Benefits and Parameters of Shadow in Augmented Reality-Environments

Andreas Jakl

BAKKALAUREATSARBEIT

Nr. 238-003-045-1

eingereicht am
Fachhochschul-Bakkalaureatsstudiengang

Multimedia Technology and -Design

in Hagenberg

im February 2004

Diese Arbeit entstand im Rahmen des Gegenstands

Computergrafik

im

Wintersemester 2003

Betreuer:

Dr. Michael Haller

Erklärung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Arbeit selbstständig und ohne fremde Hilfe verfasst, andere als die angegebenen Quellen und Hilfsmittel nicht benutzt und die aus anderen Quellen entnommenen Stellen als solche gekennzeichnet habe.

Hagenberg, am 19th January 2004

Andreas Jakl

Contents

Abstract

Integrating shadows in Augmented Reality-Environments is very important as they significantly increase the realism and through that the areas of use of the whole technology. At the moment several algorithms exist, which differ greatly both in the quality of the results they generate, but also in the amount of work that is required to get them to run. This paper introduces *Rapid Shadow Generation*, which generates highly realistic soft shadows, but needs a difficult process to recreate rooms as a virtual environment. *shadowAReality* is built for real time applications and therefore delivers only hard shadows.

As live-Images are an integral part of AR, several tests were conduced using the second algorithm. The goal was to better understand some of the most important parameters of shadow. Different scenarios were developed and subjects judged them using a questionnaire. Through that it was possible to reveal several important aspects and prove the importance of shadows.

They are of crucial necessity for a correct interpretation of 3D space. Nonetheless, it is viable to plan the scenes carefully, as they can be confusing as well. In addition the questions of how precisely the shadow has to be projected and how detailed it has to be are answered. Finally it is pointed out, that this kind of shadow does not look overly realistic, but in most cases it is not an option to deactivate it. Depending on the targeted environment it should be considered if soft shadows would not be worth the additional effort.

Overall it was possible to create a solid foundation through the results gathered using the experiments, which is important for the further development of this technology. Without theoretical and practical background knowledge, using and improving shadows in AR would not be possible.

Kurzfassung

Die Integration von Schatten in Augmented Reality-Umgebungen ist ein sehr aktuelles Thema, da diese den Realismus und dadurch die Anwendungsgebiete der generierten Bilder deutlich steigern. Zurzeit existieren verschiedene Algorithmen, die einen unterschiedlichen Aufwand bedeuten und dadurch auch jeweils eigene Anwendungsgebiete haben. Vorgestellt wird kurz ein Algorithmus namens *Rapid Shadow Generation*, der bei hohem Vorbereitungsaufwand weiche und realistische Schatten berechnet. Eher für Echtzeitumgebungen gedacht ist *shadowAReality*, das für Live-Daten optimiert ist und dafür nur harte Schatten liefert.

Da Live-Bilder ein wichtiger Bestandteil von AR sind, wurden mit dem zweiten Algorithmus einige Tests durchgeführt, um die wichtigen Faktoren und Parameter von Schatten besser verstehen zu können. Verschiedene Testszenarien wurden erarbeitet, die mittels eines Fragebogens bewertet wurden. Dadurch konnten einige wichtige Faktoren herausgearbeitet werden und die Wichtigkeit der Integration von Schatten wurde deutlich.

Für die räumliche Einschätzung sind Schatten unumgänglich. Gezeigt wurde allerdings, dass die Szenen mit Vorsicht zu Planen sind, da Schatten auch verwirrend sein können. Eine Antwort auf die Themen, wie genau diese projiziert werden müssen und wie detailliert sie ausfallen müssen, wurde ebenfalls gefunden. Schließlich und endlich wird noch hervorgehoben, dass diese Art von Schatten zwar nicht realistisch wirkt, jedoch ist er trotzdem unumgänglich. Je nach Anwendung sollte dennoch überlegt werden, ob nicht mehr Rechenzeit für realistischere, weichere Schatten aufgewendet werden sollte.

Insgesamt konnten durch die interpretierten Ergebnisse fundierte Grundlagen für die weitere Entwicklung der Technologie geschaffen werden, da die Anwendung und Verbesserung von Technologien ohne theoretisch-praktisches Hintergrundswissen nicht funktionieren würde.

Chapter 1

Introduction

This chapter will give you an overview on how the idea for this term paper evolved and sum up the contents of the individual chapters.

1.1 Motivation

Augmented Reality is a fascinating topic. While still being in development at many universities and companies around the world, it is already obvious that in the future this technology will be the foundation of many great and unique possibilities. Those will be integrated into every day's life and make many technologies more accessible and easier to use.

One of the biggest disadvantages of the current state of development of AR is certainly the artificial look of the computer-generated objects that are placed in real scenes. Most of the times it is possible to tell which objects are real and which have been added at the first glance. In some cases this will not be a problem, on the other hand there are many applications where everything should look of a piece.

Shadows are of significant importance for realistic perception. They allow a correct estimation of the placement of objects in vertical direction [5], [6] and of the distance to the camera. Furthermore, it defines the volume of the object. Overall shadows contain a big amount of additional information, which obviously lacks if shadows are missing. They play an important role in how realistic a scene looks [7].

At the moment, different algorithms are developed to integrate shadow into Augmented Reality-scenes. One of those [2] is examined in this paper and one of the main goals is to prove the importance of the shadows. Furthermore, the functionality of this algorithm is tested and shortcomings are pointed out. This paper will go more into the detail than the experiments done in [9] and will answer concrete, practical questions.

This will create a solid foundation for further development, as it is not wise to code without knowing what is important or put efforts in details that are not worth it.

1.2 Structure of the Paper

At first, chapter 2 is supposed to provide a short overview of shadow algorithms that have already been developed for Augmented Reality-environments, including the implementation developed in [2]. It also provides a brief analysis of how recently developed shadowing algorithms could be used for AR.

Chapter 3 describes the general approach that was taken. It also provides details about the experiments and how they were carried out.

In chapter 4, the execution of the testing scenarios with several subjects is described, along with the results gathered through the experiments.

Chapter 5 provides a conclusion together with a short summary of the most important results and facts.

Chapter 2

Overview of Shadow Algorithms

This chapter takes a look at new developments concerning shadow algorithms. A special emphasis is put on their possible integration in Augmented Reality environments. Moreover, existing implementations of shadows for AR are introduced.

2.1 Existing Implementations for AR

Some algorithms introducing shadows to AR have already been developed. The common goal is to provide accurate and good-looking shadows with as little processing time as possible and without too much setup-work. At the time of writing, there is no implementation that fulfils all three criteria. Nevertheless, the following two algorithms do a fairly good job, even though each of them has its own shortcomings and advantages.

2.1.1 Rapid Shadow Generation

This project was developed by Simon Gibson from the University of Manchester and Alan Chalmers from the University of Bristol [4].

Approach

This algorithm leads to nearly photorealistic pictures. A big disadvantage is that setting it up requires a lot of work. First, it is necessary to calibrate the camera. Afterwards the scene has to be approximately rebuilt in the computer using simple 3D primitives. Eventually the light sources of the room have to be recorded using a light-probe[1]. Using the data recorded

[1]A highly reflecting sphere has been used, which is photographed to get a surround view image of the room.

(a) (b)

Figure 2.1: Screenshots of the 'Rapid Shadow Generation'-Algorithm. (a) shows an interactively rendered picture with a virtual table. In (b) a real table is shown. Concerning the look of the shadows, nearly no difference is visible. This picture was taken from [4].

by this probe, a light map can be generated, which is projected on the 3D representation of the room.

Next, an efficient process determines which light rays are obscured by a virtual object. Using this information multiple *Shadow Maps* are generated, which are blended and – using an appropriate number of them – create a highly realistic soft shadow in the real scene. Fig. 2.1 shows a comparison between a rendered scene and a real photo.

Integration in AR

Because of the hardware-acceleration current graphics cards offer, this algorithm is able to generate impressing frame rates on well-equipped standard PCs. On the other hand, it was only implemented for static scenes, even though the authors state that the same algorithm should work for filmed material as well. To accomplish that, the camera position and orientation would have to be tracked.

Compared to *shadowAReality*, this algorithm is able to generate better and more realistic looking results. However, a big amount of work has to be done for setting everything up. Moreover, *shadowAReality* is focused on live video data because it uses the interactive marker system *ARToolkit* provides [3].

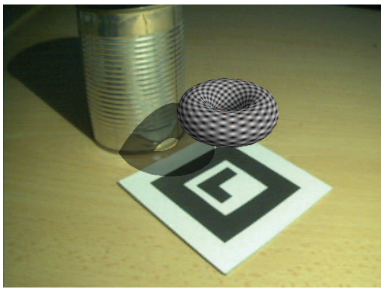

Figure 2.2: A sample image of *shadowAReality*. This picture was taken from [2].

2.1.2 shadowAReality

This study is based on *shadowAReality*, which was developed by Stephan A. Drab for his master thesis at the University of Applied Sciences for Multimedia Technology and Design in Hagenberg, Upper Austria [2].

It uses the *ARToolkit*, which automatically tracks the position and orientation of flat surfaces using special markers. Based on the data gathered through the tracking-process, it can calculate the perspective and the position of the markers and consequently place correctly projected virtual objects in the real scene, without setting up a complex scene geometry.

To produce correct shadows of virtual geometry on real objects (like a can) that are not flat planes defined by markers (like a table), these have to be approximately reconstructed using 3D-software. Only the surface is relevant, as the model is only needed to distort the shadow.

Besides a few models and the definition of the light source, nearly no further configuration work has to be done. As *shadowAReality* uses *shadow volumes* with hard edges, the resulting shadow does not look overly realistic. However, the algorithm is fast and sufficient for many applications. Fig. 2.2 illustrates a typical picture generated by this algorithm.

2.2 New Methods for Calculating Shadows

Lately existing shadow algorithms have been tweaked and improved to provide better results while requiring less computing power. They would be candidates to be included in AR environments as they have several advantages. This chapter provides a short description of them and points out how they could be useful for Augmented Reality.

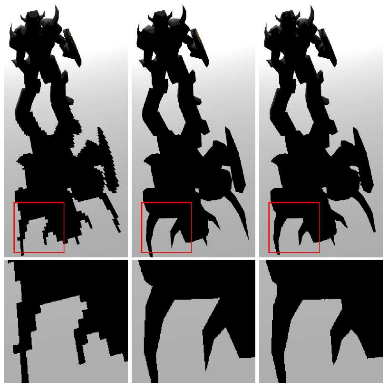

Figure 2.3: Comparison of *Silhouette Maps* to other algorithms. (Left) Normal *Shadow Map*. (Center) *Shadow Volumes* (Right) *Silhouette Map* at the same resolution as the *Shadow Map*. This picture was taken from [8].

2.2.1 Smoothies, Penumbra Maps and Silhouette Maps

The initial planning phase of *shadowAReality* already started in 2002, since then new algorithms for calculating shadows have been developed. For the prototype, *Shadow Maps* have not been chosen mainly because of their inferior visual quality, instead *Shadow volumes* have been used. This generates shadows that are more exact; however, the performance is dependent on the level of detail of the objects.

A new algorithm called *Silhouette Maps* makes it possible to eliminate the main disadvantage of *Shadow Maps* and generates non-pixelized contours. When using normal resolutions of the maps, the quality is equal to the shadows generated by *Shadow Volumes* as depicted in Fig. 2.3. To make this possible, additional information is stored about the edges of the *Shadow Maps*. This generates high quality shadows with nearly no performance reduction [8].

Even more interesting are *Penumbra Maps* [10] and *Smoothies* [1], which were developed independently from each other. They are both based on the same principles and are only different in details of the implementation.

Fig. 2.4 shows the basic principle of those algorithms. Several small extra objects are placed around the silhouettes of the objects (as seen from the light source). Next, like in a normal *Shadow Map*-algorithm, the system checks if the light ray is blocked on its way from the light source to the surface. The difference for this algorithm is that when the rays hit such an extra object called *Smoothie*, they are only reduced in brightness instead of set to just light (bright) or shadow (dark). This generates a soft gradient and is a good simulation of real *Soft Shadows*.

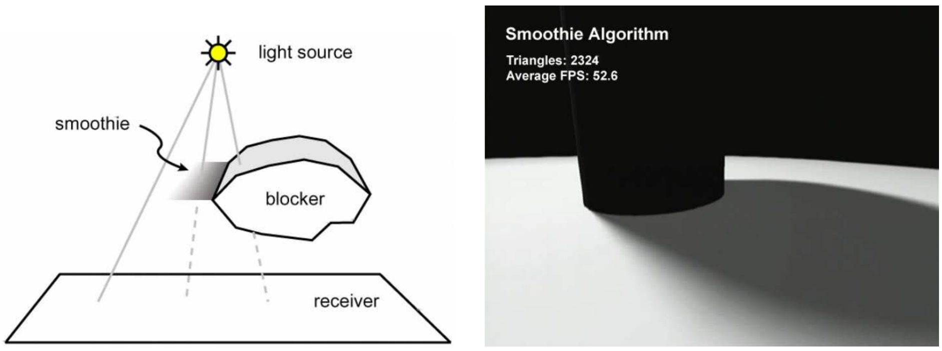

(a) (b)

Figure 2.4: The left picture shows the basic technique of the *Smoothie*-algorithm. On the right hand side there is an example of a shadow rendered using this technique. These pictures were taken from [1].

Integration in AR

Calculating real *Soft Shadows* by blending many hard *Shadow Maps* as done in the *Rapid Shadow*-algorithm (described in section 2.1.1) generates the most accurate results. As speed is the most important factor for AR-applications and *Smoothies* simulate soft shadows quite well, they would be more than sufficient in most situations and are a good compromise of speed and accuracy.

Chapter 3

Design and Approach

This chapter provides a short overview of why this project was realized, some of the considerations that were important and it explains how the tests were done.

3.1 Goal

The goal of this project is to study the effects that shadows have on Augmented Reality-environments. An important aspect that is examined is the technical limitation of current technology, based on one of the algorithms which renders virtual shadows in filmed scenes in real time. Thus, the ultimate goal is to gain a better understanding on how shadows influence and improve the perception and which factors should be considered when implementing them.

3.2 Experimental Approach

To answer the questions, the chosen approach was to let participants fill out a questionnaire. The most important aspect when designing the questions was to get significant results to be able to make clear statements. Therefore the first experiment is based on exact measurements of accuracy and time, in the second one each scenario has only one right answer (out of three), the other parts of the questionnaire asked the participants about their own judgment, like if they think a scene looks realistic or not.

By asking 23 subjects, it is possible to calculate a profound mean and through statistical tests the results are interpreted. The interesting background is that the instructions did not contain exact definitions on what's real or how something has to be to look good. These are factors that are different for every person; however, the experiments showed that there is a high correlation between the subjects, allowing an precise interpretation.

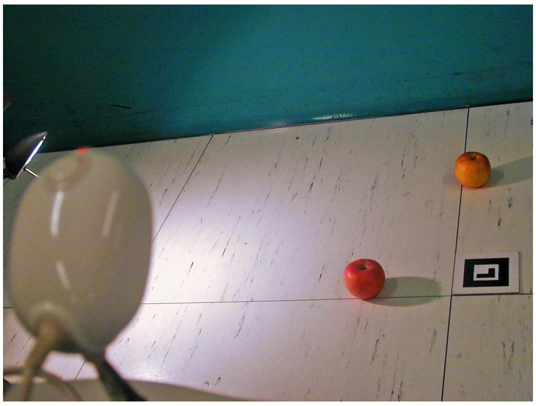

Figure 3.1: Photo of the test setup for experiment 3 concerning light position.

Figure 3.2: A picture of one of the subjects filling out the questionnaire for experiment 4 concerning shadow detail.

3.3 Visual Factors

As can be seen in Fig. 3.1, the experiments were done using a live camera image. The camera position stayed the same through all experiments. Using live pictures instead of screenshots leads to more realistic results as the study analyzes real time augmented reality, which has more inaccuracies and a worse image quality than when using photographs. For example the virtual objects always wobble a bit as the marker detection is not perfect.

A marker used by the AR Toolkit was positioned at a fixed position on the floor, which stayed the same for all experiments. A spot light source was placed 60 cm to the left of it, in a height of 30 cm. The camera was in a fixed place 40 cm in front of the marker, in a height of 1m.

For filming the scene, a *Philips PCVC740K ToUcam Pro*-Webcam was used, which produces a high image quality at a resolution of 640×480. The subjects sat at a PC with the scene setup positioned on the floor behind them and watched the filmed scene through a 17" monitor set to a resolution of 800×600 (See Fig. 3.2).

3.4 Experimental Details

Students of the university were asked to participate; most of them did not have special experience or knowledge about the topic. Each test took about 15 minutes, the participants had to fill out a two-paged questionnaire that can be found in appendix A. Everyone got a short introduction explaining the reason for the tests as well as the the tasks for each test scenario. To prevent a direct influence from one scene to the next one, there was a short pause between the individual scenarios. Furthermore, the series were presented in a randomized order rather than sequential. After the experiment was completed, the participants got a small present to thank them for their cooperation.

Chapter 4

Test Scenarios

This chapter features a description of the experiments, their results and significance. For the tests modified versions of the original *shadowAReality*-program from [2] have been used. The source code, the program and the scenario files can be found on the CD-ROM (see appendix B).

4.1 Estimation of Height

This experiment is supposed to show how important shadows are to estimate the vertical position [1] of objects. It is expected that with activated shadows, users will be able to position objects on the floor in shorter time and with higher accuracy.

4.1.1 Setup

As in all experiments, the camera was pointed at the floor with an AR-marker defining the ground plane. A virtual, rendered object floated in the 3D-space. This experiment was done with two objects, a cube and a sphere, and for each object two times – once without shadows, the other time with enabled shadows.

The objective was to position the floating object directly on the floor as fast as possible. To achieve this, the participant moved the object up and down with the cursor keys. He did not get feedback nor statistics on how near he had already moved the object to the floor – he had to estimate it. The main factor that helped for this were shadows, for the scenarios without shadows only the perspective of the virtual object compared to the real scene was of help. When he thought the object was positioned exactly on the table, he had to acknowledge the final position by pressing the ESC-key. After that, the accuracy (distance to the floor) and the time needed are printed out. Fig. 4.1 shows screenshots of both the cube and

[1] height or distance to the floor

sphere scenarios. For each object there are three pictures: one of the initial position without shadows, with shadows and one of the exact position on the floor.

The reason for using two different objects, a cube and a sphere, was the expected difference in the difficulty of positioning them as precisely as possible. While the object contact can be judged more easily for cubes when the shadow meets directly with the object borders, it is more difficult to position a round sphere without edges on the floor, when the scene is not viewed directly from the side. Additionally, for each test the object is initially positioned at a different height and position.

A test setup with some similarities was examined in [5], where accuracy and time of object positioning were measured, with / without the help of shadows as well as interreflections[2]. The most important difference is that they used a complete immersive virtual reality environment in contrast to the augmented reality setup our experiments are based on. Their result was that shadows do help a bit to position the objects accurately, especially with binocular vision in a head mounted display.

4.1.2 Results

The very first experiment was not measured to give the subjects a chance to acquaint themselves with the program and their objective. The results of the following four tests are shown in the graphs. Fig. 4.2 shows how precise the subjects were able to position the object on the floor. The accuracy is depicted for both objects, both with and without shadows. The error bars show the 95% confidence interval.

To analyze the graphs, an independent t-test was used. This method of analysis calculates if there is a significant difference between the scenarios with and without shadows for the same object. It was found that the presence of a shadow is a statistically highly significant factor that increases accuracy. p-Factors below .05 are significant, below .01 highly significant. In this case, they were far below .01 (Cube: $p < .001$, Sphere: $p < .001$). This validates the assumption that the presence of shadows is very important to estimate the distance of the object to the floor.

Nevertheless, shadows do not always have to be a perfect indicator for distance. Apart from factors concerning the position of the light source, which would require further investigation, the form of the object is important as well. While the cube has straight edges that meet with the shadow-edges when the cube is positioned exactly on the floor, the sphere has a curved surface so that the contact of the shadow edge with the object edge can no longer be used to judge if the sphere touches the floor. Therefore, most of

[2]the *shadowAReality*–algorithm used in this study is not able to calculate and render interreflections

(a)

(b)

(c)

(d)

(e)

(f)

Figure 4.1: Screenshots of the scenario *Estimation of Height*. (a) to (c) show the scenes with the cube as the interactive object, (d) to (f) used the sphere. (a) and (d) show the initial position of the test without shadows. As the results prove, it is difficult to judge the distance to the floor. (b) and (e) is what the scenes looks like with shadows enabled. Note the different starting position of every experiment. (c) and (f) are screenshots taken when the object is positioned exactly on the floor. In the case of the sphere (f), many subjects believed the sphere still has to be moved down to be positioned correctly.

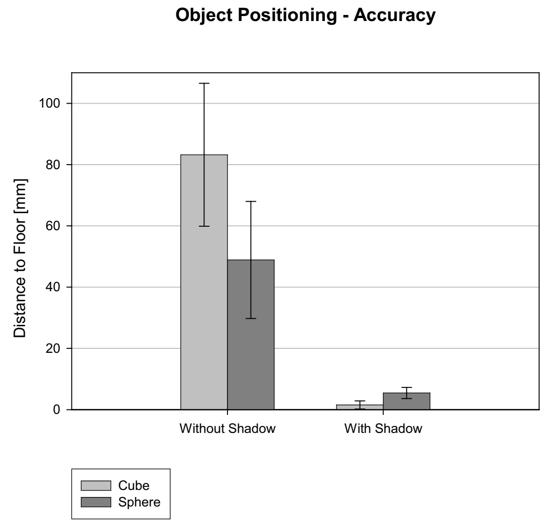

Figure 4.2: Diagram of the results of the first test. It shows the measured accuracy of both scenarios with and without the activated shadow, for cube and sphere.

the subjects positioned the sphere too far down so that a part of the sphere would be below the floor. See the screenshot Fig. 4.1(f).

The second measurement tested the time until subjects were confident that the object was positioned on the floor. The results are visualized in Fig. 4.3. The expectation that the participants felt more secure with a shadow proved to be true for the cube where the difference of the needed time between activated and deactivated shadows is again highly significant (Cube: $p < .001$). For the sphere the time needed with shadow in comparison to the same scene without shadows is again shorter, however it fails to meet the limit to be statistically significant by a tiny bit (Sphere: $p = .058$). An interpretation for this (which was also observed during the tests) would be that the subjects noticed that the shadow edge in this case does not meet with the object edge and therefore were not as sure about the correct position as they were when positioning the cube.

4.2 Estimation of Depth

The objective of this experiment is to find out if and how shadows can help to judge the depth[3] of a virtual 3D-object in a real scene. For many applications which are using Augmented Reality, it is important that its users can estimate the position exactly. The question is if shadows are

[3] distance to the camera

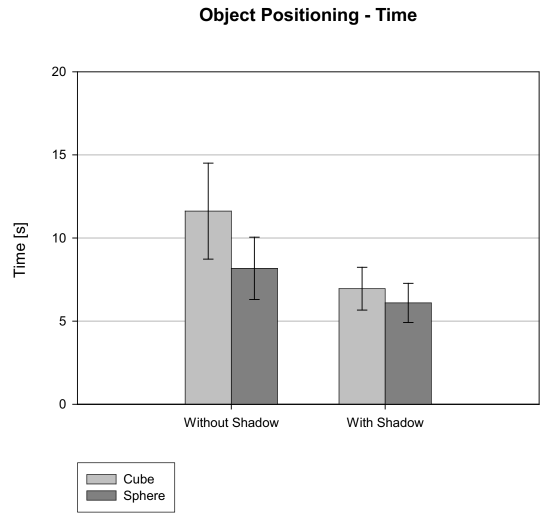

Figure 4.3: The measured time the subjects needed until they were confident that they positioned the object on the floor. The graph visualizes the results with and without shadows, for both objects.

enough to make this possible and if subjects can tell the correct result even if the objects are positioned in a way that it looks like the other way round.

4.2.1 Setup

In this experiment, two virtual spheres floated above the ground. For each scenario, they varied in size and depth[4]. Each scene was presented both with and without shadows, however not directly one after the other and in a randomized order. That way the subjects did not remember the scenes. Fig. 4.4 combines screenshots of all scenes.

For each scene the user had to tell which sphere was nearer to him and which one was farther away, or if both had the same distance to the camera. The subjects were asked to stay with their first idea and not to think that they might have been fooled and because of that decide the other way round. This was important because not only the position of the spheres changed for each experiment, but also the size. The idea was to give subjects the impression that for example the bigger sphere was nearer to the camera while in reality it was farther away and just bigger than the other sphere [6, chapter 3.1.3]. By comparing the results of this scene with and without shadows, it gets more evident if and how much shadows really help in judging the depth.

[4]distance from the camera

(a) Scene 1

(b) Scene 1

(c) Scene 5

(d) Scene 5

(e) Scene 3

(f) Scene 3

(g) Scene 4

(h) Scene 4

(i) Scene 2

(j) Scene 2

Figure 4.4: Screenshots of all experiments of *Estimation of Depth* in the order they are discussed in the text and presented in the diagram, on the left without shadow and on the right hand side with shadows enabled.

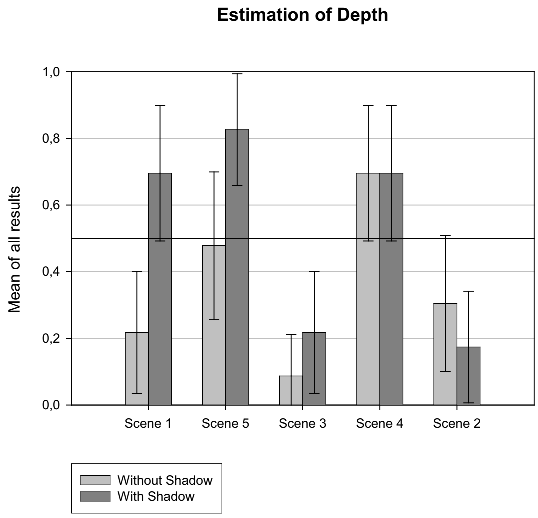

Figure 4.5: Diagram of the results of scenario 2, *Estimation of Depth*. The height of the bar indicates how many subject answered correctly, each scene was tested with and without shadows.

4.2.2 Results

Scenes 1 and 5 were associated with each other as the spheres in both scenarios had a different distance to the camera. In scene 5, the smaller sphere was farther away than the bigger sphere. This meets the expectations we normally have that objects, which are closer, are also bigger. In scenario 5, the scene was setup in a way that this situation was the other way round and the sphere being and appearing bigger was in reality farther away. The results prove that the assumption that this is a misleading setup without shadows, was right. Only 21.8% were still able to answer correctly whereas about 69.6% chose the right answer when shadows were enabled.

In general, the diagram in Fig. 4.5 shows that for both scenarios shadows helped subjects a lot to find out which sphere was nearer. In scene 1, the difference between the experiment with and without shadows is highly significant ($p = 0.001$), for scene 5 it is also significant with $p = 0.013$. Comparing both tests, the number of people answering correctly when shadows are enabled is nearly the same, and in both cases quite high (70.0%, 82.6%). Without shadows, more people chose the correct answer when the setup was not misleading. (47.8% compared to 21.7%).

Another interesting pair of tests that can be compared is scene 3 with scene 4. In both cases, the two spheres had the same distance to the camera as well as the same size. However, in scenario 3 the height is different while in scene 4 both spheres have the same distance to the floor.

The results are quite interesting. When they are at the same height, the result of the scenario with and without shadows is exactly the same – 69.6% answered correctly. In the scene where they have a different height, most subjects chose the wrong answer, even shadows did not help a lot and there is no statistically significant difference between the scenario with and without shadows.

An interpretation of those results is that shadows do not always help people to find out the correct position of objects in a three-dimensional environment. For some applications, it might be more useful to use the limited processing power of today's computers to make the object more realistic than investing it in producing a good-looking shadow.

In Scene 2, the percentage of people answering correctly is even lower when shadows are enabled. Both spheres had the same distance to the camera as well as the same height; however, they were given a different size. Therefore, the bigger sphere (left) appears to be nearer – the answer most of the subjects chose. The danger seems to be that many rely too much on the shadows and as the bigger sphere has a larger shadow, even less people chose the correct answer. There is no statistically significant difference though.

4.3 Light Position

In an Augmented Reality-scene, the virtual light position has to be defined manually. In the past, it was not always necessary that this matched the real light position closely, however now with shadows any inaccuracies are more likely to be seen. The aim of this experiment is to find out how precise the positioning has to be.

4.3.1 Setup

Two apples were positioned in the scene as a reference for the shadow. In between those two, the computer rendered a virtual apple. Two test series modified the position of the virtual light source gradually. The first decreased the height so that the shadow got longer; the second series put the virtual lamp farther away from the camera so that the direction of the shadow changed. The real lamp always stayed in the same position; therefore, the shadows of the real apples always stayed the same and the shadow of the virtual apple changed in each scene. Screenshots of this are presented in Fig. 4.6; Fig. 3.1 is a photo of this scene.

The subjects were asked to judge whether they thought the shadow was calculated correctly or if the light source was in a wrong position. The scenes were presented in a randomized order, between them there were short breaks.

In one scene (shown with the title "*Wrong*" in Fig. 4.7) the shadow was projected in the completely wrong direction. This was presented as the last question of the first page of the questionnaire before subjects had to turn the page and helped keeping them motivated.

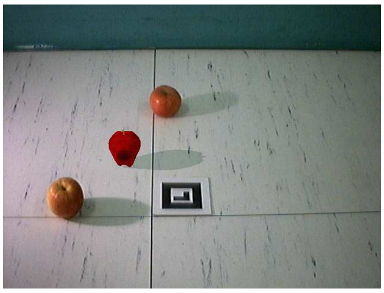

(a) Correct Position

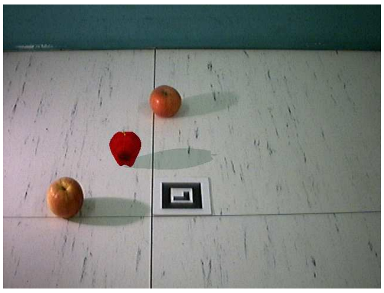

(b) Height −30 mm

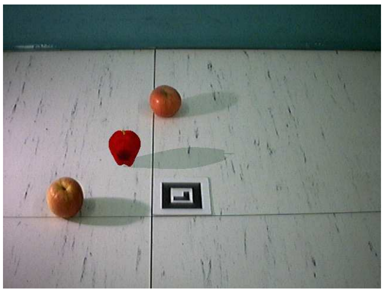

(c) Height −60 mm

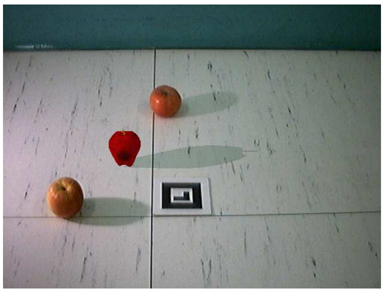

(d) Height −90 mm

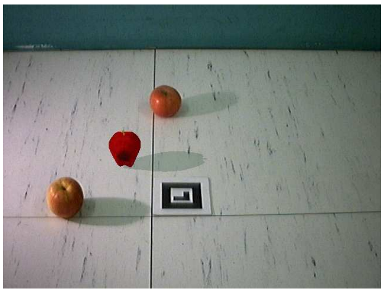

(e) Depth +30 mm

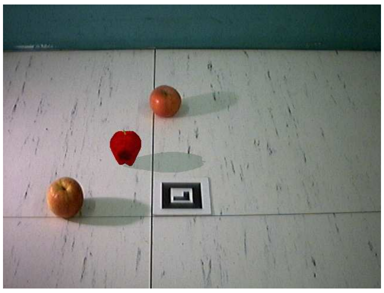

(f) Depth +60 mm

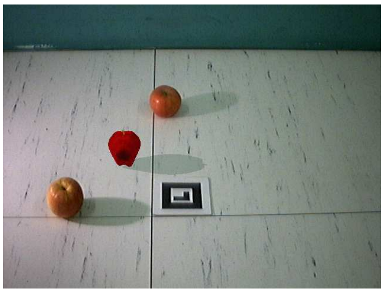

(g) Depth +90 mm

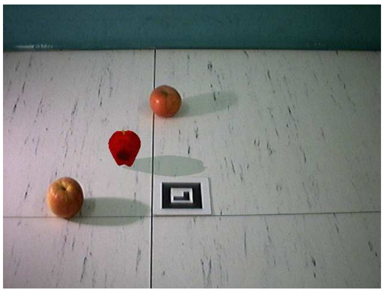

(h) Depth +120 mm

Figure 4.6: Screenshots of all scenes of the experiment *Light-Position*. (a) shows the reference picture with the correct light position. (b) to (d) are pictures of the first series where the virtual light source is moved down, making the shadow longer. In (e) to (h) the virtual light is moved backwards so that the direction of the shadow changes.

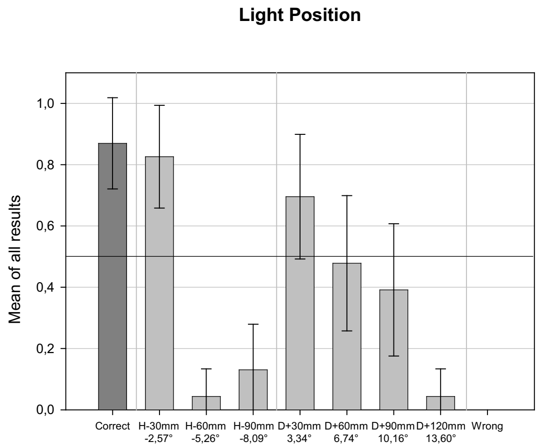

Figure 4.7: The results of scenario 3, *Light Position.* The bar on the left represents the mean of all answers for the correctly positioned virtual light source. The next three bars are of the first test series where the height of the light source was reduced so that the shadow got longer. In the next series, represented by the other four bars, the light source was moved backwards so that the direction of the shadow changed.

4.3.2 Results

The scene where the virtual light was positioned exactly where the real lamp was, was judged to be correct by most subjects (87.0%). In the diagram Fig. 4.7 this scene is represented by the bar on the far left side. The first series to be analyzed is when the height of the virtual light position is modified.

The farther down the light source is positioned, the longer the shadow gets. As can be seen in the screenshots, the effect of this is quite visible and already in the second scenario of this series, the shadow is so long that it is no longer judged as realistic by the subjects. The difference is of course highly significant with $p < 0.001$. If the light source is only 3 cm wrong (the light rays meet the apple at a 2.57° lower angle than in the correct position), there is no significant difference to the correct position.

In comparison to that, moving the light source the same amount backwards (away from the camera) in the second series leads to an initially slightly worse result, however still not statistically significant compared to the correct position ($p = 0.153$). Even more interesting is the second case, when the light source is moved a total of 6 cm backwards. About half of the subjects (47.8%) judged the shadow still as correct, which already corresponds to a highly significant difference compared to the correct position ($p = 0.005$).

When looking at the results of both series, it is obvious that a change of the direction of the shadow is less visible than a wrong length. When the light source is moved 6 cm backwards (difference to original position: 6.74°), the result is highly significant better ($p < 0.001$) than when it is moved 6 cm down (5.26° difference).

In general the virtual light source does not have to match the real light source exactly, however especially if the length of the shadow is wrong, the error gets obvious quite soon. Of course, it is important to consider that for these experiments the subjects had to take a quite close look at the shadows and had comparison objects; in most applications, the shadow will just be a supporting visual factor that will not be the center of attention. It can be assumed that in this case, everyone will be much more error-tolerant and higher inaccuracies can be allowed.

4.4 Required Detail of Shadows

shadowAReality uses volumetric shadows. Most of the times these are calculated using simplified virtual objects to reduce the required computing time. However if the source object of the shadow volume has less polygons, the resulting shadow will look less realistic. As Augmented Reality-Applications can be quite complex, it would not be possible to use much processing power for ultra-realistic shadows. This experiment will answer the question how far details can be reduced so that shadows still look realistic and/or acceptable.

4.4.1 Setup

For this experiment, a high-quality torus floated above the floor. The shadow was generated using several different levels of detail; the highest had no visible edges while the lowest was basically only a pentagon. All scenes were shown in a randomized order with short pauses between the individual live-pictures. Screenshots are presented in Fig. 4.8.

Using the questionnaire, the subjects were asked whether the shadow looked realistic (yes/no) and acceptable (yes/no). They were also told that in all scenarios the shadow was calculated correctly and absolutely fine concerning its position and general appearance.

As all the other models in this paper, the torus was generated in 3DS Max (by AutoDesk). Using this application, it is possible to specify the level of detail for basic objects. For the torus these are radial segments[5] and sides[6]. As the shadow is only a flat, projected version of the 3D model, the number of sides are not visible, however they were reduced as well to provide a more life-like test setup as in real applications one wouldn't use a torus with only

[5]number of divisions around the torus.

[6]The number of sides on the cross-sectional circle of the torus.

(a) 30 × 15

(b) 25 × 10

(c) 15 × 7

(d) 12 × 6

(e) 10 × 5

(f) 7 × 4

(g) 5 × 3

Figure 4.8: Screenshots of all scenes of the experiment *Shadow Detail*. They are ordered from the highest quality (a) to the lowest quality (g). [*radial segments* x *sides of cross-sectional circle*]

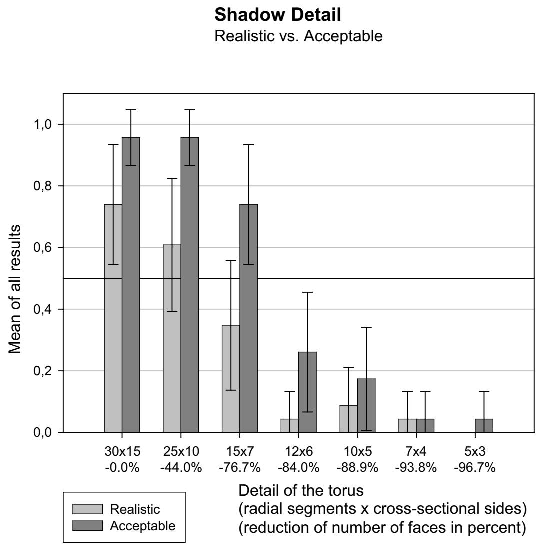

Figure 4.9: Diagram of the results of test 4, *Required Detail of Shadows*.

three radial segments and 15 sides. The shadow model was generated in seven different levels of detail.

4.4.2 Results

Several aspects of the diagram Fig. 4.9 can be analyzed. It was expected that shadows would be more acceptable than realistic. The graph proves this, and for every level of detail, more people judged the shadow as acceptable than realistic.

Realistic Shadows

Looking at the part showing the realism-answers, the lower the detail gets, there is a continuous decline in the number of people who judge the shadow as realistic. An interesting fact is that even with the highest level of detail (30×15) which matches the source object and where no edges are visible anymore, was judged as realistic by only 73.9% of the subjects. This might be because of the hard shadow borders and because the shape of this hard shadow looks uncommon even though it is calculated correctly. With the light source coming from the left, the upper and lower sides of the shadow are quite narrow while the left and right side are broad. With soft shadows, this effect would not have been less, resulting in a more familiar shadow-shape.

There is no statistically significant difference between the first two levels of detail ($p = 0.345$), however comparing the original level (30×15) with the third level (15×7), the difference is highly significant ($p = 0.008$). This means that the number of radial segments for the shadow object can be reduced by 1/6th, leading to a reduction in the number of faces by 44%, and the result will still look as realistic as the highest detail version.

Acceptable Shadows

As said before, the number of people judging a shadow as acceptable for a computer game or any other Augmented Reality-application is generally high. For the first two levels of detail, the number of people judging the shadow as acceptable is equally high (95.7%).

The first significant difference is between the second and the third scene with $p = 0.040$ and 73.9% still judging the shadow as acceptable. The fourth shadow is clearly not acceptable anymore. If the 15×7 level is good enough, depends on the application. Quite often this will be the case, resulting in a possible reduction of faces of the 3D object by -76.7%.

General Interpretation

The results and statements of the subjects indicate that the level of detail of shadows is not that important as long as the edges are not too obvious. To make a shadow look good, it has to have a shape we are used to, which can sometimes be difficult to achieve with hard shadows.

The tests show that the level of detail for the object, which is used as the source for shadow projection can, be reduced by up to 75% and it will still look acceptable. As chapter 4.3.2 describes, it might be possible that subjects are much more tolerant if the shadow of objects is not their center of attention. After all shadows in current computer games are most of the times less detailed than the level found as acceptable in this experiment.

4.5 Hard vs. Soft Shadows

The most significant disadvantage of the *shadowAReality*-algorithm that was used is that it projects shadows with hard edges. In reality, this would only be possible with a perfect point-light source that cannot exist. In all other cases, the shadow has a certain penumbra. Calculating soft shadows is much more difficult however – the questions in this section try to find out if a soft shadow would be worth the effort.

4.5.1 Setup

This scene used more complex objects to provide shadows that do not have basic shapes like a sphere or a cube. Therefore, it consisted of a real and a virtual rubber duck, which do not look exactly alike as only the appearance

Figure 4.10: Screenshot of the scene used for the two questions of the section *Hard vs. Soft Shadows*.

of the shadow mattered for this experiment. Fig. 4.10 is a screenshot of this scene.

The subjects were asked two questions. The first one being: “The virtual and the real shadow have a different appearance. Nevertheless, do you still think that the shadows of this scene are visual consistent (= monolithic)?” The second question was “Does the difference in the visual appearance of the shadows have a negative or disturbing effect on the visual quality of the picture?” The subjects were given a choice of yes and no.

4.5.2 Results

While the other sections of this paper analyzed measurements or series of scenarios, these two questions were very subjective and directly represented the opinion of the subjects. The answers are visualized in Fig. 4.11.

Only 21.7% of the participants thought that the two shadows are visually consistent. This is not a very high number and therefore there is still a long road to go until virtual objects included in Augmented Reality-environments will look truly realistic.

One of the subjects pointed out that apart from the hard shadow, another problem with the virtual shadow was the aliasing of the edges. Implementing full screen anti-aliasing improves the situation and could possibly lead to slightly better results.

The second question was whether the different appearance of the shadows had a negative effect. 56.5% of the subjects responded with yes, which is about half of the subjects. It is certainly important to know that number because it can be used as a reference for own applications. If shadows play only a supporting role, hard shadows should not be a problem for the visual appearance. If they play an integral part or if the AR-application is for

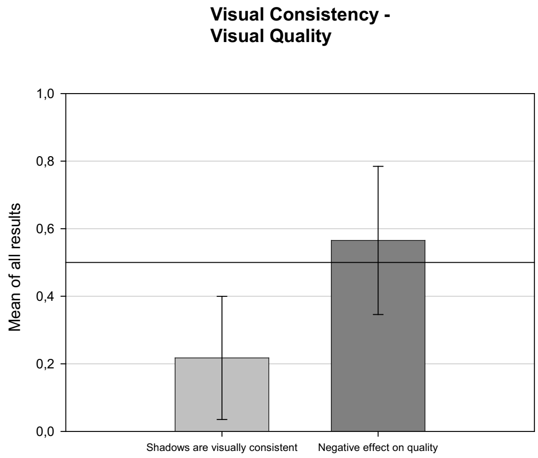

Figure 4.11: Diagram showing the mean of the answers for the two questions concerning *Hard vs. Soft Shadows*.

demonstration purposes, the additional effort of soft shadows might well be worth it.

Chapter 5

Summary and Conclusion

The graphics community is currently doing a lot of research work on how Augmented Reality can be improved. Even more important is that features, which are already implemented, are analyzed, so that they are not just existent but also useful and perform as perfect as possible.

This paper takes a look at many aspects concerning shadows, which are important for Augmented Reality. Virtual objects often do not fit in well into the filmed scene, as currently there are no implementations that render realistic objects for AR in real time. Therefore, without shadows, they look like they have been glued into the filmed image. Shadows play an important role in giving them a position in the three dimensional space.

The first experiment (chapter 4.1) proves this assumption and shows that it is rather impossible to judge the distance to the floor correctly without shadows. It is easier to do that in completely virtual environments as shown in [5]; however, in AR perspective and lightning do not perfectly match, making this task enormously harder.

Even with added shadows, it is still necessary to plan the scenes carefully. The second experiment (chapter 4.2) shows that in some situations the information shadows provide might even be misleading. As a result, just enabling shadows in AR-applications does not remove the need to think about the layout of the picture and the tasks the users have to fulfill.

When setting up the scenes, special accuracy is required as well. Experiment 3 (chapter 4.3) tested how exact the virtual light position has to match the real light situation and found out that while there is indeed a certain tolerance, it is not very big. Especially if the shadows are too long, users will quickly notice that something is wrong.

In the quest to make the scene look as good as possible, the technical barrier will soon be reached. Current hardware is not able to do everything as good as we want to in real time. Therefore, there have to be some compromises. An important part is the level of detail of the shadows, as shadow calculation needs much processing power. The fourth experiment (chapter 4.4) shows that the level of detail can be reduced from an exact calculation to a simplified model by more than 75% and the shadow will still look acceptable, which is certainly enough for most application.

The questions in experiment five (chapter 4.5) showed the limits of the technology used. While shadows help a lot to improve the scene, they do not look realistic and do not fit in well into the rest of the image. This aspect leaves much room for improvement.

This can supposedly be done with soft shadows. Using the new algorithms presented in chapter 2.2, it would be possible to do that in real time. It remains for future work to examine how much this can improve the experience of Augmented Reality.

Appendix A

Questionnaire

This is a translated version of the German questionnaire that the subjects had to fill out. The layout differs from the original ones.

A.1 Estimation of Height

It is your task to place the object as accurately as possible on the floor using the cursor keys. When you think that the object should now touch the floor, press the ESC-key as fast as possible. You will do the test two times - once with and once without shadows. However, the starting position and height of the object will be different every time.

Scene **1** (Cube) - **without** shadow:	Time: ________	Accuracy: ________
Scene **1** (Cube) - **with** shadow:	Time: ________	Accuracy: ________
Scene **2** (Sphere) - **without** shadow:	Time: ________	Accuracy: ________
Scene **2** (Sphere) - **with** shadow:	Time: ________	Accuracy: ________

A.2 Estimation of Depth

In the following tests you will see two spheres with random size placed at different positions in the room. Which sphere is do you think nearer to the camera (in the front)? Please tell your honest estimation and do not choose the other answer because you think you might have been tricked.

	Left	**Right**	**Equal**
1.	☐	☐	☐
2.	☐	☐	☐
3.	☐	☐	☐
4.	☐	☐	☐
5.	☐	☐	☐
6.	☐	☐	☐
7.	☐	☐	☐
8.	☐	☐	☐
9.	☐	☐	☐
10.	☐	☐	☐

A.3 Light Position

You will now see several scenes that closely resemble each other. Please state for each one of them if you think that the shadow has been projected correctly. This can be the case in more than one of the scenes, but it does not have to.

	Yes	**No**
1.	☐	☐
2.	☐	☐
3.	☐	☐
4.	☐	☐
5.	☐	☐
6.	☐	☐
7.	☐	☐
8.	☐	☐
9.	☐	☐

A.4 Required Detail of Shadows

for Shadows The shadows will have a different level of detail in the following scenarios. Please state for every scene, if this shadow is... ("Realistic" means that the shape of the shadow looks real. "Acceptable" means that you could live with this shadow if it looked like that in a real application).

... realistic?			**... acceptable?**		
	Yes	No		Yes	No
1.	☐	☐	1.	☐	☐
2.	☐	☐	2.	☐	☐
3.	☐	☐	3.	☐	☐
4.	☐	☐	4.	☐	☐
5.	☐	☐	5.	☐	☐
6.	☐	☐	6.	☐	☐
7.	☐	☐	7.	☐	☐

A.5 Hard vs. Soft Shadows

Shadows The virtual and the real shadow have a different appearance. Nevertheless, do you still think that the shadows of this scene are visual consistent (= monolithic)?

Yes	**No**
☐	☐

Does the difference in the visual appearance of the shadows have a negative or disturbing effect on the visual quality of the picture?

Yes	**No**
☐	☐

Appendix B

Contents of the CD-ROM

File System: Joliet

Mode: Single-Session (CD-ROM)

B.1 Term Paper

Path: /

ba.dvi	Paper (as DVI-File, without graphics)
ba.pdf	Paper (PDF-File)
ba.ps	Paper (PostScript-File)

B.2 shadowAReality

Normal version Used for all scenarios except the first. Differs from the original version from [2] in two aspects. 1. The light source is set to invisible by default. 2. The resolution is fixed to 640 × 480 and no configuration window pops up when starting the application.

Path: /shadowAReality/

build/ARToolKit/ . . .	Configuration files of ARToolkit
build/geometry/	3DS-Models and textures of the virtual objects
build/geometry/max/ .	3DS Max-scenes containing the objects
build/scenes/	Scene definition Files
include/	Headerfiles of the libraries used
lib/	Libraries
src/	Source Code

Movable Version Used in chapter 4.1. Modified in the same way as the normal version, in addition the first virtual object of the scene is movable with the cursor keys. After exiting, time and accuracy are printed out.

Path: /shadowAReality-movable/

build/ARToolKit/ . . .	Configuration files of ARToolkit
build/geometry/	3DS-Models and textures of the virtual objects
build/geometry/max/ .	3DS Max-scenes containing the objects
build/scenes/	Scene definition Files
include/	Headerfiles of the libraries used
lib/	Libraries
src/	Source Code

B.3 Pictures and Diagrams

Path: /images/

screenshots/	Screenshots of all scenarios described in chapter 4
diagrams/	Graphs visualizing survey results
examples/	Other pictures related to algorithms for shadows in Augmented Reality
photos/	Pictures taken during the tests

Bibliography

[1] Chan, E. and F. Durand: *Rendering fake soft shadows with smoothies*. In *Proceedings of the 13th Eurographics workshop on Rendering*, pp. 208–218. Eurographics Association, 2003.

[2] Drab, S. A.: *Echtzeitrendering von Schatten virtueller Objekte in realen Umgebungen*. Master's thesis, Fachhochschule Hagenberg, Medientechnik und -design, Hagenberg, Austria, Juni 2003.

[3] Gibson, S. and A. Chalmers: *ARtoolkit Manual 2.33*, December 1999.

[4] Gibson, S., J. Cook, T. Howard and R. Hubbold: *Rapid shadow generation in real-world lighting environments*. In *Proceedings of the 13th Eurographics workshop on Rendering*, pp. 219–229. Eurographics Association, 2003.

[5] Hu, H. H., A. A. Gooch, W. B. Thompson, B. E. Smits, J. J. Rieser and P. Shirley: *Visual Cues for Imminent Object Contact in Realistic Virtual Environments*. In *Visualization*, Salt Lake City, Utah, October 2000. IEEE.

[6] Interrante, V.: *Illustrating Transparency: communicating the 3D shape of layered transparent surfaces via texture*. PhD thesis, University of North Carolina at Chapel Hill, 1996.

[7] Rademacher, P., J. Lengyel, E. Cutrell and T. Whitted: *Measuring the Perception of Visual Realism in Images*. In *12th Eurographics Workshop on Rendering*, pp. 235–248, London, 2001.

[8] Sen, P., M. Cammarano and P. Hanrahan: *Shadow silhouette maps*. ACM Trans. Graph., 22(3):521–526, 2003.

[9] Sugano, N., H. Kato and K. Tachibana: *The Effects of Shadow Representation of Virtual Objects in Augmented Reality*. In *2003 IEEE / ACM International Symposium on Mixed and Augmented Reality (ISMAR 2003)*, pp. 76–83. IEEE Computer Society, October 2003.

[10] Wyman, C. and C. Hansen: *Penumbra maps: approximate soft shadows in real-time*. In *Proceedings of the 13th Eurographics workshop on Rendering*, pp. 202–207. Eurographics Association, 2003.