

Study of the Perception of 3D Spatial Relations for a Volumetric Display^a

Christoph Hoffmann¹, Zygmunt Pizlo², Voicu Popescu¹, and Paul Rosen¹

¹Computer Sciences, ²Psychology

Purdue University, 250 N University Street, West Lafayette, IN 47907-2066

ABSTRACT

We test perception of 3D spatial relations in 3D images rendered by a 3D display and compare it to that of a high-resolution flat panel display. Our 3D display is a device that renders a 3D image by displaying, in rapid succession, radial slices through the scene on a rotating screen. The image is contained in a glass globe and can be viewed from virtually any direction. We conducted a psychophysical experiment where objects with varying complexity were used as stimuli. On each trial an object or a distorted version is shown at an arbitrary orientation. The subject's task is to decide whether the object is distorted or not, under several viewing conditions (monocular/binocular, with/without motion parallax, and near/far). The subject's performance is measured by the detectability d' , a conventional dependent variable in signal detection experiments. Highest d' values were measured for the 3D display when the subject is allowed to walk around the display.

^a Work supported in part by NSF grants DMS-0138098, DCNS-0216131, DHER-0227828, DSC-0325227, DCMS-0443148, EEC-0227828, EIA-0216131, and ACI-0325227. Hoffmann is also supported in part by an IBM Faculty Award.

1. INTRODUCTION

Traditionally, computer graphics systems compute a 2D image by projecting a 3D scene from the view desired by the user onto an image plane. Here and throughout the paper we use the term scene to designate the data visualized by the graphics application. Depending on the application, the scene can be geometric or abstract, it can be a replica of an actual real-world scene or a fictive, future, or past environment, it can be at nano or macro scale, and it can be a single object or a complex ensemble that surrounds the user. The 2D image so computed is then shown on a 2D display. This approach has three fundamental problems, for which there is no perfect solution yet.

First, the system has to give the user exploring the scene an intuitive way of specifying novel desired views. Keyboards, joysticks, and trackers designed to tell the graphics system the user's position and view direction are often non-intuitive, imprecise, and/or bulky and consequently do not allow the user to navigate freely in the 3D scene. All these interfaces scale very poorly with the number of users.

Second, the 2D image has to be recomputed for each novel view. Computer graphics algorithms and their hardware implementation have progressed spectacularly. Radiosity, photon mapping, ray tracing and image-based rendering techniques reuse precomputed or pre-acquired global illumination solutions to produce compelling visual experiences at interactive rates. However, when the surface appearance depends decisively on the view point as in the case of, for example, refractive or reflective surfaces, when the scene objects are dynamic, or when the lighting conditions change, interactively computed imagery falls short of realism. In such conditions rendering images that are mistaken for photographs still require more time than available in an interactive rendering system.

Third, the user should be allowed to take advantage of binocular stereo vision. Existing head-mounted displays are bulky, have low resolution and limited field of view. Active stereo glasses and polarized passive glasses are uncomfortable, and produce little and sometimes incorrect disparity.

One way of avoiding these problems is to display a sculpture of light that is the exact replica of the 3D scene to be rendered. With such a 3D image, the user requests novel views naturally by gaze, by head movement, and by walking. For static scenes at least, once the 3D image is computed, no per-view rendering is required. Lastly, the user will perceive correctly different images with each of his eyes. Such a 3D display technology does not yet exist, but several technologies have demonstrated promising results. In this paper we analyze the perception of 3D spatial relations in images generated by a rotating-screen 3D display. We briefly review the main 3D display technologies in the next section. Then we describe the psychophysical experiment we conducted. Section 4 discusses the results of the experiment and Section 5 gives possible directions of future work.

2. THREE DIMENSIONAL DISPLAYS

The Perspecta from Actuality Systems (Perspecta 2003) is a 3D display device that allows a volumetric image to be viewed by several observers simultaneously from any direction. A globular glass housing (Figure 1), with a 60" diameter, contains a rotating, semi-transparent screen of approximately 10" diameter on which 198 images are displayed by an internal DLP projector for each rotation of the screen. The screen rotates at 24Hz. Each image has a resolution of 768 by 768 pixels and represents a slice of the 3D scene to be displayed. The eye of the observer combines these slices and perceives a complete volumetric image. The globe is free-standing, so the image is visible from all directions except from underneath (a $360^\circ \times 270^\circ$ field of view). The volumetric image comprises approximately 100 million voxels, each with 64 possible colors. Thus the technology puts great demands on the bandwidth between the host computer and the display, which limits the ability of displaying moving scenes.

There are several alternatives to this technology for stereoscopic images. By far the most common stereoscopic device is based on presenting two separate images to each eye that account for parallax. In active systems, the images are presented alternately and a shuttered eye-glass device blocks one of the eyes

accordingly. In passive systems the two images are color coded or presented with polarized light, and spectacles with suitable filters ensure that each image reaches the intended eye. Another passive technology is based on lenticular gratings [3]. When the observer is positioned at the correct distance from the device, each eye sees a separate image on an LCD screen. The lenticular grating separates alternating image columns, so separating an interleaved stereoscopic pair of images. Such stereoscopic image pairs are a simpler technology to drive since the necessary data transmission rates are equivalent to normal 2D images.

An emerging technology is based on holography. At this time, there are no commercially available devices that create moving holographic imagery. There are, however, devices for creating holographic still images. As with the Perspecta, holographic moving images demand bandwidth and computational power that appears to exceed current technology.

Given the unique technological niche the Perspecta is situated in, we investigate the question whether the 3D images so generated confer a particular cognitive advantage when judging spatial relationships. If we consider images of familiar objects, such as houses, cars, and furniture, can an observer tell easily whether the object is correctly displayed or is distorted? To investigate this question, we conducted a series of experiments on the device. Briefly, in one experiment the object is displayed and the observer is encouraged to view it from several directions of his or her choosing before deciding whether the object is distorted. In another experiment, we fix the observer and investigate whether a binocular view is superior to a monocular view, and at what distance. Finally, we run the experiment also using a nonstereoscopic LCD display (the IBM T221 monitor with a pixel resolution of 3840 by 2400 on a 12" by 19" screen (IBM T221 2003)).

The last experiment is meant to provide a comparison to today's omnipresent display technology. The limitations of the Perspecta described above prevent us from comparing the two technologies at a fundamental level. While the LCD technology is quite mature, the Perspecta technology is still in its infancy stage, equivalent to the pre-VGA stage of 2D displays. The color depth, brightness, and resolution of the Perspecta technology are likely to improve. However, an inherent limitation of the technology is that it can only display

light, and not hide it. Since opacity cannot be generated, no part of the image can truly occlude any other part. Perspecta's inability to display only the first surface further complicates the comparison to LCD displays.

We considered the option of using opacity data sets for the experiments: these are naturally shown on the Perspecta, and they can be converted to 2D images to be shown on the LCD. We dismissed this option for two reasons. First, producing an *equivalent* volume visualization on the LCD is difficult. One approach is to use a state of the art volume rendering algorithm. Another approach is to computer simulate the Perspecta display using a conventional 3D computer graphics API such as OpenGL [8]. A third approach is to photograph the Perspecta and display the photographs on the LCD. The particulars of these approaches would interfere with the comparison. The second reason for not using volume visualizations in the experiments is the need of presenting simple and/or familiar shapes to the subject. This enables formulating a clear, unambiguous task, regardless of the subject's training in the application domain that provided the stimuli.

We opted for using objects represented by their surfaces, and for showing the images on the LCD at their best, as described in the next section. As 3D display technology evolves, future experiments will zero in on the fundamental differences between 3D and 2D displays, free of the noise characteristic to the early stages of an emerging technology.

3. PSYCHOPHYSICAL EXPERIMENT

In order to evaluate capabilities and limitations of Perspecta, we designed an experiment on the visual perception of spatial relations in a 3D scene. Perception of spatial relations in 3D scenes refers to a wide range of visual abilities: perception of distances among objects, distance from the observer, motion in depth, size and shape of an object. We decided to test the subject's ability to perceive 3D shapes, because shapes, as

opposed to 3D distances or sizes of objects can be recognized and discriminated reliably by human observers (see [13] and [12]).

In order to minimize the role of mental processes that are not directly related to visual perception, we designed an experiment in which the subject was presented with one stimulus on each trial and was asked to make a judgment about the shape of this stimulus. Other shape perception tasks like shape discrimination would require the subject to compare two or more shapes. Such comparison is likely to involve memory: the subject has to first look at one object, construct a mental representation of this object, and then compare this representation with the percept of the other object. Results of such a task are therefore likely to conflate the role of memory with the role of visual perception.

3.1. Subjects

The initial subjects in this experiment were the four authors. ZP is experienced as a subject in psychophysical experiments on shape, including binocular experiments. The other three subjects did not have experience in psychophysical experiments, except for a couple of preliminary sessions of this experiment. PR received more experience with the stimuli due to the fact that he was directly involved in designing the stimuli.

Using a small number of subjects and including the authors among the subjects in psychophysical experiments is commonly accepted in the psychophysics community (the reader may verify this claim by consulting leading perception journals like *Vision Research* and *Perception & Psychophysics*). First, it is well established that all fundamental mechanisms underlying visual perception, including shape, are innate ([7], [14]). As a result, the magnitude of individual differences is extremely small. In other words, we all see things the same way, regardless of where we were born and raised. This means that results from just a few subjects who have normal vision is representative for the entire human population.

Second, testing a subject who knows the hypotheses behind the experiment, not to mention, testing the person who formulated the hypotheses and designed the experiment in the first place, leads to valid data as long as a reliable psychophysical method is used. The main problem in studying perception is related to the

fact that the subject's response is a result of two factors: the percept itself, and a decision, which mediates between the percept and the behavioral response. There are psychophysical methods, which allow measuring independently the percept and the decision. This method is called Signal Detection Experiment (SDE) (Green & Swets, 1966). The main elements of SDE are described in the next subsection.

To verify that the data collected using the authors as subjects is reliable, we reran the experiment with two subjects naïve to the task. The additional subjects were students from our computer graphics laboratory. They did not know the hypothesis being tested, and they participated in psychophysical experiments for the first time. Their performance was comparable to that of the initial 4 subjects, as shown in Section 4.

3.2. Signal Detection Experiment

In a signal detection experiment two types of stimuli are used: S_1 , called *noise* and S_2 , called *signal* (where $S_2 > S_1$). The subject produces one of two responses, R_1 and R_2 , respectively. Each of the two stimuli is presented 50% of the time in a random order: the subject does not know which of the two stimuli is presented on a given trial. Response R_1 is the correct response when S_1 is presented, and response R_2 is the correct response when S_2 is presented. The percept of a stimulus is represented by a random variable X . When S_1 is presented, X is subject to a normal distribution with the probability density function $N(\mu_1, \sigma^2)$, and when S_2 is presented, X is subject to a normal distribution with the probability density function $N(\mu_2, \sigma^2)$ ($\mu_2 > \mu_1$).

The subject's ability to detect the signal is related to $(\mu_2 - \mu_1) / \sigma$. This ratio is called detectability and is denoted by d' . For a given σ , when the difference $(\mu_2 - \mu_1)$ is greater, it is easier to tell the two stimuli apart. Similarly, for a given difference $(\mu_2 - \mu_1)$, when σ is smaller, it is easier to tell the two stimuli apart. It is important to note that d' is a measure of the percept unconfounded with the subject's bias towards either of the two responses R_1 or R_2 . This is obvious because d' is a function of the parameters that characterizes the perceptual representation of the stimuli (μ_1, μ_2, σ) , but not the actual responses.

Now the main challenge is to show how to estimate d' from the responses R_1 and R_2 . Let $h=P(R_2|S_2)$ be the *hit rate* and $f=P(R_2|S_1)$ be the *false alarm rate*. It is assumed that the subject produces responses R_1 and R_2 based on a subjectively (and arbitrarily) adopted criterion k for the magnitude of the percept X . The decision criterion is as follows: if $X>k$, respond R_2 , otherwise, respond R_1 . Let $\Phi(z)$ be the cumulative function of the standard normal distribution, and $\phi(z)$ its density function. Let $z_p=\Phi^{-1}(p)$, be the inverse of the cumulative distribution function. It is easy to show that (Green & Sweats, 1966):

$$d' = z_h - z_f \tag{1}$$

In practice, one does not know h and f , but only their estimates. So, z_h and z_f must be computed from estimated h and f . To keep the notation simple, we will use symbols h , f , and d' to represent not only the parameters, but also their estimators. This should not produce confusion. Hit rate and false alarm rate are then computed as follows.

Let N_1 be the number of trials in which S_1 was presented and N_2 the number of trials in which S_2 was presented. Let N_h be the number of trials where S_2 was presented and the subject responded R_2 (hits). Similarly, let N_f be the number of trials where S_1 was presented and the subject responded R_2 (false alarms) Then:

$$h = N_h/N_2, \quad f = N_f/N_1 \tag{2}$$

Now d' is estimated from (1) using hit and false alarm rates as estimated from (2). Interestingly, *even though the actual hit and false alarm rates strongly depend on response criterion k , and d' is computed from hit and false alarm rates, d' itself does not depend on the response criterion k [8] .*

Finally, it must be pointed out that the formula (1) provides a good estimator of detectability only if the response criterion k is stable throughout the session. Otherwise, the estimated d' will be an underestimation of the true d' . So, even though, theoretically, d' is a measure of the percept that is unconfounded by a response bias, in practice this may not be true if the response criterion k is not stable. The best way to assure

that the response criterion is stable is to use experienced subjects who are familiar with the stimuli and the experimental setup.

3.2.1. Stimuli

For each trial the subject is shown an image rendered from one of 18 objects of varying complexity, from buildings and automobiles, to simple pieces of furniture (Figure 2). For about half of the images the objects are distorted, while for the other half they are not. The procedure to generate the images is shown in Figure 3. The original object (*a*) is stretched by different random amounts along each of its three major axes (*b*). This step ensures that although the same object is used for several tests, the object has a unique normal undistorted appearance for each of the tests. If an object is to be distorted, it is rotated about an arbitrary axis (*c*), stretched 40% (*d*), and then rotated back (*e*). Then, regardless of whether it has been distorted or not, the object is rotated about an arbitrary axis to generate a new random viewing angle (*f*).

The images are generated on the Perspecta using our custom application that uses the manufacturer's implementation of OpenGL [8], a commonly used 3D computer graphics API. The application specifies the geometric data of the scene and the scene subvolume that should be mapped to the volume of the 3D display. The OpenGL driver then computes the 3D image by slicing (intersecting) the scene with the plane of the revolving disk, and transfers the slice images to the display. More details on the spatial 3D image generation can be found in the three whitepapers on Actuality's website.

All lights in the laboratory were turned off except for a 75W incandescent spot light that was pointed away from the display. The light source provided sufficient ambient light to allow the subject to move safely around the display. Figure 4 shows photographs of the stimuli as seen on the Perspecta. The display was photographed in the actual lighting conditions used during the experiments. The exposure was set to allow for a full revolution of the display's revolving screen. Figure 5 gives left-right-left 3-panel stereograms showing various stimuli used, built by photographing the display from two locations separated by a 60mm translation. Fusing these images is difficult. The images differ by more than the camera position because of

the dynamic aspect of the display (rotating screen). Figure 6 shows several stimuli, each from three different views. In our case, this seems to be the best way to convey the 3D images on paper, although the actual images on the Perspecta looked substantially better.

3.2.2. Procedure

Five experimental conditions were used, one session per condition (see Table 1). In four conditions, the stimuli so generated were shown on the Perspecta (Figure 1). In the fifth condition, the stimuli were shown on a LCD screen (Figure 7). In the case of Perspecta, the stimuli were viewed binocularly from the viewing distances of 1m, and 2m. The near distance was chosen to ensure that stereo disparity, given the size of the globe and the human interpupillary distance, was effective. The far distance was chosen to minimize the stereo disparity while still providing a view of the display with sufficiently high resolution. For the same reasons, these distances also define the range of typical viewing distances when the display is in use. In the monocular sessions, the subject was at 1m from the display. Monocular viewing is achieved using an eye patch. In these three conditions, the head was supported by a chin-forehead rest.

In the fourth condition, the subject was instructed to walk around Perspecta while making judgments. The subjects were not restricted in anyway, but most subjects chose to revolve around Perspecta on a circle of a 1m radius. In the case of the LCD, the viewing was binocular from a distance of about 1m, which was the simulated distance used in computing perspective images. In this case, however, binocular disparity was absent. Therefore, this condition is called monoscopic, as opposed to stereoscopic.

Each session started with 20 practice trials. Then, 200 experimental trials followed. On each trial, an object was shown and the subject's task was to judge whether the object was symmetric. Exposure duration was 15 sec. The subject had the option to respond at anytime during the 15 sec period, or after that. Each subject received the same set of images, but the order in which the images were displayed was randomized with each subject and with each session. The subjects typically responded while the image was on the screen. A session of 220 trials took between 25 and 45 minutes to complete. The authors received additional practice

by running practice sessions with 200 trials. This was done in order to verify as to whether substantial learning effects are present. No such effects were found.

The subject was provided with a display of the current test number, the currently selected answer, and whether the previous answer was correct or not. The verdict on the previous answer was given to help the subject maintain focus. The feedback was given on a nearby LCD (Figure 1) that had a black background so as not to interfere with the experiment.

4. RESULTS

Estimated d' are shown in Table 2. Recall that a chance performance is represented by $d'=0$ and perfect performance by $d'=\infty$. The first two rows show results of the two naïve subjects. Average d' values are shown in Figure 8. The error bars in Figure 8 represent one standard deviation of the mean.

Performance in the case of walking around Perspecta was substantially higher than that in the remaining four conditions (see Figure 8). Binocular viewing of Perspecta from a distance of 1m led to higher performance than that from a distance of 2m, and this, in turn was higher than monocular viewing of Perspecta from a distance of 1m. These results are not unexpected. Next, binocular performance with Perspecta is similar to that in monoscopic viewing of an LCD. This result is probably related to the fact that the images rendered on Perspecta have rather low contrast. Interestingly, differences in performance across binocular viewing of Perspecta from distances 1 and 2m, and monocular viewing of Perspecta are not very large. Implications of this result will be discussed in the next section.

To evaluate statistical significance of the results, we performed a one-way repeated measures analysis of variance. The main effect of the viewing condition was significant ($p<0.001$). Post hoc tests showed that free viewing of Perspecta led to better performance than any of the other four viewing conditions ($p<0.001$). The only other significant effect was found between binocular viewing of Perspecta from 1m viewing

distance and monocular viewing of Perspecta ($p < 0.05$). All other differences were statistically not significant ($p > 0.05$).

5. DISCUSSION

The main result of our experiment is that free viewing of stimuli shown in Perspecta leads to a very reliable performance and this performance is much higher than performance in monoscopic viewing of an LCD display. The fact that viewing Perspecta from one viewing direction either monocularly or binocularly does not lead to better performance than monoscopic viewing of stimuli on an LCD display should be attributed to the presence of a number of effective monocular cues that are available on LCD displays. Monocular cues are likely to be more effective with an LCD display as compared to Perspecta because of higher contrast, the presence of texture and color and clear edges, as well as self occlusions. The fact that a single perspective image of an object can lead to a reliable 3D percept is not new and is not restricted to symmetry judgments (e.g. [5]). It is known that shapes of structured objects allow the human visual system to use strong priors, which lead to 3D percept from a single 2D image ([10] and [11]). To further characterize the effectiveness of 3D images, it would be interesting to use other types of stimuli (e.g. motion) and other types of judgments (speed or 3D distances).

6. FUTURE WORK

Our Perspecta display is a second generation device. Cost trade-offs in the manufacture of the device have resulted in an image that is slightly wobbly. We estimate that the amplitude of the wobbling is approximately 5mm. A more costly, high-precision mechanical fabrication would lead to a more stable image. Furthermore, the current pixel resolution is expected to increase in future generations of the device. Both factors suggest that the full potential of the device has not yet been realized. Moreover, the bandwidth of the connection between computer and device undoubtedly will improve in the future. Today, moving complex images

cannot be so displayed. When this has changed, new experiments should be conducted to assess the effect of these improvements.

We are pursuing an alternative path to improving the stability of the 3D image by modeling the wobbling pattern. We speculate that the pattern does not change considerably over time and that it can be almost eliminated by better calibration. The model parameters will be tuned manually by displaying a calibration scene and searching for a combination that minimizes the wobbling. We plan to also investigate the possibility of using a camera to provide the feedback needed for calibration and completely automate the procedure.

There are image applications that are traditionally difficult to do on conventional displays. They include point clouds and point/line arrangements in 3D-space. Such images are easy to display and comprehend on the Perspecta, and are difficult to render well on conventional displays. A familiar strategy on conventional displays is to add as depth cue a reduced size and/or a fainter intensity of more distant parts of the arrangement. It would be interesting to compare the effectiveness of these display strategies with a straightforward rendering on the Perspecta.

A fundamental limitation of the 3D display technology employed by the Perspecta is its inability to display view dependent effects such as occlusions and reflections. For example a bright back surface cannot be hidden by a dark front surface, or a specular highlight's position is ambiguous when more than one view is considered. We will investigate how to reduce the artifacts resulting from these fundamental limitations. An a priori knowledge of the set of desired view locations could be employed to first eliminate the surfaces that are not visible from any or most of the desired views.

Work in image-based rendering has shown that stationary specular highlights are preferable to eliminating all the highlights and treating all surfaces as diffuse. It is our experience that users generally do not notice that the specular highlights do not change with the desired view, formal user studies are in order to establish which applications could tolerate this approximation. When the correct highlight is important, tracking the

user appears to be the only solution. For multiple users, each will have his own highlight but they will also see the highlight rendered for the other users, which will probably give the impression of a moving light.

3D displays are a promising technology; however, the technology is still at its infancy. Level-of-detail and occlusion culling algorithms, schemes for parallel and distributed rendering, antialiasing algorithms (another view dependent effect), interfaces (for pointing, selecting, navigation, etc.) are yet to be developed and constitute interesting and potentially very fruitful avenues for future research.

7. ACKNOWLEDGEMENTS

We would like to thank all the members of our computer graphics and visualization laboratory that put up with the dim lighting conditions required to run the numerous experiments. This research has been supported in part by NSF grants DMS-0138098, DCNS-0216131, DHER-0227828, DSC-0325227, DCMS-0443148, EEC-0227828, EIA-0216131, and ACI-0325227. Hoffmann is also supported in part by an IBM Faculty Award.

REFERENCES

- [1] Perspecta 2003, Acutality Systems Inc., URL: <http://www.actuality-systems.com/>
- [2] Perspecta whitepapers. <http://www.actuality-systems.com/index.php/actualty/products/whitepapers>
- [3] DTI 2D/3D 2003, Dimension Technologies Inc. URL: <http://www.dti3d.com/>
- [4] IBM T221 2003, IBM Corporation, URL: <http://www.ibm.com/>
- [5] Chan M.W., Pizlo Z. & Chelberg D. (1999) Binocular shape reconstruction: psychological plausibility of the 8 point algorithm. *Computer Vision & Vision Understanding*, 74, 121-37.
- [6] Green D.M. & Swets J.W. (1966) *Signal Detection Theory and Psychophysics*. NY: Wiley.
- [7] Hochberg J. & Brooks V. (1962) Pictorial recognition as an unlearned ability: a study of one childe's performance. *American Journal of Psychology*, 75, 624-8.
- [8] Macmillan N.A. & Creelman C.D. (1991) *Detection Theory: A User's Guide*. Cambridge University Press.
- [9] OpenGL URL: <http://www.opengl.org>
- [10] Pizlo Z. Li Y. & Francis G. (2005) A new look at binocular stereopsis. *Vision Research*, 45, 2244-2255.
- [11] Pizlo Z. (2001) Perception viewed as an inverse problem. *Vision Research*, 41, 3145-61.
- [12] Pizlo Z. & Stevenson A. (1999) Shape constancy from novel views. *Perception & Psychophysics*, 61, 1299-1307.
- [13] Pizlo Z. (1994) A theory of shape constancy based on perspective invariants. *Vision Research*, 34, 1637-1658.

- [14] Slater A. (1998) Visual organization and perceptual constancies in early infancy. In: Walsh V. & Kulikowski J. (Eds.), *Perceptual Constancy*, Cambridge University Press (pp. 6-30).
- [15] Wickens T.D. (2002) *Elementary Signal Detection Theory*. Oxford University Press.

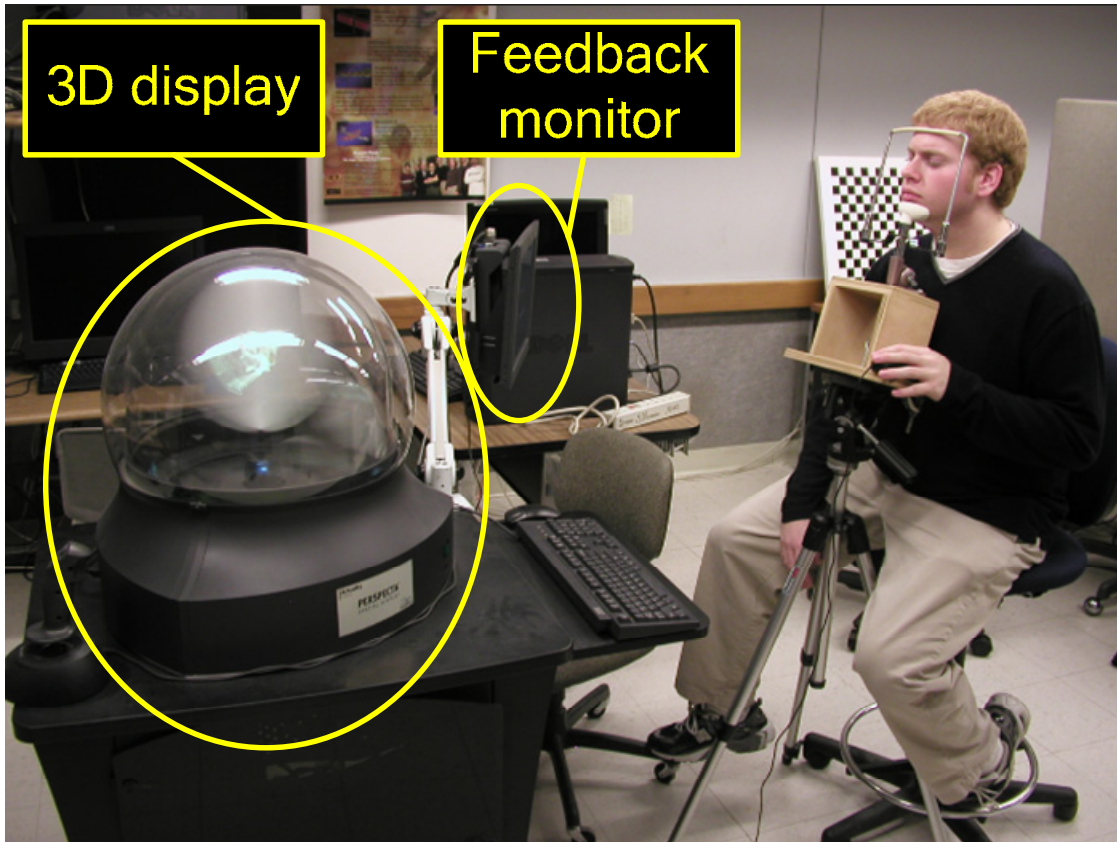


Figure 1 – Experiment setup. The lab lights were dimmed during actual experiment.

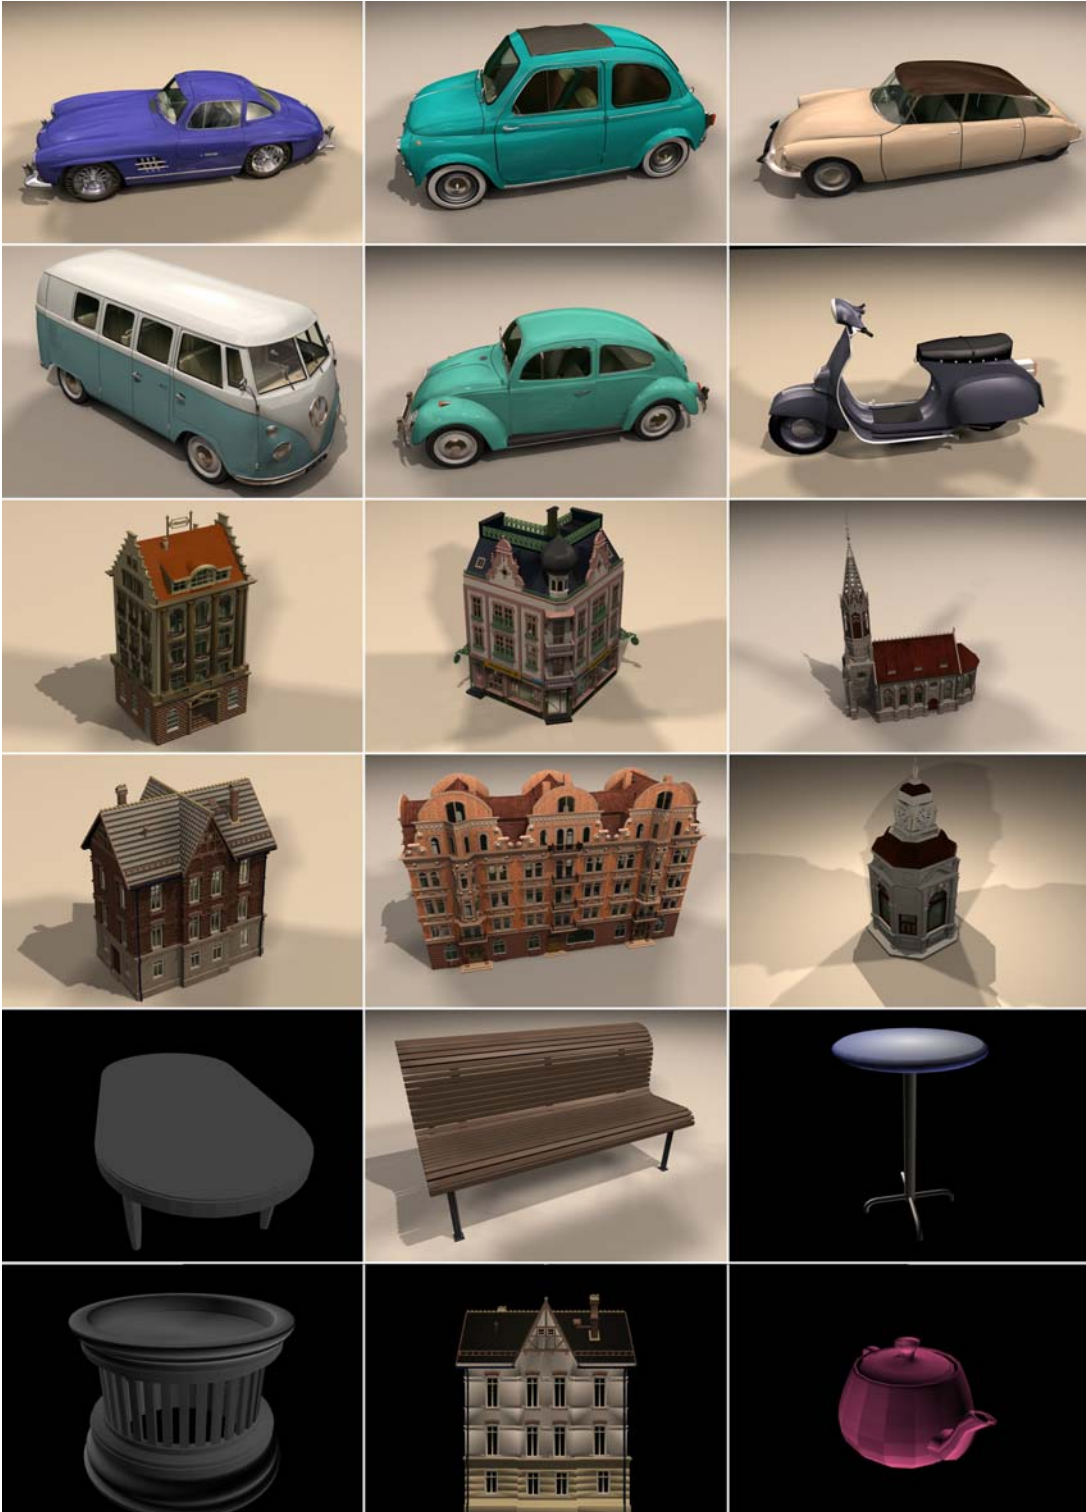


Figure 2. Objects used to generate stimuli

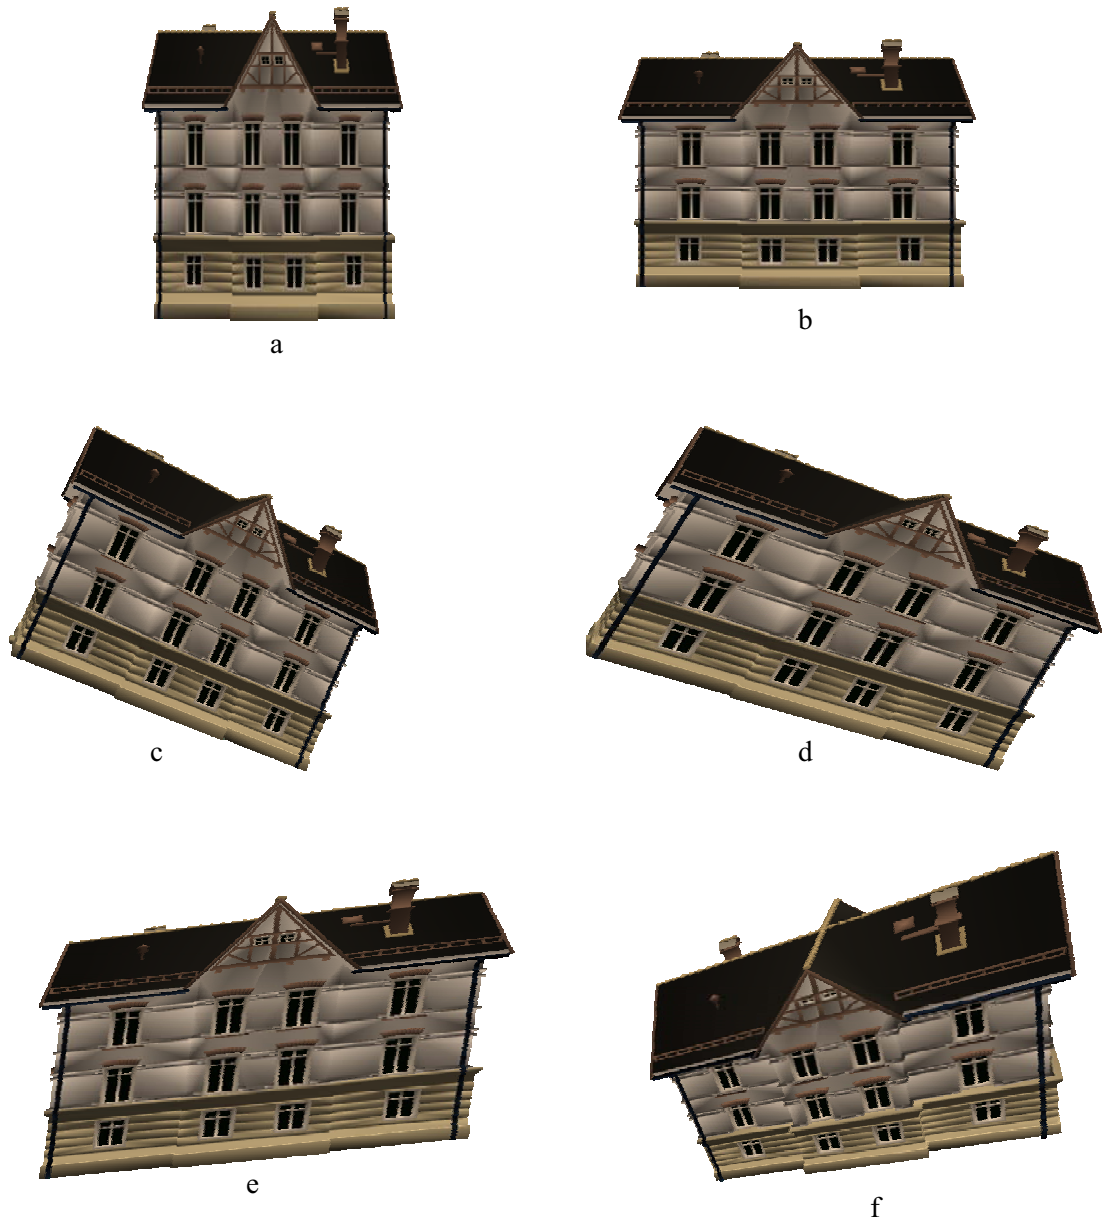


Figure 3. Distorted trial image generation

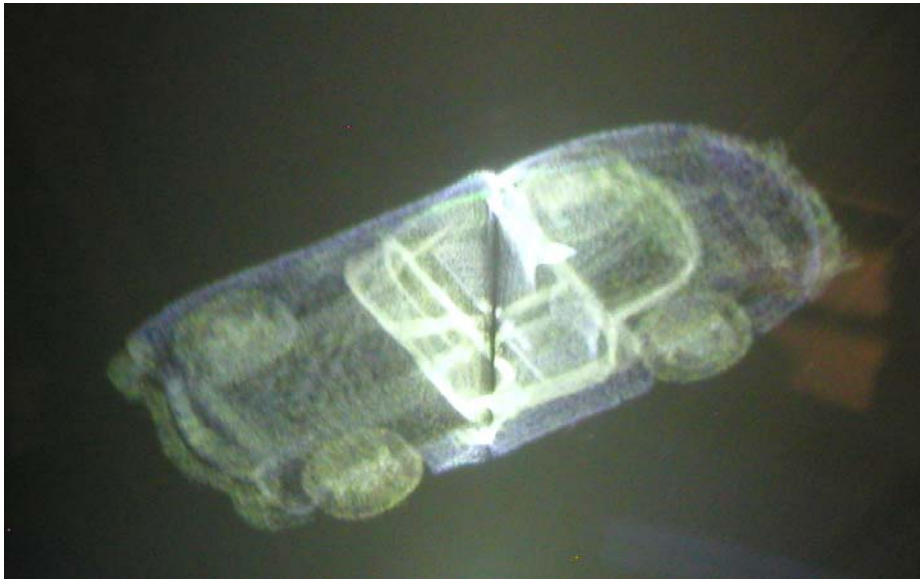
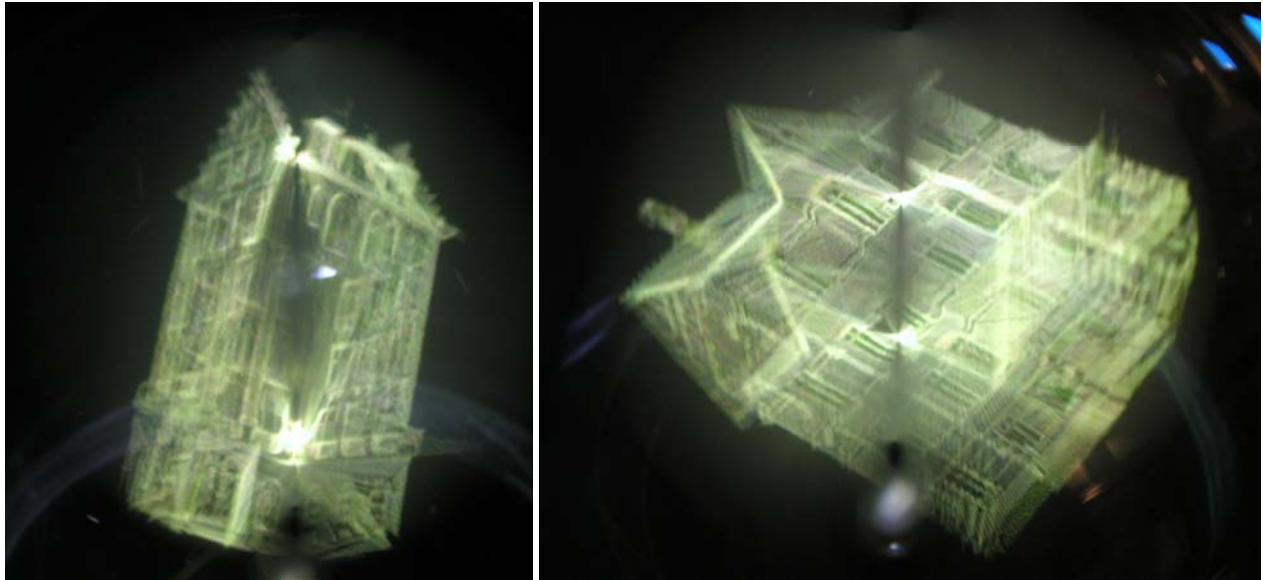
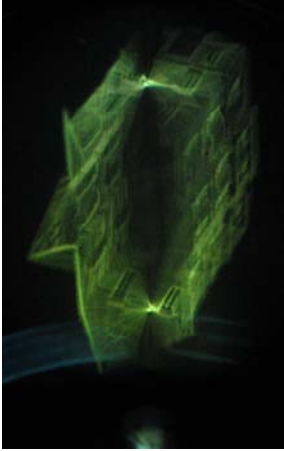
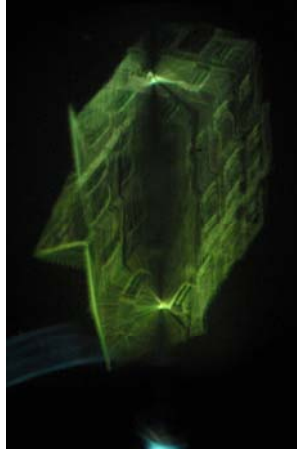


Figure 4. Examples of 3D images generated by the Perspecta display.

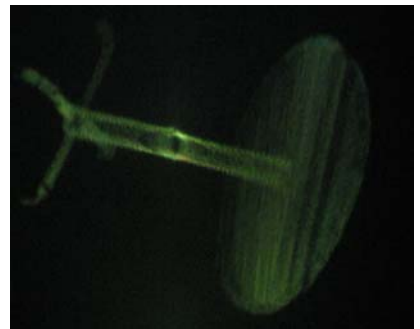
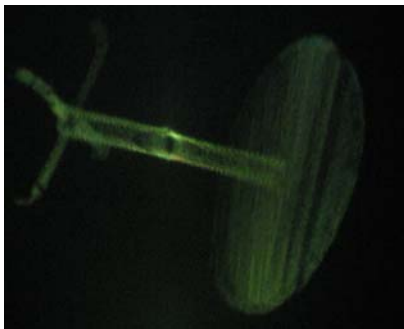
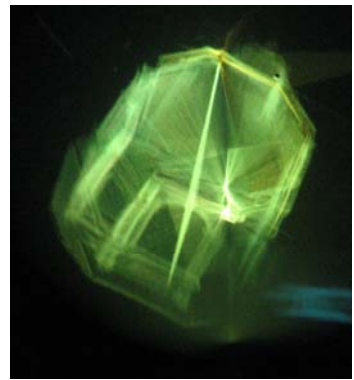
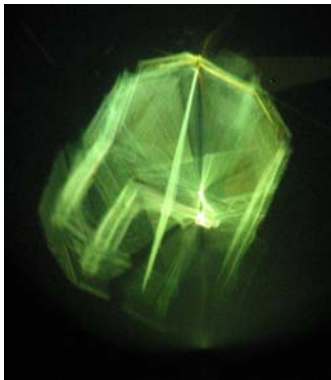
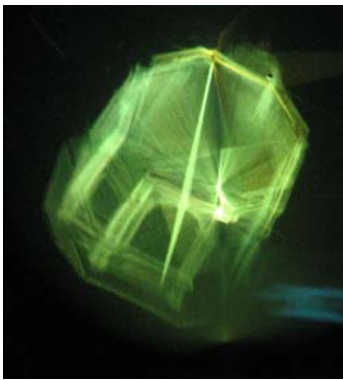
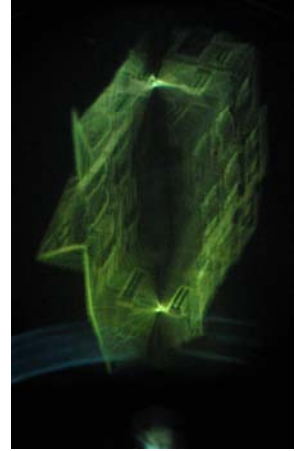
Left

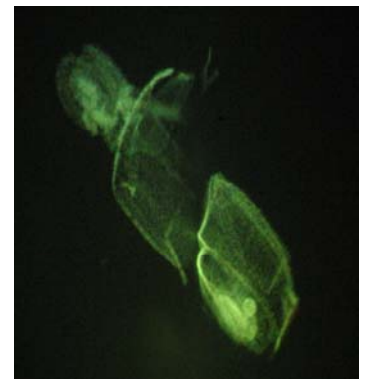
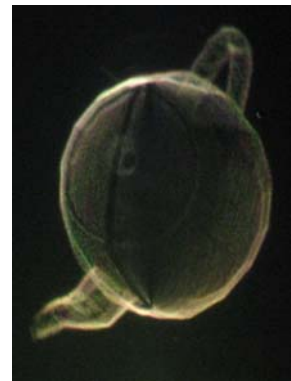
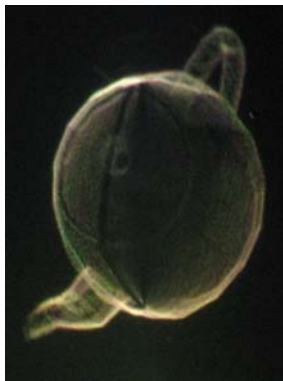
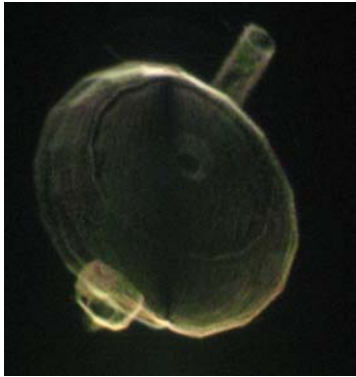


Right



Left





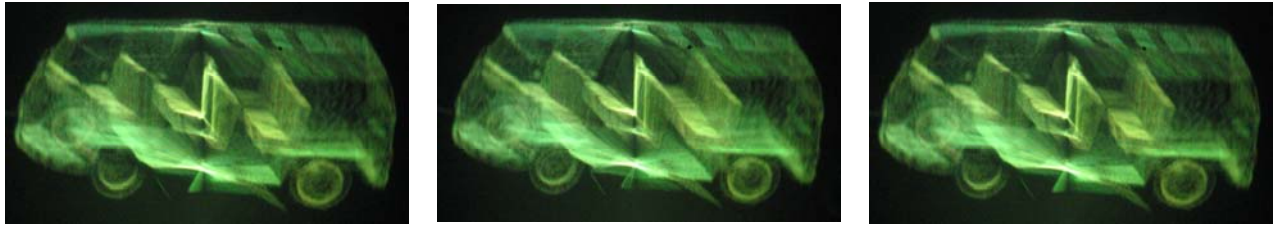
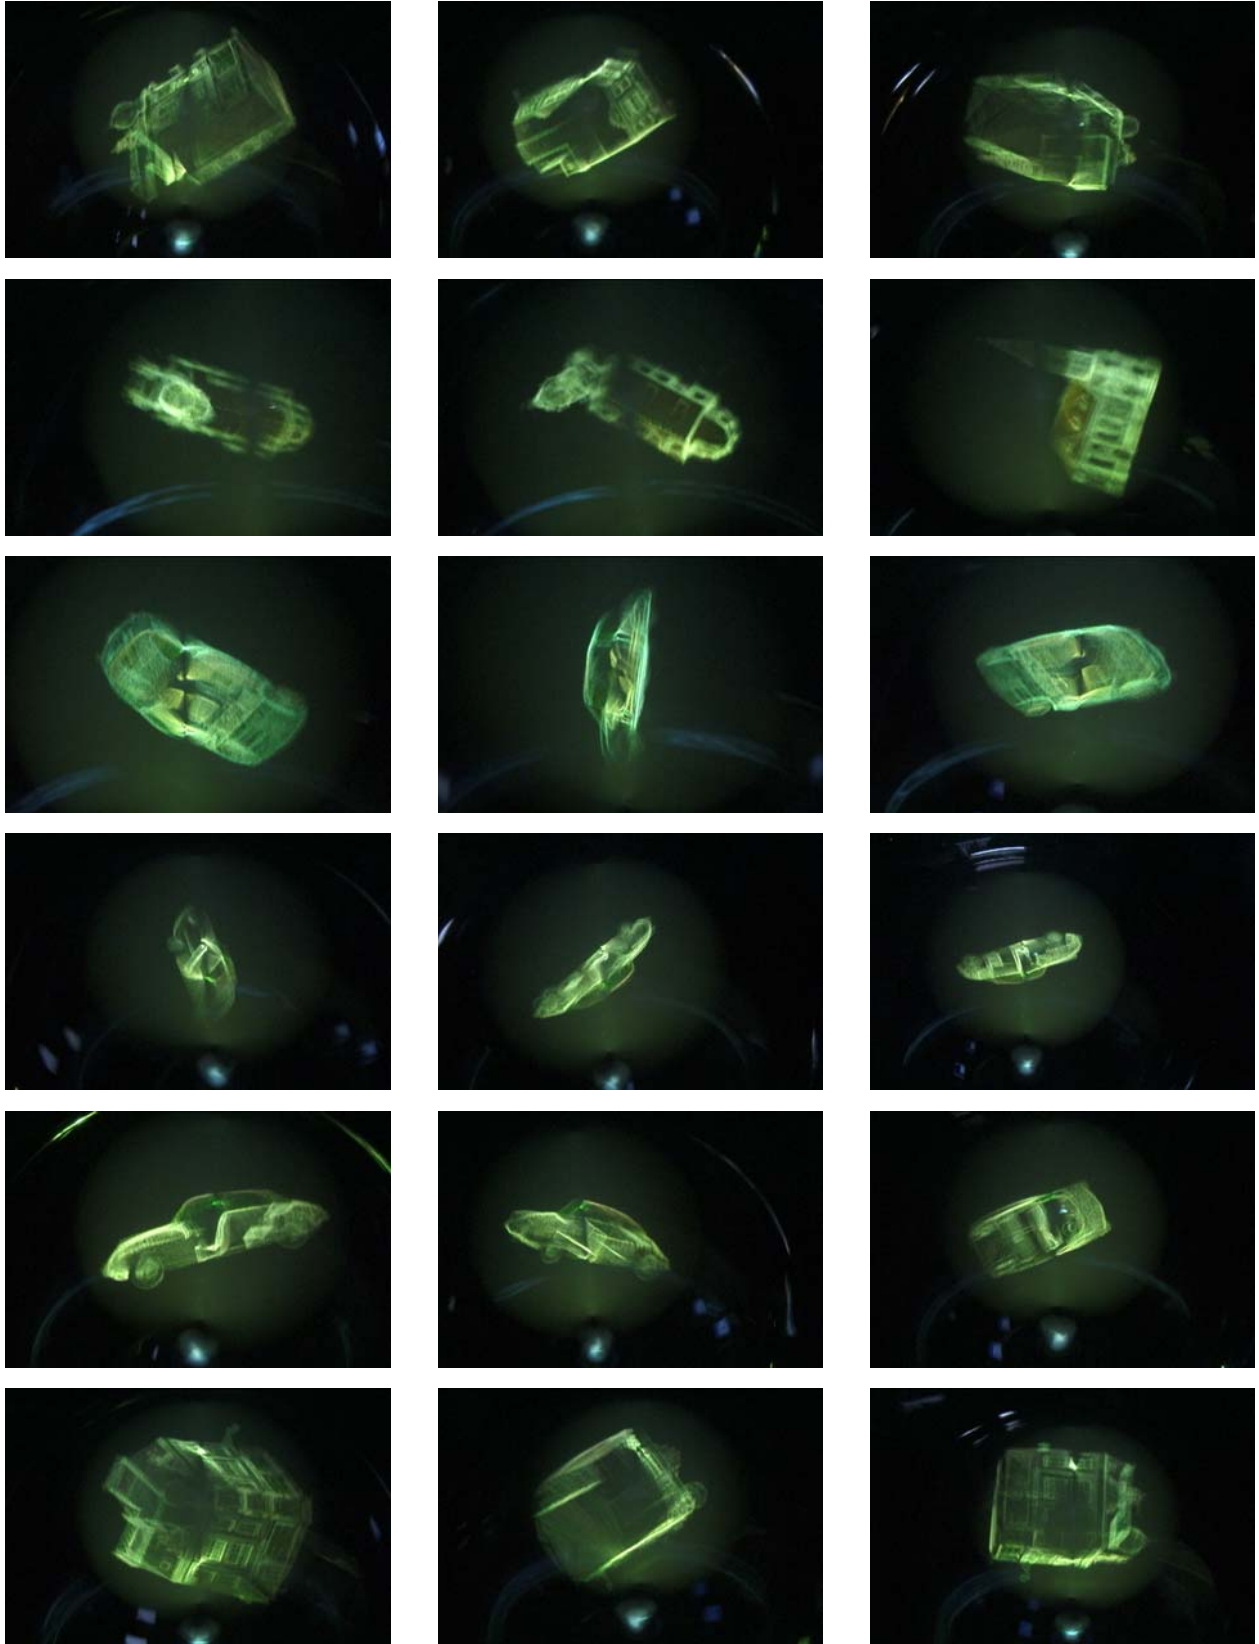


Figure 5. Left-right-left 3-panel stereograms exemplifying the stimuli used.



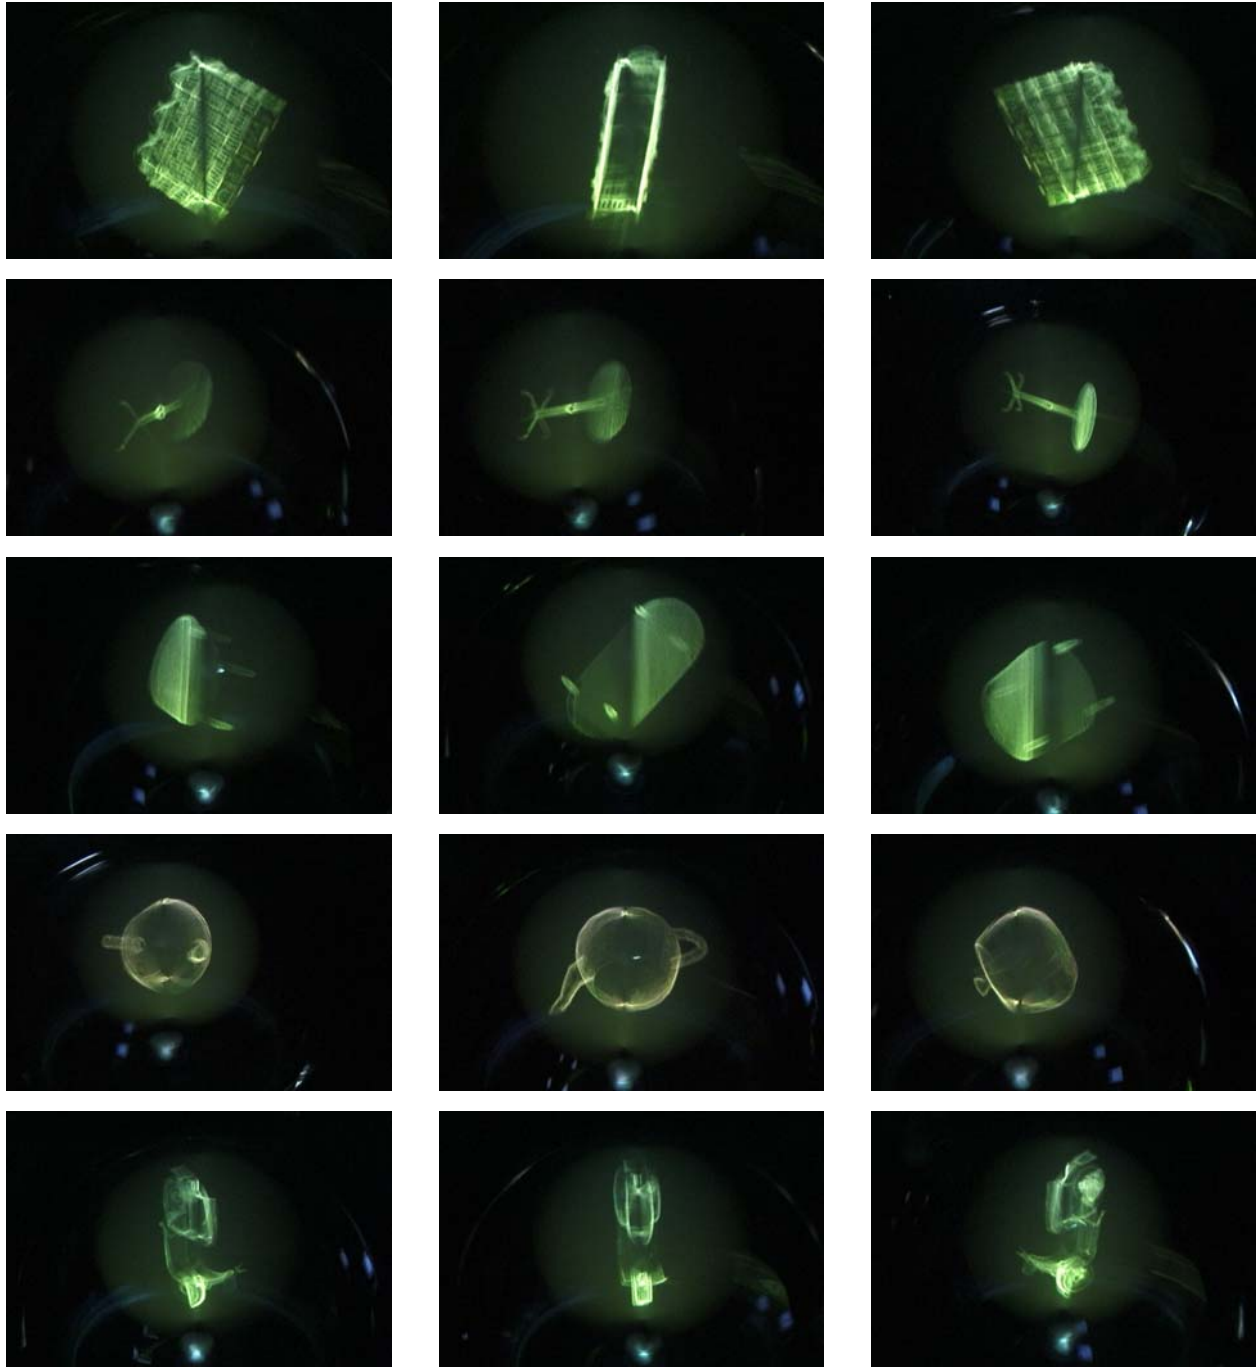


Figure 6. Photographs of various stimuli from 3 different views.



Figure 7. LCD experiment setup

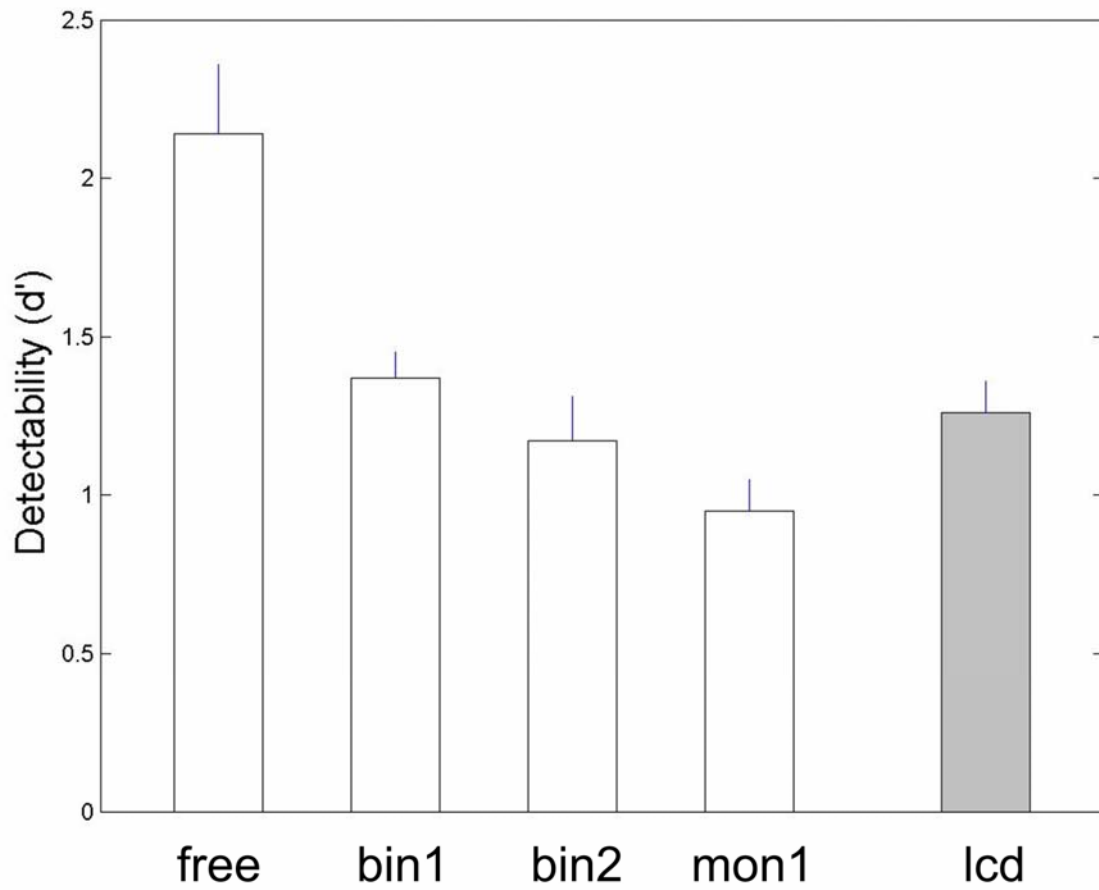


Figure 8. Detectability averaged across subjects for each of the 5 experimental conditions.

		Condition				
		Display	Monocular/Bino cular	Viewing distance	Viewpoint	Label
Experiment #	1	3D	Binocular	1m	fixed	bin1
	2	3D	Monocular	1m	fixed	mon1
	3	3D	Binocular	2m	fixed	bin2
	4	3D	Binocular	Variable, walk around freely	free	free
	5	LCD	Binocular	1m	free	lcd

Table 1 Viewing conditions used for each of the 5 experiments

		Experiment				
		free	bin1	bin2	mon1	lcd
Subject	SP	2.24	1.48	1.41	0.69	1.37
	JD	1.26	1.02	0.59	1.05	1.30
	VP	2.22	1.23	1.23	0.72	1.47
	PR	1.91	1.55	1.32	1.29	1.33
	CH	2.30	1.45	1.48	1.17	1.33
	ZP	2.90	1.51	0.97	0.77	0.75

Table 2 Detectability d' measured for the 6 subjects for each of the 5 experimental conditions used.