 related to whether the more realistic illumination affected the degree of realism of individual objects. When we asked the participants to give their scores for their second and third trials, we always reminded them of their previous score for the same question. Q7 related to whether seeing the mirror reflection would generate a feeling of ownership with respect to the virtual body—an issue that wasn't a concern of this experiment.

Physiological Responses

We fitted participants with a Nexus 4 device that recorded an electrocardiogram (ECG), placing electrodes on the left and right collar bones and the lowest left rib.

We acquired the ECG as a standard Einthoven I derivation (sampling frequency: 1,024 Hz) and performed the analysis with the g.BSanalyze biosignal analysis software package (from g.tec—Guger Technologies OEG, Graz, Austria). The software determined the QRS (ventricular contraction) complexes, which we inspected and manually corrected where necessary. These complexes are the time distances from one heart contraction to the next. We also measured *NN intervals*—that is, normal-to-normal beat distances (we excluded nonnormal beats such as extra systoles). We retained these measures:[2]

- Heart rate (HR)—the mean beats per minute (bpm) for the middle 2 minutes of each session, and
- NN50—the number of successive NN intervals greater than 50 microseconds.

Anxiety is generally exhibited by a higher HR and lower HR variability, here measured by NN50.

Additionally, we computed the heart-rate deceleration (HRD) related to the virtual boy appearing in the mirror. HRD has been shown to correlate with strong negative feelings.[8]

Results

We assessed the participants' responses through the questionnaire, heart rate, heart rate variability, and heart rate deceleration. Considering each of these in turn gives an overall picture of what happened in the different experimental conditions.

Presence Questionnaire Responses

Figure 4 shows bar charts for Q1 through Q6 (for 19 participants; we removed a participant who gave 0 scores on every question). The scores for the presence-related questions (Q1 through Q3 and Q6) were relatively high. Clearly, order had no

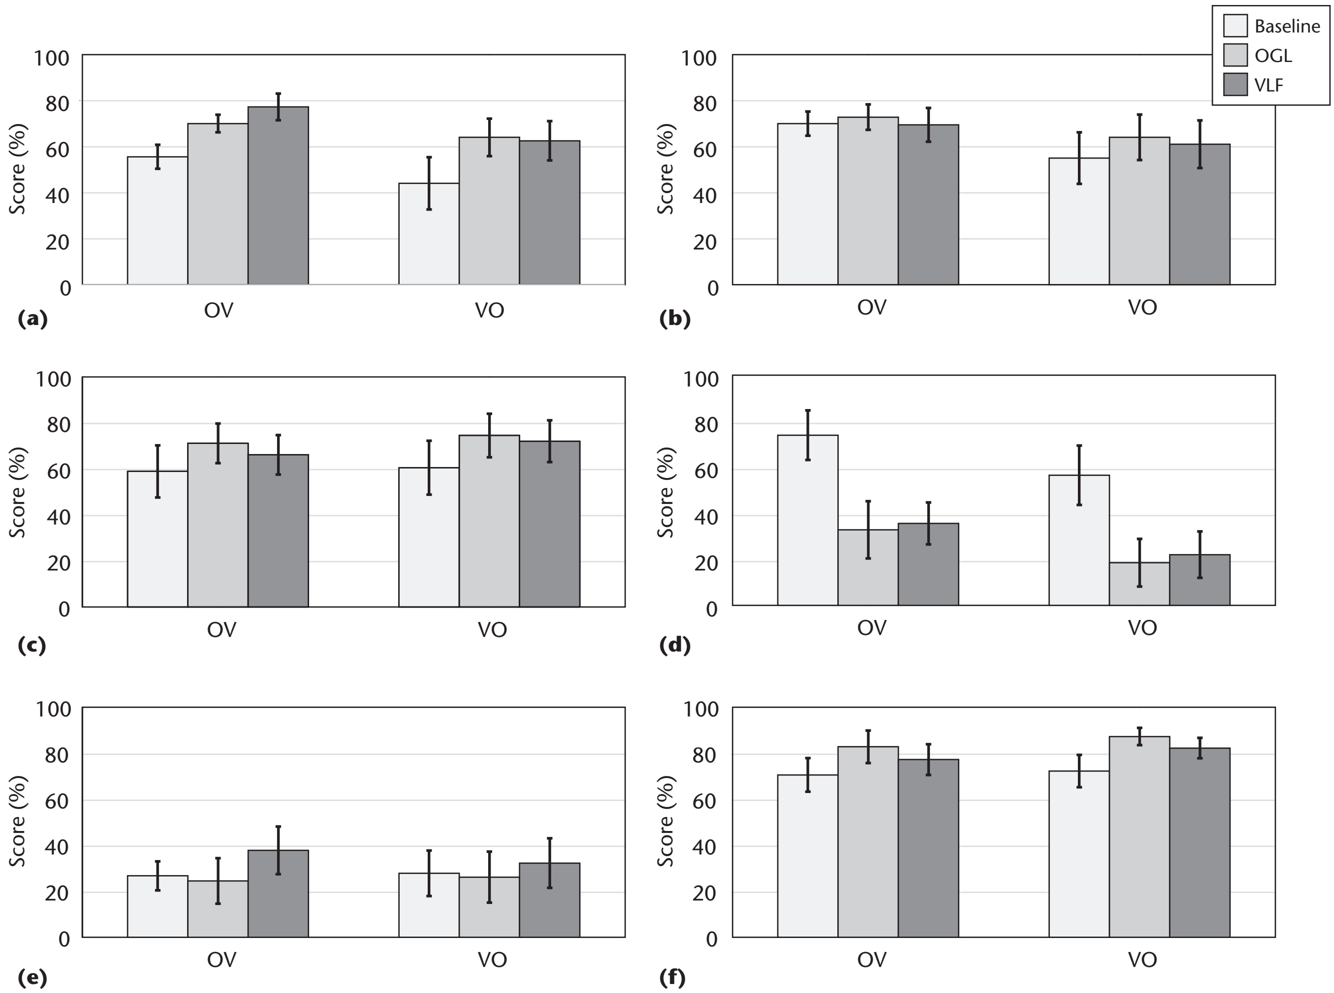

Figure 4. The mean and standard-error response variables for (a) Q1, (b) Q2, (c) Q3, (d) Q4, (e) Q5, and (f) Q6. (For the list of these questions, see the section "The Presence Questionnaire.") OV refers to the group that experienced OGL first; VO is the group that experienced VLF first. The *y*-axis is the questionnaire score on the scale of 0 to 100.

effect because the bar charts show the same pattern of responses for both the OGL-to-VLF (OV) sequence and the VLF-to-OGL (VO) sequence. The within-groups design is therefore appropriate.

It's also clear that scores for the presence-related questions increased from the baseline to OGL and VLF but showed no difference between the two latter conditions.

The idea of sitting on the chair (Q4) tended to occur to people more during the baseline but not so much in the subsequent trials. The illusion that the plant might be real (Q5) tended to occur more in VLF than in the baseline or OGL. However, the only strictly valid comparisons are between OGL and VLF because the only difference between them was the overall illumination quality, whereas the baseline included no avatar reflections or events.

Because the conditions' order had no effect, we ignored it when conducting paired tests on the combined data to compare responses to OGL and VLF. (That is, we could compare each participant's scores in one condition with his or her scores in the other.) We used the nonparametric sign test throughout because n = 19 and the data was ordinal. The questionnaire responses showed no significant differences between OGL and VLF on any question. The smallest significance levels were P = 0.11 for both Q5 and Q6.

Heart Rate

The overall mean (± standard deviation) HR was almost the same for OGL and VLF: 91.8 ± 16.9 and 91.4 ± 17.0 bpm, respectively. Similarly, the mean NN50 count was almost the same: 10.3 ± 14.1 for OGL and 8.8 ± 12.5 for VLF. So, there's no overall indication of higher stress in one condition compared to another.

However, one very interesting result is related to HR. We took a conservative measure of the overall presence level by computing the number of times each individual gave a score of at least 70 percent to Q1 through Q3 and Q6. This transformed the presence scores into a count and corresponded to what we used in an earlier study.[9] We call this the

Table 2. The mean and standard deviation for heart-rate deceleration,* for an event intended to shock the participant.

	OGL			VLF		
Trial	10 sec. before the event	During the event	P†	10 sec. before the event	During the event	P†
2	5.4 ± 5.6	9.3 ± 20.5	1	1.9 ± 4.3	11.6 ± 16.4	0.002
3	6.6 ± 10.1	4.3 ± 8.1	1	6.7 ± 10.8	5.9 ± 11.0	1

*Heart-rate deceleration is the negative of the slope of the instantaneous heart rate (in beats per minute) measured from just before the event's onset to the minimum heart rate attained in the 6 seconds after the event.

†The P values are from the sign test (n = 10 in each cell).

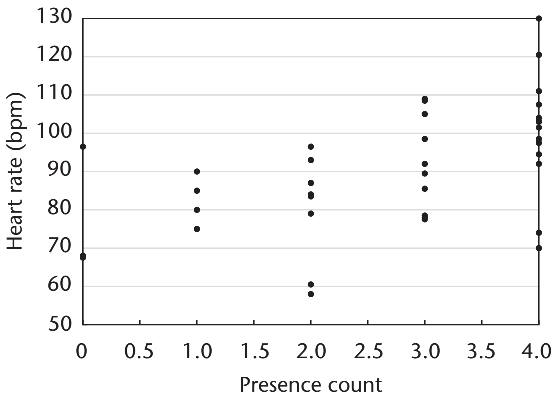

Figure 5. The heart rate in beats per minute (bpm) according to the *presence count*—the number of the four presence-related question questions that received a score of at least 70 percent. The result is a clear positive linear trend, and the correlation is high.

presence count; we plotted it on the horizontal axis against HR on the vertical axis (see Figure 5). The result is a clear positive linear trend, and the correlation is high (r = 0.51, $P < 0.001$, n = 38). This result takes into account the missing observations: 20 subjects each performing OGL and VLF yield 40 data points, of which two are lost—one each in OGL and VLF—because one participant answered 0 for all questions.

Heart Rate Deceleration

The virtual floating boy produced visible shocks in 10 subjects on their first exposure; some participants actually screamed when they first saw him. To examine this physiologically, we computed the HRD, which was the negative of the slope of instantaneous HR measured from just before the event's onset to the minimum HR attained in the 6 seconds after the event. The steeper the HRD, the greater indication of negative valence—that is, the more likely the participant reacted negatively. As a control, we compared this with the 6-second period that started 10 seconds before the boy entered the environment. We expected an ordering effect for this variable because participants could likely guess what would occur the second time the system displayed the scenario.

Because we calculated the HRD values as the negative of the slope, higher values indicate greater deceleration. Table 2 indicates that HRD was significantly higher for the event than at 10 seconds before the event, but only for VLF, and only in the trial immediately following the baseline trial. This was the case not only on average but also for each participant individually. No significant difference showed up in the participants' second exposure to this event (that is, trial 3), thus indicating adaptation.

We similarly analyzed the falling-books event. There were no significant differences at all (the lowest P value among those equivalent to Table 2 was 0.11). This illustrates that the HRD could distinguish between an event that several participants found frightening and one that—although unusual—hadn't been designed to produce a negative response.

Presence: Place Illusion and Plausibility

Summarizing the immediate results, we found

- no differences between OGL and VLF for any presence-related question, indicating that the rendering type didn't influence the level of presence;
- no differences in the overall ECG measures, indicating no observable overall physiological differences, such as stress, between OGL and VLF;
- a significant positive correlation between HR and an overall measure of presence; and
- strong HRD regarding the virtual boy entering the room for VLF but not for OGL (but only for the first exposure to the event).

If we reconsider the meaning of the experiment we described in part 1, it's now clear that the dynamic changes to shadows and reflections corresponding to body movements—not the illumination quality—accounted for the level of presence. There was no difference in presence when the dynamic changes were the same in the two illumination types.

We can put these findings into the context of a new approach to presence that distinguishes between *place illusion* (PI) and *plausibility illusion* (Psi) (see the related sidebar).

Place Illusion and Plausibility Illusion

Recent deconstructions of the concept of presence yield two orthogonal components:[1,2]

- *Place illusion* (PI) is the participants' illusion of being in the place depicted by the virtual environment.
- *Plausibility illusion* (Psi) is the participant's illusion that the events apparently happening in the virtual environment are real.

Human perceptual and cognitive systems produce these illusions, which occur in spite of high-level knowledge that the participant actually is, for example, in a laboratory wearing equipment that generates the sensory data stream leading to these illusions. The illusions are automatic—beyond conscious control.

The Physical Basis

PI's physical basis is thought to be *sensorimotor contingencies* that correspond to physical reality.[3] Such contingencies are rules that we know implicitly concerning how to use our body to perceive the world. For example, we know that to see behind an object, we would have to move our head to the side to bring obscured parts of the scene into view. When our actions in a virtual environment cause us to perceive something much as we would after performing those actions in the real world, the simplest hypothesis for our brain to adopt is that what we're perceiving is actually there. That is, we're in the place depicted by the virtual environment.

Psi's physical basis is thought to be the extent to which

- the application is programmed to produce events correlating with the participant's behavior,
- the VR events refer personally to the participant, and
- the scenario is valid with reference to a similar real-life situation in which the events depicted might occur.

The Need for Illusions

Why are these illusions necessary for successful VR applications? Imagine that, to help overcome your phobia of public speaking, you enter a virtual seminar room and stand in front of a virtual audience.[4] If you don't have the illusion of being in that place, no anxiety would arise. Similarly, if you don't have the illusion of people in that place who are responding to your behavior, it's unlikely you would feel anxious. Yet a degree of anxiety that's concomitant with what you would experience in a corresponding real place and situation is necessary for this therapy to be effective. Some automatically responding part of the brain must believe that you're standing in a room with an audience looking at you, which in turn generates the appropriate physiological and emotional responses. The fact that you know you're not actually in that place is irrelevant to the therapy's success.

To take another example, social-psychology experiments on how people respond to violent incidents, such as two people fighting in a bar,[5] can't be conducted in physical reality for ethical and practical reasons. They can be conducted in virtual environments, but for them to be valid, the test subjects must have the illusion that they're in a bar (PI) and that the fight is really happening (Psi). Without those strong illusions, the experimental studies wouldn't be useful because the appropriate physiological, emotional, and behavioral responses wouldn't arise.

Optimizing PI or Psi

Another experiment that focused on PI and Psi included two participant groups, one that had been instructed to concentrate on only their sensation of being in the virtual place (PI) and the other on only their sensation of the reality of what was happening (Psi).[2]

Participants could attempt to optimize PI or Psi by step-by-step improvements in the quality of several VR system properties: illumination rendering, field of view, type of display, and self-representation by an avatar. Participants who had been instructed to maximize Psi tended to opt more often and earlier to improve the quality of illumination than those who had been instructed to optimize PI. This points to the different underlying mechanisms behind PI and Psi.

References

1. M. Slater, "Place Illusion and Plausibility Can Lead to Realistic Behaviour in Immersive Virtual Environments," *Philosophical Trans. Royal Soc. London*, vol. 364, no. 1535, 2009, pp. 3549–3557.
2. M. Slater, B. Spanlang, and D. Corominas, "Simulating Virtual Environments within Virtual Environments as the Basis for a Psychophysics of Presence," *ACM Trans. Graphics*, vol. 29, no. 4, article 92.
3. J. O'Regan and A. Noë, "What It Is Like to See: A Sensorimotor Theory of Perceptual Experience," *Synthese*, vol. 129, no. 1, 2001, pp. 79–103.
4. D.P. Pertaub, M. Slater, and C. Barker, "An Experiment on Public Speaking Anxiety in Response to Three Different Types of Virtual Audience," *Presence: Teleoperators and Virtual Environments*, vol. 11, no. 1, 2002, pp. 68–78.
5. A. Rovira et al., "The Use of Virtual Reality in the Study of People's Responses to Violent Incidents," *Frontiers in Behavioral Neuroscience*, vol. 3, no. 59, 2009; doi:10.3389/neuro.08.059.2009.

Place illusion. PI refers to the original idea of presence as the strong illusion of being in the place represented by a VR display. The new approach suggests that PI is largely a function of the *sensorimotor contingencies* afforded by the display and tracking systems (for more on this, see the sidebar). The more that people use their whole bodies to carry out perceptual acts in a normal way, the more likely PI will occur. For example, a system employing a wide field of view and head tracking with a stereo head-mounted display will more likely lead to PI than a similar system with a narrow field of view or one that uses a joystick to navigate the environment.

With respect to OGL and VLF, the perceptual affordances were the same—in the Cave with head tracking—so we wouldn't expect to find differences in the reported PI between them, which was the case.

If our application aims to have people respond to the virtual environment realistically, we must also take plausibility illusion into account.

Consider also the strong correlation between HR and the overall presence count for OGL and VLF. The overall mean HR for the two conditions is statistically the same, and the NN50 measures indicate no change in HR variability. So, the correlation result probably reflects the participants' amount of physical activity—the greater the presence, the more activity. Furthermore, this relationship between HR and presence is visible even for just the baseline, although not so strongly (r = 0.42, P = 0.07).

So, the correlation isn't due to the necessity of participants responding to events in the virtual environment; it must instead be due to some intrinsic participant response. This finding supports the association of PI with natural sensorimotor contingencies because perception was the only reason for greater body activity, such as using the body to look and move around.

Plausibility illusion. Psi is the illusion that what's happening is real. We can now interpret the experiment reported in part 1 as showing that correlations between body movements and the corresponding changes in shadows and reflections caused the increased presence. This is in line with the factors thought to influence Psi (see the sidebar)—specifically, the correlation between participant actions and VR events.

Regarding the effect of the ghost-like boy's sudden appearance on HRD in VLF and not OGL: If participants were tending to treat the situation as real, the boy's sudden appearance would indeed have been frightening and therefore likely to induce greater HRD. In another study, we did find a greater probability of Psi under full global illumination with dynamic changes, compared with a radiosity-like solution (global illumination with no dynamic changes).[10] Our results from this new experiment support this finding.

Body ownership. Finally, six of the 20 participants did touch their own heads after they saw the book on the top of the avatar's head in the mirror. The mean score for Q7 (feeling that the avatar in the mirror was the self) was 58 ± 27, with no difference between OGL and VLF.

This experiment shows that PI isn't a function of the illumination rendering quality. PI depends on the form through which participants can perceive the virtual environment, which is a function of the degree of head and body tracking and the display factors that support natural perception (field of view, resolution, latency, frame rate, and so on). The close-to-natural visual sensorimotor contingencies that typically occur in a Cave-like system with head tracking are enough to generate high PI.

However, if our application aims to have people respond to the virtual environment realistically, we must also take Psi into account. Our results suggest two things. First, where appropriate, events in the virtual environment should correlate with the participants' actions. In our two experiments, these events were the dynamically changing shadows and mirror reflections in response to body movements. But other correlated events are also important, such as virtual characters responding appropriately to the participant's actions (not considered in this article).

Second, global illumination helps foster the illusion that the situation depicted is really happening. So, after all, we do consider real-time global illumination with dynamic changes to reflections and shadows as worth the effort.

Acknowledgments

We performed this research under the UK's Engineering and Physical Sciences Research Council Project, Presence in a Virtual Light Field (EP/C511824/1). Mel Slater's contribution was also supported through the European Research Council Project "Traverse" (contract 227985). We thank David Swapp for his invaluable help with the University College London Cave system.

References

1. M.V. Sanchez-Vives and M. Slater, "From Presence to Consciousness through Virtual Reality," *Nature Reviews Neuroscience*, vol. 6, no. 4, 2005, pp. 332–339.
2. M. Slater et al., "Visual Realism Enhances Realistic Response in an Immersive Virtual Environment," *IEEE Computer Graphics and Applications*, vol. 29, no. 3, 2009, pp. 76–84.
3. J. Mortensen et al., "Real-Time Global Illumination for Virtual Reality Applications," *IEEE Computer Graphics and Applications*, vol. 28, no. 6, 2008, pp. 56–64.
4. R. Ramamoorthi and P. Hanrahan, "An Efficient Representation for Irradiance Environment Maps," *Proc. Siggraph*, ACM, 2001, pp. 497–500.
5. G. Greger et al., "The Irradiance Volume," *IEEE Computer Graphics and Applications*, vol. 18, no. 2, 1998, pp. 32–43.
6. Y. Uralsky, "Efficient Soft-Edged Shadows Using Pixel Shader Branching," *GPU Gems*, vol. 2, Addison-Wesley Professional, 2005, pp. 269–282.
7. M.J. Kilgard, "Improving Shadows and Reflections via the Stencil Buffer," white paper, Nvidia, 2000; http://assassin.cs.rpi.edu/~cutler/classes/advanced-graphics/S11/papers/kilgard_stencil_buffer.pdf.
8. J. Cacioppo, L. Tassinary, and G. Berntson, *Handbook of Psychophysiology*, Cambridge Univ. Press, 2007, p. 588.
9. M. Slater and A. Steed, "A Virtual Presence Counter," *Presence: Teleoperators and Virtual Environments*, vol. 9, no. 5, 2000, pp. 413–434.
10. M. Slater, B. Spanlang, and D. Corominas, "Simulating Virtual Environments within Virtual Environments as the Basis for a Psychophysics of Presence," *ACM Trans. Graphics*, vol. 29, no. 4, 2010, article 92.

Insu Yu *is a postdoctoral researcher in the Télécom ParisTech's Computer Graphics Group at the Ecole Nationale Superieure des Telecoms. His research interests include real-time photorealistic rendering, high-performance GPU computing, perception-based rendering, and presence in virtual environments. Yu received a PhD in vision, imaging, and virtual environments from University College London. Contact him at insuyu@gmail.com.*

Jesper Mortensen *is a senior graphics R&D engineer at Unity Technologies, working on real-time global illumination and rendering algorithms for the games industry. Mortensen received a PhD in global illumination from University College London. Contact him at jesper@unity3d.com.*

Pankaj Khanna *is a financial services management consultant at Oliver Wyman and an honorary research associate at University College London. His research interests include global illumination, real-time physically based rendering, and VR. Khanna has an MSc in vision, imaging, and virtual environments from University College London. Contact him at p.khanna@cs.ucl.ac.uk.*

Bernhard Spanlang *is a senior research fellow in the University of Barcelona's Event Lab. His research interests include real-time animation of and interaction with virtual characters and VR applications in psychology and neuroscience. Spanlang received an engineering doctorate in vision, imaging, and virtual environments from University College London. Contact him at bspanlang@ub.edu.*

Mel Slater *is an ICREA (Catalan Institution for Research and Advanced Studies) research professor at the University of Barcelona and a professor of virtual environments at University College London. His research interests include VR and its applications in psychology and neuroscience. Slater received a DSc in computer science from the University of London. Contact him at melslater@ub.edu.*