Visuo-Haptic Systems: Half-Mirrors Considered Harmful

Christian Sandor   Shinji Uchiyama   Hiroyuki Yamamoto

Human Machine Perception Laboratory, Canon Inc.
30-2, Shimomaruko 3-chome, Ohta-ku, Tokyo 146-8501, Japan
E-mail: (sandor.christian, uchiyama.shinji, yamamoto.hiroyuki125)@canon.co.jp

Abstract

In recent years, systems that allow users to see and touch virtual objects in the same space (visuo-haptic systems) are being investigated. Most research projects are employing a half-mirror, while few use a video see-through, head-mounted display (HMD). The work presented in this paper points out advantages of the HMD-based approach. First, we present an experiment that analyzes human performance in a target acquisition task. We have compared a half-mirror system with an HMD system. Our main finding is, that a half-mirror significantly reduces performance. Second, we present an HMD-based painting application, which introduces new interaction techniques that could not be implemented with a half-mirror display. We believe that our findings could inspire other researchers, employing a half-mirror, to reconsider their approach.

1. Introduction

Human perception is multi-modal: the senses of touch and vision do not operate in isolation, but are rather closely coupled. This observation has inspired systems that allow users to see and touch virtual objects at the same location in space (visuo-haptic systems). Most visuo-haptic systems have been implemented using a half-mirror to display computer graphics in the haptic workspace [7, 9, 12]. This approach achieves a better integration of vision and touch than a conventional, screen-based display; thus, user interactions are more natural. Few research projects (for example, [2]) use a video see-through, head-mounted display (HMD) instead of a half-mirror. An obvious advantage of the HMD is that the user's view of the real world and the computer graphics are not dimmed. While this is definitely increasing the realism of the virtual objects, it is hard to present to the user a consistent scene: real-world, computer graphics, and haptic forces have to be aligned very precisely.
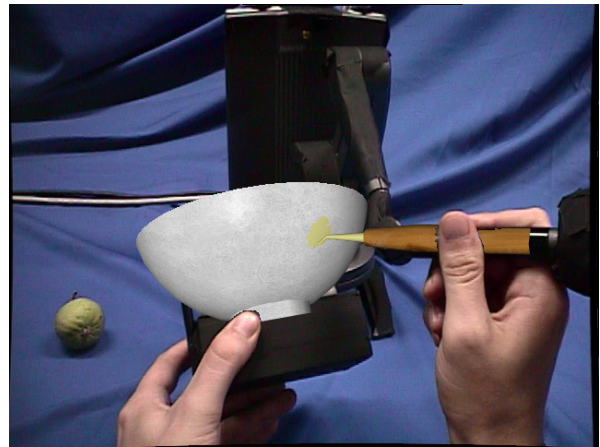

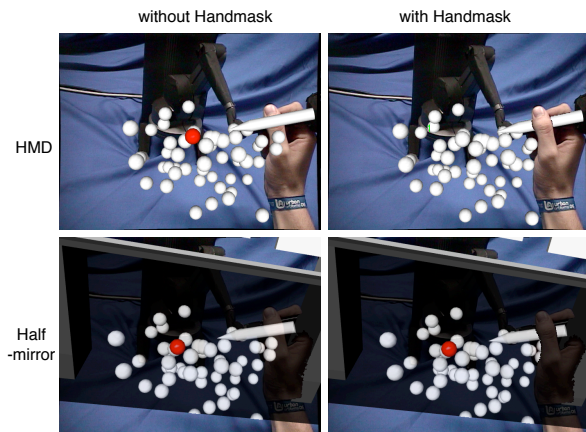Figure 1: View through an HMD in our visuo-haptic painting application.

It is our goal to show that the HMD-based approach has two significant advantages. First, novel interaction techniques (see Section 3) can be implemented. Figure 1 shows a user who paints with a virtual brush on a virtual teacup. He can see and feel the brush, as it is superimposed over a PHANTOM [11]. The user can feel that he is holding a cylindric object in his right hand. Combined with the visual sensation, he experiences a believable illusion of a real brush. To achieve this effect, two ingredients are necessary: fully opaque occlusion of real-world objects by computer graphics, and handmasking. Handmasking [8] refers to the correct occlusions between the user's hands and virtual objects. These effects could hardly be implemented with a half-mirror display. The second advantage is improved user performance, as shown by the experiment in Section 2.

2. Experiment

In our experiment, we have chosen a target acquisition task, and completion time as metric. This combination

IEEE
COMPUTER
SOCIETY

(a) The user wears an HMD and operates a PHANTOM.

(b) The user's view of the four tested conditions.

Figure 2: Experimental Platform.

has been extensively used as benchmark for 2D interaction techniques, which has led to many fundamental insights about human psychophysics; for example, Fitt's law. In our factorial design, we compare two factors on two different levels: display device and handmasking.

We had collected feedback from users of our mixed reality (MR) systems, who highly appreciate handmasking, because it makes the interaction with virtual objects more natural. In our study, we wanted to evaluate the benefits of handmasking in visuo-haptic applications.

The second factor is the display device being used. Since we did not have access to a half-mirror device, we implemented a virtual, stereoscopic, half-mirror device that users can observe via an HMD. As most half-mirror devices do not employ head-tracking, we tried to emulate this as truthfully as possible. Luciano et al. [7] do use head-tracking, however, they do not align the virtual pointer with the haptic device.

2.1. Experimental Platform

Hardware The haptic device in our experiment is a PHANTOM Desktop [11]. We used Canon's COASTAR-type HMD (Co-Optical Axis See-Through for Augmented Reality [14]) for visual augmentations. It is lightweight (327 grams) and provides a wide field of view (51 degrees in horizontal direction). It is stereoscopic with a resolution of 640x480 for each eye. A special feature of this HMD is that the axes of its two video cameras and displays are aligned. For accurate position measurements, we have used a Vicon tracker [16]. This is a high-precision optical tracker, typically used in motion capturing applications. It delivers up to 150 Hz update rate and has a high absolute accuracy (0.5 mm precision). All software was deployed

on one PC with 1GB RAM, Dual 3.6GHz Intel Xeon CPUs, GeForce 6600GT, and two Bt878 framegrabbers. The operating system was Gentoo Linux with 2. 4. 31. Kernel.

Software For rendering the computer graphics, we used plain OpenGL with an additional model loader. Furthermore, we have employed two frameworks: OpenHaptics (Version 2.0; included with the PHANTOM) and MR Platform [14] (Internal version). MR Platform provides a set of functions for implementing MR applications; for example, calibration, tracking, and handmasking. Our implementation of handmasking does color-based detection of the user's hands in the video images obtained from the HMD's cameras. This information can be used to mask this part of the computer graphics scene via OpenGL's stencil buffer. As a result, the user's hands are always visible (see Figure 2(b)).

Calibration and Tracking Procedure From a calibration perspective, we have three relevant objects (see Figure 2(a)): the user's HMD, the base of the PHANTOM and the PHANTOM pen. The relation between the attached markers and these physical objects can be calibrated using MR Platform's calibration tools.

For precise alignment of graphics and haptics, a good solution seems to be the methods proposed by Bianchi et al. [2]. We have chosen a much simpler, yet not as thorough approach. For rendering of the computer graphics, we just use the Vicon's tracking data. Its update rate is high enough, whereas the jitter is almost not visually perceivable by a user. The graphical framerate was constantly 30 Hz. For the haptic rendering, we had to chose a different

IEEE
COMPUTER
SOCIETY

approach, since OpenHaptic's HLAPI bases its force rendering on the values of the PHANTOM's encoders. However, the absolute position accuracy is bad (we measured up to 20mm error). Essentially, the PHANTOM's measurements are non-linearly distorted. To keep the haptic and MR world consistent, we determine the offset between the PHANTOM's measurements and the real pen position (as determined by the Vicon) in every haptic rendering pass. The inverse of this offset is applied to the geometry that is passed to HLAPI. Thus, the haptic experience matches the visual experience, although they happen internally in two different locations. This approach results in haptic rendering that jitters with the same amplitude as the Vicon's data. However, for the short haptic contacts in our study, this was hardly a problem.

2.2. Experimental Task and Conditions

In our experiment, the participants' task was to consecutively touch 50 balls as fast as possible. All balls were rendered with haptic feedback, and had a diameter of 10 millimeters. The ball that had to be touched was indicated by rendering it red. When the user touched that ball, another ball was randomly highlighted. After touching all 50 balls, the task completion time was displayed to the participant. This task had to be done in four different conditions (see Figure 2(b)).

To simulate the half-mirror display, we had to simulate two optical effects of such a system: the appropriate transparency of the mirror and the distortion of the graphics. The transparency was simply implemented with blending in OpenGL. We concluded from observations of commercially available half-mirror systems that their transparency is about 25%. Since most half-mirror systems do not use head-tracking, the image gets distorted, when the user moves her head from the ideal viewpoint. To achieve this effect, we have implemented a parallel camera stereo display, as described in [17]. We have placed the center of the virtual image plane at the PHANTOM's origin and rotated it 45 degrees horizontally. As distances we have chosen: 65 millimeters for left to right eye, and 450 millimeters for eyes to image plane.

2.3. Subjects

Ten naive volunteers have participated in the experiment. All subjects were screened with a Bausch & Lomb Orthorater for near/far acuity, and an additional test for stereo perception. Half of the subjects were male, half female. Their ages ranged from 25 to 40 (mean=32, standard deviation=5). They have been randomly selected from members of our research center. Eight of the subjects had previous experience with MR systems. Only two subjects had previously used a haptic device.

2.4. Experimental Design and Procedure

Each subject has participated in four separate sessions on four consecutive days. Each 10 minute session involved only one condition and was divided into two parts. First, the subjects could practice the condition. Basically, the subjects were doing the same interactions as during the measurement part. However, the experimenter was giving them advice on how to improve their results. Second, a measurement phase followed. Subjects had to do the pointing task four times consecutively. After each completed task, the completion time was shown to the participant.

A within subjects design has been chosen. To minimize skill transfer between the sessions, three preventive measures have been taken. First, the order of conditions was randomized. Second, the sessions were at least 15 hours apart. Third, the training at the start of each session served as a buffer.

2.5. Results and Discussion

Figure 3 presents the data that we have collected in the experiment. Analysis of variance for the overall mean task completion times showed that the display device ($F(1, 156) = 73.34, p < 0.0001$), and interaction between display device and handmasking ($F(1, 156) = 10.76, p = 0.0013$), were statistically significant factors. Handmasking was not statistically significant: $F(1, 156) = 1.88, p = 0.17, a = 0.05$).
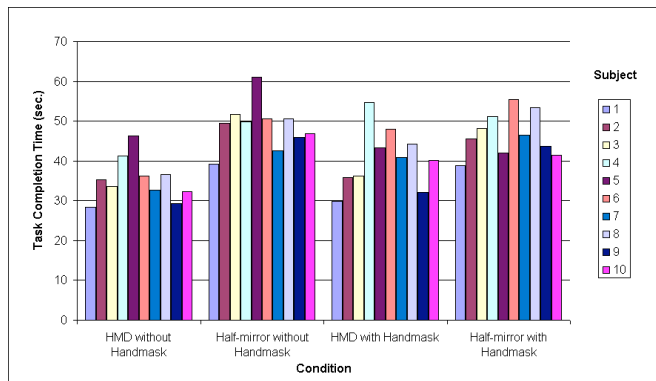
Further statistical analysis of the factor handmasking for each display device revealed: in the HMD condition, handmasking is significantly impeding performance ($p < 10^{-6}$, t-Test); in the half-mirror condition, the handmasking does not significantly change performance ($p = 0.20$, t-Test).

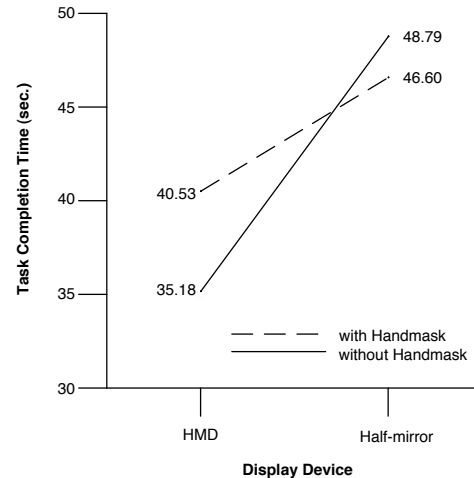The implications of our experiment are threefold:

Performance with HMD is better than with half-mirror
Task completion time in the HMD condition was significantly faster than in the half-mirror condition, which was the result we had expected. We now describe three possible reasons, which go in line with Coello's findings [3].

First, the half mirror condition impedes 3D perception due to the absence of head-tracking. Motion parallax cannot be employed by users to achieve depth perception. Even worse, head movements by the user cause distortions of the computer graphics. Second, due to the relatively low level of transparency of the half-mirror, the scene's background can hardly be perceived by the user. As described by Coello, target acquisition tasks can be performed more easily in the presence of a textured background. Third, an important issue is the visibility of the user's hand. Studies have shown that when users can see their own hand, target acquisition time is reduced. Other studies have shown that during path planning, specific brain regions respond to the stimulus of seeing one's hand [4]. The half-mirror conditions makes it hard for subjects to perceive their own hand.

(a) Mean task completion times for all subjects and all conditions.

(b) Mean task completion times for all conditions.

Figure 3: Results of the experiment.

Handmasking decreases performance with HMD   The significant decrease in user performance when using handmasking in the HMD condition was *unexpected* for us. Although our implementation is imperfect, it is generally liked by our users. Because we do not obtain depth information for the hand, we can only correctly occlude virtual objects when they coincide with real objects (for example, the fude pen in Section 3). A possible reason is that sometimes the target ball was hidden by the user's hand.
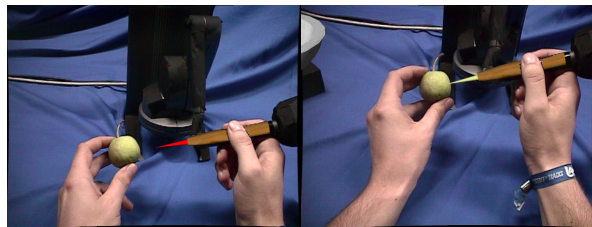
Handmasking has no effect at half-mirror   This result was *surprising* for us, since the computer graphics for the balls (not taking into account distortions) are identical for the half-mirror and the HMD condition. Thus, we would expect that handmasking also decreases performance in the half-mirror condition. From the data collected in our experiment a definite reason cannot be deducted. However, we can offer a possible explanation. Handmasking must have a positive in the half-mirror condition, that compensates for the bad effects as determined in the HMD condition. The factor that could be mitigated is the diminished display of the user's hand, as handmasking gives an indication of the user's hand.

3. Painting Application

To explore the potentials for novel interaction techniques in an HMD-based approach, we have decided to create a visuo-haptic painting application. The setup is

the same as for our experiment; however, we have added a graspable object to enable users to control the virtual teacup. An interesting idea in our application (see Figure 1) is that haptic device itself is used as a tangible [15] representation of a real-world object (the brush). The idea to mask a haptic device has already been proposed by Inami et al. [6]. However, their system does not allow to display opaque graphics on top of the PHANTOM. We have invented a new interaction technique, called cross-reality color picking, to make color selection from real objects very easy (see Figure 4). The user just touches a real object with his brush and presses the button of the PHANTOM. We know the position and orientation of the user's head and the brush; furthermore, the distortion parameters of the HMD's cameras had been determined in its calibration. From these data, we can easily obtain the (R,G,B) values of the pixel that lies on the back-projected tip of the brush in the image plane. One part of the interaction is very similar to the I/O brush [10]: acquiring colors from real world objects by touching them with a brush. However, actually using these colors is very different in our system. The I/O brush still needs a computer screen to paint on. In our system, we can paint directly on objects located in the real world, eliminating the unnatural interaction of using a computer screen as a canvas. Also, we could achieve a much higher precision (around 1 millimeter) for color selection.

The haptic and computer graphics rendering of the brush are quite simple. For the computer graphics part: the parts of the bristle that would penetrated the teacup

IEEE
COMPUTER
SOCIETY

(a) The user picks up a real apple …   (b) and selects its color.

Figure 4: Cross-reality color picking: selecting colors from real objects to apply them to a virtual object.

Figure 5: Drawings on virtual teacups, created with our painting application. The bottom right subfigure shows a drawing that imitates the appearance of a real teacup.

are projected on the tangential plane, as defined by the teacup's surface and the contact point. This simple technique is non-realistic, but provides visual feedback for the users. On the haptic side, to simulate the soft feeling of a fude brush, we use these parameters in OpenHaptics: `stiffness=0.2; damping=0.1; friction=0.1`.

We have tested our painting application by conducting a user test. We asked 14 subjects to paint teacups with our system (see Figure 5). The procedure was: First, users could familiarize themselves with the graspable object for controlling the virtual object, the colorpicking technique, and applying color to the teacup (3 minutes). For this purpose, we had put several objects such as vegetables and fruit on the table. Next, the subjects could paint whatever they felt like (5 minutes). Finally, they had to answer a questionnaire.

The feedback we have collected indicates that our system was extremely intuitive to use. Even in the short time-slots for our study, subjects had no problems to understand and use our system. It would hardly be possible to achieve similar results in this short time using a standard 3D modeling application. A common problem that was pointed out by our users is the unstable haptic experience. Our simple registration approach is not sufficient for supporting very fine-grained paintings.

4. Conclusions and Future Work

To conclude this paper, we first discuss our experiment, then the painting application, and finally the overall implications of our work and possible future directions.

Experiment The experiment we have conducted provided novel insights about human perception in visuo-haptic systems. Our experimental platform contained two unique aspects: First, the usage of handmasking in a half-mirror system. Second, in this condition, the user's hand is

moving *without* distortions, and, at the same time, the virtual representation of the pointing device is moving *with* distortions. This is an unprecedented experiment. According to Coello [3], studies have been conducted that distorted either the complete subject's view, or only the lower half of it.

The main results of our study point towards favoring an HMD over a half-mirror display. We have also discovered perceptual problems by users of our handmasking implementation. To further investigate the unexpected results for handmasking in the half-mirror condition, we want to separate the effects of missing head-tracking, and diminished visibility of hand and background. For conducting experiments about the visibility aspects, our virtual half-mirror comes in handy, since we can change the half-mirror's transparency arbitrarily. We can even set different transparency values for the user's hand and the background. We also plan to integrate Hayashi et al.'s method [5] for handmasking with depth recovery into our system.

An important question is the generalizability of our experiment. Was our implementation of the virtual half-mirror really adequate? While we cannot answer this question with certainty, we strongly believe so. We have tried to reproduce the half-mirror as concisely as possible. Of course, the only way to fully prove our claims would be to conduct studies with a real half-mirror. However, in this case, the implementation of handmasking would be very difficult.

Painting While we could not overcome all technical difficulties, our visuo-haptic painting application has shown new directions for human-computer interaction. While other systems perform better on particular aspects of the

IEEE
COMPUTER
SOCIETY

painting interaction (e.g., better computer graphics [1], or better haptic rendering [13]), the overall concept of our system contains novel points. We foster the advantages of using an HMD by allowing users to naturally interact with real world objects, as exemplified by our new cross-reality color picking technique. Also, by fully occluding the tip of the PHANTOM with a computer graphics representation of a brush, we create a virtual, tangible device. Furthermore, we support bi-manual interaction in a visuo-haptic system.

Another advantage of using an HMD is the possibility of collaborative work. In future work, we would like to exploit this advantage, by creating collaborative, visuo-haptic applications. Currently, the biggest problem in our implementation is the jitter of the haptic force rendering caused by our extremely simple registration approach. We plan to implement Bianchi et al.'s method [2] to overcome this limitation.

Overall implications   We have pointed out deficiencies of half-mirrors in visuo-haptic systems. We believe that every researcher in visuo-haptics should carefully reconsider if half-mirrors are really adequate.

We have shown that HMDs offer new possibilities for human-computer interaction in visuo-haptic systems. The feedback we have collected from the users of our painting application was encouraging. Thus, we plan to further investigate interaction techniques for the combination of haptic devices, tangible interactions and MR.

We have obtained interesting results by our experiment. We could show that, for our task, an HMD outperforms a half-mirror. However, there are other tasks, where results should be on the same level. For example, writing a virtual letter should have a similar speed in both cases, since the task is completely 2D and requires few head movements. An interesting question for future work is: what are the criteria of a task that cause problems with a half-mirror? Overall, it seems that using an HMD is closer to performing a task with real objects, than is using a half-mirror.

Another, slightly unsettling, observation is: there are phenomenas that result in significantly differing human performances for HMDs and half-mirrors. In the HMD case, hand-masking had no significant effect, whereas the effect was significant in the half-mirror condition. A lot of experiments in psycho-physics and cognitive science have used half-mirrors to investigate the influence of phenomenas on human perception. What if some experiments had also chosen a phenomena akin to handmasking? It seems that the findings in our paper suggest at least a certain fuzziness in those experiments.

Acknowledgements

References

[1] William V. Baxter, Vincent Scheib, and Ming C. Lin. dAb: Interactive haptic painting with 3D virtual brushes. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 461–468. ACM Press / ACM SIGGRAPH, 2001.

[2] G. Bianchi, B. Knörlein, G. Szkely, and M. Harders. High precision augmented reality haptics. In *Eurohaptics 2006*, pages 169–177, July 2006.

[3] Yann Coello. Spatial context and visual perception for action. *Psicológica, 26*, pages 39–59, 2005.

[4] Michael S. A. Graziano, Dylan F. Cooke, and Charlotte S. R. Taylor. Coding the location of the arm by sight. *Science*, 290(5497):1782–1786, December 2000.

[5] Kenichi Hayashi, Hirokazu Kato, and Shogo Nishida. Occlusion detection of real objects using contour based stereo matching. In *ICAT '05: Proceedings of the 2005 international conference on Augmented tele-existence*, pages 180–186, New York, NY, USA, 2005. ACM Press.

[6] Masahiko Inami, Naoki Kawakami, Dairoku Sekiguchi, Yasuyuki Yanagida, Taro Maeda, and Susumu Tachi. Visuo-haptic display using head-mounted projector. In *VR '00: Proceedings of the IEEE Virtual Reality 2000 Conference*, pages 233–240, New Brunswick, New Jersey, USA, 2000. IEEE Computer Society.

[7] Industrial Virtual Reality Inc. Immersive touch. http://www.immersivetouch.com/.

[8] Toshikazu Ohshima and Hiroyuki Yamamoto. A mixed reality styling simulator for automobile development. In *International Workshop on Potential Industrial Applications of Mixed and Augmented Reality (PIA)*, Tokyo, Japan, October 2003.

[9] Reachin Technologies AB. Reachin display. http://www.reachin.se/products/Reachindisplay/.

[10] Kimiko Ryokai, Stefan Marti, and Hiroshi Ishii. I/o brush: drawing with everyday objects as ink. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 303–310, New York, NY, USA, 2004. ACM Press.

[11] Sensable Technologies, Inc. Phantom desktop haptic device. http://www.sensable.com/products/phantom_ghost/phantom-desktop.asp.

[12] SenseGraphics AB. Sensegraphics. http://www.sensegraphics.com/products_immersive.html.

[13] Yuichi Suzuki, Yasushi Inoguchi, and Susumu Horiguchi. Brush model for calligraphy using a haptic device. *Transactions of the Virtual Reality Society of Japan*, 10(4):573–580, 2005.

[14] Shinji Uchiyama, Kazuki Takemoto, Kiyohide Satoh, Hiroyuki Yamamoto, and Hideyuki Tamura. Mr platform: A basic body on which mixed reality applications are built. In *Proceedings of the International Symposium on Mixed and Augmented Reality*, pages 246–255, Darmstadt, Germany, October 2002.

[15] Brygg Ullmer and Hiroshi Ishii. The metadesk: Models and prototypes for tangible user interfaces. In *ACM Symposium on User Interface Software and Technology*, pages 223–232, 1997.

[16] Vicon Peak. Vicon motion capturing systems. http://www.vicon.com/.

[17] Andrew Woods, Tom Docherty, and Rolf Koch. Image distortions in stereoscopic video systems. In *Proceedings of the SPIE Volume 1915*, San Jose, CA, USA, February 1993.

Second Joint EuroHaptics Conference and Symposium on Haptic
Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)
0-7695-2738-8/07 $20.00 © 2007 **IEEE**

IEEE
COMPUTER
SOCIETY