4

Detecting sharp objects

As discussed in Chapter 2, detecting sharp objects, such as knives , is a task that is relevant to x-ray baggage inspection. Especially in the hand baggage, such items are not allowed, although they may be transported in hold baggage. Usually, the inspector will recognise knives because he is familiar with most of them, but this experiment is designed merely to check the visibility of these dangerous sharp points and edges in hand-baggage. In this experiment, a sharp edge was defined as an edge sharper than 30˚. Figure 4.1a shows an object with a sharp edge, Figure 4.1b an object without such an edge. We have no x-ray scanner to make x-ray images. To match the x-ray baggage inspection task, non-familiar objects with and without sharp edges were made of transparent polyester potting resin.

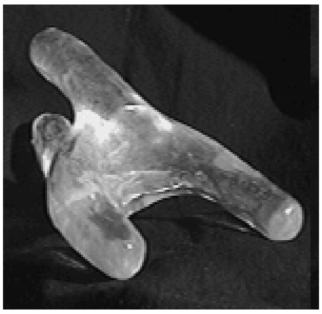

Figure 4.1a. Example of a sharp object. See Figure 2.10 for a stereoscopic depiction of this object.

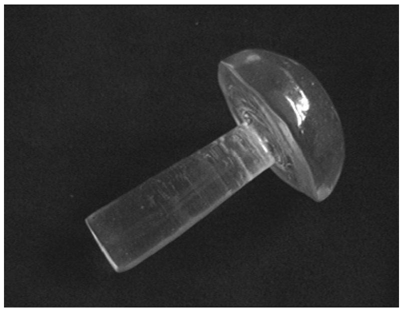

Figure 4.1b. Example of a blunt object.

If the observer can manipulate the object itself to do all the checks he wants, he is expected to be able to find such sharp edges when present. However, if the objects are inspected via a monitor he cannot manipulate the objects, feel the sharpness of the edges or choose any view he likes. Furthermore, the limited resolution and the limited number of grey levels will lower the sharpness cues, and thus his ability to see sharp edges. This expected decrease of observer performance with decreasing image quality corresponds to results from the literature (Ranadivé, 1979; Swartz, Wallace and Tkacz, 1992; Snyder, 1973).

This chapter describes three experiments. The first experiment checks the visibility of sharp edges when the objects are inspected via the DVWS, while the resolution and the number of grey levels in the views are varied between low and high settings. The second experiment tests the visibility of these edges when the objects are inspected naturally. To test the effect of the number of views on performance without the disturbing effects of varying image resolution and averaging over participants, the third experiment again tests their visibility when inspected via the DVWS while only high-resolution views are provided to the observer.

As discussed in Chapter 3, the number of available views has to be limited because there is a maximum x-ray dose to which the baggage can be exposed. Therefore, the camera will be given just one degree of freedom: the left-right movement (Figure 4.2). So movements up-down and forward-backward are allowed, but do not give him another view of the scene. The camera keeps aimed at a point in the scene: the fixation point. It moves on an arc around this point according to the eye positions of the observer. To achieve this, the eye position of the observer is continuously tracked, and the viewing angle φ_{obs} determines φ_{cam} as will be described below.

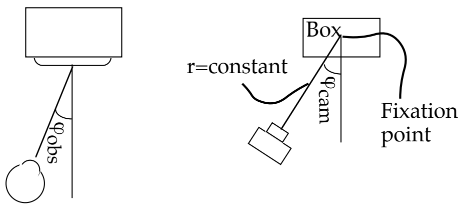

Figure 4.2. Top view of the DVWS. The angle of the camera relative to the scene φ_{cam} is adapted continuously to fit the actual φ_{obs}.

For this left-right movement only N images will be taken, with a constant angle $\Delta\varphi$ between two images. Thus, the observer will see the same view when he moves within a certain sector. This gives the situation of Figure 4.3. The lines indicate the direction from which the images were taken [illegible]

The angle between the lef[illegible] $[N - 1] \cdot \Delta\varphi$) also seems important. In ord[illegible]er of images N, a large angle betw[illegible]gle $\Delta\varphi$ gets too large, the jerkiness of the [illegible] impression of selecting one from a numb[illegible]iewing a rigid 3D scene.

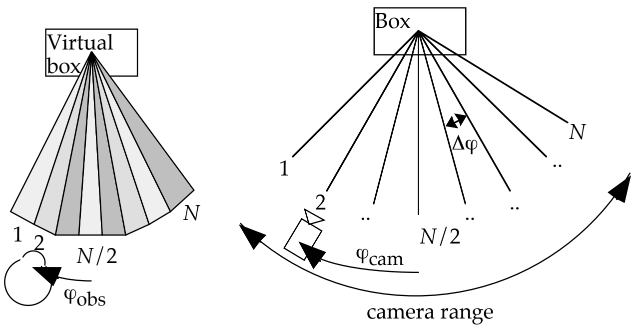

Figure 4.3. Top view of the scene. There are N available views and an angle $\Delta\varphi$ between the views.

I expect a larger camera range to be effective up to 180˚. This cannot be viewed if φ_{obs} and φ_{cam} are kept equal: to see the extreme views the observer then has to look at the display from the side. Therefore, I scaled the camera movement to a maximum comfortable head movement of ±22.5˚ (see McVey, 1970), giving a scale factor $\varphi_{cam} / \varphi_{obs} = N.\Delta\varphi / 45$. Furthermore, I expect that the observer performance will increase with the number of available views, as in earlier studies (Edelman and Bülthoff, 1992; Field, Michell, Wallis and Wilson 1995; Braunstein, Hoffman, Shapiro Andersen and Bennett, 1987).

Experiment 1- inspection via a monitor

The objective of this experiment is to test the effect of a reduced image quality and number of available viewpoints on the performance of observers in detecting sharp edges.

Method

Stimuli

The stimuli were images of mock-up baggage consisting of a transparent box with two different transparent objects as shown in Figure 4.1 in it, each possibly having sharp edges as described above. The sharpness of edges was tested during manufacturing with a wedge-shaped aperture. The box of 25 x 10.6 x 20.6 cm was made of tinted perspex. For each stimulus, two new objects were placed in new positions in the tranparent box. To get recordings of 34 mock-ups, 68 different objects were prepared. Figure 4.4 shows a sample recording of a box containing two objects.

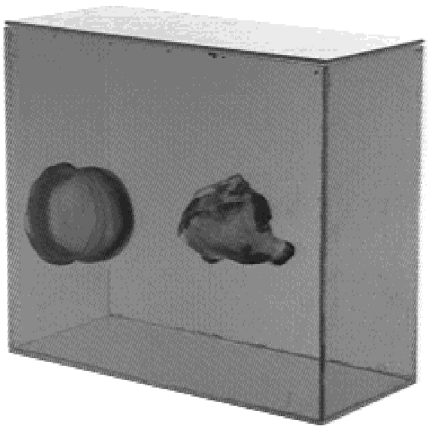

Figure 4.4. A view of a box containing two objects.

80 images of each box of objects were recorded from different viewpoints with a video camera (Sony CCDTR805E). The camera images were digitized (Archimedes real-time video digitizer from Watford Electronics) to digital images of 512x256 pixels with 16 grey levels, and stored uncompressed on a hard disc. Before recording the stimulus, the camera was tuned to use the total range of grey levels. The size of the image of the box on the screen was 12.5 x 10.9 cm in front view. The distance between camera and fixation point was 90 cm. This value is appropriate as the size of the image of the box is half of the real

box size, and the average viewing distance between observer and screen was expected to be 45 cm.

For a reduction of the number of grey levels, the original 16 grey levels were divided into 4 or 8 groups, and for each group the brightest value was taken. An informal evaluation by the experimenter indicated that the image contrast was not altered very much by this reduction. For a resolution reduction, the image pixels were grouped in 2 x 2 or 4 x 4 pixels whose intensity was averaged.

Apparatus

An Archimedes A5000 computer was used to display the images according to the eye position of the participants, and to store their responses. Preceding each trial, the appropriate viewing angles of the box were read from hard disc, reduced in number of grey levels and resolution if necessary, and stored in working memory. This caused a pause of about 10 s between trials. During the trial, the appropriate images were shown from working memory on the display (Puretek PT143D PLUS: non-glare monitor, 0.29 mm dot pitch). The screen refresh rate was 88 Hz (max. delay 11 ms), and the average light output was 150 Lux.

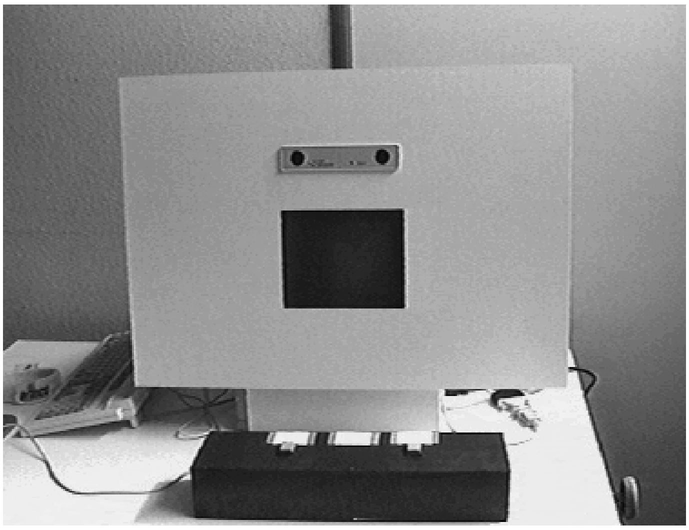

Figure 4.5. Experimental setup. The display is placed behind a reduction screen. Above the screen is the infrared eye position tracker. Below the screen is the button box with two buttons, by which the participants could make their judgements.

Figure 4.5 shows the experimental setup. To enhance the depth in the displayed scene (Gibson, 1971), a white reduction screen was placed in front of the display, making 19.3 x 16.1 cm of the display visible. The eye position was measured by a Dynasight infrared sensor (Origin Instruments,1993). This sensor was placed on top of the screen, and reported at 37 Hz (delay 27 ms) the position of a small reflector between the two eyes on a headband to the computer. This way, the eye position could be estimated with an accuracy of about 3 cm. Directly after receiving a new eye position, the corresponding image was shown on the screen (delay 16 ms). The total lag was about 11 + 27 + 16 = 54 ms. Below the screen was a button box with two buttons, by which the participants made their decision. The left button was labelled with a picture of an angle of 20˚; the right button was labelled with an angle of 40˚; between the two buttons was a label showing 30˚ angle, indicating the boundary between sharp and non-sharp.

Procedure

Participants were told that they would see a number of boxes, each with two objects in it, on the screen. Their task would be to check whether there was an object with a sharp edge in the box. It was explained that a sharp edge was a knife-like edge, sharper than 20˚: the angle indicated above the left button on the button box. They were told that they could look at most 10 seconds, and that a beep would warn for the time limit when 8 seconds had passed. If they had not made their choice after 10 seconds, the screen would turn dark, but they always had to make a choice. They made their choice by pressing a button on a button box: the left button if they detected a sharp edge, the right button if they did not. After making their choice, they had to wait for the next series of images to be loaded into main memory. After this was done, participants were warned with a beep that the next trial would start in 1 second. They were asked to respond as accurately as possible, and they were told that the participant with the most correct answers was going to be rewarded with a cake. It was explained that they could look around the box by moving their heads, and that their approximate eye positions were being tracked with an infrared tracker. The participants were asked to try the range and speed of the position tracker, in order to get used to the tracker. The sensor provided feedback by a control light which was green if the reflector was in track, and red if it was not. In this part of the training no images were shown on the display.

Prior to the experiment the participant was trained with seven trials. For the first training trial, they were allowed to look 30 seconds before the screen would go dark. Directly after the participant had made his choice, the screen showed whether he made the right choice, his response time and the range covered by his eye positions. During the actual experiment, the participant was shown 27 boxes. Overall, each experiment took about 25 minutes.

Variables, Design, Participants

The independent variables are the image resolution *R*, the number of grey levels in the image *G*, the number of available views *N* and the angle between two adjacent views $\Delta\varphi$. Table 4.1 gives the levels of the independent variables. To simplify the expressions for the angle between two views $\Delta\varphi$, I define the smallest angle between two images $\theta = 22.5˚/32$. The dependent variables were the correctness of the response *C* (right if they judged correctly about the sharpness of an object, or wrong) and the response time *T*.

Table 4.1. Independent variables and their values for experiment 1. $\theta = 22.5˚/32$.

Name of variable	Description	Possible values
R	Image resolution	256 x 128, 512 x 256
G	number of grey levels	4 , 8 , 16
N	number of available views	1 , 2 , 4 , 8 , 16 , 32
$\Delta\varphi$	angle between two adjacent views	θ, 2θ, 4θ, 8θ

The 3 (G) x 6 (N) x 4 ($\Delta\varphi$) x 3 (R) = 216 conditions were randomized over 8 participants, giving 27 responses per participant. This randomization was done 3 (repetition) times, for a total of 24 participants. All participants saw the 27 boxes in the same order. The participants were 24 students, mainly from the faculty of Industrial Design Engineering (10 women, 14 men) with normal or corrected-to-normal vision. They did not know the purpose of the experiment. Each participant received NLG 10 (USD 5) for taking part, and the best-performing participant received a cake.

Hypotheses

It is expected that an increase of the resolution R, number of grey levels G or the number of available views N, will increase the percentage of correct answers and decrease the response time. However, for each variable there will be a saturation point, where a higher value for that variable will not improve performance any more. For the angle between the views $\Delta\varphi$ the effect is less clear: a larger value for $\Delta\varphi$ will increase the jerkiness in the movement, but on the other hand it allows a larger range of available views with the same number of views. Our hypothesis here is that for this task, an appropriate viewpoint is more important than a smooth movement between adjacent viewpoints, and therefore that observer performance is expected to increase with the angle between the views $\Delta\varphi$. An alpha level of 0.05 was used for all statistical tests.

Results

An analysis of variance was done to find effects of the independent variables on the correctness of the response C. It shows that the angle between the views $\Delta\varphi$ is close to significance: $F(3,432)=2.61$, $p=0.051$. Figure 4.6 shows the effect: the percentage of correct answers is lower when the angle between adjacent images is 8θ than when it is smaller. The other main effects and the interactions were not significant.

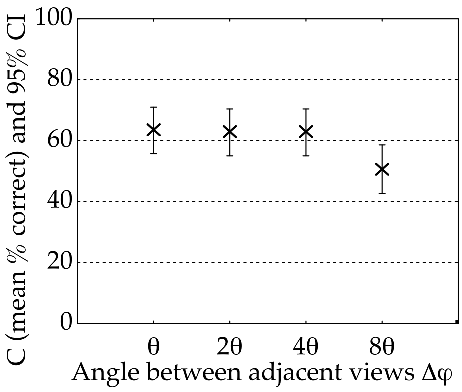

Figure 4.6. The effect of the angle between the views Δφ on the mean percentage correct answers.

An analysis of variance was done to find the effects of the independent variables on the response time T. Here, the number of available views N proved to be significant: $F(5,432)=4.51$, $p=0.001$.

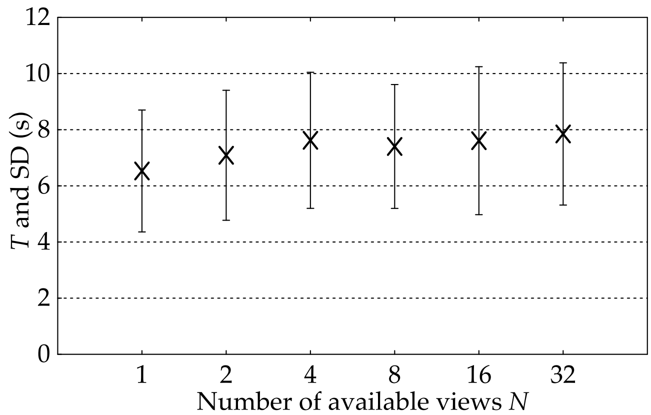

Figure 4.7. The effect of the number of available views on the mean response time.

Figure 4.7 shows the effect: response time increases with the number of views. The effect is quite small: from 1 to 32 views the response time increases from 6.5 to 8 seconds.

The two-way interaction between the number of available views and the angle between two adjacent views also proved significant: $F(15,432)=1.87$, $p=0.024$. Figure 4.8 shows the average response time for these conditions. This effect seems to be caused by the high response time in the condition with $\Delta\varphi=4\theta$ and $N=4$, and this value seems accidental.

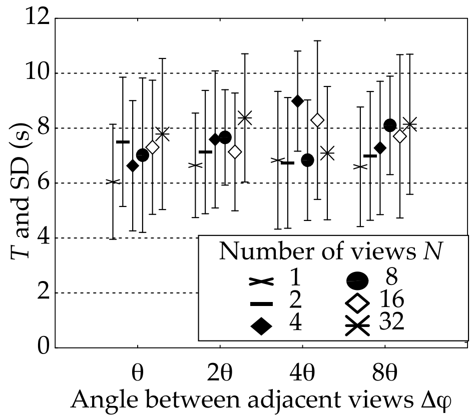

Figure 4.8. The effect of the number of available views and the angle between the views on mean response time.

A detailed analysis of the stimuli and responses indicated that there seemed to be some misinterpretation of the sharpness of edges by the participants. The objects with the highest scores for sharpness, according to the participants, are shown in Figure 4.9 and Figure 4.10. However, these objects have rounded edges, and were meant to be blunt. Instead, participants seem to judge objects as sharp if there is a thin plane at a side of the object. The last column of Table 4.3 shows for each box the percentage of participants judging one of the objects as sharp.

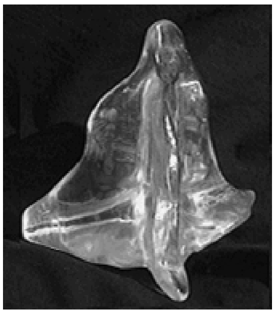

Figure 4.9. The second object from box 2. This object is judged sharp by 83% of the participants of the first experiment.

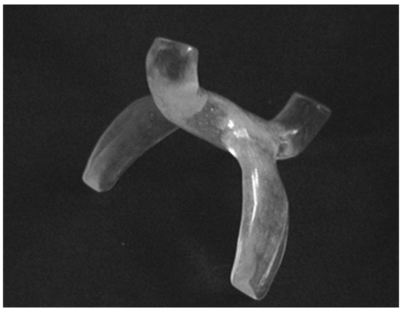

Figure 4.10. The first object from box 24. This object is judged sharp by 71% of the participants of the first experiment.

Concluding, the results are not quite as expected. First, most variables had no significant effect on the response time or percentage of correct answers. Furthermore, in contrast to our expectations, observer performance decreased with the angle between the views, and the response time increased with the number of available views. Probably, either the number of views available or the resolution was too low to perform this task. Alternatively, the task of detecting sharp edges may simply be too difficult for the participants, causing only higher response times but no performance increase when more different views are provided. Finally, the instruction may have been insufficient. I will have to pay more attention to how the participants interpret sharpness.

Experiment 2- natural inspection

Experiment 1 failed to demonstrate any advantage of the Delft Virtual Window System for finding objects with sharp edges. The resolution and number of views may have been too restricted, but alternatively the task might be too difficult, even when the participants are allowed to handle each object to explore it fully. The last hypothesis is checked in experiment 2, by giving participants the real objects and the same task.

Method

Stimuli, Procedure, Apparatus

The stimuli are real transparent objects, those that were placed in pairs in a box for the recording of the stimuli of Experiment 1. Participants was asked to detect sharp edges on the objects they would be given. It was explained that sharp edges are knife-like edges, sharper than 20˚, and that such an edge need not necessarily be straight or on the outside of the object. They were allowed to inspect each object for about 10 seconds. During inspection, they were allowed to take the objects in their hands, but they had to keep them above a cushion, as the objects break easily when dropped. Participants sat on a chair, in front of a cushion lying on a table. The room was illuminated with fluorescent light. After the inspection, they had to write down their judgement (sharp or blunt) on a form. The same three angles as the labels of the buttons of experiment 1 were printed on the form, and it was explained to participants that edges sharper than the middle angle (30˚) were to be classified as sharp, and that angles larger than 30˚ were blunt.

Prior to the experiment, the participants inspected 14 objects (the objects that also were used for training in Experiment 1). After judging an object, they were told whether their judgement was correct and why this was so.

During the experiment, the participants judged 54 objects. Now, however, they were not told about the correctness of their choice.

Variables, Participants, Design

The independent variable was the object (54 levels). The objects were in the same order as in experiment 1, but one after another instead two at once. The dependent variable was the judgement of the participant (sharp or blunt). The participants were 3 students from the faculty of Industrial Design Engineering (2 women, 1 man) with normal or corrected-to-normal vision. They did not know the purpose of the experiment. Each participant received NLG 10 (USD 5) for taking part.

Results, Discussion

Table 4.2 shows the results. Participant RH said that he felt the sharpness of the edges with his fingers, but IG and ES said they preferred looking to feeling.

There is agreement between the participants about most objects. However, the first object in box 13, 24 and 25 and the second object in box 2 seem to give problems.

There seem to be three problems that may explain the deviating answers. The first problem is misunderstanding of the instructions. For example, the first object in box 21 contained a sharp edge of a non-transparent material, but some informal talking with IG after the experiment showed that he considered this part not to belong to the object. One object in box 11 has a sharp cut in stead of a sharp edge, and ES judged the object as sharp. The second problem is caused by the rounding of the edges. For example the first object in box 24 was a bird-like object where the thin wings were rounded, but ES and RH judged them as sharp. The last problem is difficulty with estimating angles of objects. For example box 25 contains a folded starfish, and the sharpness of its edges is very difficult to estimate.

Concluding, most responses are correct, so detecting sharp edges is not too difficult if the objects themselves can be handled. A few objects cause confusion, which may be solved with more precise instructions. The poor responses of Experiment 1 must be related to inspection via the DVWS.

Table 4.2. Judgements sharp ('*') or blunt ('-') of the participants, and the physical measurements. The stimuli were presented in the same order as in Experiment 1.

	First object in box				Second object in box			
	Participants				Participants			
Box	ES	RH	IG	measured	ES	RH	IG	measured
1	-	-	*	-	*	*	*	*
2	-	-	-	-	*	*	-	-
3	*	*	*	*	-	-	-	-
4	-	-	-	-	-	-	-	-
5	-	-	-	-	-	-	-	-
6	*	-	-	-	-	-	-	-
7	-	-	-	-	*	*	*	*
8	-	-	-	-	-	-	-	-
9	-	-	-	-	-	-	-	-
10	-	-	-	-	-	-	-	-
11	-	-	-	-	-	-	-	-
12	-	-	-	-	*	-	-	-
13	*	*	-	-	-	-	-	-
14	-	-	-	-	*	*	*	*
15	-	-	-	-	-	-	-	-
16	-	-	-	-	-	-	-	-
17	-	-	-	-	-	-	-	-
18	*	*	*	*	-	-	-	-
19	-	-	-	-	-	-	-	-
20	-	-	-	-	-	-	*	-
21	*	*	-	*	-	-	-	-
22	-	-	-	-	-	-	-	-
23	-	-	-	-	-	-	-	-
24	*	*	-	-	-	-	-	-
25	*	-	*	-	-	-	-	-
26	-	-	-	-	-	-	-	-
27	-	-	-	-	*	*	*	*

Experiment 3- only high resolution

Experiment 2 revealed minor problems with the task itself and with misinterpretation of the instructions. However, these results could not explain the poor responses found in experiment 1. To find out whether the changes of image resolution and number of viewpoints for each subsequent trial caused the problem with inspection via the Delft Virtual Window System, and to avoid rounding effects due to low resolution, two participants judged the boxes of experiment 1 with 32 available views with a high image quality. To allow later comparison with experiment 1 and 2, the instructions were kept the same.

Method

Variables, Design, Participants

The stimuli, apparatus and procedure were the same as experiment 1. The independent variable was the box content (27 boxes). Again, the dependent variable was the judgement of the participant ('sharp' or 'blunt'). Each box was inspected with the Delft Virtual Window System with an image resolution of 512 x 256, 16 grey levels, 32 available views and an angle between two adjacent views of 8θ. Participants were 2 students from the faculty of Industrial Design Engineering (1 woman, 1 man) with normal or corrected-to-normal vision. They did not know the purpose of the experiment. Each participant received NLG 10 (USD 5) for taking part.

Results

Table 4.3 shows the results. The participants perform worse than in experiment 2. The wrong judgements for box 2, 13, 24 and 25 correspond to similar judgements in experiment 2.

For box 7 and 27 it seems that the sharp edge can be detected only from a viewpoint that is not present in the 32 available views. If I compare the percentage of sharp judgements from the first experiment with those of this experiment, and use a difference of 40% as level of significance, boxes 1, 7,14, 18, 20 and 26 seem to be judged more similar to the measured values in the present experiment, while boxes 3, 5, and 8 are judged less similar to the measurements. It is difficult to draw conclusions from these results, as only two measurements are available for each box, but it seems that indeed the switching of the number of available views and the image quality between the trials has lowered the performance of the participants in the first experiment.

Table 4.3. Judgements of the participants. The boxes were presented in the same order as in Experiment 1. Judgements are sharp ('*') or blunt ('-'). Shaded areas indicate boxes that contain difficult objects according to the second experiment.

Box	participant VS	participant XZ	measured	%judgements sharp in experiment 1
1	*	*	*	58
2	-	*	-	83
3	-	-	*	42
4	-	-	-	8
5	-	*	-	8
6	-	*	-	17
7	-	*	*	12
8	*	*	-	25
9	-	*	-	62
10	-	-	-	12
11	-	-	-	25
12	-	-	-	21
13	-	*	-	8
14	*	*	*	29
15	*	-	-	37
16	-	-	-	29
17	-	-	-	62
18	*	*	*	37
19	-	-	-	29
20	-	-	-	37
21	*	-	*	75
22	-	-	-	29
23	-	-	-	4
24	-	*	-	71
25	-	*	-	42
26	-	-	-	46
27	-	-	*	25

Discussion and conclusions

Many results are unexpected, and explanations are difficult to find due to the low number of tested participants in the second and third experiments. The first experiment showed that for finding sharp edges, an angle of 8θ between adjacent views gives *lower* performance of the participants as compared to smaller angles between adjacent views. Apparently, the jerkiness of the movement disturbs the participants. Furthermore, the average response time *increases* with the number of viewpoints. One explanation may be that observers need extra time to interpret the extra views. Finally, Experiment 2 showed that the task can be done when participants can handle the real objects, but that there may be misunderstanding of the instructions, misinterpretation of rounded edges and difficulty with estimating angles of objects. Results from Experiment 3 suggest that limiting the available viewpoints gives additional difficulties when a particular view is required for estimating the sharpness of an edge, and that the manipulation of the parameters may have lowered observer performance in Experiment 1.

The experiments in the following chapters will again try to show the possible trade-off between the number of available viewpoints and static image quality. In order to avoid misinterpretation, another task will be used that can be explained more clearly and allows less misinterpretation by the participants.