

Visual Occlusion in an Augmented Reality Post-Stroke Therapy Scenario

Max Allen

University of Otago
Information Science
Dunedin, New Zealand
max.allen@otago.ac.nz

Simon Hoermann

University of Otago
Medicine and
Information Science
Dunedin, New Zealand
simon.hoermann@otago.ac.nz

Thammathip
Piumsomboon

University of Canterbury
The HIT Lab NZ
Christchurch, New Zealand
thammathip.piumsomboon
@pg.canterbury.ac.nz

Holger Regenbrecht

University of Otago
Information Science
Dunedin, New Zealand
holger.regenbrecht@
otago.ac.nz

ABSTRACT

We investigate the effect that visual occlusion plays on users' perception in a prototypical augmented reality post-stroke therapy system. We implemented a reach-based therapeutic exercise using the Microsoft Kinect to enable depth sensing and correct visual occlusion of the upper-limb. Thirty participants evaluated the exercise with three different visual occlusion modes: the correct visual occlusion, a virtual-always-occludes—to date the most commonly used mode in augmented reality—and a mode with semi-transparent virtual objects. The analysis of their reported experience showed that correct occlusion was the most preferred mode for performing the reaching exercise, providing a more tangible and realistic interactive experience.

Author Keywords

Occlusion, augmented reality, therapy, rehabilitation, post stroke, upper limb, visualisation, perception

ACM Classification Keywords

H.5.1 [Multimedia Information Systems]: Artificial, augmented, and virtual realities; J.3 [Life and Medical Sciences]: Health

INTRODUCTION

Upper-limb impairments brought on by neurological conditions such as complex regional pain syndrome and stroke can severely reduce quality of life. Once widely considered to be intractable, much research has been undertaken to provide methods for reducing pain, increasing movement, and promoting the rehabilitation and general well-being of patients. “Visual illusion” is a successful method to treat these conditions. Since its first introduction in the form of mirror visual feedback by Ramachandran et al. (1995) as a method for treating patients with phantom limb pain, an abundance of research has broadened the therapeutic-application including Complex Regional Pain Syndrome (McCabe, Haigh, & Blake, 2008; Moseley, 2006) and Stroke (Altschuler et al., 1999; Dohle et al., 2009; Yavuzer et al., 2008).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHINZ'13, November 15–16, 2013, Christchurch, New Zealand.
Copyright 2013 ACM XXXXXXXXXX.

With the advancement of computer technology, systems have been designed which add extra functionality and therapeutic potential over the original mirror box design. For example, researchers have developed new ways to extend the possibilities of visual illusion using computer video technology and low-cost non-traditional gaming hardware (Regenbrecht, Franz, McGregor, Dixon, & Hoermann, 2011; Regenbrecht et al., 2012; Regenbrecht, McGregor, et al., 2011).

Virtual environments have often been used to supplement upper limb therapy, which often incorporated game-like tasks and situations (see for reviews Cameirao, Badia, & Verschure, 2008; Henderson, Korner-Bitensky, & Levin, 2007; Laver, George, Thomas, Deutsch, & Crotty, 2011). They give the user a sense of presence within an environment that can simulate the real world and real world interactions (Rizzo & Kim, 2005). Virtual reality (VR) environments can be seen as having a number of strengths when used for therapy and rehabilitation. They provide a safe testing and training environment, can enhance motivation using gaming situations, and can provide real time feedback on rehabilitation performance. VR also provides many future opportunities over traditional rehabilitation techniques, such as the ability to include real time and intelligent analysis of data, and the ability to incorporate gaming-industry devices (Rizzo & Kim, 2005; Rose, Brooks, & Rizzo, 2005).

An approach that could combine the benefits of visual illusion with the advantages of virtual reality could be the use of augmented reality technology which “allows the user to see the real world, with virtual objects superimposed upon or composited with the real world” (Azuma, 1997). AR by definition must combine real and virtual elements, in interactive real time, which are registered in three dimensions. This requires a number of technical challenges to be overcome in order for visual information to be rendered correctly in relation to the real scene. Firstly, an AR system must have knowledge of the user's position in space, and of the orientation of the viewing device, as well as knowledge of the real objects within the viewed real-scene, to correctly position virtual objects within that scene (Azuma et al., 2001). Secondly, those virtual objects must be correctly rendered in relation to the real objects in the scene.

Since 2010, the Microsoft Kinect has been used by enthusiasts and academics alike in creating novel (Izadi et al., 2011) and practical (Huang et al., 2011; Chang et al.,

2011) applications. The Microsoft Kinect is a peripheral hardware device initially intended for use with the Xbox 360 gaming console as a full body motion controller. The Kinect device consists of a RGB camera, infrared laser matrix projector and receiver (comprising the depth sensing component), as well as stereo microphones. Both the depth camera and RGB camera run at 640x480 resolution, at a frame rate of 30Hz. The depth sensing ability has especially shown its benefits in augmented and virtual reality applications. Because the Kinect provides information about the precise spatial arrangement of physical objects in the mixed reality scene, developers have now been able to elegantly solve the occlusion problem inherent in mixed reality.

The Kinect and has more recently also been used in motor rehabilitation therapy research. In Chang et al. (2011), a system called “Kinerehab” was proposed and tested in a rehabilitation setting with two participants with motor disabilities. The system was designed to assist therapists working in a public school setting, using the Kinect to determine the quality of the exercise movements performed by the participants. This was achieved by using the skeleton tracking ability of the Kinect to determine joint angle positions and movement quality and by tracking the number of repetitions performed during a session. At the completion of the study it was found that the participants performed significantly better with the Kinerehab intervention than with oral instructions and the absence of technology or visual feedback.

Huang et al. (2011) describe a rehabilitation game which uses a dataglove as well as the Kinect. The jewel thief game, where users stands in front of the Kinect and a laptop while wearing the dataglove, requires users to move their arm to grasp jewels as they approach them on the laptop screen. Once grasped, the user moves the jewel and "drops" it into a basket. Although the system was designed for arm rehabilitation, with the added dataglove, finger joint and hand gestures can also be recorded and implemented as part of the game. The Kinect camera is again used for skeletal tracking in this case. The system collects data from the performed movements for the purpose of medical examination. Although the authors claim it is low cost, they use a data glove in addition to the Kinect. The Kinect in this case is used in a way where the user must be sufficiently far from the device in order for the skeleton tracker to detect the entire upper body, even though the system only focuses on arm rehabilitation. The dataglove is then used to detect and record movements of the hand and fingers. For the purpose of upper limb rehabilitation exercises, as is the focus of this study, a dataglove for detecting small hand and finger movements performed during virtual object grasping is presumably superfluous. By relying solely on a Kinect style motion controller, we can reduce the cost and complexity of such a system even further, and remove the need for a patient with restricted upper-limb movement to don a dataglove.

In summary, the current literature shows that Kinect based systems can be potentially used in therapeutic applications. However, the occlusion problem—i.e., when

we render virtual objects in relation to real world geometry, those virtual objects are naturally superimposed on top of the real geometry, regardless of their intended position—needs to be addressed further in research. In order to allow the rendering of a real hand in front of a virtual object, for example, the system must be aware of the position of the hand in relation to the virtual object. This fundamental challenge in AR is visual occlusion, and its relationship to users’ ability to perform an AR reach-based therapy task is the focus of this study.

METHOD

An experiment was designed to investigate the effects that certain rendering conditions have on users’ experience of an augmented reality post-stroke therapy reaching exercise. These were “correct occlusion (CO),” “virtual always occludes (VO),” and “semi-transparent (ST).” These conditions are shown in Figure 1.




Figure 1. Three modes of occlusion rendering: Correct Occlusion (left), Virtual Always Occludes (centre), Semi-transparent (right)

In particular, the following hypotheses were investigated:

H1: A reach exercise for upper-limb post-stroke motor-rehabilitation can be computerized using a low cost commercially available motion controller without the use of additional data-gloves

H2: The correct visual occlusion rendering condition is superior in terms of perceived task-performance to the “virtual always occludes” and semi-transparent conditions

H3: The correct visual occlusion rendering condition is perceived to be the most correct visual experience compared to “virtual always occludes” and semi-transparent conditions

H4: users prefer to interact with the virtual objects rendered with correct occlusion, compared to “virtual always occludes” and semi-transparent

The experiment was approved by the University of Otago ethics committee at category level B.

Participants

30 participants took part in the experiment, 14 female and 16 male, between 19 and 38 years old ($M = 23.17$, $SD = 3.56$). Participants were recruited from university mailing lists, colleagues, and associates. All were compensated for their time with a \$10 grocery voucher. All participants provided written informed consent.

System

The system used in this study was a combination of two approaches. The physical setup was designed based on

previous work by Regenbrecht et al. (Regenbrecht, Franz, et al., 2011; Regenbrecht et al., 2012; Regenbrecht, McGregor, et al., 2011) which was shown to allow the participants' decoupling from the real hands and supports participants' taking ownership of the computer mediated and manipulated hands on the screen. This was extensively evaluated in (Hoermann, Franz, & Regenbrecht, 2012).

The software used to compute the occlusion effects was based on the work by Clark & Piumsomboon (2011) and is described in detail in, (Piumsomboon, Clark, & Billinghamurst, 2011). The experimental setup is shown in Figure 2. The system consists of the Kinect depth sensor which was positioned above the occlusion space of the box facing downward. A reference marker was placed in the occlusion space to calculate the transform between the depth sensor and the viewing camera.

The simulated "reach" therapy exercise was implemented based on a task described in the Graded Repetitive Arm Supplementary Program (GRASP). GRASP combines various therapeutic exercises for patients to perform during their recovery phase after stroke without the required supervision of a therapist. In a previous study GRASP was shown to be beneficial for stroke patients by improving their arm function during their early rehabilitation after stroke. (Harris, Eng, Miller, & Dawson, 2009)

The therapeutic application, a "ball reaching" task, was designed to display single red virtual spheres within the therapy box. These could be reached and interacted with through the front of the box using the upper limb. A curtain was attached to the front of the box to keep the interaction space from view, forcing the user to rely solely on the monitor display to perform the augmented reality interactions (Figure 3). When a user successfully reached a sphere it would be deflected off the hand. These virtual spheres are placed by the researcher/therapist onto the ground surface of the box at random positions within a predefined "reachable" space. Keyboard controls are used to control the rate of appearance of the spheres and to select one of the three defined occlusion conditions described above (Figure 1).

The system software processes could be divided into three components, which were marker tracking, image processing and graphics rendering.

For marker tracking, we used OPIRA (<http://www.hitlabnz.org/index.php/research/augmented-reality?view=project&task=show&id=47>), an image-based natural feature registration library. The size of the registration marker was calculated in millimetres using the Euclidian distance between corners, and the viewing coordinate system was established at this scale with the origin at the top left corner of the marker. The transformation between the corner positions in the depth sensor coordinate and the viewing coordinate was calculated and stored. With the transformation between the depth sensor and the viewing camera known through the referenced marker, 3D data from the depth sensor could be transformed into coordinate of the user's

viewpoint. Object segmentation could be achieved easily by applying a simple distance threshold on each pixel from the marker plane.

The raw depth image obtained was prone to missing values due to shadowing of the infrared data. To resolve this, missing values were identified and we applied an inpainting algorithm to estimate their values. The depth information from the depth image represented the height of the pixel relative to the referenced marker. The origin had a height equal to zero, so any points above and below the reference marker had positive and negative values respectively. We stored the depth value of any object above the marker in the occlusion space into the vertex array for rendering. OpenCV (<http://opencv.org>) library was used for image processing.

The OpenSceneGraph (<http://www.openscenegraph.org>) framework was used for graphics rendering. The input video image from the viewing camera was rendered as the background, with all the virtual objects rendered on top. At the top level of the scene graph, the viewing transformation was applied such that all virtual objects were transformed if the viewing camera needed to be moved. The hand mesh data were stored in a vertex array with an alpha channel set to zero. This allowed realistic occlusion effect of the hand and virtual objects, while not affecting the users' view of the real environment.




Figure 2. Front view of the therapy system (left) and view from above (right)

Materials

A nine question questionnaire was designed to evaluate participants' experience using the system to perform the therapy task. The first four questions (Q1: "It was easy to reach the spheres;" Q2: "I successfully reached the spheres;" Q3: "It was easy to perform the task;" and Q4: "I could complete the task to my satisfaction") were aimed at evaluating the participants' experience in performing the therapy task itself. The remaining 5 questions (Q5: "I had the impression I could reach the spheres at any time;" Q6: "The actions and reactions of the spheres seemed natural to me;" Q7: "I could tell where the spheres were positioned in space;" Q8: "I had the impression of seeing the spheres as 3D objects;" and

Q9: “I had the impression of seeing the spheres as merely flat images,”) were chosen to evaluate the participants’ perception of the visual environment. These “environmental perception” questions were selected from the Mixed Reality Experience Questionnaire (Regenbrecht et al, 2013). All questions had a 7-point Likert-like response scale, from 1 = *strongly disagree*, to 7 = *strongly agree*. Participants were also asked “Of the 3 sessions, which did you prefer?” at the conclusion of each trial.

Design

A within-subjects experiment was conducted with the three visual occlusion conditions as independent variables. The measured dependent variables were perceived task performance and users’ experience of the visualisations. In addition, the users’ overall preference of the visualization condition was evaluated. The three conditions were presented to the users in randomized and counterbalanced order to account for potential training effects.

Procedure

Participants began the experiment by reading an information sheet and task description. If they agreed to participate, they then filled in and signed a consent form and a demographic survey.

Participants started with a training exercise to get familiar with the system. This was performed under their first prescribed occlusion condition, and they would continue to reach for spheres until they were familiar with the task (generally four to five reaches).

In the main part of the experiment participants performed the following repetitive reaching task. Beginning with their hand resting on the surface of the box at the entrance, they had to wait for a sphere to appear on the surface box and then reach for it. After this they returned their hand to the starting position, and the researcher would “drop” the next sphere into the box (Figure 3). This session continued for two minutes, and was repeated three times, once for each occlusion condition.




Figure 3. Participant (left) using system while the researcher (right) controls the experiment

Participants filled out a questionnaire at the end of each session, reporting their experience on the nine 7-point Likert-like scales. These questions were related to their experience of the task and virtual environment.

The researcher made observations and took notes during the sessions and encouraged participants to vocalise their thoughts and provide feedback after they had experienced each occlusion condition.

After participants had completed all three sessions they were asked to choose their preferred session/occlusion condition for performing the task.

Once all tasks and questionnaires were completed, and feedback obtained, participants were thanked for their time and presented with the gift voucher.

Statistical Analysis

Nonparametric tests were used as the data were found to be not normally distributed according to the performed Shapiro-Wilk W tests of normality ($p > .05$). Friedman’s 2-way ANOVA by ranks was performed across all three conditions in the first instance, and if significance was found, post-hoc Wilcoxon matched-pair signed-rank tests were performed between condition pairs. If not otherwise specified the alpha-level for statistical significance was set to $p < .05$.

RESULTS

The results of the quantitative data collection are presented here. The responses range from 1 (most negative, strongly disagree) to 7 (most positive, strongly agree.) The data analysis is split into three categories: user task perception, user environmental perception, and user response overall.

Overall

Participants’ overall response averaged across all answers was highest in the correct visual occlusion (CO) condition overall ($M = 6.08$, $SD = 0.73$), followed by virtual-always-occludes (VO) ($M = 5.66$, $SD = 0.844$), and finally semi-transparent (ST) ($M = 5.41$, $SD = 0.99$). The overall difference between occlusion conditions for all questions was significant, $X^2(2) = 12.9$, $p = .002$. Post-hoc testing showed a significant difference between CO and VO ($Z = -2.285$, $p = .022$) and between CO and ST ($Z = -3.026$, $p = .002$). The difference between VO and ST was not significant ($Z = -0.765$, $p = .444$).




Figure 4. Mean responses over question sets (+/- SE)

Task Perception

The task perception section of the questionnaire asked participants four questions relating to their own perceived experience while performing the task, including ease of task, perceived success at reaching the spheres, ease of reaching the spheres, and task satisfaction.

Participants' mean responses in the task perception section of the questionnaire were: CO condition ($M = 6.16$, $SD = 0.97$) compared to VO condition ($M = 5.82$, $SD = 1.09$) and ST condition ($M = 5.82$, $SD = 1.31$). These values did not differ significantly as shown in the Friedman 2-way ANOVA by ranks ($X^2(2) = 5.264$, $p = .072$), thus no post-hoc tests were performed.

Environmental Perception

The environmental perception section of the questionnaire asked participants five questions relating to how they experienced the visualisations, including: perceived ability to reach spheres, the naturalness of the reactions of the spheres, perception of the position of the spheres in space, and the degree to which they believed the spheres to be flat or three dimensional.

Participants responses were highest in the CO condition ($M = 6.02$, $SD = 0.71$) compared to the VO ($M = 5.54$, $SD = 0.80$) and ST ($M = 5.08$, $SD = 1.08$) conditions across the environmental perception section of the questionnaire. The differences between conditions was significant ($X^2(2) = 13.78$, $p = .001$). Post-hoc tests showed a significant difference between CO and VO ($Z = -2.441$, $p = .015$), and between CO and ST ($Z = -3.621$, $p < .001$), but not between VO and ST ($Z = -1.551$, $p = .121$).

Of note here is the difference in perception for Q8: "I had the impression of seeing the spheres as 3D objects." Here CO ($M = 6.58$, $SD = 0.620$) was higher than VO ($M = 6.03$, $SD = 0.983$) and ST ($M = 4.74$, $SD = 1.632$). These differences were significant ($X^2(2) = 27.853$, $p < .001$), with CO significantly higher than VO ($Z = -2.512$, $p = .012$); CO significantly higher than ST

($Z = -4.322$, $p < .001$); and VO significantly higher than ST ($Z = -3.577$, $p < .001$).




Figure 5. Mean response for Q8: "I had the impression of seeing the spheres as 3D objects." (+/- SE)

Preference and Qualitative Data

In response to the question: "Of the three sessions, which did you prefer?" 19 said they preferred the correct occlusion condition, 9 preferred the virtual-always-occludes condition, and 2 preferred the semi-transparent condition. Participants found the CO spheres to be comparatively more tangible and lifelike, although a few described the look to be "uncanny valley," close to looking real, but far enough from it to be distracting while performing the task. Participants were often observed trying to "play" with the sphere, often by flicking and jumping it, beyond the task specifications.




Figure 6. Participant occlusion condition preferences

Participants would often notice that the VO and ST spheres would always appear to be above the hand despite knowing they were positioned on the ground surface, and would even reach searchingly within the box to investigate. One participant noted that during the CO session, their hand would fully occlude the sphere leading them to believe they had performed a successful "reach,"

however their hand was simply in front of the sphere and not in contact.

DISCUSSION AND CONCLUSION

In this work we investigated the effects of different occlusion rendering conditions on users' perceived task and environmental perception while performing an augmented reality post-stroke "reach" exercise. This rehabilitation scenario was designed and implemented with a Kinect-based system, and the three occlusion rendering conditions (correct occlusion, virtual-occludes, and semi-transparent) were evaluated during a within-subjects experiment with 30 participants, recording both quantitative and qualitative data.

Participants found that the system was easy to use overall, and that the task was easy to perform across all conditions. The data showed a significant difference between the three conditions in terms of the way participants perceived the visualisations, with correct occlusion producing the most positive response. It could not be concluded, however, that this affected participants' perceived ability to perform the task, therefore the second hypothesis was not supported. The data did support the third and fourth hypotheses.

Limitations of this study include the limited number and diversity of the experimental participants, as well as the limited time available in which to have users experience the occlusion conditions. Preferably, more tasks would have been included to showcase a wider range of therapeutic interaction scenarios, as well as demonstrating a wider range of the possible implications of incorrect visual rendering. More objective measures of task performance would also have been taken.

In particular, this system and experiment only demonstrated one possible form of visualisation, namely spheres in an empty box environment. More complex scenes with many intersecting objects might lead to more insights into the perception of occlusions.

The fact that participants were eager to "play" with the spheres suggest that the technology itself may prove to be a motivational factor for therapy. Future work may include developing "gamification" elements, extending the system to include "grasp" actions—another aspect of post-stroke therapy—and further increasing the quality of the correct visual occlusion.

ACKNOWLEDGMENTS

We thank the participants in our study, the technical and administrative staff who helped us and gratefully acknowledge the funding from the department of Information Science.

This study was part of the MSc thesis work of the first author supervised by the last author.

REFERENCES

Altschuler, E. L., Wisdom, S. B., Stone, L., Foster, C., Galasko, D., Llewellyn, D. M. E., & Ramachandran, V. (1999). Rehabilitation of hemiparesis after stroke with

a mirror. *The Lancet*, 353(9169), 2035–2036. doi:10.1016/S0140-6736(99)00920-4

Azuma, R. T. (1997). A survey of augmented reality. *Presence-Teleoperators and Virtual Environments*, 6(4), 355–385.

Azuma, R. T., Baillot, Y., Behringer, R., Feiner, S., Julier, S., & MacIntyre, B. (2001). Recent advances in augmented reality. *Computer Graphics and Applications, IEEE*, 21(6), 34–47. doi:10.1109/38.963459

Cameirao, M. S., Badia, S. B., & Verschure, P. (2008). Virtual Reality Based Upper Extremity Rehabilitation following Stroke: a Review. *Journal of CyberTherapy & Rehabilitation*, 1(1), 63–74.

Chang, Y.-J., Chen, S.-F., & Huang, J.-D. (2011). A Kinect-based system for physical rehabilitation: A pilot study for young adults with motor disabilities. *Research in Developmental Disabilities*, 32(6), 2566–2570. doi:10.1016/j.ridd.2011.07.002

Clark, A., & Piumsomboon, T. (2011). A realistic augmented reality racing game using a depth-sensing camera. In *Proceedings of the 10th International Conference on Virtual Reality Continuum and Its Applications in Industry* (pp. 499–502). New York, NY, USA: ACM. doi:10.1145/2087756.2087851

Dohle, C., Pullen, J., Nakaten, A., Kust, J., Rietz, C., & Karbe, H. (2009). Mirror Therapy Promotes Recovery From Severe Hemiparesis: A Randomized Controlled Trial. *Neurorehabilitation and Neural Repair*, 23(3), 209–217. doi:10.1177/1545968308324786

Harris, J. E., Eng, J. J., Miller, W. C., & Dawson, A. S. (2009). A Self-Administered Graded Repetitive Arm Supplementary Program (GRASP) Improves Arm Function During Inpatient Stroke Rehabilitation A Multi-Site Randomized Controlled Trial. *Stroke*, 40(6), 2123–2128. doi:10.1161/STROKEAHA.108.544585

Henderson, A., Korner-Bitensky, N., & Levin, M. (2007). Virtual Reality in Stroke Rehabilitation: A Systematic Review of its Effectiveness for Upper Limb Motor Recovery. *Topics in Stroke Rehabilitation*, 14(2), 52–61. doi:10.1310/tsr1402-52

Hoermann, S., Franz, E. A., & Regenbrecht, H. (2012). Referred Sensations Elicited by Video-Mediated Mirroring of Hands. *PLoS ONE*, 7(12), e50942. doi:10.1371/journal.pone.0050942

Huang, M.-C., Chen, E., Xu, W., & Sarrafzadeh, M. (2011). Gaming for upper extremities rehabilitation. In *Proceedings of the 2nd Conference on Wireless Health* (pp. 27:1–27:2). New York, NY, USA: ACM. doi:10.1145/2077546.2077576

Izadi, S., Kim, D., Hilliges, O., Molyneaux, D., Newcombe, R., Kohli, P., ... & Fitzgibbon, A. (2011, October). KinectFusion: real-time 3D reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on User interface software and technology* (pp. 559-568). ACM.

- Laver, K. E., George, S., Thomas, S., Deutsch, J. E., & Crotty, M. (2011). Virtual reality for stroke rehabilitation, *Computers & Graphics*, 36(7), 819–834. doi:10.1016/j.cag.2012.04.012
- McCabe, C. S., Haigh, R. C., & Blake, D. R. (2008). Mirror visual feedback for the treatment of complex regional pain syndrome (type 1). *Current Pain and Headache Reports*, 12(2), 103–107. doi:10.1007/s11916-008-0020-7
- Moseley, G. L. (2006). Graded motor imagery for pathologic pain: A randomized controlled trial. *Neurology*, 67(12), 2129–2134. doi:10.1212/01.wnl.0000249112.56935.32
- Piumsomboon, T., Clark, A., & Billingham, M. (2011). Physically-based interaction for tabletop augmented reality using a depth-sensing camera for environment mapping. In *Proceedings of the 26th International Conference on Image and Vision Computing New Zealand*.
- Ramachandran, V. S., Rogers-Ramachandran, D., & Cobb, S. (1995). Touching the phantom limb. *Nature*, 377(6549), 489–490. doi:10.1038/377489a0
- Regenbrecht, H., Franz, E. A., McGregor, G., Dixon, B. G., & Hoermann, S. (2011). Beyond the Looking Glass: Fooling the Brain with the Augmented Mirror Box. *Presence: Teleoperators and Virtual Environments*, 20(6), 559–576. doi:10.1162/PRES_a_00082
- Regenbrecht, H., Hoermann, S., McGregor, G., Dixon, B., Franz, E., Ott, C., ... Hoermann, J. (2012). Visual manipulations for motor rehabilitation. *Computers & Graphics*, 36(7), 819–834. doi:10.1016/j.cag.2012.04.012
- Regenbrecht, H., McGregor, G., Ott, C., Hoermann, S., Schubert, T., Hale, L., ... Franz, E. (2011). Out of reach? — A novel AR interface approach for motor rehabilitation. In *Mixed and Augmented Reality (ISMAR), 2011 10th IEEE International Symposium on* (pp. 219–228). Basel, Switzerland: IEEE. doi:10.1109/ISMAR.2011.6092389
- Regenbrecht, H., Botella, C., Banos, R., & Schubert, T. (2013). *Mixed Reality Experience Questionnaire v1.0*. Unpublished online document. <http://www.hci.otago.ac.nz/mreq/MixedRealityExperienceQuestionnaireWeb.htm>, last accessed 23/May/2013.
- Rizzo, A. "Skip", & Kim, G. J. (2005). A SWOT Analysis of the Field of Virtual Reality Rehabilitation and Therapy. *Presence: Teleoperators and Virtual Environments*, 14(2), 119–146.
- Rose, F. D., Brooks, B. M., & Rizzo, A. A. (2005). Virtual Reality in Brain Damage Rehabilitation: Review. *CyberPsychology & Behavior*, 8(3), 241–262. doi:10.1089/cpb.2005.8.241
- Yavuzer, G., Selles, R., Sezer, N., Sutbeyaz, S., Bussmann, J. B., Koseoglu, F., ... Stam, H. J. (2008). Mirror therapy improves hand function in subacute stroke: a randomized controlled trial. *Archives of physical medicine and rehabilitation*, 89(3), 393–398.