

Enhancing Depth Perception in Translucent Volumes

Marta A. Kersten, A. James Stewart, Niko Troje, and Randy Ellis *Member, IEEE*

Abstract— We present empirical studies that consider the effects of stereopsis and simulated aerial perspective on depth perception in translucent volumes. We consider a purely absorptive lighting model, in which light is not scattered or reflected, but is simply absorbed as it passes through the volume.

A purely absorptive lighting model is used, for example, when rendering digitally reconstructed radiographs (DRRs), which are synthetic X-ray images reconstructed from CT volumes. Surgeons make use of DRRs in planning and performing operations, so an improvement of depth perception in DRRs may help diagnosis and surgical planning.

Index Terms—Stereo, Stereopsis, X-ray, Radiograph, Volume Rendering

1 Introduction

It can be difficult to perceive three dimensional shape and structure in translucent, volume-rendered data. This paper investigates perceptual cues that might be used to enhance depth perception in such data.

The cues used to perceive depth by humans in everyday life have been thoroughly studied. For computer-generated images, however, what cues and the way in which these cues should be combined to convey depth remains an ongoing topic of research.

With *opaque* volumetric objects, shape and depth can be perceived when perceptual depth cues — such as shadows and shading, occlusion, perspective, and texture gradients — are used [9, 21]. However, very little is known about the effects of these cues on depth perception of purely absorptive media.

A *purely absorptive medium* is one through which light passes with no reflection or scattering. This is approximately what happens with X-ray images of human anatomy. Such images show no solid surfaces and have no depth cues. The lack of solid surfaces and depth cues makes it especially difficult to understand the structure of the volume [1].

To study the effects of different perceptual cues in purely absorptive media, we look at digitally reconstructed radiographs (DRRs). A DRR is a synthetically computed X-ray image that is calculated by integrating the attenuation of light as it passes through a volume toward the viewer (see Figure 1). Each pixel value of a DRR image is a function of the CT values encountered along the projection rays. DRRs are of particular interest because they accurately simulate plain radiographs [20] and, therefore, are used in many medical applications. In intensity-based 2D-to-3D intra-operative registration, DRRs are computed from many different viewpoints to find one that most closely matches a fluoroscopic image [15]. For radiotherapy treatment, DRRs provide reference images to assist in patient positioning [24]. In planning of orthopaedic surgery, DRRs can be used to show intra-articular features not visible in a surface-rendered CT image [16].

Although volume rendering has gained widespread acceptance in the medical community, a DRR rendering may be preferred by radiologists and surgeons (especially orthopaedic surgeons) to a surface (i.e. opaque) rendering of volumetric data as it can show these



Fig. 1. A DRR of a pelvis, generated from CT data. No depth cues are present. A DRR may be rendered as black-on-white, or as white-on-black to appear like a conventional X-ray, as shown above.

otherwise-hidden features. Furthermore, radiologists and surgeons have substantial training and experience interpreting X-ray images, and may be reluctant not to use that knowledge.

However, radiologists and surgeons have very few depth cues to use when working with two dimensional X-ray images. The goal of our work is to determine whether the cues of stereopsis and simulated aerial perspective can provide better depth perception in purely absorptive media. To the best of our knowledge, no studies have considered how these two cues can aid depth perception in such media.

2 Cues for Depth Perception

2.1 Stereopsis

Stereopsis is the perception of depth due to differences in the left and right retinal images. The separation of our left and right eyes (of about 6 cm) creates two slightly different images of the world; points that lie at different depths in a 3D scene are offset from the center of projection by different distances in the retinal images.

It is not obvious that the stereopsis cue is effective when viewing purely absorptive media: In an absorptive medium, the left and right eye can focus on a point within a volume and see two different images. This occurs because the retinal images are produced by the integral along two different rays that *pass through* the focal point toward each of the eyes, as shown in Figure 2.

- Marta Kersten is with the Medical Computing Lab, Queen's University, E-mail: marta@macko.ws.
- James Stewart is with the Medical Computing Lab, Queen's University, E-mail: jstewart@cs.queensu.ca.
- Niko Troje is with the Medical Computing Lab, Queen's University, E-mail: troje@post.queensu.ca.
- Randy Ellis is with the Medical Computing Lab, Queen's University, E-mail: ellis@cs.queensu.ca.

Manuscript received 31 March 2006; accepted 1 August 2006; posted online 6 November 2006.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.




Fig. 2. Does stereopsis work in absorptive media? With absorptive media (left) the two eyes see different images of a particular point because light from behind is attenuated differently along each path. With surface rendered objects (right) the left and right eye see the same feature in the same 3D position, displaced differently in the projections on the left and right retinas.

This is similar to the situation of specular reflections on transparent objects: The left and right eye see the specular spot in different positions on the surface because the spot positions are viewpoint dependent. This suggests that specular highlights alone do not provide surface depth information in either monoscopic or stereoscopic views [11].

Our experiments consider whether stereoscopic viewing can provide depth information in a purely absorptive medium, even though the images projected onto the left and right retinas are different due to the translucent nature of the medium.

Opacity and spatial frequency are two factors that might affect stereopsis in absorptive media. As overall opacity increases, distant features will become more obscured by close features, which might mimic the perceptual depth cue of occlusion. Low spatial frequency may make it easy to detect and track large clusters. On the other hand, high spatial frequency introduces edges which may aid in finding correspondences between the retinal images.

Our experiments considered how stereoscopic rendering affects *relative* depth perception in absorptive media, and how opacity and spatial frequency affect the accuracy and speed of the depth perception. We use the term “stereoscopic” to refer to viewing with 3D glasses.

2.2 Aerial Perspective

Aerial perspective is the perception of depth due to the scattering of light in the atmosphere, as shown in Figure 3. Light that reflects off of a close object does not scatter much before reaching the eye, but the same light that reflects off of a distant object undergoes much scattering, causing the darker portions of the object to be reduced in contrast [6]. Objects that exhibit more contrast are assumed to be closer than objects that do not.

In a purely absorptive illumination model, however, there is no reflectance and no incident illumination. Rather, light is attenuated as it passes through the volume. Therefore, a distant difference in translucency is indistinguishable from an equal, close difference in translucency (see Figure 4).

Our experiments considered how aerial perspective affects relative depth perception in absorptive media. For aerial perspective, we considered *monocular* rendering (i.e. without 3D glasses). The aerial perspective cue was modelled by decreasing contrast as distance from the viewer increased. This had the effect of adding an approximate “uniform scattering” effect, as is present in the atmosphere.

3 Related work

A number of studies have shown that stereoscopic viewing can aid in perceiving depth, grasping, recognizing, or understanding the shape of computer generated objects [9, 25]. Other studies, however, show that the benefit of stereoscopic viewing is task dependent [22] and for certain applications there is no benefit to stereoscopic viewing [5, 23].

3.1 Stereopsis for Diagnosis and Surgery

In the medical field, clinical application studies have shown promising results for stereoscopic volume rendering of images from the X-ray



Fig. 3. An example of perceived depth from aerial perspective. Licensed from dreamstime.com.




Fig. 4. (a) In a scattering medium, incident light reflected off a close object or feature does not scatter as much as that off a more distant object. Scattering causing dark portions of the distant object to appear lighter. (b) In a purely absorptive medium, there is no incident light and the difference in adjacent translucencies (e.g. $|t_1 - t_2|$) provides no depth cue.

domain. One study considered the effectiveness of stereoscopic imaging in finding abnormalities in computer-generated mammograms [8]. Mammograms (the standard imaging modality for screening and diagnosis of breast cancer) are devoid of depth cues, making it difficult to detect abnormalities. The study showed that stereomammography improves the detection of certain tissue abnormalities, which can lead to early detection and better diagnosis of breast cancer.

In laparoscopy, a minimally invasive surgical technique used to diagnose and treat a range of abdominal or pelvic problems, surgeons work with 2D video pictures provided on a monitor. It is known that, because of the need for video imaging, there are performance limits to traditional laparoscopic systems due to the 2D representation of the 3D anatomy. Thus, there is a need to develop some mechanism to improve depth perception [5]. Studies, however, have provided a mixed picture of the usefulness of stereoscopic imaging for laparoscopy. Some studies [7, 5] have demonstrated no statistical advantage, whereas others [3, 13] showed a significant improvement in speed and outcome of laparoscopic tasks. In monocular systems, other available cues (such as motion parallax, relative position, occlusion, perspective, and lighting) have been used to compensate for the loss of depth perception produced by the use of 2D images [5].

Mora and Ebert [18] studied whether stereoscopic rendering using

“order independent volume rendering” methods, including maximum intensity projection (MIP) and X-ray projection (DRR), can provide enough information to understand volumetric data. While no formal user studies were done, the authors’ experience was that stereoscopic rendering with these methods allowed the viewer to better understand the volumetric data.

Calhoun *et al.* [4] stated that stereoscopic viewing may provide better visualization of the 3D structure of complex anatomy, such as spinal data. Their preliminary results suggest that both radiologists and non-radiologists prefer stereoscopic viewing of medical data sets to conventional monocular display.

3.2 Transparency

Transparency is often used in medical and scientific imaging as a way to enhance shape and depth by superimposing multiple transparent layers of information [11, 14]. In medical imaging, transparency can be used to integrate images from different modalities. For example, 2D anatomical CT data can be overlaid with fMRI data in such a way that one can see through the functional image to the structural image [14]. In radiation treatment planning, transparency may help to more effectively visualize radiation dosages with respect to a patient’s anatomy. Kasrai *et al.* [14] performed a number of psychophysical experiments to study transparency with stereoscopic viewing, transparency and spatial frequency, and multiple surface transparency.

To date, research on the perception of transparency and use of transparency to enhance depth perception has focused on viewing an opaque object through a transparent layer [17, 25] where contour junctions exist [2]. In these experiments, at least three layers exist: a background; an opaque surface; and a transparent layer covering these surfaces [12]. To our knowledge, no studies on perception in purely absorptive media (where no opaque objects exist and there is no presence of partial occlusion) have been done.

4 Methodology

Our test data consisted of a vertical cylindrical volume that rotated about its vertical axis, as shown in Figure 5. Test subjects were asked to determine the direction of rotation ... that is, whether the front surface was rotating right-to-left or left-to-right. A DRR image was computed by backlighting the cylinder and computing the attenuation of light as it passed through the cylinder. The cylinder was rendered under orthographic projection to remove any perspective cues.

In a monocular orthographic view with no depth cues, the rotation direction is ambiguous because the order of the attenuation coefficients along each ray toward the eye has no effect on the image, and subjects are expected to have 50% accuracy.

Subjects pressed either the left or the right arrow key on the computer keyboard to indicate their classification. Pressing the key recorded the response and, after a time interval of eight seconds during which time subjects were shown a blank screen, the next stimulus was presented. Subjects were timed and were asked to respond as accurately, but also as quickly, as possible. Preliminary tests showed that subjects performed better with black-on-white DRR rendering, so this was used. (Note that Figure 1 is white-on-black.)

4.1 Apparatus

Images were generated at greater than 15 frames per second on a 2.8GHz Pentium 4 processor. They were displayed on a ViewSonic Professional Series P95F+ CRT display and viewed by the subjects through e-Dimensional Wireless 3D LCD shutter glasses (Florida, USA). The P95F+ display, with PerfectFlat screen technology, displays distortion-free (stereoscopic) images. The window size was 300x410 pixels and the stimulus occupied 260 by 285 pixels in the center of the window. Subjects sat approximately 50 cm away from the monitor. An interpupillary distance of 6 cm was used to create the stereo pairs.

In order to present flicker-free images to subjects, two things were done. First, the maximum monitor refresh rate 100Hz (50Hz per eye) was used. Second, each experiment was conducted in a dark room with no fluorescent lights which could have interfered with the infrared




Fig. 5. Experimental setup: The cylinder model was orthographically projected onto the screen and subjects were asked to determine the direction of rotation.

receptor of the glasses and caused flickering. No subjects reported flickering.

4.2 Illumination Models

For DRR rendering, the radiance visible to the eye coming along ray ω was defined as

$$L(\omega) = L_0 e^{-\int_{s_0}^{s_1} \tau(s) ds} \quad (1)$$

where the ray through the volume is parameterized between s_0 and s_1 , the attenuation coefficient is $\tau(s)$, and L_0 is the radiance of the backlight.

To implement contrast reduction, we used the general volume rendering integral

$$L(\omega) = \int_{s_0}^{s_1} C(s) \tau(s) ds e^{-\int_{s_0}^s \tau(u) du} \quad (2)$$

where $C(s)$ is the “radiance density.” To get the purely absorptive model of Equation 1, $C(s)$ is set to zero everywhere except where the backlight enters the volume (i.e. $C(s_1) \tau(s_1) ds = L_0$).

Contrast reduction is then achieved by shifting radiance densities, $C(s)$, toward white from their usual solid black. At distance d from the viewer, normalized so that $d \in [0 \dots 1]$ for all points in the volume, we set

$$C(s) = (1 - k) d \quad (3)$$

where k was a “contrast factor”. This is identical to the OpenGL fog model.

For $k = 0$, points near the back of the volume emit the same radiance as the backlight, resulting in complete loss of contrast at the back. For $k = 1$, points near the back emit no radiance, resulting in no change in contrast. For any value of k , points at the front did not change in contrast. To study contrast reduction, k , was varied from 0.75 to 1.0 in increments of 0.05.

4.3 Stimuli

The attenuation coefficients of the volume inside the cylinder were generated with a Perlin noise function [19], which permitted us to vary

the overall opacity and spatial frequency of the medium. The attenuation coefficient, τ , at point x inside the volume was defined by a sum of scaled harmonics of a predefined random noise function, $N(x)$:

$$\text{Perlin}(x) = \sum_{i=0}^{n-1} \frac{N(b^i x)}{a^i} \quad (4)$$

where $1/a$ is the persistence (the relative amplitude between adjacent harmonics) and b is the relative frequency between adjacent harmonics. $N(x)$ is created with a seeded random number generator. For the purpose of our experiments we set $n = 2$ and $b = 2$. To study the effect of spatial frequency, the persistence value a was varied from one to five. With higher persistence, more weight is given to higher frequencies.

The cylinder surface was modelled separately from the interior so that the interior opacity could be varied without changing the surface opacity (Figure 6). At a particular point on the interior, the opacity was varied using a parameterized transfer function, $f_v(\tau)$:

$$\begin{aligned} d\alpha &= ds f_v(\tau) \\ &= ds \begin{cases} \tau + \tau(2v-1) & \text{for } v \leq \frac{1}{2} \\ \tau + (1-\tau)(2v-1) & \text{for } v > \frac{1}{2} \end{cases} \end{aligned} \quad (5)$$

where ds is the infinitesimal distance along ω in Equation 2, and $d\alpha$ is the opacity of that infinitesimal segment.




Fig. 6. The cylinder surface (a) was modelled separately so that the interior opacity could be varied without changing the surface opacity. The cylinder surface was texture mapped with Perlin noise and the attenuation coefficients of the volume inside the cylinder (b) were generated with a Perlin noise function.

To study the effect of opacity, we used values of v from 0.0 (completely transparent) to 0.9 (almost opaque). The same Perlin noise was used for the surface and for the interior volume.

Perlin noise and the simple cylindrical stimulus were used to avoid bias from any domain-specific knowledge that the test subjects might have. Subjects with a medical background, for example, might perform better with anatomical shapes than would other subjects.

4.4 Stereoscopic Experiments

Fifteen subjects were each presented with 100 instances of rotating cylinders. Five opacity values, $v = 0.0, 0.1, 0.3, 0.6,$ and 0.9 (from equation 5), and five persistence values, $a = 1, 2, 3, 4,$ and 5 (from equation 4), were considered. Opacity was varied by altering the density of Perlin noise within the cylinder. Each combination of opacity and spatial frequency were shown to the subjects, twice in mono and twice in stereo (one time in each direction, left-to-right and right-to-left, in order to avoid personal bias). The stimuli were presented randomly, interleaving stereoscopic and monoscopic images and the order

in which instances were presented was randomized for each subject. Subjects wore the stereoscopic glasses at all times so that the absence of glasses alone did not allow them to distinguish whether they were viewing monoscopic or stereoscopic images. Subjects were evaluated on correctness and decision response time.

4.5 Simulated Aerial Perspective Experiments

Fifteen different subjects were presented with 24 cases of rotating cylinders where contrast on the back surface and within the cylinder was reduced, as described in Equation 3. The cylinder interior was rendered with a low opacity ($v = 0.1$) and the same low frequency Perlin noise ($a = 1, b = 2, n = 2$) was used. Six contrast values (k in Equation 3), each shown four times, ranging from 0.75 to 1.0 in 0.05 increments, were evaluated.

5 Results and Discussion

For the experiments, we measured each subject's classification correctness and decision response time, then analyzed the data using an analysis of variance (ANOVA) test.

5.1 Stereoscopic Experiments

A one-way ANOVA test showed that viewing mode (monoscopic or stereoscopic) significantly affected classification correctness ($p < 0.001$). Mean classification correctness over all persistences and opacities for monocular viewing was 51.7% (SE 1.5%), which, as expected, was near the chance value of 50%; stereopsis improved overall classification correctness to 80.1% (SE 1.6%). The results for mean classification correctness are plotted in Figure 7, where correctness is a function of opacity. The results for mean classifications where correctness is a function of persistence are plotted in Figure 8.

These results are strong evidence that stereoscopic rendering provides good relative depth perception in purely absorptive media.

5.1.1 Opacity

A one-way ANOVA test showed that opacity has no significant effect on classification correctness within absorptive materials ($p = 0.80$). This is apparent in Figure 7, where the mean classification correctness is essentially unchanged for various opacity values. We analyzed monocular viewing and stereoscopic viewing separately, to determine if there are differences for the two presentation methods.

For monocular viewing, a one-way ANOVA test showed that opacity has no significant effect on classification correctness, regardless of the persistence values used ($p = 0.98$). These results are plotted in Figure 9.

For stereoscopic viewing, a one-way ANOVA test showed that opacity also has no significant effect on classification correctness, regardless of the persistence values used ($p = 0.37$). These results are plotted in Figure 10.

5.1.2 Spatial Frequency

Spatial frequency has no significant effect on classification correctness ($p = 0.83$). These results are presented in Figure 8, where the mean classification correctness as a function of persistence are plotted. We analyzed the presentation modes separately and determined that there were no differences for either monocular viewing ($p = 0.41$) or for stereoscopic viewing ($p = 0.46$).

Plots of classification correctness as a function of persistence, for all values of opacity, are shown in Figure 11 for monocular viewing and in Figure 12 for stereoscopic viewing.

5.1.3 Response Times

Interestingly, stereoscopic response times (mean = 6.23 secs, SE = 0.24) were similar to those of monocular response times (mean = 5.93 secs, SE = 0.26) although a reasonable assumption would be that stereoscopic viewing would reduce response times, as it should make the decision task easier. ANOVA test results showed that there was no significant difference between response times for the two viewing groups ($p = 0.64$). We posit two possible explanations for similar response times. One possibility is that there was a latency in fusing the




Fig. 7. Mean classification correctness as a function of opacity. As expected, mean classification correctness for monocular viewing was near 50% (SE 1.5%). Mean classification correctness increased to 80.1% (SE 1.6%) with stereoscopic viewing. The standard errors, 1.5% and 1.6% respectively, are too small to appear clearly in the plots.




Fig. 8. Mean classification correctness as a function of the persistence of spatial frequencies. Mean classification correctness increased from 51.7% (SE 1.5%) for monocular viewing to 80.1% (SE 1.6%) for stereoscopic viewing.




Fig. 9. Mean classification correctness as a function of opacity, for various persistence values, in **monocular viewing**. Persistence has no significant effect on classification correctness.




Fig. 10. Mean classification correctness as a function of opacity, for various persistence values, in **stereoscopic viewing**. As for monocular viewing, opacity has no significant effects on classification correctness.




Fig. 11. Mean classification correctness as a function of persistence, for various opacity values, in **monocular viewing**. Persistence has no significant effect on classification correctness.




Fig. 12. Mean classification correctness as a function of persistence, for various opacity values, in **stereoscopic viewing**. Persistence has no significant effect on classification correctness.




Fig. 13. Mean response time as a function of opacity. Opacity significantly affected response times of the two viewing modes.




Fig. 14. Mean response times as a function of spatial frequency. There was an unrecognizable effect of response times with changing persistence.

stereoscopic pairs when they first appeared. Another possibility is that the monocular cases were ambiguous and thus led subjects, upon recognizing the presentation mode as being monocular, to quickly choose a direction based on personal bias.

The ANOVA test showed that both opacity and spatial frequency had an effect on response time. Opacity had a significant effect on response time ($p = 0.011$). Increasing opacity was associated with longer response times, as plotted in Figure 13. We speculate that this was because it was more difficult, with increased opacity, for the subject to find features to follow, either in the volume or on the surface, in order to determine the direction of rotation.

There was a marginal effect of spatial frequency on response time ($p = 0.065$). Response times change with persistence, as shown in Figure 14, but not in a recognizable way.

5.2 Simulated Aerial Perspective Experiments

Results showed that at contrast factors between 0.75 and 0.85, the effectiveness of contrast reduction matched that of stereopsis as a depth cue. Mean classification correctness with contrast reduction between factors of 0.75 and 0.85 was 81.1% (SE=4.9%), and with stereoscopic viewing for a low opacity cylinder, $v = 0.1$ was 81.3% (SE=0.8%). Furthermore, at a contrast factor of 0.75 the effectiveness of this depth cue exceeded that of stereopsis; classification correctness was 91.7% (SE=3.135%).

A one-way ANOVA test showed a significant effect of contrast reduction on classification correctness ($p < 0.001$). Classification correctness was about 50% with no contrast reduction ($k = 1.0$) but

quickly improved with even slightly reduced contrast. Figure 15 shows a clearly increasing trend in classification accuracy as the contrast of the cylinder was decreased. Lower contrast factors, however, resulted in a perceptible loss of information in the distant parts of the volume.




Fig. 15. Mean classification correctness as a function of contrast. Even slightly reduced contrast improved classification correctness.

6 Conclusion and Future Work

We have presented a discussion of stereoscopic volume rendering of purely absorptive media, an examination of the perception of relative depth in DRRs using stereopsis and simulated aerial perspective, and psychophysical experiments to investigate the effectiveness of these cues.

The experimental results strongly suggest that relative depth can be perceived using stereoscopic viewing in purely absorptive media. Interestingly, the data's opacity and spatial frequency have little effect on accuracy with stereoscopic rendering. It was also interesting to find that simulated aerial perspective provides a cue that can equal that of stereopsis, at levels of contrast reduction that seem not to obscure parts of the data. But we have not determined experimentally the extent of data loss that comes with contrast reduction.

To extend the stereoscopic work presented above and to evaluate the scalability of our results, we proposed a quick and easy stereoscopic DRR technique to measure acetabular coverage (i.e. how much of the acetabulum or hip socket is covered by the femoral head) in order to diagnose hip dysplasia [10]. In this study, one nonexpert observer viewed 20 pelvic CT scans (10 preoperative and 10 postoperative) of patients treated using computer-assisted periacetabular osteotomy. The observer picked points on the DRRs, separately outlining the femoral head and the acetabulum. The outlining was performed separately in monocular and stereoscopic rendering modes for each of the 20 CT data sets. A two-way independent sample t -test showed that whereas the stereoscopic technique correlated to a currently accepted method for measuring acetabular coverage ($t = 2.148$, $p = 0.038$), the monocular equivalent did not correlate ($t = 0.825$, $p = 0.415$). These results suggest that stereoscopic viewing of DRRs is a viable technique for measuring acetabular coverage. Furthermore, these results support our initial experiments in showing that stereoscopic rendering of DRRs provides more information than monocular rendering.

Future work will consist of studies with surgeons to determine how depth perception improves their performance at particular tasks. We will also measure experimentally the loss of data that comes with contrast reduction, and will test different contrast reduction functions.

Acknowledgements

This work has been supported by the Ontario Center of Excellence and the Natural Sciences and Engineering Research Council of Canada. We would like to thank the reviewers for their helpful comments.

References

- [1] Stephen J. Adelson and Charles D. Hansen. Fast stereoscopic images with ray-traced volume rendering. In *Proceedings of the 1994 symposium on Volume visualization*, pages 3–9. ACM Press, 1994.
- [2] B. L. Anderson. A theory of illusory lightness and transparency in monocular and binocular images: the role of contour junctions. *Perception*, 26:419–453, 1997.
- [3] G.F. Buess, P. van Bergen, W. Kunert, and M.O. Schurr. Comparative study of various 2-D and 3-D vision systems in minimally invasive surgery. *Chirurg*, 67(10):1041–6, Oct 1996.
- [4] P. Calhoun, B. Kuszyk, D. Heath, J. Calrey, and E. Fishman. Three-dimensional volume rendering of spiral ct data: Theory and method. *RadioGraphics*, 19(3):745–764, 1999.
- [5] L. Way F. Tendick, S. Bhojwala. Comparison of laparoscopic imaging systems and conditions using a knot-tying task. *Computer Aided Surgery*, 2:24–33, 1997.
- [6] James J. Gibson. *The Perception of the Visual World*. The Riverside Press, 1950.
- [7] G. B. Hanna, S. M. Shimi, and A. Cuschieri. Randomized study of influence of two-dimensional versus three-dimensional imaging on performance of laparoscopic cholecystectomy. *Lancet*, 351:248–251, 1998.
- [8] J Hsu, C. F. Babbs, D. M. Chelberg, Z. Pizlo, and E. J. Delp. Preclinical roc studies of digital stereomammography. *IEEE Transactions on Medical Imaging*, 14(2):318–327, 1995.
- [9] G. Hubona, P. Wheeler, G. Shirah, and M. Brandt. The relative contributions of stereo, lightin and background scenes in promoting 3d depth visualization. *ACM Transaction on Computer-Human Interaction*, 6(3):214–242, 1999.
- [10] J. Inoue, M. Kersten, B. Ma, J. Stewart, J. Rudan, and R. Ellis. Fast assessment of acetabular coverage using stereoscopic drs. *Medicine Meets Virtual Reality 14: Accelerating Change in Healthcare: Next Medical Toolkit*, 219:225–227, 2006.
- [11] V. Interrante, H. Fuchs, and S. M. Pizer. Conveying the 3d shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2), 1997.
- [12] S. P. Johnson and R. N. Aslin. Infants' perception of transparency. *Developmental Psychology*, 36(6):808–816, 2000.
- [13] I. C. Jourdan, E. Dutson, A. Garcia, T. Vleugels, J. Leroy, D. Mutter, and J. Marescaux. Stereoscopic vision provides a significant advantage for precision robotic laparoscopy. *British Journal of Surgery*, 91:879–885, 2004.
- [14] R. Kasrai, F. A. A. Kingdom, and T. Peters. The psychophysics of transparency in medical images. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI'99, Second International Conference, Cambridge, UK Proceedings*, volume 1679, pages 726–733. Springer, 1999.
- [15] David LaRose. *Iterative X-ray/CT Registration Using Accelerated Volume Rendering*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2001.
- [16] P. Messmer, G. Long, N. Suhm, M. Hehli, J. Wirth, P. Pregazzoni, and A. L. Jacob. Three-dimensional fracture simulation for preoperative planning and education. *European Journal of Trauma*, 27(4):171–177, August 2001.
- [17] F. Metelli. The perception of transparency. *Scientific American*, 230:90–98, 1974.
- [18] B. Mora and D. S. Evert. Instant volumetric understanding with order-independent volume rendering. *Computer Graphics Forum*, 23(3), 2004.
- [19] K. Perlin. An image synthesizer. *Computer Graphics (SIGGRAPH Proceedings)*, 19(3):287–296, July 1995.
- [20] C. Robertson, R. E. Ellis, T. Goetz, W. Gofton, P. V. Fenton, C. F. Small, and D. R. Pichora. The sensitivity of carpal bone indices to rotational malpositioning. *J Hand Surg*, 27A(3):435–442, 2002.
- [21] Leonard C. Wanger, James A. Ferwerda, and Donald P. Greenberg. Perceiving spatial relationships in computer-generated images. *IEEE Comput. Graph. Appl.*, 12(3):44–51, 54–58, 1992.
- [22] C. Wickens and Y. Liu. Use of computer graphics and cluster analysis in aiding relational judgement. *Human Factors*, 34(2):165–178, 1992.
- [23] C. Wickens and L. Thomas. Effects of CDTI display dimensionality and conflict geometry on conflict resolution performance. In *Proceedings of the 13th International Symposium on Aviation Psychology*, 2005.
- [24] C. Yang, M. Guiney, P. Hughes, S. Leung, K.H. Liew, J. Matar, and G. Quong. Use of digitally reconstructed radiographs in radiotherapy treatment and verification. *Australasian Radiology*, 44(4):439–443, November 2000.
- [25] Shumin Zhai, William Buxton, and Paul Milgram. The "silk cursor": Investigating transparency for 3D target acquisition. In *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, volume 1, pages 459–464, 1994.

